JFP **35**, e25, 61 pages, 2025. © The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

doi:10.1017/S0956796825100142

Hefty algebras: Modular elaboration of higher-order effects

CAS VAN DER REST
Shielded Technologies, London, UK
(*e-mail:* cas.vanderrest@shielded.io)

CASPER BACH
University of Southern Denmark, Odense, Denmark
Delft University of Technology, Delft, Netherlands
(*e-mail:* casperbach@imada.sdu.dk)

Abstract

Algebraic effects and handlers is an increasingly popular paradigm for programming with effects. A key benefit is modularity: programs with effects are defined against an interface of operations, allowing the implementation of effects to be defined and refined without changing or recompiling programs. The behavior of effects is specified using equational theories, with equational proofs inheriting the same modularity. However, higher-order operations (that take computations as arguments) break this modularity: while they can often be encoded in terms of algebraic effects, this typically breaks modularity as operations defined this way are not encapsulated in an interface, inducing changes to programs and proofs upon refinement of the implementation. In this paper, we show that *syntactic overloading* is a viable solution to this modularity problem by defining *hefty algebras*: a formal framework that captures an overloading-based semantics of higher-order effects by defining modular elaborations from higher-order effect trees into primitive algebraic effects. We demonstrate how this approach scales to define a wide range of known higher-order effects from the literature and develop modular higher-order effect theories and modular reasoning principles that build on and extend the state of the art in modular algebraic effect theories. We formalize our contributions in Agda.

1 Introduction

Defining abstractions that support both programming with and reasoning about side effects is a research question with a long and rich history. The goal is to define an abstract interface of (possibly) side-effecting operations together with equations describing their behavior, where the interface hides operational details about the operations and their side effects that are irrelevant for defining or reasoning about a program. Such encapsulation makes it easy to refactor, optimize, or even change the behavior of a program while preserving proofs, by changing the implementation of the interface.

Monads (Moggi, 1989*b*) have long been the preferred solution to this research question, but they lack modularity: given two computations defined in different monads, there is

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

no canonical way to combine them that is both universally applicable and preserves modular reasoning. This presents a problem for scalability since, in practice, programs, and therefore proofs, are developed incrementally. *Algebraic effects and handlers* (Plotkin & Power, 2002; Plotkin & Pretnar, 2009) provide a solution for this problem by defining a syntactic class of monads, which permits composition of syntax, equational theories, and proofs. Algebraic effects and handlers maintain a strict separation of *syntax* and *semantics*, where programs are only syntax, and semantics is assigned later on a per-effect basis using handlers.

Many common operations, however, cannot be expressed as syntax in this framework. Specifically, *higher-order operations* that take computational arguments, such as exception, catching or modifying environments in the reader monad. While it is possible to express higher-order operations by inlining handler applications within the definition of the operation itself, this effectively relinquishes key modularity benefits. The syntax, equational theories, and proofs of such inlined operations do not compose.

In this paper, we propose to address this problem by appealing to an *overloading* mechanism which postpones the choice of handlers to inline. As we demonstrate, this approach provides a syntax and semantics of higher-order operations with similar modularity benefits as traditional algebraic effects; namely syntax, equational theories, and proofs that compose. To realize this, we use a syntactic class of monads that supports higher-order operations (which we dub *hefty trees*). Algebras over this syntax (*hefty algebras*) let us modularly *elaborate* this syntax into standard algebraic effects and handlers. We show that a wide variety of higher-order operations can be defined and assigned a semantics this way. Crucially, program definitions using hefty trees enjoy the same modularity properties as programs defined with algebraic effects and handlers. Specifically, they support the composition of syntax, semantics, equational theories, and proofs. This demonstrates that overloading is not only syntactically viable but also supports the same modular reasoning as algebraic effects for programs with side effects that involve higher-order operations.

1.1 Background: Algebraic effects and handlers

To understand the modularity benefits of algebraic effects and handlers, and why this modularity breaks when defining operations that take computations as parameters, we give a brief introduction to algebraic effects. To this end, we will use informal examples using a simple calculus inspired by Pretnar's calculus for algebraic effects (Pretnar, 2015). Section 2 of this paper provides a semantics for algebraic effects and handlers in Agda which corresponds to this calculus.

1.1.1 Effect signatures

Say we want an effectful operation *out* for printing output. Besides its side effect of printing output, the operation should take a string as an argument and return the unit value. Using algebraic effects, we can declare this operation using the following *effect signature*:

$$\textbf{effect}\ \mathit{Output} = \mathit{out} : \mathit{String} \to ()$$

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

We can use this operation in any program that has the *Output* effect. For example, the following *hello* program:

$$\begin{aligned} &\mathit{hello} : () \,!\, \mathit{Output} \\ &\mathit{hello} = \mathit{out}\ \text{“Hello”};\ \mathit{out}\ \text{“ world!”} \end{aligned}$$

The type $() \,!\, \mathit{Output}$ indicates that *hello* is an effectful computation which returns a unit value, and which is allowed to call the operations in *Output* (i.e., only the *out* operation).

More generally, computations of type $A \,!\, \Delta$ are allowed (but not required) to call any operation of any effect in Δ, where Δ is a *row* (i.e., unordered sequence) of effects. An *effect* is essentially a label associated with a set of operations. The association of labels to operations is declared using effect signatures, akin to the signature for *Output* above.

1.1.2 Effect theories

A crucial feature of algebraic effects and handlers is that it permits abstract reasoning about programs containing effects, such as *hello* above. That is, each effect is associated with a set of laws that characterizes the behavior of its operations. Their purpose is to constrain an effect's behavior without appealing to any specifics of the implementation of the effects. Consequently, program proofs derived from these equations remain valid for all handler implementations satisfying the laws of its equational theory.

Importantly, these laws are purely *syntactic*, in the sense that they are part of the effect's specification rather than representing universal truths about the behavior of effectful computation. Whether a law is "valid" depends entirely on how we handle the effects, and different handlers may satisfy different laws. Figuring out a suitable set of laws is part of the development process of (new) effects. Typically, the final specification of an effect is the result of a back-and-forth refinement between an effect's specification and its handler implementations. This process ultimately converges to a definition that matches our intuitive understanding of what an effect should do.

For example, the *Output* effect has a single law that characterizes the behavior of *out*:

$$\mathit{out}\ s_1; \mathit{out}\ s_2 \equiv \mathit{out}\ (s_1 +\!\!+ s_2)$$

Here, $+\!\!+$ is string concatenation. Using this law, we can prove that our *hello* program will print the string "Hello world!". Crucially, this proof does not depend on operational implementation details of the *Output* effect. Instead, it uses the laws of the equational theory of the effect. While the program and effect discussed so far has been deliberately simple, the approach illustrates how algebraic effects let us reason about effectful programs in a way that abstracts from the concrete implementation of the underlying effects.

1.1.3 Effect handlers

An alternative perspective is to view effects as interfaces that specify the parameter, return type, and laws of each operation. Implementations of such interfaces are given by *effect handlers*. An effect handler essentially defines how to interpret operations in the execution context they occur in. This interpretation must be consistent with the laws of the effect. (We will not dwell further on this consistency here; we return to this in Section 5.6.)

The type of an effect handler is $A \,!\, \Delta \Rightarrow B \,!\, \Delta'$, where Δ is the row of effects before applying the handler and Δ' is the row after. For example, here is a specific type of an

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

effect handler for *Output*:[1]

$$hOut : A \,!\, Output, \Delta \Rightarrow (A \times String) \,!\, \Delta$$

The *Output* effect is being handled, so it is only present in the effect row on the left.[2] As the type suggests, this handler handles *out* operations by accumulating a string of output. Below is an example implementation of this handler:

$$\begin{aligned} hOut = \textbf{handler}\ \{\ & (\textbf{return}\ x) \mapsto \textbf{return}\ (x, \text{“”}) \\ & (out\ s; k) \mapsto \textbf{do}\ (y, s') \leftarrow k\ (); \ \textbf{return}\ (y, s +\!\!+ s')\ \} \end{aligned}$$

The **return** case of the handler says that, if the computation being handled terminates normally with a value x, then we return a pair of x and the empty string. The case for *out* binds a variable s for the string argument of the operation, but also a variable k representing the *execution context* (or *continuation*). Invoking an operation suspends the program and its execution context up-to the nearest handler of the operation. The handler can choose to re-invoke the suspended execution context (possibly multiple times). The handler case for *out* above always invokes k once. Since k represents an execution context that includes the current handler, calling k gives a pair of a value y and a string s', representing the final value and output of the execution context. The result of handling *out* s is then y and the current output (s) plus the output of the rest of the program (s').

In general, a computation $m : A \,!\, \Delta$ can only be run in a context that provides handlers for each effect in Δ. To this end, the expression **with** h **handle** m represents applying the handler h to handle a subset of effects of m. For example, we can run the *hello* program from earlier with the handler *hOut* to compute the following result:

$$(\textbf{with}\ hOut\ \textbf{handle}\ hello) \equiv ((), \text{“Hello world!”})$$

The key benefit of algebraic effects and handlers is that programs such as *hello* are defined *independently* of how the effectful operations they use are implemented. This makes it possible to reason about programs independently of how the underlying effects are implemented and also makes it possible to refine, refactor, and optimize the semantics of operations, without having to modify the programs that use them. For example, we could refine the meaning of *out without modifying the hello program or proofs derived from equations of the Output effect*, by using a different handler which prints output to the console. However, some operations are challenging to express while retaining the same modularity benefits.

1.2 The modularity problem with higher-order operations

We discuss the problem with defining higher-order operations using effect signatures (Section 1.2.1) and potential workarounds (Sections 1.2.2 and 1.2.3).

[1] Here and throughout the rest of this paper, type variables that are not explicitly bound elsewhere are implicitly universally quantified in prenex position of the type in which they occur.

[2] *Output* could occur in the universally quantified Δ too. This raises the question: which *Output* effect does a given handler actually handle? We refer to the literature for answers to this question; see, e.g., the row treatment of Morris & McKinna (2019), the *effect lifting* of Biernacki *et al*. (2018), and the *effect tunneling* of Zhang & Myers (2019).

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

1.2.1 The problem

Say we want to declare an operation *censor f m*, which applies a censoring function $f : \mathit{String} \to \mathit{String}$ to the side-effectful output of the computation m. We might try to declare an effect *Censor* with a *censor* operation by the following type:

$$\mathit{censor} : (\mathit{String} \to \mathit{String}) \to A \,!\, \mathit{Censor}, \Delta \to A \,!\, \mathit{Censor}, \Delta$$

However, using algebraic effects, we cannot declare *censor* as an operation.

The problem is that effect signatures do not offer direct support for declaring operations with computation parameters. Effect signatures have the following shape:

$$\textbf{effect}\ E = \mathit{op}_1 : A_1 \to B_1 \mid \cdots \mid \mathit{op}_n : A_n \to B_n$$

Here, each operation parameter type A_i is going to be typed as a value. While we may pass around computations as values, passing around computations as arguments of computations is not a desirable approach to defining higher-order operations in general. We will return to this point in Section 1.2.2.

The fact that effect signatures do not directly support operations with computational arguments is also evident from how handler cases are typed (Pretnar, 2015, Fig. 6):

$$\textbf{handler}\ \{\ \cdots\ (\mathit{op}\ \underbrace{v}_{A}\ ;\ \underbrace{k}_{B \to C \,!\, \Delta'}\ ) \mapsto\ \underbrace{c}_{C \,!\, \Delta'},\ \cdots \} \qquad (*)$$

Here, A is the argument type of an operation, and B is the return type of an operation. The term c represents the code of the handler case, which must have type $C \,!\, \Delta'$.

Observe how only the continuation k that is statically known to have computation type which matches the effects of the context in which the handler is applied. While the argument type A could be instantiated with a computation type $A \,!\, \Delta''$, the effects of this computation are hardcoded in the definition of the operation. Because handlers are agnostic to the row Δ' of effects that they do not handle, and since in general $\Delta' \neq \Delta''$, we are forced, in a clear violation of modularity, to hardcode the handler for Δ'' as well. As a result, the only option for defining operations with computation parameters that preserves modularity is to encode them in the continuation k.

A consequence of this observation is that we can only define and modularly handle higher-order operations whose computation parameters are *continuation-like*. Following Plotkin & Power (2003), such operations satisfy the following law, known as the *algebraicity property*. For any operation $\mathit{op} : A \,!\, \Delta \to \cdots \to A \,!\, \Delta \to A \,!\, \Delta$ and any $m_1, \ldots, m_n$ and k,

$$\textbf{do}\ x \leftarrow (\mathit{op}\ m_1 \ldots m_n); k\ x \ \equiv\ \mathit{op}\ (\textbf{do}\ x \leftarrow m_1; k\ x) \ldots (\textbf{do}\ x \leftarrow m_n; k\ x) \qquad (\dagger)$$

The law says that the computation parameter values $m_1, \ldots, m_n$ are only ever run in a way that *directly* passes control to k. Such operations can without loss of generality or modularity be encoded as operations *without* computation parameters;[3] i.e., as algebraic operations that match the handler typing in $(*)$ above.

[3] Concretely, we can represent the operation in question as $\mathit{op}\ m_1 \ldots m_n = \textbf{do}\ x \leftarrow \mathit{op}'\ (); \mathit{select}\ x$ where $\mathit{op}' : () \to D^n \,!\, \Delta$ and $\mathit{select} : D^n \to A \,!\, \Delta$ is a function that chooses between n different computations using a data type D^n whose constructors are $d_1, \ldots, d_n$ such that $\mathit{select}\ d_i = m_i$ for $i \in \{1, \ldots, n\}$.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

Some higher-order operations obey the algebraicity property; many do not. Examples that do not obey algebraicity include:

- Exception handling: let *catch* m_1 m_2 be an operation that handles exceptions thrown during evaluation of computation m_1 by running m_2 instead, and *throw* be an operation that throws an exception. These operations are not algebraic. For example,

$$\textbf{do}\ (\mathit{catch}\ m_1\ m_2); \mathit{throw} \not\equiv \mathit{catch}\ (\textbf{do}\ m_1; \mathit{throw})\ (\textbf{do}\ m_2; \mathit{throw})$$

- Local binding (the *reader monad* Jones, 1995): let *ask* be an operation that reads a local binding, and *local* r m be an operation that makes r the current binding in computation m. Observe:

$$\textbf{do}\ (\mathit{local}\ r\ m); \mathit{ask} \not\equiv \mathit{local}\ r\ (\textbf{do}\ m; \mathit{ask})$$

- Logging with censoring (an extension of the *writer monad* Jones, 1995): let *out* s be an operation for logging a string, and *censor* f m be an operation for post-processing the output of computation m by applying $f : \mathit{String} \to \mathit{String}$.[4] Observe:

$$\textbf{do}\ (\mathit{censor}\ f\ m); \mathit{out}\ s \not\equiv \mathit{censor}\ f\ (\textbf{do}\ m; \mathit{out}\ s)$$

- Function abstraction as an effect: let *abs* x m be an operation that constructs a function value binding x in computation m, *app* v m be an operation that applies a function value v to an argument computation m, and *var* x be an operation that dereferences a bound x. Observe:

$$\textbf{do}\ (\mathit{abs}\ x\ m); \mathit{var}\ x \not\equiv \mathit{abs}\ x\ (\textbf{do}\ m; \mathit{var}\ x)$$

1.2.2 Potential workaround: Computations as arguments of operations

A tempting possible workaround to the issues summarized in Section 1.2.1 is to declare an effect signature with a parameter type A_i that carries effectful computations. However, this workaround can cause operations to escape their handlers. Following Pretnar (2015), the semantics of effect handling obeys the following law.[5] If h handles operations other than *op*, then:

$$\textbf{with}\ h\ \textbf{handle}\ (\textbf{do}\ x \leftarrow \mathit{op}\ v; k\ x) \equiv \textbf{do}\ x \leftarrow \mathit{op}\ v; (\textbf{with}\ h\ \textbf{handle}\ k\ x) \qquad (\dagger)$$

This law tells us that effects in v will not be handled by h. This is problematic if h is the intended handler for one or more effects in v. The solution we describe in Section 1.3 does not suffer from this problem.

Nevertheless, for applications where it is known exactly which effects v contains, we can work around the issue by encoding computations as argument values of operations. We consider how and discuss the modularity problems that this workaround suffers from. The following *Censor* effect declares the type of an operation *censorOp* (f, m) where f is

[4] The *censor* operation is a variant of the function by the same name the widely used Haskell `mtl` library: `https://hackage.haskell.org/package/mtl-2.2.2/docs/Control-Monad-Writer-Lazy.html`.

[5] This law concerns so-called *deep handlers*. However, the semantics of so-called *shallow handlers* (Lindley *et al*., 2017; Hillerström & Lindley, 2018) exhibit similar behavior.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

a censoring function and m is a computation encoded as a value argument:[6]

$$\textbf{effect}\ \mathit{Censor} = \mathit{censorOp} : (\mathit{String} \to \mathit{String}) \times (A\ !\ \mathit{Censor}, \mathit{Output}) \to A$$

This effect can be handled as follows:

$$\begin{aligned}
&\mathit{hCensor} : A\ !\ \mathit{Censor}, \mathit{Output} \Rightarrow A\ !\ \mathit{Output}\\
&\mathit{hCensor} = \textbf{handler}\ \{\ (\mathit{censorOp}\ (f, m); k)\ \mapsto\ \textbf{do}\\
&\qquad\qquad (x, s) \leftarrow \textbf{with}\ \mathit{hOut}\ \textbf{handle}\ (\textbf{with}\ \mathit{hCensor}\ \textbf{handle}\ m)\\
&\qquad\qquad \mathit{out}\ (f\ s)\\
&\qquad\qquad k\ x\\
&\qquad (\textbf{return}\ x)\ \mapsto\ \textbf{return}\ x\ \}
\end{aligned}$$

The handler case for *censorOp* runs m with handlers for both the *Output* and *Censor* effects, which yields a pair (x, s) where x represents the value returned by m, and s represents the (possibly sub-censored) output of m. We then output the result of applying the censoring function f to s, before passing x to the continuation k.

This handler lets us run programs such as the following:

$$\begin{aligned}
&\mathit{censorHello} : () \,!\, \mathit{Censor}, \mathit{Output}\\
&\mathit{censorHello} = \mathit{censorOp}\ ((\lambda s.\ \textbf{if}\ (s \equiv \text{“Hello”})\ \textbf{then}\ \text{“Goodbye”}\ \textbf{else}\ s), \mathit{hello})
\end{aligned}$$

Applying *hCensor* and *hOut* yields the printed output “Hello world!”, because “Hello world!” $\not\equiv$ “Hello”:

$$\textbf{with}\ \mathit{hOut}\ \textbf{handle}\ (\textbf{with}\ \mathit{hCensor}\ \textbf{handle}\ \mathit{censorHello}) \equiv ((), \text{“Hello world!”})$$

As the example above illustrates, it is sometimes possible to encode higher-order effects as algebraic operations. However, encoding higher-order operations in this way suffers from a modularity problem. Say we want to extend our program with a new effect for throwing exceptions—i.e., an effect with a single operation *throw*—and a new effect for catching exceptions—i.e., an effect with a single operation *catch* $m_1\ m_2$ where exceptions thrown by m_1 are handled by running m_2. The *Throw* effect is a plain algebraic effect, defined as follows.[7]

$$\textbf{effect}\ \mathit{Throw} = \mathit{throw} : () \to \bot$$

The *Catch* effect is higher-order. We can again encode it as an operation with computations as value arguments.

$$\begin{aligned}
\textbf{effect}\ \mathit{Catch} = \mathit{catchOp}\ \ &: A\ !\ \mathit{Catch}, \mathit{Throw}, \mathit{Censor}, \mathit{Output}\\
&\times A\ !\ \mathit{Catch}, \mathit{Throw}, \mathit{Censor}, \mathit{Output}\\
&\to A
\end{aligned}$$

To support subcomputations with exception catching, we need to refine the computation type we use for *Censor*. (This refinement could have been done modularly if we had used a more polymorphic type.)

$$\textbf{effect}\ \mathit{Censor} = \mathit{censorOp} : (\mathit{String} \to \mathit{String}) \times (A\ !\ \mathit{Catch}, \mathit{Throw}, \mathit{Censor}, \mathit{Output}) \to A$$

[6] The self-reference to *Censor* in the computation parameter requires type-level recursion that is challenging to express in, e.g., the Agda formalization of algebraic effects we present in Section 2. However, such type-level recursion is supported by, e.g., Frank (Lindley *et al.*, 2017), Koka (Leijen, 2017), and in a Haskell embeddings of algebraic effects (Kammar *et al.*, 2013; Wu *et al.*, 2014).

[7] Here $\bot$ is the empty type.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

The modularity problem arises when we consider whether to handle *Catch* or *Censor* first. If we handle *Censor* first, then we get exactly the problem described earlier in connection with the law (†): the handler *hCensor* is not applied to sub-computations of *catchOp* operations. Let us consider the consequences of this. Say we want to define a handler for *catchOp* of the following type:

$$hCatch : A \,!\, Catch, Throw, Output \Rightarrow A \,!\, Throw, Output$$

Any such handler which runs the sub-computation m_1 of an operation $catchOp\ m_1\ m_2$ must hardcode a handler for the *Censor* effect. Otherwise the handler would allow effects to escape in a way that breaks the typing discipline of algebraic effects and handlers (Pretnar, 2015). To illustrate why this is the case, consider the following program.

$$\textbf{with}\ hCatch\ \textbf{handle}\ (\textbf{with}\ hCensor\ \textbf{handle}\ catchOp\ censorHello\ m_2)$$

Per equation 12 of Pretnar (2015, Figure 7), this program is equivalent to:

$$\textbf{with}\ hCatch\ \textbf{handle}\ catchOp \quad \underbrace{censorHello}_{\text{Still has the } Censor \text{ effect!}} \quad m_2$$

Which means that *hCatch* is now responsible for handling the remaining *Censor* effect in the first sub-computation of *catchOp*, otherwise it violates the promise of its type that the resulting computation does not have the *Censor* effect. While for some applications it may be acceptable to hardcode handler invokations this way, it is non-modular and should not be—and, indeed, is not–necessary.

Before showing the solution, we propose which avoids this, we first consider a different workaround (Section 1.2.3) and previous solutions proposed in the literature (Section 1.2.4).

1.2.3 Potential workaround: Define higher-order operations as handlers

It is possible to define many higher-order operations in terms of algebraic effects and handlers. For example, instead of defining *censor* as an operation, we could define it as a function which uses an inline handler application of *hOut*:

$$\begin{aligned} &censor : (String \to String) \to A \,!\, Output, \Delta \to A \,!\, Output, \Delta \\ &censor\ f\ m = \textbf{do}\ (x, s) \leftarrow (\textbf{with}\ hOut\ \textbf{handle}\ m);\ out\ (f\ s);\ \textbf{return}\ x \end{aligned}$$

Defining higher-order operations in terms of standard algebraic effects and handlers in this way is a key use case of effect handlers (Plotkin & Pretnar, 2009). In fact, all other higher-order operations above (with the exception of function abstraction) can be defined in a similar manner.

However, it is unclear what the semantics is of composing syntax, equational theories, and proofs of programs with such functions occuring inline in programs. We address this gap by proposing notions of syntax and equational theories for programs with higher-order operations. The semantics of such programs and theories is given by elaborating them into standard algebraic effects and handlers.

1.2.4 Previous approaches to solving the modularity problem

The modularity problem of higher-order effects, summarized in Section 1.2.1, was first observed by Wu *et al*. (2014) who proposed *scoped effects and handlers* (Wu *et al*., 2014;

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

Piróg *et al*., 2018; Yang *et al*., 2022) as a solution. Scoped effects and handlers work for a wide class of effects, including many higher-order effects, providing similar modularity benefits as algebraic effects and handlers when writing programs. Using *parameterized algebraic theories* (Lindley *et al*., 2024; Matache *et al*., 2025), it is possible reason about programs with scoped effects independently of how their effects are implemented.

Van den Berg *et al*. (2021) recently observed, however, that operations that defer computation, such as evaluation strategies for λ application or *(multi-)staging* (Taha & Sheard, 2000), are beyond the expressiveness of scoped effects. Therefore, van den Berg *et al*. (2021) introduced another flavor of effects and handlers that they call *latent effects and handlers*. It remains an open question how to reason about latent effects and handlers independently of how effects are implemented.

In this paper, we demonstrate that an overloading-based approach provides a semantics of composition for effect theories that is comparable to standard algebraic effects and handlers, and, we expect, to parameterized algebraic theories. Furthermore, we demonstrate that the approach lets us model the syntax and semantics of key examples from the literature: optionally transactional exception catching, akin to scoped effect handlers (Wu *et al*., 2014), and evaluation strategies for effectful λ application (van den Berg *et al*., 2021). Formally relating the expressiveness of our approach with, e.g., scoped effects and parameterized algebraic theories, is future work.

1.3 Solution: Elaboration algebras

We propose to define operations such as *censor* from Section 1.2 as *modular elaborations* from a syntax of higher-order effects into algebraic effects and handlers. To this end, we introduce a new type of *computations with higher-order effects*, which can be modularly translated into computations with only standard algebraic effects:

$$A \,!!\, H \xrightarrow{elaborate} A \,!\, \Delta \xrightarrow{handle} Result$$

Here $A \,!!\, H$ is a computation type where A is a return type and H is a row comprising both algebraic and higher-order effects. The key idea is that the higher-order effects in the row H are modularly elaborated into a computation with effects given by the row Δ. To achieve this, we define *elaborate* such that it can be modularly composed from separately defined elaboration cases, called elaboration *algebras* (for reasons explained in Section 3). $A \,!!\, H \Rrightarrow A \,!\, \Delta$ as the type of elaboration algebras that elaborate the higher-order effects in H to Δ, we can modularly compose any pair of elaboration algebras $e_1 : A \,!!\, H_1 \Rrightarrow A \,!\, \Delta$ and $e_2 : A \,!!\, H_2 \Rrightarrow A \,!\, \Delta$ into an algebra $e_{12} : A \,!!\, H_1, H_2 \Rrightarrow A \,!\, \Delta$.[8]

Elaboration algebras are as simple to define as non-modular elaborations such as *censor* (Section 1.2.3). For example, here is the elaboration algebra for the higher-order *Censor* effect whose only associated operation is the higher-order operation $censor_{op} : (String \to String) \to A \,!!\, H \to A \,!!\, H$:

$$eCensor : A \,!!\, Censor \Rrightarrow A \,!\, Output, \Delta$$
$$eCensor\ (censor_{op}\ f\ m;\ k) = \mathbf{do}\ (x, s) \leftarrow (\mathbf{with}\ hOut\ \mathbf{handle}\ m);\ out\ (f\ s);\ k\ x$$

[8] Readers familiar with data types à la carte (Swierstra, 2008) may recognize this as the usual closure of algebras under functor coproducts.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

The implementation of *eCensor* is essentially the same as *censor*, with two main differences. First, elaboration happens in-context, so the value yielded by the elaboration is passed to the context (or continuation) k. Second, and most importantly, programs that use the $censor_{op}$ operation are now programmed against the interface given by *Censor*, meaning programs do not (and *cannot*) make assumptions about how $censor_{op}$ is elaborated. As a consequence, we can modularly refine the elaboration of higher-order operations such as $censor_{op}$, without modifying the programs that use the operations. Similarly, we can define equational theories that constrain the behavior elaborations of higher-order operations and write proofs about programs using higher-order operations that are valid for any elaboration that satisfies these equations.

For example, here is again a program which censors and replaces "Hello" with "Goodbye":[9]

$$\begin{aligned} &censorHello : () \,!!\, Censor, Output \\ &censorHello = censor_{op}\ (\lambda s.\ \textbf{if}\ (s \equiv \text{“Hello”})\ \textbf{then}\ \text{“Goodbye”}\ \textbf{else}\ s)\ hello \end{aligned}$$

Say we have a handler $hOut' : (String \to String) \to A\,!\,Output, \Delta \Rightarrow (A \times String)\,!\,\Delta$ which handles each operation *out s* by pre-applying a censor function $(String \to String)$ to s before emitting it. Using this handler, we can give an alternative elaboration of $censor_{op}$ which post-processes output strings *individually*:

$$\begin{aligned} &eCensor' : A \,!!\, Censor \Rrightarrow A\,!\,Output, \Delta \\ &eCensor'\ (censor_{op}\ f\ m;\ k) = \textbf{do}\ (x, s) \leftarrow (\textbf{with}\ hOut'\ f\ \textbf{handle}\ m);\ out\ s;\ k\ x \end{aligned}$$

In contrast, *eCensor* applies the censoring function $(String \to String)$ to the batch output of the computation argument of a $censor_{op}$ operation. The batch output of *hello* is "Hello world!" which is unequal to "Hello", so *eCensor* leaves the string unchanged. On the other hand, $eCensor'$ censors the individually output "Hello":

$$\begin{aligned} &\textbf{with}\ hOut\ \textbf{handle}\ (\textbf{with}\ eCensor\ \textbf{elaborate}\ censorHello) \equiv ((), \text{“Hello world!”}) \\ &\textbf{with}\ hOut\ \textbf{handle}\ (\textbf{with}\ eCensor'\ \textbf{elaborate}\ censorHello) \equiv ((), \text{“Goodbye world!”}) \end{aligned}$$

This gives higher-order operations the same modularity benefits as algebraic operations for defining programs. In Section 5, we show that these modularity benefits extend to program reasoning as well.

1.4 Contributions

This paper formalizes the ideas sketched in this introduction by shallowly embedding them in Agda. However, the ideas transcend Agda. Similar shallow embeddings can be implemented in other dependently typed languages, such as Idris (Brady, 2013*a*); but also in less dependently typed languages like Haskell, OCaml, or Scala.[10] By working in a dependently typed language, we can state algebraic laws about interfaces of effectful operations,

[9] This program relies on the fact that it is generally possible to lift computation $A\,!\,\Delta$ to $A\,!!\,H$ when $\Delta \subseteq H$.

[10] The artifact accompanying this paper (van der Rest & Poulsen, 2024) contains a shallow embedding of elaboration algebras in Haskell.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

and prove that implementations of the interfaces respect the laws. We make the following technical contributions:

- Section 2 describes how to encode algebraic effects in Agda, revisits the modularity problem with higher-order operations, and summarizes how scoped effects and handlers address the modularity problem, for some (*scoped* operations) but not all higher-order operations.
- Section 3 presents our solution to the modularity problem with higher-order operations. Our solution is to (1) type programs as *higher-order effect trees* (which we dub *hefty trees*) and (2) build modular elaboration algebras for folding hefty trees into algebraic effect trees and handlers. The computations of type $A \mathbin{!!} H$ discussed in Section 1.3 correspond to hefty trees, and the elaborations of type $A \mathbin{!!} H \Rrightarrow A \mathbin{!} \Delta$ correspond to hefty algebras.
- Section 4 presents examples of how to define hefty algebras for common higher-order effects from the literature on effect handlers.
- Section 5 shows that hefty algebras support formal and modular reasoning on a par with algebraic effects and handlers, by developing reasoning infrastructure that supports verification of equational laws for higher-order effects such as exception catching. Crucially, proofs of correctness of elaborations are compositional. When composing two proven correct elaboration, correctness of the combined elaboration follows immediately without requiring further proof work.

Section 6 discusses related work and Section 7 concludes. The paper assumes a passing familiarity with dependent types. We do not assume familiarity with Agda: we explain Agda-specific syntax and features when we use them.

An artifact containing the code of the paper and a Haskell embedding of the same ideas is available online (van der Rest & Poulsen, 2024). A subset of the contributions of this paper was previously published in a conference paper (Poulsen & van der Rest, 2023). While that version of the paper too discusses reasoning about higher-order effects, the correctness proofs were non-modular, in that they make assumptions about the order in which the algebraic effects implementing a higher-order effect are handled. When combining elaborations, these assumptions are often incompatible, meaning that correctness proofs for the individual elaborations do not transfer to the combined elaboration. As a result, one would have to re-prove correctness for every combination of elaborations. For this extended version, we developed reasoning infrastructure to support modular reasoning about higher-order effects in Section 5 and proved that correctness of elaborations is preserved under composition of elaborations.

2 Algebraic effects and handlers in Agda

This section describes how to encode algebraic effects and handlers in Agda. We do not assume familiarity with Agda and explain Agda specific notation in footnotes. Sections 2.1 to 2.4 defines algebraic effects and handlers; Section 2.5 revisits the problem of defining higher-order effects using algebraic effects and handlers; and Section 2.6 discusses how scoped effects (Wu *et al*., 2014; Piróg *et al*., 2018; Yang *et al*., 2022) solves the problem for *scoped* operations but not all higher-order operations.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

2.1 Algebraic effects and the free monad

We encode algebraic effects in Agda by representing computations as an abstract syntax tree given by the *free monad* over an *effect signature*. Such effect signatures are traditionally (Swierstra, 2008; Awodey, 2010; Kammar *et al.*, 2013; Wu *et al.*, 2014; Kiselyov & Ishii, 2015) given by a *functor*; i.e., a type of kind Set → Set together with a (lawful) mapping function.[11] In our Agda implementation, effect signature functors are defined by giving a *container* (Abbott *et al.*, 2003, 2005). Each container corresponds to a value of type Set → Set that is both *strictly positive*[12] and *universe consistent*[13] (Martin-Löf, 1984), meaning they are a constructive approximation of endofunctors on Set. Effect signatures are given by a (dependent) record type:[14,15]

```
record Effect : Set₁ where
  field Op  : Set
        Ret : Op → Set
```

Here, Op is the set of operations, and Ret defines the *return type* for each operation in the set Op. The extension of an effect signature, ⟦_⟧, reflects its input of type Effect as a value of type Set → Set:[16]

```
⟦_⟧ : Effect → Set → Set
⟦ Δ ⟧ X = Σ (Op Δ) λ op → Ret Δ op → X
```

The extension of an effect Δ into Set → Set is indeed a functor, as witnessed by the following function:[17]

```
map-sig : (X → Y) → ⟦ Δ ⟧ X → ⟦ Δ ⟧ Y
map-sig f (op , k) = ( op , f ∘ k )
```

As discussed in the introduction, computations may use multiple different effects. Effect signatures are closed under co-products:[18,19]

```
_⊕_ : Effect → Effect → Effect
Op  (Δ₁ ⊕ Δ₂) = Op Δ₁ ⊎ Op Δ₂
Ret (Δ₁ ⊕ Δ₂) = [ Ret Δ₁ , Ret Δ₂ ]
```

[11] Set is the type of types in Agda. More generally, functors mediate between different *categories*. For simplicity, this paper only considers *endofunctors* on Set, where an endofunctor is a functor whose domain and codomain coincides; e.g., Set → Set.

[12] https://agda.readthedocs.io/en/v2.6.2.2/language/positivity-checking.html.

[13] https://agda.readthedocs.io/en/v2.6.2.2/language/universe-levels.html.

[14] https://agda.readthedocs.io/en/v2.6.2.2/language/record-types.html.

[15] The type of effect rows has type Set₁ instead of Set. To prevent logical inconsistencies, Agda has a hierarchy of types where Set : Set₁, Set₁ : Set₂, etc.

[16] Here, Σ : (A : Set) → (A → Set) → Set is a *dependen sum*.

[17] To show that this is truly a functor, we should also prove that map-sig satisfies the *functor laws*. We will not make use of these functor laws in this paper, so we omit them.

[18] The _⊕_ function uses *copattern matching*: https://agda.readthedocs.io/en/v2.6.2.2/language/copatterns.html. The Op line defines how to compute the Op field of the record produced by the function; and similarly for the Ret line.

[19] _⊎_ is a *disjoint sum* type from the Agda standard library. It has two constructors, inj₁ : A → A ⊎ B and inj₂ : B → A ⊎ B. The [_,_] function (also from the Agda standard library) is the *eliminator* for the disjoint sum type. Its type is [_,_] : (A → X) → (B → X) → (A ⊎ B) → X.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

We compute the co-product of two effect signatures by taking the disjoint sum of their operations and combining the return type mappings pointwise. We use co-products to encode effect rows. For example, the effect $\Delta_1 \oplus \Delta_2$ corresponds to the row union denoted as Δ_1, Δ_2 in the introduction.

The syntax of computations with effects Δ is given by the free monad over Δ. We encode the free monad as follows:

```
data Free (Δ : Effect) (A : Set) : Set where
  pure   : A                  → Free Δ A
  impure : ⟦ Δ ⟧ (Free Δ A) → Free Δ A
```

Here, pure is a computation with no side-effects, whereas impure is an operation whose syntax is given by the functor $[\![\Delta]\!]$. By applying this functor to Free Δ A, we encode an operation whose *continuation* may contain more effectful operations.[20] To see in what sense, let us consider an example.

Example. The data type on the left below defines an operation for outputting a string. On the right is its corresponding effect signature.

```
data OutOp : Set where
  out : String → OutOp
```

```
Output : Effect
Op  Output          = OutOp
Ret Output (out s) = ⊤
```

The effect signature on the right says that out returns a unit value ($\top$ is the unit type). Using this, we can write a simple hello world corresponding to the *hello* program from Section 1:

```
hello : Free Output ⊤
hello = impure (out "Hello" , λ _ → impure (out " world!" , λ x → pure x))
```

Section 2.1 shows how to make this program more readable by using monadic **do** notation.

The hello program above makes use of just a single effect. Say we want to use another effect, Throw, with a single operation, throw, which represents throwing an exception (therefore having the empty type $\bot$ as its return type):

```
data ThrowOp : Set where
  throw : ThrowOp
```

```
Throw : Effect
Op  Throw = ThrowOp
Ret Throw throw = ⊥
```

Programs that use multiple effects, such as Output and Throw, are unnecessarily verbose. For example, consider the following program which prints two strings before throwing an exception:[21]

```
hello-throw : Free (Output ⊕ Throw) A
hello-throw = impure (inj₁ (out "Hello") , λ _ →
                impure (inj₁ (out " world!") , λ _ →
                  impure (inj₂ throw , ⊥-elim)))
```

[20] By unfolding the definition of $[\![_]\!]$ one can see that our definition of the free monad is identical to the I/O trees of Hancock & Setzer (2000), or the so-called *freer monad* of Kiselyov & Ishii (2015).

[21] ⊥-elim is the eliminator for the empty type, encoding the *principle of explosion*: $\bot\text{-elim} : \bot \to A$.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

To reduce syntactic overhead, we use *row insertions* and *smart constructors* (Swierstra, 2008).

2.2 Row insertions and smart constructors

A *smart constructor* constructs an effectful computation comprising a single operation. The type of this computation is polymorphic in what other effects the computation has. For example, the type of a smart constructor for the out effect is

```
'out : ⦃ Output ≲ Δ ⦄ → String → Free Δ ⊤
```

Here, the ⦃ Output ≲ Δ ⦄ type declares the row insertion witness as an *instance argument* of 'out. Instance arguments in Agda are conceptually similar to type class constraints in Haskell: when we call 'out, Agda will attempt to automatically find a witness of the right type, and implicitly pass this as an argument.[22] Thus, calling 'out will automatically inject the Output effect into some larger effect row Δ.

We define the ≲ order on effect rows in terms of a different $\Delta_1 \bullet \Delta_2 \approx \Delta$ which witnesses that any operation of Δ is isomorphic to *either* an operation of Δ_1 *or* an operation of Δ_2:[23,24]

```
record _•_≈_ (Δ₁ Δ₂ Δ : Effect) : Set₁ where
  field reorder : ∀ {X} → ⟦ Δ₁ ⊕ Δ₂ ⟧ X ↔ ⟦ Δ ⟧ X
```

Using this, the ≲ order is defined as follows:

```
_≲_ : (Δ₁ Δ₂ : Effect) → Set₁
Δ₁ ≲ Δ₂ = Σ Effect (λ Δ′ → Δ₁ • Δ′ ≈ Δ₂)
```

It is straightforward to show that ≲ is a *preorder*; i.e., that it is a *reflexive* and *transitive* relation.

We can also define the following function, which uses a $\Delta_1 \lesssim \Delta_2$ witness to coerce an operation of effect type Δ_1 into an operation of some larger effect type Δ_2.[25]

```
inj : ⦃ Δ₁ ≲ Δ₂ ⦄ → ⟦ Δ₁ ⟧ A → ⟦ Δ₂ ⟧ A
inj ⦃ _ , w ⦄ (c , k) = w .reorder .to (inj₁ c , k)
```

Furthermore, we can freely coerce the operations of a computation from one effect row to a different effect row:[26,27]

[22] For more details on how instance argument resolution works, see the Agda documentation: https://agda.readthedocs.io/en/v2.6.2.2/language/instance-arguments.html.

[23] Here ∀ {X} is implicit universal quantification over an X : Set: https://agda.readthedocs.io/en/v2.6.2.2/language/implicit-arguments.html

[24] ↔ is the type of an *isomorphism* on Set from the Agda Standard Library. It is given by a record with two fields: the to field represents the → direction of the isomorphism, and from field represents the ← direction of the isomorphism.

[25] The dot notation w .reorder projects the reorder field of the record w.

[26] The notation ∀[_] is from the Agda Standard library, and is defined as follows: ∀[P] = ∀ x → P x.

[27] We can think of the hmap-free function as a "higher-order" map for Free: given a natural transformation between (the extension of) signatures, we can can transform the signature of a computation. This amounts to the observation that Free is a functor over the category of containers and container morphisms; assuming hmap-free preserves naturality.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
hmap-free : ∀[ ⟦ Δ₁ ⟧ ⇒ ⟦ Δ₂ ⟧ ] → ∀[ Free Δ₁ ⇒ Free Δ₂ ]
hmap-free θ (pure x)          = pure x
hmap-free θ (impure (c , k)) = impure (θ (c , hmap-free θ ∘ k))
```

Using this infrastructure, we can now implement a generic inject function which lets us define smart constructors for operations such as the out operation discussed in the previous subsection.

```
inject : ⦃ Δ₁ ≲ Δ₂ ⦄ → Free Δ₁ A → Free Δ₂ A
inject = hmap-free inj

`out : ⦃ Output ≲ Δ ⦄ → String → Free Δ ⊤
`out s = inject (impure (out s , pure))
```

2.3 Fold and monadic bind for Free

Since Free Δ is a monad, we can sequence computations using *monadic bind*, which is naturally defined in terms of the fold over Free.

```
fold : (A → B) → Alg Δ B → Free Δ A → B
fold g a (pure x) = g x
fold g a (impure (op , k)) = a (op , fold g a ∘ k)

Alg : (Δ : Effect) (A : Set) → Set
Alg Δ A = ⟦ Δ ⟧ A → A
```

Besides the input computation to be folded (last parameter), the fold is parameterized by a function $A \to B$ (first parameter) which folds a pure computation, and an *algebra* Alg Δ A (second parameter) which folds an impure computation. We call the latter an algebra because it corresponds to an F-algebra (Arbib & Manes, 1975; Pierce, 1991) over the signature functor of Δ, denoted F_Δ. That is, a tuple (A, α) where A is an object called the *carrier* of the algebra, and α a morphism $F_\Delta(A) \to A$. Using fold, monadic bind for the free monad is defined as follows:

```
_>>=_ : Free Δ A → (A → Free Δ B) → Free Δ B
m >>= g = fold g impure m
```

Intuitively, $m \gg\!= g$ concatenates g to all the leaves in the computation m.

Example. The following defines a smart constructor for throw:

```
`throw : ⦃ Throw ≲ Δ ⦄ → Free Δ A
```

Using this and the definition of E>>= above, we can use **do**-notation in Agda to make the hello-throw program from Section 2.1 more readable:

```
hello-throw₁ : ⦃ Output ≲ Δ ⦄ → ⦃ Throw ≲ Δ ⦄ → Free Δ A
hello-throw₁ = do `out "Hello"; `out " world!"; `throw
```

This illustrates how we use the free monad to write effectful programs against an interface given by an effect signature. Next, we define *effect handlers*.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

2.4 Effect handlers

An effect handler implements the interface given by an effect signature, interpreting the syntactic operations associated with an effect. Like monadic bind, effect handlers can be defined as a fold over the free monad. The following type of *parameterized handlers* (Leijen, 2017, §2.2) defines how to fold, respectively, pure and impure computations:[28]

```
record ⟨ _ ! _ ⇒ _ ⇒ _ ! _ ⟩ (A : Set) (Δ : Effect) (P : Set) (B : Set) (Δ′ : Effect) : Set₁ where
  field ret : A → P → Free Δ′ B
        hdl : Alg Δ (P → Free Δ′ B)
```

A handler of type $\langle\ A\ !\ \Delta \Rightarrow P \Rightarrow B\ !\ \Delta'\ \rangle$ is parameterized in the sense that it turns a computation of type Free $\Delta\ A$ into a parameterized computation of type $P \to$ Free $\Delta'\ B$. The following function does so by folding using ret, hdl, and a to-front function:[29]

```
to-front : ⦃ Δ₁ • Δ₂ ≈ Δ ⦄ → Free Δ A → Free (Δ₁ ⊕ Δ₂) A
to-front ⦃ w ⦄ = hmap-free (w .reorder .from)

given_handle_ : ⦃ w : Δ₁ • Δ₂ ≈ Δ ⦄
                → ⟨ A ! Δ₁ ⇒ P ⇒ B ! Δ₂ ⟩ → Free Δ A → (P → Free Δ₂ B)
given h handle m = fold
  (ret h)
  ( λ where (inj₁ c , k) p → hdl h (c , k) p
            (inj₂ c , k) p → impure (c , flip k p) )
  (to-front m)
```

Comparing with the syntax, we used to explain algebraic effects and handlers in the introduction, the ret field corresponds to the **return** case of the handlers from the introduction, and hdl corresponds to the cases that define how operations are handled. The parameterized handler type $\langle\ A\ !\ \Delta \Rightarrow P \Rightarrow B\ !\ \Delta'\ \rangle$ corresponds to the type $A\ !\ \Delta, \Delta' \Rightarrow P \to B\ !\ \Delta'$, and given h handle m corresponds to **with** h **handle** m.

Using this type of handler, the *hOut* handler from the introduction can be defined as follows:

```
hOut : ⟨ A ! Output ⇒ ⊤ ⇒ (A × String) ! Δ ⟩
ret hOut x _ = pure (x , "")
hdl hOut (out s , k) p = do (x , s′) ← k tt p; pure (x , s ++ s′)
```

The handler *hOut* in Section 1.1 did not bind any parameters. However, since we are encoding it as a *parameterized* handler, hOut now binds a unit-typed parameter. Besides this difference, the handler is the same as in Section 1.1. We can use the hOut handler to run computations. To this end, we introduce a Nil effect with no associated operations which we will use to indicate where an effect row ends:

[28] A simpler type of handler could omit the parameter; i.e., $\langle\ A\ !\ \Delta \Rightarrow B\ !\ \Delta'\ \rangle$, for some A,B : Set and Δ,Δ' : Effect. As demonstrated in, e.g., the work of Pretnar (2015, §2.4), this type of handler can leverage host language binding to handle, e.g., the *state effect* which we discuss later. The style of parameterized handler we use here follows the exposition of, e.g., Wu *et al.* (2014), Wu & Schrijvers (2015).

[29] The syntax λ **where** ... is a *pattern-matching* lambda in Agda. The function flip has the following type: $(A \to B \to C) \to (B \to A \to C)$.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
data StateOp : Set where
  get :       StateOp
  put : ℕ → StateOp
```

```
State : Effect
Op State = StateOp
Ret State get      = ℕ
Ret State (put n) = ⊤
```

```
hSt : ⟨ A ! State ⇒ ℕ ⇒ (A × ℕ) ! Δ′ ⟩
ret hSt x s = pure (x , s)
hdl hSt (put m , k) n = k tt m
hdl hSt (get , k) n = k n      n

`incr : ❴ State ≲ Δ ❵ → Free Δ ⊤
`incr = do n ← `get; `put (n + 1)

incr-test : un ((given hSt handle `incr) 0) ≡ (tt , 1)
incr-test = refl
```

Fig. 1. A state effect (upper), its handler (hSt below), and a simple test (incr-test, also below) which uses (the elided) smart constructors for get and put.

```
Nil : Effect
Op  Nil = ⊥
Ret Nil = ⊥-elim
```

```
un : Free Nil A → A
un (pure x) = x
```

Using these, we can run a simple hello world program:[30]

```
hello′ : ❴ Output ≲ Δ ❵ → Free Δ ⊤
hello′ = do
  `out "Hello"; `out " world!"
```

```
test-hello : un (given hOut handle hello′ $ tt)
                ≡ (tt , "Hello world!")
test-hello = refl
```

An example of parameterized (as opposed to unparameterized) handlers is the state effect. Figure 1 declares and illustrates how to handle such an effect with operations for reading (get) and changing (put) the state of a memory cell holding a natural number.

2.5 The modularity problem with higher-order effects, revisited

Section 1.2 described the modularity problem with higher-order effects, using a higher-order operation that interacts with output as an example. In this section, we revisit the problem, framing it in terms of the definitions introduced in the previous section. To this end, we use a different effect whose interface is summarized by the CatchM record below. The record asserts that the computation type $M : \mathsf{Set} \to \mathsf{Set}$ has at least a higher-order operation catch and a first-order operation throw:

```
record CatchM (M : Set → Set) : Set₁ where
  field catch : M A → M A → M A
        throw :             M A
```

[30] The refl constructor is from the Agda standard library, and witnesses that a propositional equality (≡) holds.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

The idea is that throw throws an exception, and catch m_1 m_2 handles any exception thrown during evaluation of m_1 by running m_2 instead. The problem is that we cannot give a modular definition of operations such as catch using algebraic effects and handlers alone. As discussed in Section 1.2, the crux of the problem is that algebraic effects and handlers provide limited support for higher-order operations. However, as also discussed in Section 1.2, we can encode catch in terms of more primitive effects and handlers, such as the following handler for the Throw effect:

```
hThrow : ⟨ A ! Throw ⇒ ⊤ ⇒ (Maybe A) ! Δ′ ⟩
ret hThrow x _ = pure (just x)
hdl hThrow (throw , k) _ = pure nothing
```

The handler modifies the return type of the computation by decorating it with a Maybe. If no exception is thrown, ret wraps the yielded value in a just constructor. If an exception is thrown, the handler never invokes the continuation k and aborts the computation by returning nothing instead. We can elaborate catch into an inline application of hThrow. To do so, we make use of *effect masking* which lets us "weaken" the type of a computation by inserting extra effects in an effect row:

```
♯_ : ⦃ Δ₁ ≲ Δ₂ ⦄ → Free Δ₁ A → Free Δ₂ A
```

Using this, the following elaboration defines a semantics for the catch operation:[31,32]

```
catch : ⦃ Throw ≲ Δ ⦄ → Free Δ A → Free Δ A → Free Δ A
catch m₁ m₂ = (♯ (given hThrow handle m₁) tt) >>= maybe pure m₂
```

If m_1 does not throw an exception, we return the produced value. If it does, m_2 is run.

As observed by Wu *et al*. (2014), programs that use elaborations such as catch are less modular than programs that only use plain algebraic operations. In particular, the effect row of computations no longer represents the interface of operations that we use to write programs, since the catch elaboration is not represented in the effect type at all. So we have to rely on different machinery if we want to refactor, optimize, or change the semantics of catch without having to change programs that use it.

In the next subsection, we describe how to define effectful operations such as catch modularly using scoped effects and handlers and discuss how this is not possible for, e.g., operations representing λ-abstraction.

2.6 Scoped effects and handlers

This subsection gives an overview of scoped effects and handlers. While the rest of the paper can be read and understood without a deep understanding of scoped effects and handlers, we include this overview to facilitate comparison with the alternative solution that we introduce in Section 3.

[31] The maybe function is the eliminator for the Maybe type. Its first parameter is for eliminating a just; the second for nothing. Its type is maybe : $(A \to B) \to B \to \text{Maybe}\, A \to B$.

[32] The instance resolution machinery of Agda requires some help to resolve the instance argument of ♯ here. We provide a hint to Agda's instance resolution machinery in an implicit instance argument that we omit for readability in the paper. In the rest of this paper, we will occasionally follow the same convention.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

Scoped effects extend the expressiveness of algebraic effects to support a class of higher-order operations that Wu *et al*. (2014), Piróg *et al*. (2018), Yang *et al*. (2022) call *scoped operations*. We illustrate how scoped effects work, using a freer monad encoding of the endofunctor algebra approach of Yang *et al*. (2022). The work of Yang *et al*. (2022) does not include examples of modular handlers, but the original paper on scoped effects and handlers by Wu *et al*. (2014) does. We describe an adaptation of the modular handling techniques due to Wu *et al*. (2014) to the endofunctor algebra approach of Yang *et al*. (2022).

2.6.1 Scoped programs

Scoped effects extend the free monad data type with an additional row for scoped operations. The return and call constructors of Prog below correspond to the pure and impure constructors of the free monad, whereas enter is new:

```
data Prog (Δ γ : Effect) (A : Set) : Set where
  return : A                                  → Prog Δ γ A
  call   : ⟦ Δ ⟧ (Prog Δ γ A)                 → Prog Δ γ A
  enter  : ⟦ γ ⟧ (Prog Δ γ (Prog Δ γ A))      → Prog Δ γ A
```

Here, the enter constructor represents a higher-order operation with *sub-scopes*; i.e., computations that themselves return computations:

$$\underbrace{\mathsf{Prog}\ \Delta\ \gamma}_{\text{outer}}\ (\underbrace{\mathsf{Prog}\ \Delta\ \gamma}_{\text{inner}}\ A)$$

This type represents *scoped* computations in the sense that outer computations can be handled independently of inner ones, as we illustrate in Section 2.6.2. One way to think of inner computations is as continuations (or join-points) of sub-scopes.

Using Prog, the catch operation can be defined as a scoped operation:

```
data CatchOp : Set where
  catch : CatchOp
```

```
Catch : Effect
Op  Catch = CatchOp
Ret Catch catch = Bool
```

The effect signature indicates that Catch has two scopes since Bool has two inhabitants. Following Yang *et al*. (2022), scoped operations are handled using a structure-preserving fold over Prog:

```
hcata :  (∀ {X} → X → G X)
      → CallAlg  Δ G
      → EnterAlg γ G
      → Prog Δ γ A → G A
```

```
CallAlg : (Δ : Effect) (G : Set → Set) → Set₁
CallAlg Δ G =
  {A : Set} → ⟦ Δ ⟧ (G A) → G A

EnterAlg : (γ : Effect) (G : Set → Set) → Set₁
EnterAlg γ G =
  {A B : Set} → ⟦ γ ⟧ (G (G A)) → G A
```

The first argument represents the case where we are folding a return node; the second and third correspond to, respectively, call and enter.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

2.6.2 Scoped effect handlers

The following defines a type of parameterized scoped effect handlers:

```
record ⟨•!_!_⇒_⇒_•!_!_⟩ (Δ γ : Effect) (P : Set) (G : Set → Set)
                        (Δ′ γ′ : Effect) : Set₁ where
  field ret   : X → P → Prog Δ′ γ′ (G X)
        hcall : CallAlg   Δ (λ X → P → Prog Δ′ γ′ (G X))
        henter : EnterAlg γ (λ X → P → Prog Δ′ γ′ (G X))
        glue  : (k : C → P → Prog Δ′ γ′ (G X)) (r : G C) → P → Prog Δ′ γ′ (G X)
```

A handler of type $\langle \bullet\ !\ \Delta\ !\ \gamma \Rightarrow P \Rightarrow G\ \bullet !\ \Delta'\ !\ \gamma\ \rangle$ handles operations of Δ and γ *simultaneously* and turns a computation Prog $\Delta\ \gamma\ A$ into a parameterized computation of type $P \to$ Prog $\Delta'\ \gamma'\ (G\ A)$. The ret and hcall cases are similar to the ret and hdl cases from Section 2.4. The crucial addition that adds support for higher-order operations is the henter case.

The henter field is given by an EnterAlg case. This case takes as input a scoped operation whose outer and inner computation have been folded into a parameterized computation of type $P \to$ Prog $\Delta'\ \gamma'\ (G\ X)$; and returns as output an interpretation of that operation as a computation of type $P \to$ Prog $\Delta'\ \gamma'\ (G\ X)$. The glue function is used for modularly *weaving* (Wu *et al.*, 2014) side effects of handlers through sub-scopes of yet-unhandled operations.

2.6.3 Weaving

To see why glue is needed, it is instructional to look at how the fields in the record type above are used to fold over Prog:

```
given_handle-scoped_ :  ⦃ w₁ : Δ₁ • Δ₂ ≈ Δ ⦄ ⦃ w₂ : γ₁ • γ₂ ≈ γ ⦄
                        → ⟨•! Δ₁ ! γ₁ ⇒ P ⇒ G •! Δ₂ ! γ₂ ⟩
                        → Prog Δ γ A → P → Prog Δ₂ γ₂ (G A)
given h handle-scoped m = hcata (ret h)
  ⊕[ hcall h
   , (λ where (c , k) p → call (c , flip k p)) ]
  ⊕[ (λ {A} → henter h {A})
   , (λ where (c , k) p → enter (c , λ x → map-prog (λ y → glue h id y p) (k x p))) ]′
   (to-frontΔ (to-frontγ m))
```

The second to last line above shows how glue is used. Because hcata eagerly folds the current handler over scopes (*sc*), there is a mismatch between the type that the continuation expects (B) and the type that the scoped computation returns ($G\ B$). The glue function fixes this mismatch for the particular return type modification G : Set → Set of a parameterized scoped effect handler.

The scoped effect handler for exception catching is thus:

```
hCatch : ⟨•! Throw ! Catch ⇒ ⊤ ⇒ Maybe •! Δ ! γ ⟩
ret      hCatch x _ = return (just x)
```

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
hcall   hCatch (throw , k) _ = return nothing
henter  hCatch (catch , k) _ = let m₁ = k true
                                   m₂ = k false in
  m₁ tt >>= λ where
    (just f)  → f tt
    nothing → m₂ tt >>= maybe (_$ tt) (return nothing)
glue hCatch k x _ = maybe (flip k tt) (return nothing) x
```

The henter field for the catch operation first runs m_1. If no exception is thrown, the value produced by m_1 is forwarded to k. Otherwise, m_2 is run and its value is forwarded to k, or its exception is propagated. The glue field of hCatch says that, if an unhandled exception is thrown during evaluation of a scope, the continuation is discarded and the exception is propagated; and if no exception is thrown the continuation proceeds normally.

2.6.4 Discussion and limitations

As observed by van den Berg *et al*. (2021), some higher-order effects do not correspond to scoped operations. In particular, the LambdaM record shown below is not a scoped operation:

```
record LambdaM (V : Set) (M : Set → Set) : Set₁ where
  field lam : (V → M V) → M V
        app : V → M V   → M V
```

The lam field represents an operation that constructs a λ value. The app field represents an operation that will apply the function value in the first parameter position to the argument computation in the second parameter position. The app operation has a computation as its second parameter so that it remains compatible with different evaluation strategies.

To see why the operations summarized by the LambdaM record above are not scoped operations, let us revisit the enter constructor of Prog:

$$\mathsf{enter} : [\![\ \gamma\ ]\!]\ (\underbrace{\mathsf{Prog}\ \Delta\ \gamma}_{\text{outer}}\ (\underbrace{\mathsf{Prog}\ \Delta\ \gamma}_{\text{inner}}\ A)) \to \mathit{adProg}\ \Delta\ \gamma\ A$$

As summarized earlier in this subsection, enter lets us represent higher-order operations (specifically, *scoped operations*), whereas call does not (only *algebraic operations*). Just like we defined the computational parameters as scopes (given by the outer Prog in the type of enter), we might try to define the body of a lambda as a scope in a similar way. However, whereas the catch operation always passes control to its continuation (the inner Prog), the lam effect is supposed to package the body of the lambda into a value and pass this value to the continuation (the inner computation). Because the inner computation is nested within the outer computation, *the only way to gain access to the inner computation (the continuation) is by first running the outer computation (the body of the lambda)*. This does not give us the right semantics.

It is possible to elaborate the LambdaM operations into more primitive effects and handlers, but as discussed in Sections 1.2 and 2.5, such elaborations are not modular. In the next section, we show how to make such elaborations modular.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

3 Hefty trees and algebras

As observed in Section 2.5, operations such as catch can be elaborated into more primitive effects and handlers. However, these elaborations are not modular. We solve this problem by factoring elaborations into interfaces of their own to make them modular.

To this end, we first introduce a new type of abstract syntax trees (Sections 3.1–3.3) representing computations with higher-order operations, which we dub *hefty trees* (an acronymic pun on *h*igher-order *ef*fec*t*s). We then define elaborations as algebras (*hefty algebras*; Section 3.4) over these trees. The following pipeline summarizes the idea, where H is a *higher-order effect signature*:

$$\mathsf{Hefty}\ H\ A \xrightarrow{elaborate} \mathsf{Free}\ \Delta\ A \xrightarrow{handle} Result$$

For the categorically inclined reader, Hefty conceptually corresponds to the initial algebra of the functor $HeftyF\ H\ A\ R = A + H\ R\ (R\ A)$ where $H : (\mathsf{Set} \to \mathsf{Set}) \to (\mathsf{Set} \to \mathsf{Set})$ defines the signature of higher-order operations and is a *higher-order functor*, meaning we have both the usual functorial $map : (X \to Y) \to H\ F\ X \to H\ F\ Y$ for any functor F as well as a function $hmap : \mathrm{Nat}(F, G) \to \mathrm{Nat}(H\ F, H\ G)$ which lifts natural transformations between any F and G to a natural transformation between $H\ F$ and $H\ G$. A hefty algebra is then an F-algebra over a higher-order signature functor H. The notion of elaboration that we introduce in Section 3.4 is an F-algebra whose carrier is a "first-order" effect tree ($\mathsf{Free}\ \Delta$).

In this section, we encode this conceptual framework in Agda using a container-inspired approach to represent higher-order signature functors H as a strictly positive type. We discuss and compare our approach with previous work in Section 3.5.

3.1 Generalizing Free to support higher-order operations

As summarized in Section 2.1, $\mathsf{Free}\ \Delta\ A$ is the type of abstract syntax trees representing computations over the effect signature Δ. Our objective is to arrive at a more general type of abstract syntax trees representing computations involving (possibly) higher-order operations. To realize this objective, let us consider how to syntactically represent this variant of the *censor* operation (Section 1.2), where M is the type of abstract syntax trees whose type we wish to define:

$$\mathsf{censor}_{op} : (\mathsf{String} \to \mathsf{String}) \to M\ \top \to M\ \top$$

We call the second parameter of this operation a *computation parameter*. Using Free, computation parameters can only be encoded as continuations. But the computation parameter of censor_{op} is *not* a continuation, since

$$\mathbf{do}\ (\mathsf{censor}_{op}\ f\ m); \text{`}\mathsf{out}\ s \not\equiv \mathsf{censor}_{op}\ f\ (\mathbf{do}\ m; \text{`}\mathsf{out}\ s).$$

The crux of the issue is how signature functors $[\![\ \Delta\ ]\!] : \mathsf{Set} \to \mathsf{Set}$ are defined. Since this is an endofunctor on Set, the only suitable option in the impure constructor is to apply the functor to the type of *continuations*:

$$\mathsf{impure} : [\![\ \Delta\ ]\!]\ (\underbrace{\mathsf{Free}\ \Delta\ A}_{\text{continuation}}) \to \mathsf{Free}\ \Delta\ A$$

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

A more flexible approach would be to allow signature functors to build computation trees with an *arbitrary return type*, including the return type of the continuation. A *higher-order* signature functor of some higher-order signature H, written $[\![H]\!]^{\mathsf{H}} : (\mathsf{Set} \to \mathsf{Set}) \to \mathsf{Set} \to \mathsf{Set}$, would fit that bill. Using such a signature functor, we could define the impure case as follows:

$$\mathsf{impure} : [\![H]\!]^{\mathsf{H}} \; (\underbrace{\mathsf{Hefty}\; H}_{\substack{\text{computation}\\ \text{type}}}) \; \overbrace{A}^{\substack{\text{continuation}\\ \text{return type}}} \to \mathsf{Hefty}\; H\; A$$

Here, Hefty is the type of the free monad using higher-order signature functors instead. In the rest of this subsection, we consider how to define higher-order signature functors H, their higher-order functor extensions $[\![_]\!]^{\mathsf{H}}$, and the type of Hefty trees.

Recall how we defined plain algebraic effects in terms of *containers*:

```
record Effect : Set₁ where
  field Op  : Set
        Ret : Op → Set
```

Here, Op is the type of operations, and Ret defines the return type of each operation. In order to allow operations to have sub-computations, we generalize this type to allow each operation to be associated with a number of sub-computations, where each sub-computation can have a different return type. The following record provides this generalization:

```
record Effect^H : Set₁ where
  field Op^H  : Set                                  – As before
        Ret^H : Op^H → Set                           – As before
        Fork  : Op^H → Set                           – New
        Ty    : {op : Op^H} (φ : Fork op) → Set      – New
```

The set of operations is still given by a type field (Op^{H}), and each operation still has a return type ($\mathsf{Ret}^{\mathsf{H}}$). Fork associates each operation with a type that indicates how many sub-computations (or *forks*) an operation has, and Ty indicates the return type of each such fork. For example, say we want to encode an operation *op* with two sub-computations with different return types, and whose return type is of a unit type. That is, using M as our type of computations:

$$op : M\,\mathbb{Z} \to M\,\mathbb{N} \to M\,\top$$

The following signature declares a higher-order effect signature with a single operation with return type $\top$, and with two forks (we use Bool to encode this fact), and where each fork has, respectively, $\mathbb{Z}$ and $\mathbb{N}$ as return types.

```
example-op : Effect^H
Op^H  example-op    = ⊤   – A single operation
Ret^H example-op tt = ⊤   – with return type ⊤
```

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
data CensorOp : Set where
  censor :  (String → String)
         → CensorOp
```

```
Censor : Effect^H
Op^H  Censor          = CensorOp
Ret^H Censor (censor f) = ⊤
Fork  Censor (censor f) = ⊤
Ty    Censor {censor f} tt = ⊤
```

```
censor_op : (String → String) → Hefty (Censor + H) ⊤ → Hefty (Censor + H) ⊤
censor_op f m = impure (inj₁ (censor f) , (λ where tt → m) , pure)
```

Fig. 2. A higher-order censor effect and operation, with a single computation parameter (declared with $\mathsf{Op} = \top$ in the effect signature top right) with return type $\top$ (declared with $\mathsf{Ret} = \lambda _ \to \top$ top right).

```
Fork example-op tt    = Bool – with two forks
Ty   example-op false = ℤ    – one fork has return type ℤ
Ty   example-op true  = ℕ    – the other has return type ℕ
```

The extension of higher-order effect signatures implements the intuition explained above:

```
⟦_⟧^H : Effect^H → (Set → Set) → Set → Set
⟦ H ⟧^H M X =
  Σ (Op^H H) λ op → (Ret^H H op → M X) × ((ϕ : Fork H op) → M (Ty H ϕ))
```

Let us unpack this definition.

$$\underbrace{\Sigma\ (\mathsf{Op}^{\mathsf{H}}\ H)\ \lambda\ op \to}_{(1)}\ \underbrace{(\mathsf{Ret}^{\mathsf{H}}\ H\ op \to M\ X)}_{(2)} \times \underbrace{(\underbrace{(\phi : \mathsf{Fork}\ H\ op)}_{(3)} \to \underbrace{M\ (\mathsf{Ty}\ H\ \phi)}_{(4)})}$$

The extension of a higher-order signature functor is given by (1) the sum of operations of the signature, where each operation has (2) a continuation (of type $M\ X$) that expects to be passed a value of the operation's return type, and (3) a set of forks where each fork is (4) a computation that returns the expected type for each fork.

Using the higher-order signature functor and its extension above, our generalized free monad becomes:

```
data Hefty (H : Effect^H) (A : Set) : Set where
  pure   : A → Hefty H A
  impure : ⟦ H ⟧^H (Hefty H) A → Hefty H A
```

This type of Hefty trees can be used to define higher-order operations with an arbitrary number of computation parameters, with arbitrary return types. Using this type, and using a co-product for higher-order effect signatures ($_ + _$) which is analogous to the co-product for algebraic effect signatures in Section 2.2, Figure 2 represents the syntax of the censor_{op} operation.

Just like Free, Hefty trees can be sequenced using monadic bind. Unlike for Free, the monadic bind of Hefty is not expressible in terms of the standard fold over Hefty trees. The difference between Free and Hefty is that Free is a regular data type, whereas Hefty is a *nested datatype* (Bird & Paterson, 1999). The fold of a nested data type is limited to

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

describe *natural transformations*. As Bird & Paterson (1999) show, this limitation can be overcome by using a *generalized fold*, but for the purpose of this paper, it suffices to define monadic bind as a recursive function:

$$\begin{aligned}
&_ \gg\!= _ : \mathsf{Hefty}\ H\ A \to (A \to \mathsf{Hefty}\ H\ B) \to \mathsf{Hefty}\ H\ B \\
&\mathsf{pure}\ x \qquad\qquad\ \gg\!= g = g\ x \\
&\mathsf{impure}\ (op\ ,\ k\ ,\ \psi) \gg\!= g = \mathsf{impure}\ (op\ ,\ (_ \gg\!= g) \circ k\ ,\ \psi)
\end{aligned}$$

The bind behaves similarly to the bind for Free; i.e., $m \gg\!= g$ concatenates g to all the leaves in the continuations (but not computation parameters) of m.

In Section 3.4, we show how to modularly elaborate higher-order operations into more primitive algebraic effects and handlers (i.e., computations over Free), by folding modular elaboration algebras (*hefty algebras*) over Hefty trees. First, we show (in Section 3.2) how Hefty trees support programming against an interface of both algebraic and higher-order operations. We also address (in Section 3.3) the question of how to encode effect signatures for higher-order operations whose computation parameters have polymorphic return types, such as the highlighted A below:

$$\text{`}\mathsf{catch} : \mathsf{Hefty}\ H\ A \to \mathsf{Hefty}\ H\ A \to \mathsf{Hefty}\ H\ A$$

3.2 Programs with algebraic and higher-order effects

Any algebraic effect signature can be lifted to a higher-order effect signature with no fork (i.e., no computation parameters):

$$\begin{aligned}
&\mathsf{Lift} : \mathsf{Effect} \to \mathsf{Effect}^{\mathsf{H}} \\
&\mathsf{Op}^{\mathsf{H}}\ (\mathsf{Lift}\ \Delta) = \mathsf{Op}\ \Delta \\
&\mathsf{Ret}^{\mathsf{H}}\ (\mathsf{Lift}\ \Delta) = \mathsf{Ret}\ \Delta \\
&\mathsf{Fork}\ \ (\mathsf{Lift}\ \Delta) = \lambda\ _ \to \bot \\
&\mathsf{Ty}\quad\ (\mathsf{Lift}\ \Delta) = \lambda()
\end{aligned}$$

Using this effect signature and using higher-order effect row insertion witnesses analogous to the ones we defined and used in Section 2.2, the following smart constructor lets us represent any algebraic operation as a Hefty computation:

$$\uparrow _ : \{\!|\ w : \mathsf{Lift}\ \Delta \lesssim^{\mathsf{H}} H\ |\!\} \to (op : \mathsf{Op}\ \Delta) \to \mathsf{Hefty}\ H\ (\mathsf{Ret}\ \Delta\ op)$$

Using this notion of lifting, Hefty trees can be used to program against interfaces of both higher-order and plain algebraic effects.

3.3 Higher-order operations with polymorphic return types

Let us consider how to define Catch as a higher-order effect. Ideally, we would define an operation that is parameterized by a return type of the branches of a particular catch operation, as shown on the left, such that we can define the higher-order effect signature on the right:[33]

[33] *d* is for *dubious*.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
data CatchOp^d : Set₁ where
  catch^d : Set → CatchOp^d
```

```
Catch^d : Effect^H
Op^H  Catch^d = CatchOp^d
Ret^H Catch^d (catch^d A) = A
Fork  Catch^d (catch^d A) = Bool
Ty    Catch^d {catch^d A} _ = A
```

The Fork field on the right says that Catch has two sub-computations (since Bool has two constructors), and that each computation parameter has some return type A. However, the signature on the right above is not well defined!

The problem is that, because $\mathsf{CatchOp}^d$ has a constructor that quantifies over a type (Set), the $\mathsf{CatchOp}^d$ type lives in Set_1. Consequently, it does not fit the definition of Effect^H, whose operations live in Set. There are two potential solutions to this problem: (1) increase the universe level of Effect^H to allow Op^H to live in Set_1 or (2) use a *universe of types* (Martin-Löf, 1984). Either solution is applicable here; we choose type universes.

A universe of types is a (dependent) pair of a syntax of types ($\mathsf{Ty} : \mathsf{Set}$) and a semantic function ($[\![_]\!]^T : \mathsf{Ty} \to \mathsf{Set}$) defining the meaning of the syntax by reflecting it into Agda's Set:

```
record Univ : Set₁ where
  field Type : Set
        [[_]]^T : Type → Set
```

Section 4.1 contains a concrete example usage this notion of type universe. Using type universes, we can parameterize the catch constructor on the left below by a *syntactic type* Ty of some universe u and use the *meaning of this type* ($[\![\ t\ ]\!]^T$) as the return type of the computation parameters in the effect signature on the right below:

```
data CatchOp {{ u : Univ }} : Set where
  catch : Type → CatchOp
```

```
Catch : {{ u : Univ }} → Effect^H
Op^H  Catch           = CatchOp
Ret^H Catch (catch t) = [[ t ]]^T
Fork  Catch (catch t) = Bool
Ty    Catch {catch t} = λ _ → [[ t ]]^T
```

While the universe of types encoding restricts the kind of type that catch can have as a return type, the effect signature is parametric in the universe. Thus the implementer of the Catch effect signature (or interface) is free to choose a sufficiently expressive universe of types.

3.4 Hefty algebras

As shown in Section 2.5, the higher-order catch operation can be encoded as a non-modular elaboration:

```
catch m₁ m₂ = (♯ ((given hThrow handle m₁) tt)) >>= (maybe pure m₂)
```

We can make this elaboration modular by expressing it as an *algebra* over Hefty trees containing operations of the Catch signature. To this end, we will use the following notion of hefty algebra (Alg^H) and fold (or *catamorphism* Meijer *et al.*, 1991, cata^H) for Hefty:

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
record Alg^H (H : Effect^H) (F : Set → Set) : Set₁ where
  field alg : ⟦ H ⟧^H F A → F A

cata^H : (∀ {A} → A → F A) → Alg^H H F → Hefty H A → F A
cata^H g a (pure x)                = g x
cata^H g a (impure (op , k , ψ)) = alg a (op , ((cata^H g a ∘ k) , (cata^H g a ∘ ψ)))
```

Here, $\mathsf{Alg}^{\mathsf{H}}$ defines how to transform an impure node of type Hefty $H\ A$ into a value of type $F\ A$, assuming we have already folded the computation parameters and continuation into F values. A nice property of algebras is that they are closed under higher-order effect signature sums:

```
_⋎_ : Alg^H H₁ F → Alg^H H₂ F → Alg^H (H₁ ∔ H₂) F
alg (A₁ ⋎ A₂) (inj₁ op , k , ψ) = alg A₁ (op , k , ψ)
alg (A₁ ⋎ A₂) (inj₂ op , k , ψ) = alg A₂ (op , k , ψ)
```

By defining elaborations as hefty algebras (below) we can compose them using $_ \curlyvee _$.

```
Elaboration : Effect^H → Effect → Set₁
Elaboration H Δ = Alg^H H (Free Δ)
```

An Elaboration $H\ \Delta$ elaborates higher-order operations of signature H into algebraic operations of signature Δ. Given an elaboration, we can generically transform any hefty tree into more primitive algebraic effects and handlers:

```
elaborate : Elaboration H Δ → Hefty H A → Free Δ A
elaborate = cata^H pure
```

Example 1 (Elaboration for Output Censoring)**.** Let us return to the example from the introduction. Here is the elaboration of the Censor effect from Figure 2.

```
eCensor : ⦃ w : Output ≲ Δ ⦄ → Elaboration Censor Δ
alg eCensor (censor f , k , ψ) = do
  (x , s) ← ♯ ((given hOut handle ψ tt) tt)
  `out (f s)
  k x
```

This elaboration matches the *eCensor* elaboration discussed in Section 1.3.

Example 2 (Elaboration for Exception Catching)**.** We can also elaborate exception catching analogously to the non-modular catch elaboration discussed in Section 2.5 and in the beginning of this subsection:

```
eCatch : ⦃ u : Univ ⦄ ⦃ w : Throw ≲ Δ ⦄ → Elaboration Catch Δ
alg eCatch (catch t , k , ψ) =
  (♯ ((given hThrow handle ψ true) tt)) >>= maybe k (ψ false >>= k)
```

The elaboration is essentially the same as its non-modular counterpart, except that it now uses the universe of types encoding discussed in Section 3.3, and that it now transforms syntactic representations of higher-order operations instead. Using this elaboration, we can,

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

for example, run the following example program involving the state effect from Figure 1, the throw effect from Section 2.1, and the catch effect defined here:

```
transact : ⦃ w_s : Lift State ≲H H ⦄ ⦃ w_t : Lift Throw ≲H H ⦄ ⦃ w : Catch ≲H H ⦄
         → Hefty H ℕ
transact = do
  ↑ put 1
  `catch (do ↑ (put 2); (↑ throw) >>= ⊥-elim) (pure tt)
  ↑ get
```

The program first sets the state to 1; then to 2; and then throws an exception. The exception is caught, and control is immediately passed to the final operation in the program which gets the state. By also defining elaborations for Lift and Nil, we can elaborate and run the program:

```
eTransact : ⦃ _ : Throw ≲ Δ ⦄ ⦃ _ : State ≲ Δ ⦄
          → Elaboration (Catch + Lift Throw + Lift State + Lift Nil) Δ
eTransact = eCatch ϒ eLift ϒ eLift ϒ eNil


test-transact : un ( ( given hSt
                         handle ( ( given hThrow
                                      handle (elaborate eTransact transact) )
                                  tt ) )
                    0 ) ≡ (just 2 , 2)
test-transact = refl
```

The program in Example 2 uses a so-called *global* interpretation of state, where the put operation in the "try block" of `catch causes the state to be updated globally. In Section 4.2.2, we return to this example and show how we can modularly change the elaboration of the higher-order effect Catch to yield a so-called *transactional* interpretation of state where the put operation in the try block is rolled back when an exception is thrown.

3.5 Discussion and limitations

Which (higher-order) effects can we describe using hefty trees and algebras? Since the core mechanism of our approach is modular elaboration of higher-order operations into more primitive effects and handlers, it is clear that hefty trees and algebras are at least as expressive as standard algebraic effects. The crucial benefit of hefty algebras over algebraic effects is that higher-order operations can be declared and implemented modularly. In this sense, hefty algebras provide a modular abstraction layer over standard algebraic effects that, although it adds an extra layer of indirection by requiring both elaborations and handlers to give a semantics to hefty trees, is comparatively cheap and implemented using only standard techniques such as F-algebras. As we show in Section 5, hefty algebras also let us define higher-order effect theories, akin to algebraic effect theories.

Conceptually, we expect that hefty trees can capture any *monadic* higher-order effect whose signature is given by a higher-order functor on Set → Set. Filinski (1999) showed

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

that any monadic effect can be represented using continuations, and given that we can encode the continuation monad using algebraic effects (Schrijvers *et al*., 2019) in terms of the *sub/jump* operations (Section 4.2.2) by Thielecke (1997), Fiore & Staton (2014), it is possible to elaborate any monadic effect into algebraic effects using hefty algebras. The current Agda implementation, though, is slightly more restrictive. The type of effect signatures, $\mathsf{Effect}^{\mathsf{H}}$, approximates the set of higher-order functors by constructively enforcing that all occurrences of the computation type are strictly positive. Hence, there may be higher-order effects that are well-defined semantically, but which cannot be captured in the Agda encoding presented here.

Recent work by van den Berg & Schrijvers (2023) introduced a higher-order free monad that coincides with our Hefty type. Their work shows that hefty trees are, in fact, a free monad. Furthermore, they demonstrate that a range of existing effects frameworks from the literature can be viewed as instances of hefty trees.

When comparing hefty trees to scoped effects, we observe two important differences. The first difference is that the syntax of programs with higher-order effects is fundamentally more restrictive when using scoped effects. Specifically, as discussed at the end of Section 2.6.4, scoped effects impose a restriction on operations that their computation parameters must pass control directly to the continuation of the operation. Hefty trees, on the other hand, do not restrict the control flow of computation parameters, meaning that they can be used to define a broader class of operations. For instance, in Section 4.1, we define a higher-order effect for function abstraction, which is an example of an operation where control does not flow from the computation parameter to the continuation.

The second difference is that hefty algebras and scoped effects and handlers are modular in different ways. Scoped effects are modular because we can modularly define, compose, and handle scoped operations, by applying scoped effect handlers in sequence; i.e.:

$$\mathsf{Prog}\ \Delta_0\ \gamma_0\ A_0 \xrightarrow{h'_1} \mathsf{Prog}\ \Delta_1\ \gamma_1\ A_1 \xrightarrow{h'_2} \cdots \xrightarrow{h'_n} \mathsf{Prog}\ \mathsf{Nil}\ \mathsf{Nil}\ A_n \qquad (\ddagger)$$

As discussed in Section 2.6.3, each handler application modularly "weaves" effects through sub-computations, using a dedicated glue function.applying different scoped effect handlers in different orders.

Hefty algebras, on the other hand, work by applying an elaboration algebra assembled from modular components in one go. The program resulting from elaboration can then be handled using standard algebraic effect handlers; i.e.:

$$\mathsf{Hefty}\ (H_0 \dotplus \cdots \dotplus H_m)\ A \xrightarrow{\mathsf{elaborate}\ (E_0\ \curlyvee \cdots \curlyvee\ E_m)} \mathsf{Free}\ \Delta\ A \xrightarrow{h_1} \cdots \xrightarrow{h_k} \mathsf{Free}\ \mathsf{Nil}\ A_k \qquad (\S)$$

The algebraic effect handlers $h_1, \ldots, h_k$ in (§) serve the same purpose as the scoped effect handlers $h'_1, \ldots, h'_n$ in (†); namely, to provide a semantics of operations. The order of handling is significant for both algebraic effect handlers and for scoped effect handlers: applying the same handlers in different orders may give a different semantics.

In contrast, the order that elaborations $(E_1, \ldots, E_m)$ are composed in (§) does not matter. Hefty algebras merely mediate higher-order operations into "first-order" effect trees that then must be handled, using standard effect handlers. While scoped effects supports "weaving", standard algebraic effect handlers do not. This might suggest that scoped effects and handlers are generally more expressive. However, many scoped effects and handlers can

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

be emulated using algebraic effects and hanlders, by encoding scoped operations as algebraic operations whose continuations encode a kind of scoped syntax, inspired by Wu *et al.* (2014, §7-9).[34] We illustrate how in Section 4.2.2.

4 Examples

As discussed in Section 2.5, there is a wide range of examples of higher-order effects that cannot be defined as algebraic operations directly and are typically defined as non-modular elaborations instead. In this section, we give examples of such effects and show to define them modularly using hefty algebras. The artifact (van der Rest & Poulsen, 2024) contains the full examples.

4.1 λ as a higher-order operation

As recently observed by van den Berg *et al.* (2021), the (higher-order) operations for λ abstraction and application are neither algebraic nor scoped effects. We demonstrate how hefty algebras allow us to modularly define and elaborate an interface of higher-order operations for λ abstraction and application, inspired by Levy's call-by-push-value (Levy, 2006). The interface we will consider is parametric in a universe of types given by the following record:

```
record LamUniv : Set₁ where
  field ⦃ u ⦄   : Univ
        _↣_ : Type → Type → Type
        c     : Type → Type
```

The `_↣_` field represents a function type, whereas `c` is the type of *thunk values*. Distinguishing thunks in this way allows us to assign either a call-by-value or call-by-name semantics to the interface for λ abstraction, given by the higher-order effect signature in Figure 3 and summarized by the following smart constructors:

```
`lam : {t₁ t₂ : Type} → (⟦ c t₁ ⟧ᵀ → Hefty H ⟦ t₂ ⟧ᵀ)      → Hefty H ⟦ (c t₁) ↣ t₂ ⟧ᵀ
`var : {t : Type}     → ⟦ c t ⟧ᵀ                          → Hefty H ⟦ t ⟧ᵀ
`app : {t₁ t₂ : Type} → ⟦ (c t₁) ↣ t₂ ⟧ᵀ → Hefty H ⟦ t₁ ⟧ᵀ → Hefty H ⟦ t₂ ⟧ᵀ
```

Here, `` `lam `` is a higher-order operation with a function typed computation parameter and whose return type is a function value (⟦ c t_1 ↣ t_2 ⟧ᵀ). The `` `var `` operation accepts a thunk value as argument and yields a value of a matching type. The `` `app `` operation is also a higher-order operation: its first parameter is a function value type, whereas its second parameter is a computation parameter whose return type matches that of the function value parameter type.

The interface above defines a kind of *higher-order abstract syntax* (Pfenning & Elliott, 1988), which piggy-backs on Agda functions for name binding. However, unlike most Agda functions, the constructors above represent functions with side effects. The

[34] We suspect that it is generally possible to encode scoped syntax and handlers in terms of algebraic operations and handlers, but verifying this is future work.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
data LamOp ❴ l : LamUniv ❵ : Set where
  lam : {t₁ t₂ : Type} → LamOp
  var : {t : Type}     → ⟦ c t ⟧ᵀ          → LamOp
  app : {t₁ t₂ : Type} → ⟦ (c t₁) ↣ t₂ ⟧ᵀ → LamOp

Lam : ❴ l : LamUniv ❵ → Effectᴴ
Opᴴ  Lam                      = LamOp
Retᴴ Lam (lam {t₁} {t₂})      = ⟦ (c t₁) ↣ t₂ ⟧ᵀ
Retᴴ Lam (var {t} _)          = ⟦ t ⟧ᵀ
Retᴴ Lam (app {t₁} {t₂} _)    = ⟦ t₂ ⟧ᵀ
Fork Lam (lam {t₁} {t₂})      = ⟦ c t₁ ⟧ᵀ
Fork Lam (var _)              = ⊥
Fork Lam (app {t₁} {t₂} _)    = ⊤
Ty   Lam {lam {t₁} {t₂}} _    = ⟦ t₂ ⟧ᵀ
Ty   Lam {var _} ()
Ty   Lam {app {t₁} {t₂} _} _  = ⟦ t₁ ⟧ᵀ
```

Fig. 3. Higher-order effect signature of λ abstraction and application.

representation in principle supports functions with arbitrary side effects since it is parametric in what the higher-order effect signature H is. Furthermore, we can assign different operational interpretations to the operations in the interface without having to change the interface or programs written against the interface. To illustrate we give two different implementations of the interface: one that implements a call-by-value evaluation strategy, and one that implements call-by-name.

4.1.1 Call-by-value

We give a call-by-value interpretation of `lam by generically elaborating to algebraic effect trees with any set of effects Δ. Our interpretation is parametric in proof witnesses that the following isomorphisms hold for value types (↔ is the type of isomorphisms from the Agda standard library):

```
iso₁ : {t₁ t₂ : Type} → ⟦ t₁ ↣ t₂ ⟧ᵀ ↔ (⟦ t₁ ⟧ᵀ → Free Δ ⟦ t₂ ⟧ᵀ)
iso₂ : {t : Type}     → ⟦ c t      ⟧ᵀ ↔ ⟦ t ⟧ᵀ
```

The first isomorphism says that a function value type corresponds to a function which accepts a value of type t_1 and produces a computation whose return type matches that of the function type. The second says that thunk types coincide with value types. Using these isomorphisms, the following defines a call-by-value elaboration of functions:

```
eLamCBV : Elaboration Lam Δ
alg eLamCBV (lam , k , ψ) = k (from ψ)
alg eLamCBV (var x , k , _) = k (to x)
alg eLamCBV (app f , k , ψ) = do
  a ← ψ tt
  v ← to f (from a)
  k v
```

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

The lam case passes the function body given by the sub-tree ψ as a value to the continuation, where the from function mediates the sub-tree of type $[\![\ \mathsf{c}\ t_1\ ]\!]^{\mathsf{T}} \to \mathsf{Free}\ \Delta\ [\![\ t_2\ ]\!]^{\mathsf{T}}$ to a value type $[\![\ (\mathsf{c}\ t_1) \rightarrowtail t_2\ ]\!]^{\mathsf{T}}$, using the isomorphism iso_1. The var case uses the to function to mediate a $[\![\ \mathsf{c}\ t\ ]\!]^{\mathsf{T}}$ value to a $[\![\ t\ ]\!]^{\mathsf{T}}$ value, using the isomorphism iso_2. The app case first eagerly evaluates the argument expression of the application (in the sub-tree ψ) to an argument value and then passes the resulting value to the function value of the application. The resulting value is passed to the continuation.

Using the elaboration above, we can evaluate programs such as the following which uses both the higher-order lambda effect, the algebraic state effect, and assumes that our universe has a number type $[\![\ num\ ]\!]^{\mathsf{T}} \leftrightarrow \mathbb{N}$:

```
ex : Hefty (Lam ∔ Lift State ∔ Lift Nil) ℕ
ex = do
  ↑ put 1
  f ← `lam (λ x → do
        n₁ ← `var x
        n₂ ← `var x
        pure (from ((to n₁) + (to n₂))))
  v ← `app f incr
  pure (to v)
  where incr = do s₀ ← ↑ get; ↑ put (s₀ + 1); s₁ ← ↑ get; pure (from s₁)
```

The program first sets the state to 1. Then it constructs a function that binds a variable x, dereferences the variable twice, and adds the two resulting values together. Finally, the application in the second-to-last line applies the function with an argument expression which increments the state by 1 and returns the resulting value. Running the program produces 4 since the state increment expression is eagerly evaluated before the function is applied.

```
elab-cbv : Elaboration (Lam ∔ Lift State ∔ Lift Nil) (State ⊕ Nil)
elab-cbv = eLamCBV ϒ eLift ϒ eNil

test-ex-cbv : un ((given hSt handle (elaborate elab-cbv ex)) 0) ≡ (4 , 2)
test-ex-cbv = refl
```

4.1.2 Call-by-name

The key difference between the call-by-value and the call-by-name interpretation of our λ operations is that we now assume that thunks are computations. That is, we assume that the following isomorphisms hold for value types:

```
iso₁ : {t₁ t₂ : Type} → [[ t₁ ⇀ t₂ ]]ᵀ ↔ ([[ t₁ ]]ᵀ → Free Δ [[ t₂ ]]ᵀ)
iso₂ : {t : Type}     → [[ c t ]]ᵀ      ↔ Free Δ [[ t ]]ᵀ
```

Using these isomorphisms, the following defines a call-by-name elaboration of functions:

```
eLamCBN : Elaboration Lam Δ
alg eLamCBN (lam , k , ψ) = k (from ψ)
```

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
alg eLamCBN (var x , k , _) = to x ≫= k
alg eLamCBN (app f , k , ψ) = to f (from (ψ tt)) ≫= k
```

The case for lam is the same as the call-by-value elaboration. The case for var now needs to force the thunk by running the computation and passing its result to k. The case for app passes the argument sub-tree (ψ) as an argument to the function f, runs the computation resulting from doing so, and then passes its result to k. Running the example program ex from above now produces 5 as result, since the state increment expression in the argument of `app is thunked and run twice during the evaluation of the called function.

```
elab-cbn : Elaboration (Lam ∔ Lift State ∔ Lift Nil) (State ⊕ Nil)
elab-cbn = eLamCBN ϒ eLift ϒ eNil

test-ex-cbn : un ((given hSt handle (elaborate elab-cbn ex)) 0) ≡ (5 , 3)
test-ex-cbn = refl
```

4.2 Optionally transactional exception catching

A feature of scoped effect handlers (Wu *et al.*, 2014; Piróg *et al.*, 2018; Yang *et al.*, 2022) is that changing the order of handlers makes it possible to obtain different semantics of *effect interaction*. A classical example of effect interaction is the interaction between state and exception catching that we briefly considered at the end of Section 3.4 in connection with this transact program:

```
transact : ⦃ w_s : Lift State ≲H H ⦄ ⦃ w_t : Lift Throw ≲H H ⦄ ⦃ w : Catch ≲H H ⦄
         → Hefty H ℕ
transact = do
  ↑ put 1
  `catch (do ↑ put 2; (↑ throw) ≫= ⊥-elim) (pure tt)
  ↑ get
```

The state and exception catching effect can interact to give either of these two semantics:

1. *Global* interpretation of state, where the transact program returns 2 since the put operation in the "try" block causes the state to be updated globally.
2. *Transactional* interpretation of state, where the transact program returns 1 since the state changes of the put operation are *rolled back* when the "try" block throws an exception.

With monad transformers (Cenciarelli & Moggi, 1993; Liang *et al.*, 1995), we can recover either of these semantics by permuting the order of monad transformers. With scoped effect handlers, we can also recover either by permuting the order of handlers. However, the eCatch elaboration in Section 3.4 always gives us a global interpretation of state. In this section, we demonstrate how we can recover a transactional interpretation of state by using a different elaboration of the catch operation into an algebraically effectful program with the throw operation and the off-the-shelf *sub/jump* control effects due to Thielecke (1997), Fiore & Staton (2014). The different elaboration is modular in the sense that we do not have to change the interface of the catch operation nor any programs written against the interface.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
data CCOp {| u : Univ |} (Ref : Type → Set) : Set where
  sub  : {t : Type} → CCOp Ref
  jump : {t : Type} (ref : Ref t) (x : ⟦ t ⟧ᵀ) → CCOp Ref

CC : {| u : Univ |} (Ref : Type → Set) → Effect
Op (CC Ref) = CCOp Ref
Ret (CC Ref) (sub {t})      = Ref t ⊎ ⟦ t ⟧ᵀ
Ret (CC Ref) (jump ref x) = ⊥
```

Fig. 4. Effect signature of the sub/jump effect.

4.2.1 Sub/Jump

We recall how to define two operations, sub and jump, due to Thielecke (1997), Fiore & Staton (2014). We define these operations as algebraic effects following Schrijvers *et al.* (2019). The algebraic effect signature of CC *Ref* is given in Figure 4 and is summarized by the following smart constructors:

```
'sub  : {| w : CC Ref ≲ Δ |} (b : Ref t → Free Δ A) (k : ⟦ t ⟧ᵀ → Free Δ A) → Free Δ A
'jump : {| w : CC Ref ≲ Δ |} (ref : Ref t) (x : ⟦ t ⟧ᵀ) → Free Δ B
```

An operation 'sub *f g* gives a computation *f* access to the continuation *g* via a reference value *Ref t* which represents a continuation expecting a value of type ⟦ *t* ⟧ᵀ. The 'jump operation invokes such continuations.

The operations and their handler (abbreviated to h) satisfy the following laws:

```
                h ('sub (λ _ → p) k) ≡ h p
h ('sub (λ r → m >>= 'jump r) k) ≡ h (m >>= k)
          h ('sub p ('jump r')) ≡ h (p r')
    h ('sub p q >>= k) ≡ h ('sub (λ x → p x >>= k ) (λ x → q x >>= k))
```

The last law asserts that 'sub and 'jump are *algebraic* operations, since their computational sub-terms behave as continuations. Thus, we encode 'sub and its handler as an algebraic effect.

4.2.2 Optionally transactional exception catching

By using the 'sub and 'jump operations in our elaboration of catch, we get a semantics of exception catching whose interaction with state depends on the order that the state effect and sub/jump effect is handled.

```
eCatchOT : {| w₁ : CC Ref ≲ Δ |} {| w₂ : Throw ≲ Δ |} → Elaboration Catch Δ
alg eCatchOT (catch x , k , ψ) = let m₁ = ψ true; m₂ = ψ false in
  'sub (λ r → (♯ ((given hThrow handle m₁) tt)) >>= maybe k ('jump r (from tt)))
       (λ _ → m₂ >>= k)
```

The elaboration uses 'sub to capture the continuation of a higher-order catch operation. If no exception is raised, then control flows to the continuation *k* without invoking the

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
data ChoiceOp : Set where
  or   : ChoiceOp
  fail : ChoiceOp

Choice : Effect
Op Choice = ChoiceOp
Ret Choice or = Bool
Ret Choice fail = ⊥
```

Fig. 5. Effect signature of the choice effect.

```
data OnceOp ⦃ u : Univ ⦄ : Set where
  once : {t : Type} → OnceOp

Once : ⦃ u : Univ ⦄ → Effect^H
Op^H  Once = OnceOp
Ret^H Once (once {t}) = ⟦ t ⟧^T
Fork  Once (once {t}) = ⊤
Ty    Once {once {t}} _ = ⟦ t ⟧^T
```

Fig. 6. Higher-order effect signature of the once effect.

continuation of `sub. Otherwise, we jump to the continuation of `sub, which runs m_2 before passing control to k. Capturing the continuation in this way interacts with state because the continuation of `sub may have been pre-applied to a state handler that only knows about the "old" state. This happens when we handle the state effect before the sub/jump effect: in this case, we get the transactional interpretation of state, so running transact gives 1. Otherwise, if we run the sub/jump handler before the state handler, we get the global interpretation of state and the result 2.

The sub/jump elaboration above is more involved than the scoped effect handler that we considered in Section 2.6. However, the complicated encoding does not pollute the higher-order effect interface and only turns up if we strictly want or need effect interaction.

4.3 Logic programming

Following Schrijvers *et al.* (2014), Wu *et al.* (2014), Yang *et al.* (2022), we can define a nondeterministic choice operation (_`or_) as an algebraic effect, to provide support for expressing the kind of non-deterministic search for solutions that is common in logic programming. We can also define a `fail operation that indicates that the search in the current branch was unsuccessful. The effect signature for Choice is given in Figure 5. The following smart constructors are the lifted higher-order counterparts to the `or and `fail operations:

```
_`or^H_ : ⦃ Lift Choice ≲^H H ⦄ → Hefty H A → Hefty H A → Hefty H A
`fail^H  : ⦃ Lift Choice ≲^H H ⦄                        → Hefty H A
```

A useful operator for cutting non-deterministic search short when a solution is found is the `once operator. The `once operator, whose higher-order effect signature is given in Figure 6, is not an algebraic effect, but a scoped (and thus higher-order) effect.

```
`once : ⦃ w : Once ≲^H H ⦄ {t : Type} → Hefty H ⟦ t ⟧^T → Hefty H ⟦ t ⟧^T
```

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
data ConcurOp ❴ u : Univ ❵ : Set where
  spawn  : (t : Type) → ConcurOp
  atomic : (t : Type) → ConcurOp
```

```
Concur : ❴ u : Univ ❵ → Effect^H
Op^H Concur = ConcurOp
Ret^H Concur (spawn t) = ⟦ t ⟧^T
Ret^H Concur (atomic t) = ⟦ t ⟧^T

Fork Concur (spawn t) = Bool
Fork Concur (atomic t) = ⊤
Ty Concur {spawn t} _ = ⟦ t ⟧^T
Ty Concur {atomic t} _ = ⟦ t ⟧^T
```

Fig. 7. Higher-order effect signature of the concur effect.

We can define the meaning of the once operator as the following elaboration:

```
eOnce : ❴ Choice ≲ Δ ❵ → Elaboration Once Δ
alg eOnce (once , k , ψ) = do
  l ← ♯ ((given hChoice handle (ψ tt)) tt)
  maybe k `fail (head l)
```

The elaboration runs the branch (ψ) of once under the hChoice handler to compute a list of all solutions of ψ. It then tries to choose the first solution and pass that to the continuation k. If the branch has no solutions, we fail. Under a strict evaluation order, the elaboration computes all possible solutions which is doing more work than needed. Agda 2.6.2.2 does not have a specified evaluation strategy, but does compile to Haskell which is lazy. In Haskell, the solutions would be lazily computed, such that the once operator cuts search short as intended.

4.4 Concurrency

Finally, we consider how to define higher-order operations for concurrency, inspired by Yang et al.'s (2022) *resumption monad* (Schmidt, 1986; Claessen, 1999; Piróg & Gibbons, 2014) defined using scoped effects. We summarize our encoding and compare it with the resumption monad. The goal is to define the two operations, whose higher-order effect signature is given in Figure 7, and summarized by these smart constructors:

```
`spawn  : {t : Type} → (m₁ m₂ : Hefty H ⟦ t ⟧^T) → Hefty H ⟦ t ⟧^T
`atomic : {t : Type} → Hefty H ⟦ t ⟧^T           → Hefty H ⟦ t ⟧^T
```

The operation `spawn $m_1$ $m_2$ spawns two threads that run concurrently, and returns the value produced by $m_1$ after both have finished. The operation `atomic m represents a block to be executed atomically; i.e., no other threads run before the block finishes executing.

We elaborate `spawn by interleaving the sub-trees of its computations. To this end, we use a dedicated function that interleaves the operations in two trees and yields as output the value of the left input tree (the first computation parameter):

```
interleave_l :  {Ref : Type → Set} → Free (CC Ref ⊕ Δ) A → Free (CC Ref ⊕ Δ) B
                → Free (CC Ref ⊕ Δ) A
```

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

Here, the CC effect is the sub/jump effect that we also used in Section 4.2.2. The interleave$_l$ function ensures atomic execution by only interleaving code that is not wrapped in a `sub operation. We elaborate Concur into CC as follows, where the to-front and from-front functions use the row insertion witness w_a to move the CC effect to the front of the row and back again:

```
eConcur : ⦃ w : CC Ref ≲ Δ ⦄ → Elaboration Concur Δ
alg eConcur (spawn t , k , ψ) =
  from-front (interleave_l (to-front (ψ true)) (to-front (ψ false))) >>= k
alg eConcur (atomic t , k , ψ) = `sub (λ ref → ψ tt >>= `jump ref) k
```

The elaboration uses `sub as a delimiter for blocks that should not be interleaved, such that the interleave$_l$ function only interleaves code that does not reside in atomic blocks. At the end of an atomic block, we `jump to the (possibly interleaved) computation context, k. By using `sub to explicitly delimit blocks that should not be interleaved, we have encoded what Wu *et al*. (2014, § 7) call *scoped syntax*.

Example. Below is an example program that spawns two threads that use the Output effect. The first thread prints 0, 1, and 2; the second prints 3 and 4.

```
ex-01234 : Hefty (Lift Output + Concur + Lift Nil) ℕ
ex-01234 = `spawn (do ↑ out "0"; ↑ out "1"; ↑ out "2"; pure 0)
                  (do ↑ out "3"; ↑ out "4"; pure 0)
```

Since the Concur effect is elaborated to interleave the effects of the two threads, the printed output appears in interleaved order:

```
test-ex-01234 : un ( ( given hOut
                       handle ( ( given hCC
                                  handle (elaborate concur-elab ex-01234)
                                ) tt ) ) tt ) ≡ (0 , "03142")
test-ex-01234 = refl
```

The following program spawns an additional thread with an `atomic block

```
ex-01234567 : Hefty (Lift Output + Concur + Lift Nil) ℕ
ex-01234567 = `spawn ex-01234
                     (`atomic (do ↑ out "5"; ↑ out "6"; ↑ out "7"; pure 0))
```

Inspecting the output, we see that the additional thread indeed computes atomically:

```
test-ex-01234567 : un ( ( given hOut
                          handle ( ( given hCC
                                     handle (elaborate concur-elab ex-01234567)
                                   ) tt ) ) tt ) ≡ (0 , "05673142")
test-ex-01234567 = refl
```

The example above is inspired by the resumption monad, and in particular by the scoped effects definition of concurrency due to Yang *et al*. (2022). Yang *et al*. do not (explicitly) consider how to define the concurrency operations in a modular style. Instead, they give a

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

direct semantics that translates to the resumption monad which we can encode as follows in Agda (assuming resumptions are given by the free monad):

```
data Resumption Δ A : Set where
  done : A                                → Resumption Δ A
  more : Free Δ (Resumption Δ A) → Resumption Δ A
```

We could elaborate into this type using a hefty algebra $\mathsf{Alg}^{\mathsf{H}}$ Concur (Resumption Δ) but that would be incompatible with our other elaborations which use the free monad. For that reason, we emulate the resumption monad using the free monad instead of using the Resumption type directly.

5 Modular reasoning for higher-order effects

A key aspect of algebraic effects and handlers is the ability to state and prove *equational laws* as part of an effect's specification that characterize correct implementations. Usually, an effect is associated with several laws that govern its behavior. An effect, together with its laws, constitutes an *effect theory* (Plotkin & Power, 2002, 2003; Hyland *et al*., 2006; Yang & Wu, 2021). Equational reasoning within an effect theory can be used to derive syntactic equalities between effectful programs without appealing to the effect's implementation. Such equalities remain true in the semantic domain for any handler that respects the laws of the theory.

The concept of effect theory extends to *higher-order effect theories*, which describe the intended behavior of higher-order effects. In this section, we first discuss how to define theories for algebraic effects in Agda by adapting the exposition of Yang & Wu (2021), and show how correctness of implementations with respect to a given theory can be stated and proved. We then extend this reasoning infrastructure to higher-order effects, which allows for the derivation of syntactic equalities between programs with higher-order effects and modular reasoning about the correctness of elaborations.

5.1 Effect theories and implementation correctness

Let us consider the state effect as an example, which comprises the get and put operations. With the state effect, we typically associate a set of equations (or laws) that specify how its implementations ought to behave. One such law is the *get-get* law, which captures the intuition that the state returned by two subsequent get operations does not change if we do not use the put operation in between:

$$\text{`get} \gg\!= \lambda s \to \text{`get} \gg\!= \lambda s' \to k\ s\ s' \equiv \text{`get} \gg\!= \lambda s \to k\ s\ s$$

We can define equational laws for higher-order effects in a similar fashion. For example, the following *catch-return* law for the `catch operation of the Catch effect, stating that catching exceptions in a computation that only returns a value does nothing.

$$\text{`catch}\ (\mathsf{pure}\ x)\ m \equiv \mathsf{pure}\ x$$

Correctness of an implementation of an algebraic effect with respect to a given theory is defined by comparing the implementations of programs that are equal under that theory.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

That is, if we can show that two programs are equal using the equations of a theory for its effects, handling the effects should produce equal results. For instance, a way to implement the state effect is by mapping programs to functions of the form $S \to S \times A$. Such an implementation would be correct if programs that are equal with respect to a theory of the state effect are mapped to functions that give the same value and output state for every input state.

For higher-order effects, correctness is defined in a similar manner. However, since we define higher-order effects by elaborating them into algebraic effects, correctness of elaborations with respect to a higher-order effect theory is defined by comparing the elaborated programs. Crucially, the elaborated programs do not have to be syntactically equal, but rather we should be able to prove them equal using a theory of the algebraic effects used to implement a higher-order effect.

Effect theories are known to be closed under the co-product of effects, by combining the equations into a new theory that contains all equations for both effects (Hyland *et al*., 2006). Similarly, theories of higher-order effects are closed under sums of higher-order effect signatures. In Section 5.9, we show that composing two elaborations preserves their correctness, with respect to the sum of their respective theories.

5.2 Theories of algebraic effects

Theories of effects are collections of equations, so we start defining the type of equations in Agda. At its core, an equation for an effect Δ is given by a pair of effect trees of type Free Δ A, that define the left- and right-hand side of the equation. However, looking at the *get-get* law above, we see that this equation contains a *term metavariable*; i.e., k. Furthermore, when considering the type of k, which is $S \to S \to$ Free Δ A, we see that it refers to a *type metavariable*; i.e., A. Generally speaking, an equation may refer to any number of term metavariables, which, in turn, may depend on any number of type metavariables. Moreover, the type of the value returned by the left-hand side and right-hand side of an equation may depend on these type metavariables as well, as is the case for the *get-get* law above. This motivates the following definition of equations in Agda.

```
record Equation (Δ : Effect) : Set₁ where
  field
    V       : ℕ
    Γ       : Vec Set V → Set
    R       : Vec Set V → Set
    lhs rhs : (vs : Vec Set V) → Γ vs → Free Δ (R vs)
```

An equation consists of five components. The field V defines the number of type metavariables used in the equation. Then, the fields Γ and R, respectively, define the term metavariables (Vec Set V → Set) and return type (Vec Set V → Set) of the equation.

Example. To illustrate how the Equation record captures equational laws of effects, we consider how to define the *get-get* as a value of type Equation State.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
get-get : Equation State
V   get-get = 1
Γ   get-get = λ where (A :: []) → ℕ → ℕ → Free State A
R   get-get = λ where (A :: []) → A
lhs get-get (A :: []) k = 'get >>= λ s → 'get >>= λ s' → k s s'
rhs get-get (A :: []) k = 'get >>= λ s → k s s
```

The fields lhs and rhs define the left- and right-hand sides of the equation. Both sides only use a single term metavariable, representing a continuation of type $\mathbb{N} \to \mathbb{N} \to$ Free State A. The field Γ declares this term meta-variable. For equations with more than $n > 1$ metavariables, we would define Γ as an n-ary product instead.

5.3 Modal necessity

The current definition of equations is too weak, in the sense that it does not apply in many situations where it should. The issue is that it fixes the set of effects that can be used in the left- and right-hand side. To illustrate why this is problematic, consider the following equality:

$$\text{get} >\!\!>\!\!= \lambda s \to \text{get} >\!\!>\!\!= \lambda s' \to \text{throw} \equiv \text{get} >\!\!>\!\!= \lambda s \to \text{throw} \tag{5.1}$$

We might expect to be able to prove this equality using the *get-get* law, but using the embedding of the law defined above—i.e., get-get—this is not possible. The reason for this is that we cannot pick an appropriate instantiation for the term metavariable k: it ranges over values of type $S \to S \to$ Free State A, inhibiting all references to effectful operation that are not part of the state effect, such as throw.

Given an equation for the effect Δ, the solution to this problem is to view Δ as a *lower bound* on the effects that might occur in the left-hand and right-hand side of the equation, rather than an exact specification. Effectively, this means that we close over all posible contexts of effects in which the equation can occur. This pattern of closing over all possible extensions of a type index is well-known (Allais *et al.*, 2021; van der Rest *et al.*, 2022), and corresponds to a shallow embedding of the Kripke semantics of the necessity modality from modal logic. We can define it in Agda as follows.[35]

```
record □ (P : Effect → Set₁) (Δ : Effect) : Set₁ where
  constructor necessary
  field
    □⟨_ER⟩ : ∀ {Δ'} → ⦃ Δ ≲ Δ' ⦄ → P Δ'
```

Intuitively, the □ modality transforms, for any effect-indexed type (P : Effect → Set_1), an *exact* specification of the set of effects to a *lower bound* on the set of effects. For equations, the difference between terms of type Equation Δ and □ Equation Δ amounts to the former defining an equation relating programs that have exactly effects Δ, while the latter defines an equation relating programs that have at least the effects Δ but potentially more. The □ modality is a *comonad*: the counit (extract below) witnesses that we can always transform

[35] The **constructor** keyword declares a function that we can call to construct an instance of a record; and that we can pattern match on to destruct record instances.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

a lower bound on effects to an exact specification, by instantiating the extension witness with a proof of reflexivity.

```
extract : {P : Effect → Set₁} → □ P Δ → P Δ
extract px = □⟨ px ⟩ ⦇ ≲-refl ⦈
```

We can now redefine the *get-get* law such that it applies to all programs that have the State effect, but potentially other effects too.

```
get-get : □ Equation State
V   □⟨ get-get ⟩ = 1
Γ   □⟨ get-get ⟩ (A :: [])   = ℕ → ℕ → Free _ A
R   □⟨ get-get ⟩ (A :: [])   = A
lhs □⟨ get-get ⟩ (A :: []) k = 'get >>= λ s → 'get >>= λ s′ → k s s′
rhs □⟨ get-get ⟩ (A :: []) k = 'get >>= λ s → k s s
```

The above definition of the *get-get* law now lets us prove the equality in Equation (5.1); the term metavariable k ranges ranges over all continuations that return a tree of type Free $\Delta' A$, for all Δ' such that State $\lesssim \Delta'$. This way, we can instantiate Δ' with an effect signature that subsumes both the State and the Throw, which in turn allows us to instantiate k with throw.

5.4 Effect theories

Equations for an effect Δ can be combined into a *theory* for Δ. A theory for the effect Δ is simply a collection of equations, transformed using the $\Box$ modality to ensure that term metavariables can range over programs that include more effects than just Δ.

```
record Theory (Δ : Effect) : Set₁ where
  field
    arity     : Set
    equations : arity → □ Equation Δ
```

An effect theory consists of an arity that defines the number of equations in the theory, and a function that maps arities to equations. We can think of effect theories as defining a specification for how implementations of an effect ought to behave. Although implementations may vary, depending for example on whether they are tailored to readability or efficiency, they should at least respect the equations of the theory of the effect they implement. We will make precise what it means for an implementation to respect an equation in Section 5.6.

Effect theories are closed under several composition operations that allow us to combine the equations of different theories into single theory. The most basic way of combining effect theories is by summing their arities.

```
_⟨+⟩_ : Theory Δ → Theory Δ → Theory Δ
arity     (T₁ ⟨+⟩ T₂) = arity T₁ ⊎ arity T₂
equations (T₁ ⟨+⟩ T₂) (inj₁ a) = equations T₁ a
equations (T₁ ⟨+⟩ T₂) (inj₂ a) = equations T₂ a
```

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

This way of combining effects is somewhat limiting, as it imposes that the theories we are combining are theories for the exact same effect. It is more likely, however, that we would want to combine theories for different effects. This requires that we can *weaken* effect theories with respect to the $_\lesssim_$ relation.

```
weaken-□ : {P : Effect → Set₁} → ⦃ Δ₁ ≲ Δ₂ ⦄ → □ P Δ₁ → □ P Δ₂
□⟨ weaken-□ ⦃ w ⦄ px ⟩ ⦃ w′ ⦄ = □⟨ px ⟩ ⦃ ≲-trans w w′ ⦄

weaken-theory : ⦃ Δ₁ ≲ Δ₂ ⦄ → Theory Δ₁ → Theory Δ₂
arity (weaken-theory T) = arity T
equations (weaken-theory T) = λ a → weaken-□ (T .equations a)
```

Categorically speaking, the observation that for a given effect-indexed type P we can transform a value of type $P\ \Delta_1$ to a value of type $P\ \Delta_2$ if we know that $\Delta_1 \lesssim \Delta_2$ is equivalent to saying that P is a functor from the category of containers and container morphisms to the category of sets. From this perspective, the existence of weakening for free Free, as witnessed by the $\sharp$ operation discussed in Section 3, implies that it too is a such a functor.

With weakening for theories at our disposal, we can combine effect theories for different effects into a theory of the coproduct of their respective effects. This requires us to first define appropriate witnesses relating coproducts to effect inclusion.

```
≲-⊕-left  : Δ₁ ≲ (Δ₁ ⊕ Δ₂)
≲-⊕-right : Δ₂ ≲ (Δ₁ ⊕ Δ₂)
```

It is now straightforward to show that effect theories are closed under the coproduct of effect signatures, by summing the weakened theories.

```
_[+]_ : Theory Δ₁ → Theory Δ₂ → Theory (Δ₁ ⊕ Δ₂)
T₁ [+] T₂ = weaken-theory ⦃ ≲-⊕-left ⦄ T₁ ⟨+⟩ weaken-theory ⦃ ≲-⊕-right ⦄ T₂
```

While this operation is in principle sufficient for our purposes, it forces a specific order on the effects of the combined theories. We can further generalize the operation above to allow for the effects of the combined theory to appear in any order. This requires the following instances.

```
≲-•-left  : ⦃ Δ₁ • Δ₂ ≈ Δ ⦄ → Δ₁ ≲ Δ
≲-•-right : ⦃ Δ₁ • Δ₂ ≈ Δ ⦄ → Δ₂ ≲ Δ
```

We show that effect theories are closed under coproducts up to reordering by, again, summing the weakened theories.

```
compose-theory : ⦃ Δ₁ • Δ₂ ≈ Δ ⦄ → Theory Δ₁ → Theory Δ₂ → Theory Δ
compose-theory T₁ T₂
  = weaken-theory ⦃ ≲-•-left ⦄ T₁ ⟨+⟩ weaken-theory ⦃ ≲-•-right ⦄ T₂
```

Since equations are defined by storing the syntax trees that define their left-hand and right-hand side, and effect trees are weakenable, we would expect equations to be weakenable too. Indeed, we can define the following function witnessing weakenability of equations.

```
weaken-eq : ⦃ Δ₁ ≲ Δ₂ ⦄ → Equation Δ₁ → Equation Δ₂
```

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

This begs the question: why would we opt to use weakenability of the $\Box$ modality (or, bother with the $\Box$ modality at all) to show that theories are weakenable, rather than using weaken-eq directly? Although the latter approach would indeed allow us to define the composition operations for effect theories defined above, the possible ways in which we can instantiate term metavariables remains too restrictive. That is, we would still not be able to prove the equality in Equation (5.1), despite the fact that we can weaken the *get-get* law so that it applies to programs that use the Throw effect as well. Instantiations of the term metavariable k will be limited to weakened effect trees, precluding any instantiation that use operations of effects other than State, such as throw.

Finally, we define the following predicate to witness that an equation is part of a theory.

```
_◀_ : □ Equation Δ → Theory Δ → Set₁
eq ◀ T = ∃ λ a → T .equations a ≡ eq
```

We construct a proof $eq \blacktriangleleft T$ that an equation eq is part of a theory T by providing an arity together with a proof that T maps to eq for that arity.

5.5 Syntactic equivalence of effectful programs

Propositional equality of effectful programs is too strict, as it precludes us from proving equalities that rely on a semantic understanding of the effects involved, such as the equality in Equation (5.1). The solution is to define an inductive relation that captures syntactic equivalence modulo some effect theory. We base our definition of syntactic equality of effectful programs on the relation defining equivalent computations by Yang & Wu (2021), Definition 3.1, adapting their definition where necessary to account for the use of modal necessity in the definition of Theory.

```
data _≈⟨_⟩_ {Δ Δ′} ⦃ _ : Δ ≲ Δ′ ⦄
     : (m₁ : Free Δ′ A) → Theory Δ → (m₂ : Free Δ′ A) → Set₁ where
```

A value of type $m_1 \approx\langle\ T\ \rangle\ m_2$ witnesses that programs m_1 and m_2 are equal modulo the equations of theory T. The first three constructors ensure that it is an equivalence relation.

```
≈-refl   : m ≈⟨ T ⟩ m
≈-sym    : m₁ ≈⟨ T ⟩ m₂ → m₂ ≈⟨ T ⟩ m₁
≈-trans  : m₁ ≈⟨ T ⟩ m₂ → m₂ ≈⟨ T ⟩ m₃ → m₁ ≈⟨ T ⟩ m₃
```

Then, we add the following congruence rule, which establishes that we can prove equality of two programs starting with the same operation by proving that the continuations yield equal programs for every possible value.

```
≈-cong :  (op : Op Δ′)
       → (k₁ k₂ : Ret Δ′ op → Free Δ′ A)
       → (∀ x → k₁ x ≈⟨ T ⟩ k₂ x)
       → impure (op , k₁) ≈⟨ T ⟩ impure (op , k₂)
```

The final constructor allows to prove equality of programs by reifying equations of an effect theory.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
≈-eq :  (eq : □ Equation Δ)
     → (px : eq ◀ T)
     → (vs : Vec Set (V (□⟨ eq ⟩))))
     → (γ : Γ (□⟨ eq ⟩)) vs)
     → (k : R (□⟨ eq ⟩)) vs → Free Δ′ A)
     → (lhs (□⟨ eq ⟩)) vs γ >>= k) ≈⟨ T ⟩ (rhs (□⟨ eq ⟩)) vs γ >>= k)
```

Since the equations of a theory are wrapped in the □ modality, we cannot refer to its components directly, but we must first provide a suitable extension witness.

With the ≈-eq constructor, we can prove equivalence between the left-hand and right-hand side of an equation, sequenced with an arbitrary continuation k. For convenience, we define the following lemma that allows us to apply an equation where the sides of the equation do not have a continuation.

```
use-equation :  ⦃ _ : Δ ≲ Δ′ ⦄
             → {T : Theory Δ}
             → (eq : □ Equation Δ)
             → eq ◀ T
             → (vs : Vec Set (V □⟨ eq ⟩))
             → {γ : Γ (□⟨ eq ⟩)) vs}
             → lhs (□⟨ eq ⟩)) vs γ ≈⟨ T ⟩ rhs (□⟨ eq ⟩)) vs γ
```

The definition of use-equation follows readily from the right-identity law for monads, i.e., $m >\!\!>= \mathsf{pure} \equiv m$, which allows us to instantiate ≈-eq with pure.

To construct proofs of equality, it is convenient to use the following set of combinators to write proof terms in an equational style. They are completely analogous to the combinators commonly used to construct proofs of Agda's propositional equality, for example, as found in PLFA (Wadler *et al.*, 2020).

```
module ≈-Reasoning (T : Theory Δ) ⦃ _ : Δ ≲ Δ′ ⦄ where
  begin_ : {m₁ m₂ : Free Δ′ A} → m₁ ≈⟨ T ⟩ m₂ → m₁ ≈⟨ T ⟩ m₂
  begin eq = eq

  _■ : (m : Free Δ′ A) → m ≈⟨ T ⟩ m
  m ■ = ≈-refl

  _≈⟪⟫_ : (m₁ : Free Δ′ A) {m₂ : Free Δ′ A} → m₁ ≈⟨ T ⟩ m₂ → m₁ ≈⟨ T ⟩ m₂
  m₁ ≈⟪⟫ eq = eq

  _≈⟪_⟫_ : (m₁ {m₂ m₃} : Free Δ′ A) → m₁ ≈⟨ T ⟩ m₂ → m₂ ≈⟨ T ⟩ m₃ → m₁ ≈⟨ T ⟩ m₃
  m₁ ≈⟪ eq₁ ⟫ eq₂ = ≈-trans eq₁ eq₂
```

We now have all the necessary tools to prove syntactic equality of programs modulo a theory of their effect. To illustrate, we consider how to prove the equation in Equation (5.1). First, we define a theory for the State effect containing the get-get◀ law. While this is not the only law typically associated with State, for this example it is enough to only have the get-get law.

```
StateTheory : Theory State
arity StateTheory        = ⊤
equations StateTheory tt = get-get
```

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

Now to prove the equality in Equation (5.1) is simply a matter of invoking the get-get law.

```
get-get-throw :
      ⦃ _ : Throw ≲ Δ ⦄ ⦃ _ : State ≲ Δ ⦄
  → (`get >>= λ s → `get >>= λ s′ → `throw {A = A})
    ≈⟨ StateTheory ⟩ (`get >>= λ s → `throw)
get-get-throw {A = A} = begin
   `get >>= (λ s → `get >>= (λ s′ → `throw))
 ≈⟪ use-equation get-get (tt , refl) (A :: []) ⟫
   `get >>= (λ s → `throw)
 ∎
 where open ≈-Reasoning StateTheory
```

5.6 Handler correctness

A handler is correct with respect to a given theory if handling syntactically equal programs yields equal results. Since handlers are defined as algebras over effect signatures, we start by defining what it means for an algebra of an effect Δ to respect an equation of the same effect, adapting Definition 2.1 from the exposition of Yang & Wu (2021).

```
Respects : Alg Δ A → Equation Δ → Set₁
Respects alg eq = ∀ {vs γ k} →
  fold k alg (lhs eq vs γ) ≡ fold k alg (rhs eq vs γ)
```

An algebra *alg* respects an equation *eq* if folding with that algebra produces propositionally equal results for the left- and right-hand side of the equation, for all possible instantiations of its type and term metavariables, and continuations k.

A handler H is correct with respect to a given theory T if its algebra respects all equations of T (Yang & Wu, 2021, Definition 4.3).

```
Correct : {P : Set} → Theory Δ → ⟨ A ! Δ ⇒ P ⇒ B ! Δ′ ⟩ → Set₁
Correct T H = ∀ {eq} → eq ◀ T → Respects (H .hdl) (extract eq)
```

We can now show that the handler for the State effect defined in Figure 1 is correct with respect to StateTheory. The proof follows immediately by reflexivity.

```
hStCorrect : Correct {A = A} {Δ′ = Δ} StateTheory hSt
hStCorrect (tt , refl) {_ :: []} {γ = k} = refl
```

5.7 Theories of higher-order effects

For the most part, equations and theories for higher-order effects are defined in the same way as for first-order effects and support many of the same operations. Indeed, the definition of equations ranging over higher-order effects is exactly the same as its first-order counterpart, the most major difference being that the left-hand and right-hand side are now defined as Hefty trees. To ensure compatibility with the use of type universes to avoid size-issues, we must also allow type metavariables to range over the types in a universe

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

in addition to Set. For this reason, the set of type metavariables is no longer described by a natural number, but rather by a list of kinds, which stores for each type metavariable whether it ranges over a types in a universe, or an Agda Set.

```
data Kind : Set where set type : Kind
```

A TypeContext carries unapplied substitutions for a given set of type metavariables and is defined by induction over a list of kinds.[36]

```
TypeContext : List Kind → Set₁
TypeContext []           = Level.Lift _ ⊤
TypeContext (set  :: vs) = Set × TypeContext vs
TypeContext (type :: vs) = Level.Lift (sℓ 0ℓ) Type × TypeContext vs
```

Equations of higher-order effects are then defined as follows.

```
record Equation^H (H : Effect^H) : Set₁ where
  field
    V       : List Kind
    Γ       : TypeContext V → Set
    R       : TypeContext V → Set
    lhs rhs : (vs : TypeContext V) → Γ vs → Hefty H (R vs)
```

This definition of equations suffers the same problem when it comes to term metavariables, which here too can only range over programs that exhibit the exact effect that the equation is defined for. Again, we address the issue using an embedding of modal necessity to close over all possible extensions of this effect. The definition is analogous to the one in Section 5.3, but this time we use higher-order effect subtyping as the modal accessibility relation:

```
record □ (P : Effect^H → Set₁) (H : Effect^H) : Set₁ where
  constructor necessary
  field □⟨_⟩ : ∀ {H'} → ⦃ H ≲^H H' ⦄ → P H'
```

To illustrate: we can define the *catch-return* law from the introduction of this section as a value of type □ Equation^H Catch a follows. Since the `catch operation relies on a type universe to avoid size issues, the sole type metavariable of this equation must range over the types in this universe as well.

```
catch-return : □ Equation^H Catch
V   □⟨ catch-return ⟩               = type :: []
Γ   □⟨ catch-return ⟩ (lift t , _) = ⟦ t ⟧^T × Hefty _ ⟦ t ⟧^T
R   □⟨ catch-return ⟩ (lift t , _) = ⟦ t ⟧^T
lhs □⟨ catch-return ⟩ _ (x , m)    = `catch (pure x) m
rhs □⟨ catch-return ⟩ _ (x , m)    = pure x
```

Theories of higher-order effects bundle extensible equations. The setup is the same as for theories of first-order effects.

[36] Level.Lift lifts a type in Set to a type in Set₁. The constructor of Level.Lift is lift.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
record Theory^H (H : Effect^H) : Set₁ where
  field
    arity : Set
    equations : arity → □ Equation^H H
```

The following predicate establishes that an equation is part of a theory. We prove this fact by providing an arity whose corresponding equation is equal to *eq*.

```
_◀^H_ : □ Equation^H H → Theory^H H → Set₁
eq ◀^H Th = ∃ λ a → eq ≡ equations Th a
```

Weakenability of theories of higher-order effects then follows from weakenability of its equations.

```
weaken-□ : ∀ {P} → ⦃ H₁ ≲^H H₂ ⦄ → □ P H₁ → □ P H₂
□⟨ weaken-□ ⦃ w ⦄ px ⟩ ⦃ w′ ⦄ = □⟨ px ⟩ ⦃ ≲^H-trans w w′ ⦄

weaken-theory^H : ⦃ H₁ ≲^H H₂ ⦄ → Theory^H H₁ → Theory^H H₂
arity     (weaken-theory^H Th)   = Th .arity
equations (weaken-theory^H Th) a = weaken-□ (Th .equations a)
```

Theories of higher-order effects can be combined using the following sum operation. The resulting theory contains all equations of both argument theories.

```
_⟨+⟩^H_ : ∀[ Theory^H ⇒ Theory^H ⇒ Theory^H ]
arity     (Th₁ ⟨+⟩^H Th₂)          = arity Th₁ ⊎ arity Th₂
equations (Th₁ ⟨+⟩^H Th₂) (inj₁ a) = equations Th₁ a
equations (Th₁ ⟨+⟩^H Th₂) (inj₂ a) = equations Th₂ a
```

Theories of higher-order effects are closed under sums of higher-order effect theories as well. This operation is defined by appropriately weakening the respective theories, for which we need the following lemmas witnessing that higher-order effect signatures can be injected in a sum of signatures.

```
≲-∔-left  : H₁ ≲^H (H₁ ∔ H₂)
≲-∔-right : H₂ ≲^H (H₁ ∔ H₂)
```

The operation that combines theories under signature sums is then defined like so.

```
_[+]^H_ : Theory^H H₁ → Theory^H H₂ → Theory^H (H₁ ∔ H₂)
Th₁ [+]^H Th₂
  = weaken-theory^H ⦃ ≲-∔-left ⦄ Th₁ ⟨+⟩^H weaken-theory^H ⦃ ≲-∔-right ⦄ Th₂
```

5.8 Equivalence of programs with higher-order effects

We define the following inductive relation to capture equivalence of programs with higher-order effects modulo the equations of a given theory.

```
data _≅⟨_⟩_ ⦃ _ : H₁ ≲^H H₂ ⦄
       : (m₁ : Hefty H₂ A) → Theory^H H₁ → (m₂ : Hefty H₂ A) → Set₁ where
```

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

To ensure that it is indeed an equivalence relation, we include constructors for reflexivity, symmetry, and transitivity.

```
≅-refl  : ∀ {m : Hefty H₂ A}
        → m ≅⟨ Th ⟩ m

≅-sym   : ∀ {m₁ : Hefty H₂ A} {m₂}
        → m₁ ≅⟨ Th ⟩ m₂
        → m₂ ≅⟨ Th ⟩ m₁

≅-trans : ∀ {m₁ : Hefty H₂ A} {m₂ m₃}
        → m₁ ≅⟨ Th ⟩ m₂ → m₂ ≅⟨ Th ⟩ m₃
        → m₁ ≅⟨ Th ⟩ m₃
```

Furthermore, we include the following congruence rule that equates two program trees that have the same operation at the root, if their continuations are equivalent for all inputs.

```
≅-cong :    (op : Op^H H₂)
        → (k₁ k₂ : Ret^H H₂ op → Hefty H₂ A)
        → (ψ₁ ψ₂ : (ϕ : Fork H₂ op) → Hefty H₂ (Ty H₂ ϕ))
        → (∀ {x} → k₁ x ≅⟨ Th ⟩ k₂ x)
        → (∀ {ϕ} → ψ₁ ϕ ≅⟨ Th ⟩ ψ₂ ϕ)
        → impure (op , k₁ , ψ₁) ≅⟨ Th ⟩ impure ( op , k₂ , ψ₂ )
```

Finally, we include a constructor that equates two programs using an equation of the theory.

```
≅-eq    :    (eq : □ Equation^H H₁)
        → eq ◀^H Th
        → (vs : TypeContext (V □⟨ eq ⟩))
        → (γ : Γ □⟨ eq ⟩ vs)
        → (k : R □⟨ eq ⟩ vs → Hefty H₂ A)
        → (lhs □⟨ eq ⟩ vs γ >>= k) ≅⟨ Th ⟩ (rhs □⟨ eq ⟩ vs γ >>= k)
```

We can define the same reasoning combinators as in Section 5.5 to construct proofs of equivalence for programs with higher-order effects.

```
module ≅-Reasoning ⦃ _ : H₁ ≲^H H₂ ⦄ (Th : Theory^H H₁) where

  begin_ : {m₁ m₂ : Hefty H₂ A} → m₁ ≅⟨ Th ⟩ m₂ → m₁ ≅⟨ Th ⟩ m₂
  begin eq = eq

  _■ : (c : Hefty H₂ A) → c ≅⟨ Th ⟩ c
  c ■ = ≅-refl

  _≅⟪⟫_ : (m₁ : Hefty H₂ A) {m₂ : Hefty H₂ A} → m₁ ≅⟨ Th ⟩ m₂ → m₁ ≅⟨ Th ⟩ m₂
  c₁ ≅⟪⟫ eq = eq

  _≅⟪_⟫_ : (c₁ {c₂ c₃} : Hefty H₂ A) → c₁ ≅⟨ Th ⟩ c₂ → c₂ ≅⟨ Th ⟩ c₃ → c₁ ≅⟨ Th ⟩ c₃
  c₁ ≅⟪ eq₁ ⟫ eq₂ = ≅-trans eq₁ eq₂
```

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

To illustrate, we can prove that the programs catch throw (censor f m) and censor f m are equal under a theory for the *af Catch* effect that contains the *catch-return* law.

```
catch-return-censor : ∀ {t : Type} {f} {x : ⟦ t ⟧ᵀ} {m : Hefty H ⟦ t ⟧ᵀ}
                      → ⦃ _ : Catch ≲ᴴ H ⦄ → ⦃ _ : Censor ≲ᴴ H ⦄
                      → `catch (pure x) (`censor f m)
                        ≅⟨ CatchTheory ⟩ pure x
catch-return-censor {f = f} {x = x} {m = m} =
  begin
    `catch (pure x) (`censor f m)
  ≅⟪ use-equationᴴ catch-return (tt , refl) _ ⟫
    pure x
  ∎
  where open ≅-Reasoning _
```

The equivalence proof above makes, again, essential use of modal necessity. That is, by closing over all possible extensions of the Catch effe, the term metavariable in the *catch-return* law to range over programs that have higher-order effects other than Catch, which is needed to apply the law if the second branch of the catch operation contains the censor operation.

5.9 Correctness of elaborations

As the first step toward defining correctness of elaborations, we must specify what it means for an algebra over a higher-order effect signature H to respect an equation. The definition is broadly similar to its counterpart for first-order effects in Section 5.6, with the crucial difference that the definition of "being equation respecting" for algebras over higher-order effect signatures is parameterized over a binary relation _≈_ between first-order effect trees. In practice, this binary relation will be instantiated with the inductive equivalence relation defined in Section 5.5; propositional equality would be too restrictive, since that does not allow us prove equivalence of programs using equations of the first-order effect(s) that we elaborate into.

```
Respectsᴴ : (_≈_ : ∀ {A} → Free Δ A → Free Δ A → Set₁)
            → Algᴴ H (Free Δ) → Equationᴴ H → Set₁
Respectsᴴ _≈_ alg eq =
  ∀ {vs γ} → cataᴴ pure alg (lhs eq vs γ) ≈ cataᴴ pure alg (rhs eq vs γ)
```

Since elaborations are composed in parallel, the use of necessity in the definition of equations has additional consequences for the definiton of elaboration correctness. That is, correctness of an elaboration is defined with respect to a theory whose equations have left-hand and right-hand sides that may contain term metavariables that range over programs with more higher-order effects than those the elaboration is defined for. Therefore, to state correctness, we must also close over all possible ways these additional effects are elaborated. For this, we define the following binary relation on extensible elaborations.[37]

[37] Here, injᴴ is the higher-order counterpart to the inj function discussed in Section 2.2.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
record _⊑_ (e₁ : □ (Elaboration H₁) Δ₁) (e₂ : □ (Elaboration H₂) Δ₂) : Set₁ where
  field
    ⦃ ≲-eff   ⦄ : Δ₁ ≲ Δ₂
    ⦃ ≲ᴴ-eff ⦄ : H₁ ≲ᴴ H₂
    preserves-cases
      : ∀ {M} (m : ⟦ H₁ ⟧ᴴ M A)
      → (e′ : ∀[ M ⇒ Free Δ₂ ])
      → □⟨ e₁ ⟩ .alg (map-sigᴴ (λ {x} → e′ {x}) m)
        ≡ extract e₂ .alg (map-sigᴴ (λ {x} → e′ {x}) (injᴴ {X = A} m))
```

A proof of the form $e_1 \sqsubseteq e_2$ witnesses that the elaboration e_1 is included in e_2. Informally, this means that e_2 may elaborate a bigger set of higher-order effects, for which it may need to refer to a bigger set of first-order effects, but for those higher-order effects that both e_1 and e_2 know how to elaborate, they should agree on how those effects are elaborated.

We then define correctness of elaborations as follows.

```
Correctᴴ : Theoryᴴ H → Theory Δ → □ (Elaboration H) Δ → Set₁
Correctᴴ Th T e =
  ∀ {Δ′ H′}
  → (e′ : □ (Elaboration H′) Δ′)
  → ⦃ _ : e ⊑ e′ ⦄
  → {eq : □ Equationᴴ _}
  → eq ◀ᴴ Th
  → Respectsᴴ (_≈⟨ T ⟩_) (extract e′) □⟨ eq ⟩
```

Which is to say that an elaboration is correct with respect to a theory of the higher-order effects it elaborates (*Th*) and a theory of the first-order effects it elaborates into (*T*), if all possible extensions of said elaboration respect all equations of the higher-order theory, modulo the equations of the first-order theory.

Crucially, correctness of elaborations is preserved under composition of elaborations. Figure 8 shows the type of the corresponding correctness theorem in Agda; for the full details of the proof we refer to the Agda formalization accompanying this paper (van der Rest & Poulsen, 2024). We remark that correctness of a composed elaboration is defined with respect to the composition of the theories of the first-order effects that the respective elaborations use. Constructing a handler that is correct with respect to this composed first-order effect theory is a separate concern; a solution based on *fusion* is detailed in the work by Yang & Wu (2021).

5.10 Proving correctness of elaborations

To illustrate how the reasoning infrastructure build up in this section can be applied to verify correctness of elaborations, we show how to verify the *catch-return* law for the elaboration eCatch defined in Section 3.4. First, we define the following syntax for invoking a known elaboration.

```
module Elab (e : □ (Elaboration H) Δ) where
  ℰ⟦_⟧ : Hefty H A → Free Δ A
  ℰ⟦ m ⟧ = elaborate (extract e) m
```

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

```
compose-elab-correct :  ⦃ _ : Δ₁ • Δ₂ ≈ Δ ⦄
                      → (e₁ : □ (Elaboration H₁) Δ₁)
                      → (e₂ : □ (Elaboration H₂) Δ₂)
                      → (T₁ : Theory Δ₁)
                      → (T₂ : Theory Δ₂)
                      → (Th₁ : Theory^H H₁)
                      → (Th₂ : Theory^H H₂)
                      → Correct^H Th₁ T₁ e₁
                      → Correct^H Th₂ T₂ e₂
                      → Correct^H (Th₁ [+]^H Th₂) (compose-theory T₁ T₂)
                          (compose-elab e₁ e₂)
```

Fig. 8. The central correctness theorem, which establishes that correctness of elaborations is preserved under composition.

When opening the module *Elab*, we can use the syntax $\mathcal{E}[\![\ m\ ]\!]$ for elaborating m with some known elaboration, which helps to simplify and improve readability of equational proofs for higher-order effects.

Now, to prove that eCatch is correct with respect to a higher-order theory for the Catch effect containing the *catch-return* law, we must produce a proof that the programs $\mathcal{E}[\![\ \text{`catch}\ (\text{return}\ x)\ m\ ]\!]$ and $\mathcal{E}[\![\ \text{return}\ ]\!]$ are equal (in the sense of the inductive equivalence relation defined in Section 5.5) with respect to some first-order theory for the Throw effect. In this instance, we do not need any equations from this underlying theory to prove the equality, but sometimes it is necessary to invoke equations of the underlying first-order effects to prove correctness of an elaboration.

```
eCatchCorrect : {T : Theory Throw} → Correct^H CatchTheory T eCatch
eCatchCorrect {Δ′ = Δ′} e′ ⦃ ζ ⦄ (tt , refl) {γ = x , m} =
  begin
    ℰ[[ `catch (pure x) m ]]
  ≈⟨ from-≡ (sym $ ζ .preserves-cases _ ℰ[[_]]) ⟩
    (♯ (given hThrow handle (pure x) $ tt)) >>= maybe′ pure (ℰ[[ m ]])
  ≈⟨⟩ {- By definition of hThrow -}
    (pure (just x) >>= maybe′ pure ((ℰ[[ m ]] >>= pure)))
  ≈⟨⟩ {- By definition of >>= -}
    ℰ[[ pure x ]]
  ■
  where
    open ≈-Reasoning _
    open Elab e′
```

In the Agda formalization accompanying this paper (van der Rest & Poulsen, 2024), we verify correctness of elaborations for the higher-order operations that are part of the 3MT library by Delaware *et al.* (2013). Table 1 shows an overview of first-order and higher-order effects included in the development, and the laws which we prove about their handlers respectively elaborations.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

Table 1. Overview of effects, their operations, and verified laws in the Agda code

Effect	Laws	
Throw	$\text{`throw} \gg\!= k \equiv k$	*bind-throw*
State	$\text{`get} \gg\!= \lambda\, s \to \text{`get} \gg\!= k\, s \equiv \text{`get} \gg\!= k\, s\, s$	*get-get*
	$\text{`get} \gg\!= \text{`put} \equiv \text{pure}\, x$	*get-put*
	$\text{`put}\, s \gg \text{`get} \equiv \text{`put}\, s \gg \text{pure}\, s$	*put-get*
	$\text{`put}\, s \gg \text{`put}\, s' \equiv \text{`put}\, s'$	*put-put*
Reader	$\text{`ask} \gg m \equiv m$	*ask-query*
	$\text{`ask} \gg\!= \lambda\, r \to \text{`ask} \gg\!= k\, r \equiv \text{`ask} \gg\!= \lambda\, r \to k\, r\, r$	*ask-ask*
	$m \gg\!= \lambda\, x \to \text{`ask} \gg\!= \lambda\, r \to k\, x\, r \equiv \text{`ask} \gg\!= \lambda\, r \to m \gg\!= \lambda\, x \to k\, x\, r$	*ask-bind*
LocalReader	$\text{`local}\, f\, (\text{pure}\, x) \equiv \text{pure}\, x$	*local-pure*
	$\text{`local}\, f\, (m \gg\!= k) \equiv \text{`local}\, f\, m \gg\!= \text{`local}\, f \circ k$	*local-bind*
	$\text{`local f `ask} \equiv \text{pure} \circ f$	*local-ask*
	$\text{`local}\, (f \circ g)\, m \equiv \text{`local}\, g\, (\text{`local}\, f\, m)$	*local-local*
Catch	$\text{`catch}\, (\text{pure}\, x)\, m \equiv \text{pure}\, x$	*catch-pure*
	$\text{`catch `throw}\, m \equiv m$	$\textit{catch-throw}_1$
	$\text{`catch}\, m\, \text{`throw} \equiv m$	$\textit{catch-throw}_2$
Lambda	$\text{`abs}\, f \gg\!= \lambda\, f' \to \text{`app}\, f'\, m \equiv m \gg\!= f$	*beta*
	$\text{pure}\, f \equiv \text{`abs}\, (\lambda\, x \to \text{`app}\, f\, (\text{`var}\, x))$	*eta*

5.11 Discussion

In the introduction, we discussed the desired degrees of modularity that hefty algebras and their reasoning infrastructure should support. That is, they should support the modular composition of syntax, semantics, equational theories, and proofs.

Composability of the syntax and semantics of higher order effects follows readily from the fact that we define higher-order effect signatures and elaborations as higher-order functors and their algebras respectively (Section 3.1), which are closed under coproducts.

For proofs, we demonstrate in Section 5.3 that equational proofs for programs with higher-order effects are similarly modular. For instance, the proof of the catch-return-censor law can be reused to reason about larger programs even if they involve additional (higher-order) effects. Crucially, correctness proofs of elaborations also compose: combining two correct elaborations automatically yields a proof of correctness for the composite elaboration (Section 5.9). This demonstrates that hefty trees enjoy the same, if not more, modularity properties than algebraic effects, which require sophisticated reasoning about fusion to compose proofs of handler correctness (Yang & Wu, 2021).

Finally, we remark that elaborations and their correctness proofs support a more fine-grained composition than coproducts. Building on techniques developped by van der Rest *et al*. (2022), composition operators consume a separation witness, which act as a specification of which equations to identify across theories. This way, we avoid indiscriminate summation of effect theories. For more details we refer to the artifact accompanying this paper (van der Rest & Poulsen, 2024).

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

6 Related work

As stated in the introduction of this paper, defining abstractions for programming constructs with side effects is a research question with a long and rich history, which we briefly summarize here. Moggi (1989*a*) introduced monads as a means of modeling side effects and structuring programs with side effects; an idea which Wadler (1992) helped popularize. A problem with monads is that they do not naturally compose. A range of different solutions have been developed to address this issue (Cenciarelli & Moggi, 1993; Jones & Duponcheel, 1993; Steele, 1994; Filinski, 1999). Of these solutions, monad transformers (Cenciarelli & Moggi, 1993; Liang *et al*., 1995; Jaskelioff, 2008) is the more widely adopted solution. However, more recently, algebraic effects (Plotkin & Power, 2002) was proposed as an alternative solution which offers some modularity benefits over monads and monad transformers. In particular, whereas monads and monad transformers may "leak" information about the implementation of operations, algebraic effects enforce a strict separation between the interface and implementation of operations. Furthermore, monad transformers commonly require glue code to "lift" operations between layers of monad transformer stacks. While the latter problem is addressed by the Monatron framework of Jaskelioff (2008), algebraic effects have a simple composition semantics that does not require intricate liftings.

However, some effects, such as exception catching, did not fit into the framework of algebraic effects. *Effect handlers* (Plotkin & Pretnar, 2009) were introduced to address this problem. Algebraic effects and handlers has since been gaining traction as a framework for modeling and structuring programs with side effects in a modular way. Several libraries have been developed based on the idea such as *Handlers in Action* (Kammar *et al*., 2013), the freer monad (Kiselyov & Ishii, 2015), or Idris' `Effects` DSL (Brady, 2013*b*); but also standalone languages such as Eff (Bauer & Pretnar, 2015), Koka (Leijen, 2017), Frank (Lindley *et al*., 2017), and Effekt (Brachthäuser *et al*., 2020).[38]

As discussed in Sections 1.2 and 2.5, some modularity benefits of algebraic effects and handlers do not carry over to higher-order effects. Scoped effects and handlers (Wu *et al*., 2014; Piróg *et al*., 2018; Yang *et al*., 2022) address this shortcoming for *scoped operations*, as we summarized in Section 2.6. This paper provides a different solution to the modularity problem with higher-order effects. Our solution is to provide modular elaborations of higher-order effects into more primitive effects and handlers. We can, in theory, encode any effect in terms of algebraic effects and handlers. However, for some effects, the encodings may be complicated. While the complicated encodings are hidden behind a higher-order effect interface, complicated encodings may hinder understanding the operational semantics of higher-order effects, and may make it hard to verify algebraic laws about implementations of the interface. Our framework would also support elaborating higher-order effects into scoped effects and handlers, which might provide benefits for verification. We leave this as a question to explore in future work.

Existing languages for algebraic effects and handlers, such as Eff (Bauer & Pretnar, 2015), Frank (Lindley *et al*., 2017), Koka (Leijen, 2017), Effekt (Brachthäuser *et al*., 2020), or Flix (Lutze & Madsen, 2024) offer indirect support for higher-order effects, via

[38] A more extensive list of applications and frameworks can be found in Jeremy Yallop's Effects Bibliography: `https://github.com/yallop/effects-bibliography`.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

the encoding discussed in Section 1.2.2. As also discussed in Section 1.2.2, this encoding suffers from a modularity problem. Nevertheless, the encoding may suffice for applications in practice.

Whereas most languages (e.g., Eff, Koka, Flix, and Effekt) use so-called *deep handlers*, Frank (Lindley *et al*., 2017) uses *shallow handlers* (Hillerström & Lindley, 2018). The difference between shallow effect and deep effect handlers is in how continuations are typed. A deep handler of type $X \,!\, \Delta \Rightarrow C \,!\, \Delta'$ is typed as follows, where $op : A \to B$ is an operation of the effect row Δ:

$$\textbf{handler}\,\{\ \cdots\ (op\ \underbrace{v}_{A};\ \underbrace{k}_{B \to C \,!\, \Delta'}\ ) \mapsto\ \underbrace{c}_{C \,!\, \Delta'},\ \cdots\}$$

In contrast, shallow handlers are typed as follows:

$$\textbf{handler}\,\{\ \cdots\ (op\ \underbrace{v}_{A};\ \underbrace{k}_{B \to X \,!\, \Delta}\ ) \mapsto\ \underbrace{c}_{C \,!\, \Delta'},\ \cdots\}$$

Following Hillerström & Lindley (2018), shallow handlers can emulate deep handlers by always invoking their continuations in the scope of a recursive call to the handler being defined (assuming a language with recursive functions). Hillerström & Lindley (2018) also shows how deep handlers can emulate shallow handlers. As far as we are aware, shallow handlers support higher-order effects on a par with deep handlers, using the same encoding as we discussed in Section 1.2.2.

A recent paper by van den Berg *et al*. (2021) introduced a generalization of scoped effects that they call *latent effects*, which supports a broader class of effects, including λ abstraction. While the framework appears powerful, it currently lacks a denotational model, and seems to require similar weaving glue code as scoped effects. The solution we present in this paper does not require weaving glue code and is given by a modular but simple mapping onto algebraic effects and handlers.

Another recent paper by van den Berg & Schrijvers (2023) presents a unified framework for describing higher-order effects, which can be specialized to recover several instances such as Scoped Effects (Wu *et al*., 2014) or Latent Effects (van den Berg *et al*., 2021). They present a generic free monad generated from higher-order signatures that coincides with the type of Hefty trees that we present in Section 3. Their approach relies on a *Generalized Fold* (Bird & Paterson, 1999) for describing semantics of handling operations, in contrast to the approach in this paper, where we adopt a two-stage process of elaboration and handling that can be expressed using the standard folds of first-order and higher-order free monads. To explore how the use of generalized folds versus standard folds affects the relative expressivity of approaches to higher-order effects is a subject of further study.

The equational framework we present in Section 5 is inspired by the work of Yang & Wu (2021). Specifically, the notion of higher-order effect theory we formalized in Agda is an extension of the notion of (first-order) effect theory they use. In closely related recent work by Kidney *et al*. (2024), they present a formalization of first-order effect theories in *Cubical Agda* (Vezzosi *et al*., 2021). Whereas our formalization requires extrinsic verification of the equalities of an effect theory, they use *quotient types* as provided by homotopy type theory (Program, 2013) and cubical type theory (Angiuli *et al*., 2021; Cohen *et al*., 2017) to verify that handlers intrinsically respect their effect theories. They also present a Hoare

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

logic for verifying pre- and post-conditions. An interesting question for future work is whether this logic and the framework of Kidney *et al*. (2024) could be extended to higher-order effect theories.

In other recent work, Matache *et al*. (2025) developed an equational reasoning system for scoped effects. The work builds on previous work by Staton on *parameterized algebraic theories* (Staton, 2013a,b) which provide a syntactic framework for modeling computational effects with notions of locality (or, in scoped effects terminology, *scope*). Matache *et al*. (2025) show that scoped effects translate into a variant of parameterized algebraic theories and demonstrate that such theories provide algebraic characterizations of key examples from the literature on scoped effects: nondeterminism with semi-determinism, catching exceptions, and local state.

Whereas Matache *et al*. use parameterized algebraic theories as their underlying abstraction, Section 5 of this paper develops a notion of algebraic theory ($\mathsf{Theory}^{\mathsf{H}}$ in Section 5.7) over the *higher-order free monad*—i.e., a free monad construction that uses *higher-order functors*, given by a suitably generalized notion of container, instead of usual plain functors and containers (Abbott *et al*., 2005)—in Agda's Set. The equations of our higher-order effect theories are validated by elaborations into free ordinary effect theories. An interesting question for future work is to study the relationship between and compare the expressiveness of our proposed notion of higher-order effect theory and parameterized algebraic theories+scoped effects.

As discussed in the introduction, this paper explores a formal semantics for overloading-based definitions of higher-order effects. We formalized this semantics using an initial algebra semantics. An alternative approach would have been to use a so-called *final tagless* (Carette *et al*., 2009) encoding. That is, instead of declaring syntax as an inductive datatype, we declare it as a record type, and program against that record. A benefit of the final tagless approach is that we do not have to explicitly fold over syntax. The idea is to program against interfaces given by record types; e.g.:

```
record NumSymantics (Repr : Set → Set) : Set₁ where
  field num : ℕ → Repr ℕ

record LamSymantics (Repr : Set → Set) : Set₁ where
  field lam : (Repr A → Repr B) → Repr (A → B)
        app : Repr (A → B) → Repr A → Repr B

symantics-ex : ∀ {R} → NumSymantics R → LamSymantics R → R ℕ
symantics-ex n l = app (lam (λ x → x)) (num 42)
  where open NumSymantics n; open LamSymantics l
```

Using this final tagless encoding, the semantics of symantics-ex will be given by passing two concrete implementations of NumSymantics and LamSymantics. In contrast, with the initial algebra semantics approach we use in Section 5, we would define symantics-ex in terms of an inductive data type for app, lam, and num; and then give its semantics by folding algebras over the abstract syntax tree. A benefit of final tagless is that it tends to have a have a lower interpretive overhead (Carette *et al*., 2009), since it avoids the need to iterate over syntax trees. These benefits extend to effects (Devriese, 2019). On the other hand, the inductive data types of initial encodings support induction, whereas final tagless

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

encodings generally do not. We do not make extensive use of inductive reasoning in this paper, and we expect that it should be possible to port most of the definitions in our paper to use final tagless encodings. Our main reason for using an initial encoding for our hefty trees and algebras is that it follows the tradition of modeling algebraic effects and handlers using initial encodings, and that we expect induction to be useful for some applications.

Looking beyond purely functional models of semantics and effects, there are also lines of work on modular support for side effects in operational semantics (Plotkin, 2004). Mosses' Modular Structural Operational Semantics (Mosses, 2004) (MSOS) defines small-step rules that implicitly propagate an open-ended set of *auxiliary entities* which encode common classes of effects, such as reading or emitting data, stateful mutation, and even control effects (Sculthorpe *et al*., 2015). The K Framework (Rosu & Serbanuta, 2010) takes a different approach but provides many of the same benefits. These frameworks do not encapsulate operational details but instead make it notationally convenient to program (or specify semantics) with side-effects.

7 Conclusion

In this paper, we presented a semantics for higher-order effects based on *overloading*, by defining higher-order effects in terms of *elaborations* to algebraic effect trees. In this setup, we program against an interface of higher-order effects in a way that provides effect encapsulation. This means we can modularly change the implementation of effects without changing programs written against the interface, and without changing the definition of any interface implementations.

Crucially, hefty trees and their elaborations support modular reasoning. Equational proofs about programs with higher-order effects inherit this modularity: they can be reused in the context of larger programs, even if those rely on additional effects. Most significantly, correctness proofs of elaborations are themselves modular. As a result, correctness proofs can be lifted to proofs over composite elaborations, something which is generally not possible for algebraic effect handlers without appealing to fusion theorems.

While we have made use of Agda and dependent types throughout this paper, the framework should be portable to less dependently-typed functional languages, such as Haskell, OCaml, or Scala. An interesting direction for future work is to explore whether the framework could provide a denotational model for handling higher-order effects in standalone languages with support for effect handlers.

Acknowledgments

We thank the anonymous reviewers for their comments that helped improve the exposition of the paper. Furthermore, we thank Nicolas Wu, Andrew Tolmach, Peter Mosses, and Jaro Reinders for feedback on earlier drafts. This research was partially funded by the NWO VENI Composable and Safe-by-Construction Programming Language Definitions project (VI.Veni.192.259).

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

Data availability statement

The code and proofs supporting this paper are publically available (van der Rest & Bach Poulsen, 2025).

Competing interests

The authors report no conflict of interest.

References

Abbott, M. G., Altenkirch, T. & Ghani, N. (2003) Categories of containers. In 6th International Conference on Foundations of Software Science and Computational Structures, FOSSACS 2003 Held as Part of the Joint European Conference on Theory and Practice of Software, ETAPS 2003, Warsaw, Poland, April 7–11, 2003, Proceedings. Springer, pp. 23–38.

Abbott, M. G., Altenkirch, T. & Ghani, N. (2005) Containers: Constructing strictly positive types. *Theor. Comput. Sci.* **342**(1), 3–27.

Allais, G., Atkey, R., Chapman, J., McBride, C. & McKinna, J. (2021) A type- and scope-safe universe of syntaxes with binding: Their semantics and proofs. *J. Funct. Program.* **31**, e22.

Angiuli, C., Brunerie, G., Coquand, T., Harper, R., (Favonia), K. H. & Licata, D. R. (2021) Syntax and models of cartesian cubical type theory. *Math. Struct. Comput. Sci.* **31**(4), 424–468.

Arbib, M. A. & Manes, E. G. (1975) *Arrows, Structures, and Functors: The Categorical Imperative*. Academic Press.

Awodey, S. (2010) *Category Theory*, 2nd ed. USA: Oxford University Press, Inc.

Bauer, A. & Pretnar, M. (2015) Programming with algebraic effects and handlers. *J. Log. Algebraic Methods Program.* **84**(1), 108–123.

Biernacki, D., Piróg, M., Polesiuk, P. & Sieczkowski, F. (2018) Handle with care: Relational interpretation of algebraic effects and handlers. *Proc. ACM Program. Lang.* **2**(POPL), 8:1–8:30.

Bird, R. S. & Paterson, R. (1999) Generalised folds for nested datatypes. *Formal Aspects Comput.* **11**(2), 200–222.

Brachthäuser, J. I., Schuster, P. & Ostermann, K. (2020) Effects as capabilities: Effect handlers and lightweight effect polymorphism. *Proc. ACM Program. Lang.* **4**(OOPSLA), 126:1–126:30.

Brady, E. C. (2013a) Idris, a general-purpose dependently typed programming language: Design and implementation. *J. Funct. Program.* **23**(5), 552–593.

Brady, E. C. (2013b) Programming and reasoning with algebraic effects and dependent types. In ACM SIGPLAN International Conference on Functional Programming, ICFP'13, Boston, MA, USA - September 25–27, 2013, Morrisett, G. & Uustalu, T. (eds). ACM, pp. 133–144.

Carette, J., Kiselyov, O. & Shan, C. (2009) Finally tagless, partially evaluated: Tagless staged interpreters for simpler typed languages. *J. Funct. Program.* **19**(5), 509–543.

Castagna, G. & Gordon, A. D. (eds) (2017) Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18–20, 2017. ACM.

Cenciarelli, P. & Moggi, E. (1993) A syntactic approach to modularity in denotational semantics.

Claessen, K. (1999) A poor man's concurrency monad. *J. Funct. Program.* **9**(3), 313–323.

Cohen, C., Coquand, T., Huber, S. & Mörtberg, A. (2017) Cubical type theory: A constructive interpretation of the univalence axiom. *FLAP* **4**(10), 3127–3170.

Delaware, B., Keuchel, S., Schrijvers, T. & d. S. Oliveira, B. C. (2013) *Modular Monadic Meta-Theory*, pp. 319–330.

Devriese, D. (2019) Modular effects in haskell through effect polymorphism and explicit dictionary applications: A new approach and the μverifast verifier as a case study. In Proceedings of the 12th ACM SIGPLAN International Symposium on Haskell, Haskell@ICFP 2019, Berlin, Germany, August 18–23, 2019, Eisenberg, R. A. (ed.). ACM, pp. 1–14.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

Eisenberg, R. A. (ed.) (2019) Proceedings of the 12th ACM SIGPLAN International Symposium on Haskell, Haskell@ICFP 2019, Berlin, Germany, August 18–23, 2019. ACM.

Filinski, A. (1999) Representing layered monads. In POPL'99, Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Antonio, TX, USA, January 20–22, 1999. ACM, pp. 175–188.

Fiore, M. P. & Staton, S. (2014) Substitution, jumps, and algebraic effects. In Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS'14, Vienna, Austria, July 14–18, 2014. ACM, pp. 41:1–41:10.

Hancock, P. G. & Setzer, A. (2000) Interactive programs in dependent type theory. In Computer Science Logic, 14th Annual Conference of the EACSL, Fischbachau, Germany, August 21–26, 2000, Proceedings. Springer, pp. 317–331.

Hillerström, D. & Lindley, S. (2018) Shallow effect handlers. In Programming Languages and Systems – 16th Asian Symposium, APLAS 2018, Wellington, New Zealand, December 2–6, 2018, Proceedings. Springer, pp. 415–435.

Hyland, M., Plotkin, G. D. & Power, J. (2006) Combining effects: Sum and tensor. *Theor. Comput. Sci.* **357**(1-3), 70–99.

Jaskelioff, M. (2008) Monatron: An extensible monad transformer library. In Implementation and Application of Functional Languages - 20th International Symposium, IFL 2008, Hatfield, UK, September 10–12, 2008. Revised Selected Papers. Springer, pp. 233–248.

Jones, M. P. (1995) Functional programming with overloading and higher-order polymorphism. In Advanced Functional Programming, First International Spring School on Advanced Functional Programming Techniques, Båstad, Sweden, May 24–30, 1995, Tutorial Text. Springer, pp. 97–136.

Jones, M. P. & Duponcheel, L. (1993) *Composing Monads*. Research Report YALEU/DCS/RR-1004. Yale University. New Haven, Connecticut, USA.

Kammar, O., Lindley, S. & Oury, N. (2013) Handlers in action. In ACM SIGPLAN International Conference on Functional Programming, ICFP'13, Boston, MA, USA - September 25–27, 2013, Morrisett, G. & Uustalu, T. (eds). ACM, pp. 145–158.

Kidney, D. O., Yang, Z. & Wu, N. (2024) Algebraic effects meet Hoare logic in Cubical Agda. *Proc. ACM Program. Lang.* **8**(POPL), 1663–1695.

Kiselyov, O. & Ishii, H. (2015) Freer monads, more extensible effects. In Proceedings of the 8th ACM SIGPLAN Symposium on Haskell, Haskell 2015, Vancouver, BC, Canada, September 3–4, 2015. ACM, pp. 94–105.

Leijen, D. (2017) 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. ACM, pp. 486–499.

Levy, P. B. (2006) Call-by-push-value: Decomposing call-by-value and call-by-name. *High. Order Symb. Comput.* **19**(4), 377–414.

Liang, S., Hudak, P. & Jones, M. P. (1995) Monad transformers and modular interpreters. In Conference Record of POPL'95: 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Francisco, California, USA, January 23–25, 1995. ACM Press, pp. 333–343.

Lindley, S., Matache, C., Moss, S. K., Staton, S., Wu, N. & Yang, Z. (2024) Scoped effects as parameterized algebraic theories. In Programming Languages and Systems - 33rd European Symposium on Programming, ESOP 2024, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2024, Luxembourg City, Luxembourg, April 6–11, 2024, Proceedings, Part I. Springer, pp. 3–21.

Lindley, S., McBride, C. & McLaughlin, C. (2017) Do be do be do. In Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18–20, 2017, Castagna, G. & Gordon, A. D. (eds). ACM, pp. 500–514.

Lutze, M. & Madsen, M. (2024) Associated effects: Flexible abstractions for effectful programming. *Proc. ACM Program. Lang.* **8**(PLDI), 394–416.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

Martin-Löf, P. (1984) *Intuitionistic Type Theory*. Studies in Proof Theory, vol. 1. Bibliopolis.

Matache, C., Lindley, S., Moss, S. K., Staton, S., Wu, N. & Yang, Z. (2025) Scoped effects, scoped operations, and parameterized algebraic theories. *ACM Trans. Program. Lang. Syst.* **47**(2), 8:1–8:33.

Meijer, E., Fokkinga, M. M. & Paterson, R. (1991) Functional programming with bananas, lenses, envelopes and barbed wire. In Functional Programming Languages and Computer Architecture, 5th ACM Conference, Cambridge, MA, USA, August 26–30, 1991, Proceedings. Springer, pp. 124–144.

Moggi, E. (1989a) *An Abstract View of Programming Languages*. Technical Report ECS-LFCS-90-113. Edinburgh University, Department of Computer Science.

Moggi, E. (1989b) Computational lambda-calculus and monads. In Proceedings of the Fourth Annual Symposium on Logic in Computer Science (LICS'89), Pacific Grove, California, USA, June 5–8, 1989. IEEE Computer Society, pp. 14–23.

Morris, J. G. & McKinna, J. (2019) Abstracting extensible data types: or, rows by any other name. *Proc. ACM Program. Lang.* **3**(POPL), 12:1–12:28.

Morrisett, G. & Uustalu, T. (eds) (2013) ACM SIGPLAN International Conference on Functional Programming, ICFP'13, Boston, MA, USA - September 25–27, 2013. ACM.

Mosses, P. D. (2004) Modular structural operational semantics. *J. Log. Algebraic Methods Program.* **60-61**, 195–228.

Pfenning, F. & Elliott, C. (1988) Higher-order abstract syntax. In Proceedings of the ACM SIGPLAN'88 Conference on Programming Language Design and Implementation (PLDI), Atlanta, Georgia, USA, June 22–24, 1988. ACM, pp. 199–208.

Pierce, B. C. (1991) *Basic Category Theory for Computer Scientists*. Foundations of Computing. MIT Press.

Piróg, M. & Gibbons, J. (2014) The coinductive resumption monad. In Proceedings of the 30th Conference on the Mathematical Foundations of Programming Semantics, MFPS 2014, Ithaca, NY, USA, June 12–15, 2014. Elsevier, pp. 273–288.

Piróg, M., Schrijvers, T., Wu, N. & Jaskelioff, M. (2018) Syntax and semantics for operations with scopes. In Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09–12, 2018. ACM, pp. 809–818.

Plotkin, G. D. (2004) A structural approach to operational semantics. *J. Log. Algebraic Methods Program.* **60-61**, 17–139.

Plotkin, G. D. & Power, J. (2002) Notions of computation determine monads. In Foundations of Software Science and Computation Structures, 5th International Conference, FOSSACS 2002. Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2002 Grenoble, France, April 8–12, 2002, Proceedings. Springer, pp. 342–356.

Plotkin, G. D. & Power, J. (2003) Algebraic operations and generic effects. *Appl. Categorical Struct.* **11**(1), 69–94.

Plotkin, G. D. & Pretnar, M. (2009) Handlers of algebraic effects. In Programming Languages and Systems, 18th European Symposium on Programming, ESOP 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22–29, 2009. Proceedings. Springer, pp. 80–94.

Poulsen, C. B. & van der Rest, C. (2023) Hefty algebras: Modular elaboration of higher-order algebraic effects. *Proc. ACM Program. Lang.* **7**(POPL), 1801–1831.

Pretnar, M. (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. Elsevier, pp. 19–35.

Program, T. U. F. (2013) *Homotopy Type Theory: Univalent Foundations of Mathematics*. Institute for Advanced Study.

Rosu, G. & Serbanuta, T. (2010) An overview of the K semantic framework. *J. Log. Algebraic Methods Program.* **79**(6), 397–434.

Schmidt, D. (1986) *Denotational Semantics*. Allyn and Bacon.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

Schrijvers, T., Piróg, M., Wu, N. & Jaskelioff, M. (2019) Monad transformers and modular algebraic effects: What binds them together. In Proceedings of the 12th ACM SIGPLAN International Symposium on Haskell, Haskell@ICFP 2019, Berlin, Germany, August 18–23, 2019, Eisenberg, R. A. (ed.) ACM, pp. 98–113.

Schrijvers, T., Wu, N., Desouter, B. & Demoen, B. (2014) Heuristics entwined with handlers combined: From functional specification to logic programming implementation. In Proceedings of the 16th International Symposium on Principles and Practice of Declarative Programming, Kent, Canterbury, UK, September 8–10, 2014. ACM, pp. 259–270.

Sculthorpe, N., Torrini, P. & Mosses, P. D. (2015) A modular structural operational semantics for delimited continuations. In Proceedings of the Workshop on Continuations, WoC 2016, London, UK, April 12th 2015, pp. 63–80.

Staton, S. (2013a) An algebraic presentation of predicate logic - (extended abstract). In Foundations of Software Science and Computation Structures - 16th International Conference, FOSSACS 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16–24, 2013. Proceedings. Springer, pp. 401–417.

Staton, S. (2013b) Instances of computational effects: An algebraic perspective. In 28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2013, New Orleans, LA, USA, June 25–28, 2013. IEEE Computer Society, p. 519.

Steele Jr., G. L. (1994) Building interpreters by composing monads. In Conference Record of POPL'94: 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Portland, Oregon, USA, January 17–21, 1994. ACM Press, pp. 472–492.

Swierstra, W. (2008) Data types à la carte. *J. Funct. Program.* **18**(4), 423–436.

Taha, W. & Sheard, T. (2000) Metaml and multi-stage programming with explicit annotations. *Theor. Comput. Sci.* **248**(1-2), 211–242.

Thielecke, H. (1997) *Categorical Structure of Continuation Passing Style*. PhD thesis. University of Edinburgh.

van den Berg, B. & Schrijvers, T. (2023) A framework for higher-order effects & handlers. CoRR. abs/2302.01415.

van den Berg, B., Schrijvers, T., Poulsen, C. B. & Wu, N. (2021) Latent effects for reusable language components. In Programming Languages and Systems - 19th Asian Symposium, APLAS 2021, Chicago, IL, USA, October 17–18, 2021, Proceedings. Springer, pp. 182–201.

van der Rest, C. & Bach Poulsen, C. (2025) Supporting data. https://doi.org/10.5281/zenodo.17415938.

van der Rest, C. & Poulsen, C. B. (2024) GitHub - heft-lang/hefty-equations: Modular reasoning about (elaborations of) higher-order effects — github.com. https://github.com/heft-lang/hefty-equations.

van der Rest, C., Poulsen, C. B., Rouvoet, A., Visser, E. & Mosses, P. D. (2022) Intrinsically-typed definitional interpreters à la carte. *Proc. ACM Program. Lang.* **6**(OOPSLA2), 1903–1932.

Vezzosi, A., Mörtberg, A. & Abel, A. (2021) Cubical agda: A dependently typed programming language with univalence and higher inductive types. *J. Funct. Program.* **31**, e8.

Wadler, P. (1992) The essence of functional programming. In Conference Record of the Nineteenth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Albuquerque, New Mexico, USA, January 19–22, 1992. ACM Press, pp. 1–14.

Wadler, P., Kokke, W. & Siek, J. G. (2020) *Programming Language Foundations in Agda*.

Wu, N. & Schrijvers, T. (2015) Fusion for free - efficient algebraic effect handlers. In Mathematics of Program Construction - 12th International Conference, MPC 2015, Königswinter, Germany, June 29–July 1, 2015. Proceedings. Springer, pp. 302–322.

Wu, N., Schrijvers, T. & Hinze, R. (2014) Effect handlers in scope. In Proceedings of the 2014 ACM SIGPLAN Symposium on Haskell, Gothenburg, Sweden, September 4–5, 2014. ACM, pp. 1–12.

Yang, Z., Paviotti, M., Wu, N., van den Berg, B. & Schrijvers, T. (2022) Structured handling of scoped effects. In Programming Languages and Systems - 31st European Symposium on Programming, ESOP 2022, Held as Part of the European Joint Conferences on Theory and

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press

Practice of Software, ETAPS 2022, Munich, Germany, April 2–7, 2022, Proceedings. Springer, pp. 462–491.

Yang, Z. & Wu, N. (2021) Reasoning about effect interaction by fusion. *Proc. ACM Program. Lang.* **5**(ICFP), 1–29.

Zhang, Y. & Myers, A. C. (2019) Abstraction-safe effect handlers via tunneling. *Proc. ACM Program. Lang.* **3**(POPL), 5:1–5:29.

https://doi.org/10.1017/S0956796825100142 Published online by Cambridge University Press