On the Expressive Power of User-Defined Effects: Effect Handlers, Monadic Reflection, Delimited Control

YANNICK FORSTER, Saarland University, Germany and University of Cambridge, England

OHAD KAMMAR, University of Oxford, England and University of Cambridge, England

SAM LINDLEY, University of Edinburgh, Scotland

MATIJA PRETNAR, University of Ljubljana, Slovenia

We compare the expressive power of three programming abstractions for user-defined computational effects: Plotkin and Pretnar's effect handlers, Filinski's monadic reflection, and delimited control without answer-type-modification. This comparison allows a precise discussion about the relative expressiveness of each programming abstraction. It also demonstrates the sensitivity of the relative expressiveness of user-defined effects to seemingly orthogonal language features.

We present three calculi, one per abstraction, extending Levy's call-by-push-value. For each calculus, we present syntax, operational semantics, a natural type-and-effect system, and, for effect handlers and monadic reflection, a set-theoretic denotational semantics. We establish their basic metatheoretic properties: safety, termination, and, where applicable, soundness and adequacy. Using Felleisen's notion of a macro translation, we show that these abstractions can macro-express each other, and show which translations preserve typeability. We use the adequate finitary set-theoretic denotational semantics for the monadic calculus to show that effect handlers cannot be macro-expressed while preserving typeability either by monadic reflection or by delimited control. Our argument fails with simple changes to the type system such as polymorphism and inductive types. We supplement our development with a mechanised Abella formalisation.

CCS Concepts: • **Theory of computation** → **Control primitives; Functional constructs; Type structures; Denotational semantics; Operational semantics; Categorical semantics;**

Additional Key Words and Phrases: algebraic effects and handlers, monads, delimited control, computational effects, shift and reset, monadic reflection, reify and reflect, macro expressiveness, type-and-effect systems, denotational semantics, language extension, call-by-push-value, lambda calculus

ACM Reference Format:

Yannick Forster, Ohad Kammar, Sam Lindley, and Matija Pretnar. 2017. On the Expressive Power of User-Defined Effects: Effect Handlers, Monadic Reflection, Delimited Control. *Proc. ACM Program. Lang.* 1, ICFP, Article 13 (September 2017), 29 pages.

<https://doi.org/10.1145/3110257>

13

1 INTRODUCTION

How should we compare abstractions for user-defined effects?

The use of computational effects, such as file, terminal, and network I/O, random-number generation, and memory allocation and mutation, is controversial in functional programming. While languages like Scheme and ML allow these effects to occur everywhere, pure languages like Haskell restrict the use of effects. One reason to be wary of incorporating computational effects

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

© 2017 Copyright held by the owner/author(s).

2475-1421/2017/9-ART13

<https://doi.org/10.1145/3110257>

into a language is that doing so can mean giving up some of the most basic properties of the lambda calculus, like β -equality, referential transparency, and confluence. The loss of these properties leads to unpredictable behaviour in lazy languages, makes it harder to reason about program behaviour, and limits the applicability of correctness preserving transformations like common subexpression elimination or code motion.

Monads [Moggi 1989; Spivey 1990; Wadler 1990] are the established abstraction for incorporating effects into pure languages. Recently, Bauer and Pretnar [2015] proposed the use of *algebraic effects and handlers* [Plotkin and Pretnar 2009] to structure programs with user-defined effects. In this approach, the programmer first declares *algebraic operations* as the syntactic constructs she will use to cause the effects, in analogy with declaring new exceptions. Then, she defines *effect handlers* that describe how to handle these operations, in analogy with exception handlers. While exceptions immediately transfer control to the enclosing handler without resumption, a computation may continue in the same position following an effect operation. In order to support resumption, an effect handler has access to the *continuation* at the point of effect invocation. Thus algebraic effects and handlers provide a form of *delimited control*. Delimited control operators have long been used to encode effects [Danvy 2006]. There are many variants of such control operators, and their inter-relationships are subtle [Shan 2007], and often appear only in folklore. Here we focus on a specific pair of operators: *shift-zero* and *dollar* [Materzok and Biernacki 2012] without answer-type-modification, whose operational semantics and type system are the closest to effect handlers and monads.

We study the three different abstractions for user-defined effects: effect handlers, monads, and delimited control operators. Our goal is to enable language designers to conduct a precise and informed discussion about the relative expressiveness of each abstraction. In order to compare them, we build on an idealised calculus for functional-imperative programming, namely call-by-push-value [Levy 2004], and extend it with each of the three abstractions and their corresponding natural type systems. We then assess the expressive power of each abstraction by rigorously comparing and analysing these calculi.

We use Felleisen’s notion of macro expressibility [Felleisen 1991]: when a programming language $\mathcal{L}$ is extended by some feature, we say that the extended language $\mathcal{L}_+$ is *macro expressible* when there is a local syntax-directed translation (a *macro translation*) from $\mathcal{L}_+$ to $\mathcal{L}$ that keeps the features in $\mathcal{L}$ fixed. Felleisen introduces this notion to study the relative expressive power of Turing-complete calculi, as macro expressivity is more sensitive in these contexts than notions of expressivity based on computability. We adapt Felleisen’s approach to the situation where one extension $\mathcal{L}_+^1$ of a base calculus $\mathcal{L}$ is macro expressible in another extension $\mathcal{L}_+^2$ of the same base calculus $\mathcal{L}$. Doing so allows us to formally compare the expressive power of each of the different abstractions for user-defined effects.

In the first instance, we show that, disregarding types, all three abstractions are macro-expressible in terms of one another, giving six macro translations. Some of these translations are known in less rigorous forms, either published, or in folklore. One translation, macro-expressing effect-handlers in delimited control, improves on previous concrete implementations [Kammar et al. 2013], which rely on the existence of a global higher-order memory cell storing a stack of effect-handlers. The translation from monadic reflection to effect handlers is new.

We also examine whether these translations preserve typeability: the translations of some well-typed programs are untypeable. This untypeability is sensitive to the precise choice of features of the type system. We show that the translation from delimited control to monadic reflection preserves typeability. A potential difference between the expressive power of handler type systems and between monadic reflection and delimited control type systems was recently suggested by Kammar and Pretnar [2017], who give a straightforward typeability preserving macro-translation

Fig. 1. Existing and conjectured macro translations

of *delimited dynamic state* into a calculus of effect handlers, whereas existing translations using monads and delimited control require more sophistication [Kiselyov et al. 2006]. We show that there exists no macro translation from effect handlers to monadic reflection that preserves typeability. The proof relies on the denotational semantics for the monadic calculus. This set-theoretic denotational semantics and its adequacy for Filinski’s multi-monadic metalanguage [2010] is another piece of folklore which we formalise here. We conjecture that a similar proof, though with more mathematical sophistication, can be used to prove the non-existence of a typeability-preserving macro translation from the monadic calculus to effect handlers. To this end, we give adequate set-theoretic semantics to the effect handler calculus with its type-and-effect system, and highlight the critical semantic invariant a monadic calculus will invalidate.

Fig. 1 summarises our contributions and conjectured results. Untyped calculi appear on the left and their typed fragments on the right. Unlabelled arrows between the typed calculi signify that the corresponding macro translation between the untyped calculi preserves typeability. Arrows labelled by $*$ are new direct untyped direct translations. Arrows labelled by $\#$ signify that *no* macro translation exists between the calculi, not even a partial macro translation that is only defined for well-typed programs.

The non-expressivity results are sensitive to the precise collection of features in each calculus. For example, extending the base calculus with inductive types and primitive recursion would create gaps in our non-existence arguments, and we conjecture that extending the calculi with various forms of polymorphism would make our untyped translations typeability-preserving. Adding more features to each calculus blurs the distinction between each abstraction. This sensitivity means that in a realistic programming language, such as Haskell, OCaml, or Scheme, the different abstractions are often practically equivalent [Schrijvers et al. 2016]. It also teaches us that meaningful relative expressivity results *must* be stated within a rigorous framework such as a formal calculus, where the exact assumptions and features are made explicit. The full picture is still far from complete, but our work lays the foundation for drawing it.

We supplement our pencil-and-paper proofs with a mechanised formalisation in the Abella proof assistant [Gacek 2008, 2009] of the more syntactic aspects of our work. Specifically, for each calculus, we formalise a Wright and Felleisen style progress-and-preservation safety theorem [1994], and correctness theorems for our translations.

We make the following contributions:

- syntax and semantics of formal calculi for effect handlers, monadic reflection, and delimited control, where each calculus extends a shared call-by-push-value core, and their metatheory:
 - set-theoretic denotational semantics for effect handlers and monadic reflection;
 - denotational soundness and adequacy proofs for effect handlers and monadic reflection;
 - a termination proof for monadic reflection (termination proofs for the other calculi appear in existing work);

$V, W ::=$ values	$M, N ::=$	computations	return V	returner
x	variable	case V of	$x \leftarrow M; N$	sequencing
$()$	unit value	$(x_1, x_2) \rightarrow M$	$\lambda x. M$	abstraction
(V_1, V_2)	pairing	case V of $\{$	$M V$	application
$\text{inj}_\ell V$	variant	$(\text{inj}_{\ell_i} x_i \rightarrow M_i)_i \}$	$\langle M_1, M_2 \rangle$	pairing
$\{M\}$	thunk	$V!$	$\text{prj}_i M$	projection

Fig. 2. MAM syntax

- six macro-translations between the three untyped calculi, and variations on three of those translations;
- formally mechanised metatheory in Abella¹ comprising:
 - progress and preservation theorems;
 - the translations between the untyped calculi; and
 - their correctness proofs in terms of formal simulation results;
- typeability preservation of the macro translation from delimited control to monadic reflection; and
- a proof that there exists no typeability-preserving macro translation from effect handlers to either monadic reflection or delimited control.

We structure the remainder of the paper as follows. Sections 2–5 present the core calculus and its extensions with effect handlers, monadic reflection, and delimited control, in this order, along with their metatheoretic properties. Section 6 presents the macro translations between these calculi, their correctness, and typeability-preservation. Our positive translation results appear in § 6.1–6.6, which only depend on §1–4 of Section 2–5. Section 7 concludes and outlines further work.

2 THE CORE-CALCULUS: MAM

We are interested in a functional-imperative calculus where effects and higher-order features interact well. Levy’s call-by-push-value (CBPV) calculus fits the bill [2004]. It allows us to uniformly deal with call-by-value and call-by-name evaluation strategies, making the theoretical development relevant to both ML-like and Haskell-like languages. In CBPV evaluation order is explicit, and the way it combines computational effects with higher-order features yields simpler program logic reasoning principles [Plotkin and Pretnar 2008; Kammar and Plotkin 2012]. We extend it with a type-and-effect system. It is a variant of Kammar and Plotkin’s multi-adjunctive intermediate language [2012] without effect operations or coercions. We call the resulting calculus the *multi-adjunctive metalanguage* (MAM).

2.1 Syntax

Fig. 2 presents MAM’s raw term syntax, which distinguishes between values (data) and computations (programs). We assume a countable set of variables ranged over by $x, y, \dots$, and a countable set of variant constructor literals ranged over by ℓ . The unit value, products, and finite variants/sums are standard. A computation can be suspended as a thunk $\{M\}$, which may be passed around. Products and variants are eliminated with standard pattern matching constructs. Thunks can be forced to resume their execution. A computation may simply return a value, and two computations can be sequenced, as in Haskell’s `do` notation. A function computation abstracts over values to which it may be applied. In order to pass a function $\lambda x. M$ as data, it must first be suspended as a thunk

¹<https://github.com/matijapretnar/user-defined-effects-formalization>

Frames and contexts

$\mathcal{P} ::= x \leftarrow [] ; N \mid [] V \mid \text{prj}_i []$ pure frames
 $\mathcal{F} ::= \mathcal{P}$ computation frames
 $C ::= [] \mid C[\mathcal{F} []]$ evaluation context
 $\mathcal{H} ::= [] \mid \mathcal{H}[\mathcal{P} []]$ pure context

Reduction $\boxed{M \rightsquigarrow M'}$
 $M \rightsquigarrow_\beta M'$
 $\frac{M \rightsquigarrow_\beta M'}{C[M] \rightsquigarrow C[M']}$

Beta reduction $\boxed{M \rightsquigarrow_\beta M'}$

$(\times)$ $\text{case } (V_1, V_2) \text{ of } (x_1, x_2) \rightarrow M \rightsquigarrow_\beta M[V_1/x_1, V_2/x_2]$ (U) $\{M\}! \rightsquigarrow_\beta M$
 $(+)$ $\text{case inj}_\ell V \text{ of } \{ \dots \text{inj}_\ell x \rightarrow M \dots \} \rightsquigarrow_\beta M[V/x]$ ($\rightarrow$) $(\lambda x. M) V \rightsquigarrow_\beta M[V/x]$
 (F) $x \leftarrow \text{return } V ; M \rightsquigarrow_\beta M[V/x]$ (&) $\text{prj}_i \langle M_1, M_2 \rangle \rightsquigarrow_\beta M_i$

Fig. 3. MAM operational semantics

$\{\lambda x. M\}$. For completeness, we also include CBPV's binary computation products, which subsume projections from products in call-by-name languages.

Example 2.1. Using the boolean values $\text{inj}_{\text{True}} ()$ and $\text{inj}_{\text{False}} ()$, we define a logical *not* operation:

$$\text{not} = \{\lambda b. \text{case } b \text{ of } \{\text{inj}_{\text{True}} x \rightarrow \text{return } \text{inj}_{\text{False}} () \\ \text{inj}_{\text{False}} x \rightarrow \text{return } \text{inj}_{\text{True}} ()\}\}$$

2.2 Operational Semantics

Fig. 3 presents MAM's standard structural operational semantics, in the style of Felleisen and Friedman [1987]. In order to reuse the core definitions as much as possible, we refactor the semantics into β -reduction rules and a single congruence rule. As usual, a β -reduction reduces a matching pair of introduction and elimination forms.

We factor the definition of evaluation contexts through *computation frames*. In MAM these consist of *pure frames*, the elimination frames for pure computation. For each extension we will add another kind of effectful computation frame. We use $[]$ to denote the hole in each frame or context, which signifies which term should evaluate first, and define substitution frames and terms for holes ($C[\mathcal{F} []]$, $C[M]$) in the standard way. Later, in each calculus we will make use of *pure contexts* in order to capture continuations, stacks of pure frames, extending from a control operator to the nearest delimiter. A reducible term can be decomposed into at most one pair of evaluation context and β -reducible term, making the semantics deterministic.

Example 2.2. With this semantics we have $\text{not}! (\text{inj}_{\text{True}} ()) \rightsquigarrow^+ \text{return } \text{inj}_{\text{False}} ()$.

We use the following standard syntactic sugar. We use nested patterns in our pattern matching constructs. We abbreviate the variant constructors to their labels, and omit the unit value, e.g., True desugars to $\text{inj}_{\text{True}} ()$. We allow the application of functions and the elimination constructs to apply to arbitrary computations, and not just values, by setting for example $M N := x \leftarrow N ; M x$ for some fresh x , giving a more readable, albeit call-by-value, appearance.

Example 2.3. As a running example, we express boolean state in each of our calculi. Fig. 4(a) shows the code, which toggles the state and returns the value of the original state, as we would like to write it. Fig. 4(b) shows how we do so in MAM, via a standard state-passing transformation. We may then run *toggle* with the initial value True to get the expected result $\text{runState! toggle True} \rightsquigarrow^* (\text{True}, \text{False})$. This transformation is *not* a macro translation. In addition to the definition of *put* and *get*, it globally threads the state through *toggle*'s structure. Each user-defined effect abstraction in Sections 3–5 provides a different means for macro-expressing state.

2.3 Type-and-Effect System

Fig. 5 presents MAM's types and effects. As a core calculus for three calculi with different notions of effect, MAM is pure, and the only shared effect is the empty effect $\emptyset$.

We include a kind system, unneeded in traditional CBPV where a context-free distinction between values and computations forces types to be well-formed. The two points of difference from CBPV are the kind of effects, and the refinement of the computation kind by well-kinded effects E . The other available kinds are the standard value kind and a kind for well-formed environments (without type dependencies).

Our type system includes value-type variables (which in Section 4 we use for defining monads parametrically). The simple types, finite products and variants, are the standard CBPV value types. Thunk types are annotated with effect annotations. Computation types include returners FA , which are computations that return a value of type A , similar to the monadic type $\mathbf{Monad} \, m \Rightarrow m \, a$ in Haskell. Functions are computations and only take values as arguments. We include CBPV's computation products, which account for product elimination via projection in call-by-name languages.

To ensure well-kindedness of types, which may contain type variables, we use type environments in a list notation that denotes sets of type variables. Similarly, we use a list notation for value environments, which are functions from a finite set of variable names to the set of *value* types.

Example 2.4. The type of booleans **bit** is given by $\{\mathbf{inj}_{\text{False}} \, 1, \mathbf{inj}_{\text{True}} \, 1\}$.

Fig. 6 presents the kind and type systems. The only effect ($\emptyset$) is well-kinded. Type variables must appear in the current type environment, and they are always value types. The remaining value and computation types and environments have straightforward structural kinding conditions. Thunks of E -computations of type C require the type C to be well-kinded, which includes the side-condition that E is a well-kinded effect. This kind system has the property that each valid kinding judgement has a unique derivation. Value type judgements assert that a value term has a well-formed value type under a well-formed environment in some type variable environment.

The rules for simple types are straightforward. Observe how the effect annotation moves between the E -computation type judgement and the type of E -thunks. The side condition for computation

$toggle = \{ x \leftarrow get!;$	$get = \{ \lambda s. (s, s) \}$	$toggle = \{ \lambda s. (x, s) \leftarrow get! \, s;$
$y \leftarrow not! \, x;$	$put = \{ \lambda s'. \lambda _ . ((s'), s') \}$	$y \leftarrow not! \, x;$
$put! \, y;$	$runState = \lambda c. \lambda s. c! \, s$	$(_, s) \leftarrow put! \, y \, s;$
$x \}$		$(x, s) \}$
(a) Direct style		(b) State-passing style

Fig. 4. User-defined boolean state

$E ::=$	effects	$A, B ::=$	value types	$C, D ::=$	computation types
$\emptyset$	pure effect	α	type variable	FA	returners
$K ::=$	kinds	$ 1$	unit	$ A \rightarrow C$	functions
$ \mathbf{Eff}$	effects	$ A_1 \times A_2$	products	$ C_1 \& C_2$	products
$ \mathbf{Val}$	values	$ \{\mathbf{inj}_{\ell_i} A_i\}_i$	variants	Environments:	
$ \mathbf{Comp}_E$	computations	$ U_E C$	thunks	$\Theta ::= \alpha_1, \dots, \alpha_n$	
$ \mathbf{Context}$	environments			$\Gamma, \Delta ::= x_1 : A_1, \dots, x_n : A_n$	

Fig. 5. MAM kinds and types

Effect kinding $\boxed{\Theta \vdash_k E : \text{Eff}}$

Context kinding $\boxed{\Theta \vdash_k \Gamma : \text{Context}}$

$\Theta \vdash_k \emptyset : \text{Eff}$

$\frac{[\Theta \vdash_k \Gamma(x) : \text{Val}]_{x \in \text{Dom}(\Gamma)}}{\Theta \vdash_k \Gamma : \text{Context}}$

Value kinding $\boxed{\Theta \vdash_k A : \text{Val}}$

$\frac{\alpha \in \Theta}{\Theta \vdash_k \alpha : \text{Val}} \quad \frac{}{\Theta \vdash_k 1 : \text{Val}}$
 $\frac{[\Theta \vdash_k A_i : \text{Val}]_i}{\Theta \vdash_k \{\text{inj}_{\ell_i} A_i\}_i : \text{Val}}$

$\frac{\Theta \vdash_k A_1 : \text{Val} \quad \Theta \vdash_k A_2 : \text{Val}}{\Theta \vdash_k A_1 \times A_2 : \text{Val}}$
 $\frac{\Theta \vdash_k C : \text{Comp}_E}{\Theta \vdash_k U_E C : \text{Val}}$

Computation kinding $\boxed{\Theta \vdash_k C : \text{Comp}_E} \quad (\Theta \vdash_k E : \text{Eff})$

$\frac{\Theta \vdash_k A : \text{Val}}{\Theta \vdash_k FA : \text{Comp}_E} \quad \frac{\Theta \vdash_k A : \text{Val} \quad \Theta \vdash_k C : \text{Comp}_E}{\Theta \vdash_k A \rightarrow C : \text{Comp}_E} \quad \frac{\Theta \vdash_k C_1 : \text{Comp}_E \quad \Theta \vdash_k C_2 : \text{Comp}_E}{\Theta \vdash_k C_1 \& C_2 : \text{Comp}_E}$

Value typing $\boxed{\Theta; \Gamma \vdash V : A} \quad (\Theta \vdash_k \Gamma : \text{Context}, A : \text{Val})$

$\frac{(x : A) \in \Gamma}{\Theta; \Gamma \vdash x : A} \quad \frac{}{\Theta; \Gamma \vdash () : 1} \quad \frac{\Theta; \Gamma \vdash V_1 : A_1 \quad \Theta; \Gamma \vdash V_2 : A_2}{\Theta; \Gamma \vdash (V_1, V_2) : A_1 \times A_2}$
 $\frac{\Theta; \Gamma \vdash V : A_i}{\Theta; \Gamma \vdash \text{inj}_{\ell_i} V : \{\text{inj}_{\ell_i} A_i\}_i} \quad \frac{\Theta; \Gamma \vdash_E M : C}{\Theta; \Gamma \vdash \{M\} : U_E C}$

Computation typing $\boxed{\Theta; \Gamma \vdash_E M : C} \quad (\Theta \vdash_k \Gamma : \text{Context}, E : \text{Eff}, C : \text{Comp}_E)$

$\frac{\Theta; \Gamma \vdash V : A_1 \times A_2 \quad \Theta; \Gamma, x_1 : A_1, x_2 : A_2 \vdash_E M : C}{\Theta; \Gamma \vdash_E \text{case } V \text{ of } (x_1, x_2) \rightarrow M : C} \quad \frac{\Theta; \Gamma \vdash V : \{\text{inj}_{\ell_i} A_i\}_i \quad [\Theta; \Gamma, x_i : A_i \vdash_E M_i : C]_i}{\Theta; \Gamma \vdash_E \text{case } V \text{ of } \{\text{inj}_{\ell_i} x_i \rightarrow M_i\}_i : C}$
 $\frac{\Theta; \Gamma \vdash V : U_E C}{\Theta; \Gamma \vdash_E V! : C} \quad \frac{\Theta; \Gamma \vdash V : A}{\Theta; \Gamma \vdash_E \text{return } V : FA} \quad \frac{\Theta; \Gamma \vdash_E M : FA \quad \Theta; \Gamma, x : A \vdash_E N : C}{\Theta; \Gamma \vdash_E x \leftarrow M; N : C}$
 $\frac{\Theta; \Gamma, x : A \vdash_E M : C}{\Theta; \Gamma \vdash_E \lambda x. M : A \rightarrow C} \quad \frac{\Theta; \Gamma \vdash_E M : A \rightarrow C \quad \Theta; \Gamma \vdash V : A}{\Theta; \Gamma \vdash_E M V : C}$
 $\frac{\Theta; \Gamma \vdash_E M_1 : C_1 \quad \Theta; \Gamma \vdash_E M_2 : C_2}{\Theta; \Gamma \vdash_E \langle M_1, M_2 \rangle : C_1 \& C_2} \quad \frac{\Theta; \Gamma \vdash_E M : C_1 \& C_2}{\Theta; \Gamma \vdash_E \text{prj}_i M : C_i}$

Fig. 6. MAM kind and type system

type judgements asserts that a computation term has a well-formed E -computation type under a well-formed environment for some well-formed effect E under some type variable environment. The rules for variables, value and computation products, variants, and functions are straightforward. The rules for thunking and forcing ensure that the computation's effect annotation agrees with the effect annotation of the thunk. The rule for **return** allows us to return a value at any effect annotation, reflecting the fact that this is a *may*-effect system: the effect annotations track which effects may be caused, without prescribing that any effect *must* occur. The rule for sequencing reflects our

choice to omit any form of effect coercion, subeffecting, or effect polymorphism: the three effect annotations must agree. More sophisticated effect systems allow greater flexibility [Katsumata 2014]. We leave the precise treatment of such extensions to later work.

Example 2.5. The values from Fig. 4(b) have the following types:

$not : U_0(\text{bit} \rightarrow F\text{bit}) \quad get : U_0(\text{bit} \rightarrow F(\text{bit} \times \text{bit})) \quad put : U_0(\text{bit} \rightarrow \text{bit} \rightarrow F(\text{bit} \times \text{bit}))$
 $toggle : U_0(\text{bit} \rightarrow F(\text{bit} \times \text{bit})) \quad runState : U_0(U_0(\text{bit} \rightarrow F(\text{bit} \times \text{bit})) \rightarrow \text{bit} \rightarrow F(\text{bit} \times \text{bit}))$

2.4 Operational Metatheory

We establish the basic properties of MAM.

THEOREM 2.6 (MAM SAFETY). *Well-typed programs don't go wrong: for all closed MAM returners $\Theta; \vdash_\emptyset M : FA$, either $M \rightsquigarrow N$ for some $\Theta; \vdash_\emptyset N : FA$ or else $M = \text{return } V$ for some $\Theta; \vdash V : A$.*

The standard inductive progress-and-preservation proof is in the Abella formalisation.

We extend existing termination results for CBPV [Doczkal 2007; Doczkal and Schwinghammer 2009]. We say that a term M *diverges*, and write $M \rightsquigarrow^\infty$ if for every $n \in \mathbb{N}$ there exists some N such that $M \rightsquigarrow^n N$. We say that M *does not diverge* when $M \not\rightsquigarrow^\infty$.

THEOREM 2.7 (MAM TERMINATION). *There are no infinite reduction sequences: for all MAM terms $\Theta; \vdash_\emptyset M : FA$, we have $M \not\rightsquigarrow^\infty$, and there exists some unique $\Theta; \vdash V : A$ such that $M \rightsquigarrow^* \text{return } V$.*

The proof uses Tait's method [1967] to establish totality. Explicitly, we define a (unary) relational interpretation to types and establish a basic lemma. To interpret returners FA , we need a monadic lifting. We use the lifting from Hermida's thesis [1993], defined to contains the returners that reduce to a return value for all closed substitutions. The remainder of the proof is immediate as the semantics is deterministic.

We now define contextual equivalence of MAM terms. We define the subclass of *ground types*:

$$(\text{ground values}) \ G ::= 1 \mid G_1 \times G_2 \mid \{\text{inj}_{\ell_i} G_i\}_i$$

The standard next step is to define well-typed *program contexts* $X[\]$ – terms with zero, one, or more occurrences of a *hole*, denoted by $[\]$, not to be confused with evaluation contexts $C[\]$, which always contain exactly one hole. Defining program contexts and their type judgements directly is straightforward but tedious and lengthy, with four kinds of judgements, and so we take a different approach. Informally, given two computation terms M_1 and M_2 , in order to define their contextual equivalence, we need to quantify over the set of all the pairs of contexts plugged with M_1 and M_2 :

$$\Xi[M_1, M_2] := \{\langle X[M_1], X[M_2] \rangle \mid X[\] \text{ is a well-typed enclosing context}\}$$

Once we define this set, we do not need contexts, their type system, nor their semantics in the remainder of the development, and so we will define this set directly.

When defining this set, we need to know the form of the typing judgement, and so it will contain, apart from the two terms $\langle X[M_1], X[M_2] \rangle$, their shared environments, Θ' and Γ' , and shared type, A or C . When this shared type is a computation type C , we also need to know the effect annotation E . So this set will contain quintuples $\langle \Theta', \Gamma', V_1, V_2, A \rangle$ representing the simultaneous value judgement $\Theta'; \Gamma' \vdash V_1, V_2 : A$ and sextuples $\langle \Theta', \Gamma', E', N_1, N_2, C \rangle$ representing the simultaneous computation judgement $\Theta'; \Gamma' \vdash_{E'} N_1, N_2 : C$. Because values can contain computations through thunking, and vice versa through forcing, this set will contain both such quintuples and sextuples. Because the program context can introduce identifiers through function abstraction, we should also allow for environment extension.

Effects	Value types	
$\llbracket \emptyset \rrbracket_\theta := \langle \text{Id}, \text{id}, \text{id} \rangle$	$\llbracket \alpha \rrbracket_\theta := \theta(\alpha)$	$\llbracket A_1 \times A_2 \rrbracket_\theta := \llbracket A_1 \rrbracket_\theta \times \llbracket A_2 \rrbracket_\theta$
	$\llbracket 1 \rrbracket_\theta := \{\star\}$	$\llbracket \{\text{inj}_{\ell_i} A_i\}_i \rrbracket_\theta := \bigcup_i \{\ell_i\} \times \llbracket A_i \rrbracket_\theta$
		$\llbracket UEC \rrbracket_\theta := \llbracket C \rrbracket_\theta$
Computation types		
$\llbracket FA \rrbracket_\theta := F\llbracket A \rrbracket_\theta$	$\llbracket A \rightarrow C \rrbracket_\theta := \langle \llbracket C \rrbracket_\theta \mid \llbracket A \rrbracket_\theta, \lambda f_s. \lambda x. c(\text{fmap}(\lambda f. f(x)) f_s) \rangle$	
$\llbracket C_1 \& C_2 \rrbracket_\theta := \langle \llbracket C_1 \rrbracket_\theta \times \llbracket C_2 \rrbracket_\theta, \lambda c_s. \langle c_1(\text{fmap} \pi_1 c_s), c_2(\text{fmap} \pi_2 c_s) \rangle \rangle$		

Fig. 7. MAM denotational semantics for types

The full definition is as follows. We say that an environment Γ' *extends* an environment Γ , and write $\Gamma' \geq \Gamma$ if Γ' extends Γ as a partial function from identifiers to value types. Consider any two computations with the same type C_0 under the same environments Θ_0, G_0 , that is, $\Theta_0; \Gamma_0 \vdash_{E_0} M_1 : C_0$ and $\Theta_0; \Gamma_0 \vdash_{E_0} M_2 : C_0$. Define the set $\Xi[\Theta_0; \Gamma_0 \vdash_{E_0} M_1, M_2 : C_0]$ to be the smallest set of tuples $\langle \Theta', \Gamma', V_1, V_2, A \rangle$ and $\langle \Theta', \Gamma', E', N_1, N_2, C \rangle$ that is compatible with the typing rules and contains all the tuples $\langle \Theta, \Gamma, E_0, M_1, M_2, C_0 \rangle$, where $\Theta \supseteq \Theta_0$ and $\Gamma \geq \Gamma_0$. The compatibility with the rules means, for example, that if $\langle \Theta', \Gamma', V_1, V_2, A \rangle$ is in $\Xi[\Theta_0; \Gamma_0 \vdash_{E_0} M_1, M_2 : C_0]$, then so is $\langle \Theta', \Gamma', \emptyset, \text{return } V_1, \text{return } V_2, FA \rangle$. Define the set $\Xi[\Theta_0; \Gamma_0 \vdash V_1, V_2 : A]$ for contexts plugged with values analogously.

For uniformity's sake, we let types X range over both value and E -computation types, and phrases P range over both value and computation terms. Judgements of the form $\Theta; \Gamma \vdash_E P : X$ are meta judgements, ranging over value judgements $\Theta; \Gamma \vdash P : X$ and E -computation judgement $\Theta; \Gamma \vdash_E P : X$.

Let $\Theta; \Gamma \vdash_E P, Q : X$ be two MAM phrases. We say that P and Q are *contextually equivalent* and write $\Theta; E \vdash_\Gamma P \simeq Q : X$ when, for all pairs of plugged *closed ground-returner pure* contexts $\langle \emptyset, \emptyset, \emptyset, M_P, M_Q, FG \rangle$ in $\Xi[\Theta; \Gamma \vdash_E P, Q : X]$ and for all closed ground value terms $; \vdash V : G$, we have $M_P \rightsquigarrow^* \text{return } V$ if and only if $M_Q \rightsquigarrow^* \text{return } V$.

2.5 Denotational Semantics

MAM has a straightforward set-theoretic denotational semantics. Presenting the semantics for the core calculus will simplify our later presentation. To do so, we first recall the following established facts about monads, specialised and concretised to the set-theoretic setting.

A monad is a triple $\langle T, \text{return}, \gg \rangle$ where T assigns to each set X a set TX , return assigns to each set X a function $\text{return}^X : X \rightarrow TX$ and $\gg$ assigns to each function $f : X \rightarrow TY$ its *Kleisli extension*: a function $\gg f : TX \rightarrow TY$, and the three assignments satisfy the *monad laws*:

$$(\text{return } x) \gg f = f(x), \quad a \gg \text{return } x = a, \quad (a \gg f) \gg g = a \gg (\lambda x. (f x \gg g))$$

for all $f : X \rightarrow TY$, $x \in X$, $a \in TX$, and $g : Y \rightarrow TZ$. A *T-algebra* for a monad $\langle T, \text{return}, \gg \rangle$, following [Marmolejo and Wood \[2010\]](#), is a pair $C = \langle |C|, \gg^C \rangle$ where $|C|$ is a set, called the *carrier*, and $\gg^C$ assigns to every function $f : X \rightarrow |C|$ its *Kleisli extension* $\gg f : TX \rightarrow |C|$ satisfying:

$$(\text{return } x) \gg^C f = f(x), \quad (a \gg g) \gg^C f = a \gg^C (\lambda y. (gy \gg^C f))$$

for all $x \in X$, $f : X \rightarrow |C|$, $a \in TY$, and $g : Y \rightarrow TX$ of the appropriate types. For each set X , the pair $FX := \langle TX, \gg \rangle$ forms a *T-algebra* called the *free T-algebra over X*.

We parameterise MAM's semantics function $\llbracket \Theta \vdash_k E : \text{Eff} \rrbracket$ by an assignment θ of sets $\theta(\alpha)$ to each of the type variables α in Θ . Given such a type variable assignment θ , we assign to each

- effect: a monad $\llbracket \Theta \vdash_k E : \text{Eff} \rrbracket_\theta$, denoted by $\langle T_{\llbracket E \rrbracket_\theta}, \text{return}^{\llbracket E \rrbracket_\theta}, \gg^{\llbracket E \rrbracket_\theta} \rangle$;
- value type: a set $\llbracket \Theta \vdash_k A : \text{Val} \rrbracket_\theta$;

Value terms

$$\begin{aligned} \llbracket x \rrbracket_\theta(y) &:= \pi_x(y) & \llbracket () \rrbracket_\theta(y) &:= \star & \llbracket (V_1, V_2) \rrbracket_\theta(y) &:= \langle \llbracket V_1 \rrbracket_\theta(y), \llbracket V_2 \rrbracket_\theta(y) \rangle \\ \llbracket \text{inj}_\ell V \rrbracket_\theta(y) &:= \langle \ell, \llbracket V \rrbracket_\theta(y) \rangle & \llbracket \{M\} \rrbracket_\theta(y) &:= \llbracket M \rrbracket_\theta(y) \end{aligned}$$

Computation terms

$$\begin{aligned} \llbracket \text{case } V \text{ of } (x_1, x_2) \rightarrow M \rrbracket_\theta(y) &:= \llbracket M \rrbracket_\theta(y[x_1 \mapsto a_1, x_2 \mapsto a_2]) \quad \text{where } \llbracket V \rrbracket_\theta(y) = \langle a_1, a_2 \rangle \\ \llbracket \text{case } V \text{ of } \{\text{inj}_{\ell_i} x_i \rightarrow M_i\}_i \rrbracket_\theta(y) &:= \llbracket M_j \rrbracket_\theta(y[x_j \mapsto a_j]) \quad \text{where } \llbracket V \rrbracket_\theta(y) = \langle \ell_j, a_j \rangle \\ \llbracket V! \rrbracket_\theta(y) &:= \llbracket V \rrbracket_\theta(y) \\ \llbracket \text{return } V \rrbracket_\theta(y) &:= \text{return } (\llbracket V \rrbracket_\theta(y)) & \llbracket x \leftarrow M; N \rrbracket_\theta(y) &:= \llbracket M \rrbracket_\theta(y) \gg \lambda a. \llbracket N \rrbracket_\theta(y[x \mapsto a]) \\ \llbracket \lambda x. M \rrbracket_\theta(y) &:= \lambda a. \llbracket M \rrbracket_\theta(y[x \mapsto a]) & \llbracket M V \rrbracket_\theta(y) &:= (\llbracket M \rrbracket_\theta(y))(\llbracket V \rrbracket_\theta(y)) \\ \llbracket \langle M_1, M_2 \rangle \rrbracket_\theta(y) &:= \langle \llbracket M_1 \rrbracket_\theta(y), \llbracket M_2 \rrbracket_\theta(y) \rangle & \llbracket \text{prj}_i M \rrbracket_\theta(y) &:= \pi_i(\llbracket M \rrbracket_\theta(y)) \end{aligned}$$

Fig. 8. MAM denotational semantics for terms

- E -computation type: a $T_{[E]_\theta}$ -algebra $\llbracket \Theta \vdash_k C : \mathbf{Comp}_E \rrbracket_\theta$; and
- context: the set $\llbracket \Theta \vdash_k \Gamma : \mathbf{Context} \rrbracket_\theta := \prod_{x \in \text{Dom}(\Gamma)} \llbracket \Gamma(x) \rrbracket_\theta$.

Fig. 7 defines the standard set-theoretic semantics function over the structure of types. The pure effect denotes the identity monad, which sends each set to itself, and extends a function by doing nothing. The extended languages in the following sections will assign more sophisticated monads to other effects. The semantics of type variables uses the type assignment given as parameter. The unit type always denotes the singleton set. Product types and variants denote the corresponding set-theoretic operations of cartesian product and disjoint union, and thus the empty variant type $0 := \{\}$ denotes the empty set. The type of thunked E -computations of type C denotes the carrier of the $T_{[E]_\theta}$ -algebra $\llbracket C \rrbracket_\theta$. The E -computation type of A returners denotes the free $[E]_\theta$ -algebra. Function and product types denote well-known algebra structures over the sets of functions and pairs, respectively [Barr and Wells 1985, Theorem 4.2].

Terms can have multiple types, for example the function $\lambda x. \text{return } x$ has the types $1 \rightarrow F1$ and $0 \rightarrow F0$, and type judgements can have multiple type derivations. We thus give a Curry-style semantics [Reynolds 2009] by defining the semantic function for type judgement derivations rather than for terms. To increase readability, we write $\llbracket P \rrbracket$ and omit typing derivation for P .

The semantic function for terms is parameterised by an assignment θ of sets to type variables. It assigns to each well-typed derivation for a :

- value term: a function $\llbracket \Theta; \Gamma \vdash V : A \rrbracket_\theta : \llbracket \Gamma \rrbracket_\theta \rightarrow \llbracket A \rrbracket_\theta$; and
- E -computation term: a function $\llbracket \Theta; \Gamma \vdash_E M : C \rrbracket_\theta : \llbracket \Gamma \rrbracket_\theta \rightarrow \llbracket C \rrbracket_\theta$.

Fig. 8 defines the standard set-theoretic semantics over the structure of derivations. The semantics of sequencing uses the Kleisli extended function $(\gg^{[C]} f) : TX \rightarrow \llbracket C \rrbracket$ for functions into non-free algebras $f : X \rightarrow \llbracket C \rrbracket$, given by the algebra structure.

2.6 Denotational Metatheory

We develop the basic properties of our denotational semantics. Our goal is to establish the *adequacy* of the semantics: terms with equivalent denotations are observationally equivalent. In our set-theoretic setting, the proof-recipe is well-established, using the following compositionality and soundness theorems:

THEOREM 2.8 (MAM COMPOSITIONALITY). *The meaning of a term depends only on the meaning of its sub-terms: for all pairs of well-typed plugged MAM contexts M_P, M_Q in $\Xi[\Theta; \Gamma \vdash_E P, Q : X]$, if $\llbracket P \rrbracket = \llbracket Q \rrbracket$ then $\llbracket M_P \rrbracket = \llbracket M_Q \rrbracket$.*

The proof is a straightforward induction on the set of plugged contexts.

To phrase our simulation results in later development, we adopt a relaxed variant of simulation: let $\sim_{\text{cong}}$ be the smallest relation containing $\sim_{\beta}$ that is closed under the term formation constructs, and so contains $\sim$ as well, and let $\simeq_{\text{cong}}$ be the smallest congruence relation containing $\sim_{\beta}$.

THEOREM 2.9 (MAM SOUNDNESS). *Reduction preserves the semantics: for every pair of well-typed MAM terms $\Theta; \Gamma \vdash_E P, Q : X$, if $P \simeq_{\text{cong}} Q$ then $\llbracket P \rrbracket = \llbracket Q \rrbracket$. In particular, for every well-typed closed term of ground type $;\vdash_{\emptyset} P : FG$, if $P \sim^* \text{return } V$ then $\llbracket P \rrbracket = \llbracket V \rrbracket$.*

The proof is standard: check that $\sim_{\beta}$ preserves the semantics via calculation, and appeal to compositionality. It now follows that the semantics is adequate:

THEOREM 2.10 (MAM ADEQUACY). *Denotational equivalence implies contextual equivalence: for all well-typed MAM terms $\Theta; \Gamma \vdash_E P, Q : X$, if $\llbracket P \rrbracket = \llbracket Q \rrbracket$ then $P \simeq Q$.*

To see how the different pieces fit together, consider any two denotationally equivalent terms P and Q , a closed ground context plugged with them $\mathcal{X}[P]$, $\mathcal{X}[Q]$, and assume $\mathcal{X}[P] \sim^* \text{return } V$. By the Compositionality Theorem 2.8, $\mathcal{X}[P]$ and $\mathcal{X}[Q]$ have equal denotations. By the Safety Theorem 2.6 and Termination Theorem 2.7, $\mathcal{X}[Q] \sim^* \text{return } V'$ for some value $\text{return } V'$. And so by the Soundness Theorem 2.9:

$$\llbracket \text{return } V \rrbracket = \llbracket \mathcal{X}[P] \rrbracket = \llbracket \mathcal{X}[Q] \rrbracket = \llbracket \text{return } V' \rrbracket$$

Conclude by verifying that, for ground returners, denotational equality implies syntactic equality.

As a consequence, we deduce that our operational semantics is very well-behaved: for all well-typed computations $\Theta; \Gamma \vdash_E M, M' : C$, if $M \sim_{\text{cong}} M'$ then $M \simeq M'$.

3 EFFECT HANDLERS: EFF

Algebraic effects and handlers provide a basis for modular programming with user-defined effects [Kammar et al. 2013; Kiselyov et al. 2013; Bauer and Pretnar 2015; Hillerström and Lindley 2016; Leijen 2017; Lindley et al. 2017]. Programmable effect handlers arose as part of Plotkin and Power’s denotational theory of computational effects [2002], which investigates the consequences of using the additional structure in algebraic presentations of monadic models of effects. This account refines Moggi’s monadic account [1989] by incorporating into the theory the syntactic constructs that generate effects as *algebraic operations for a monad* [Plotkin and Power 2003]: each monad is accompanied by a collection of syntactic operations, whose interaction is specified by a collection of equations, i.e., an algebraic theory, which fully determines the monad. To fit exception handlers into this account, Plotkin and Pretnar [2009] generalised exception handlers to effect handlers, handling arbitrary algebraic effects and, following Levy’s CBPV, give a computational interpretation of algebras for a monad. By allowing the user to declare operations, effects can be described in a composable manner. Bauer and Pretnar [2015] demonstrate how, by defining algebras for the free monad with these operations, users can give the abstract operations different meanings, in similar fashion to Swierstra’s use of free monads [2008].

3.1 Syntax

Fig. 9(a) presents the extension EFF, Kammar et al.’s core calculus of effect handlers [2013]. We assume a countable set of elements of a separate syntactic class, called *operation names* and ranged over by op . For each operation name op , EFF’s operation call construct allows the programmer to invoke the effect associated with op by passing it a value as an argument. Operation names are the only interface to effects the language has. The handling construct allows the programmer to use a handler to interpret the operation calls of a given returner computation. As the given computation

$M, N ::= \dots$	computations	$H ::=$	handlers
$\text{op } V$	operation call	$\{\text{return } x \mapsto M\}$	return clause
$\text{handle } M \text{ with } H$	handling construct	$ H \uplus \{\text{op } p \ k \mapsto N\}$	operation clause

(a) Syntax extensions to Fig. 2

Frames and contexts $\mathcal{F} ::= \dots | \text{handle } [\] \text{ with } H$ computation frames

Beta reduction

- (ret) $\text{handle } (\text{return } V) \text{ with } H \rightsquigarrow_{\beta} H^{\text{return}}[V/x]$
 (op) $\text{handle } \mathcal{H}[\text{op } V] \text{ with } H \rightsquigarrow_{\beta} H^{\text{op}}[V/p, \{\lambda x. \text{handle } \mathcal{H}[\text{return } x] \text{ with } H\}/k]$

(b) Operational semantics extensions to Fig. 3

Fig. 9. EFF

$\text{toggle} = \{x \leftarrow \text{get } (); y \leftarrow \text{not! } x; \text{put } y; x\}$	$\text{toggle} : U_{\text{State}} \text{Fbit}$
$H_{ST} = \{\text{return } x \mapsto \lambda s. \text{return } x$ $\quad \text{get } _ \ k \mapsto \lambda s. k! \ s \ s$ $\quad \text{put } s' \ k \mapsto \lambda _ . k! \ () \ s'\}$	$H_{ST} : \text{bit}^{\text{State} \Rightarrow^0} \text{bit} \rightarrow \text{Fbit}$
$\text{runState} = \{\lambda c. \text{handle } c! \text{ with } H_{ST}\}$	$\text{State} = \{\text{get} : 1 \rightarrow \text{bit}, \text{put} : \text{bit} \rightarrow 1\} : \text{Eff}$ $\text{runState} : U_{\emptyset}((U_{\text{State}} \text{Fbit}) \rightarrow \text{bit} \rightarrow \text{Fbit})$

Fig. 10. User-defined boolean state in EFF

may call thunks returned by functions, the decision which handler will handle a given operation call is dynamic. Handlers are specified by two kinds of clauses. A *return clause* describes how to proceed when returning a value. An *operation clause* describes how to proceed when invoking an operation op . The body of an operation clause can access the value passed in the operation call using the first bound variable p , which is similar to the bounding occurrence of an exception variable when handling exceptions. But unlike exceptions, we expect arbitrary effects like reading from or writing to memory to resume. Therefore the body of an operation clause can also access the continuation k at the operation's calling point.

Example 3.1. The left column of Fig. 10 expresses user-defined boolean state in EFF. The handler H_{ST} is parameterised by the current state. When the computation terminates, we discard this state. When the program calls get , the handler returns the current state and leaves it unchanged. When the program calls put , the handler returns the unit value, and instates the newly given state.

3.2 Operational Semantics

Fig. 9(b) presents EFF's extension to MAM's operational semantics. Computation frames $\mathcal{F}$ now include the handling construct, whereas the pure frames $\mathcal{P}$ do not, allowing a handled computation to β -reduce under the handler. We add two β -reduction cases for the added construct. When the returner computation inside a handler is fully evaluated, the return clause proceeds with the return value. When the returner computation inside a handler needs to evaluate an operation call, the definition of pure contexts $\mathcal{H}$ ensures $\mathcal{H}$ is precisely the continuation of the operation call delimited by the handler. Put differently, it ensures that the handler in the root of the reduct is the closest handler to the operation call in the call stack. The operation clause corresponding to the operation called then proceeds with the supplied parameter and current continuation. Rewrapping the handler around this continuation ensures that all operation calls invoked in the continuation are handled in the same way.

Kinds and types

$E ::= \dots$ effects
 $|\ \{\text{op} : A \rightarrow B\} \uplus E$ arity assignment
 $K ::= \dots$ kinds
 $|\ \text{Handler}$ handlers
 $R ::= A \xRightarrow{E} E' C$ handler types

Computation typing

$(\text{op} : A \rightarrow B) \in E \quad \Theta; \Gamma \vdash V : A$
 $\frac{\Theta; \Gamma \vdash_E \text{op } V : FB \quad \Theta; \Gamma \vdash H : A \xRightarrow{E} E' C}{\Theta; \Gamma \vdash_E \text{handle } M \text{ with } H : C}$

Handler typing $\boxed{\Theta; \Gamma \vdash H : R} \quad (\Theta \vdash_k \Gamma : \text{Context}, R : \text{Handler})$

$\frac{\Theta; \Gamma, x : A \vdash_E M : C \quad [\Theta; \Gamma, p : A_i, k : U_E(B_i \rightarrow C) \vdash_E N_i : C]_i}{\Theta; \Gamma \vdash \{\text{return } x \mapsto M\} \uplus \{\text{op}_i \ p \ k \mapsto N_i\}_i : A^{\{\text{op}_i : A_i \rightarrow B_i\}_i} \xRightarrow{E} C}$

Effect kinding

$\Theta \vdash_k A : \text{Val}$
 $\Theta \vdash_k B : \text{Val}$
 $\frac{\text{op} \notin E \quad \Theta \vdash_k E : \text{Eff}}{\Theta \vdash_k \{\text{op} : A \rightarrow B\} \uplus E : \text{Eff}}$

Handler kinding

$\boxed{\Theta \vdash_k R : \text{Handler}}$

$\frac{\Theta \vdash_k A : \text{Val} \quad \Theta \vdash_k E, E' : \text{Eff} \quad \Theta \vdash_k C : \text{Comp}_{E'}}{\Theta \vdash_k A \xRightarrow{E} E' C : \text{Handler}}$

Fig. 11. EFF's kinding and typing (extending Fig. 5 and 6)

Example 3.2. With this semantics, the user-defined state from Fig. 10 behaves as expected:

$\text{runState! toggle True} \rightsquigarrow^* (\text{handle True with } H_{ST}) \text{ False} \rightsquigarrow^* \text{True}$

More generally, the handler H_{ST} expresses *dynamically scoped* state [Kammar and Pretnar 2017]. For additional handlers for state and other effects, see Pretnar's tutorial [2015].

3.3 Type-and-Effect System

Fig. 11 presents EFF's extension to the kind and type system. The effect annotations in EFF are functions from finite signatures, assigning to each operation name its parameter type A and its return type B . We add a new kind for handler types, which describes the kind and the returner type the handler can handle, and the kind and computation type of the handling clause.

In the kinding judgement for effects, the types in each operation's arity assignment must be value types. The kinding judgement for handlers requires all the types and effects involved to be well-kinded.

Computation type judgements now include two additional rules for each new computation construct. An operation call is well-typed when the parameter and return type agree with the arity assignment in the effect annotation. An instance of the handling construct is well-typed when the type and effect of the handled computation and the type-and-effect of the construct agree with the types and effects in the handler type. The set of handled operations must strictly agree with the set of operations in the effect annotation. The variable bound to the return value has the returner type in the handler type. In each operation clause, the bound parameter variable has the parameter type from the arity assignment for this operation, and the continuation variable's input type matches the return type in the operation's arity assignment. The overall type of all operation clauses agrees with the computation type of the handler. The second effect annotation on the handler type matches the effect annotations on the continuation and the body of the operation and return clauses.

Example 3.3. The boolean state terms are assigned the types given in the right column of Fig. 10.

3.4 Operational Metatheory

We follow MAM's development, formalising EFF's safety theorem in Abella:

THEOREM 3.4 (EFF SAFETY). *Well-typed programs don't go wrong: for all closed EFF returners $\Theta; \vdash_\emptyset M : FA$, either $M \rightsquigarrow N$ for some $\Theta; \vdash_\emptyset N : FA$ or else $M = \text{return } V$ for some $\Theta; \vdash V : A$.*

Using the monadic lifting from Kammar's thesis [2014], we obtain termination for EFF [Kammar et al. 2013]:

THEOREM 3.5 (EFF TERMINATION). *There are no infinite reduction sequences: for all EFF terms $; \vdash_\emptyset M : FA$, we have $M \rightsquigarrow^\infty$, and there exists a unique $; \vdash V : A$ such that $M \rightsquigarrow^* \text{return } V$.*

EFF shares MAM's ground types, and we define plugged contexts, $\simeq$ and $\simeq_{\text{cong}}$ as in MAM.

3.5 Denotational Semantics

We now give a set-theoretic denotational semantics for EFF. First, recall the following concepts in universal and categorical algebra. A *signature* Σ is a pair consisting of a set $|\Sigma|$ whose elements we call *operation symbols*, and a function arity_Σ from $|\Sigma|$ assigning to each operation symbol $\varphi \in |\Sigma|$ a (possibly infinite) set $\text{arity}(\varphi)$. We write $(\varphi : A) \in \Sigma$ when $\varphi \in |\Sigma|$ and $\text{arity}_\Sigma(\varphi) = A$. Given a signature Σ and a set X , we inductively form the set $T_\Sigma X$ of Σ -terms over X :

$$t ::= x \mid \varphi \langle t_a \rangle_{a \in A} \quad (x \in X, (\varphi : A) \in \Sigma)$$

The assignment T_Σ together with the following assignments form a monad

$$\text{return } x := x \quad t \gg f := t[f(x)/x]_{x \in X} \quad (f : X \rightarrow T_\Sigma Y)$$

The T_Σ -algebras $\langle C, \gg^C \rangle$ are in bijective correspondence with Σ -algebras on the same carrier. These are pairs $\langle C, \llbracket - \rrbracket \rangle$ where $\llbracket - \rrbracket$ assigns to each $(\varphi : A) \in \Sigma$ a function $\llbracket \varphi \rrbracket : C^A \rightarrow C$ from A -ary tuples of C elements to C . The bijection is given by setting $\gg^C f$ to be the Σ -homomorphic extension of $f : X \rightarrow |C|$ to $T_\Sigma X$.

EFF's denotational semantics is given by extending MAM's semantics as follows. Given a type variable assignment θ , we assign to each handler type a pair $\llbracket \Theta \vdash_k R : \text{Handler} \rrbracket_\theta = \langle C, f \rangle$ consisting of an algebra C and a function f into the carrier $|C|$ of this algebra.

Fig. 12 presents how EFF extends MAM's denotational semantics. Each effect E gives rise to a signature whose operation symbols are the operation names in E tagged by an element of the denotation of the corresponding parameter type. This signature gives rise to the monad E denotes. When $E = \emptyset$, the induced signature is empty, and gives rise to the identity monad, and so this semantic function extends MAM's semantics. Handlers of E -computations returning A -values using E' -computations of type C denote a pair. Its first component is an $\llbracket E \rrbracket_\theta$ -algebra structure over the carrier $\llbracket C \rrbracket_\theta$, which may have nothing to do with the $\llbracket E' \rrbracket_\theta$ -algebra structure $\llbracket C \rrbracket_\theta$ already possesses. The second component is a function from $\llbracket A \rrbracket_\theta$ to the carrier $\llbracket C \rrbracket_\theta$.

The denotation of $\text{op } V$ at effect type E , where $\text{op} : A \rightarrow B \in E$, is $\text{op}_{\llbracket V \rrbracket_\theta(Y)} \langle a \rangle_{a \in \llbracket B \rrbracket_\theta}$. The denotation of the handling construct uses the Kleisli extension of the second component in the denotation of the handler. The denotation of a handler term defines the T_Σ -algebras by defining a Σ -algebra for the associated signature Σ . The operation clause for op allows us to interpret each of the operation symbols associated to op . The denotation of the return clause gives the second component of the handler.

3.6 Denotational Metatheory

We repeat the recipe for proving adequacy.

Effects

Handler types

$\llbracket E \rrbracket_\theta := T_{\{\text{op}_p : \llbracket A \rrbracket_\theta \mid (\text{op} : A \rightarrow B) \in E, p \in \llbracket A \rrbracket_\theta\}}$ $\llbracket A \xRightarrow{E} C \rrbracket := \{\llbracket E \rrbracket\text{-algebras with carrier } \llbracket C \rrbracket \mid \times \llbracket C \rrbracket^{\llbracket A \rrbracket}\}$

Computation terms

$\llbracket \text{op } V \rrbracket_\theta (Y) := \text{op}_{\llbracket V \rrbracket_\theta Y} \langle \text{return } a \rangle_{a \in \llbracket B \rrbracket_\theta}$
 $\llbracket \text{handle } M \text{ with } H \rrbracket_\theta (Y) := \llbracket M \rrbracket_\theta (Y) \ggg f \quad \text{where } \llbracket H \rrbracket (Y) = \langle D, f : \llbracket A \rrbracket \rightarrow \llbracket C \rrbracket \rangle$

Handler terms

$\llbracket \{\text{return } x \mapsto M\} \uplus \{\text{op}_i p k \mapsto N_i\}_i \rrbracket_\theta (Y) := \langle D, f \rangle$ where D 's algebra structure and f given by:
 $\llbracket \text{op}_q \rrbracket_D \langle \xi_a \rangle_a := \llbracket N_{\text{op}} \rrbracket_\theta (Y[q/p, \langle \xi_a \rangle_a / k]) f(a) := \llbracket M \rrbracket_\theta (Y[a/x])$

Fig. 12. EFF denotational semantics (extending Fig. 7 and 8)

THEOREM 3.6 (EFF COMPOSITIONALITY). *The meaning of a term depends only on the meaning of its sub-terms: for all pairs of well-typed plugged EFF contexts M_P, M_Q in $\Xi[\Theta; \Gamma \vdash_E P, Q : X]$, if $\llbracket P \rrbracket = \llbracket Q \rrbracket$ then $\llbracket M_P \rrbracket = \llbracket M_Q \rrbracket$.*

The proof is identical to MAM, with two more cases for $\sim_\beta$.

THEOREM 3.7 (EFF SOUNDNESS). *Reduction preserves the semantics: for every pair of well-typed EFF terms $\Theta; \Gamma \vdash_E P, Q : X$, if $P \simeq_{\text{cong}} Q$ then $\llbracket P \rrbracket = \llbracket Q \rrbracket$. In particular, for every well-typed closed term of ground type $;\vdash_\emptyset P : FG$, if $P \sim^* \text{return } V$ then $\llbracket P \rrbracket = \llbracket V \rrbracket$.*

We combine the previous results, as with MAM:

THEOREM 3.8 (EFF ADEQUACY). *Denotational equivalence implies contextual equivalence: for all well-typed EFF terms $\Theta; \Gamma \vdash_E P, Q : X$, if $\llbracket P \rrbracket = \llbracket Q \rrbracket$ then $P \simeq Q$.*

Therefore, EFF also has a well-behaved operational semantics: for all well-typed computations $\Theta; \Gamma \vdash_E M, M' : C$, if $M \sim_{\text{cong}} M'$ then $M \simeq M'$.

4 MONADIC REFLECTION: mon

Moggi [1989] conceptualises computational effects as monads, which he uses to give a uniform denotational semantics for a wide range of different effects. Spivey [1990] and Wadler [1990] introduce programming abstractions based on monads, allowing new effects to be declared and used as if they are native. Examples include parsing [Hutton and Meijer 1998], backtracking and constraint solving [Schrijvers et al. 2013], and mechanised reasoning [Ziliani et al. 2015; Bulwahn et al. 2008]. Libraries now exist for monadic programming even in impure languages such as OCaml², Scheme³, and C++ [Sinkovics and Porkoláb 2013].

Languages that use monads as an abstraction for user-defined effects typically employ other mechanisms to support them—usually an overloading resolution mechanism, such as type-classes in Haskell and Coq, and functors/implicits in OCaml. As a consequence, such accounts do not study monads as an abstraction in their own right, and are intertwined with implementation details and concepts stemming from the added mechanism. Filinski's work on monadic reflection [1994; 1996; 1999; 2010] provides a more canonical abstraction for incorporating monads into a programming language. In his calculi, user-defined monads stand independently.

$M, N ::= \dots$	computations	Frames and contexts
$ \mu(N)$	reflect	$\mathcal{F} ::= \mathcal{P} \mid [\]^T$ computation frames
$ [N]^T$	reify	Beta reduction
$T ::=$	monads	for every $T = \mathbf{where} \{ \lambda x. N_u; \lambda y. \lambda f. N_b \} :$
$\mathbf{where} \{ \mathbf{return} \ x = M;$	return clause	(<i>ret</i>) $[\mathbf{return} \ V]^T \rightsquigarrow_{\beta} N_u[V/x]$
$\quad y \gg f = N \}$	bind clause	(<i>reflection</i>) $[\mathcal{H}[\mu(N)]]^T \rightsquigarrow_{\beta}$ $N_b[\{N\}/y, \{(\lambda x. [\mathcal{H}[\mathbf{return} \ x]]^T)/f\}]$
(a) Syntax (extending Fig. 2)		(b) Operational semantics (extending Fig. 3)

Fig. 13. MON

$toggle = \{x \leftarrow get!; y \leftarrow not! \ x; put! \ y; x\}$	$toggle : U_{State}Fbit$
$get = \{ \mu(\lambda s.(s, s)) \}$	$get : U_{State}Fbit$
$put = \{ \lambda s'. \mu(\lambda_.((), s')) \}$	$put : U_{State}(bit \rightarrow F1)$
$State = \mathbf{where} \{$	$\emptyset < \mathbf{instance} \ \mathbf{monad}$
$\quad \mathbf{return} \ x = \lambda s.(x, s);$	$(\alpha.bit \rightarrow F(\alpha \times bit)) \ State : \mathbf{Eff}$
$\quad f \gg k = \lambda s.(x, s') \leftarrow f \ s;$	
$\quad k! \ x \ s'\}$	
$runState = \{ \lambda c.[c!]^{State} \}$	$runState : U_{\emptyset}((U_{State}Fbit) \rightarrow bit \rightarrow F(bit \times bit))$

Fig. 14. User-defined boolean state in MON

4.1 Syntax

Fig. 13(a) presents MON's syntax. The $\mathbf{where} \{ \mathbf{return} \ x = N_u; y \gg f = N_b \}$ construct binds x in the term N_u and y and f in N_b . The term N_u describes the unit and the term N_b describes the Kleisli extension/bind operation. We elaborate on the choice of the keyword **where** when we describe MON's type system. Using monads, the programmer can write programs as if the new effect was native to the language. We call the mode of programming when the effect appears native the *opaque* view of the effect. In contrast, the *transparent* mode occurs when the code can access the implementation of the effect directly in terms of its defined monad. The *reflect* construct $\mu(N)$ allows the programmer to graft code executing in transparent mode into a block of code executing in opaque mode. The *reify* construct $[N]^T$ turns a block of opaque code into the result obtained by the implementation of the effect.

Example 4.1. The left column of Fig. 14 expresses user-defined boolean state in MON using the standard *State* monad. To express *get* and *put*, we reflect the concrete definition of the corresponding operations of the state monad. To run a computation, we use reification to get the monadic representation of the computation as a state transformer, and apply it to the initial state.

4.2 Operational Semantics

Fig. 13(b) describes the extension to the operational semantics. The *ret* transition uses the user-defined monadic return to reify a value. To explain the *reflection* transition, note that the pure context $\mathcal{H}$ captures the continuation at the point of reflection delimited by an enclosing reification, with an opaque view of the effect T . The reflected computation N views this effect transparently. By reifying $\mathcal{H}$, we can use the user-defined monadic bind to graft the two together.

²http://www.cas.mcmaster.ca/~carette/pa_monad/

³<http://okmij.org/ftp/Scheme/monad-in-Scheme.html>

Example 4.2. With this semantics we have $\text{runState! toggle True} \rightsquigarrow^* \text{return (True, False)}$.

The example we have given here fits with the way in which monadic reflection is often used, but is not as flexible as the effect handler version because *get* and *put* are concrete functions rather than abstract operations, which means we cannot abstract over how to interpret them. To write a version of *toggle* that can be interpreted in different ways is possible using monadic reflection but requires more sophistication.

4.3 Type-and-Effect System

Fig. 15 presents the natural extension to MAM's kind and type system for monadic reflection. Effects are a stack of monads. The empty effect is the identity monad. A monad T can be *layered* on top of an existing stack E by $E < \text{instance monad } (\alpha.C) \text{ where } \{\text{return } x = M; y \gg f = N\}$. The intention is that the type constructor $C[-/\alpha]$ has an associated monad structure given by the bodies of the return M and the bind N , and can use effects from the rest of the stack E . To be well-kinded, C must be an E -computation, and T must be a well-typed monad: return should be typed $C[A/\alpha]$ when substituted for some value $V : A$, and $\gg$ typed as a Kleisli extension operation.

Example 4.3. The boolean state terms are assigned the types given in the right column of Fig. 14.

The choice of keywords for monads and their types follows their syntax in Haskell. We stress that our calculus does not, however, include a type-class mechanism. The *type* of a monad contains the return and bind *terms*, which means that we must check for equality of terms during type-checking, for example, to ensure that we are sequencing two computations with compatible effect annotations (for our purposes α -equivalence suffices). The need to check equality of terms arises from our choice of structural, anonymous, monads—in Haskell monads are given *nominally*, and two monads are compatible if they have exactly the same name. As our monads are structural, the bodies of the return and the bind operations must be closed, apart from their immediate arguments. If layered monad definitions were allowed to contain open terms, types in type contexts would contain these open terms through the effect annotations in thunks, requiring us to support dependently-typed contexts. The monad abstraction is parametric, so naturally requires the use of type variables, and for this reason we include type variables in the base calculus MAM. We choose monads to be structural and closed primarily in order to keep them closer to the other abstractions.

Kinds and types

$E ::= \dots$ effects
 $| E < \text{instance monad } (\alpha.C) T$ layered monad

Effect kinding

$$\frac{\Theta, \alpha \vdash_k C : \text{Comp}_E \quad \vdash_m T : E < \text{instance monad } (\alpha.C) T}{\Theta \vdash_k E < \text{instance monad } (\alpha.C) T : \text{Eff}}$$

Monad typing

$$\frac{\Theta, \alpha; x : \alpha \vdash_E N_u : C \quad \Theta, \alpha, \beta; y : U_E C, f : U_E(\alpha \rightarrow C[\beta/\alpha]) \vdash_E N_b : C[\beta/\alpha]}{\Theta \vdash_m \text{where } \{\text{return } x = N_u; y \gg f = N_b\} : E < \text{instance monad } (\alpha.C) \text{ where } \{\text{return } x = N_u; y \gg f = N_b\}}$$

Computation typing

$$\frac{\Theta; \Gamma \vdash_E N : C[A/\alpha]}{\Theta; \Gamma \vdash_{E < \text{instance monad } (\alpha.C) T} \mu(N) : FA} \quad \frac{\Theta \vdash_m T : E < \text{instance monad } (\alpha.C) T \quad \Theta; \Gamma \vdash_{E < \text{instance monad } (\alpha.C) T} N : FA}{\Theta; \Gamma \vdash_E [N]^T : C[A/\alpha]}$$

Fig. 15. MON's kinding and typing (extending Fig. 5 and 6)

Effects

$\llbracket E < \text{instance monad } (\alpha.C) N_u N_b \rrbracket_\theta := \langle T, \text{return}, \gg \rangle$
 where $TX := \left\llbracket C \right\rrbracket_{(\theta[\alpha \mapsto X])}$ $\text{return}^X := \llbracket N_u \rrbracket_{(\theta[\alpha \mapsto X])} : X \rightarrow TX$
 $\gg^{X,Y} := \llbracket N_b \rrbracket_{(\theta[\alpha_1 \mapsto X, \alpha_2 \mapsto Y])} : TX \rightarrow (X \rightarrow TY) \rightarrow TY$
 (provided these form a monad)

Monads

$\llbracket \Theta \vdash_m T : E \rrbracket := \llbracket E \rrbracket$

Computation terms

$\llbracket [N]^T \rrbracket(\gamma) := \llbracket N \rrbracket(\gamma)$
 $\llbracket \mu(N) \rrbracket(\gamma) := \llbracket N \rrbracket(\gamma)$

Fig. 16. MON denotational semantics (extending Fig. 7 and 8)

Our calculus differs from Filinski's [2010] in that our effect definitions are local and structural, whereas his allow nominal declarations of new effects only at the top level. Because we do not allow the bodies of the return and the bind to contain open terms, this distinction between the two calculi is minor. As a consequence, effect definitions in both calculi are *static*, and the monadic bindings can be resolved at compile time. Filinski's calculus also includes a sophisticated *effect-basing* mechanism, that allows a computation to immediately use, via reflection, effects from any layer in the hierarchy below it, whereas our calculus only allows reflecting effects from the layer immediately below. However, effect-basing does not significantly change the expressiveness of the calculus: the monad stack is statically known, and, having access to the type information, we can insert multiple reflection operators and lift effects from lower levels into the current level.

4.4 Operational Metatheory

We prove MON's Felleisen-Wright safety in our Abella formalisation:

THEOREM 4.4 (MON SAFETY). *Well-typed programs don't go wrong: for all closed MON returners $\Theta; \vdash_\emptyset M : FA$, either $M \rightsquigarrow N$ for some $\Theta; \vdash_\emptyset N : FA$ or else $M = \text{return } V$ for some $\Theta; \vdash V : A$.*

As with EFF, MON's ground types are the same as MAM's. While we can define an observational equivalence relation in the same way as for MAM and EFF, we will not do so. Monads as a programming abstraction have a well-known conceptual complication — user-defined monads must obey the *monad laws*. These laws are a syntactic counterpart to the three equations in the definition of (set-theoretic/categorical) monads. The difficulty involves deciding what equality between such terms means. The natural candidate is observational equivalence, but as the contexts can themselves define additional monads, it is not straightforward to do so. Giving an acceptable operational interpretation to the monad laws is an open problem. We avoid the issue by giving a *partial* denotational semantics to MON.

4.5 Denotational Semantics

We extend MAM's denotational semantics to MON as follows. Given a type variable assignment θ , we assign to each monad type and effect a monad $\llbracket \Theta \vdash_m T : E \rrbracket_\theta = \llbracket \Theta \vdash_k E : \text{Eff} \rrbracket_\theta$, if the sub-derivations have well-defined denotations, and this data does indeed form a set-theoretic monad. Consequently, the denotation of any derivation is undefined if at least one of its sub-derivations has undefined semantics. Moreover, the definition of kinding judgement denotations now depends on term denotations.

Fig. 16 extends the denotational semantics of MAM to MON. The denotation of the layered monad construct is only well-defined if the user-defined type constructor, return, and bind, form a monad. For the denotation of computation terms, recall that $T_{\llbracket E < \text{instance monad } (\alpha.C) T \rrbracket} X = \left\llbracket C \right\rrbracket_{(\theta[\alpha \mapsto X])}$ and therefore, semantically, we can view any computation of type FA subject to the kinding judgement $\Theta \vdash_k FA : \text{Comp}_{E < \text{instance monad } (\alpha.C) T}$ as an E -computation of type $C[A/\alpha]$.

Compare this semantics with Filinski's original semantics [1994], in which $\llbracket \mu(N) \rrbracket = \llbracket N \rrbracket \gg \text{id}$, $\llbracket [N]^T \rrbracket = \text{return}^T \llbracket N \rrbracket$. To explain the difference, bear in mind that our calculus is based on CBPV, whereas Filinski's original calculus is based on a pure λ -calculus. Specifically, Filinski interprets the judgement $M : A$ as $M : TA$. The corresponding judgement for us is $M : FA$. The semantics of the pure λ -calculus does not insert monadic returns and binds in the appropriate places, and so Filinski's translation inserts them explicitly. In contrast, CBPV inserts returns and binds (and if the term is pure, they cancel out), and so MON's semantics need not add them.

4.6 Denotational Metatheory

We define a *proper derivation* to be a derivation whose semantics is well-defined for all type variable assignments, and a *proper term or type* to be a term or type that has a proper derivation. Thus, a term is proper when all the syntactic monads it contains denote semantic set-theoretic monads. When dealing with the typed fragment of MON, we restrict our attention to such proper terms as they reflect the intended meaning of monads. Doing so allows us to mirror the metatheory of MAM and EFF for proper terms.

We define *plugged proper contexts* as with MAM and EFF with the additional requirement that all terms are proper. The definitions of the equivalences $\simeq$ and $\simeq_{\text{cong}}$ are then identical to those of MAM and EFF.

THEOREM 4.5 (MON TERMINATION). *There are no infinite reduction sequences: for all proper MON terms $\vdash_{\emptyset} M : FA$, we have $M \not\rightarrow^{\infty}$, and there exists some unique $\vdash V : A$ such that $M \rightarrow^* \text{return } V$.*

Our proof uses Lindley and Stark's $\top\top$ -lifting [2005].

THEOREM 4.6 (MON COMPOSITIONALITY). *The semantics depends only on the semantics of sub-terms: for all pairs of well-typed plugged proper MON contexts M_P, M_Q in $\Xi[\Theta; \Gamma \vdash_E P, Q : X]$, if $\llbracket P \rrbracket = \llbracket Q \rrbracket$ then $\llbracket M_P \rrbracket = \llbracket M_Q \rrbracket$.*

The proof is identical to MAM, with two more cases for $\rightarrow_{\beta}$. Similarly, we have:

THEOREM 4.7 (MON SOUNDNESS). *Reduction preserves the semantics: for every pair of well-typed proper MON terms $\Theta; \Gamma \vdash_E P, Q : X$, if $P \simeq_{\text{cong}} Q$ then $\llbracket P \rrbracket = \llbracket Q \rrbracket$. In particular, for every well-typed proper closed term of ground type $\vdash_{\emptyset} P : FG$, if $P \rightarrow^* \text{return } V$ then $\llbracket P \rrbracket = \llbracket V \rrbracket$.*

We combine the previous results, as with MAM and EFF:

THEOREM 4.8 (MON ADEQUACY). *Denotational equivalence implies contextual equivalence: for all well-typed proper MON terms $\Theta; \Gamma \vdash_E P, Q : X$, if $\llbracket P \rrbracket = \llbracket Q \rrbracket$ then $P \simeq Q$.*

Therefore, the *proper* fragment of MON also has a well-behaved operational semantics: for all well-typed proper computations $\Theta; \Gamma \vdash_E M, M' : C$, if $M \rightarrow_{\text{cong}} M'$ then $M \simeq M'$.

In contrast to EFF the semantics for MON is finite:

LEMMA 4.9 (FINITE DENOTATION PROPERTY). *For every type variable assignment $\theta = \langle X_{\alpha} \rangle_{\alpha \in \Theta}$ of finite sets, every proper MON value type $\Theta \vdash_k A :$ and computation type $\Theta \vdash_k C :$ denote finite sets $\llbracket A \rrbracket_{\theta}$ and $\llbracket C \rrbracket_{\theta}$.*

5 DELIMITED CONTROL: DEL

Control operators have a long history of expressing both user-defined effects [Danvy 2006] and algorithms with sophisticated control flow [Felleisen et al. 1988] such as tree-fringe comparison, and other control mechanisms, such as coroutines. The delimited operators enjoy an improved metatheory in comparison with their undelimited counterparts [Felleisen et al. 1988]. The operator closest in spirit to handlers is S_0 , pronounced "shift zero". It was introduced by Danvy and Filinski [1990] as part of a systematic study of continuation-passing-style conversion.

$M, N ::= \dots$	computations	Frames and contexts
$ S_0 k.M$	shift-0	$\mathcal{F} ::= \dots \langle [\] x.N \rangle$ computation frames
$ \langle M x.N \rangle$	reset	Beta reduction
		(<i>ret</i>) $\langle (\text{return } V) x.M \rangle \rightarrow_\beta M[V/x]$
		(<i>capture</i>) $\langle \mathcal{H}[S_0 k.M] x.N \rangle \rightarrow_\beta$ $M[\lambda y. \langle \mathcal{H}[\text{return } y] x.N \rangle / k]$
(a) Syntax (extending Fig. 2)		(b) Operational semantics (extending Fig. 3)

Fig. 17. DEL

5.1 Syntax

Fig. 17(a) presents the extension DEL. The construct $S_0 k.M$, which we abbreviate to “shift”, captures the current continuation and binds it to k , and replaces it with M . The construct $\langle M | x.N \rangle$, which we will call “reset”, delimits any continuations captured by shift inside M . Once M runs its course and returns a value, this value is bound to x and N executes. For delimited control cognoscenti this construct is sometimes called “dollar”, and can macro express the entire CPS hierarchy [Materzok and Biernacki 2012; Kiselyov and Shan 2007].

Example 5.1. The left column of Fig. 18 expresses user-defined boolean state in DEL [Danvy 2006, Section 1.4]. The code assumes the environment outside the closest reset will apply it to the currently stored state. By shifting and abstracting over this state, *get* and *put* can access this state and return the appropriate result to the continuation. When running a stateful computation, we discard the state when we reach the final return value.

5.2 Operational Semantics

The extension to the operational semantics in Fig. 17(b) reflects our informal description. The *ret* rule states that once the delimited computation returns a value, this value is substituted in the remainder of the reset computation. For the *capture* rule, the definition of pure contexts guarantees that in the reduct $\langle \mathcal{H}[S_0 k.M] | x.N \rangle$ there are no intervening resets in $\mathcal{H}$, and as a consequence $\mathcal{H}$ is the delimited continuation of the evaluated shift. After the reduction takes place, the continuation is re-wrapped with the reset, while the body of the shift has access to the enclosing continuation. If we were to, instead, not re-wrap the continuation with a reset, we would obtain the control/prompt-zero operators, (cf. Shan’s [2007] and Kiselyov et al.’s [2005] analyses of macro expressivity relationships between these two, and other, variations on untyped delimited control).

Example 5.2. We have: $\text{runState! toggle True} \leadsto^* \langle \text{True} | x.\lambda s.x \rangle \text{ False} \leadsto^* \text{return True}$.

$\text{toggle} = \{x \leftarrow \text{get!}; y \leftarrow \text{not! } x; \text{put! } y; x\}$	$\text{toggle} : U_{\text{State}} \text{Fbit}$
$\text{get} = \{S_0 k.\lambda s.k! s\}$	$\text{get} : U_{\text{State}} \text{Fbit}$
$\text{put} = \{\lambda s'.S_0 k.\lambda_.k! () s'\}$	$\text{put} : U_{\text{State}} (\text{bit} \rightarrow \text{F1})$
$\text{runState} = \{\lambda c. \langle c! x.\lambda s.x \rangle\}$	$\text{runState} : U_0((U_{\text{State}} \text{Fbit}) \rightarrow \text{bit} \rightarrow \text{Fbit})$
	$\text{State} = \emptyset, \text{bit} \rightarrow \text{Fbit} : \text{Eff}$

Fig. 18. User-defined boolean state in DEL

Kinds and types

$E ::= \dots$ effects
 $| E, C$ enclosing continuation type

Effect kinding

$$\frac{\Theta \vdash_k E : \mathbf{Eff} \quad \Theta \vdash_k C : \mathbf{Comp}_E}{\Theta \vdash_k E, C : \mathbf{Eff}}$$

Computation typing

$$\frac{\Theta; \Gamma, k : U_E(A \rightarrow C) \vdash_E M : C}{\Theta; \Gamma \vdash_{E,C} S_0 k.M : FA} \quad \frac{\Theta; \Gamma \vdash_{E,C} M : FA \quad \Theta; \Gamma, x : A \vdash_E N : C}{\Theta; \Gamma \vdash_E \langle M | x.N \rangle : C}$$

Fig. 19. DEL's kinding and typing (extending Fig. 5 and 6)

5.3 Type-and-Effect System

Fig. 19 presents the natural extension to MAM's kind and type system for delimited control. It is based on Danvy and Filinski's description [Danvy and Filinski 1989]; they were the first to propose a type system for delimited control. Effects are now a stack of computation types, with the empty effect standing for the empty stack. The top of this stack is the return type of the currently delimited continuation. Thus, as Fig. 19 presents, a shift pops the top-most type off this stack and uses it to type the current continuation, and a reset pushes the type of the delimited return typed onto it.

Example 5.3. The boolean state terms are assigned the types given in the right column of Fig. 18.

In this type system, the return type of the continuation remains fixed inside every reset. Existing work on type systems for delimited control (Kiselyov and Shan [2007] provide a substantial list of references) focuses on type systems that allow *answer-type modification*, as these can express typed printf and type-state computation (as in Asai's analysis [2009]). We exclude answer-type modification to keep the fundamental account clearer and simpler: the type system with answer-type modification is further removed from the well-known abstractions for effect-handlers and monadic reflection. We conjecture that the relative expressiveness of delimited control does not change even with answer-type modification, once we add analogous capabilities to effect handlers [Brady 2013; Kiselyov 2016] and monadic reflection [Atkey 2009].

5.4 Operational Metatheory

Our Abella formalisation establishes:

THEOREM 5.4 (DEL SAFETY). *Well-typed programs don't go wrong: for all closed DEL returners $\Theta; \vdash_0 M : FG$, either $M \rightsquigarrow N$ for some $\Theta; \vdash_0 N : FG$ or else $M = \mathbf{return} V$ for some $\Theta; \vdash V : G$.*

In the next section, we extend DEL's metatheory using the translation from DEL to MON.

We define DEL's ground types, plugged contexts and the equivalences $\simeq$ and $\simeq_{\text{cong}}$ as in MAM.

6 MACRO TRANSLATIONS

Felleisen [1991] argues that the usual notions of computability and complexity reduction do not capture the expressiveness of general-purpose programming languages. The Church-Turing thesis and its extensions assert that any reasonably expressive model of computation can be efficiently reduced to any other reasonably expressive model of computation. Thus the notion of a polynomial-time reduction with a Turing-machine is too crude to differentiate expressive power of two general-purpose programming languages. As an alternative, Felleisen introduces *macro translation*: a *local* reduction of a language extension, in the sense that it is homomorphic with respect to the syntactic constructs, and *conservative*, in the sense that it does not change the core

language. We adapt this concept to local translations between conservative extensions of a shared core.

Translation Notation. We define translations $S \rightarrow T$ from each source calculus S to each target calculus T . By default we assume untyped translations, writing EFF , MON , and DEL in translations that disregard typeability. In typeability preserving translations, which must also respect the monad laws where MON is concerned, we explicitly write TYPED EFF , TYPED MON , and TYPED DEL . We allow translations to be *hygienic* and introduce fresh binding occurrences. We write $M \mapsto \underline{M}$ for the translation at hand. We include only the non-core cases in the definition of each translation.

Out of the six possible untyped macro-translations, the ideas behind the following four already appear in the literature: $\text{DEL} \rightarrow \text{MON}$ [Wadler 1994], $\text{MON} \rightarrow \text{DEL}$ [Filinski 1994], $\text{DEL} \rightarrow \text{EFF}$ [Bauer and Pretnar 2015], and $\text{EFF} \rightarrow \text{MON}$ [Kammar et al. 2013]. The Abella formalisation contains the proofs of the simulation results for each of the six translations. Three translations formally simulate the source calculus by the target calculus: $\text{MON} \rightarrow \text{DEL}$, $\text{DEL} \rightarrow \text{EFF}$, and $\text{MON} \rightarrow \text{EFF}$. The other translations, $\text{DEL} \rightarrow \text{MON}$, $\text{EFF} \rightarrow \text{DEL}$, and $\text{EFF} \rightarrow \text{MON}$, introduce suspended redexes during reduction that invalidate simulation on the nose.

For the translations that introduce suspended redexes, we use a relaxed variant of simulation, namely the relations $\sim_{\text{cong}}$, which are the smallest relations containing $\sim$ that are closed under the term formation constructs. We say that a translation $M \mapsto \underline{M}$ is a simulation *up to congruence* if for every reduction $M \sim N$ in the source calculus we have $\underline{M} \sim_{\text{cong}}^+ \underline{N}$ in the target calculus. In fact, the suspended redexes always β -reduce by substituting a variable, i.e., $\{\lambda x.M\}! x \sim_{\text{cong}}^+ \lambda x.M$, thus only performing simple rewiring.

6.1 Delimited Continuations as Monadic Reflection ($\text{DEL} \rightarrow \text{MON}$)

We adapt Wadler’s analysis of delimited control [1994], using the continuation monad [Moggi 1989]:

LEMMA 6.1. *For all $\Theta \vdash_k E : \text{Eff}$, $\Theta \vdash_k C : \text{Comp}_E$, we have the following proper monad Cont:*

$$\Theta \vdash_k E < \text{instance monad } (\alpha.U_E (\alpha \rightarrow C) \rightarrow C) \text{ where } \{\text{return } x = \lambda c.c! x; \\ m \gg f = \lambda c.m! \{\lambda y.f! y c\}\} : \text{Eff}$$

Using Cont we define the macro translation $\text{DEL} \rightarrow \text{MON}$ as follows:

$$\underline{S_0 k.M} := \mu(\lambda k.M) \quad \underline{\langle M | x.N \rangle} := [\underline{M}]^{\text{Cont}} \{\lambda x.N\}$$

Shift is interpreted as reflection and reset as reification in the continuation monad.

THEOREM 6.2 ($\text{DEL} \rightarrow \text{MON}$ CORRECTNESS). *MON simulates DEL up to congruence:*

$$M \sim N \implies \underline{M} \sim_{\text{cong}}^+ \underline{N}$$

The only suspended redex arises in simulating the reflection rule, where we substitute a continuation into the bind of the continuation monad yielding a term of the form $\{\lambda y.\{\lambda y.M\} y c\}$ which we must reduce to $\{\lambda y.M c\}$.

$\text{DEL} \rightarrow \text{MON}$ extends to a macro translation at the type level:

$$\underline{E}, \underline{C} := \underline{E} < \text{instance monad } (\alpha.U_E (\alpha \rightarrow \underline{C}) \rightarrow \underline{C}) \text{ Cont}$$

THEOREM 6.3 ($\text{DEL} \rightarrow \text{MON}$ PRESERVES TYPEABILITY). *Every well-typed DEL phrase $\Theta; \Gamma \vdash_E P : X$ translates into a proper well-typed MON phrase: $\Theta; \underline{\Gamma} \vdash_{\underline{E}} \underline{P} : \underline{X}$.*

We use this result to extend the metatheory of DEL:

COROLLARY 6.4 (DEL TERMINATION). *All well-typed closed ground returners in DEL must reduce to a unique normal form: if $\vdash_0 M : FG$ then there exists V such that $\vdash V : G$ and $M \sim^* \text{return } V$.*

6.2 Monadic Reflection as Delimited Continuations ($\text{MON} \rightarrow \text{DEL}$)

We define the macro translation $\text{MON} \rightarrow \text{DEL}$ as follows:

$$\begin{aligned} \underline{\mu}(M) &:= S_0 k. \lambda b. b! (\{ \underline{M} \}, \{ \lambda x. k! x \ b \}) \\ \underline{[M]}^{\text{where } \{ \text{return } x=N_u; y \gg f=N_b \}} &:= \left\langle \underline{M} \middle| x. \lambda b. \underline{N_u} \right\rangle \{ \lambda (y, f). \underline{N_b} \} \end{aligned}$$

Reflection is interpreted by capturing the current continuation and abstracting over the bind operator which is then invoked with the reflected computation and a function that wraps the continuation in order to ensure it uses the same bind operator. Reification is interpreted as an application of a reset. The continuation of the reset contains the unit of the monad. We apply this reset to the bind of the monad.

THEOREM 6.5 ($\text{MON} \rightarrow \text{DEL}$ CORRECTNESS). *DEL simulates MON up to congruence:*

$$M \rightsquigarrow N \implies \underline{M} \rightsquigarrow_{\text{cong}}^+ \underline{N}$$

This translation does not preserve typeability because the bind operator can be used at different types. We conjecture that a) any other macro translation will suffer from the same issue and b) adding (predicative) polymorphism to the base calculus is sufficient to adapt this translation to one that does preserve typeability.

Filinski's translation from monadic reflection to delimited continuations [1994] does preserve typeability, but it is a global translation. It is much like our translation except each instance of bind is inlined (hence it does not need to be polymorphic).

6.2.1 Alternative Translation with Nested Delimited Continuations. An alternative to $\text{MON} \rightarrow \text{DEL}$ is to use two nested shifts for reflection and two nested resets for reification:

$$\begin{aligned} \underline{\mu}(M) &:= S_0 k. S_0 b. b! (\{ \underline{M} \}, \{ \lambda x. \langle k! x | z. z! b \rangle \}) \\ \underline{[M]}^{\text{where } \{ \text{return } x=N_u; y \gg f=N_b \}} &:= \left\langle \left\langle \underline{M} \middle| x. S_0 b. \underline{N_u} \right\rangle (y, f). \underline{N_b} \right\rangle \end{aligned}$$

In the translation of reflection, the reset inside the wrapped continuation ensures that any further reflections in the continuation are interpreted appropriately: the first shift, which binds k , has popped one continuation off the stack so we need to add one back on. In the translation of reification, the shift guarding the unit garbage collects the bind once it is no longer needed.

6.3 Delimited Continuations as Effect Handlers ($\text{DEL} \rightarrow \text{EFF}$)

We define $\text{DEL} \rightarrow \text{EFF}$ as follows:

$$\underline{S_0 k. N} := \text{shift0 } \{ \lambda k. \underline{N} \} \quad \underline{\langle M | x. N \rangle} := \text{handle } \underline{M} \text{ with } \{ \text{return } x \mapsto \underline{N} \} \uplus \{ \text{shift0 } y \ f \mapsto f! y \}$$

Shift is interpreted as an operation and reset is interpreted as a straightforward handler.

THEOREM 6.6 ($\text{DEL} \rightarrow \text{EFF}$ CORRECTNESS). *EFF simulates DEL on the nose: $M \rightsquigarrow N \implies \underline{M} \rightsquigarrow^+ \underline{N}$.*

This translation does not preserve typeability because inside a single reset shifts can be used at different types. We conjecture that a) any other macro translation will suffer from the same issue and b) adding polymorphic operations [Kammar et al. 2013] to EFF is sufficient to ensure this translation does preserve typeability.

One can adapt our translation to a global translation in which every static instance of a shift is interpreted as a separate operation, thus avoiding the need for polymorphic operations.

6.4 Effect Handlers as Delimited Continuations ($\text{EFF} \rightarrow \text{DEL}$)

We define $\text{EFF} \rightarrow \text{DEL}$ as follows:

$$\begin{aligned} \text{op } V &:= S_0 k. \lambda h. h! (\text{inj}_{\text{op}} (V, \{\lambda y. k! y h\})) & \underline{\text{handle } M \text{ with } H} &:= \langle \underline{M} | H^{\text{ret}} \rangle \{H^{\text{ops}}\} \\ \left(\text{handle } M \text{ with } \begin{array}{l} \{\text{return } x \mapsto N_{\text{ret}}\} \\ \uplus \{\text{op}_i p k \mapsto N_i\}_i \end{array} \right)^{\text{ret}} &:= x. \lambda h. \underline{N_{\text{ret}}} & \left(\text{handle } M \text{ with } \begin{array}{l} \{\text{return } x \mapsto N_{\text{ret}}\} \\ \uplus \{\text{op}_i p k \mapsto N_i\}_i \end{array} \right)^{\text{ops}} &:= \lambda y. \text{case } y \text{ of } \{ \\ & & & (\text{inj}_{\text{op}_i} (p, k) \rightarrow \underline{N_i})_i \end{aligned}$$

Operation invocation is interpreted by capturing the current continuation and abstracting over a dispatcher which is passed an encoding of the operation. The encoded operation is an injection whose label is the name of the operation containing a pair of the operation parameter and a wrapped version of the captured continuation, which ensures the same dispatcher is threaded through the continuation.

Handling is interpreted as an application of a reset whose continuation contains the return clause. The reset is applied to a dispatcher function that encodes the operation clauses.

THEOREM 6.7 ($\text{EFF} \rightarrow \text{DEL}$ CORRECTNESS). *DEL simulates EFF up to congruence:*

$$M \rightsquigarrow N \implies \underline{M} \rightsquigarrow_{\text{cong}}^+ \underline{N}$$

The $\text{EFF} \rightarrow \text{DEL}$ translation is simpler than Kammar et al.'s which uses a global higher-order memory cell storing the handler stack [2013].

This translation does not preserve typeability because the interpretation of operations needs to be polymorphic in the return type of the dispatcher over which it abstracts. We conjecture that a) any other macro translation will suffer from the same issue and b) adding polymorphism to the base calculus is sufficient to adapt this translation to one that preserves typeability.

6.4.1 Alternative Translation with Nested Delimited Continuations. Similarly to the $\text{MON} \rightarrow \text{DEL}$ translation there is an alternative to $\text{EFF} \rightarrow \text{DEL}$ which uses two nested shifts for operations and two nested resets for handlers:

$$\begin{aligned} \text{op } V &:= S_0 k. S_0 h. h! (\text{inj}_{\text{op}} (V, \{\lambda x. \langle k! x | y. y! h \rangle\})) & \underline{\text{handle } M \text{ with } H} &:= \langle \langle \underline{M} | H^{\text{ret}} \rangle | H^{\text{ops}} \rangle \\ \left(\begin{array}{l} \{\text{return } x \mapsto N_{\text{ret}}\} \\ \uplus \{\text{op}_i p k \mapsto N_i\}_i \end{array} \right)^{\text{ret}} &:= x. S_0 h. \underline{N_{\text{ret}}} & \left(\begin{array}{l} \{\text{return } x \mapsto N_{\text{ret}}\} \\ \uplus \{\text{op}_i p k \mapsto N_i\}_i \end{array} \right)^{\text{ops}} &:= y. \text{case } y \text{ of } \{ \\ & & & (\text{inj}_{\text{op}_i} (p, k) \rightarrow \underline{N_i}) \end{aligned}$$

6.5 Monadic Reflection as Effect Handlers ($\text{MON} \rightarrow \text{EFF}$)

We simulate reflection with an operation and reification with a handler. Formally, for every anonymous monad T given by **where** $\{\text{return } x = N_u; y \gg f = N_b\}$ we define $\text{MON} \rightarrow \text{EFF}$ as follows:

$$\begin{aligned} \underline{\mu(N)} &:= \text{reflect } \{ \underline{N} \} & \underline{[M]^T} &:= \text{handle } \underline{M} \text{ with } \underline{T} \\ \underline{T} &:= \{\text{return } x \mapsto \underline{N_u}\} \uplus \{\text{reflect } y f \mapsto \underline{N_b}\} \end{aligned}$$

Reflection is interpreted as a reflect operation and reification as a handler with the unit of the monad as a handler and the bind of the handler as the implementation of the reflect operation.

THEOREM 6.8 ($\text{MON} \rightarrow \text{EFF}$ CORRECTNESS). *EFF simulates MON on the nose: $M \rightsquigarrow N \implies \underline{M} \rightsquigarrow^+ \underline{N}$.*

$\text{MON} \rightarrow \text{EFF}$ does not preserve typeability. For instance, consider the following computation of type $F\text{bit}$ using the environment monad Reader given on the right:

$$\begin{aligned} [b \leftarrow \mu(\{\lambda(b, f). b\}); & \quad \vdash_k \emptyset < \text{instance monad } (\alpha. \text{bit} \times U_0 (\text{bit} \rightarrow F \text{bit}) \rightarrow F \alpha) \\ f \leftarrow \mu(\{\lambda(b, f). f\}); & \quad \text{where } \{\text{return } x = \lambda e. \text{return } x; \\ f! b\}^{\text{Reader}} (\text{inj}_{\text{true}} (), \{\lambda b. \text{return } b\}) & \quad m \gg f = \lambda e. x \leftarrow m! e; f! x e\} : \text{Eff} \end{aligned}$$

Its translation into EFF is not typeable: reflection can appear at any type, whereas a single operation is monomorphic. We conjecture that a) this observation can be used to prove that *no* macro translation $\text{TYPED MON} \rightarrow \text{TYPED EFF}$ exists and that b) adding polymorphic operations [Kammar et al. 2013] to EFF is sufficient for typing this translation.

6.6 Effect Handlers as Monadic Reflection (EFF $\rightarrow$ MON)

We define EFF $\rightarrow$ MON as follows:

$$\begin{aligned} \underline{\text{op}} V &:= \mu(\lambda k. \lambda h. h! (\text{inj}_{\text{op}} (V, \{\lambda y. k! y h\}))) & \underline{\text{handle}} M \text{ with } H &:= [\underline{M}]^{\text{Cont}} \{H^{\text{ret}}\} \{H^{\text{ops}}\} \\ \left(\begin{array}{c} \text{handle } M \text{ with} \\ \{\text{return } x \mapsto N_{\text{ret}}\} \\ \uplus \{\text{op}_i p k \mapsto N_i\}_i \end{array} \right)^{\text{ret}} &:= \lambda x. \lambda h. \underline{N}_{\text{ret}} & \left(\begin{array}{c} \text{handle } M \text{ with} \\ \{\text{return } x \mapsto N_{\text{ret}}\} \\ \uplus \{\text{op}_i p k \mapsto N_i\}_i \end{array} \right)^{\text{ops}} &:= \lambda y. \text{case } y \text{ of } \{ \\ & & (\text{inj}_{\text{op}_i} (p, k) \rightarrow \underline{N}_i)_i \end{aligned}$$

The translation is much like EFF $\rightarrow$ DEL, using the continuation monad in place of first class continuations.

Operation invocation is interpreted by using reflection to capture the current continuation and abstracting over a dispatcher which is passed an encoding of the operation. The encoded operation is an injection whose label is the name of the operation containing a pair of the operation parameter and a wrapped version of the captured continuation, which ensures the same dispatcher is threaded through the continuation.

Handling is interpreted as an application of a reified continuation monad computation to the return clause and a dispatcher function that encodes the operation clauses.

THEOREM 6.9 (EFF $\rightarrow$ MON CORRECTNESS). *MON simulates EFF up to congruence:*

$$M \rightsquigarrow N \implies \underline{M} \rightsquigarrow_{\text{cong}}^+ \underline{N}$$

This translation does not preserve typeability for the same reason as the EFF $\rightarrow$ DEL translations: the interpretation of operations needs to be polymorphic in the return type of the dispatcher over which it abstracts. We conjecture that a) any other macro translation will suffer from the same issue and b) adding polymorphism to the base calculus is sufficient to adapt this translation to one that does preserve typeability.

6.6.1 Alternative Translation Using a Free Monad. An alternative to interpreting effect handlers using a continuation monad is to use a free monad:

$$\begin{aligned} \underline{\text{op}} V &:= \mu(\text{return} (\text{inj}_{\text{op}} (V, \lambda x. \text{return } x))) & \underline{\text{handle}} M \text{ with } H &:= H^{\star} [\underline{M}]^{H^{\dagger}} \\ & \text{where } \{ & \text{return } x &= \text{return} (\text{inj}_{\text{ret}} x); \\ \left(\begin{array}{c} \{\text{return } x \mapsto N_{\text{ret}}\} \\ \uplus \{\text{op}_i p k \mapsto N_i\}_i \end{array} \right)^{\dagger} &:= & y \gg f = \text{case } y \text{ of } \{ \text{inj}_{\text{ret}} x \rightarrow k! x \\ & & (\text{inj}_{\text{op}_i} (p, k) \rightarrow \text{return} (\text{inj}_{\text{op}_i} (p, \lambda x. k! x \gg f)))_i \} \\ & \} \\ \left(\begin{array}{c} \{\text{return } x \mapsto N_{\text{ret}}\} \\ \uplus \{\text{op}_i p k \mapsto N_i\}_i \end{array} \right)^{\star} &:= & h = \lambda y. \text{case } y \text{ of } \{ \text{inj}_{\text{ret}} x \rightarrow \underline{N}_{\text{ret}} \\ & & (\text{inj}_{\text{op}_i} (p, k') \rightarrow k \leftarrow \text{return} \{ \lambda x. y \leftarrow k'! x; h! y \}; \underline{N}_i \}_i \end{aligned}$$

Both the bind operation for the free monad $H^{\dagger}$ and the function h that interprets the free monad $H^{\star}$ are recursive. Given that we are in an untyped setting we can straightforwardly implement the recursion using a suitable variation of the Y combinator. This translation does not extend to the typed calculi as they do not support recursion. Nevertheless, we conjecture that it can be adapted to a typed translation if we extend our base calculus to include inductive data types, as the recursive functions are structurally recursive.

6.7 Nonexistence Results

THEOREM 6.10. *The following macro translations do not exist:*

- *TYPED EFF* $\rightarrow$ *TYPED MON* satisfying: $M \rightsquigarrow N \implies \underline{M} \simeq \underline{N}$.
- *TYPED EFF* $\rightarrow$ *TYPED DEL* satisfying: $M \rightsquigarrow N \implies \underline{M} \simeq \underline{N}$.

Our proof of the first part hinges on the finite denotation property (Lemma 4.9). Briefly, assume to the contrary that there was such a translation. Consider a single effect operation symbol $\text{tick} : 1 \rightarrow 1$ and set $\text{tick}^0 := \text{return } ()$, and $\text{tick}^{n+1} := \text{tick}(); \text{tick}^n$. All these terms have the same type, and by the homomorphic property of the hypothesised translation, their translations all have the same type. By the finite denotation property two of them are observationally equivalent and by virtue of a macro translation the two source terms are observationally equivalent in *EFF*. But every distinct pair of tick^n terms is observationally distinguishable using an appropriate handler. See Forster’s thesis [2016] for the full details. The second part follows from Theorem 6.3.

Regarding the remaining four possibilities, we have seen that there is a typeability-preserving macro translation *TYPED DEL* $\rightarrow$ *TYPED MON* (Theorem 6.3), but we conjecture that there are no typeability-preserving translations *TYPED MON* $\rightarrow$ *TYPED DEL*, *TYPED DEL* $\rightarrow$ *TYPED EFF*, or *TYPED MON* $\rightarrow$ *TYPED EFF*.

7 CONCLUSION AND FURTHER WORK

We have given a uniform family of formal calculi expressing the common abstractions for user-defined effects: effect handlers (*EFF*), monadic reflection (*MON*), and delimited control (*DEL*), together with their natural type-and-effect systems. We have used these calculi to formally analyse the relative expressive power of these abstractions. Effect handlers, monadic reflection, and delimited control have equivalent expressivity when types are not taken into consideration. However, neither monadic reflection nor delimited control can macro-express effect handlers whilst preserving typeability. We have formalised the more syntactic aspects of our work in the Abella proof assistant, and have used set-theoretic denotational semantics to establish inexpressivity results.

Our work has already born unexpected if not entirely unsurprising fruit. By composing our translation from effect handlers to delimited continuations with a CPS translation for delimited continuations Hillerström et al. [2017] derived a CPS translation for effect handlers, which they then used as the basis for an implementation.

Further work abounds. We want to extend each type system until each translation preserves typeability. We conjecture that adding polymorphic operations to *EFF* would allow it to macro express *DEL* and *MON*, and that adding polymorphism to *MON* and *DEL* would allow them to macro express *EFF*. We conjecture polymorphism would also allow *DEL* to macro express *MON*, and inductive data types with primitive recursion would also allow *MON* to macro express *EFF*.

We are also interested in analysing *global* translations between these abstractions. In particular, whereas *MON* and *DEL* allow reflection/shifts to appear anywhere inside a piece of code, in practice, library designers define a fixed set of primitives using reflection/shifts and only expose those primitives to users. This observation suggests calculi in which each *reify/reset* is accompanied by declarations of this fixed set of primitives. We conjecture that *MON* and *DEL* can be simulated on the nose via a global translation into the corresponding restricted calculus, and that the restricted calculi can be macro translated into *EFF* whilst preserving typeability. Such two-stage translations would give a deeper reason why so many examples typically used for monadic reflection and delimited control can be directly recast using effect handlers. Other global pre-processing may also eliminate administrative reductions from our translations and establish simulation on the nose.

We hope the basic calculi we have analysed will form a foundation for systematic further investigation. Supporting answer-type modification [Asai 2009; Kobori et al. 2015] can inform

more expressive type system design for effect handlers and monadic reflection, and account for type-state [Atkey 2009] and session types [Kiselyov 2016]. In practice, effect systems are often extended with sub-effecting or effect polymorphism [Lucassen and Gifford 1988; Bauer and Pretnar 2014; Pretnar 2014; Leijen 2017; Hillerström and Lindley 2016; Lindley et al. 2017]. To these we add effect forwarding [Kammar et al. 2013] and rebasing [Filinski 2010].

We have taken the perspective of a programming language designer deciding which programming abstraction to select for expressing user-defined effects. In contrast, Schrijvers et al. [2016] take the perspective of a library designer for a specific programming language, Haskell, and compare the abstractions provided by libraries based on monads with those provided by effect handlers. They argue that both libraries converge on the same interface for user-defined effects via Haskell’s type-class mechanism.

Relative expressiveness results are subtle, and the potentially negative results that are hard to establish make them a risky line of research. We view denotational models as providing a fruitful method for establishing such inexpressivity results. It would be interesting to connect our work with that of Laird [2002, 2013, 2017], who analyses the macro-expressiveness of a hierarchy of combinations of control operators and exceptions using game semantics, and in particular uses such denotational techniques to show certain combinations cannot macro express other combinations. We would like to apply similar techniques to compare the expressive power of local effects such as ML-style reference cells with effect handlers.

ACKNOWLEDGMENTS

Supported by the European Research Council grant ‘events causality and symmetry — the next-generation semantics’, and the Engineering and Physical Sciences Research Council grants EP/H005633/1 ‘Semantic Foundations for Real-World Systems’, EP/K034413/1 ‘From Data Types to Session Types—A Basis for Concurrency and Distribution’, and EP/N007387/1 ‘Quantum computation as a programming language’. The material is based upon work supported by the Air Force Office of Scientific Research, Air Force Materiel Command, USAF under Award No. FA9550-14-1-0096. We thank Bob Atkey, Andrej Bauer, Paul Downen, Marcelo Fiore, Tamara von Glehn, Mathieu Huot, Daan Leijen, Craig McLaughlin, Kayvan Memarian, Sean Moss, Alan Mycroft, Ian Orton, Hugo Paquet, Jean Pichon-Pharabod, Matthew Pickering, Didier Remy, Reuben Rowe, Philip Saville, Ian Stark, Sam Staton, Philip Wadler, and Jeremy Yallop for useful suggestions and discussions.

REFERENCES

- Kenichi Asai. 2009. On typing delimited continuations: three new solutions to the printf problem. *Higher-Order and Symbolic Computation* 22, 3 (2009), 275–291.
- Robert Atkey. 2009. Parameterised notions of computation. *J. Funct. Program.* 19, 3-4 (2009), 335–376.
- M. Barr and C. Wells. 1985. *Toposes, triples, and theories*. Springer-Verlag.
- Andrej Bauer and Matija Pretnar. 2014. An Effect System for Algebraic Effects and Handlers. *Logical Methods in Computer Science* 10, 4 (2014).
- Andrej Bauer and Matija Pretnar. 2015. Programming with algebraic effects and handlers. *J. Log. Algebr. Meth. Program.* 84, 1 (2015), 108–123.
- Edwin Brady. 2013. Programming and reasoning with algebraic effects and dependent types. In *ICFP*. ACM, 133–144.
- Lukas Bulwahn, Alexander Krauss, Florian Haftmann, Levent Erkök, and John Matthews. 2008. Imperative Functional Programming with Isabelle/HOL. In *TPHOLs (Lecture Notes in Computer Science)*, Vol. 5170. Springer, 134–149.
- Olivier Danvy. 2006. *An Analytical Approach to Programs as Data Objects*. DSc dissertation. Department of Computer Science, University of Aarhus.
- Olivier Danvy and Andrzej Filinski. 1989. *A Functional Abstraction of Typed Contexts*. Technical Report 89/12. DIKU.
- Olivier Danvy and Andrzej Filinski. 1990. Abstracting Control. In *LISP and Functional Programming*, 151–160.
- Christian Doczkal. 2007. *Strong Normalization of CBPV*. Technical Report. Saarland University.
- Christian Doczkal and Jan Schwinghammer. 2009. Formalizing a Strong Normalization Proof for Moggi’s Computational Metalanguage. In *LFMTP*. ACM, 57–63.

- Matthias Felleisen. 1991. On the Expressive Power of Programming Languages. *Sci. Comput. Program.* 17, 1-3 (1991), 35–75.
- Matthias Felleisen and Daniel P. Friedman. 1987. A Reduction Semantics for Imperative Higher-Order Languages. In *PARLE (2) (Lecture Notes in Computer Science)*, Vol. 259. Springer, 206–223.
- Matthias Felleisen, Mitchell Wand, Daniel P. Friedman, and Bruce F. Duba. 1988. Abstract Continuations: A Mathematical Semantics for Handling Full Jumps. In *LISP and Functional Programming*, 52–62.
- Andrzej Filinski. 1994. Representing Monads. In *POPL*. ACM.
- Andrzej Filinski. 1996. *Controlling effects*. Ph.D. Dissertation. School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania.
- Andrzej Filinski. 1999. Representing Layered Monads. In *POPL*. ACM.
- Andrzej Filinski. 2010. Monads in Action. *SIGPLAN Not.* 45, 1 (Jan. 2010), 483–494.
- Yannick Forster. 2016. *On the expressive power of effect handlers and monadic reflection*. Technical Report. University of Cambridge.
- Andrew Gacek. 2008. The Abella Interactive Theorem Prover (System Description). In *IJCAR*, Vol. 5195. Springer, 154–161.
- Andrew Gacek. 2009. *A Framework for Specifying, Prototyping, and Reasoning about Computational Systems*. Ph.D. Dissertation. University of Minnesota.
- Claudio Hermida. 1993. *Fibrations, logical predicates and related topics*. Ph.D. Dissertation. University of Edinburgh.
- Daniel Hillerström and Sam Lindley. 2016. Liberating effects with rows and handlers. In *TyDe@ICFP*. ACM, 15–27.
- Daniel Hillerström, Sam Lindley, Robert Atkey, and KC Sivaramakrishnan. 2017. Continuation Passing Style for Effect Handlers. In *FSCD (LIPIcs)*, Vol. 84. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Article 18.
- Graham Hutton and Erik Meijer. 1998. Monadic Parsing in Haskell. *J. Funct. Program.* 8, 4 (1998), 437–444.
- Ohad Kammar. 2014. *An Algebraic Theory of Type-and-Effect Systems*. Ph.D. Dissertation. University of Edinburgh.
- Ohad Kammar, Sam Lindley, and Nicolas Oury. 2013. Handlers in action. In *ICFP*. ACM, 145–158.
- Ohad Kammar and Gordon D. Plotkin. 2012. Algebraic foundations for effect-dependent optimisations. In *POPL*. ACM.
- Ohad Kammar and Matija Pretnar. 2017. No value restriction is needed for algebraic effects and handlers. *J. Funct. Program.* 27 (2017), e7.
- Shin-ya Katsumata. 2014. Parametric Effect Monads and Semantics of Effect Systems. *SIGPLAN Not.* 49, 1 (Jan. 2014), 633–645.
- Oleg Kiselyov. 2016. Parameterized extensible effects and session types (extended abstract). In *TyDe@ICFP*. ACM, 41–42.
- Oleg Kiselyov, Daniel P. Friedman, and Amr A. Sabry. 2005. *How to remove a dynamic prompt: static and dynamic delimited continuation operators are equally expressible*. Technical Report. 16 pages. Technical Report TR611.
- Oleg Kiselyov, Amr Sabry, and Cameron Swords. 2013. Extensible effects: an alternative to monad transformers. In *Haskell*. ACM, 59–70.
- Oleg Kiselyov and Chung-chieh Shan. 2007. A Substructural Type System for Delimited Continuations. In *TLCA*. 223–239.
- Oleg Kiselyov, Chung-chieh Shan, and Amr Sabry. 2006. Delimited dynamic binding. In *ICFP*. ACM, 26–37.
- Ikuo Kobori, Yuki Yoshi Kameyama, and Oleg Kiselyov. 2015. Answer-Type Modification without Tears: Prompt-Passing Style Translation for Typed Delimited-Control Operators. In *WoC 2015 (EPTCS)*, Vol. 212. 36–52.
- James Laird. 2002. Exceptions, Continuations and Macro-expressiveness. In *ESOP*. 133–146.
- James Laird. 2013. Combining and Relating Control Effects and their Semantics. In *COS*. 113–129.
- J. Laird. 2017. Combining control effects and their models: Game semantics for a hierarchy of static, dynamic and delimited control effects. *Annals of Pure and Applied Logic* 168, 2 (2017), 470–500. Eighth Games for Logic and Programming Languages Workshop (GaLoP).
- Daan Leijen. 2017. Type directed compilation of row-typed algebraic effects. In *POPL*. ACM, 486–499.
- Paul Blain Levy. 2004. *Call-By-Push-Value: A Functional/Imperative Synthesis*. Semantics Structures in Computation, Vol. 2. Springer.
- Sam Lindley, Conor McBride, and Craig McLaughlin. 2017. Do be do be do. In *POPL*. ACM, 500–514.
- Sam Lindley and Ian Stark. 2005. Reducibility and $\tau\tau$ -Lifting for Computation Types. In *TLCA (Lecture Notes in Computer Science)*, Vol. 3461. Springer, 262–277.
- John M. Lucassen and David K. Gifford. 1988. Polymorphic Effect Systems. In *POPL*. ACM, 47–57.
- Francisco Marmolejo and Richard J. Wood. 2010. Monads as extension systems — no iteration is necessary. *Theory and Applications of Categories* 24, 4 (2010), 84–113.
- Marek Materzok and Dariusz Biernacki. 2012. A Dynamic Interpretation of the CPS Hierarchy. In *APLAS (LNCS)*, Vol. 7705. Springer, 296–311.
- Eugenio Moggi. 1989. Computational Lambda-Calculus and Monads. In *LICS*. IEEE Computer Society, 14–23.
- Gordon D. Plotkin and John Power. 2002. Notions of Computation Determine Monads. In *FoSSaCS*. Springer-Verlag.
- Gordon D. Plotkin and John Power. 2003. Algebraic Operations and Generic Effects. *Appl. Categ. Structures* 11, 1 (2003), 69–94.
- Gordon D. Plotkin and Matija Pretnar. 2008. A Logic for Algebraic Effects. In *LICS*. IEEE Computer Society, 118–129.

- Gordon D. Plotkin and Matija Pretnar. 2009. Handlers of Algebraic Effects. In *ESOP*. Springer-Verlag.
- Matija Pretnar. 2014. Inferring Algebraic Effects. *Logical Methods in Computer Science* 10, 3 (2014).
- Matija Pretnar. 2015. An Introduction to Algebraic Effects and Handlers. Invited tutorial paper. *Electr. Notes Theor. Comput. Sci.* 319 (2015), 19–35.
- John C. Reynolds. 2009. *Theories of Programming Languages*. Cambridge University Press.
- Tom Schrijvers and others. 2016. *Monad transformers and modular algebraic effects*. Technical Report. University of Leuven.
- Tom Schrijvers, Guido Tack, Pieter Wuille, Horst Samulowitz, and Peter J. Stuckey. 2013. Search combinators. *Constraints* 18, 2 (2013), 269–305.
- Chung-chieh Shan. 2007. A static simulation of dynamic delimited control. *Higher-Order and Symbolic Computation* 20, 4 (2007), 371–401.
- Ábel Sinkovics and Zoltán Porkoláb. 2013. Implementing monads for C++ template metaprograms. *Sci. Comput. Program.* 78, 9 (2013), 1600–1621.
- J. Michael Spivey. 1990. A Functional Theory of Exceptions. *Sci. Comput. Program.* 14, 1 (1990), 25–42.
- Wouter Swierstra. 2008. Data types à la carte. *J. Funct. Program.* 18, 4 (2008), 423–436.
- William W Tait. 1967. Intensional interpretations of functionals of finite type I. *The journal of symbolic logic* 32, 02 (1967), 198–212.
- Philip Wadler. 1990. Comprehending Monads. In *LISP and Functional Programming*. 61–78.
- Philip Wadler. 1994. Monads and Composable Continuations. *Lisp and Symbolic Computation* 7, 1 (1994), 39–56.
- Andrew K. Wright and Matthias Felleisen. 1994. A Syntactic Approach to Type Soundness. *Inf. Comput.* 115, 1 (1994), 38–94.
- Beta Ziliani, Derek Dreyer, Neelakantan R. Krishnaswami, Aleksandar Nanevski, and Viktor Vafeiadis. 2015. Mtac: A monad for typed tactic programming in Coq. *J. Funct. Program.* 25 (2015).