No value restriction is needed for algebraic effects and handlers

Ohad Kammar[1], Sean K. Moss[2], and Matija Pretnar[3]

[1] University of Cambridge Computer Laboratory
ohad.kammar@cl.cam.ac.uk
[2] University of Cambridge Department of Pure Mathematics
and Mathematical Statistics
s.k.moss@dpmms.cam.ac.uk
[3] University of Ljubljana, Faculty of Mathematics and Physics
matija.pretnar@fmf.uni-lj.si

The proposed talk describes submitted [7] (attached) and ongoing [6] work about the interaction of computational effects and predicative polymorphism. We have previously presented this work at the Types 2016 conference, and we would like to expose it to LOLA participants as well.

ML-style reference cells are known to be hard to combine with polymorphism [11, 4, 14], where a naïve type system is unsound (cf. Figure 1). The working solution, the value restriction [18] and its relaxation [3], are ad-hoc and restrict the programmer unnecessarily. We reexamine this problem in the context of algebraic effects [12] which extend the monadic account of computational effects [10] (e.g., the state monad) with the syntactic operations involving them (e.g., memory look-up and update). Bauer and Pretnar [2] use effect handlers, a generalisation of exception handlers that allows to handle arbitrary user-defined algebraic effects, to structure impure functional code, in analogy with monads [17]. The smooth integration of algebraic effects with polymorphism is surprising as effect handlers can implement local-state-like programming examples by manipulating continuations (k below) [13]:

$$
\begin{aligned}
&(\textbf{with } H_{ST} \textbf{ handle set true;} && \text{where: } H_{ST} := \textbf{handler } \{ \\
&\qquad \textbf{let } y = \mathsf{get} \; () \textbf{ in} && \qquad \textbf{return } x \mapsto \textbf{fun } \_ \mapsto \textbf{return } x \\
&\qquad \textbf{return } y) \textbf{ false} && \qquad \mathsf{get}(\_;k) \mapsto \textbf{fun } s \mapsto k \; s \; s \\
&\leadsto^* \; \textbf{return true} && \qquad \mathsf{set}(s';k) \mapsto \textbf{fun } \_ \mapsto k \; () \; s' \}
\end{aligned}
$$

Figure 1:

> **let** $r = $ **ref** $[\,]$ **in**
> $r := [()]$;
> **true** $::\, !r$

Figure 1: unsound polymorphism and references [3]

In this work, we extend Bauer and Pretnar's [2] calculus for algebraic effects and handlers with Hindley-Milner polymorphism, in a standard way, without any value restriction:

- We add local effect signatures [8] Σ as finite mappings from operations op to pairs of value types A, B, which we denote by $(\mathsf{op} : A \to B) \in \Sigma$.
- We extend types with *type variables* α.
- We introduce *schemes* $\forall \vec{\alpha}.A$, where $\vec{\alpha}$ denotes a finite set of $|\vec{\alpha}|$-many type variables ranged over by α_i.

Our main result concerns the soundness of the type system w.r.t. the reduction relation $\leadsto$:

Theorem (Safety). *If $\vdash c : A\,!\,\Sigma$ holds, then either: (i) $c \leadsto c'$ for some $\vdash c' : A\,!\,\Sigma$; (ii) $c = $ **return** v for some $\vdash v : A$; or (iii) $c = \mathsf{op}(v; y.\,c')$ for some $(\mathsf{op} : A_{\mathsf{op}} \to B_{\mathsf{op}}) \in \Sigma$, $\vdash v : A_{\mathsf{op}}$, and $y : B_{\mathsf{op}} \vdash c' : A\,!\,\Sigma$. In particular, when $\Sigma = \emptyset$, evaluation will not get stuck before returning a value.*

We use Leroy's [9] benchmarks for evaluating the interaction of effects and polymorphism. For example, if we extend the language with lists and bounded iteration, we can integrate effects in polymorphic functions, as for any Σ:

let imp_map = **fun** $f\,xs \mapsto$
    with H_{ST} **handle** $\Big(\big($**foldl** $\;\big($**fun** $x \mapsto \mathsf{set}(f\;x :: \mathsf{get}\;())\big)\;()\;xs\big)$; $\mathsf{reverse}(\mathsf{get}\;())$
    [] (* initial state *) **in** ...   (* imp_map : $\forall\alpha\beta.(\alpha \to \beta\,!\,\Sigma) \to (\alpha\;\mathsf{list} \to \beta\;\mathsf{list}\,!\,\Sigma)\,!\,\emptyset$ *)

These benchmarks also highlight the limited expressiveness of effect handlers — we do not know how to implement, using effect handlers, even Leroy's basic benchmark, in which we return a newly allocated reference cell. The advantage is that, when trying to express the problematic program in Figure 1, we cannot express the first line, and the type system forbids handling the last two lines using the same state handler, as they refer to memory cells of different types.

A deeper soundness result comes from a denotational model for algebraic effects and polymorphism [6]. We modify Seely's models of impredicative polymorphism [15] by separating the fibred category of types into a fibred embedding of a fibred category of types into a fibred category of schemes. The universal quantifier $\forall$, previously right adjoint to structural weakening, is now replaced by a *relative* right adjoint [16, 1] along the inclusion of types in schemes. Using this relativisation, we can construct a parametric version of Harper and Mitchell's [5] set-theoretic models relative to a universal set $\mathcal{U}$. To add computational effects to this model, we construct a free fibred monad T_Δ and prove the following theorem, which allows us to interpret the calculus of algebraic effects and handlers, establishing soundness via a denotational model.

Theorem. *If $\mathcal{U} \neq \emptyset$, then the canonical morphism $T_{\Delta'}\forall\Delta.\tau \to \forall\Delta.T_{\Delta'\times\Delta}\tau$ is invertible.*

References

[1] Altenkirch, T., Chapman, J., and Uustalu, T. (2014). Monads need not be endofunctors. *CoRR*, **abs/1412.7148**.

[2] Bauer, A. and Pretnar, M. (2015). Programming with algebraic effects and handlers. *J. Log. Algebr. Meth. Program.*, **84**(1), 108–123.

[3] Garrigue, J. (2004). Relaxing the value restriction. In Y. Kameyama and P. J. Stuckey, editors, *Functional and Logic Programming*, volume 2998 of *Lecture Notes in Computer Science*, pages 196–213. Springer Berlin Heidelberg.

[4] Harper, R. and Lillibridge, M. (1993). Polymorphic type assignment and CPS conversion. *Lisp and Symbolic Computation*, **6**(3-4), 361–380.

[5] Harper, R. and Mitchell, J. C. (1993). On the type structure of standard ml. *ACM Trans. Program. Lang. Syst.*, **15**(2), 211–252.

[6] Kammar, O. and Moss, S. K. (2015). A denotational semantics for Hindley-Milner polymorphism. talk given at the 4th ACM SIGPLAN Workshop on Higher-Order Programming with Effects.

[7] Kammar, O. and Pretnar, M. (2015). No value restriction is needed for algebraic effects and handlers. *Journal of Functional Programming*. Submitted.

[8] Kammar, O., Lindley, S., and Oury, N. (2013). Handlers in action. *SIGPLAN Not.*, **48**(9), 145–158.

[9] Leroy, X. (1992). *Typage polymorphe d'un langage algorithmique*. Phd thesis (in french), Université Paris 7.

[10] Moggi, E. (1991). Notions of computation and monads. *Information and Computation*, **93**(1), 55–92.

[11] Pierce, B. C. (2002). *Types and Programming Languages*. MIT Press, Cambridge, MA, USA.

[12] Plotkin, G. D. and Power, J. (2002). Notions of computation determine monads. In M. Nielsen and U. Engberg, editors, *Foundations of Software Science and Computation Structures*, volume 2303 of *Lecture Notes in Computer Science*, pages 342–356. Springer Berlin Heidelberg.

[13] Plotkin, G. D. and Pretnar, M. (2013). Handling algebraic effects. *Logical Methods in Computer Science*, **9**(4).

[14] Rémy, D. (2015). Type systems. Lecture notes, Parisian Master of Research in Computer Science.

[15] Seely, R. A. G. (1987). Categorical semantics for higher order polymorphic lambda calculus. *J. Symb. Log.*, **52**(4), 969–989.

[16] Ulmer, F. (1968). Properties of dense and relative adjoint functors. *Journal of Algebra*, **8**(1), 77 – 95.

[17] Wadler, P. (1990). Comprehending monads. In *Proceedings of the 1990 ACM conference on LISP and functional programming*, LFP '90, pages 61–78, New York, NY, USA. ACM.

[18] Wright, A. K. (1995). Simple imperative polymorphism. *LISP and Symbolic Computation*, **8**(4), 343–355.

1

No value restriction is needed
for algebraic effects and handlers

Ohad Kammar·†
University of Cambridge Computer Laboratory and
University of Oxford Department of Computer Science, England
and
Matija Pretnar†
Faculty of Mathematics and Physics, University of Ljubljana, Slovenia

Abstract

We present a straightforward, sound Hindley-Milner polymorphic type system for algebraic effects
and handlers in a call-by-value calculus, which allows type variable generalisation of arbitrary com-
putations, not just values. This result is surprising. On the one hand, the soundness of unrestricted
call-by-value Hindley-Milner polymorphism is known to fail in the presence of computational effects
such as reference cells and continuations. On the other hand, many programming examples can be
recast to use effect handlers instead of these effects. Analysing the expressive power of effect handlers
with respect to state effects, we claim handlers cannot express reference cells, and show they can
simulate dynamically scoped state.

1 Introduction

The following *OCaml* example (Garrigue, 2004) demonstrates the problematic interaction
between Hindley-Milner polymorphism, which increases code reuse, and computational
effects, such as reference cells, in a call-by-value language:

> **let** $r =$ **ref** $[\,]$ **in** $(*$ generalise $r : \forall \alpha.\alpha$ list ref $*)$
>
> $r := [()];$ $(*$ specialise $\alpha :=$ unit $*)$
>
> **true** $:: !r$ $(*$ specialise $\alpha :=$ bool $*)$

A naïve type inference algorithm would assign the type α list ref to the term **ref** $[\,]$.
Unrestricted, it would assign to r the *type scheme* $\forall \alpha.\alpha$ list ref. But doing so allows us
to instantiate r with the unit type $\alpha :=$ unit to store the singleton list with the unit value,
and then to instantiate r with the boolean type $\alpha :=$ bool. The result is a list whose second
element is the unit value, but appears to the type system as a list of booleans.

The current way to avoid this well-known unsound behaviour (Pierce, 2002; Harper & Lillibridge, 1993; Rémy, 2015) is to enforce a *value restriction*: the inference algorithm will generalise the type variables only in value terms that cannot be reduced further (Wright, 1995). While this restriction can be weakened to allow some computation (Garrigue, 2004), it still rules out sound pure programs:

$$\textbf{let } id = (\textbf{fun } f \mapsto f) \, (\textbf{fun } x \mapsto x) \textbf{ in} \qquad (*\text{ id is } not \text{ polymorphic } *)$$
$$id \, (id)$$

The problem only arises when all three components are present: computational effects, polymorphism, and call-by-value. Without effects, Milner's original calculus soundly integrates call-by-value with type inference (Milner, 1978). Without polymorphism, computational effects behave predictably in call-by-value languages. Without call-by-value, Leroy (1993) combines computational effects with polymorphism without restriction. Leroy's language has two constructs for sequencing: a call-by-name polymorphic $\textbf{let } x = c_1 \textbf{ in } c_2$ construct in which c_1 is re-executed whenever it is specialised in c_2, and a call-by-value monomorphic $\textbf{do } x \leftarrow c_1 \textbf{ in } c_2$ construct in which c_1 is only evaluated once, but its type is not generalised. The situation is identical in the *Haskell* programming language, from which we borrowed this notation.

Programming with algebraic effects and handlers (Bauer & Pretnar, 2015) is a new approach to structuring functional programs with computational effects. The programmer declares a collection of *algebraic effect operations* with which she structures her effectful code. Then, separately, she defines *effect handlers* that implement these abstract operations. Bauer & Pretnar's *Eff* programming language is a strict (i.e., call-by-value) functional language with Hindley-Milner polymorphism, in which all computational effects are treated as algebraic effects that can be handled. There is a pre-defined collection of effects that receive special treatment: runtime errors and memory accesses. If these effects are not handled by the program, the runtime will handle them, invoking the corresponding real computational effects. As *Eff* combines the three problematic components (strictness, polymorphism, effects), it currently imposes the standard value restriction on the programmer.

In this paper, we show that if only algebraic effects and handlers are present, the language does not need a value restriction. We present a straightforward sound Hindley-Milner polymorphic type system for a call-by-value language that incorporates computational effects in the form of algebraic effects and their handlers. In order to simplify the presentation, we present a type system without its associated complete inference algorithm. Doing so decouples the algorithmic concerns of finding principal types and complexity from the semantic concern for soundness. As first-class polymorphism typically makes type inference undecidable (Wells, 1999), our type system uses ML-style polymorphism.

The rest of the paper is structured as follows. In Sec. 2, we give a short recap of handlers and show how they may be used to simulate global state. Next, in Sec. 3, we give a type and effect system and sketch the proof of its soundness. We formalized the proof in the Twelf proof assistant (Pfenning & Schürmann, 1999), extending Bauer & Pretnar's (2014) existing formalization of *Eff*'s core calculus. In Sec. 4 we evaluate our type system and discuss its expressivity with respect to mutable references and dynamically scoped state. Sec. 5 concludes.

Syntax

$v ::=$		value
	x	variable
	true $\mid$ **false**	boolean constants
	fun $x \mapsto c$	function
	h	handler
$h ::=$		handler
	handler $\{$ **return** $x \mapsto c_r,$	return clause
	$\quad \mathsf{op}_1(x;k) \mapsto c_1, \ldots, \mathsf{op}_n(x;k) \mapsto c_n\}$	operation clauses
$c ::=$		computation
	return v	return
	do $x \leftarrow c_1$ **in** c_2	sequencing
	$\mathsf{op}(v;y.c)$	operation call
	if v **then** c_1 **else** c_2	conditional
	$v_1 \, v_2$	application
	with v **handle** c	handling

Syntactic sugar

Sugar	*Elaboration*
$\_$	fresh variable binding
$c_1 ; c_2$	**do** $\_ \leftarrow c_1$ **in** c_2
$c_1 \, c_2$	**do** $f \leftarrow c_1$ **in do** $a \leftarrow c_2$ **in** $f \, a$
if c **then** c_1 **else** c_2	**do** $b \leftarrow c$ **in if** b **then** c_1 **else** c_2
$\mathsf{op}(c_p;y.c_k)$	**do** $p \leftarrow c_p$ **in** $\mathsf{op}(p;y.c_k)$
fun $x_1 \, x_2 \ldots x_n \mapsto c$	**fun** $x_1 \mapsto$ **fun** $x_2 \mapsto \ldots$ **fun** $x_n \mapsto c$
op	**fun** $x \mapsto \mathsf{op}(x;y.\,\textbf{return } y)$

Fig. 1. an idealised calculus of effect handlers

2 Handlers of algebraic effects

Algebraic effects are an approach to computational effects based on a premise that impure behaviour arises from a set of *operations* such as get and set for mutable store, read and print for interactive input and output, or raise for exceptions (Plotkin & Power, 2003). This naturally gives rise to *handlers* not only of exceptions, but of any other effect, yielding a novel concept that, amongst others, can capture backtracking, co-operative multithreading, Unix-style stream redirection, and delimited continuations (Plotkin & Pretnar, 2013; Bauer & Pretnar, 2015).

2.1 Language

We base our development on the calculus (Fig. 1) given in Pretnar's (2015) tutorial. The language is a variant of the fine-grained call-by-value λ-calculus of Levy *et al.* (2003), in which terms are split into inert *values* and potentially effectful *computations*.

Programmers introduce effects with the construct $\mathsf{op}(v; y. c)$, which calls the operation op with the parameter v. The effect invocation may yield a value to the continuation c using the bound variable y. Programmers define the meaning of such operation calls by enclosing them in effect handlers. A handler specifies a return clause, used when the computation returns a final value, and a collection of operation clauses $\mathsf{op}(x; k) \mapsto c$, which specify how we should execute an invocation of the operation op called with the parameter x and a continuation k. The underlying idea is that operation calls behave as signals that propagate outwards until they reach a handler with a matching clause.

Our handlers are *deep*: the additional effects in the continuation are also handled by the current handler. Our handlers are also *forwarding*: unhandled operations propagate through each handler until they are handled or reach the top level. None of these design choices is essential to the development below, but we make them to mirror *Eff*'s design choices.

We use the following syntactic sugar (Fig. 1): semicolons elaborate to binding fresh (dummy) variables; function calls, conditionals, and operation calls are elaborated to call-by-value evaluation order; function introduction may abstract over multiple arguments; and bare operations without a parameter or a continuation argument elaborate to the corresponding *generic effect* (Plotkin & Power, 2003). In our examples, we further assume to have the type unit with the sole inhabitant ( ).

2.2 State handlers

We represent state with an operation set, which sets the state contents to a given parameter and returns ( ), and get, which takes a unit parameter and returns the state contents. For example, here is a computation that toggles the state and returns the old value:

$$T \stackrel{\text{def}}{=} \textbf{if } \mathsf{get} \, () \textbf{ then}$$
$$\mathsf{set} \, \textbf{false}; \textbf{return true}$$
$$\textbf{else}$$
$$\mathsf{set} \, \textbf{true}; \textbf{return false}$$

As mentioned above, the runtime of *Eff* (Bauer & Pretnar, 2015) deals with unhandled primitive effects, but in our calculus, the behaviour of operations will be determined exclusively by handlers, and the computation T gets stuck when evaluated.

A simple example of a handler that can handle a stateful computation is one that sets the state to a fixed value, say **true**, and ignores all its modifications:

$$H_C := \textbf{handler } \{ \mathsf{get}(\_; k) \mapsto k \, \textbf{true}$$
$$\mathsf{set}(\, s; k) \mapsto k \, ()$$
$$\textbf{return } x \mapsto \textbf{return } x \}$$

Whenever a get operation is called, we yield **true** to the continuation, whereas all set calls are silently ignored by yielding the expected unit value ( ) and doing nothing else. The return clause of a handler states that the returned values are kept unmodified. When we handle T with H_C, we get back the result **true**, no matter how many times we call T.

A more useful handler is one that handles get and set in a way that results in the expected stateful behaviour. It uses a technique called *parameter-passing* (Plotkin & Pretnar,

Semantics

$$\frac{c_1 \rightsquigarrow c_1'}{\textbf{do } x \leftarrow c_1 \textbf{ in } c_2 \rightsquigarrow \textbf{do } x \leftarrow c_1' \textbf{ in } c_2} \qquad \frac{}{\textbf{do } x \leftarrow \textbf{return } v \textbf{ in } c \rightsquigarrow c[v/x]}$$

$$\frac{}{\textbf{do } x \leftarrow \textsf{op}(v;y.c_1) \textbf{ in } c_2 \rightsquigarrow \textsf{op}(v;y.\textbf{do } x \leftarrow c_1 \textbf{ in } c_2)}(\text{DO-OP})$$

$$\frac{}{\textbf{if true then } c_1 \textbf{ else } c_2 \rightsquigarrow c_1} \qquad \frac{}{\textbf{if false then } c_1 \textbf{ else } c_2 \rightsquigarrow c_2}$$

$$\frac{}{(\textbf{fun } x \mapsto c)\, v \rightsquigarrow c[v/x]}$$

For every $h = \textbf{handler } \{\textbf{return } x \mapsto c_r, \textsf{op}_1(x;k) \mapsto c_1, \ldots, \textsf{op}_n(x;k) \mapsto c_n\}$, define:

$$\frac{c \rightsquigarrow c'}{\textbf{with } h \textbf{ handle } c \rightsquigarrow \textbf{with } h \textbf{ handle } c'} \qquad \frac{}{\textbf{with } h \textbf{ handle } (\textbf{return } v) \rightsquigarrow c_r[v/x]}$$

$$\frac{}{\textbf{with } h \textbf{ handle } \textsf{op}_i(v;y.c) \rightsquigarrow c_i[v/x, (\textbf{fun } y \mapsto \textbf{with } h \textbf{ handle } c)/k] \quad (1 \le i \le n)}(\text{HANDLED-OP})$$

$$\frac{}{\textbf{with } h \textbf{ handle } \textsf{op}(v;y.c) \rightsquigarrow \textsf{op}(v;y.\textbf{with } h \textbf{ handle } c) \quad (\textsf{op} \notin \{\textsf{op}_1, \ldots, \textsf{op}_n\})}(\text{UNHANDLED-OP})$$

Fig. 2. operational semantics

2013), where we transform the handled computation into a function that passes around a parameter, in our case the state contents:

$$\begin{aligned} H_{ST} := \textbf{handler } \{\, & \textsf{get}(\_;k) \mapsto \textbf{return } (\textbf{fun } s \mapsto (k\, s)\, s) \\ & \textsf{set}(s';k) \mapsto \textbf{return } (\textbf{fun } \_ \mapsto (k\, ())\, s') \\ & \textbf{return } x \mapsto \textbf{return } (\textbf{fun } \_ \mapsto \textbf{return } x)\} \end{aligned}$$

We handle get with a function that takes the current state contents s and in the first application, passes them as a result of get to the continuation. As our handlers are deep, the continuation is further handled into a function, which we again need to supply with the state contents. Since reading does not modify the state, we again pass s. We handle set by first passing the unit result, and then applying the handled continuation to the new state s' as given by the parameter of set. The return clause of H_{ST} also needs to produce a function that depends on the given state, in particular, a function that returns the given value regardless of the state contents.

2.3 Operational semantics

To see how exactly H_{ST} can be used to simulate state, consider the operational semantics of the calculus, also copied verbatim from Pretnar's (2015) tutorial. The semantics is given in terms of the small-step relation $c \rightsquigarrow c'$, defined in Fig. 2. As expected, there is no such relation for values, as these are inert.

The rules for conditionals and function application are standard. For the sequencing construct, $\mathbf{do}\ x \leftarrow c_1\ \mathbf{in}\ c_2$, we start by evaluating c_1. If this returns some value v, we bind it to x and evaluate c_2. But if c_1 calls an operation, we propagate the call outwards and defer further evaluation to the continuation of the call, for example:

$$\mathbf{do}\ x_1 \leftarrow (\mathbf{do}\ x_2 \leftarrow \mathsf{op}(x; y.\, c_2)\ \mathbf{in}\ c_1)\ \mathbf{in}\ c \rightsquigarrow$$
$$\mathbf{do}\ x_1 \leftarrow \mathsf{op}(x; y.\, \mathbf{do}\ x_2 \leftarrow c_2\ \mathbf{in}\ c_1)\ \mathbf{in}\ c \rightsquigarrow$$
$$\mathsf{op}(x; y.\, \mathbf{do}\ x_1 \leftarrow (\mathbf{do}\ x_2 \leftarrow c_2\ \mathbf{in}\ c_1)\ \mathbf{in}\ c)$$

In our account, we gloss over the standard issues with capture-avoiding substitution and implicitly assume the appropriate freshness conditions. For example, in this case, that y is fresh for c_1.

To evaluate $\mathbf{with}\ h\ \mathbf{handle}\ c$, we start by evaluating c. If it returns a value, we continue by evaluating the return clause of h. If c calls an operation op, there are two options. If h has a matching clause for op, we start evaluating that, passing in the parameter and the continuation. Recall that our handlers are deep, thus the continuation k are also handled by the current handler, see HANDLED-OP. If h does not have a matching clause, we forward the call outwards just like in sequencing, see UNHANDLED-OP.

Let us return to the state handler H_{ST}. If we use it on a stateful computation, no effects occur as the handled computation returns a function waiting for an initial state. To run it, we need to apply this function to the initial state. Let us abbreviate such an application by:

$$\langle c, s \rangle := (\mathbf{with}\ H_{ST}\ \mathbf{handle}\ c)\ s$$

(note that we use the syntactic sugar for call-by-value function calls from Fig. 1).

Even though our calculus is pure, we can show the handler H_{ST} simulates global state in the following way. Let $\overset{st}{\rightsquigarrow}$ be the usual small-step semantics for global state, i.e.:

$$\langle \mathsf{get}(), s \rangle \overset{st}{\rightsquigarrow} \langle \mathbf{return}\ s, s \rangle \qquad\qquad \langle \mathsf{set}(s'), s \rangle \overset{st}{\rightsquigarrow} \langle \mathbf{return}\ (), s' \rangle$$

$$\frac{\langle c_1, s \rangle \overset{st}{\rightsquigarrow} \langle c_1', s' \rangle}{\langle \mathbf{do}\ x \leftarrow c_1\ \mathbf{in}\ c_2, s \rangle \overset{st}{\rightsquigarrow} \langle \mathbf{do}\ x \leftarrow c_1'\ \mathbf{in}\ c_2, s' \rangle}$$

etc.

We can prove that for each $\langle c_1, s \rangle \overset{st}{\rightsquigarrow} \langle c_1', s' \rangle$, we have $\langle c_1, s \rangle \rightsquigarrow^+ \langle c_1', s' \rangle$, and therefore effect handlers simulate the operational semantics for global state. For example:

$$\langle \mathsf{get}(), s \rangle \rightsquigarrow (\mathbf{with}\ H_{ST}\ \mathbf{handle}\ (\mathsf{get}((); y.\, \mathbf{return}\ y)))\ s$$
$$\rightsquigarrow (\mathbf{fun}\ s' \mapsto ((\mathbf{fun}\ y \mapsto \mathbf{with}\ H_{ST}\ \mathbf{handle}\ (\mathbf{return}\ y))\ s')\ s')\ s$$
$$\rightsquigarrow ((\mathbf{fun}\ y \mapsto \mathbf{with}\ H_{ST}\ \mathbf{handle}\ (\mathbf{return}\ y))\ s)\ s$$
$$\rightsquigarrow (\mathbf{with}\ H_{ST}\ \mathbf{handle}\ (\mathbf{return}\ s))\ s$$
$$= \langle \mathbf{return}\ s, s \rangle$$

Similarly, we can prove:

$$\langle \mathsf{set}(s'), s \rangle \rightsquigarrow^+ \langle \mathbf{return}\ (), s' \rangle$$

Types

$A, B ::=$		value type
	$\boxed{\alpha}$	type variable
	$\mid$ bool	boolean type
	$\mid$ $A \rightarrow \underline{C}$	function type
	$\mid$ $\underline{C} \Rightarrow \underline{D}$	handler type
$\underline{C}, \underline{D} ::= A\,!\,\boxed{\Sigma}$		computation type
	$\forall \vec{\alpha}.A$	scheme
Σ	$::= \{op_1 : A_1 \rightarrow B_1, \ldots, op_n : A_n \rightarrow B_n\}$	effect signature
Θ	$::= \{\alpha_1, \ldots, \alpha_n\}$	type variable environment
Γ	$::= \varnothing \mid \Gamma, x : A$	monomorphic environment
Ξ	$::= \varnothing \mid \Xi, x : \forall \vec{\alpha}.A$	polymorphic environment

Fig. 3. types and effects

For the third transition, case-split on the possible transitions $\langle c_1, s \rangle \rightsquigarrow \langle c_1^I, s^I \rangle$.

In summary, the H_{ST} handler faithfully simulates state. For more details on simulating state, see Bauer & Pretnar (2014) and Danvy (2006). Therefore, even though our calculus is pure, it faithfully simulates impure computation. By giving an unrestricted Hindley-Milner type system to this calculus, we now show that the effects expressible by effect handlers interact well with polymorphism.

3 Type system

The type and effect system (Figs. 3–4) closely follows Pretnar (2015). It comprises two kinds of types: values are typed with simple types A, while the types of computations are additionally annotated with finite sets of operations Σ like in an effect system of Lucassen & Gifford (1988).

We modify Pretnar's system in two ways. The first modification is minor. We generalise the type system to allow for more flexible *local* operation signatures Σ, where operations may have different types when handled by different handlers, as in Kammar *et al.* (2013). In contrast, Pretnar's account posits a global assignment of predefined types to the effect operations, and the effect annotations Σ only list which operations may be present. Local signatures allow the same operation symbol to appear in disjoint parts of the program with different types. Local signatures also give the calculus stronger theoretical properties, such as strong normalisation and simpler denotational semantics, cf. Kammar *et al.*.

The second modification is our main contribution. We incorporate Hindley-Milner polymorphism in a standard way, without any value restriction. We indicate these latter modifications by *shading* in the figures. Amongst these:

- Local effect signatures Σ are finite mappings from operations op to pairs of value types A, B, whose action we denote by $(op : A \rightarrow B) \in \Sigma$. We denote the restriction of a signature Σ to the set of operations disjoint from a given set $\Delta = \{op_i \mid 1 \le i \le n\}$ by $\Sigma \setminus \Delta$.

Well-formed value types:

$$\frac{\alpha \in \Theta}{\Theta \vdash \alpha} \qquad \frac{}{\Theta \vdash \mathsf{bool}} \qquad \frac{\Theta \vdash A \qquad \Theta \vdash \underline{C}}{\Theta \vdash A \to \underline{C}} \qquad \frac{\Theta \vdash \underline{C} \qquad \Theta \vdash \underline{D}}{\Theta \vdash \underline{C} \Rightarrow \underline{D}}$$

Well-formed computation types, schemes, and effect signatures:

$$\frac{\Theta \vdash A \qquad \Theta \vdash \Sigma}{\Theta \vdash A!\Sigma} \qquad \frac{\Theta, \vec{\alpha} \vdash A}{\Theta \vdash \forall \vec{\alpha}.A} \qquad \frac{[\Theta \vdash A_i \qquad \Theta \vdash B_i]_{1 \le i \le n}}{\Theta \vdash \{\mathsf{op}_1 : A_1 \to B_1, \dots, \mathsf{op}_n : A_n \to B_n\}}$$

Well-formed monomorphic and polymorphic contexts:

$$\frac{[\Theta \vdash A]_{(x:A) \in \Gamma}}{\Theta \vdash \Gamma} \qquad \frac{[\Theta \vdash \forall \vec{\alpha}.A]_{(x:\forall \vec{\alpha}.A) \in \Xi}}{\Theta \vdash \Xi}$$

Value judgements $\boxed{\Theta; \Xi; \Gamma \vdash v : A}$***, assuming*** $\Theta \vdash \Xi, \Gamma, A$***:***

$$\frac{(x:A) \in \Gamma}{\Theta; \Xi; \Gamma \vdash x : A} \qquad \frac{(x : \forall \vec{\alpha}.B) \in \Xi \qquad [\Theta \vdash A_i]_{1 \le i \le |\vec{\alpha}|}}{\Theta; \Xi; \Gamma \vdash x : B[A_i/\alpha_i]_{1 \le i \le |\vec{\alpha}|}} \qquad \frac{}{\Theta; \Xi; \Gamma \vdash \mathbf{true} : \mathsf{bool}}$$

$$\frac{}{\Theta; \Xi; \Gamma \vdash \mathbf{false} : \mathsf{bool}} \qquad \frac{\Theta; \Xi; \Gamma, x : A \vdash c : \underline{C}}{\Theta; \Xi; \Gamma \vdash \mathbf{fun}\, x \mapsto c : A \to \underline{C}}$$

$$\frac{\left[(\mathsf{op}_i : A_i \to B_i) \in \Sigma \quad \begin{array}{c} \Theta; \Xi; \Gamma, x : A \vdash c_r : B!\Sigma' \\ \Theta; \Xi; \Gamma, x : A_i, k : B_i \to B!\Sigma' \vdash c_i : B!\Sigma' \end{array}\right]_{1 \le i \le n}}{\Theta; \Xi; \Gamma \vdash \mathbf{handler}\, \{\mathbf{return}\, x \mapsto c_r, \mathsf{op}_1(x;k) \mapsto c_1, \dots, \mathsf{op}_n(x;k) \mapsto c_n\} : A!\Sigma \Rightarrow B!\Sigma'}$$

with side condition $\Sigma \setminus \{\mathsf{op}_i \mid 1 \le i \le n\} \subseteq \Sigma'$.

Computation judgements $\boxed{\Theta; \Xi; \Gamma \vdash c : A!\Sigma}$***, assuming*** $\Theta \vdash \Xi, \Gamma, A$***:***

$$\frac{\Theta; \Xi; \Gamma \vdash v : A}{\Theta; \Xi; \Gamma \vdash \mathbf{return}\, v : A!\Sigma} \qquad \frac{\Theta; \Xi; \Gamma \vdash c_1 : (\forall \vec{\alpha}.A)\,!\Sigma \qquad \Theta; \Xi, x : \forall \vec{\alpha}.A; \Gamma \vdash c_2 : B!\Sigma}{\Theta; \Xi; \Gamma \vdash \mathbf{do}\, x \leftarrow c_1\, \mathbf{in}\, c_2 : B!\Sigma}$$

$$\frac{(\mathsf{op} : A_{\mathsf{op}} \to B_{\mathsf{op}}) \in \Sigma \qquad \Theta; \Xi; \Gamma \vdash v : A_{\mathsf{op}} \qquad \Theta; \Xi; \Gamma, y : B_{\mathsf{op}} \vdash c : A!\Sigma}{\Theta; \Xi; \Gamma \vdash \mathsf{op}(v; y.c) : A!\Sigma}$$

$$\frac{\Theta; \Xi; \Gamma \vdash v : \mathsf{bool} \qquad \Theta; \Xi; \Gamma \vdash c_1 : \underline{C} \qquad \Theta; \Xi; \Gamma \vdash c_2 : \underline{C}}{\Theta; \Xi; \Gamma \vdash \mathbf{if}\, v\, \mathbf{then}\, c_1\, \mathbf{else}\, c_2 : \underline{C}}$$

$$\frac{\Theta; \Xi; \Gamma \vdash v_1 : A \to \underline{C} \qquad \Theta; \Xi; \Gamma \vdash v_2 : A}{\Theta; \Xi; \Gamma \vdash v_1\, v_2 : \underline{C}} \qquad \frac{\Theta; \Xi; \Gamma \vdash v : \underline{C} \Rightarrow \underline{D} \qquad \Theta; \Xi; \Gamma \vdash c : \underline{C}}{\Theta; \Xi; \Gamma \vdash \mathbf{with}\, v\, \mathbf{handle}\, c : \underline{D}}$$

Scheme judgement $\boxed{\Theta; \Xi; \Gamma \vdash c : (\forall \vec{\alpha}.A)!\Sigma}$***, assuming*** $\Theta \vdash \Xi, \Gamma, (\forall \vec{\alpha}.A), \Sigma$***:***

$$\frac{\Theta, \vec{\alpha}; \Xi; \Gamma \vdash c : A!\Sigma}{\Theta; \Xi; \Gamma \vdash c : (\forall \vec{\alpha}.A)!\Sigma} \text{(GEN)}$$

Fig. 4. a polymorphic type and effect system

- We extend types with *type variables* α and add *type variable environments* Θ, which are just finite sets of type variables.
- We introduce *schemes* $\forall \vec{\alpha}.A$, where $\vec{\alpha}$ denotes a finite set of $|\vec{\alpha}|$-many type variables ranged over by α_i.
- We introduce *kinding judgements* $\Theta \vdash X$ to explicitly keep track of the free type variables in X. The shorthand $\Theta \vdash X, Y, Z$ stands for the conjunction $\Theta \vdash X, \Theta \vdash Y$, and $\Theta \vdash Z$.
- Typing judgements $\Theta; \Xi; \Gamma \vdash M : X$ include the standard monomorphic environments Γ which are a unique assignment of types to variables. We extend those with type variable environments Θ and *polymorphic environments* Ξ, which are a unique assignment of schemes to variables. We assume that no variable can appear in both Γ and Ξ.[1] These polymorphic variables can be specialised at any type.
- We add *scheme judgements* whose effect annotation is outside the scope of the quantifier. The kinding assumption $\Theta \vdash \Sigma$ ensures that none of the type variables $\vec{\alpha}$ appears in Σ. It is at this point that the decision of the inference algorithm which type variables $\vec{\alpha}$ to generalise over takes effect. Our choice to separate scheme judgements from type judgements simplifies the let-rule, and makes it very similar to its standard, monomorphic counterpart.

The remaining kinding and typing rules are standard. Fine-grained call-by-value functions take values and perform computations. An operation invocation is well-typed if the type assigned to it by the local signature must agree with the type of the given parameter value v, and with the type of argument the continuation c expects. A handler is well-typed if the type of result the return clause expects matches with the type of computation the handler can handle, and each operation clause is well-typed when the parameter type and continuation type match the local signature the handler can handle. Both clauses can cause additional effects, and their effect annotation must include these operations, as well as any effect operations the handler does not explicitly handle, reflecting the fact that our handlers are forwarding. Thus, the rule also requires the type and effect of both clauses to agree. The fact that our handlers are deep is reflected by the type of the continuation: the effects the continuation may cause have already been handled, and so the continuation may cause effects in the resulting signature and of the resulting return type.

For the given effect system, we then have:

Theorem (Safety). *If $\vdash c : A \mathbin{!} \Sigma$ holds, then either:*

 (i) $c \rightsquigarrow c'$ *for some* $\vdash c' : A \mathbin{!} \Sigma$;
 (ii) $c = \textbf{return } v$ *for some* $\vdash v : A$; *or*
 (iii) $c = \mathsf{op}(v; y.c')$ *for some* $(\mathsf{op} : A_{\mathsf{op}} \to B_{\mathsf{op}}) \in \Sigma$, $\vdash v : A_{\mathsf{op}}$, *and* $y : B_{\mathsf{op}} \vdash c' : A \mathbin{!} \Sigma$.

In particular, when $\Sigma = \varnothing$, evaluation will not get stuck before returning a value.

Proof

---

[1] This separation into two environments is not strictly necessary, as a monomorphic environment Γ may be identified with a polymorphic environment where each quantifier ranges over an empty tuple of type variables. We choose to separate the two to highlight which parts of the language interact with polymorphism.

We prove progress and preservation lemmata separately by induction. We formalized[2] the calculus and the safety theorem in the Twelf proof assistant (Pfenning & Schürmann, 1999). Our formalization extends Bauer & Pretnar's (2014) existing formalization of Eff's core calculus with type schemes and polymorphism. The code is compatible with version 1.7.1 of Twelf. We summarise the crucial step, namely proving type and effect preservation under the DO-OP transition.

Assume that the reduct in DO-OP is well-typed, and invert its type derivation:

$$
\cfrac{
\cfrac{
(\mathsf{op}:A_{\mathsf{op}} \to B_{\mathsf{op}}) \in \Sigma \quad
\cfrac{\vdots}{\Theta, \vec{\alpha} \vdash v : A_{\mathsf{op}}} \quad
\cfrac{\vdots}{\Theta, \vec{\alpha}; y : B_{\mathsf{op}} \vdash c_1 : A\,!\,\Sigma}
}{
\cfrac{\Theta, \vec{\alpha} \vdash \mathsf{op}(v; y.\,c_1) : A\,!\,\Sigma}{\Theta \vdash \mathsf{op}(v; y.\,c_1) : (\forall \vec{\alpha}.A)\,!\,\Sigma}
} \quad
\cfrac{\vdots}{\Theta; x : \forall \vec{\alpha}.A \vdash c_2 : B\,!\,\Sigma}
}{
\Theta \vdash \mathbf{do}\ x \leftarrow \mathsf{op}(v; y.\,c_1)\ \mathbf{in}\ c_2 : B\,!\,\Sigma
}
$$

The GEN rule ensures that none of the type variables in $\vec{\alpha}$ appear in Σ. Because Σ includes $\mathsf{op} : A_{\mathsf{op}} \to B_{\mathsf{op}}$, none of these variables appear in A_{op}, and we may strengthen the derivation of $\Theta, \vec{\alpha} \vdash v : A_{\mathsf{op}}$ to a derivation of $\Theta \vdash v : A_{\mathsf{op}}$. As a consequence, the following derivation is valid:

$$
\cfrac{
(\mathsf{op}:A_{\mathsf{op}} \to B_{\mathsf{op}}) \in \Sigma \quad
\cfrac{\vdots}{\Theta \vdash v : A_{\mathsf{op}}} \quad
\cfrac{
\cfrac{
\cfrac{\vdots}{\Theta, \vec{\alpha}; y : B_{\mathsf{op}} \vdash c_1 : A\,!\,\Sigma}
}{
\Theta; y : B_{\mathsf{op}} \vdash c_1 : (\forall \vec{\alpha}.A)\,!\,\Sigma
} \quad
\cfrac{\vdots}{\Theta; x : \forall \vec{\alpha}.A \vdash c_2 : B\,!\,\Sigma}
}{
\Theta; y : B_{\mathsf{op}} \vdash \mathbf{do}\ x \leftarrow c_1\ \mathbf{in}\ c_2 : B\,!\,\Sigma
}
}{
\Theta \vdash \mathsf{op}(v; y.\,\mathbf{do}\ x \leftarrow c_1\ \mathbf{in}\ c_2) : B\,!\,\Sigma
}
$$

Therefore, the reduction in DO-OP preserves both the type and the effect annotation. ∎

The Safety Theorem is robust under the following standard variations in the calculus:

coarse annotations. We can make the signature Σ global, and only keep track of which operations are used, as in Pretnar (2015). The types in this global signature cannot use any type variables. The soundness proof remains essentially unchanged[3]. Due to the lack of type variables in the global signature, there is no need to impose a side-condition on the well-formedness of the effect annotation in the GEN rule

It may seem this coarser system is a restriction of our current system, where the type information for each operation has to agree in all effect annotations, and hence it is sound by the Safety Theorem. This is not the case. In this coarser system, the signatures on function types are not annotated with the types of the operations. If those types were fully written out, they would involve the global signature, leading to potential mutual recursion between signatures and function types. For example, if we elaborate the global signature $\Sigma = \{\mathsf{op} : \mathsf{unit} \to (\mathsf{unit} \to \mathsf{unit})\}$, we would get:

$$
\Sigma = \{\mathsf{op} : \mathsf{unit} \to (\mathsf{unit} \to (\mathsf{unit}\,!\,\Sigma))\}
$$

---

[2] `https://github.com/matijapretnar/twelf-eff/tree/val-restriction-local-sig`

[3] `https://github.com/matijapretnar/twelf-eff/tree/val-restriction-global-sig`

where left arrow is part of the signature syntax and receives no effect annotation on the co-domain. This recursion is not a mere formality. The type-and-effect system with local signatures we have described ensures well-typed terms terminate, cf. Kammar *et al.* (2013). When we switch to a global signature, we can use effect operations with higher-order return types to express well-typed diverging computations. With the above global signature $\Sigma = \{\mathsf{op} : \mathsf{unit} \to (\mathsf{unit} \to \mathsf{unit})\}$, consider the handler

$$H := \mathbf{handler}\,\{\mathbf{return}\,x \mapsto \mathbf{return}\,x,$$
$$\mathsf{op}(\_;k) \mapsto k(\mathbf{fun}\,\_ \mapsto \mathsf{op}()())\}$$

In the coarse type system, we can derive the judgement:

$$\vdash H : (\mathsf{unit}\,!\,\{\mathsf{op}\}) \Rightarrow (\mathsf{unit}\,!\,\varnothing)$$

If we handle the simple looking computation $\vdash \mathsf{op}()() : \mathsf{unit}\,!\,\{\mathsf{op}\}$ with H, we get a diverging computation:

$$\mathbf{with}\,H\,\mathbf{handle}\,\mathsf{op}()() \rightsquigarrow^{+} \mathbf{with}\,H\,\mathbf{handle}\,\big(\mathbf{fun}\,\_ \mapsto \mathsf{op}()()\big)()$$
$$\rightsquigarrow \mathbf{with}\,H\,\mathbf{handle}\,\mathsf{op}()()$$

In fact, by a variation on Landin's (1964) knot, we can express a variant of the *Y*-combinator, such that for a function f that is pure, $Y f$ behaves like the fixed-point of f when invoked on pure arguments.

no annotations. We can remove all the effect annotations Σ from type judgements and fix a single, global signature Σ. The advantage of having an effect system is the additional guarantee in clause *(iii)* of the Safety Theorem, which ensures that any unhandled operation must appear in Σ. Without annotations, any operation may be called. This system is a restriction of the coarse variation, where each effect annotation is the entire signature. Consequently, it is sound.

additional language features. To the calculus with coarse annotations, we can add *structural subtyping* and static *effect instances* (further discussed in Sec. 4.2). The soundness proof remains essentially unchanged[4] as these modifications are orthogonal to polymorphism. Similarly, we can replace deep handlers with shallow ones, as in Kammar *et al.* (2013) and Kiselyov *et al.* (2013). As the changes[5] are again orthogonal to polymorphism, we may reasonably assume a similar soundness result to hold for a calculus that incorporates all of the above: subtyping, instances, and, through two separate syntactic constructs, both deep and shallow handlers.

4 Expressivity

There is currently no simple type system integrating reference cells with polymorphism without the value restriction. This non-existence contrasts the simplicity of our type system, and calls into question both its degree of feature integration and its expressiveness. First, we evaluate the degree and smoothness of the interaction between polymorphism and

---

[4] `https://github.com/matijapretnar/twelf-eff/tree/val-restriction-instances`
[5] `https://github.com/matijapretnar/twelf-eff/tree/val-restriction-shallow`

other features in our calculus. Then, we highlight the difference in expressiveness between effect handlers and reference cells. As a basis for our evaluation and comparison, we use Leroy's (1992) set of example programs for analysing the usefulness of a polymorphic type system for reference cells.

4.1 Evaluation

Algebraic effects allow us to lace a piece of code with operations in the signature

$$\{\mathsf{get} : \mathsf{unit} \to \alpha, \mathsf{set} : \alpha \to \mathsf{unit}\}$$

The scheme assigned to the handler H_{ST}, which handles them away, is

$$H_{ST} : \forall \alpha, \beta.\alpha \, ! \, \{\mathsf{get} : \mathsf{unit} \to \beta, \mathsf{set} : \beta \to \mathsf{unit}\} \Rightarrow (\beta \to \alpha \, ! \, \varnothing) \, ! \, \varnothing$$

It takes a computation of type α that interacts with a state of type β, and handles it to a pure function of type $\beta \to \alpha \, ! \, \varnothing$. The rightmost $\varnothing$ indicates that no effects can occur when producing the function.

This handler can handle computations with different types of state, for example:

$$(\textbf{with } H_{ST} \textbf{ handle set } ()) \, ();$$
$$(\textbf{with } H_{ST} \textbf{ handle get } ()) \, \textbf{true}$$

We can also use effects in polymorphic code:

$$\textbf{do } f \leftarrow \textbf{if } \mathsf{get} \, () \textbf{ then return fun } x \, y \mapsto \textbf{return } x$$
$$\textbf{else return fun } x \, y \mapsto \textbf{return } y$$
$$\textbf{in } (f \, (\textbf{fun } b \mapsto \textbf{return } b)$$
$$(\textbf{fun } b \mapsto \mathsf{set} \, b; \textbf{return } b))$$
$$(f \, \textbf{true false})$$

In our call-by-value semantics, if we wrap this computation with the state handler, the memory look-up in f's definition will only occur once.

To demonstrate that the polymorphic, effectful, and high-order features interact well, we hypothetically extend our calculus with pairs and lists. The hypothesised extension may include primitives such as the empty list $[\,]$, a list cons $(::)$ and tail-recursive iteration **foldl**, which we expect to interact smoothly with polymorphism. Thus we can use H_{ST} to implement functional features in an imperative style.

$$\textbf{do } \mathsf{imp\_map} \leftarrow \textbf{fun } f \, xs \mapsto$$
$$\textbf{with } H_{ST} \textbf{ handle } (\textbf{foldl } \, (\textbf{fun } x \mapsto \mathsf{set} \, (f \, x :: \mathsf{get} \, ()))$$
$$()$$
$$xs;$$
$$\mathsf{reverse} \, (\mathsf{get} \, ())$$
$$[\,] \, (* \text{ initial state } *) \textbf{ in } \dots$$

The scheme assigned to imp_map is

$$\mathsf{imp\_map} : \forall \alpha \beta.(\alpha \to \beta \, ! \, \Sigma) \to (\alpha \, \mathsf{list} \to \beta \, \mathsf{list} \, ! \, \Sigma) \, ! \, \varnothing$$

for any Σ. This implementation is imperative in style, but not imperative *per se*, as all operations are handled by high-order functions. The function imp_map can also be partially applied and retain its polymorphism, for example, in

$$\mathbf{do}\ \text{list\_id} \leftarrow \text{imp\_map id}\ \mathbf{in}$$
$$\mathbf{do}\ \text{nil} \quad \leftarrow \text{list\_id}\ [\,]\ \mathbf{in} \dots$$

we have the scheme assignments:

$$\text{list\_id} : \forall \alpha.\alpha\ \text{list} \rightarrow \alpha\ \text{list}\,!\,\varnothing$$
$$\text{nil} \quad : \forall \alpha.\alpha\ \text{list}$$

Most importantly, the following program is well-typed:

$$\mathbf{do}\ id \leftarrow (\mathbf{fun}\ f \mapsto f)\ (\mathbf{fun}\ x \mapsto x)\ \mathbf{in}$$
$$\mathbf{do}\ id' \leftarrow id\,(id)\ \mathbf{in} \dots$$

and both functions are assigned the polymorphic type $\forall \alpha.\alpha \rightarrow \alpha\,!\,\Sigma$. Such mixed-variance polymorphism is ruled out by all current value restrictions.

4.2 Reference cells

We believe it is impossible to implement full blown reference cells using effect handlers without other language features. We can increase modularity by introducing instances (Bauer & Pretnar, 2015, 2014; Pretnar, 2014). These may be thought of as first class atomic names. With instances, each effect instance ι and an operation symbol op determine an operation ι#op. In handlers, each operation clause v#op$(x;k) \mapsto c$ specifies which instance, dynamically given by the value v, of the statically chosen effect operation symbol op the handler handles. At runtime, invocations of the same operation op but with different instances will not be caught by this handler and will be forwarded.

Instances allow us to pass a cell around by passing an instance, but they are still less expressive than having the ability to allocate arbitrarily many new cells dynamically. For example, we do not know how to implement even the simplest of Leroy's (1992) benchmarks:

$$\mathbf{do}\ \text{make\_ref} \leftarrow \mathbf{fun}\ x \mapsto \mathbf{ref}\ x\ \mathbf{in} \dots$$

We believe it is impossible to encode general references without additional language features. *Eff* provides such a mechanism, which can both generate *fresh* instances and attach them to a stateful *resource* (Bauer & Pretnar, 2015), allowing one to directly implement a make_ref analogue: make_ref creates a fresh instances that has get and set operations associated with it. Only code that knows what the instance is, can handle these effects. However, it is not easy to find a corresponding type and effect system for fresh instances (Bauer & Pretnar, 2014; Pretnar, 2014), let alone a polymorphic one.

As a final example, recall the problematic reference cell example which cannot be directly expressed in our calculus:

$$\mathbf{do}\ r \leftarrow \mathbf{ref}\ [\,]\ \mathbf{in}$$
$$r := [()];$$
$$\mathbf{true} :: !r$$

We can express a computation that writes a unit list value and reads a bool list value:

$$\mathsf{set}\,[()];$$
$$\mathbf{true} :: \mathsf{get}\,()$$

However, this computation has the effect annotation

$$\{\mathsf{set} : \mathsf{unit\,list} \to \mathsf{unit}, \mathsf{get} : \mathsf{unit} \to \mathsf{bool\,list}\}$$

which is incompatible with the type of the state handler H_{ST}. Other handlers for the state operations may have a compatible type. For example, the read-only state handler H_{RO} which ignores any memory updates:

$$
\begin{aligned}
H_{RO} := \mathbf{handler}\,\{&\mathbf{return}\,x \mapsto \mathbf{fun}\,\_ \mapsto \mathbf{return}\,x \\
&\mathsf{get}(\_;k) \mapsto \mathbf{fun}\,s \mapsto k\,s\,s \\
&\mathsf{set}(\_;k) \mapsto \mathbf{fun}\,s \mapsto k\,()\,s\}
\end{aligned}
$$

It has the scheme

$$H_{RO} : \forall \alpha, \beta, \gamma.\,\alpha\,!\{\mathsf{get} : \mathsf{unit} \to \beta, \mathsf{set} : \gamma \to \mathsf{unit}\} \Rightarrow (\beta \to \alpha\,!\varnothing)\,!\varnothing$$

and can be applied to the above computation without run-time errors.

4.3 Dynamically scoped state

As we saw in Sec. 2.2, we can simulate global state using the handler H_{ST}, and this state can be handled locally to give a pure computation. While we do not know whether effect handlers can simulate reference cells or not, we will now characterise the handler H_{ST} as expressing the notion of *dynamically scoped* state.

In order to explain what we mean by dynamically scoped state, and to make the discussion precise, we consider the calculus presented in Fig. 5. It is a fine-grained call-by-value variation on the dynamic scope calculi of Kiselyov *et al.* (2006) and Moreau (1998).

We assume a set of parameters ranged over by p that name dynamically scoped memory cells. These cells can be dereferenced, $!p$, or assigned to, $p := v$, just like ref cells. The rebinding construct **dlet** $p \leftarrow v$ **in** c declares that in executing c, all references to p will be bound to this occurrence of p, and shadow other binding declarations that may be in place.

For example, assuming we have a type of integers the following code will evaluate to **return** 2.

$$
\begin{aligned}
\mathbf{do}\,f \leftarrow\,&\mathbf{dlet}\,\mathsf{p} \leftarrow 0\,\mathbf{in} \\
&\quad\mathbf{return}\,(\mathbf{fun}\,\_ \mapsto \\
&\qquad\qquad \mathsf{p} := 1 + !\mathsf{p})\,\mathbf{in} \\
\mathbf{dlet}\,\mathsf{p} \leftarrow\,&1\,\mathbf{in} \\
&f\,(); \\
&!\mathsf{p}
\end{aligned}
$$

The reason is that the state changes inside the function bind dynamically to the closest enclosing rebinding, which is the second one.

Fig. 6 describes the (Felleisen-style) operational semantics for this calculus. We kept the style of semantics as close as possible to Kiselyov *et al.*'s (2006) to make it clear we use

Syntax

p	$::=$	$\texttt{p} \mid \texttt{q} \mid \texttt{r} \mid \dots$	parameter
v	$::=$		value
		x	variable
	$\mid$	**true** $\mid$ **false**	boolean constants
	$\mid$	$()$	unit value
	$\mid$	**fun** $x \mapsto c$	function
c	$::=$		computation
		return v	return
	$\mid$	**do** $x \leftarrow c_1$ **in** c_2	sequencing
	$\mid$	**if** v **then** c_1 **else** c_2	conditional
	$\mid$	$v_1 \, v_2$	application
	$\mid$	$!p$	dereferencing
	$\mid$	$p := v$	assignment
	$\mid$	**dlet** $p \leftarrow v$ **in** c	rebinding

Fig. 5. a calculus for dynamically scoped state

Auxiliary definitions

Evaluation contexts:

$$E \quad ::= \quad [\ ] \mid E[\textbf{do } x \leftarrow [\ ] \textbf{ in } c] \mid E[\textbf{dlet } p \leftarrow v \textbf{ in } [\ ]]$$

Parameter binding:

$$\text{bp}([\ ]) := \varnothing \qquad \text{bp}(E[\textbf{do } x \leftarrow [\ ] \textbf{ in } c]) := \text{bp}(E) \qquad \text{bp}(E[\textbf{dlet } p \leftarrow v \textbf{ in } [\ ]]) := \text{bp}(E) \cup \{p\}$$

Semantics

$$E[\textbf{do } x \leftarrow \textbf{return } v \textbf{ in } c] \overset{dyn}{\rightsquigarrow} E[c[v/x]] \qquad E[\textbf{if true then } c_1 \textbf{ else } c_2] \overset{dyn}{\rightsquigarrow} E[c_1]$$

$$E[\textbf{if false then } c_1 \textbf{ else } c_2] \overset{dyn}{\rightsquigarrow} E[c_2] \qquad E[(\textbf{fun } x \mapsto c) \, v] \overset{dyn}{\rightsquigarrow} E[c[v/x]]$$

$$E[\textbf{dlet } p \leftarrow v \textbf{ in return } v'] \overset{dyn}{\rightsquigarrow} E[\textbf{return } v']$$

$$\frac{}{E[\textbf{dlet } p \leftarrow v \textbf{ in } E'[!p]] \overset{dyn}{\rightsquigarrow} E[\textbf{dlet } p \leftarrow v \textbf{ in } E'[\textbf{return } v]]} (p \notin \text{bp}(E'))$$

$$\frac{}{E[\textbf{dlet } p \leftarrow v \textbf{ in } E'[p := v']] \overset{dyn}{\rightsquigarrow} E[\textbf{dlet } p \leftarrow v' \textbf{ in } E'[\textbf{return } ()]]} (p \notin \text{bp}(E'))$$

Fig. 6. semantics for dynamically scoped state

Term-level translation

$$\lceil x \rceil := x \qquad \lceil \mathbf{true} \rceil := \mathbf{true} \qquad \lceil \mathbf{false} \rceil := \mathbf{false} \qquad \lceil \mathbf{fun}\ x \mapsto c \rceil := \mathbf{fun}\ x \mapsto \lceil c \rceil$$

$$\lceil v_1\ v_2 \rceil := \lceil v_1 \rceil \lceil v_2 \rceil \qquad \lceil \mathbf{return}\ v \rceil := \mathbf{return}\ \lceil v \rceil \qquad \lceil \mathbf{do}\ x \leftarrow c_1\ \mathbf{in}\ c_2 \rceil := \mathbf{do}\ x \leftarrow \lceil c_1 \rceil\ \mathbf{in}\ \lceil c_2 \rceil$$

$$\lceil !p \rceil := \mathtt{get\_}p() \qquad \lceil p := v \rceil := \mathtt{set\_}p(\lceil v \rceil) \qquad \lceil \mathbf{dlet}\ p \leftarrow v\ \mathbf{in}\ c \rceil := (\mathbf{with}\ H^p_{ST} \lceil c \rceil\ \mathbf{handle}\ )\ \lceil v \rceil$$

where:

$$H^p_{ST} := \mathbf{handler}\ \{\mathtt{get\_}p(\_;k) \mapsto \mathbf{return}\ (\mathbf{fun}\ s \mapsto (k\ s)\ s)$$
$$\mathtt{set\_}p(s';k) \mapsto \mathbf{return}\ (\mathbf{fun}\ \_ \mapsto (k\ ())\ s')$$
$$\mathbf{return}\ x \qquad \mapsto \mathbf{return}\ (\mathbf{fun}\ \_ \mapsto \mathbf{return}\ x)\}$$

Fig. 7. handlers expressing dynamically scoped state

the same notion of dynamic scope, and our theoretical treatment closely mirrors their own. The semantics use the set of parameters bound in a given context E, denoted by $\mathrm{bp}(E)$. The three transitions specific to dynamic scope are shaded. First, a fully evaluated computation removes a preceding parameter binding, as it will no longer be used. For the other two transitions, the side condition $p \notin \mathrm{bp}(E')$ ensures the uniqueness of the decomposition into the context E' by locating the closest rebinding of p. The semantics of dereferencing returns the value associated to this closest rebinding, while the semantics of assignment modifies it. In our design, assignment evaluates to the unit value, deviating from Kiselyov *et al.*'s semantics. This purely cosmetic change does not alter the nature of dynamically scope state we are dealing with, and makes the relationship with H_{ST} tighter.

The example above evaluates as follows:

$$
\begin{aligned}
&\mathbf{do}\ f \leftarrow \mathbf{dlet}\ \mathrm{p} \leftarrow 0\ \mathbf{in} \\
&\qquad\qquad \mathbf{return}\ (\mathbf{fun}\ \_ \mapsto \\
&\qquad\qquad\qquad \mathrm{p} := 1 + !\mathrm{p})\ \mathbf{in} \quad \overset{dyn}{\rightsquigarrow} \\
&\quad \mathbf{dlet}\ \mathrm{p} \leftarrow 1\ \mathbf{in} \\
&\qquad f\ (); \\
&\qquad !\mathrm{p}
\end{aligned}
\qquad
\begin{aligned}
&\mathbf{do}\ f \leftarrow \mathbf{return}\ (\mathbf{fun}\ \_ \mapsto \\
&\qquad\qquad\qquad \mathrm{p} := 1 + !\mathrm{p})\ \mathbf{in} \\
&\quad \mathbf{dlet}\ \mathrm{p} \leftarrow 1\ \mathbf{in} \\
&\qquad f\ (); \\
&\qquad !\mathrm{p}
\end{aligned}
$$

$$
\overset{dyn}{\rightsquigarrow}
\begin{aligned}
&\mathbf{dlet}\ \mathrm{p} \leftarrow 1\ \mathbf{in} \\
&\quad (\mathbf{fun}\ \_ \mapsto \\
&\qquad \mathrm{p} := 1 + !\mathrm{p})\ (); \\
&\quad !\mathrm{p}
\end{aligned}
\quad \overset{dyn}{\rightsquigarrow}
\begin{aligned}
&\mathbf{dlet}\ \mathrm{p} \leftarrow 1\ \mathbf{in} \\
&\quad \mathrm{p} := 1 + !\mathrm{p}; \\
&\quad !\mathrm{p}
\end{aligned}
\quad \overset{dyn+}{\rightsquigarrow} \quad \mathbf{return}\ 2
$$

Fig. 7 shows how effect handlers express dynamically scoped state. Using Felleisen's (1990) terminology, it is a *macro* translation. First, it does not use any information collected globally as it is defined homomorphically over the syntax of the language. Second, it keeps the common core of the two languages unchanged, translating a boolean value to itself, a function to a function, and so forth. The translation is straightforward: it translates dereferencing and assignments to p as specially named effects, $\mathtt{get\_}p$ and $\mathtt{set\_}p$. Rebinding

amounts to handling with H_{ST}, and passing the translated rebinding value as the initial value.

This translation simulates dynamic allocation:

Theorem (Simulation). *For all $c \overset{dyn}{\leadsto} c'$, we have $\lceil c \rceil \leadsto^+ \lceil c' \rceil$.*

Proof
First, extend the translations to evaluation contexts, and show that $\lceil E[c] \rceil = \lceil E \rceil[\lceil c \rceil]$. Then, show the translation respects capture avoiding substitution: $\lceil c[v/x] \rceil = \lceil c \rceil[\lceil v \rceil/x]$. To deal with the mismatch between Felleisen-style and small-step semantics, show that for all evaluation contexts E, if $c \overset{dyn}{\leadsto} c'$ then $\lceil E \rceil[c] \leadsto^+ \lceil E \rceil[c']$. It therefore suffices to prove the theorem for each of the transitions in Fig. 6 specialised to $E := [\ ]$.

For each of the common constructs of the two calculi, the proof is immediate, for example:

$$\lceil \mathbf{do}\ x \leftarrow \mathbf{return}\ v\ \mathbf{in}\ c \rceil = \mathbf{do}\ x \leftarrow \mathbf{return}\ \lceil v \rceil\ \mathbf{in}\ \lceil c \rceil \leadsto \lceil c \rceil[\lceil v \rceil/x] = \lceil c[v/x] \rceil$$

The next remaining transition amounts to handling a terminal computation:

$$\lceil \mathbf{dlet}\ \mathrm{p} \leftarrow v\ \mathbf{in}\ \mathbf{return}\ v' \rceil = (\mathbf{with}\ H_{ST}^p\ \mathbf{handle}\ \mathbf{return}\ \lceil v' \rceil)\ \lceil v \rceil$$
$$\leadsto^+ (\mathbf{fun}\ \_ \mapsto \mathbf{return}\ \lceil v' \rceil)\ \lceil v \rceil \leadsto \mathbf{return}\ \lceil v' \rceil$$

For the final two transition, show that, for all contexts E, parameters $p \notin \mathrm{bp}(E)$, operations op that is either $\mathtt{get\_}p$ or $\mathtt{set\_}p$, and x fresh for E, we have:

$$\lceil E \rceil[\mathsf{op}(v; x.c)] \leadsto^* \mathsf{op}(v; x.\lceil E \rceil[c])$$

And finally, calculate:

$$\lceil \mathbf{dlet}\ p \leftarrow v\ \mathbf{in}\ E[!p] \rceil = (\mathbf{with}\ H_{ST}^p\ \mathbf{handle}\ \lceil E \rceil[\mathtt{get\_}p((); x.\mathbf{return}\ x))\ \lceil v \rceil$$
$$\leadsto^* (\mathbf{with}\ H_{ST}^p\ \mathbf{handle}\ \mathtt{get\_}p((); x.\lceil E \rceil[\mathbf{return}\ x]))\ \lceil v \rceil$$
$$\leadsto^+ (\mathbf{fun}\ s \mapsto ((\mathbf{fun}\ x \mapsto \mathbf{with}\ H_{ST}^p\ \mathbf{handle}\ \lceil E \rceil[\mathbf{return}\ x])\ s)\ s)\ \lceil v \rceil$$
$$\leadsto^+ \mathbf{with}\ H_{ST}^p\ \mathbf{handle}\ \lceil E \rceil[\mathbf{return}\ x]\ \lceil v \rceil$$
$$= \lceil \mathbf{dlet}\ p \leftarrow v\ \mathbf{in}\ E[\mathbf{return}\ v] \rceil$$

A similar calculation for assignment completes the proof. ∎

This translation, while being straightforward, also preserves the type system. Fig. 8 presents the types for the calculus. The only notable feature is that, like Kiselyov *et al.*, we assume a global signature assigning to each parameter a type. As the signature is global, these (monomorphic) types do not contain any type variables.

Fig. 9 presents the kind and (Hindley-Milner polymorphic) type system for the calculus. The kind system ensures well-kinded signatures assign types without type variables. Typing judgements $\Theta; \Xi; \Gamma \vdash_{\Sigma}^{\mathrm{dyn}} c : A$ refer to the fixed, ambient, well-kinded parameter signature Σ. The typing rules specific to dynamically scoped state (shaded) ensure that we may only dereference, assign to, and rebind a parameter in accordance with the ambient signature. The assignment rule also highlights our decision to ascribe the unit type to assignment, in a minor deviation from Kiselyov *et al.*. The (GEN) rule is now completely

---

Types

A, B	$::=$		value type
		α	type variable
	$\mid$	bool	boolean type
	$\mid$	unit	unit type
	$\mid$	$A \to B$	function type
		$\forall \vec{\alpha}.A$	scheme
Σ	$::=$	$\{p_1 : A_1, \ldots, p_n : A_n\}$	parameter signature
Θ	$::=$	$\{\alpha_1, \ldots, \alpha_n\}$	type variable environment
Γ	$::=$	$\varnothing \mid \Gamma, x : A$	monomorphic environment
Ξ	$::=$	$\varnothing \mid \Xi, x : \forall \vec{\alpha}.A$	polymorphic environment

---

Fig. 8. polymorphic types for dynamically scoped state

unrestricted, ensured by the assumption that the type signature does not involve type variables.

Fig. 10 extends the translation to types. The parameter signature Σ translates into an effect signature containing the distinct pair of effects corresponding to this parameter, namely get_p and set_p, with the appropriate type. Function types may cause any effect in this translated signature $\lceil \Sigma \rceil$. This translation is therefore not-well-defined: if Σ contains any function types, then $\lceil \Sigma \rceil$ refers to $\lceil A \to B \rceil$, which refers to $\lceil \Sigma \rceil$ again.

There are at least three ways around this issue. The simplest solution, presented in the top half of Fig. 10 is to restrict Σ to *ground* types, i.e., prohibit storing functions.

A less restrictive solution is to use the coarser type system for effect handlers that does not track effect annotations at all, and define $\lfloor A \to B \rfloor := \lfloor A \rfloor \to \lfloor B \rfloor$, as in the bottom half of Fig. 10. This solution works well, as the effects get_p and put_p maintain their type.

A more sophisticated potential solution is to use equi-recursive effect signatures. At this point in time, such a type-and-effect system has not been developed, but we do not foresee any serious obstacles in developing it: its denotational semantics would involve a recursive domain equation in the same spirit as in Bauer & Pretnar (2014).

The fact that higher-order parameter types merit domain-theoretic semantics is not surprising, as such parameters allow non-terminating programs. We say that a type A is *inhabited* if there exists a closed value $\vdash_{\Sigma}^{\mathrm{dyn}} v : A$.

Proposition. *If Σ contains a higher-order type parameter $(p : A \to B) \in \Sigma$ for some inhabited type A, then there is a term c satisfying:*

$$c \overset{dyn+}{\rightsquigarrow} c$$

Proof
Let $\vdash_{\Sigma}^{\mathrm{dyn}} v : A$ be an inhabitant of *A*, and take:

$$c := \mathbf{dlet}\, p \leftarrow (\mathbf{fun}\, a \mapsto (!p)\, a)\ \mathbf{in}$$
$$(!p)\, v$$

Well-formed types, parameter signatures, and schemes:

$$\frac{\alpha \in \Theta}{\Theta \vdash^{\mathrm{dyn}} \alpha} \qquad \frac{}{\Theta \vdash^{\mathrm{dyn}} \mathsf{bool}} \qquad \frac{}{\Theta \vdash^{\mathrm{dyn}} \mathsf{unit}} \qquad \frac{\Theta \vdash^{\mathrm{dyn}} A \quad \Theta \vdash^{\mathrm{dyn}} C}{\Theta \vdash^{\mathrm{dyn}} A \to C}$$

$$\frac{[\vdash^{\mathrm{dyn}} A_i]_{1 \le i \le n}}{\Theta \vdash^{\mathrm{dyn}} \{p_1 : A_1, \ldots, p_n : A_n\}} \qquad \frac{\Theta, \vec{\alpha} \vdash^{\mathrm{dyn}} A}{\Theta \vdash^{\mathrm{dyn}} \forall \vec{\alpha}.A}$$

Well-formed polymorphic and monomorphic environments:

$$\frac{[\Theta \vdash^{\mathrm{dyn}} \forall \vec{\alpha}.A]_{(x : \forall \vec{\alpha}.A) \in \Xi}}{\Theta \vdash^{\mathrm{dyn}} \Xi} \qquad \frac{[\Theta \vdash^{\mathrm{dyn}} A]_{(x : A) \in \Gamma}}{\Theta \vdash^{\mathrm{dyn}} \Gamma}$$

Value judgements $\boxed{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} v : A}$, **assuming** $\Theta \vdash^{\mathrm{dyn}} \Xi, \Gamma, A, \Sigma$:

$$\frac{(x : A) \in \Gamma}{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} x : A} \qquad \frac{(x : \forall \vec{\alpha}.B) \in \Xi \quad [\Theta \vdash^{\mathrm{dyn}} A_i]_{1 \le i \le |\vec{\alpha}|}}{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} x : B[A_i/\alpha_i]_{1 \le i \le |\vec{\alpha}|}} \qquad \frac{}{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} \mathbf{true} : \mathsf{bool}}$$

$$\frac{}{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} \mathbf{false} : \mathsf{bool}} \qquad \frac{}{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} () : \mathsf{unit}} \qquad \frac{\Theta; \Xi; \Gamma, x : A \vdash^{\mathrm{dyn}}_{\Sigma} c : B}{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} \mathbf{fun}\, x \mapsto c : A \to B}$$

Computation judgements $\boxed{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} c : A}$, **assuming** $\Theta \vdash^{\mathrm{dyn}}_{\Sigma} \Xi, \Gamma, A, \Sigma$:

$$\frac{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} v : A}{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} \mathbf{return}\, v : A} \qquad \frac{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} c_1 : (\forall \vec{\alpha}.A) \quad \Theta; \Xi, x : \forall \vec{\alpha}.A; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} c_2 : B}{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} \mathbf{do}\, x \leftarrow c_1 \mathbf{\, in\,} c_2 : B}$$

$$\frac{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} v : \mathsf{bool} \quad \Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} c_1 : C \quad \Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} c_2 : C}{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} \mathbf{if}\, v \mathbf{\, then\,} c_1 \mathbf{\, else\,} c_2 : C}$$

$$\frac{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} v_1 : A \to B \quad \Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} v_2 : A}{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} v_1\, v_2 : B} \qquad \frac{(p : A) \in \Sigma}{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} !p : A}$$

$$\frac{(p : A) \in \Sigma \quad \Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} v : A}{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} p := v : \mathsf{unit}}$$

$$\frac{(p : A) \in \Sigma \quad \Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} v : A \quad \Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} c : B}{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} \mathbf{dlet}\, p \leftarrow v \mathbf{\, in\,} c : B}$$

Scheme judgement $\boxed{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} c : (\forall \vec{\alpha}.A)}$, **assuming** $\Theta \vdash^{\mathrm{dyn}}_{\Sigma} \Xi, \Gamma, (\forall \vec{\alpha}.A), \Sigma$:

$$\frac{\Theta, \vec{\alpha}; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} c : A}{\Theta; \Xi; \Gamma \vdash^{\mathrm{dyn}}_{\Sigma} c : (\forall \vec{\alpha}.A)} (\mathrm{GEN})$$

Fig. 9. a polymorphic type system for dynamically scoped state

---

Type-level translation with effect annotations

$$\lceil \alpha \rceil := \alpha \qquad \lceil \text{bool} \rceil := \text{bool} \qquad \boxed{\lceil A \to B \rceil := \lceil A \rceil \to \lceil B \rceil ! \lceil \Sigma \rceil} \qquad \lceil \forall \vec{\alpha}.A \rceil := \forall \vec{\alpha}.\lceil A \rceil$$

$$\lceil \Theta \rceil := \Theta \qquad \lceil \Gamma \rceil := \{x : \lceil A \rceil \mid (x : A) \in \Gamma\} \qquad \lceil \Xi \rceil := \{x : \forall \vec{\alpha}.\lceil A \rceil \mid (x : \forall \vec{\alpha}.A) \in \Gamma\}$$

$$\boxed{\lceil \Sigma \rceil := \{\text{get\_}p : \text{unit} \to \lceil A \rceil, \text{set\_}p : \lceil A \rceil \to \text{unit} \mid (p : A) \in \Sigma\}}$$

provided Σ is ground.

Type-level translation without effect annotations

$$\lfloor \alpha \rfloor := \alpha \qquad \lfloor \text{bool} \rfloor := \text{bool} \qquad \boxed{\lfloor A \to B \rfloor := \lfloor A \rfloor \to \lfloor B \rfloor} \qquad \lfloor \forall \vec{\alpha}.A \rfloor := \forall \vec{\alpha}.\lfloor A \rfloor \qquad \lfloor \Theta \rfloor := \Theta$$

$$\lfloor \Gamma \rfloor := \{x : \lfloor A \rfloor \mid (x : A) \in \Gamma\} \qquad \lfloor \Xi \rfloor := \{x : \forall \vec{\alpha}.\lfloor A \rfloor \mid (x : \forall \vec{\alpha}.A) \in \Gamma\}$$

for the ambient effect signature:

$$\boxed{\lfloor \Sigma \rfloor := \{\text{get\_}p : \text{unit} \to \lfloor A \rfloor, \text{set\_}p : \lfloor A \rfloor \to \text{unit} \mid (p : A) \in \Sigma\}}$$

Fig. 10. handlers types system expressing dynamically scoped state

Then:

$$c \overset{dyn}{\rightsquigarrow} \begin{array}{c} \mathbf{dlet}\, p \leftarrow (\mathbf{fun}\, a \mapsto (!p)\, a)\, \mathbf{in} \\ (\mathbf{fun}\, a \mapsto (!p)\, a)\, v \end{array} \overset{dyn_+}{\rightsquigarrow} \begin{array}{c} \mathbf{dlet}\, p \leftarrow (\mathbf{fun}\, a \mapsto (!p)\, a)\, \mathbf{in} \\ (!p)\, v \end{array} = c$$

as required.   ∎

Moreover, every parameter $(p : A \to B)$ lets us define a form of a fixed-point combinator $Y : ((A \to B) \to A \to B) \to (A \to B)$ by a variant of Landin's knot, provided the functions passed to this combinator and their arguments do not involve p.

The two proposed translations are correct:

Theorem (Type Preservation). *For every* $\Theta; \Xi; \Gamma \vdash^{\text{dyn}}_\Sigma c : A$ *and* $\Theta; \Xi; \Gamma \vdash^{\text{dyn}}_\Sigma v : A,$ *we have:*

- *If* Σ *is ground, then* $\lceil \Theta \rceil; \lceil \Xi \rceil; \lceil \Gamma \rceil \vdash \lceil c \rceil : \lceil A \rceil ! \lceil \Sigma \rceil$ *and* $\lceil \Theta \rceil; \lceil \Xi \rceil; \lceil \Gamma \rceil \vdash \lceil v \rceil : \lceil A \rceil.$
- $\lfloor \Theta \rfloor; \lfloor \Xi \rfloor; \lfloor \Gamma \rfloor \vdash \lceil c \rceil : \lfloor A \rfloor$ *and* $\lfloor \Theta \rfloor; \lfloor \Xi \rfloor; \lfloor \Gamma \rfloor \vdash \lceil v \rceil : \lfloor A \rfloor.$

Proof

For the first part only, first show that if A is ground, then $\lceil A \rceil = A$, and so if Σ is a well-kinded ground signature, then $\lceil \Sigma \rceil$ is well-defined and well-kinded.

Then the proofs of both parts follow the same lines. By mutual induction on the kinding judgements, show that well-kinded types, schemes, and contexts translate into well-kinded types, schemes, and contexts, respectively. Then show that both translations respect type-level substitution:

$$\lceil B[A_i/\alpha_i]_{1 \le i \le n} \rceil = \lceil B \rceil [\lceil A_i \rceil / \alpha_i]_{1 \le i \le n}$$

and similarly for the coarse translation.

Finally, by mutual induction on typing judgements for values and computations, and on scheming judgements, show the hypothesis. We mention only the interesting cases.

For dereferencing a cell $(p : A) \in \Sigma$, by the translation's definition,

$$(\texttt{get\_} p : \texttt{unit} \to \lceil A \rceil) \in \lceil \Sigma \rceil$$

Use this fact to derive that $\lceil !p \rceil$ has the type $\lceil A \rceil$. Use a similar argument for assignment.

Next, show that for all $(p : A) \in \Sigma$:

$$\lceil \Theta \rceil ; \lceil \Xi \rceil ; \lceil \Gamma \rceil \vdash H_{ST}^{p} : (B \,!\, \lceil \Sigma \rceil) \Rightarrow ((\lceil A \rceil \to (B \,!\, \lceil \Sigma \rceil)) \,!\, \lceil \Sigma \rceil)$$

and use this fact, together with the induction hypotheses, to give a valid derivation for $\lceil \textbf{dlet } p \leftarrow v \textbf{ in } c \rceil$. ∎

In summary, the handler H_{ST} expresses dynamically scoped state, in both terms and types.

5 Conclusion and further work

Unexpectedly, Hindley-Milner polymorphism integrates smoothly and robustly with existing type and effect systems for algebraic effects and handlers. However, combining reference cell allocation with polymorphism remains an open problem, as does incorporating dynamic generation of instances as used in *Eff*. Consequently, *Eff* still uses the value restriction. Our contribution is to identify a larger class of languages in which effects and polymorphism coexist naturally.

For type-system cognoscenti, these results may not come as a complete surprise. First, using effect systems to ensure soundness has been proposed (Leroy & Weis, 1991) before Wright's value restriction. Second, even if we consider the non-effect-annotated safety result, we do not believe the type system can encode the problematic effects: local reference cells and continuations. Nonetheless, previous solutions require a *specialised*, and sometimes subtle, type system. In the algebraic setting, adding polymorphism to existing systems is strikingly natural.

This result arose as part of a broader (denotational) semantic investigation of effects and polymorphism, which does not yet account for reference cells. We hope that an algebraic understanding of locality (Staton, 2013; Fiore & Staton, 2014) and scope and polymorphic arities (Wu *et al.*, 2014) will explain the interaction between reference cells and polymorphism. The robustness of type safety leads us to believe standard extensions, such as type inference, principal types, and impredicative and row polymorphism will not pose problems. The latter is particularly interesting, as it can serve as an effect system with effect variables (Lindley & Cheney, 2012; Leijen, 2014; Pretnar, 2014).

We want to investigate the expressive difference between effect handlers and delimited control, and polymorphism forms another comparison axis. We defer a thorough comparison, as there are several notions of delimited control (shift, shift0, with or without answer-type modification) and several proposals for polymorphic type systems (Asai & Kameyama, 2007; Gunter *et al.*, 1995; Kiselyov *et al.*, 2006), and as delimited control is subtle. That said, there are two immediate points of comparison between delimited control and effect handlers.

First, Kiselyov *et al.*'s translation of dynamic scope into delimited control requires some complication in order to preserve the type. This complication is caused by their effect system for delimited control tracking, the return type of the computation enclosed by the

nearest rebinding. When an access to a dynamically scoped cell escapes the current binding in scope the type expected in the nearest rebinding may change, resulting in a type error of their translated program. The example on page 14 demonstrates such a shift from a function type to an integer type. In contrast, our effect system only tracks the local type for each effect operation, and the translation from dynamically scoped state to effect handlers extends smoothly to types.

Second, these systems include a form of a purity restriction or value restriction. As a consequence, they cannot type purely functional computations like the final example in the Evaluation Subsection 4.1. In contrast, the type system proposed here allows unrestricted Hindley-Milner polymorphism in both purely functional and effectful code.

Acknowledgments *Omitted for peer review.*

Bibliography

Asai, Kenichi, & Kameyama, Yukiyoshi. (2007). *Programming languages and systems: 5th asian symposium, aplas 2007, singapore, november 29-december 1, 2007. proceedings*. Berlin, Heidelberg: Springer Berlin Heidelberg. Chap. Polymorphic Delimited Continuations, pages 239–254.

Bauer, Andrej, & Pretnar, Matija. (2014). An effect system for algebraic effects and handlers. *Logical methods in computer science*, **10**(4).

Bauer, Andrej, & Pretnar, Matija. (2015). Programming with algebraic effects and handlers. *J. log. algebr. meth. program.*, **84**(1), 108–123.

Danvy, Olivier. (2006). *An analytical approach to programs as data objects*. Dsc dissertation, Department of Computer Science, University of Aarhus.

Felleisen, Matthias. (1990). *Esop '90: 3rd european symposium on programming copenhagen, denmark, may 15–18, 1990 proceedings*. Berlin, Heidelberg: Springer Berlin Heidelberg. Chap. On the expressive power of programming languages, pages 134–151.

Fiore, Marcelo, & Staton, Sam. (2014). Substitution, jumps, and algebraic effects. *Pages 41:1–41:10 of: Proceedings of the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. CSL-LICS '14. New York, NY, USA: ACM.

Garrigue, Jacques. (2004). Relaxing the value restriction. *Pages 196–213 of:* Kameyama, Yukiyoshi, & Stuckey, Peter J. (eds), *Functional and Logic Programming*. Lecture Notes in Computer Science, vol. 2998. Springer Berlin Heidelberg.

Gunter, Carl A., Rémy, Didier, & Riecke, Jon G. 1995 (June). A generalization of exceptions and control in ML. *Proc. ACM Conf. on Functional Programming and Computer Architecture*.

Harper, Robert, & Lillibridge, Mark. (1993). Polymorphic type assignment and CPS conversion. *Lisp and symbolic computation*, **6**(3-4), 361–380.

Kammar, Ohad, Lindley, Sam, & Oury, Nicolas. (2013). Handlers in action. *Sigplan not.*, **48**(9), 145–158.

Kiselyov, Oleg, Shan, Chung-chieh, & Sabry, Amr. (2006). Delimited dynamic binding. *Sigplan not.*, **41**(9), 26–37.

Kiselyov, Oleg, Sabry, Amr, & Swords, Cameron. (2013).   Extensible effects: An alternative to monad transformers. *Sigplan not.*, **48**(12), 59–70.

Landin, P. J. (1964). The mechanical evaluation of expressions. *The computer journal*, **6**(4), 308–320.

Leijen, Daan. (2014). Koka: Programming with row polymorphic effect types. *Pages 100–126 of:* Levy, Paul, & Krishnaswami, Neel (eds), *Proceedings 5th Workshop on Mathematically Structured Functional Programming, MSFP 2014, Grenoble, France, 12 April 2014.* EPTCS, vol. 153.

Leroy, Xavier. (1992). *Typage polymorphe d'un langage algorithmique.* Phd thesis (in french), Université Paris 7.

Leroy, Xavier. (1993). Polymorphism by name for references and continuations. *Pages 220–231 of: 20th symposium Principles of Programming Languages.* ACM Press.

Leroy, Xavier, & Weis, Pierre. (1991). Polymorphic type inference and assignment. *Pages 291–302 of: Proceedings of the 18th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages.* POPL '91. New York, NY, USA: ACM.

Levy, Paul Blain, Power, John, & Thielecke, Hayo. (2003). Modelling environments in call-by-value programming languages. *Inf. comput.*, **185**(2), 182–210.

Lindley, Sam, & Cheney, James. (2012). Row-based effect types for database integration. *Pages 91–102 of: Proceedings of the 8th ACM SIGPLAN Workshop on Types in Language Design and Implementation.* TLDI '12. New York, NY, USA: ACM.

Lucassen, John M., & Gifford, David K. (1988). Polymorphic effect systems. *Pages 47–57 of: Proceedings of the 15th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages.* POPL '88. New York, NY, USA: ACM.

Milner, Robin. (1978). A theory of type polymorphism in programming. *Journal of computer and system sciences*, **17**(3), 348 – 375.

Moreau, Luc. (1998). A syntactic theory of dynamic binding. *Higher-order and symbolic computation*, **11**(3), 233–279.

Pfenning, Frank, & Schürmann, Carsten. (1999). System description: Twelf — a meta-logical framework for deductive systems. *Pages 202–206 of: Automated Deduction — CADE-16.* Lecture Notes in Computer Science, vol. 1632. Springer Berlin Heidelberg.

Pierce, Benjamin C. (2002). *Types and programming languages.* Cambridge, MA, USA: MIT Press.

Plotkin, Gordon D., & Power, John. (2003). Algebraic operations and generic effects. *Applied categorical structures*, **11**(1), 69–94.

Plotkin, Gordon D., & Pretnar, Matija. (2013). Handling algebraic effects. *Logical methods in computer science*, **9**(4).

Pretnar, Matija. (2014). Inferring algebraic effects. *Logical methods in computer science*, **10**(3).

Pretnar, Matija. (2015). An introduction to algebraic effects and handlers. invited tutorial paper. *Electr. notes theor. comput. sci.*, **319**, 19–35.

Rémy, Didier. (2015). *Type systems.* Lecture notes. Parisian Master of Research in Computer Science.

Staton, Sam. 2013 (June). Instances of computational effects: An algebraic perspective. *Pages 519–519 of: Logic in Computer Science (LICS), 2013 28th Annual IEEE/ACM Symposium on.*

24   *O. Kammar and M. Pretnar*

Wells, J.B. (1999).   Typability and type checking in system f are equivalent and undecidable. *Annals of pure and applied logic*, **98**(13), 111 – 156.

Wright, Andrew K. (1995).   Simple imperative polymorphism.   *Lisp and symbolic computation*, **8**(4), 343–355.

Wu, Nicolas, Schrijvers, Tom, & Hinze, Ralf. (2014). Effect handlers in scope. *Pages 1–12 of: Proceedings of the 2014 ACM SIGPLAN Symposium on Haskell*. Haskell '14. New York, NY, USA: ACM.