Handling Polymorphic Algebraic Effects

Taro Sekiyama¹(✉) and Atsushi Igarashi²(✉)

¹ National Institute of Informatics, Tokyo, Japan
tsekiyama@acm.org

² Kyoto University, Kyoto, Japan
igarashi@kuis.kyoto-u.ac.jp

Abstract. Algebraic effects and handlers are a powerful abstraction mechanism to represent and implement control effects. In this work, we study their extension with parametric polymorphism that allows abstracting not only expressions but also effects and handlers. Although polymorphism makes it possible to reuse and reason about effect implementations more effectively, it has long been known that a naive combination of polymorphic effects and let-polymorphism breaks type safety. Although type safety can often be gained by restricting let-bound expressions—e.g., by adopting value restriction or weak polymorphism—we propose a complementary approach that restricts handlers instead of let-bound expressions. Our key observation is that, informally speaking, a handler is safe if resumptions from the handler do not interfere with each other. To formalize our idea, we define a call-by-value lambda calculus $\lambda_{\text{eff}}^{\text{let}}$ that supports let-polymorphism and polymorphic algebraic effects and handlers, design a type system that rejects interfering handlers, and prove type safety of our calculus.

1 Introduction

Algebraic effects [20] and handlers [21] are a powerful abstraction mechanism to represent and implement control effects, such as exceptions, interactive I/O, mutable states, and nondeterminism. They are growing in popularity, thanks to their success in achieving modularity of effects, especially the clear separation between their interfaces and their implementations. An interface of effects is given as a set of *operations*—e.g., an interface of mutable states consists of two operations, namely, **put** and **get**—with their signatures. An implementation is given by a *handler* H , which provides a set of interpretations of the operations (called *operation clauses*), and a **handle-with** expression **handle** M with H associates effects invoked during the computation of M with handler H . Algebraic effects and handlers work as *resumable exceptions*: when an effect operation is invoked, the run-time system tries to find the nearest handler that handles the invoked operation; if it is found, the corresponding operation clause is evaluated by using the argument to the operation invocation and the continuation up to the handler. The continuation gives the ability to resume the computation from the point where the operation was invoked, using the result from the operation

clause. Another modularity that algebraic effects provide is flexible composition: multiple algebraic effects can be combined freely [13].

In this work, we study an extension of algebraic effects and handlers with another type-based abstraction mechanism—parametric polymorphism [22]. In general, parametric polymorphism is a basis of generic programming and enhance code reusability by abstracting expressions over types. This work allows abstracting not only expressions but also effect operations and handlers, which makes it possible to reuse and reason about effect implementations that are independent of concrete type representations. Like in many functional languages, we introduce polymorphism in the form of *let-polymorphism* for its practically desirable properties such as decidable typechecking and type inference.

As is well known, however, a naive combination of polymorphic effects and let-polymorphism breaks type safety [11, 23]. Many researchers have attacked this classical problem [1, 2, 10, 12, 14, 17, 23, 24], and their common idea is to restrict the form of let-bound expressions. For example, value restriction [23, 24], which is the standard way to make ML-like languages with imperative features and let-polymorphism type safe, allows only syntactic values to be polymorphic.

In this work, we propose a new approach to achieving type safety in a language with let-polymorphic and polymorphic effects and handlers: the idea is to restrict handlers instead of let-bound expressions. Since a handler gives an implementation of an effect, our work can be viewed as giving a criterion that suggests what effects can cooperate safely with (unrestricted) let-polymorphism and what effects cannot. Our key observation for type safety is that, informally speaking, an invocation of a polymorphic effect in a let-bound expression is safe if resumptions in the corresponding operation clause do not interfere with each other. We formalize this discipline into a type system and show that typeable programs do not get stuck.

Our contributions are summarized as follows.

- We introduce a call-by-value, statically typed lambda calculus $\lambda_{\text{eff}}^{\text{let}}$ that supports let-polymorphism and polymorphic algebraic effects and handlers. The type system of $\lambda_{\text{eff}}^{\text{let}}$ allows any let-bound expressions involving effects to be polymorphic, but, instead, disallows handlers where resumptions interfere with each other.
- To give the semantics of $\lambda_{\text{eff}}^{\text{let}}$, we formalize an intermediate language λ_{eff}^A wherein type information is made explicit and define a formal elaboration from $\lambda_{\text{eff}}^{\text{let}}$ to λ_{eff}^A .
- We prove type safety of $\lambda_{\text{eff}}^{\text{let}}$ by type preservation of the elaboration and type soundness of λ_{eff}^A .

We believe that our approach is complementary to the usual approach of restricting let-bound expressions: for handlers that are considered unsafe by our criterion, the value restriction can still be used.

The rest of this paper is organized as follows. Section 2 provides an overview of our work, giving motivating examples of polymorphic effects and handlers, a problem in naive combination of polymorphic effects and let-polymorphism,

and our solution to gain type safety with those features. Section 3 defines the surface language $\lambda_{\text{eff}}^{\text{let}}$, and Sect. 4 defines the intermediate language λ_{eff}^A and the elaboration from $\lambda_{\text{eff}}^{\text{let}}$ to λ_{eff}^A . We also state that the elaboration is type-preserving and that λ_{eff}^A is type sound in Sect. 4. Finally, we discuss related work in Sect. 5 and conclude in Sect. 6. The proofs of the stated properties and the full definition of the elaboration are given in the full version at <https://arxiv.org/abs/1811.07332>.

2 Overview

We start with reviewing how monomorphic algebraic effects and handlers work through examples and then extend them to a polymorphic version. We also explain why polymorphic effects are inconsistent with let-polymorphism, if naively combined, and how we resolve it.

2.1 Monomorphic Algebraic Effects and Handlers

Exception. Our first example is exception handling, shown in an ML-like language below.

```

1  effect fail : unit  $\hookrightarrow$  unit
2
3  let div100 (x:int) : int =
4    if x = 0 then (#fail()); -1)
5    else 100 / x
6
7  let f (y:int) : int option =
8    handle (div_100 y) with
9      return z  $\rightarrow$  Some z
10     fail    z  $\rightarrow$  None

```

Some and None are constructors of datatype α option. Line 1 declares an effect operation **fail**, which signals that an anomaly happens, with its signature $\text{unit} \hookrightarrow \text{unit}$, which means that the operation is invoked with the unit value $()$, causes some effect, and may return the unit value. The function **div100**, defined in Lines 3–5, is an example that uses **fail**; it returns the number obtained by dividing 100 by argument **x** if **x** is not zero; otherwise, if **x** is zero, it raises an exception by calling effect operation **fail**.¹ In general, we write $\#op(M)$ for invoking effect operation **op** with argument *M*. The function **f** (Lines 7–10) calls **div_100** inside a **handle-with** expression, which returns **Some** *n* if **div_100** returns integer *n* normally and returns **None** if it invokes **fail**.

An expression of the form **handle** *M* **with** *H* handles effect operations invoked in *M* (which we call *handled expression*) according to the effect interpretations given by handler *H*. A handler *H* consists of two parts: a single *return*

¹ Here, “; -1” is necessary to make the types of both branches the same; it becomes unnecessary when we introduce polymorphic effects.

clause and zero or more *operation clauses*. A return clause **return** $x \rightarrow M'$ will be executed if the evaluation of M results in a value v . Then, the value of M' (where x is bound to v) will be the value of the entire **handle-with** expression. For example, in the program above, if a nonzero number n is passed to **f**, the **handle-with** expression would return **Some** $(100/n)$ because **div100** n returns $100/n$. An operation clause **op** $x \rightarrow M'$ defines an implementation of effect **op**: if the evaluation of handled expression M invokes effect **op** with argument v , expression M' will be evaluated after substituting v for x and the value of M' will be the value of the entire **handle-with** expression. In the program example above, if zero is given to **f**, then **None** will be returned because **div100** 0 invokes **fail**.

As shown above, algebraic effect handling is similar to exception handling. However, a distinctive feature of algebraic effect handling is that it allows *resumption* of the computation from the point where an effect operation was invoked. The next example demonstrates such an ability of algebraic effect handlers.

Choice. The next example is effect **choose**, which returns one of the given two arguments.

```

1  effect choose : int × int ↦ int
2
3  handle (#choose(1,2) + #choose(10,20)) with
4      return x → x
5      choose x → resume (fst x)

```

As usual, $A_1 \times A_2$ is a product type, (M_1, M_2) is a pair expression, and **fst** is the first projection function. The first line declares that effect **choose** is for choosing integers. The handled expression **#choose**(1,2) + **#choose**(10,20) intuitively suggests that there would be four possible results—11, 21, 12, and 22—depending on which value each invocation of **choose** returns. The handler in this example always chooses the first element of a given pair² and returns it by using a **resume** expression, and, as a result, the expression in Lines 3–5 evaluates to 11.

A resumption expression **resume** M in an operation clause makes it possible to return a value of M to the point where an effect operation was invoked. This behavior is realized by constructing a *delimited continuation* from the point of the effect invocation up to the **handle-with** expression that deals with the effect and passing the value of M to the continuation. We illustrate it by using the program above. When the handled expression **#choose**(1,2) + **#choose**(10,20) is evaluated, continuation $c \stackrel{\text{def}}{=} [] + \text{\#choose}(10,20)$ is constructed. Then, the body **resume** (fst x) of the operation clause is evaluated after binding x to the invocation argument (1,2). Receiving the value 1 of **fst** (1,2), the resumption

² We can think of more practical implementations, which choose one of the two arguments by other means, say, random values.

expression passes it to the continuation c and $c[1] = 1 + \text{\#choose}(10, 20)$ is evaluated under the same handler. Next, **choose** is invoked with argument $(10, 20)$. Similarly, continuation $c' \stackrel{\text{def}}{=} 1 + []$ is constructed and the operation clause for **choose** is executed again. Since **fst** $(10, 20)$ evaluates to 10, $c'[10] = 1 + 10$ is evaluated under the same handler. Since the **return** clause returns what it receives, the entire expression evaluates to 11.

Finally, we briefly review how an operation clause involving resumption expressions is typechecked [3, 13, 16]. Let us consider operation clause $\text{op}(\mathbf{x}) \rightarrow M$ for **op** of type signature $A \hookrightarrow B$. The typechecking is performed as follows. First, argument $\mathbf{x}$ is assigned the domain type A of the signature as it will be bound to an argument of an effect invocation. Second, for resumption expression **resume** M' in M , (1) M' is required to have the codomain type B of the signature because its value will be passed to the continuation as the result of the invocation and (2) the resumption expression is assigned the same type as the return clause. Third, the type of the body M has to be the same as that of the return clause because the value of M is the result of the entire **handle-with** expression. For example, the above operation clause for **choose** is typechecked as follows: first, argument $\mathbf{x}$ is assigned type $\text{int} \times \text{int}$; second, it is checked whether the argument **fst** $\mathbf{x}$ of the resumption expression has **int**, the codomain type of **choose**; third, it is checked whether the body **resume** (**fst** $\mathbf{x}$) of the clause has the same type as the return clause, i.e., **int**. If all the requirements are satisfied, the clause is well typed.

2.2 Polymorphic Algebraic Effects and Handlers

This section discusses motivation for polymorphism in algebraic effects and handlers. There are two ways to introduce polymorphism: by *parameterized effects* and by *polymorphic effects*.

The former is used to parameterize the declaration of an effect by types. For example, one might declare:

```
effect  $\alpha$  choose :  $\alpha \times \alpha \hookrightarrow \alpha$ 
```

An invocation **#choose** involves a parameterized effect of the form A **choose** (where A denotes a type), according to the type of arguments: For example, **#choose**(**true**, **false**) has the effect **bool choose** and **#choose**(1, -1) has **int choose**. Handlers are required for each effect A **choose**.

The latter is used to give a polymorphic type to an effect. For example, one may declare

```
effect choose :  $\forall \alpha. \alpha \times \alpha \hookrightarrow \alpha$ 
```

In this case, the effect can be invoked with different types, but all invocations have the same effect **choose**. One can implement a single operation clause that can handle all invocations of **choose**, regardless of argument types. Koka supports both styles [16] (with the value restriction); we focus, however, on the latter in this paper. A type system for parameterized effects lifting the value restriction is studied by Kammar and Pretnar [14] (see Sect. 5 for comparison).

In what follows, we show a polymorphic version of the examples we have seen, along with brief discussions on how polymorphic effects help with reasoning about effect implementations. Other practical examples of polymorphic effects can be found in Leijen’s work [16].

Polymorphic Exception. First, we extend the exception effect `fail` with polymorphism.

```

1  effect fail∀ : ∀α. unit ↦ α
2
3  let div100∀ (x:int) : int =
4    if x = 0 then #fail∀()
5    else 100 / x

```

The polymorphic type signature of effect `fail∀`, given in Line 1, means that the codomain type α can be any. Thus, we do not need to append the dummy value `-1` to the invocation of `fail∀` by instantiating the bound type variable α with `int` (the shaded part).

Choice. Next, let us make `choose` polymorphic.

```

1  effect choose∀ : ∀α. α × α ↦ α
2
3  let rec random_walk (x:int) : int =
4    let b = #choose∀(true,false) in
5    if b then random_walk (x + #choose∀(1,-1))
6    else x
7
8  let f (s:int) =
9    handle random_walk s with
10     return x → x
11     choose∀ y → if rand() < 0.0 then resume (fst y)
12                  else resume (snd y)

```

The function `random_walk` implements random walk; it takes the current coordinate `x`, chooses whether it stops, and, if it decides to continue, recursively calls itself with a new coordinate. In the definition, `choose∀` is used twice with different types: `bool` and `int`. Lines 11–12 give `choose∀` an interpretation, which calls `rand` to obtain a random `float`,³ and returns either the first or the second element of `y`.

Typechecking of operation clauses could be extended in a straightforward manner. That is, an operation clause `op(x) → M` for an effect operation of signature $\forall\alpha. A \hookrightarrow B$ would be typechecked as follows: first, α is locally bound in the clause and x is assigned type A ; second, an argument of a resumption

³ One might implement `rand` as another effect operation.

expression must have type B (which may contain type variable α); third, M must have the same type as that of the return clause (its type cannot contain α as α is local) under the assumption that resumption expressions have the same type as the return clause. For example, let us consider typechecking of the above operation clause for **choose**[∇]. First, the typechecking algorithm allocates a local type variable α and assigns type $\alpha \times \alpha$ to **y**. The body has two resumption expressions, and it is checked whether the arguments **fst y** and **snd y** have the codomain type α of the signature. Finally, it is checked whether the body is typed at **int** assuming that the resumption expressions have type **int**. The operation clause meets all the requirements, and, therefore, it would be well typed.

An obvious advantage of polymorphic effects is reusability. Without polymorphism, one has to declare many versions of **choose** for different types.

Another pleasant effect of polymorphic effects is that, thanks to parametricity, inappropriate implementations for an effect operation can be excluded. For example, it is not possible for an implementation of **choose**[∇] to resume with values other than the first or second element of **y**. In the monomorphic version, however, it is possible to resume with any integer, as opposed to what the name of the operation suggests. A similar argument applies to **fail**[∇]; since the codomain type is α , which does not appear in the domain type, it is not possible to resume! In other words, the signature $\forall\alpha. \text{unit} \hookrightarrow \alpha$ enforces that no invocation of **fail**[∇] will return.

2.3 Problem in Naive Combination with Let-Polymorphism

Although polymorphic effects and handlers provide an ability to abstract and restrict effect implementations, one may easily expect that their unrestricted use with naive *let-polymorphism*, which allows any let-bound expressions to be polymorphic, breaks type safety. Indeed, it does.

We develop a counterexample, inspired by Harper and Lillibridge [11], below.

```
effect get_id : ∀α. unit ↪ (α → α)

let f () : int =
  let g = #get_id() in (* g : ∀α. α → α *)
  if (g true) then ((g 0) + 1) else 2
```

The function **f** first binds **g** to the invocation result of **op**. The expression **#get_id()** is given type $\alpha \rightarrow \alpha$ and the naive let-polymorphism would assign type scheme $\forall\alpha. \alpha \rightarrow \alpha$ to **g**, which makes both **g true** and **g 0** (and thus the definition of **f**) well typed.

An intended use of **f** is as follows:

```
handle f () with
  return x → x
  get_id y → resume (λz. z)
```

The operation clause for `get_id` resumes with the identity function $\lambda z.z$. It would be well typed under the typechecking procedure described in Sect. 2.2 and it safely returns 1.

However, the following strange expression

```
handle f () with
  return x → x
  get_id y → resume (λz1. (resume (λz2. z1))); z1)
```

will get stuck, although this expression would be well typed: both $\lambda z1. \dots ; z1$ and $\lambda z2. z1$ could be given type $\alpha \rightarrow \alpha$ by assigning both $z1$ and $z2$ type α , which is the type variable local to this clause. Let us see how the evaluation gets stuck in detail. When the handled expression `f ()` invokes effect `get_id`, the following continuation will be constructed:

$$c \stackrel{\text{def}}{=} \text{let } g = [] \text{ in if } (g \text{ true}) \text{ then } ((g \ 0) + 1) \text{ else } 2.$$

Next, the body of the operation clause `get_id` is evaluated. It immediately resumes and reduces to

$$c'[(\lambda z1. c'[(\lambda z2.z1)]; z1)]$$

where

```
handle c with
  c'  $\stackrel{\text{def}}{=}$  return x → x
  get_id y → resume (λz1. (resume (λz2.z1))); z1),
```

which is the continuation c under the same handler. The evaluation proceeds as follows (here, $k \stackrel{\text{def}}{=} \lambda z1. c'[(\lambda z2.z1)]; z1$):

```
c'[(λz1. c'[(λz2.z1)]; z1)]
= handle let g = k in if (g true) then ((g 0) + 1) else 2 with ...
→ handle if (k true) then ((k 0) + 1) else 2 with ...
→ handle if c'[(λz2.true)]; true then ((k 0) + 1) else 2 with ...
```

Here, the hole in c' is filled by function $(\lambda z2.\text{true})$, which returns a Boolean value, *though the hole is supposed to be filled by a function of $\forall \alpha. \alpha \rightarrow \alpha$* . This weird gap triggers a run-time error:

```
c'[(λz2.true)]
handle
= let g = λz2.true in if (g true) then ((g 0) + 1) else 2
  with ...
→* handle if true then (((λz2.true) 0) + 1) else 2 with ...
→ handle ((λz2.true) 0) + 1 with ...
→ handle true + 1 with ...
```

We stop here because `true + 1` cannot reduce.

2.4 Our Solution

A standard approach to this problem is to restrict the form of let-bound expressions by some means such as the (relaxed) value restriction [10, 23, 24] or weak polymorphism [1, 12]. This approach amounts to restricting how effect operations can be *used*.

In this paper, we seek for a complementary approach, which is to restrict how effect operations can be *implemented*.⁴ More concretely, we develop a type system such that let-bound expressions are polymorphic as long as they invoke only “safe” polymorphic effects and the notion of safe polymorphic effects is formalized in terms of typing rules (for handlers).

To see what are “safe” effects, let us examine the above counterexample to type safety. The crux of the counterexample is that

1. continuation c uses g polymorphically, namely, as $\text{bool} \rightarrow \text{bool}$ in $g \text{ true}$ and as $\text{int} \rightarrow \text{int}$ in $g \text{ 1}$;
2. c is invoked twice; and
3. the use of g as $\text{bool} \rightarrow \text{bool}$ in the first invocation of c —where g is bound to $\lambda z1. \dots; z1$ —“alters” the type of $\lambda z2. z1$ (passed to **resume**) from $\alpha \rightarrow \alpha$ to $\alpha \rightarrow \text{bool}$, contradicting the second use of g as $\text{int} \rightarrow \text{int}$ in the second invocation of c .

The last point is crucial—if $\lambda z2. z1$ were, say, $\lambda z2. z2$, there would be no influence from the first invocation of c and the evaluation would succeed. The problem we see here is that the naive type system mistakenly allows *interference* between the arguments to the two resumptions by assuming that $z1$ and $z2$ share the same type.

Based on this observation, the typing rule for resumption is revised to disallow interference between different resumptions by separating their types: for each **resume** M in the operation clause for $\text{op} : \forall \alpha_1 \dots \alpha_n. A \hookrightarrow B$, M has to have type B' obtained by renaming all type variables α_i in B with *fresh* type variables α'_i . In the case of **get_id**, the two resumptions should be called with $\beta \rightarrow \beta$ and $\gamma \rightarrow \gamma$ for fresh β and γ ; for the first **resume** to be well typed, $z1$ has to be of type β , although it means that the return type of $\lambda z2. z1$ (given to the second resumption) is β , making the entire clause ill typed, as we expect. If a clause does not have interfering resumptions like

$$\text{get\_id } y \rightarrow \text{resume } (\lambda z1. z1)$$

or

$$\text{get\_id } y \rightarrow \text{resume } (\lambda z1. (\text{resume } (\lambda z2. z2)); z1),$$

it will be well typed.

---

⁴ We compare our approach with the standard approaches in Sect. 5 in detail.

3 Surface Language: $\lambda_{\text{eff}}^{\text{let}}$

We define a lambda calculus $\lambda_{\text{eff}}^{\text{let}}$ that supports let-polymorphism, polymorphic algebraic effects, and handlers without interfering resumptions. This section introduces the syntax and the type system of $\lambda_{\text{eff}}^{\text{let}}$. The semantics is given by a formal elaboration to intermediate calculus λ_{eff}^A , which will be introduced in Sect. 4.

Effect operations	op	Type variables	α, β, γ
Effects	ϵ	::=	sets of effect operations
Base types	ι	::=	bool int ...
Types	A, B, C, D	::=	$\alpha \mid \iota \mid A \rightarrow_{\epsilon} B$
Type schemes	σ	::=	$A \mid \forall \alpha. \sigma$
Constants	c	::=	true false 0 + ...
Terms	M	::=	$x \mid c \mid \lambda x. M \mid M_1 M_2 \mid \text{let } x = M_1 \text{ in } M_2 \mid$ $\# \text{op}(M) \mid \text{handle } M \text{ with } H \mid \text{resume } M$
Handlers	H	::=	return $x \rightarrow M \mid H; \text{op}(x) \rightarrow M$
Typing contexts	Γ	::=	$\emptyset \mid \Gamma, x : \sigma \mid \Gamma, \alpha$

Fig. 1. Syntax of $\lambda_{\text{eff}}^{\text{let}}$.

3.1 Syntax

The syntax of $\lambda_{\text{eff}}^{\text{let}}$ is given in Fig. 1. Effect operations are denoted by **op** and type variables by α, β , and γ . An effect, denoted by ϵ , is a finite set of effect operations. We write $\langle \rangle$ for the empty effect set. A type, denoted by A, B, C , and D , is a type variable; a base type ι , which includes, e.g., **bool** and **int**; or a function type $A \rightarrow_{\epsilon} B$, which is given to functions that take an argument of type A and compute a value of type B possibly with effect ϵ . A type scheme σ is obtained by abstracting type variables. Terms, denoted by M , consist of variables; constants (including primitive operations); lambda abstractions $\lambda x. M$, which bind x in M ; function applications; let-expressions $\text{let } x = M_1 \text{ in } M_2$, which bind x in M_2 ; effect invocations $\# \text{op}(M)$; handle-with expressions $\text{handle } M \text{ with } H$; and resumption expressions $\text{resume } M$. All type information in $\lambda_{\text{eff}}^{\text{let}}$ is implicit; thus the terms have no type annotations. A handler H has a single return clause $\text{return } x \rightarrow M$, where x is bound in M , and zero or more operation clauses of the form $\text{op}(x) \rightarrow M$, where x is bound in M . A typing context Γ binds a sequence of variable declarations $x : \sigma$ and type variable declarations α .

We introduce the following notations used throughout this paper. We write $\forall \alpha^{i \in I}. A$ for $\forall \alpha_1 \dots \forall \alpha_n. A$ where $I = \{1, \dots, n\}$. We often omit indices (i and j) and index sets (I and J) if they are not important: e.g., we often abbreviate $\forall \alpha^{i \in I}. A$ to $\forall \alpha^I. A$ or even to $\forall \alpha. A$. Similarly, we use a bold font for other sequences ($\mathbf{A}^{i \in I}$ for a sequence of types, $\mathbf{v}^{i \in I}$ for a sequence of values, etc.).

We sometimes write $\{\alpha\}$ to view the sequence α as a set by ignoring the order. Free type variables $ftv(\sigma)$ in a type scheme σ and type substitution $B[A/\alpha]$ of A for type variables α in B are defined as usual (with the understanding that the omitted index sets for A and α are the same).

We suppose that each constant c is assigned a first-order closed type $ty(c)$ of the form $\iota_1 \rightarrow \langle \rangle \cdots \rightarrow \langle \rangle \iota_n$ and that each effect operation op is assigned a signature of the form $\forall \alpha. A \hookrightarrow B$, which means that an invocation of op with type instantiation C takes an argument of $A[C/\alpha]$ and returns a value of $B[C/\alpha]$. We also assume that, for $ty(\text{op}) = \forall \alpha. A \hookrightarrow B$, $ftv(A) \subseteq \{\alpha\}$ and $ftv(B) \subseteq \{\alpha\}$.

3.2 Type System

The type system of $\lambda_{\text{eff}}^{\text{let}}$ consists of four judgments: well-formedness of typing contexts $\vdash \Gamma$; well formedness of type schemes $\Gamma \vdash \sigma$; term typing judgment $\Gamma; R \vdash M : A \mid \epsilon$, which means that M computes a value of A possibly with effect ϵ under typing context Γ and resumption type R (discussed below); and handler typing judgment $\Gamma; R \vdash H : A \mid \epsilon \Rightarrow B \mid \epsilon'$, which means that H handles a computation that produces a value of A with effect ϵ and that the clauses in H compute a value of B possibly with effect ϵ' under Γ and R .

A resumption type R contains type information for resumption.

Definition 1 (Resumption type). *Resumption types in $\lambda_{\text{eff}}^{\text{let}}$, denoted by R , are defined as follows:*

$$R ::= \text{none} \mid (\alpha, A, B \rightarrow \epsilon C) \\ \text{(if } ftv(A) \cup ftv(B) \subseteq \{\alpha\} \text{ and } ftv(C) \cap \{\alpha\} = \emptyset \text{)}$$

If M is not a subterm of an operation clause, it is typechecked under $R = \text{none}$, which means that M cannot contain resumption expressions. Otherwise, suppose that M is a subterm of an operation clause $\text{op}(x) \rightarrow M'$ that handles effect op of signature $\forall \alpha. A \hookrightarrow B$ and computes a value of C possibly with effect ϵ . Then, M is typechecked under $R = (\alpha, x : A, B \rightarrow \epsilon C)$, which means that argument x to the operation clause has type A and that resumptions in M are effectful functions from B to C with effect ϵ . Note that type variables α occur free only in A and B but not in C .

Figure 2 shows the inference rules of the judgments (except for $\Gamma \vdash \sigma$, which is defined by: $\Gamma \vdash \sigma$ if and only if all free type variables in σ are bound by Γ). For a sequence of type schemes σ , we write $\Gamma \vdash \sigma$ if and only if every type scheme in σ is well formed under Γ .

Well-formedness rules for typing contexts, shown at the top of Fig. 2, are standard. A typing context is well formed if it is empty (WF_EMPTY) or a variable in the typing context is associated with a type scheme that is well formed in the remaining typing context (WF_VAR) and a type variable in the typing context is not declared (WF_TVAR). For typing context Γ , $\text{dom}(\Gamma)$ denotes the set of type and term variables declared in Γ .

Well-formed rules for typing contexts

$$\boxed{\vdash \Gamma}$$

$$\frac{}{\vdash \emptyset} \text{WF\_EMPTY} \qquad \frac{\vdash \Gamma \quad x \notin \text{dom}(\Gamma) \quad \Gamma \vdash \sigma}{\vdash \Gamma, x : \sigma} \text{WF\_VAR}$$

$$\frac{\vdash \Gamma \quad \alpha \notin \text{dom}(\Gamma)}{\vdash \Gamma, \alpha} \text{WF\_TYVAR}$$

Typing rules

$$\boxed{\Gamma; R \vdash M : A \mid \epsilon}$$

$$\frac{\vdash \Gamma \quad x : \forall \alpha. A \in \Gamma \quad \Gamma \vdash B}{\Gamma; R \vdash x : A[B/\alpha] \mid \epsilon} \text{TS\_VAR} \qquad \frac{\vdash \Gamma}{\Gamma; R \vdash c : \text{ty}(c) \mid \epsilon} \text{TS\_CONST}$$

$$\frac{\Gamma, x : A; R \vdash M : B \mid \epsilon'}{\Gamma; R \vdash \lambda x. M : A \rightarrow \epsilon' B \mid \epsilon} \text{TS\_ABS}$$

$$\frac{\Gamma; R \vdash M_1 : A \rightarrow \epsilon' B \mid \epsilon \quad \Gamma; R \vdash M_2 : A \mid \epsilon \quad \epsilon' \subseteq \epsilon}{\Gamma; R \vdash M_1 M_2 : B \mid \epsilon} \text{TS\_APP}$$

$$\frac{\Gamma, \alpha; R \vdash M_1 : A \mid \epsilon \quad \Gamma, x : \forall \alpha. A; R \vdash M_2 : B \mid \epsilon}{\Gamma; R \vdash \text{let } x = M_1 \text{ in } M_2 : B \mid \epsilon} \text{TS\_LET}$$

$$\frac{\Gamma; R \vdash M : A \mid \epsilon' \quad \epsilon' \subseteq \epsilon}{\Gamma; R \vdash M : A \mid \epsilon} \text{TS\_WEAK}$$

$$\frac{\text{ty}(\text{op}) = \forall \alpha. A \hookrightarrow B \quad \text{op} \in \epsilon \quad \Gamma; R \vdash M : A[C/\alpha] \mid \epsilon \quad \Gamma \vdash C}{\Gamma; R \vdash \# \text{op}(M) : B[C/\alpha] \mid \epsilon} \text{TS\_OP}$$

$$\frac{\Gamma; R \vdash M : A \mid \epsilon \quad \Gamma; R \vdash H : A \mid \epsilon \Rightarrow B \mid \epsilon'}{\Gamma; R \vdash \text{handle } M \text{ with } H : B \mid \epsilon'} \text{TS\_HANDLE}$$

$$\frac{\vdash \Gamma_1, x : D, \Gamma_2 \quad \alpha \in \Gamma_1 \quad \epsilon \subseteq \epsilon' \quad \Gamma_1, \Gamma_2, \beta, x : A[\beta/\alpha]; (\alpha, x : A, B \rightarrow \epsilon C) \vdash M : B[\beta/\alpha] \mid \epsilon'}{\Gamma_1, x : D, \Gamma_2; (\alpha, x : A, B \rightarrow \epsilon C) \vdash \text{resume } M : C \mid \epsilon'} \text{TS\_RESUME}$$

$$\boxed{\Gamma; R \vdash H : A \mid \epsilon \Rightarrow B \mid \epsilon'}$$

$$\frac{\Gamma, x : A; R \vdash M : B \mid \epsilon' \quad \epsilon \subseteq \epsilon'}{\Gamma; R \vdash \text{return } x \rightarrow M : A \mid \epsilon \Rightarrow B \mid \epsilon'} \text{THS\_RETURN}$$

$$\frac{\Gamma; R \vdash H : A \mid \epsilon \Rightarrow B \mid \epsilon' \quad \text{ty}(\text{op}) = \forall \alpha. C \hookrightarrow D \quad \Gamma, \alpha, x : C; (\alpha, x : C, D \rightarrow \epsilon' B) \vdash M : B \mid \epsilon'}{\Gamma; R \vdash H; \text{op}(x) \rightarrow M : A \mid \epsilon \uplus \{\text{op}\} \Rightarrow B \mid \epsilon'} \text{THS\_OP}$$

Fig. 2. Typing rules.

Typing rules for terms are given in the middle of Fig. 2. The first six rules are standard for the lambda calculus with let-polymorphism and a type-and-effect system. If a variable x is introduced by a let-expression and has type scheme $\forall \alpha. A$ in Γ , it is given type $A[\mathbf{B}/\alpha]$, obtained by instantiating type variables α with well-formed types $\mathbf{B}$. If x is bound by other constructors (e.g., a lambda abstraction), x is always bound to a monomorphic type and both α and $\mathbf{B}$ are the empty sequence. Note that (TS_VAR) gives any effect ϵ to the typing judgment for x . In general, ϵ in judgment $\Gamma; R \vdash M : A \mid \epsilon$ means that the evaluation of M may invoke effect operations in ϵ . Since a reference to a variable involves no effect, it is given any effect; for the same reason, value constructors are also given any effect. The rule (TS_CONST) means that the type of a constant is given by (meta-level) function ty . The typing rules for lambda abstractions and function applications are standard in the lambda calculus equipped with a type-and-effect system. The rule (TS_ABS) gives lambda abstraction $\lambda x. M$ function type $A \rightarrow \epsilon' B$ if M computes a value of B possibly with effect ϵ' by using x of type A . The rule (TS_APP) requires that (1) the argument type of function part M_1 be equivalent to the type of actual argument M_2 and (2) effect ϵ' invoked by function M_1 be contained in the whole effect ϵ . The rule (TS_WEAK) allows weakening of effects.

The next two rules are mostly standard for algebraic effects and handlers. The rule (TS_OP) is applied to effect invocations. Since $\lambda_{\text{eff}}^{\text{let}}$ supports implicit polymorphism, an invocation $\# \text{op}(M)$ of polymorphic effect op of signature $\forall \alpha. A \hookrightarrow B$ also accompanies implicit type substitution of well-formed types $\mathbf{C}$ for α . Thus, the type of argument M has to be $A[\mathbf{C}/\alpha]$ and the result of the invocation is given type $B[\mathbf{C}/\alpha]$. In addition, effect ϵ contains op . The typeability of handle-with expressions depends on the typing of handlers (TS_HANDLE), which will be explained below shortly.

The last typing rule (TS_RESUME) is the key to gaining type safety in this work. Suppose that we are given resumption type $(\alpha, x : A, B \rightarrow \epsilon C)$. Intuitively, $B \rightarrow \epsilon C$ is the type of the continuation for resumption and, therefore, argument M to resume is required to have type B . As we have discussed in Sect. 2, we avoid interference between different resumptions by renaming α , the type parameters to the effect operation, to fresh type variables β , in typechecking M . Freshness of β will be ensured when well-formedness of typing contexts $\Gamma_1, \Gamma_2, \beta, \dots$ is checked at the leaves of the type derivation. The type variables α in the type of x , the parameter to the operation, are also renamed for x to be useful in M . To see why this renaming is useful, let us consider an extension of the calculus with pairs and typechecking of an operation clause for $\text{choose}^\forall$ of signature $\forall \alpha. \alpha \times \alpha \hookrightarrow \alpha$:

$$\text{choose}^\forall(x) \rightarrow \text{resume}(\text{fst } x)$$

Variable x is assigned product type $\alpha \times \alpha$ for fresh type variable α and the body $\text{resume}(\text{fst } x)$ is typechecked under the resumption type $(\alpha, x : \alpha \times \alpha, \alpha \rightarrow \epsilon A)$ for some ϵ and A (see the typing rules for handlers for details). To typecheck $\text{resume}(\text{fst } x)$, the argument $\text{fst } x$ is required to have type β , freshly generated for this resume . Without applying renaming also to x , the clause would not

typecheck. Finally, (TS_RESUME) also requires that (1) the typing context contains α , which should have been declared at an application of the typing rule for the operation clause that surrounds this **resume** and (2) effect ϵ , which may be invoked by resumption of a continuation, be contained in the whole effect ϵ' . The binding $x : D$ in the conclusion means that parameter x to the operation clause is declared outside the resumption expression.

The typing rules for handlers are standard [3, 13, 16]. The rule (THS_RETURN) for a return clause **return** $x \rightarrow M$ checks that the body M is given a type under the assumption that argument x has type A , which is the type of the handled expression. The effect ϵ stands for effects that are not handled by the operation clauses that follow the return clause and it must be a subset of the effect ϵ' that M may cause.⁵ A handler having operation clauses is typechecked by (THS_OP), which checks that the body of the operation clause $\text{op}(x) \rightarrow M$ for op of signature $\forall \alpha. C \hookrightarrow D$ is typed at the result type B , which is the same as the type of the return clause, under the typing context extended with fresh assigned type variables α and argument x of type C , together with the resumption type $(\alpha, x : C, D \rightarrow \epsilon' B)$. The effect $\epsilon \uplus \{\text{op}\}$ in the conclusion means that the effect operation op is handled by this clause and no other clauses (in the present handler) handle it. Our semantics adopts deep handlers [13], i.e., when a handled expression invokes an effect operation, the continuation, which passed to the operation clause, is wrapped by the same handler. Thus, resumption may invoke the same effect ϵ' as the one possibly invoked by the clauses of the handler, hence $D \rightarrow \epsilon' B$ in the resumption type.

Finally, we show how the type system rejects the counterexample given in Sect. 2. The problem is in the following operation clause.

$$\text{op}(y) \rightarrow \text{resume } \lambda z_1. (\text{resume } \lambda z_2. z_1); z_1$$

where op has effect signature $\forall \alpha. \text{unit} \hookrightarrow (\alpha \rightarrow \langle \rangle \alpha)$. This clause is typechecked under resumption type $(\alpha, y : \text{unit}, \alpha \rightarrow \epsilon \alpha)$ for some ϵ . By (TS_RESUME), the two resumption expressions are assigned two different type variables γ_1 and γ_2 , and the arguments $\lambda z_1. (\text{resume } \lambda z_2. z_1); z_1$ and $\lambda z_2. z_1$ are required to have $\gamma_1 \rightarrow \epsilon \gamma_1$ and $\gamma_2 \rightarrow \epsilon \gamma_2$, respectively. However, $\lambda z_2. z_1$ cannot because z_1 is associated with γ_1 but not with γ_2 .

Remark. The rule (TS_RESUME) allows only the type of the argument to an operation clause to be renamed. Thus, other variables bound by, e.g., lambda abstractions and let-expressions outside the resumption expression cannot be used as such a type. As a result, more care may be required as to where to introduce a new variable. For example, let us consider the following operation clause (which is a variant of the example of $\text{choose}^\forall$ above).

$$\text{choose}^\forall(x) \rightarrow \text{let } y = \text{fst } x \text{ in resume } y$$

The variable x is assigned $\alpha \times \alpha$ first and the resumption requires y to be typed at fresh type variable β . This clause would be rejected in the current type system

⁵ Thus, handlers in $\lambda_{\text{eff}}^{\text{let}}$ are open [13] in the sense that a **handle-with** expression does not have to handle *all* effects caused by the handled expression.

because $\text{fst } x$ appears outside resume and, therefore, y is given type α , not β . This inconvenience may be addressed by moving down the let -binding in some cases: e.g., $\text{resume}(\text{let } y = \text{fst } x \text{ in } y)$ is well typed.

4 Intermediate Language: λ_{eff}^A

The semantics of $\lambda_{\text{eff}}^{\text{let}}$ is given by a formal elaboration to an intermediate language λ_{eff}^A , wherein type abstraction and type application appear explicitly. We define the syntax, operational semantics, and type system of $\lambda_{\text{eff}}^{\text{let}}$ and the formal elaboration from $\lambda_{\text{eff}}^{\text{let}}$ to λ_{eff}^A . Finally, we show type safety of $\lambda_{\text{eff}}^{\text{let}}$ via type preservation of the elaboration and type soundness of λ_{eff}^A .

Values	v	$::=$	$c \mid \lambda x. e$
Polymorphic values	w	$::=$	$v \mid \Lambda \alpha. w$
Terms	e	$::=$	$x \ A \mid c \mid \lambda x. e \mid e_1 \ e_2 \mid \text{let } x = \Lambda \alpha. e_1 \text{ in } e_2 \mid$ $\# \text{op}(\mathbf{A}, e) \mid \# \text{op}(\sigma, w, E) \mid \text{handle } e \text{ with } h \mid$ $\text{resume } \alpha \ x. e$
Handlers	h	$::=$	$\text{return } x \rightarrow e \mid h; \Lambda \alpha. \text{op}(x) \rightarrow e$
Evaluation contexts	E^{α^I}	$::=$	$[] \text{ (if } \alpha^I = \emptyset) \mid E^{\alpha^I} e_2 \mid v_1 E^{\alpha^I} \mid$ $\text{let } x = \Lambda \beta^{J_1}. E^{\gamma^{J_2}} \text{ in } e_2 \text{ (if } \alpha^I = \beta^{J_1}, \gamma^{J_2}) \mid$ $\# \text{op}(\mathbf{A}^J, E^{\alpha^I}) \mid \text{handle } E^{\alpha^I} \text{ with } h$

Fig. 3. Syntax of λ_{eff}^A .

4.1 Syntax

The syntax of λ_{eff}^A is shown in Fig. 3. Values, denoted by v , consist of constants and lambda abstractions. Polymorphic values, denoted by w , are values abstracted over types. Terms, denoted by e , and handlers, denoted by h , are the same as those of $\lambda_{\text{eff}}^{\text{let}}$ except for the following three points. First, type abstraction and type arguments are explicit in λ_{eff}^A : variables and effect invocations are accompanied by a sequence of types and let-bound expressions, resumption expressions, and operation clauses bind type variables. Second, a new term constructor of the form $\# \text{op}(\sigma, w, E)$ is added. It represents an intermediate state in which an effect invocation is capturing the continuation up to the closest handler for op . Here, E is an evaluation context [6] and denotes a continuation to be resumed by an operation clause handling op . In the operational semantics, an operation invocation $\# \text{op}(\mathbf{A}, v)$ is first transformed to $\# \text{op}(\mathbf{A}, v, [])$ (where $[]$ denotes the empty context or the identity continuation) and then it bubbles up by capturing its context and pushing it onto the third argument. Note that σ and w of $\# \text{op}(\sigma, w, E)$ become polymorphic when it bubbles up from the body of a type abstraction. Third, each resumption expression $\text{resume } \alpha \ x. e$ declares distinct (type) variables α and x to denote the (type) argument to an operation

Reduction rules $\boxed{e_1 \rightsquigarrow e_2}$

$$\begin{array}{ll}
c_1 \ c_2 \rightsquigarrow \zeta(c_1, c_2) & (\text{R\_CONST}) \qquad (\lambda x.e) \ v \rightsquigarrow e[v/x] \quad (\text{R\_BETA}) \\
\text{let } x = \Lambda\alpha.v \text{ in } e \rightsquigarrow e[\Lambda\alpha.v/x] & (\text{R\_LET}) \qquad \text{handle } v \text{ with } h \rightsquigarrow e[v/x] \quad (\text{R\_RETURN}) \\
& \text{(where } h^{\text{return}} = \text{return } x \rightarrow e) \\
\#op(\mathbf{A}, v) \rightsquigarrow \#op(\mathbf{A}, v, []) & (\text{R\_OP}) \\
\#op(\sigma, w, E) \ e_2 \rightsquigarrow \#op(\sigma, w, E \ e_2) & (\text{R\_OPAPP1}) \\
v_1 \ \#op(\sigma, w, E) \rightsquigarrow \#op(\sigma, w, v_1 \ E) & (\text{R\_OPAPP2}) \\
\#op'(\mathbf{A}^I, \#op(\sigma^J, w, E)) \rightsquigarrow \#op(\sigma^J, w, \#op'(\mathbf{A}^I, E)) & (\text{R\_OPOP}) \\
\text{handle } \#op(\sigma, w, E) \text{ with } h \rightsquigarrow \#op(\sigma, w, \text{handle } E \text{ with } h) & (\text{R\_OPHANDLE}) \\
& \text{(where } op \notin ops(h)) \\
\text{let } x = \Lambda\alpha^I.\#op(\sigma^J, w, E) \text{ in } e_2 \rightsquigarrow & \\
\#op(\forall \alpha^I.\sigma^J, \Lambda\alpha^I.w, \text{let } x = \Lambda\alpha^I.E \text{ in } e_2) & (\text{R\_OPLET}) \\
\text{handle } \#op(\forall \beta^J.\mathbf{A}^I, \Lambda\beta^J.v, E^{\beta^J}) \text{ with } h \rightsquigarrow & \\
e[\text{handle } E^{\beta^J} \text{ with } h/\text{resume}]_{\Lambda\beta^J.v}^{\forall \beta^J.\mathbf{A}^I} [\mathbf{A}^I[\perp/\beta^J]/\alpha^I][v[\perp/\beta^J]/x] & (\text{R\_HANDLE}) \\
& \text{(where } h^{\text{op}} = \Lambda\alpha^I.op(x) \rightarrow e)
\end{array}$$

Evaluation rules $\boxed{e_1 \longrightarrow e_2}$

$$\frac{e_1 \rightsquigarrow e_2}{E[e_1] \longrightarrow E[e_2]} \text{ E\_EVAL}$$

Fig. 4. Semantics of λ_{eff}^A .

clause, whereas a single variable declared at $op(x) \rightarrow M$ and implicit type variables are used for the same purpose in $\lambda_{\text{eff}}^{\text{let}}$. For example, the $\lambda_{\text{eff}}^{\text{let}}$ operation clause $\text{choose}^\forall(x) \rightarrow \text{resume}(\text{fst } x)$ is translated to $\Lambda\alpha.\text{choose}^\forall(x) \rightarrow \text{resume } \beta y.(\text{fst } y)$. This change simplifies the semantics.

Evaluation contexts, denoted by E^α , are standard for the lambda calculus with call-by-value, left-to-right evaluation except for two points. First, they contain the form $\text{let } x = \Lambda\alpha.E^\beta \text{ in } e_2$, which allows the body of a type abstraction to be evaluated. Second, the metavariable E for evaluation contexts is indexed by type variables α , meaning that the hole in the context appears under type abstractions binding α . For example, $\text{let } x = \Lambda\alpha.\text{let } y = \Lambda\beta.[] \text{ in } e_2 \text{ in } e_1$ is denoted by $E^{\alpha, \beta}$ and, more generally, $\text{let } x = \Lambda\beta^{J_1}.E^{\gamma^{J_2}} \text{ in } e$ is denoted by $E^{\beta^{J_1}, \gamma^{J_2}}$. (Here, $\beta^{J_1}, \gamma^{J_2}$ stands for the concatenation of the two sequences β^{J_1} and γ^{J_2} .) If α is not important, we simply write E for E^α . We often use the term “continuation” to mean “evaluation context,” especially when it is expected to be resumed.

As usual, substitution $e[w/x]$ of w for x in e is defined in a capture-avoiding manner. Since variables come along with type arguments, the case for variables is defined as follows:

$$(x \mathbf{A})[\Lambda\alpha.v/x] \stackrel{\text{def}}{=} v[\mathbf{A}/\alpha]$$

Application of substitution $[\Lambda\alpha^I.v/x]$ to $x \mathbf{A}^J$, where $I \neq J$, is undefined. We define free type variables $ftv(e)$ and $ftv(E)$ in e and E , respectively, as usual.

4.2 Semantics

The semantics of λ_{eff}^A is given in the small-step style and consists of two relations: the reduction relation $\rightsquigarrow$, which is for basic computation, and the evaluation relation $\longrightarrow$, which is for top-level execution. Figure 4 shows the rules for these relations. In what follows, we write h^{return} for the return clause of handler h , $\text{ops}(h)$ for the set of effect operations handled by h , and h^{op} for the operation clause for op in h .

Most of the reduction rules are standard [13, 16]. A constant application $c_1 c_2$ reduces to $\zeta(c_1, c_2)$ (R_CONST), where function ζ maps a pair of constants to another constant. A function application $(\lambda x.e) v$ and a let-expression $\text{let } x = \Lambda\alpha.v \text{ in } e$ reduce to $e[v/x]$ (R_BETA) and $e[\Lambda\alpha.v/x]$ (R_LET), respectively. If a handled expression is a value v , the **handle-with** expression reduces to the body of the return clause where v is substituted for the parameter x (R_RETURN). An effect invocation $\#\text{op}(\mathbf{A}, v)$ reduces to $\#\text{op}(\mathbf{A}, v, [])$ with the identity continuation, as explained above (R_OP); the process of capturing its evaluation context is expressed by the rules (R_OPAPP1), (R_OPAPP2), (R_OPOP), (R_OPHANDLE), and (R_OPLET). The rule (R_OPHANDLE) can be applied only if the handler h does *not* handle op . The rule (R_OPLET) is applied to a let-expression where $\#\text{op}(\sigma^J, w, E)$ appears under a type abstraction with bound type variables α^I . Since σ^J and w may refer to α^I , the reduction result binds α^I in both σ^J and w . We write $\forall \alpha^I.\sigma^J$ for a sequence $\forall \alpha^I.\sigma_{j_1}, \dots, \forall \alpha^I.\sigma_{j_n}$ of type schemes (where $J = \{j_1, \dots, j_n\}$).

The crux of the semantics is (R_HANDLE): it is applied when $\#\text{op}(\sigma^I, w, E)$ reaches the handler h that handles op . Since the handled term $\#\text{op}(\sigma^I, w, E)$ is constructed from an effect invocation $\#\text{op}(\mathbf{A}^I, v)$, if the captured continuation E binds type variables β^J , the same type variables β^J should have been added to $\mathbf{A}^I$ and v along the capture. Thus, the handled expression on the left-hand side of the rule takes the form $\#\text{op}(\forall \beta^J.\mathbf{A}^I, \Lambda\beta^J.v, E^{\beta^J})$ (with the same type variables β^J).

The right-hand side of (R_HANDLE) involves three types of substitution: continuation substitution $[\text{handle } E^{\beta^J} \text{ with } h/\text{resume}]_{\Lambda\beta^J.v}^{\forall \beta^J.\mathbf{A}^I}$ for resumptions, type substitution for α^I , and value substitution for x . We explain them one by one below. In the following, let $h^{\text{op}} = \Lambda\alpha^I.\text{op}(x) \rightarrow e$ and $E'^{\beta^J} = \text{handle } E^{\beta^J} \text{ with } h$.

Continuation Substitution. Let us start with a simple case where the sequence β^J is empty. Intuitively, continuation substitution $[E'/\text{resume}]_v^{A^I}$ replaces a resumption expression $\text{resume } \gamma^I z.e'$ in the body e with $E'[v']$, where v' is the value of e' , and substitutes A^I and v (arguments to the invocation of op) for γ^I and z , respectively. Therefore, assuming resume does not appear in e' , we define $(\text{resume } \gamma^I z.e')[E'/\text{resume}]_v^{A^I}$ to be $\text{let } y = e'[\mathbf{A}^I/\gamma^I][v/z] \text{ in } E'[y]$ (for fresh y). Note that the evaluation of e' takes place outside of E so that an invocation of an effect in e' is *not* handled by handlers in E . When β^J is not empty,

$$\begin{aligned} (\text{resume } \gamma^I z.e')[E^{\beta^J}/\text{resume}]_{\Lambda\beta^J.v}^{\forall\beta^J.A^I} &\stackrel{\text{def}}{=} \\ \text{let } y = \Lambda\beta^J.e'[\mathbf{A}^I/\gamma^I][v/z] &\text{ in } E^{\beta^J}[y\beta^J]. \end{aligned}$$

(The differences from the simple case are shaded.) The idea is to bind β^J that appear free in A^I and v by type abstraction at let and to instantiate with the same variables at $y\beta^J$, where β^J are bound by type abstractions in E^{β^J} .

Continuation substitution is formally defined as follows:

Definition 2 (Continuation substitution). *Substitution of continuation E^{β^J} for resumptions in e , written $e[E^{\beta^J}/\text{resume}]_{\Lambda\beta^J.v}^{\forall\beta^J.A^I}$, is defined in a capture-avoiding manner, as follows (we describe only the important cases):*

$$\begin{aligned} (\text{resume } \gamma^I z.e)[E^{\beta^J}/\text{resume}]_{\Lambda\beta^J.v}^{\forall\beta^J.A^I} &\stackrel{\text{def}}{=} \\ \text{let } y = \Lambda\beta^J.e[E^{\beta^J}/\text{resume}]_{\Lambda\beta^J.v}^{\forall\beta^J.A^I} [\mathbf{A}^I/\gamma^I][v/z] &\text{ in } E^{\beta^J}[y\beta^J] \\ (\text{if } (ftv(e) \cup ftv(E^{\beta^J})) \cap \{\beta^J\} = \emptyset \text{ and } y \text{ is fresh}) & \\ (\text{return } x \rightarrow e)[E/\text{resume}]_w^\sigma &\stackrel{\text{def}}{=} \text{return } x \rightarrow e[E/\text{resume}]_w^\sigma \\ (h'; \Lambda\gamma^J.\text{op}(x) \rightarrow e)[E/\text{resume}]_w^{\sigma^I} &\stackrel{\text{def}}{=} h'[E/\text{resume}]_w^{\sigma^I}; \Lambda\gamma^J.\text{op}(x) \rightarrow e \end{aligned}$$

The second and third clauses (for a handler) mean that continuation substitution is applied only to return clauses.

Type and Value Substitution. The type and value substitutions $A^I[\perp^J/\beta^J]$ and $v[\perp^J/\beta^J]$, respectively, in (R_HANDLE) are for (type) parameters in $h^{\text{op}} = \Lambda\alpha^I.\text{op}(x) \rightarrow e$. The basic idea is to substitute A^I for β^I and v for x —similarly to continuation substitution. We erase free type variables β^J in A^I and v by substituting the designated base type $\perp$ for all of them. (We write $A^I[\perp^J/\beta^J]$ and $v[\perp^J/\beta^J]$ for the types and value, respectively, after the erasure.)

The evaluation rule is ordinary: Evaluation of a term proceeds by reducing a subterm under an evaluation context.

4.3 Type System

The type system of λ_{eff}^A is similar to that of $\lambda_{\text{eff}}^{\text{let}}$ and has five judgments: well-formedness of typing contexts $\vdash \Gamma$; well formedness of type schemes $\Gamma \vdash \sigma$; term

typing judgment $\Gamma; r \vdash e : A \mid \epsilon$; handler typing judgment $\Gamma; r \vdash h : A \mid \epsilon \Rightarrow B \mid \epsilon'$; and continuation typing judgment $\Gamma \vdash E : \forall \alpha. A \multimap B \mid \epsilon$. The first two are defined in the same way as those of $\lambda_{\text{eff}}^{\text{let}}$. The last judgment means that a term obtained by filling the hole of E with a term having A under Γ, α is typed at B under Γ and possibly involves effect ϵ . A resumption type r is similar to R but does not contain an argument variable.

Definition 3 (Resumption type). *Resumption types in λ_{eff}^A , denoted by r , are defined as follows:*

$$r ::= \text{none} \mid (\alpha, A, B \rightarrow \epsilon C) \\ (\text{if } \text{ftv}(A) \cup \text{ftv}(B) \subseteq \{\alpha\} \text{ and } \text{ftv}(C) \cap \{\alpha\} = \emptyset)$$

Typing rules

$\Gamma; r \vdash e : A \mid \epsilon$	
$\frac{\vdash \Gamma \quad x : \forall \alpha. A \in \Gamma \quad \Gamma \vdash B}{\Gamma; r \vdash x B : A[B/\alpha] \mid \epsilon} \text{T\_VAR} \qquad \frac{\vdash \Gamma}{\Gamma; r \vdash c : \text{ty}(c) \mid \epsilon} \text{T\_CONST}$	
$\frac{\Gamma, x : A; r \vdash e : B \mid \epsilon'}{\Gamma; r \vdash \lambda x. e : A \rightarrow \epsilon' B \mid \epsilon} \text{T\_ABS}$	
$\frac{\Gamma; r \vdash e_1 : A \rightarrow \epsilon' B \mid \epsilon \quad \Gamma; r \vdash e_2 : A \mid \epsilon \quad \epsilon' \subseteq \epsilon}{\Gamma; r \vdash e_1 e_2 : B \mid \epsilon} \text{T\_APP}$	
$\frac{\text{ty}(\text{op}) = \forall \alpha. A \multimap B \quad \text{op} \in \epsilon \quad \Gamma; r \vdash e : A[C/\alpha] \mid \epsilon \quad \Gamma \vdash C}{\Gamma; r \vdash \# \text{op}(C, e) : B[C/\alpha] \mid \epsilon} \text{T\_OP}$	
$\frac{\text{ty}(\text{op}) = \forall \alpha^I. A \multimap B \quad \text{op} \in \epsilon \quad \Gamma \vdash \forall \beta^J. C^I \quad \Gamma, \beta^J; r \vdash v : A[C^I/\alpha^I] \mid \epsilon \quad \Gamma \vdash E^{\beta^J} : \forall \beta^J. (B[C^I/\alpha^I]) \multimap D \mid \epsilon}{\Gamma; r \vdash \# \text{op}(\forall \beta^J. C^I, \Lambda \beta^J. v, E^{\beta^J}) : D \mid \epsilon} \text{T\_OPCONT}$	
$\frac{\Gamma; r \vdash e : A \mid \epsilon' \quad \epsilon' \subseteq \epsilon}{\Gamma; r \vdash e : A \mid \epsilon} \text{T\_WEAK}$	
$\frac{\Gamma; r \vdash e : A \mid \epsilon \quad \Gamma; r \vdash h : A \mid \epsilon \Rightarrow B \mid \epsilon'}{\Gamma; r \vdash \text{handle } e \text{ with } h : B \mid \epsilon'} \text{T\_HANDLE}$	
$\frac{\Gamma, \alpha; r \vdash e_1 : A \mid \epsilon \quad \Gamma, x : \forall \alpha. A; r \vdash e_2 : B \mid \epsilon}{\Gamma; r \vdash \text{let } x = \Lambda \alpha. e_1 \text{ in } e_2 : B \mid \epsilon} \text{T\_LET}$	
$\frac{\alpha \in \Gamma \quad \Gamma, \beta, x : A[\beta/\alpha]; (\alpha, A, B \rightarrow \epsilon C) \vdash e : B[\beta/\alpha] \mid \epsilon' \quad \epsilon \subseteq \epsilon'}{\Gamma; (\alpha, A, B \rightarrow \epsilon C) \vdash \text{resume } \beta x. e : C \mid \epsilon'} \text{T\_RESUME}$	

Fig. 5. Typing rules for terms in λ_{eff}^A .

The typing rules for terms, shown in Fig. 5, and handlers, shown in the upper half of Fig. 6, are similar to those of $\lambda_{\text{eff}}^{\text{let}}$ except for a new rule (T_OPCONT), which is applied to an effect invocation $\# \text{op}(\forall \beta^J. C^I, \Lambda \beta^J. v, E^{\beta^J})$ with a continuation. Let $ty(\text{op}) = \forall \alpha^I. A \hookrightarrow B$. Since op should have been invoked with C^I and v under type abstractions with bound type variables β^J , the argument v has type $A[C^I/\alpha^I]$ under the typing context extended with β^J . Similarly, the hole of E^{β^J} expects to be filled with the result of the invocation, i.e., a value of $B[C^I/\alpha^I]$. Since the continuation denotes the context before the evaluation, its result type matches with the type of the whole term.

The typing rules for continuations are shown in the lower half of Fig. 6. They are similar to the corresponding typing rules for terms except that a subterm is replaced with a continuation. In (TE_LET), the continuation $\text{let } x = \Lambda \alpha. E \text{ in } e$ has type $\forall \alpha. \sigma \multimap B$ because the hole of E appears inside the scope of α .

$$\boxed{\Gamma; r \vdash h : A \mid \epsilon \Rightarrow B \mid \epsilon'}$$

$$\frac{\Gamma, x : A; r \vdash e : B \mid \epsilon' \quad \epsilon \subseteq \epsilon'}{\Gamma; r \vdash \text{return } x \rightarrow e : A \mid \epsilon \Rightarrow B \mid \epsilon'} \text{ TH\_RETURN}$$

$$\frac{\Gamma; r \vdash h : A \mid \epsilon \Rightarrow B \mid \epsilon' \quad ty(\text{op}) = \forall \alpha. C \hookrightarrow D \quad \Gamma, \alpha, x : C; (\alpha, C, D \rightarrow \epsilon' B) \vdash e : B \mid \epsilon'}{\Gamma; r \vdash h; \Lambda \alpha. \text{op}(x) \rightarrow e : A \mid \epsilon \uplus \{\text{op}\} \Rightarrow B \mid \epsilon'} \text{ TH\_OP}$$

$$\boxed{\Gamma \vdash E : \sigma \multimap A \mid \epsilon}$$

$$\frac{}{\Gamma \vdash [] : A \multimap A \mid \epsilon} \text{ TE\_HOLE}$$

$$\frac{\Gamma \vdash E : \sigma \multimap (A \rightarrow \epsilon' B) \mid \epsilon \quad \Gamma; \text{none} \vdash e_2 : A \mid \epsilon \quad \epsilon' \subseteq \epsilon}{\Gamma \vdash E e_2 : \sigma \multimap B \mid \epsilon} \text{ TE\_APP1}$$

$$\frac{\Gamma; \text{none} \vdash v_1 : (A \rightarrow \epsilon' B) \mid \epsilon \quad \Gamma \vdash E : \sigma \multimap A \mid \epsilon \quad \epsilon' \subseteq \epsilon}{\Gamma \vdash v_1 E : \sigma \multimap B \mid \epsilon} \text{ TE\_APP2}$$

$$\frac{ty(\text{op}) = \forall \alpha. A \hookrightarrow B \quad \text{op} \in \epsilon \quad \Gamma \vdash E : \sigma \multimap A[C/\alpha] \mid \epsilon \quad \Gamma \vdash C}{\Gamma \vdash \# \text{op}(C, E) : \sigma \multimap B[C/\alpha] \mid \epsilon} \text{ TE\_OP}$$

$$\frac{\Gamma \vdash E : \sigma \multimap A \mid \epsilon \quad \Gamma; \text{none} \vdash h : A \mid \epsilon \Rightarrow B \mid \epsilon'}{\Gamma \vdash \text{handle } E \text{ with } h : \sigma \multimap B \mid \epsilon'} \text{ TE\_HANDLE}$$

$$\frac{\Gamma \vdash E : \sigma \multimap A \mid \epsilon' \quad \epsilon' \subseteq \epsilon}{\Gamma \vdash E : \sigma \multimap A \mid \epsilon} \text{ TE\_WEAK}$$

$$\frac{\Gamma, \alpha \vdash E : \sigma \multimap A \mid \epsilon \quad \Gamma, x : \forall \alpha. A; \text{none} \vdash e : B \mid \epsilon}{\Gamma \vdash \text{let } x = \Lambda \alpha. E \text{ in } e : \forall \alpha. \sigma \multimap B \mid \epsilon} \text{ TE\_LET}$$

Fig. 6. Typing rules for handlers and continuations in λ_{eff}^A .

4.4 Elaboration

This section defines the elaboration from $\lambda_{\text{eff}}^{\text{let}}$ to λ_{eff}^A . The important difference between the two languages from the viewpoint of elaboration is that, whereas the parameter of an operation clause is referred to by a single variable in $\lambda_{\text{eff}}^{\text{let}}$, it is done by one or more variables in λ_{eff}^A . Therefore, one variable in $\lambda_{\text{eff}}^{\text{let}}$ is represented by multiple variables (required for each **resume**) in λ_{eff}^A . We use S , a mapping from variables to variables, to make the correspondence between variable names. We write $S \circ \{x \mapsto y\}$ for the same mapping as S except that x is mapped to y .

Elaboration is defined by two judgments: term elaboration judgment $\Gamma; R \vdash M : A \mid \epsilon \triangleright^S e$, which denotes elaboration from a typing derivation of judgment $\Gamma; R \vdash M : A \mid \epsilon$ to e with S , and handler elaboration judgment $\Gamma; R \vdash H : A \mid \epsilon \Rightarrow B \mid \epsilon' \triangleright^S h$, which denotes elaboration from a typing derivation of judgment $\Gamma; R \vdash H : A \mid \epsilon \Rightarrow B \mid \epsilon'$ to h with S .

Term elaboration rules

$$\boxed{\Gamma; R \vdash M : A \mid \epsilon \triangleright^S e}$$

$$\frac{\vdash \Gamma \quad x : \forall \alpha. A \in \Gamma \quad \Gamma \vdash B}{\Gamma; R \vdash x : A[B/\alpha] \mid \epsilon \triangleright^S S(x) B} \text{ELAB\_VAR}$$

$$\frac{\Gamma, x : A; R \vdash M : B \mid \epsilon' \triangleright^{S \circ \{x \mapsto x\}} e}{\Gamma; R \vdash \lambda x. M : A \rightarrow \epsilon' B \mid \epsilon \triangleright^S \lambda x. e} \text{ELAB\_ABS}$$

$$\frac{\Gamma; R \vdash M : A \mid \epsilon \triangleright^S e \quad \Gamma; R \vdash H : A \mid \epsilon \Rightarrow B \mid \epsilon' \triangleright^S h}{\Gamma; R \vdash \text{handle } M \text{ with } H : B \mid \epsilon' \triangleright^S \text{handle } e \text{ with } h} \text{ELAB\_HANDLE}$$

$$\frac{\Gamma, \alpha; R \vdash M_1 : A \mid \epsilon \triangleright^S e_1 \quad S' = S \circ \{x \mapsto x\}}{\Gamma, x : \forall \alpha. A; R \vdash M_2 : B \mid \epsilon \triangleright^{S'} e_2} \text{ELAB\_LET}$$

$$\frac{\Gamma, \alpha; R \vdash M_1 : A \mid \epsilon \triangleright^S e_1 \quad S' = S \circ \{x \mapsto x\}}{\Gamma; R \vdash \text{let } x = M_1 \text{ in } M_2 : B \mid \epsilon \triangleright^S \text{let } x = \Lambda \alpha. e_1 \text{ in } e_2} \text{ELAB\_LET}$$

$$\frac{R = (\alpha, x : A, B \rightarrow \epsilon C) \quad \vdash \Gamma_1, x : D, \Gamma_2 \quad \alpha \in \Gamma_1 \quad \epsilon \subseteq \epsilon' \quad y \text{ is fresh} \quad S' = S \circ \{x \mapsto y\}}{\Gamma_1, \Gamma_2, \beta, x : A[\beta/\alpha]; R \vdash M : B[\beta/\alpha] \mid \epsilon' \triangleright^{S'} e} \text{ELAB\_RESUME}$$

$$\frac{\Gamma_1, \Gamma_2, \beta, x : A[\beta/\alpha]; R \vdash M : B[\beta/\alpha] \mid \epsilon' \triangleright^{S'} e}{\Gamma_1, x : D, \Gamma_2; R \vdash \text{resume } M : C \mid \epsilon' \triangleright^S \text{resume } \beta y. e} \text{ELAB\_RESUME}$$

Handler elaboration rules

$$\boxed{\Gamma; R \vdash H : A \mid \epsilon \Rightarrow B \mid \epsilon' \triangleright^S h}$$

$$\frac{\Gamma, x : A; R \vdash M : B \mid \epsilon' \triangleright^{S \circ \{x \mapsto x\}} e \quad \epsilon \subseteq \epsilon'}{\Gamma; R \vdash \text{return } x \rightarrow M : A \mid \epsilon \Rightarrow B \mid \epsilon' \triangleright^S \text{return } x \rightarrow e} \text{ELABH\_RETURN}$$

$$\frac{ty(\text{op}) = \forall \alpha. C \hookrightarrow D \quad \Gamma; R \vdash H : A \mid \epsilon \Rightarrow B \mid \epsilon' \triangleright^S h \quad \Gamma, \alpha, x : C; (\alpha, x : C, D \rightarrow \epsilon' B) \vdash M : B \mid \epsilon' \triangleright^{S \circ \{x \mapsto x\}} e}{\Gamma; R \vdash H; \text{op}(x) \rightarrow M : A \mid \epsilon \uplus \{\text{op}\} \Rightarrow B \mid \epsilon' \triangleright^S h; \Lambda \alpha. \text{op}(x) \rightarrow e} \text{ELABH\_OP}$$

Fig. 7. Elaboration rules (excerpt).

Selected elaboration rules are shown in Fig. 7; the complete set of the rules is found in the full version of the paper. The elaboration rules are straightforward except for the use of S . A variable x is translated to $S(x)$ (ELAB_VAR) and, every time a new variable is introduced, S is extended: see the rules other than (ELAB_VAR) and (ELAB_HANDLE).

4.5 Properties

We show type safety of $\lambda_{\text{eff}}^{\text{let}}$, i.e., a well-typed program in $\lambda_{\text{eff}}^{\text{let}}$ does not get stuck, by proving (1) type preservation of the elaboration from $\lambda_{\text{eff}}^{\text{let}}$ to λ_{eff}^A and (2) type soundness of λ_{eff}^A . Term M is a well-typed program of A if and only if $\emptyset; \text{none} \vdash M : A \mid \langle \rangle$.

The first can be shown easily. We write $\emptyset$ also for the identity mapping for variables.

Theorem 1 (Elaboration is type-preserving). *If M is a well-typed program of A , then $\emptyset; \text{none} \vdash M : A \mid \langle \rangle \triangleright^\emptyset e$ and $\emptyset; \text{none} \vdash e : A \mid \langle \rangle$ for some e .*

We show the second—type soundness of λ_{eff}^A —via progress and subject reduction [25]. We write Δ for a typing context that consists only of type variables. Progress can be shown as usual.

Lemma 1 (Progress). *If $\Delta; \text{none} \vdash e : A \mid \epsilon$, then (1) $e \longrightarrow e'$ for some e' , (2) e is a value, or (3) $e = \# \text{op}(\sigma, w, E)$ for some $\text{op} \in \epsilon$, σ , w , and E .*

A key lemma to show subject reduction is type preservation of continuation substitution.

Lemma 2 (Continuation substitution). *Suppose that $\Gamma \vdash \forall \beta^J. C^I$ and $\Gamma \vdash E^{\beta^J} : \forall \beta^J. (B[C^I/\alpha^I]) \multimap D \mid \epsilon$ and $\Gamma, \beta^J \vdash v : A[C^I/\alpha^I]$.*

1. *If $\Gamma; (\alpha^I, A, B \rightarrow \epsilon \ D) \vdash e : D' \mid \epsilon'$, then $\Gamma; \text{none} \vdash e[E^{\beta^J}/\text{resume}]_{\Lambda\beta^J.v}^{\forall\beta^J.C^I} : D' \mid \epsilon'$.*
2. *If $\Gamma; (\alpha^I, A, B \rightarrow \epsilon \ D) \vdash h : D_1 \mid \epsilon_1 \Rightarrow D_2 \mid \epsilon_2$, then $\Gamma; \text{none} \vdash h[E^{\beta^J}/\text{resume}]_{\Lambda\beta^J.v}^{\forall\beta^J.C^I} : D_1 \mid \epsilon_1 \Rightarrow D_2 \mid \epsilon_2$.*

Using the continuation substitution lemma as well as other lemmas, we show subject reduction.

Lemma 3 (Subject reduction)

1. *If $\Delta; \text{none} \vdash e_1 : A \mid \epsilon$ and $e_1 \rightsquigarrow e_2$, then $\Delta; \text{none} \vdash e_2 : A \mid \epsilon$.*
2. *If $\Delta; \text{none} \vdash e_1 : A \mid \epsilon$ and $e_1 \longrightarrow e_2$, then $\Delta; \text{none} \vdash e_2 : A \mid \epsilon$.*

We write $e \not\rightarrow$ if and only if e cannot evaluate further. Moreover, $\longrightarrow^*$ denotes the reflexive and transitive closure of the evaluation relation $\longrightarrow$.

Theorem 2 (Type soundness of λ_{eff}^A). *If $\Delta; \text{none} \vdash e : A \mid \epsilon$ and $e \longrightarrow^* e'$ and $e' \not\rightarrow$, then (1) e' is a value or (2) $e' = \# \text{op}(\sigma, w, E)$ for some $\text{op} \in \epsilon, \sigma, w$, and E .*

Now, type safety of $\lambda_{\text{eff}}^{\text{let}}$ is obtained as a corollary of Theorems 1 and 2.

Corollary 1 (Type safety of $\lambda_{\text{eff}}^{\text{let}}$). *If M is a well-typed program of A , there exists some e such that $\emptyset; \text{none} \vdash M : A \mid \langle \rangle \triangleright^\emptyset e$ and e does not get stuck.*

5 Related Work

5.1 Polymorphic Effects and Let-Polymorphism

Many researchers have attacked the problem of combining effects—not necessarily algebraic—and let-polymorphism so far [1, 2, 10, 12, 14, 17, 23, 24]. In particular, most of them have focused on ML-style polymorphic references. The algebraic effect handlers dealt with in this paper seem to be unable to implement general ML-style references—i.e., give an appropriate implementation to a set of effect operations **new** with the signature $\forall \alpha. \alpha \hookrightarrow \alpha \text{ ref}$, **get** with $\forall \alpha. \alpha \text{ ref} \hookrightarrow \alpha$, and **put** with $\forall \alpha. \alpha \times \alpha \text{ ref} \hookrightarrow \text{unit}$ for abstract datatype $\alpha \text{ ref}$ —even without the restriction on handlers because each operation clause in a handler assigns type variables locally and it is impossible to share such type variables between operation clauses.⁶ Nevertheless, their approaches would be applicable to algebraic effects and handlers.

A common idea in the literature is to restrict the form of expressions bound by polymorphic let. Thus, they are complementary to our approach in that they restrict how effect operations are used whereas we restrict how effect operations are implemented.

Value restriction [23, 24], a standard way adopted in ML-like languages, restricts polymorphic let-bound expressions to syntactic values. Garrigue [10] relaxes the value restriction so that, if a let-bound expression is not a syntactic value, type variables that appear only at positive positions in the type of the expression can be generalized. Although the (relaxed) value restriction is a quite clear criterion that indicates what let-bound expressions can be polymorphic safely and it even accepts interfering handlers, it is too restrictive in some cases. We give an example for such a case below.

effect choose[∀] : $\forall \alpha. \alpha \times \alpha \hookrightarrow \alpha$

```
let f1 () =
  let g = #choose∀(fst, snd) in
  if g (true, false) then g (-1, 1) else g (1, -1)
```

⁶ One possible approach to dealing with ML-style references is to extend algebraic effects and handlers so that a handler for *parameterized* effects can be connected with dynamic resources [3].

In the definition of function `f1`, variable `g` is used polymorphically. Execution of this function under an appropriate handler would succeed, and in fact our calculus accepts it. By contrast, the (relaxed) value restriction rejects it because the let-bound expression `#choose∀(fst, snd)` is not a syntactic value and the type variable `g` appear in both positive and negative positions, and so `g` is assigned a monomorphic type. A workaround for this problem is to make a function wrapper that calls either of `fst` or `snd` depending on the Boolean value chosen by `choose∀`:

```
let f2 () =
  let b = #choose∀(true, false) in
  let g = λx. if b then (fst x) else (snd x) in
  if g (true, false) then g (-1, 1) else g (1, -1)
```

However, this workaround makes the program complicated and incurs additional run-time cost for the branching and an extra call to the wrapper function.

Asai and Kameyama [2] study a combination of let-polymorphism with delimited control operators `shift/reset` [4]. They allow a let-bound expression to be polymorphic if it invokes no control operation. Thus, the function `f1` above would be rejected in their approach.

Another research line to restrict the use of effects is to allow only type variables unrelated to effect invocations to be generalized. Tofte [23] distinguishes between applicative type variables, which cannot be used for effect invocations, and imperative ones, which can be used, and proposes a type system that enforces restrictions that (1) type variables of imperative operations can be instantiated only with types wherein all type variables are imperative and (2) if a let-bound expression is not a syntactic value, only applicative type variables can be generalized. Leroy and Weis [17] allow generalization only of type variables that do not appear in a parameter type to the reference type in the type of a let-expression. To detect the hidden use of references, their type system gives a term not only a type but also the types of free variables used in the term. Standard ML of New Jersey (before ML97) adopted weak polymorphism [1], which was later formalized and investigated deeply by Hoang et al. [12]. Weak polymorphism equips a type variable with the number of function calls after which a value of a type containing the type variable will be passed to an imperative operation. The type system ensures that type variables with positive numbers are not related to imperative constructs, and so such type variables can be generalized safely. In this line of research, the function `f1` above would not typecheck because generalized type variables are used to instantiate those of the effect signature, although it could be rewritten to an acceptable one by taking care not to involve type variables in effect invocation.

```
let f3 () =
  let g = if #choose∀(true, false) then fst then snd in
  if g (true, false) then g (-1, 1) else g (1, -1)
```

More recently, Kammar and Pretnar [14] show that *parameterized* algebraic effects and handlers do not need the value restriction *if* the type variables used

in an effect invocation are not generalized. Thus, as the other work that restricts generalized type variables, their approach would reject function `f1` but would accept `f3`.

5.2 Algebraic Effects and Handlers

Algebraic effects [20] are a way to represent the denotation of an effect by giving a set of operations and an equational theory that capture their properties. Algebraic effect handlers, introduced by Plotkin and Pretnar [21], make it possible to provide user-defined effects. Algebraic effect handlers have been gaining popularity owing to their flexibility and have been made available as libraries [13, 15, 26] or as primitive features of languages, such as Eff [3], Koka [16], Frank [18], and Multicore OCaml [5]. In these languages, let-bound expressions that can be polymorphic are restricted to values or pure expressions.

Recently, Forster et al. [9] investigate the relationships between algebraic effect handlers and other mechanisms for user-defined effects—delimited control shift0 [19] and monadic reflection [7, 8]—conjecturing that there would be no type-preserving translation from a language with delimited control or monadic reflection to one with algebraic effect handlers. It would be an interesting direction to export our idea to delimited control and monadic reflection.

6 Conclusion

There has been a long history of collaboration between effects and let-polymorphism. This work focuses on polymorphic algebraic effects and handlers, wherein the type signature of an effect operation can be polymorphic and an operation clause has a type binder, and shows that a naive combination of polymorphic effects and let-polymorphism breaks type safety. Our novel observation to address this problem is that any let-bound expression can be polymorphic safely if resumptions from a handler do not interfere with each other. We formalized this idea by developing a type system that requires the argument of each resumption expression to have a type obtained by renaming the type variables assigned in the operation clause to those assigned in the resumption. We have proven that a well-typed program in our type system does not get stuck via elaboration to an intermediate language wherein type information appears explicitly.

There are many directions for future work. The first is to address the problem, described at the end of Sect. 3, that renaming the type variables assigned in an operation clause to those assigned in a resumption expression is allowed for the argument of the clause but not for variables bound by lambda abstractions and let-expressions outside the resumption expression. Second, we are interested in incorporating other features from the literature on algebraic effect handlers, such as dynamic resources [3] and parameterized algebraic effects, and restriction techniques that have been developed for type-safe imperative programming with let-polymorphism such as (relaxed) value restriction [10, 23, 24]. For example, we

would like to develop a type system that enforces the non-interfering restriction only to handlers implementing effect operations invoked in polymorphic computation. We also expect that it is possible to determine whether implementations of an effect operation have no interfering resumption from the type signature of the operation, as relaxed value restriction makes it possible to find safely generalizable type variables from the type of a let-bound expression [10]. Finally, we are also interested in implementing our idea for a language with effect handlers such as Koka [16] and in applying the idea of analyzing handlers to other settings such as dependent typing.

Acknowledgments. We would like to thank the anonymous reviewers for their valuable comments. This work was supported in part by ERATO HASUO Metamathematics for Systems Design Project (No. JPMJER1603), JST (Sekiyama), and JSPS KAKENHI Grant Number JP15H05706 (Igarashi).

References

1. Appel, A.W., MacQueen, D.B.: Standard ML of New Jersey. In: Maluszyński, J., Wirsing, M. (eds.) PLILP 1991. LNCS, vol. 528, pp. 1–13. Springer, Heidelberg (1991). https://doi.org/10.1007/3-540-54444-5_83
2. Asai, K., Kameyama, Y.: Polymorphic delimited continuations. In: Shao, Z. (ed.) APLAS 2007. LNCS, vol. 4807, pp. 239–254. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-76637-7_16
3. Bauer, A., Pretnar, M.: Programming with algebraic effects and handlers. *J. Log. Algebr. Methods Program.* **84**(1), 108–123 (2015). <https://doi.org/10.1016/j.jlamp.2014.02.001>
4. Danvy, O., Filinski, A.: Abstracting control. In: LISP and Functional Programming, pp. 151–160 (1990). <https://doi.org/10.1145/91556.91622>
5. Dolan, S., Eliopoulos, S., Hillerström, D., Madhavapeddy, A., Sivaramakrishnan, K.C., White, L.: Concurrent system programming with effect handlers. In: Wang, M., Owens, S. (eds.) TFP 2017. LNCS, vol. 10788, pp. 98–117. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-89719-6_6
6. Felleisen, M., Hieb, R.: The revised report on the syntactic theories of sequential control and state. *Theor. Comput. Sci.* **103**(2), 235–271 (1992). [https://doi.org/10.1016/0304-3975\(92\)90014-7](https://doi.org/10.1016/0304-3975(92)90014-7)
7. Filinski, A.: Representing monads. In: Proceedings of the 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, pp. 446–457 (1994). <https://doi.org/10.1145/174675.178047>
8. Filinski, A.: Monads in action. In: Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, pp. 483–494 (2010). <https://doi.org/10.1145/1706299.1706354>
9. Forster, Y., Kammar, O., Lindley, S., Pretnar, M.: On the expressive power of user-defined effects: effect handlers, monadic reflection, delimited control. *PACMPL* **1**(ICFP), 13:1–13:29 (2017). <https://doi.org/10.1145/3110257>
10. Garrigue, J.: Relaxing the value restriction. In: Kameyama, Y., Stuckey, P.J. (eds.) FLOPS 2004. LNCS, vol. 2998, pp. 196–213. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24754-8_15
11. Harper, R., Lillibridge, M.: Polymorphic type assignment and CPS conversion. *Lisp Symb. Comput.* **6**(3–4), 361–380 (1993)

12. Hoang, M., Mitchell, J.C., Viswanathan, R.: Standard ML-NJ weak polymorphism and imperative constructs. In: *Proceedings of the Eighth Annual Symposium on Logic in Computer Science, LICS 1993* (1993)
13. Kammar, O., Lindley, S., Oury, N.: Handlers in action. In: *ACM SIGPLAN International Conference on Functional Programming, ICFP 2013*, pp. 145–158 (2013). <https://doi.org/10.1145/2500365.2500590>
14. Kammar, O., Pretnar, M.: No value restriction is needed for algebraic effects and handlers. *J. Funct. Program.* **27**, e7 (2017). <https://doi.org/10.1017/S0956796816000320>
15. Kiselyov, O., Ishii, H.: Freer monads, more extensible effects. In: *Proceedings of the 8th ACM SIGPLAN Symposium on Haskell, Haskell 2015*, pp. 94–105 (2015). <https://doi.org/10.1145/2804302.2804319>
16. Leijen, D.: Type directed compilation of row-typed algebraic effects. In: *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017*, pp. 486–499 (2017). <http://dl.acm.org/citation.cfm?id=3009872>
17. Leroy, X., Weis, P.: Polymorphic type inference and assignment. In: *Proceedings of the 18th Annual ACM Symposium on Principles of Programming Languages*, pp. 291–302 (1991). <https://doi.org/10.1145/99583.99622>
18. Lindley, S., McBride, C., McLaughlin, C.: Do be do be do. In: *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017*, pp. 500–514 (2017). <http://dl.acm.org/citation.cfm?id=3009897>
19. Materzok, M., Biernacki, D.: A dynamic interpretation of the CPS hierarchy. In: Jhala, R., Igarashi, A. (eds.) *APLAS 2012*. LNCS, vol. 7705, pp. 296–311. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-35182-2_21
20. Plotkin, G.D., Power, J.: Algebraic operations and generic effects. *Appl. Categ. Struct.* **11**(1), 69–94 (2003). <https://doi.org/10.1023/A:1023064908962>
21. Plotkin, G.D., Pretnar, M.: Handlers of algebraic effects. In: Castagna, G. (ed.) *ESOP 2009*. LNCS, vol. 5502, pp. 80–94. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-00590-9_7
22. Reynolds, J.C.: Types, abstraction and parametric polymorphism. In: *IFIP Congress*, pp. 513–523 (1983)
23. Tofte, M.: Type inference for polymorphic references. *Inf. Comput.* **89**(1), 1–34 (1990). [https://doi.org/10.1016/0890-5401\(90\)90018-D](https://doi.org/10.1016/0890-5401(90)90018-D)
24. Wright, A.K.: Simple imperative polymorphism. *Lisp Symb. Comput.* **8**(4), 343–355 (1995)
25. Wright, A.K., Felleisen, M.: A syntactic approach to type soundness. *Inf. Comput.* **115**(1), 38–94 (1994). <https://doi.org/10.1006/inco.1994.1093>
26. Wu, N., Schrijvers, T., Hinze, R.: Effect handlers in scope. In: *Proceedings of the 2014 ACM SIGPLAN Symposium on Haskell, Haskell 2014*, pp. 1–12 (2014). <https://doi.org/10.1145/2633357.2633358>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.