



Answer Refinement Modification: Refinement Type System for Algebraic Effects and Handlers

FUGA KAWAMATA, Waseda University, Japan

HIROSHI UNNO, University of Tsukuba, Japan

TARO SEKIYAMA, National Institute of Informatics, Japan

TACHIO TERAUCHI, Waseda University, Japan

Algebraic effects and handlers are a mechanism to structure programs with computational effects in a modular way. They are recently gaining popularity and being adopted in practical languages, such as OCaml. Meanwhile, there has been substantial progress in program verification via *refinement type systems*. While a variety of refinement type systems have been proposed, thus far there has not been a satisfactory refinement type system for algebraic effects and handlers. In this paper, we fill the void by proposing a novel refinement type system for languages with algebraic effects and handlers. The expressivity and usefulness of algebraic effects and handlers come from their ability to manipulate *delimited continuations*, but delimited continuations also complicate programs' control flow and make their verification harder. To address the complexity, we introduce a novel concept that we call *answer refinement modification* (ARM for short), which allows the refinement type system to precisely track what effects occur and in what order when a program is executed, and reflect such information as modifications to the refinements in the types of delimited continuations. We formalize our type system that supports ARM (as well as answer *type* modification, or ATM) and prove its soundness. Additionally, as a proof of concept, we have extended the refinement type system to a subset of OCaml 5 which comes with a built-in support for effect handlers, implemented a type checking and inference algorithm for the extension, and evaluated it on a number of benchmark programs that use algebraic effects and handlers. The evaluation demonstrates that ARM is conceptually simple and practically useful.

Finally, a natural alternative to directly reasoning about a program with delimited continuations is to apply a *continuation passing style* (CPS) transformation that transforms the program to a pure program without delimited continuations. We investigate this alternative in the paper, and show that the approach is indeed possible by proposing a novel CPS transformation for algebraic effects and handlers that enjoys bidirectional (refinement-)type-preservation. We show that there are pros and cons with this approach, namely, while one can use an existing refinement type checking and inference algorithm that can only (directly) handle pure programs, there are issues such as needing type annotations in source programs and making the inferred types less informative to a user.

CCS Concepts: • **Theory of computation** → **Type theory**; **Program verification**; **Control primitives**; • **Software and its engineering** → **Functional languages**; **Control structures**.

Additional Key Words and Phrases: algebraic effects and handlers, type-and-effect system, refinement type system, answer type modification, answer refinement modification, CPS transformation

Authors' addresses: [Fuga Kawamata](mailto:fuga.kawamata@waseda.ac.jp), Waseda University, Tokyo, Japan, maple-river@fuji.waseda.jp; [Hiroshi Unno](mailto:hiroshi.unno@tsukuba.ac.jp), University of Tsukuba, Tsukuba, Japan, uhiro@cs.tsukuba.ac.jp; [Taro Sekiyama](mailto:taro.sekiyama@nii.ac.jp), National Institute of Informatics, Tokyo, Japan, ryukilon@gmail.com; [Tachio Terauchi](mailto:tachio.terauchi@waseda.jp), Waseda University, Tokyo, Japan, terauchi@waseda.jp.



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

© 2024 Copyright held by the owner/author(s).

ACM 2475-1421/2024/1-ART5

<https://doi.org/10.1145/3633280>

ACM Reference Format:

Fuga Kawamata, Hiroshi Unno, Taro Sekiyama, and Tachio Terauchi. 2024. Answer Refinement Modification: Refinement Type System for Algebraic Effects and Handlers. *Proc. ACM Program. Lang.* 8, POPL, Article 5 (January 2024), 33 pages. <https://doi.org/10.1145/3633280>

1 INTRODUCTION

Algebraic effects [Plotkin and Power 2003] and handlers [Plotkin and Pretnar 2009, 2013] are a mechanism to structure programs with computational effects in a modular way. Algebraic effects represent abstracted computational effects and handlers specify their behaviors using delimited continuations. The ability to use delimited continuations makes algebraic effects and handlers highly expressive, allowing them to describe prominent computational effects such as exceptions, nondeterminism, mutable states, backtracking, and cooperative multithreading. Additionally, algebraic effects and handlers are recently gaining quite a recognition in practice and are adopted in popular programming languages, such as OCaml [Sivaramakrishnan et al. 2021].

Meanwhile, there has been substantial progress in program verification via *refinement type systems* [Bengtson et al. 2011; Freeman and Pfenning 1991; Nanjo et al. 2018; Rondon et al. 2008; Sekiyama and Unno 2023; Swamy et al. 2016; Terauchi 2010; Unno and Kobayashi 2009; Unno et al. 2018; Vazou et al. 2014; Vekris et al. 2016; Xi and Pfenning 1999; Zhu and Jagannathan 2013]. Such type systems allow the user to express a precise specification for a program as a type embedding logic formulas and their type checking (sometimes even type inference) (semi-)algorithms (semi-)automatically check whether the program conforms to the specification. While a variety of refinement type systems have been proposed for various classes of programming languages and features, including functional languages [Freeman and Pfenning 1991; Rondon et al. 2008; Vazou et al. 2014], object-oriented languages [Vekris et al. 2016], and delimited control operators [Sekiyama and Unno 2023], there has not been a satisfactory refinement type system for programming languages with algebraic effects and handlers.

In this work, we propose a new refinement type system for algebraic effects and handlers. A challenge with the precise verification in the presence of algebraic effects and handlers is the presence of the *delimited continuations*: they are the key ingredient of algebraic effects and handlers to implement a variety of computational effects, but they also complicate programs' control flow and make it difficult to statically discern what effects occur in what order. To address this challenge, we propose a novel concept that we call *answer refinement modification* (ARM for short), inspired by *answer type modification* (ATM) employed in certain type systems for delimited control operators such as `shift` and `reset` [Asai 2009; Danvy and Filinski 1990]. Similarly to ATM that can statically track how the use of delimited control operators influence the types of expressions, ARM can statically track how the use of algebraic effect operations (and the execution of the corresponding handlers) influence the refinements in the types of expressions, where the latter, as in prior refinement type systems, is used to precisely describe the *values*, rather than just their ordinary (i.e., non-refinement) types, computed by the expressions. Thus, our novel refinement type system supporting ARM can be used to precisely reason about programs with algebraic effects and handlers.

ATM and ARM are closely related: in fact, our refinement type system supports ATM, that is, our system allows the whole types and not just the refinements in them to be modified. As far as we know, the only prior (ordinary or refinement) type system for algebraic effects and handlers that supports ATM or ARM is a recent system of Cong and Asai [2022]. However, their system does not support refinement types (and so, obviously, no ARM), and moreover, even when compared as mechanisms for ordinary type systems, their ATM is less expressive than ours. We refer to Section 6 for detailed comparison.

While our system supports the full ATM, from the perspective of program verification, ARM alone, that is, allowing only modification in type refinements, is useful. Indeed, as in other refinement-type-based approaches, our aim is verification of programs typed in *ordinary* background type systems (such as the type systems of OCaml 5 and Koka [Leijen 2014] that do not support ATM), not to make more programs typable by extending the background type systems with ATM. As a proof of concept, we have extended the refinement type system and implemented a corresponding type checking and inference algorithm for a subset of OCaml 5 which comes with a built-in support for effect handlers, and evaluated it on a number of benchmark programs that use algebraic effects and handlers. The evaluation demonstrates that ARM is conceptually simple and practically useful.

Finally, a natural alternative to directly reasoning about a program with delimited continuations is to apply a continuation passing style (CPS) transformation that transforms the program to a pure program without delimited continuations. We investigate this alternative in the paper, and show that the approach is indeed possible by proposing a novel CPS transformation for algebraic effects and handlers that enjoys bidirectional (refinement-)type-preservation. Bidirectional type-preservation means that an expression is well-typed in the source language if and only if its CPS-transformed result is well-typed in the target language. This implies that we can use existing refinement type systems without support for effect handlers to verify programs with effect handlers by applying our CPS transformation. However, like other CPS transformations [Appel 1992; Cong and Asai 2018; Danvy and Filinski 1990; Hillerström et al. 2017; Plotkin 1975], ours makes global changes to the program and can radically change its structure, making it difficult for the programmer to recast the type checking and inference results back to the original program. Also, the CPS transformation is type directed and requires the program to be annotated by types conforming to our new type system, albeit only needing type “structures” without concrete refinement predicates. Moreover, in some cases, CPS-transformed expressions need extra parameters or higher-order predicate polymorphism to be typed as precisely as the source expressions, because the CPS transformation introduces higher-order continuation arguments. Nonetheless, our CPS transformation is novel, and we foresee that it would provide new interesting insights, as CPS transformations often do [Danvy and Filinski 1990], and be a useful tool for future studies on refinement type systems and effect handlers.

Our main contributions are summarized as follows.

- We show a sound refinement type system for algebraic effects and handlers, where ARM plays an important role.
- We have implemented the refinement type system for a subset of OCaml language with effect handlers, and evaluate it on a number of programs that use effect handlers.
- We define a bidirectionally-type-preserving CPS transformation which can be used to verify programs with effect handlers, and discuss pros and cons between direct type checking using our system and indirect type checking via the CPS transformation.

The rest of the paper is organized as follows. In Section 2, we briefly explain algebraic effects and handlers and ATM, and then describe the motivation for ARM and our system. Section 3 presents our language. We define its syntax, semantics and type system, present some typing examples, and show type safety of the language. Section 4 explains the implementation of the system. In Section 5, we provide the CPS transformation and discuss pros and cons between the direct type checking via our type system and the indirect type checking via CPS transformation. Finally, we describe related works in Section 6 and conclude the paper in Section 7.

2 OVERVIEW

We briefly overview algebraic effects and handlers, ATM, and ARM.

2.1 Algebraic Effects and Handlers

Algebraic effects and handlers enable users to define their own effects in a modular way. The modularity stems from separating the use of effects from their implementations: effects are performed via *operations* and implemented via *effect handlers* (or handlers for short). For example, consider the following program where $h_d \triangleq \{x_r \mapsto x_r, \text{decide}(x, k) \mapsto \max(k \text{ true}) (k \text{ false})\}$:

with h_d **handle** **let** $a = \text{if } \text{decide } () \text{ then } 10 \text{ else } 20 \text{ in } \text{let } b = \text{if } \text{decide } () \text{ then } 1 \text{ else } 2 \text{ in } a - b$

It calls an operation `decide`, which takes the unit value `()` and returns a Boolean value, to choose one of two integer values and then calculates the difference between the chosen values. Because operation calls invoke effects in algebraic effects, the operations work as interfaces of the effects.

An implementation of an effect is given by an effect handler. The program installs the handler h_d for `decide` using the handling construct. In general, a handling construct takes the form **with** h **handle** e , which means that a handler h defines interpretations of operations performed during the evaluation of the expression e ; we call the expression e a *handled expression*. A handler consists of a single *return clause* and zero or more *operation clauses*. A return clause takes the form $x_r \mapsto e_r$, which determines the value of the handling construct by evaluating expression e_r with variable x_r that denotes the value of the handled expression. In the example, because the return clause is $x_r \mapsto x_r$, the handling construct simply returns the value of the handled expression. An operation clause takes the form $\text{op}(x, k) \mapsto e$. It defines the interpretation of the operation `op` to be expression e with variable x that denotes the arguments to the operation. When the handled expression calls the operation `op`, the remaining computation up to the handling construct is suspended and instead the body e of the operation clause evaluates. Therefore, effect handlers behave like exception handlers by regarding operation calls as raising exceptions. However, effect handlers are equipped with the additional ability to resume the suspended computation. The suspended remaining computation, called a *delimited continuation*, is functionalized, and the body e of the operation clause can refer to it via the variable k .

Let us take a closer look at the behavior of the above example. Because the handled expression starts with the call to `decide`, the operation clause for `decide` given by h_d evaluates. The delimited continuation K of the first call to `decide` is

with h_d **handle** (**let** $a = \text{if } [] \text{ then } 10 \text{ else } 20 \text{ in } \text{let } b = \text{if } \text{decide } () \text{ then } 1 \text{ else } 2 \text{ in } a - b$)

where $[]$ denotes the hole of the continuation. We write $K[e]$ for the expression obtained by filling the hole in K with expression e . Then, the functional representation of the delimited continuation K takes the form $\lambda y.K[y]$, and it is substituted for k in the body of the operation clause. Namely, the handling construct evaluates to $\max(v \text{ true}) (v \text{ false})$ where $v = \lambda y.K[y]$. Note that the variable x of the operation clause for `decide` is replaced by the unit value `()`, but it is not referenced. The first argument $v \text{ true}$ to `max` reduces to $K[\text{true}]$, that is,

with h_d **handle** (**let** $a = \text{if } \text{true} \text{ then } 10 \text{ else } 20 \text{ in } \text{let } b = \text{if } \text{decide } () \text{ then } 1 \text{ else } 2 \text{ in } a - b$)

(the grayed part represents the value by which the hole in K is replaced). Therefore, the expression $v \text{ true}$ evaluates to **with** h_d **handle** (**let** $b = \text{if } \text{decide } () \text{ then } 1 \text{ else } 2 \text{ in } 10 - b$). Again, `decide` is called and the continuation $K' \triangleq \text{with } h_d \text{ handle } (\text{let } b = \text{if } [] \text{ then } 1 \text{ else } 2 \text{ in } 10 - b)$ is captured. Then, the operation clause for `decide` evaluates after substituting $\lambda y.K'[y]$ for k . The expression $(\lambda y.K'[y]) \text{ true}$ evaluates to $K'[\text{true}]$, that is, **with** h_d **handle** (**let** $b = \text{if } \text{true} \text{ then } 1 \text{ else } 2 \text{ in } 10 - b$) and then to **with** h_d **handle** 9. Here, the handled expression is a value. Therefore, the return clause in the handler evaluates after substituting the value 9 for variable x_r . Because the return clause in h_d just returns x_r , the evaluation of $(\lambda y.K'[y]) \text{ true}$ results in 9. Similarly, $(\lambda y.K'[y]) \text{ false}$ evaluates to 8 (which is the result of binding b to 2). Therefore, $\max((\lambda y.K'[y]) \text{ true}) ((\lambda y.K'[y]) \text{ false})$

evaluates to $\max 9\ 8$ and then to 9. In a similar way, v **false** calculates $\max (20 - 1) (20 - 2)$, that is, evaluates to 19, because a is bound to 20 and b is bound to each of 1 and 2 depending on the result of the second invocation of **decide**. Finally, the entire program evaluates to 19, which is the result of $\max (v$ **true**) $(v$ **false**), that is, $\max 9\ 19$.

2.2 Answer Type Modification and Answer Refinement Modification

An *answer type* is the type of the closest enclosing delimiter, or the return type of a delimited continuation. In the setting of algebraic effects and handlers, delimiters are handling constructs. For example, consider the following expression:

$$\text{let } x = \text{with } \{x_r \mapsto x_r, \text{op}(\cdot, k) \mapsto k \mid 0 < k < 1\} \text{ handle } 1 + \text{op}(\cdot) \text{ in } c.$$

The delimited continuation of $\text{op}(\cdot)$ is $K'' \triangleq \text{with } \{x_r \mapsto x_r, \text{op}(\cdot, k) \mapsto k \mid 0 < k < 1\} \text{ handle } 1 + [\cdot]$. At first glance, the answer type of $\text{op}(\cdot)$ seems to be the integer type `int` since the handled computation in the continuation returns the integer $1 + n$ for an integer n given to fill the hole, and the return clause returns given values as they are. In other words, from the perspective of $\text{op}(\cdot)$, the handling construct seems to give an integer value to the outer context $\text{let } x = [\cdot] \text{ in } c$. However, after the operation call, the entire expression evaluates to $\text{let } x = v'' \mid 0 < v'' < 1 \text{ in } c$ where $v'' \triangleq \lambda y. K''[y]$. Now the handling construct becomes the expression $v'' \mid 0 < v'' < 1$, which gives a Boolean value to the outer context. That is, the answer type changes to the Boolean type `bool`. *Answer type modification* (ATM) is a mechanism to track this dynamic change on answer types.

ATM is not supported in existing type systems for effect handlers [Bauer and Pretnar 2013, 2015; Brady 2013; Kammar et al. 2013; Leijen 2017; Lindley et al. 2017; Plotkin and Pretnar 2013], with the exception of the one recently proposed by Cong and Asai [2022] (see Section 6 for comparison with their work). Such type systems require the answer types before and after an operation call to be unified (and so the example above will be rejected as ill-typed). Nonetheless, useful programming with effect handlers is still possible without ATM (which is why they are implemented in popular languages like OCaml without ATM).¹ For instance, the program in Section 2.1 is well-typed in existing (non-refinement) type systems for algebraic effects and handlers without ATM, since the return type of the continuation k in the **decide** clause (i.e., the answer type before the execution) is `int` and the return type of the **decide** clause (i.e., the answer type after the execution) is also `int`.

However, even if answer types are not modified, *actual values returned by delimited continuations usually change*. Let us see the program in Section 2.1 again. Focus on the first call to **decide**. When this is called, the operation clause receives the continuation $v = \lambda y. K[y]$, which returns 9 if applied to **true** and returns 19 if applied to **false**, as described previously. Therefore, v can be assigned the refinement type $(y : \text{bool}) \rightarrow \{z : \text{int} \mid z = (y ? 9 : 19)\}$, and thus the precise answer type before the execution is $\{z : \text{int} \mid z = (y ? 9 : 19)\}$ where y is the Boolean value passed to the continuation. On the other hand, the clause for **decide** returns integer 19. Thus, the precise answer type after the operation call is $\{z : \text{int} \mid z = 19\}$. Now the refinement in the answer type becomes different before and after the operation call. The same phenomenon happens in the second call to **decide**. When the second call evaluates, the handler receives the continuation $\lambda y. K'[y]$. It returns $a - 1$ if applied to **true** and returns $a - 2$ if applied to **false** (where a is either 10 or 20 depending on the result of the first call to **decide**). Thus, the answer type before the execution is $\{z : \text{int} \mid z = (y ? a - 1 : a - 2)\}$. In contrast, the return value of the clause for **decide** is $\max (a - 1) (a - 2) = a - 1$, so the answer type after the execution is $\{z : \text{int} \mid z = a - 1\}$. Here again, the refinement in the answer type changed by the operation call. We call this change *answer refinement modification* (ARM). Armed with ARM

¹One could also argue that the absence of ATM is natural for algebraic effects and handlers because they are designed after concepts from universal algebra [Bauer 2018; Plotkin and Power 2001], and there, (algebraic) operations are usually expected to preserve types.

(pun intended), the refinement type system that we propose in this paper is able to assign the precise refinement type $\{z : \text{int} \mid z = 19\}$ to the program, and more generally, the type $\{z : \text{int} \mid z = v - x\}$ when the constants 10, 20, 1, and 2 are replaced by variables u , v , x , and y respectively with the assumption $u \leq v \wedge x \leq y$ (such an assumption on free variables can be given by refinement types in the top-level type environment). The example demonstrates that ARM is useful for precisely reasoning about programs with algebraic effects and handlers in refinement type systems. Indeed, without ARM, the most precise refinement type that a type system could assign to the example would be $\{z : \text{int} \mid z \in \{8, 9, 18, 19\}\}$.

As another illuminating example, we show that ARM provides a new approach to the classic *strong update* problem [Foster et al. 2002]. It is well known that algebraic effects and handlers can implement mutable references by operations `set` and `get`, that respectively destructively updates and reads a mutable reference, and a handler that implements the operations by state-passing (see, e.g., [Pretnar 2015]). On programs with such a standard implementation of mutable references by algebraic effects and handlers, our refinement type system is able to reason flow-sensitively and derive refinement types that cannot be obtained with ordinary flow-insensitive reasoning. For instance, consider the following program where $h \triangleq \{x_r \mapsto \lambda s.x_r, \text{set}(x, k) \mapsto \lambda s.k \ () \ x, \text{get}(x, k) \mapsto \lambda s.k \ s \ s\}$:

(with h handle (set 3; let $n = \text{get} \ ()$ in set 5; let $m = \text{get} \ ()$ in $n + m$)) 0

Thanks to ARM, our type system can give the program the most precise type $\{z : \text{int} \mid z = 8\}$, which would not be possible in a type system without ARM as it would conflate the two calls to `set` and fail to reason that the first `get` () returns 3 whereas the second `get` () returns 5. Roughly, ARM accomplishes the flow-sensitive reasoning about the changes in the state by tracking changes in the refinements in the answer types, albeit in a *backward* fashion as shown in Section 3.3.

Using this ability of ARM, we can also verify that effectful operations are used in a specific order. For example, consider operations `open`, `close`, `read`, and `write` for file manipulation being implemented using effect handlers. The use of these operations should conform to the regular scheme $(\text{open} (\text{read} \mid \text{write})^* \text{close})^*$. Our refinement type system can check if a program meets this requirement. For instance, consider the following recursive function:

$\lambda x.$ while (\star) {open x ; while (\star) {let $y = \text{read} \ ()$ in write ($y \wedge "X"$)}; close ()}

where `while` (\star) { c } loops computation c and terminates nondeterministically, and the binary operation (\wedge) concatenates given strings (operation `read` is supposed to return a string).² The function repeats opening the specified file x and closing it after reading from and writing to the file zero or more times. Thus, this function follows the discipline of the file manipulation operations. We will show in Section 3.3 how ARM enables us to check it formally and detect the invalid use of the operations if any.

3 LANGUAGE

This section presents our language with algebraic effects and handlers. The semantics is formalized using evaluation contexts like in Leijen [2017], and the type system is a novel refinement type system with ARM (and ATM).

3.1 Syntax and Semantics

The upper half of Figure 1 shows the syntax of our language. It indicates that expressions are split into values, ranged over by v , and computations, ranged over by c , as in the fine-grain call-by-value style of Levy et al. [2003]. Values, which are effect-free expressions in a canonical form, consist

²For simplicity, we assume that the clause of `open` creates an object for a specified file and stores it in a reference implemented by an effect handler, and the clauses of the other operations refer to the stored object to manipulate the file.

Syntax

$$\begin{aligned}
 p &::= \mathbf{true} \mid \mathbf{false} \mid \dots \quad v ::= x \mid p \mid \mathbf{rec}(f, x).c \quad K ::= [] \mid \mathbf{let} \ x = K \ \mathbf{in} \ c \\
 c &::= \mathbf{return} \ v \mid \mathbf{op} \ v \mid v_1 \ v_2 \mid \mathbf{if} \ v \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2 \mid \mathbf{let} \ x = c_1 \ \mathbf{in} \ c_2 \mid \mathbf{with} \ h \ \mathbf{handle} \ c \\
 h &::= \{\mathbf{return} \ x_r \mapsto c_r, (\mathbf{op}_i(x_i, k_i) \mapsto c_i)_i\}
 \end{aligned}$$

Evaluation rules

$$\begin{array}{c}
 \boxed{c \longrightarrow c'} \\
 \\
 \frac{c_1 \longrightarrow c'_1}{\mathbf{let} \ x = c_1 \ \mathbf{in} \ c_2 \longrightarrow \mathbf{let} \ x = c'_1 \ \mathbf{in} \ c_2} \text{(E-LET)} \quad \frac{}{\mathbf{let} \ x = \mathbf{return} \ v \ \mathbf{in} \ c_2 \longrightarrow c_2[v/x]} \text{(E-LETRET)} \\
 \\
 \frac{}{\mathbf{if} \ \mathbf{true} \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2 \longrightarrow c_1} \text{(E-IFT)} \quad \frac{}{\mathbf{if} \ \mathbf{false} \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2 \longrightarrow c_2} \text{(E-IFF)} \\
 \\
 \frac{}{(\mathbf{rec}(f, x).c) \ v \longrightarrow c[v/x][(\mathbf{rec}(f, x).c)/f]} \text{(E-APP)} \quad \frac{}{p \ v \longrightarrow \zeta(p, v)} \text{(E-PRIM)} \\
 \\
 \text{below, let } h = \{\mathbf{return} \ x_r \mapsto c_r, (\mathbf{op}_i(x_i, k_i) \mapsto c_i)_i\} \\
 \\
 \frac{c \longrightarrow c'}{\mathbf{with} \ h \ \mathbf{handle} \ c \longrightarrow \mathbf{with} \ h \ \mathbf{handle} \ c'} \text{(E-HNDL)} \\
 \\
 \frac{}{\mathbf{with} \ h \ \mathbf{handle} \ \mathbf{return} \ v \longrightarrow c_r[v/x_r]} \text{(E-HNDLRET)} \\
 \\
 \frac{}{\mathbf{with} \ h \ \mathbf{handle} \ K[\mathbf{op}_i \ v] \longrightarrow c_i[v/x_i][(\lambda y. \mathbf{with} \ h \ \mathbf{handle} \ K[\mathbf{return} \ y])/k_i]} \text{(E-HNDLOP)}
 \end{array}$$

Fig. 1. Syntax and evaluation rules.

of variables x , primitive values p , and (recursive) functions $\mathbf{rec}(f, x).c$ where variable f denotes the function itself for recursive calls in the body c . If f does not occur in c , we simply write $\lambda x.c$. Computations, which are possibly effectful expressions, consist of six kinds of constructs. A value-return $\mathbf{return} \ v$ lifts a value v to a computation. An operation call $\mathbf{op} \ v$ performs the operation \mathbf{op} with the argument v . A function application $v_1 \ v_2$, conditional branch $\mathbf{if} \ v \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2$, and let-expression $\mathbf{let} \ x = c_1 \ \mathbf{in} \ c_2$ are standard. Note that functions, arguments, and conditional expressions are restricted to values, but this does not reduce expressivity because, e.g., a conditional branch $\mathbf{if} \ c \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2$ can be expressed as $\mathbf{let} \ x = c \ \mathbf{in} \ \mathbf{if} \ x \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2$ using a fresh variable x . A handling construct $\mathbf{with} \ h \ \mathbf{handle} \ c$ handles operations performed during the evaluation of the handled computation c using the clauses in the handler h . A handler $\{\mathbf{return} \ x_r \mapsto c_r, (\mathbf{op}_i(x_i, k_i) \mapsto c_i)_i\}$ has a return clause $\mathbf{return} \ x_r \mapsto c_r$ where the variable x_r denotes the value of the handled computation c , and an operation clause $\mathbf{op}_i(x_i, k_i) \mapsto c_i$ for each operation \mathbf{op}_i where the variables x_i and k_i denote the argument to \mathbf{op}_i and the continuation from the invocation of \mathbf{op}_i , respectively. The notions of free variables and substitution are defined as usual. We write $c[v/x]$ for the computation obtained by substituting the value v for the variable x in the computation c . We use similar notation to substitute values for variables in types and substitute types for type variables.

The semantics of the language is defined by the evaluation relation \longrightarrow , which is the smallest binary relation over computations satisfying the evaluation rules in the lower half of Figure 1. The evaluation of a let-expression $\mathbf{let} \ x = c_1 \ \mathbf{in} \ c_2$ begins by evaluating the computation c_1 . When c_1 returns a value, the computation c_2 evaluates after substituting the return value for x . The evaluation rules for conditional branching and function application are standard. The result of applying a primitive value relies on the metafunction ζ , which maps pairs of a primitive value

and an argument value to computations. For a handling construct **with** h **handle** c , the handled computation c evaluates first. When c returns a value, the body of the return clause in the handler h evaluates with the return value. If the evaluation of c encounters an operation call $\text{op}_i v$, its delimited continuation, which is represented as a pure evaluation context K defined in Figure 1, is captured. Then, the body c_i of the operation clause $\text{op}_i(x_i, k_i) \mapsto c_i$ for op_i in the handler h evaluates after substituting the argument v and the function $\lambda y. \text{with } h \text{ handle } K[\text{return } y]$ for variables x_i and k_i , respectively. Note that the function substituted for k_i wraps the delimited continuation $K[\text{return } y]$ by the handling construct with the handler h . It means that the operation calls in $K[\text{return } y]$ are handled by the handler h . Our semantics assumes that the handler h provides operation clauses for all the operations performed by the handled computation c . Our type system ensures that this assumption holds on any well-typed computations. However, our language can also implement the forwarding semantics by encoding: given a handler that does not contain an operation clause for op , we add to the handler an operation clause $\text{op}(x, k) \mapsto \text{let } y = \text{op } x \text{ in } k y$.³

3.2 Type System

Figure 2 shows the syntax of types. As in prior refinement type systems [Bengtson et al. 2011; Rondon et al. 2008; Unno and Kobayashi 2009], our type system allows a type specification for values of base types, ranged over by B , such as `bool` and `int`, to be refined using logic formulas, ranged over by ϕ . Our type system is parameterized over a logic. We assume that the logic is a predicate logic where: terms, denoted by t , include variables x ; predicates, denoted by A , include predicate variables X ; and each primitive value p can be represented as a term. Throughout the paper, we use the over-tilde notation to denote a sequence of entities. For example, \tilde{t} represents a sequence t_1, \dots, t_n of some terms t_1, \dots, t_n , and then $A(\tilde{t})$ represents a formula $A(t_1, \dots, t_n)$. We also assume that base types include at least the Boolean type `bool`.

Types consist of value and computation types, which are assigned to values and computations, respectively. A value type, denoted by T , is either a refinement type $\{x : B \mid \phi\}$, which is assigned to a value v of base type B such that the formula $\phi[v/x]$ is true, or a dependent function type $(x : T) \rightarrow C$, which is assigned to a function that, given an argument v of the type T , performs the computation specified by the type $C[v/x]$. We abbreviate $(x : T) \rightarrow C$ as $T \rightarrow C$ if x does not occur in C , and $\{z : B \mid \text{true}\}$ as B .

A computation type is formed by three components: an operation signature, which specifies operations that a computation may perform; a value type, which specifies the value that the computation returns if any; and a control effect, which specifies how the computation modifies the answer type via operation call.

Control effects, denoted by S , are inspired by the formalism of Sekiyama and Unno [2023] who extended control effects in simple typing [Materzok and Biernacki 2011] to dependent typing. A control effect is either pure or impure. The pure control effect \square means that a computation calls

Control effects, denoted by S , are inspired by the formalism of Sekiyama and Unno [2023] who extended control effects in simple typing [Materzok and Biernacki 2011] to dependent typing. A control effect is either pure or impure. The pure control effect \square means that a computation calls

term	$t ::= x \mid \dots$	formula	$\phi ::= A(\tilde{t}) \mid \dots$
predicate	$A ::= X \mid \dots$	base type	$B ::= \text{bool} \mid \dots$
value type	$T ::= \{x : B \mid \phi\} \mid (x : T) \rightarrow C$		
computation type	$C ::= \Sigma \triangleright T / S$		
operation signature	$\Sigma ::= \{(\text{op}_i : \forall X_i : \tilde{B}_i. F_i)_i\}$		
	$F ::= (x : T_1) \rightarrow ((y : T_2) \rightarrow C_1) \rightarrow C_2$		
control effect	$S ::= \square \mid (\forall x. C_1) \Rightarrow C_2$		
typing context	$\Gamma ::= \emptyset \mid \Gamma, x : T \mid \Gamma, X : \tilde{B}$		

Fig. 2. Type syntax.

³We employ the semantics without forwarding in the body of the paper to simplify the typing rule for handling constructs. The supplementary material shows an extended typing rule for handling constructs that natively supports forwarding.

no operation. An impure control effect is given to a computation that may perform operations, specifying how the execution of the computation modifies its answer type. Impure control effects take the form $(\forall x.C_1) \Rightarrow C_2$ where variable x is bound in computation type C_1 . We write $C_1 \Rightarrow C_2$ when x does not occur in C_1 . In what follows, we first illustrate impure control effects in the simple, nondependent form $C_1 \Rightarrow C_2$ and then extend to the fully dependent form $(\forall x.C_1) \Rightarrow C_2$ that can specify the behavior of captured continuations using the input (denoted by x) to the continuations.

A control effect $C_1 \Rightarrow C_2$ represents the answer type of a program changes from type C_1 to type C_2 . When it is assigned to a computation c , the initial answer type C_1 specifies how the continuation of the computation c up to the closest handling construct behaves, and the final answer type C_2 specifies what can be guaranteed for the *meta-context*, i.e., the context of the closest handling construct. To see the idea more concretely, revisit the first example in Section 2.2:

$$\mathbf{let } x = \mathbf{with } \{x_r \mapsto x_r, \text{op}(\cdot, k) \mapsto k \ 0 < k \ 1\} \mathbf{handle } 1 + \text{op } (\cdot) \mathbf{in } c .$$

Let h be the handler in the example. Focusing on the operation call $\text{op } (\cdot)$, we can find that it captures the continuation $\mathbf{with } h \mathbf{handle } 1 + [\cdot]$. Because the continuation behaves as if it is a pure function returning integers, the initial answer type of $\text{op } (\cdot)$ turns out to be the computation type int / \square (we omit Σ for a while; it will be explained shortly). Furthermore, by the operation call, the handling construct $\mathbf{with } h \mathbf{handle } 1 + \text{op } (\cdot)$ is replaced with the body $k \ 0 < k \ 1$ of op 's clause in h and the functional representation v of the continuation is substituted for k . It means that the meta-context $\mathbf{let } x = [\cdot] \mathbf{in } c$ of the operation call takes the computation $v \ 0 < v \ 1$, which is of type bool / \square (note that $v \ 0 < v \ 1$ is pure because v is a pure function). Therefore, the final answer type of $\text{op } (\cdot)$ is bool / \square . As a result, the impure control effect of $\text{op } (\cdot)$ is $\text{int} / \square \Rightarrow \text{bool} / \square$.

Sekiyama and Unno [2023] extended the simple form of impure control effects to a dependent form $(\forall x.C_1) \Rightarrow C_2$, where the initial answer type C_1 can depend on inputs, denoted by variable x , to continuations. For instance, consider the continuation $\mathbf{with } h \mathbf{handle } 1 + [\cdot]$ captured in the above example. When passed an integer n , it returns $1 + n$. Using the dependent form of impure control effects, we can describe such behavior by the control effect $(\forall x.\{y : \text{int} \mid y = x + 1\} / \square) \Rightarrow \text{bool} / \square$, where x represents the input to the continuation and the refinement type $\{y : \text{int} \mid y = x + 1\}$ precisely specifies the return value of the continuation for input x . The type of x is matched with the continuation's input type. Since the continuation of $\text{op } (\cdot)$ takes integers, the type assigned to x is int . In general, given a computation type $T / (\forall x.C_1) \Rightarrow C_2$, the type T is assigned to the variable x because it corresponds to the input type of the continuations of computations given that computation type. The type information refined by dependent impure control effects is exploited in typechecking operation clauses. In the example, our type system typechecks the body of op 's clause by assigning the function type $(x : \text{int}) \rightarrow \{y : \text{int} \mid y = x + 1\} / \square$ to the continuation variable k . Then, since the body is $k \ 0 < k \ 1$, its type—i.e., the final answer type—can be refined to $\{z : \text{bool} \mid z = \mathbf{true}\} / \square$. Hence, the type system can assign control effect $(\forall x.\{y : \text{int} \mid y = x + 1\} / \square) \Rightarrow \{z : \text{bool} \mid z = \mathbf{true}\} / \square$ to the operation call and ensure that the meta-context takes \mathbf{true} finally (if the handling construct terminates). We will demonstrate the expressivity and usefulness of dependent control effects in more detail in Section 3.3.

Operation signatures, denoted by Σ , are sets of pairs of an operation name and a type scheme. We write $(\cdot)_i$ to denote a sequence of entities indexed by i . The type scheme associated with an operation op is in the form $\forall X : \widetilde{B}.(x : T_1) \rightarrow ((y : T_2) \rightarrow C_1) \rightarrow C_2$, where the types T_1 and T_2 are the input and output types, respectively, of the operation and the types C_1 and C_2 are the initial and final answer types, respectively, of the operation call for op . Recall that the initial answer type C_1 corresponds to the return type of delimited continuations captured by the call to op , and that the continuations take the return values of the operation call. Therefore, the function type

$(y : T_2) \rightarrow C_1$ represents the type of the captured delimited continuations. Note that the variable y denotes values passed to the continuations. Furthermore, the final answer type C_2 corresponds to the type of the operation clause for `op` in the closest enclosing handler. Therefore, the operation clause `op(x, k) ↦ c` in the handler is typed by checking that the body c is of the type C_2 with the assumption that argument variable x is of the type T_1 and the continuation variable k is of the type $(y : T_2) \rightarrow C_1$. A notable point of the type scheme is that it can be parameterized over predicates. The predicate variables \tilde{X} abstract over the predicates, and the annotations \tilde{B} represent the (base) types of the arguments to the predicates. This allows calls to the same operation in different contexts to have different control effects, which is crucial for precisely verifying programs with algebraic effects and handlers as we will show in Section 3.3. It is also noteworthy that operation signatures include not only operation names but also type schemes as in Kammar et al. [2013] and Kammar and Pretnar [2017]. It allows an operation to have different types depending on the contexts where it is used. Another approach is to include only operation names and assumes that unique types are assigned to them globally as in, e.g., Bauer and Pretnar [2013] and Leijen [2017]. We decided to assign types to operations locally because it makes the type system more flexible in that the types of operations can be refined depending on contexts if needed.

Typing contexts Γ are lists of variable bindings $x : T$ and predicate variable bindings $X : \tilde{B}$. We write Γ, ϕ for $\Gamma, x : \{z : B \mid \phi\}$ where x and z are fresh. The notions of free variables, free predicate variables, and predicate substitution are defined as usual.

Well-formedness of typing contexts, value types, and computation types, whose judgments are in the forms $\vdash \Gamma$, $\Gamma \vdash T$, and $\Gamma \vdash C$, respectively, are defined straightforwardly by following Sekiyama and Unno [2023]. We refer to the supplementary material for detail.

Typing judgements for values and computations are in the forms $\Gamma \vdash v : T$ and $\Gamma \vdash c : C$, respectively. Figure 3 shows the typing rules. By (T-CVAR), a variable x of a refinement type is assigned a type which states that the value of this type is exactly x . For a variable of a non-refinement type (i.e., a function type in our language), the rule (T-VAR) assigns the type associated with the variable in the typing context. The rule (T-PRIM) uses the mapping ty to type primitive values p . We assume that ty assigns an appropriate value type to every primitive value. We refer to the supplementary material for the formalization of the assumption. The rule (T-FUN) for functions, (T-APP) for function applications, and (T-IF) for conditional branches are standard in refinement type systems (with support for value-dependent refinements). The rules (T-VSUB) and (T-CSUB) allow values and computations, respectively, to be typed at supertypes of their types. We will define subtyping shortly. By (T-RET), a value-return `return v` has a computation type where the operation signature is empty, the return value type is the type of v , and the control effect is pure.

To type a let-expression `let x = c1 in c2`, either the rule (T-LETP) or (T-LETIP) is used. Both of them require that the types of the sub-expressions c_1 and c_2 have the same operation signature Σ and then assign Σ to the type of the entire let-expression. The typing context for c_2 is extended by $x : T_1$ with the value type T_1 of c_1 , but x cannot occur in Σ and T_2 (as well as C_{21} in (T-LETIP)) to prevent the leakage of x from its scope. On the other hand, the two rules differ in how they treat control effects. When both of the control effects of c_1 and c_2 are pure, the rule (T-LETP) is used. It states that the control effect of the entire let-expression is also pure. When both are impure, the rule (T-LETIP) is used. It states that the control effect of the let-expression results in an impure control effect that is composed of the control effects of c_1 and c_2 . Note that, even when one of the control effects of c_1 and c_2 is pure and the other is impure, we can view both of them as impure effects via subtyping because it allows converting a pure control effect to an impure control effect, as shown later. We first explain how the composition works in the non-dependent form. Let the control effect of c_1 be $C_{11} \Rightarrow C_{12}$ and that of c_2 be $C_{21} \Rightarrow C_{22}$, and assume that a control effect

Typing rules for values

$$\boxed{\Gamma \vdash v : T}$$

$$\frac{\vdash \Gamma \quad \Gamma(x) = \{z : B \mid \phi\}}{\Gamma \vdash x : \{z : B \mid z = x\}} \text{(T-CVAR)} \quad \frac{\vdash \Gamma \quad \forall y, B, \phi. \Gamma(x) \neq \{y : B \mid \phi\}}{\Gamma \vdash x : \Gamma(x)} \text{(T-VAR)} \quad \frac{\vdash \Gamma}{\Gamma \vdash p : ty(p)} \text{(T-PRIM)}$$

$$\frac{\Gamma, f : (x : T) \rightarrow C, x : T \vdash c : C}{\Gamma \vdash \text{rec}(f, x).c : (x : T) \rightarrow C} \text{(T-FUN)} \quad \frac{\Gamma \vdash v : T_1 \quad \Gamma \vdash T_1 <: T_2 \quad \Gamma \vdash T_2}{\Gamma \vdash v : T_2} \text{(T-VSUB)}$$

Typing rules for computations

$$\boxed{\Gamma \vdash c : C}$$

$$\frac{\Gamma \vdash v : T}{\Gamma \vdash \text{return } v : \emptyset \triangleright T / \square} \text{(T-RET)} \quad \frac{\Gamma \vdash v_1 : (x : T) \rightarrow C \quad \Gamma \vdash v_2 : T}{\Gamma \vdash v_1 v_2 : C[v_2/x]} \text{(T-APP)}$$

$$\frac{\Gamma, v = \text{true} \vdash c_1 : C \quad \Gamma, v = \text{false} \vdash c_2 : C}{\Gamma \vdash \text{if } v \text{ then } c_1 \text{ else } c_2 : C} \text{(T-IF)} \quad \frac{\Gamma \vdash c : C_1 \quad \Gamma \vdash C_1 <: C_2 \quad \Gamma \vdash C_2}{\Gamma \vdash c : C_2} \text{(T-CSUB)}$$

$$\frac{\Gamma \vdash c_1 : \Sigma \triangleright T_1 / \square \quad \Gamma, x : T_1 \vdash c_2 : \Sigma \triangleright T_2 / \square \quad x \notin fv(T_2) \cup fv(\Sigma)}{\Gamma \vdash \text{let } x = c_1 \text{ in } c_2 : \Sigma \triangleright T_2 / \square} \text{(T-LETP)} \quad \frac{\Gamma \vdash c_1 : \Sigma \triangleright T_1 / (\forall x. C) \Rightarrow C_{12} \quad \Gamma, x : T_1 \vdash c_2 : \Sigma \triangleright T_2 / (\forall y. C_{21}) \Rightarrow C \quad x \notin fv(T_2) \cup fv(\Sigma) \cup (fv(C_{21}) \setminus \{y\})}{\Gamma \vdash \text{let } x = c_1 \text{ in } c_2 : \Sigma \triangleright T_2 / (\forall y. C_{21}) \Rightarrow C_{12}} \text{(T-LETIP)}$$

$$\frac{\Sigma \ni \text{op} : \forall X : \widetilde{B}. (x : T_1) \rightarrow ((y : T_2) \rightarrow C_1) \rightarrow C_2 \quad \Gamma \vdash \Sigma \quad \Gamma \vdash A : \widetilde{B} \quad \Gamma \vdash v : T_1[A/X]}{\Gamma \vdash \text{op } v : \Sigma \triangleright T_2[\widetilde{A/X}][v/x] / ((\forall y. C_1) \Rightarrow C_2)[\widetilde{A/X}][v/x]} \text{(T-OP)}$$

$$\frac{h = \{\text{return } x_r \mapsto c_r, (\text{op}_i(x_i, k_i) \mapsto c_i)_i\} \quad \Gamma \vdash c : \Sigma \triangleright T / (\forall x_r. C_1) \Rightarrow C_2 \quad \Gamma, x_r : T \vdash c_r : C_1 \quad \left(\Gamma, X_i : \widetilde{B}_i, x_i : T_{1i}, k_i : (y_i : T_{2i}) \rightarrow C_{1i} \vdash c_i : C_{2i} \right)_i \quad \Sigma = \{(\text{op}_i : \forall X_i : \widetilde{B}_i. (x_i : T_{1i}) \rightarrow ((y_i : T_{2i}) \rightarrow C_{1i}) \rightarrow C_{2i})_i\}}{\Gamma \vdash \text{with } h \text{ handle } c : C_2} \text{(T-HNDL)}$$

Fig. 3. Typing rules.

$C_1 \Rightarrow C_2$ is assigned to the let-expression. First, recall that the type C_1 expresses the return type of the continuation of the let-expression up to the closest handling construct and that the closest handling construct is replaced by a computation of the type C_2 . Based on this idea, the types C_1 and C_2 can be determined as follows. First, because the delimited continuation of the let-expression is matched with that of the computation c_2 , the initial answer type C_{21} of c_2 expresses the return type of the delimited continuation of the let-expression. Therefore, the type C_1 should be matched with the type C_{21} . Second, because the closest handling construct enclosing the let-expression is the same as the one enclosing the sub-computation c_1 , the type C_2 should be matched with the final answer type C_{12} of c_1 . Therefore, the control effect $C_1 \Rightarrow C_2$ should be matched with $C_{21} \Rightarrow C_{12}$, as stated in (T-LETIP). Furthermore, the rule (T-LETIP) requires that the initial answer type C_{11} of c_1 to be the same as the final answer type C_{22} of c_2 . This requirement is explained as follows. First, the computation c_1 expects its delimited continuation to behave as specified by the type C_{11} . The delimited continuation of c_1 first evaluates the succeeding computation c_2 . The final answer type C_{22} of c_2 expresses that the closest handling construct enclosing c_2 behaves as specified by the

type C_{22} . Because the closest handling construct enclosing c_2 corresponds to the top-level handling construct in the delimited continuation of c_1 , the type C_{11} should be matched with the type C_{22} . We now extend to the fully dependent form. From the discussion thus far, we can let the control effects of c_1 , c_2 , and the let-expression be $(\forall x_1.C) \Rightarrow C_{12}$, $(\forall x_2.C_{21}) \Rightarrow C$, and $(\forall y.C_{21}) \Rightarrow C_{12}$ respectively, for some variables x_1 , x_2 , and y . Then, the constraints on the names of these variables are determined as follows. First, the input to the delimited continuation of c_1 , which is denoted by the variable x_1 , should be matched with the evaluation result of c_1 . Then, since the let-expression binds the variable x to the evaluation result of c_1 , the variable x_1 is matched with x . Second, because the delimited continuation of c_2 is matched with that of the let-expression, the inputs to them should be matched with each other. They are denoted by the variables x_2 and y respectively, and hence the variable x_2 is matched with y .

The rule (T-HNDL) for handling constructs **with h handle c** is one of the most important rules of our system. It assumes that the handled computation c is of a type $\Sigma \triangleright T / (\forall x_r.C_1) \Rightarrow C_2$, where the control effect is impure. Even when c is pure (i.e., performs no operation), it can have an impure control effect via subtyping. Because the type of the handling construct represents how the expression is viewed from the context, it should be matched with the final answer type C_2 of the handled computation c . The premises in the second line define typing disciplines that the clauses in the installed handler h have to satisfy. First, let us consider the return clause **return $x_r \mapsto c_r$** . Because the variable x_r denotes the return value of the handled computation c , the value type T of c is assigned to x_r . Moreover, since the return clause is executed after evaluating c , the body c_r is the delimited continuation of c . Therefore, the type of c_r should be matched with the initial answer type C_1 of c . Because the variable x_r bound in the return clause can be viewed as the input to the delimited continuation c_r , it should be matched with the variable x_r bound in the impure control effect $(\forall x_r.C_1) \Rightarrow C_2$. Operation clauses are typed using the corresponding type schemes in the operation signature Σ , as explained above. Note that the rule also requires the installed handler h to include a clause for each of the operations in Σ , i.e., those that c may perform.

The rule (T-OP) for operation calls is another important rule. Consider an operation call $\text{op } v$. The rule assumes that an enclosing handler addresses the operation op by requiring that an operation signature Σ assigned to the operation call include the operation op with a type scheme $\forall X : \widetilde{B}.(x : T_1) \rightarrow ((y : T_2) \rightarrow C_1) \rightarrow C_2$, and instantiates the predicate variables \widetilde{X} in the type scheme with well-formed predicates \widetilde{A} to reflect the contextual information of the operation call. Then, it checks that the argument v has the input type $T_1[A/\widetilde{X}]$ of the operation. Finally, the rule assigns the output type $T_2[\widetilde{A}/\widetilde{X}][v/x]$ of the operation as the value type of the operation call, and $C_1[\widetilde{A}/\widetilde{X}][v/x]$ and $C_2[\widetilde{A}/\widetilde{X}][v/x]$ as the initial and final answer types of the operation call, respectively (note that the types T_2 , C_1 , and C_2 are parameterized over predicates and arguments).

The type system defines four kinds of subtyping judgments: $\Gamma \vdash T_1 <: T_2$ for value types, $\Gamma \vdash C_1 <: C_2$ for computations types, $\Gamma \vdash \Sigma_1 <: \Sigma_2$ for operation signatures, and $\Gamma \mid T \vdash S_1 <: S_2$ for control effects. Figure 4 shows the subtyping rules. The subtyping rules for control effects are adopted from the work of Sekiyama and Unno [2023], which extends subtyping for control effects given by Materzok and Biernacki [2011] to dependent typing. The rules (S-RFN) and (S-FUN) for value types are standard. The judgement $\Gamma \vDash \phi$ in (S-RFN) means the semantic validity of the formula ϕ under the assumption Γ . Subtyping between operation signatures is determined by (S-SIG). This rule is based on the observation that an operation signature Σ represents the types of operation clauses in handlers, as seen in (T-HNDL). Then, the rule (S-SIG) can be viewed as defining a subtyping relation between the types of handlers (except for return clauses): a handler for operations in Σ_1 can be used as one for operations in Σ_2 if every operation op in Σ_2 is included in Σ_1 (i.e., the handler has an operation clause for every op in Σ_2) and the type scheme of op in Σ_1

$$\begin{array}{c}
\text{Subtyping rules} \quad \boxed{\Gamma \vdash T_1 <: T_2} \quad \boxed{\Gamma \vdash \Sigma_1 <: \Sigma_2} \quad \boxed{\Gamma \vdash C_1 <: C_2} \quad \boxed{\Gamma \mid T \vdash S_1 <: S_2} \\
\frac{\Gamma, x : B \vDash \phi_1 \implies \phi_2}{\Gamma \vdash \{x : B \mid \phi_1\} <: \{x : B \mid \phi_2\}} \text{(S-RFN)} \quad \frac{\Gamma \vdash T_2 <: T_1 \quad \Gamma, x : T_2 \vdash C_1 <: C_2}{\Gamma \vdash (x : T_1) \rightarrow C_1 <: (x : T_2) \rightarrow C_2} \text{(S-FUN)} \\
\frac{\overline{(\Gamma, X_i : \widetilde{B}_i \vdash F_{1i} <: F_{2i})_i}}{\Gamma \vdash \{(\text{op}_i : \forall X_i : \widetilde{B}_i. F_{1i})_i, (\text{op}'_i : \forall X'_i : \widetilde{B}'_i. F'_{1i})_i\} <: \{(\text{op}_i : \forall X_i : \widetilde{B}_i. F_{2i})_i\}} \text{(S-SIG)} \\
\frac{\Gamma \vdash \Sigma_2 <: \Sigma_1 \quad \Gamma \vdash T_1 <: T_2 \quad \Gamma \mid T_1 \vdash S_1 <: S_2}{\Gamma \vdash \Sigma_1 \triangleright T_1 / S_1 <: \Sigma_2 \triangleright T_2 / S_2} \text{(S-COMP)} \quad \frac{}{\Gamma \mid T \vdash \square <: \square} \text{(S-PURE)} \\
\frac{\Gamma, x : T \vdash C_{21} <: C_{11} \quad \Gamma \vdash C_{12} <: C_{22}}{\Gamma \mid T \vdash (\forall x. C_{11}) \Rightarrow C_{12} <: (\forall x. C_{21}) \Rightarrow C_{22}} \text{(S-ATM)} \quad \frac{\Gamma, x : T \vdash C_1 <: C_2 \quad x \notin \text{fv}(C_2)}{\Gamma \mid T \vdash \square <: (\forall x. C_1) \Rightarrow C_2} \text{(S-EMBED)}
\end{array}$$

Fig. 4. Subtyping rules.

is a subtype of the type scheme of op in Σ_2 (i.e., the operation clause for op in the handler works as one for op in Σ_2). Given a computation type $C_1 \triangleq \Sigma_1 \triangleright T_1 / S_1$ and its supertype $C_2 \triangleq \Sigma_2 \triangleright T_2 / S_2$, a handler for operations performed by the computations of the type C_2 (i.e., the operations in Σ_2) is required to be able to handle operations performed by the computations of the type C_1 (i.e., the operations in Σ_1) because the subtyping allows deeming the computations of C_1 to be of C_2 . The safety of such handling is ensured by requiring $\Sigma_2 <: \Sigma_1$. In the rule (S-COMP), the first premise represents this requirement. The second premise $\Gamma \vdash T_1 <: T_2$ in (S-COMP) allows viewing the return values of the computations of the type C_1 as those of the type C_2 . The third premise $\Gamma \mid T_1 \vdash S_1 <: S_2$ expresses that the use of effects by the computations of the type C_1 is subsumed by the use of effects allowed by the type C_2 . It is derived by the last three rules: (S-PURE), (S-ATM), and (S-EMBED). The rule (S-PURE) just states reflexivity of the pure control effect. If both S_1 and S_2 are impure, the rule (S-ATM) is applied. Because initial answer types represent the assumptions of computations on their contexts, (S-ATM) allows strengthening the assumptions by being contravariant in them. By contrast, because final answer types represent the guarantees of how enclosing handling constructs behave, (S-ATM) allows weakening the guarantees by being covariant in them. Note that the typing context for the initial answer types is extended with the binding $x : T_1$ because they may reference the inputs to the continuations via the variable x and the inputs are of the type T_1 . Finally, the rule (S-EMBED) allows converting the pure control effect to an impure control effect $(\forall x. C_1) \Rightarrow C_2$. Because a computation c with the pure control effect performs no operation, what is guaranteed for the behavior of the handling construct enclosing c coincides with what is assumed on c 's delimited continuation. Because the guarantee and assumption are specified by the types C_2 and C_1 , respectively, if C_1 is matched with C_2 —more generally, the “assumption” C_1 implies the “guarantee” C_2 —the pure computation c can be viewed as the computation with the impure control effect $(\forall x. C_1) \Rightarrow C_2$. The first premise in (S-EMBED) formalizes this idea. Note that, because the variable x is bound in the type C_1 , the rule (S-EMBED) disallows x to occur in the type C_2 .

Finally, we state the type safety of our system. Its proof, via progress and subject reduction, is given in the supplementary material. We define \longrightarrow^* as the reflexive, transitive closure of the one-step evaluation relation \longrightarrow .

THEOREM 3.1 (TYPE SAFETY). *If $\emptyset \vdash c : \Sigma \triangleright T / S$ and $c \longrightarrow^* c'$, then one of the following holds: (1) $c' = \text{return } v$ for some v such that $\emptyset \vdash v : T$; (2) $c' = K[\text{op } v]$ for some K, op , and v such that $\text{op} \in \text{dom}(\Sigma)$; or (3) $c' \longrightarrow c''$ for some c'' such that $\emptyset \vdash c'' : \Sigma \triangleright T / S$.*

3.3 Examples

In this section, we demonstrate how our type system verifies programs with algebraic effects and handlers by showing typing derivations of a few examples. Here, we abbreviate a pure computation type $\{\} \triangleright T / \square$ to T and omit the empty typing context from typing and subtyping judgments. For simplicity, we often write $c_1 c_2$ for an expression **let** $x_1 = c_1$ **in** **let** $x_2 = c_2$ **in** $x_1 x_2$ where x_1 does not occur in c_2 . Furthermore, we deal with a pure computation as if it is a value. For example, we write **return** c for a computation **let** $x = c$ **in** **return** x if c is pure (e.g., as **return** $a - b$).

3.3.1 Example 1: Nondeterministic Computation. We first revisit the example presented in Section 2.1. In our language, it can be expressed as follows:

with h **handle** (**let** $a =$ (**let** $y =$ **decide** $()$ **in** **if** y **then** **return** 10 **else** **return** 20) **in**
let $b =$ (**let** $y' =$ **decide** $()$ **in** **if** y' **then** **return** 1 **else** **return** 2) **in** **return** $a - b$)

where $h \triangleq \{\mathbf{return} \ x_r \mapsto \mathbf{return} \ x_r, \mathbf{decide}(x, k) \mapsto \mathbf{let} \ r_t = k \ \mathbf{true} \ \mathbf{in} \ \mathbf{let} \ r_f = k \ \mathbf{false} \ \mathbf{in} \ \max \ r_t \ r_f\}$. As seen before, executing this program results in 19. Our system can assign the most precise type $\{z : \text{int} \mid z = 19\}$ to this program. We now show the typing process to achieve this. In what follows, we write $\Gamma_{\tilde{x}}$ for the typing context binding the variables \tilde{x} with some appropriate types \tilde{B} . In particular, these variables have these base types: $x_r : \text{int}$, $a : \text{int}$, $b : \text{int}$, $y : \text{bool}$, and $y' : \text{bool}$.

First, consider the types assigned to the clauses in the handler h . The return clause can be typed as $\Gamma_{x_r} \vdash \mathbf{return} \ x_r : \{z : \text{int} \mid z = x_r\}$. The clause for **decide** can be typed as follows:

$$\Gamma \vdash \mathbf{let} \ r_t = k \ \mathbf{true} \ \mathbf{in} \ \mathbf{let} \ r_f = k \ \mathbf{false} \ \mathbf{in} \ \max \ r_t \ r_f : \{z : \text{int} \mid \phi\} \quad (1)$$

where $\Gamma \triangleq X : (\text{int}, \text{bool}), x : \text{unit}, k : (y : \text{bool}) \rightarrow \{z : \text{int} \mid X(z, y)\}$, $\phi \triangleq \forall r_t r_f. X(r_t, \mathbf{true}) \wedge X(r_f, \mathbf{false}) \implies z = \max(r_t, r_f)$, and \max is a term-level function that returns the larger of given two integers. In this typing judgment, the predicate variable X abstracts over relationships between inputs y and outputs z of delimited continuations captured by calls to **decide**, and the refinement formula ϕ summarizes what the operation clause computes. Therefore, the operation signature Σ of the type of the handled computation, c_{body} in what follows, can be given as follows:

$$\Sigma \triangleq \{\mathbf{decide} : \forall X : (\text{int}, \text{bool}). (x : \text{unit}) \rightarrow ((y : \text{bool}) \rightarrow \{z : \text{int} \mid X(z, y)\}) \rightarrow \{z : \text{int} \mid \phi\}\}.$$

Therefore, we can conclude that the program is typable as desired by the following derivation

$$\frac{\Gamma_{x_r} \vdash \mathbf{return} \ x_r : \{z : \text{int} \mid z = x_r\} \quad (\text{Judgment (1)}) \quad (I) \vdash c_{\text{body}} : \Sigma \triangleright \text{int} / (\forall x_r. \{z : \text{int} \mid z = x_r\}) \Rightarrow \{z : \text{int} \mid z = 19\}}{\vdash \mathbf{with} \ h \ \mathbf{handle} \ c_{\text{body}} : \{z : \text{int} \mid z = 19\}} \quad (\text{T-HNDL})$$

if the premise (I) for c_{body} holds. We derive it by (T-LETIP), obtaining a derivation of the form

$$\frac{(II) \vdash (\mathbf{let} \ y = \mathbf{decide} \ () \ \mathbf{in} \ \mathbf{if} \ y \ \cdots) : \Sigma \triangleright \text{int} / (\forall a. C_1) \Rightarrow \{z : \text{int} \mid z = 19\} \quad (III) \Gamma_a \vdash \mathbf{let} \ b = \cdots \ \mathbf{in} \ \mathbf{return} \ a - b : \Sigma \triangleright \text{int} / (\forall x_r. \{z : \text{int} \mid z = x_r\}) \Rightarrow C_1}{(I) \vdash c_{\text{body}} : \Sigma \triangleright \text{int} / (\forall x_r. \{z : \text{int} \mid z = x_r\}) \Rightarrow \{z : \text{int} \mid z = 19\}} \quad (\text{T-LETIP})$$

for some type C_1 .

We start by examining judgement (III) because its derivation gives the constraints to identify the type C_1 . By (T-LETIP) again, we can derive

$$\frac{(III-1) \Gamma_a \vdash (\mathbf{let} \ y' = \mathbf{decide} \ () \ \mathbf{in} \ \mathbf{if} \ y' \ \cdots) : \Sigma \triangleright \text{int} / (\forall b. C_2) \Rightarrow C_1 \quad (III-2) \Gamma_{a,b} \vdash \mathbf{return} \ a - b : \Sigma \triangleright \text{int} / (\forall x_r. \{z : \text{int} \mid z = x_r\}) \Rightarrow C_2}{(III) \Gamma_a \vdash \mathbf{let} \ b = \cdots \ \mathbf{in} \ \mathbf{return} \ a - b : \Sigma \triangleright \text{int} / (\forall x_r. \{z : \text{int} \mid z = x_r\}) \Rightarrow C_1} \quad (\text{T-LETIP})$$

with the premises (III-1) and (III-2) and some type C_2 . Judgment (III-2) is derivable by

$$\frac{\Gamma_{a,b} \vdash \mathbf{return} \ a - b : \{z : \text{int} \mid z = a - b\}}{\text{(III-2-S)} \ \Gamma_{a,b} \vdash \{z : \text{int} \mid z = a - b\} <: \Sigma \triangleright \text{int} / (\forall x_r. \{z : \text{int} \mid z = x_r\}) \Rightarrow C_2} \text{(T-SUB)}$$

with the derivation of the subtyping judgment (III-2-S):

$$\frac{\Gamma_{a,b} \vdash \Sigma <: \emptyset \quad \Gamma_{a,b} \vdash \{z : \text{int} \mid z = a - b\} <: \text{int}}{\text{(III-2-S)} \ \Gamma_{a,b} \vdash \{z : \text{int} \mid z = a - b\} \vdash \square <: (\forall x_r. \{z : \text{int} \mid z = x_r\}) \Rightarrow C_2} \text{(S-COMP)}$$

The first two subtyping premises are derivable trivially. We can derive the third one by letting $C_2 \triangleq \{z : \text{int} \mid z = a - b\}$ because:

$$\frac{\Gamma_{a,b}, x_r : \{z : \text{int} \mid z = a - b\}, z : \text{int} \vDash (z = x_r) \implies (z = a - b)}{\Gamma_{a,b}, x_r : \{z : \text{int} \mid z = a - b\} \vdash \{z : \text{int} \mid z = x_r\} <: \{z : \text{int} \mid z = a - b\}} \text{(S-RFN)}$$

$$\frac{\Gamma_{a,b} \vdash \{z : \text{int} \mid z = a - b\} \vdash \square <: (\forall x_r. \{z : \text{int} \mid z = x_r\}) \Rightarrow \{z : \text{int} \mid z = a - b\}}{\Gamma_{a,b} \vdash \{z : \text{int} \mid z = a - b\} \vdash \square <: (\forall x_r. \{z : \text{int} \mid z = x_r\}) \Rightarrow C_2} \text{(S-EMBED)}$$

where the grayed part is denoted by C_2 in the original premise. We note that our type inference algorithm automatically infers such a type by constraint solving (cf. Section 4). Next, judgment (III-1) is derivable by

$$\frac{\text{(III-1-1)} \ \Gamma_a \vdash \mathbf{decide} \ () : \Sigma \triangleright \text{bool} / (\forall y'. C_3) \Rightarrow C_1}{\text{(III-1-2)} \ \Gamma_{a,y'} \vdash \mathbf{if} \ y' \ \cdots : \Sigma \triangleright \text{int} / (\forall b. C_2) \Rightarrow C_3} \text{(T-LETIP)}$$

with the premises (III-1-1) and (III-1-2) and some type C_3 . By letting $C_3 \triangleq \{z : \text{int} \mid z = (y' ? (a - 1) : (a - 2))\}$, we can derive judgment (III-1-2):

$$\frac{\Gamma_{a,y'} \vdash \mathbf{if} \ y' \ \cdots : \{z : \text{int} \mid z = (y' ? 1 : 2)\}}{\Gamma_{a,y'} \vdash \{z : \text{int} \mid z = (y' ? 1 : 2)\} <: \Sigma \triangleright \text{int} / (\forall b. C_2) \Rightarrow C_3} \text{(T-SUB)}$$

$$\text{(III-1-2)} \ \Gamma_{a,y'} \vdash \mathbf{if} \ y' \ \cdots : \Sigma \triangleright \text{int} / (\forall b. C_2) \Rightarrow C_3$$

It is easy to see that the first typing premise holds. We can derive the second subtyping premise similarly to subtyping judgment (III-2-S), namely, by (S-COMP) with the following derivation for the subtyping on control effects:

$$\frac{\Gamma_{a,y'}, b : \{z : \text{int} \mid z = (y' ? 1 : 2)\} \vDash (z = a - b) \implies (z = (y' ? (a - 1) : (a - 2)))}{\Gamma_{a,y'}, b : \{z : \text{int} \mid z = (y' ? 1 : 2)\} \vdash C_2 <: C_3} \text{(S-RFN)}$$

$$\frac{\Gamma_{a,y'} \vdash \{z : \text{int} \mid z = (y' ? 1 : 2)\} \vdash \square <: (\forall b. C_2) \Rightarrow C_3}{\Gamma_{a,y'} \vdash \{z : \text{int} \mid z = (y' ? 1 : 2)\} \vdash \square <: (\forall b. C_2) \Rightarrow C_3} \text{(S-EMBED)}$$

Judgment (III-1-1) is derived by (T-OP), but for that, we need to instantiate the predicate variable X in the type scheme of \mathbf{decide} in Σ with a predicate A such that the constraint $C_3 = \{z : \text{int} \mid A(z, y')\}$ imposed by (T-OP) is met. Let $A \triangleq \lambda(z, y). z = (y ? (a - 1) : (a - 2))$, which satisfies the constraint trivially. Then, by letting $C_1 \triangleq \{z : \text{int} \mid \phi\}[A/X]$, we have the following derivation:

$$\frac{\Gamma_a \vdash () : \text{unit}}{\text{(III-1-1)} \ \Gamma_a \vdash \mathbf{decide} \ () : \Sigma \triangleright \text{bool} / (\forall y'. \{z : \text{int} \mid A(z, y')\}) \Rightarrow \{z : \text{int} \mid \phi\}[A/X]} \text{(T-OP)}$$

Finally, we examine judgment (II). It is derivable by

$$\frac{\text{(II-1)} \ \vdash \mathbf{decide} \ () : \Sigma \triangleright \text{bool} / (\forall y. C_4) \Rightarrow \{z : \text{int} \mid z = 19\}}{\text{(II-2)} \ \Gamma_y \vdash \mathbf{if} \ y \ \cdots : \Sigma \triangleright \text{int} / (\forall a. C_1) \Rightarrow C_4} \text{(T-LETIP)}$$

$$\text{(II)} \ \vdash (\mathbf{let} \ y = \mathbf{decide} \ () \ \mathbf{in} \ \mathbf{if} \ y \ \cdots) : \Sigma \triangleright \text{int} / (\forall a. C_1) \Rightarrow \{z : \text{int} \mid z = 19\}$$

with the premises (II-1) and (II-2) and some type C_4 . Judgement (II-2) is derivable similarly to (III-1-2) by letting $C_4 \triangleq \{z : \text{int} \mid z = (y ? 9 : 19)\}$. For judgment (II-1), we instantiate the predicate variable X in the first call to \mathbf{decide} with the predicate $A' \triangleq \lambda(z, y). z = (y ? 9 : 19)$. Then, we can derive the judgment by the following derivation:

$$\frac{\begin{array}{l} \vdash \text{decide } () : \Sigma \triangleright \text{bool} / (\forall y. \{z : \text{int} \mid A'(z, y)\}) \Rightarrow \{z : \text{int} \mid \phi[A'/X]\} \\ \vdash \Sigma \triangleright \text{bool} / (\forall y. C_4) \Rightarrow \{z : \text{int} \mid \phi[A'/X]\} <: \Sigma \triangleright \text{bool} / (\forall y. C_4) \Rightarrow \{z : \text{int} \mid z = 19\} \end{array}}{(\text{II-1}) \vdash \text{decide } () : \Sigma \triangleright \text{bool} / (\forall y. C_4) \Rightarrow \{z : \text{int} \mid z = 19\}} \quad (\text{T-SUB})$$

(note that $C_4 = \{z : \text{int} \mid A'(z, y)\}$) where the first premise is derived by (T-OP) and the second one holds because the formula $\phi[A'/X]$ is semantically equivalent to the formula $z = 19$.

We note that the predicate variable in the type scheme of `decide` is important to typing this example. The delimited continuations captured by the two calls to `decide` behave differently. Namely, they respectively behave according to the predicates $A(u, v)$ and $A'(u, v)$ where u is the integer output given the Boolean input v . By using predicate variables, our type system gives a single type scheme to an operation that abstracts over such different behaviors.⁴

3.3.2 Example 2: State. We next revisit the second example from Section 2.1. Recall the example, which is the following program:

(with h handle (set 3; let $n = \text{get } ()$ in set 5; let $m = \text{get } ()$ in $n + m$)) 0

where $h \triangleq \{x_r \mapsto \lambda s. x_r, \text{set}(x, k) \mapsto \lambda s. k () x, \text{get}(x, k) \mapsto \lambda s. k s s\}$. For this example, we use the following syntactic sugars: $c_1; c_2 \triangleq \mathbf{let } x = c_1 \mathbf{ in } c_2$ (where x does not occur in c_2) and $\lambda x. v \triangleq \lambda x. \mathbf{return } v$. Then, the program is in our language. This program uses two operations: `set`, which updates the state value, and `get`, which returns the current state value. The handling construct returns a function that maps any integer value to the value 8; arguments to the function are initial state values, but they are not used because the function begins by initializing the state. Applying the function to the initial state value 0, the whole program returns 8.

This program is expected to be of the type $\{z : \text{int} \mid z = 8\}$. The rest of this section explains how the type system assigns this type to the program. First, the operation signature Σ for the handler h can be defined as follows:

$$\begin{aligned} \Sigma \triangleq \{ & \text{set} : \forall X : (\text{int}, \text{int}). (x : \text{int}) \rightarrow (\text{unit} \rightarrow ((s : \text{int}) \rightarrow \{z : \text{int} \mid X(z, s)\})) \\ & \rightarrow ((s : \text{int}) \rightarrow \{z : \text{int} \mid X(z, x)\}), \\ & \text{get} : \forall X : (\text{int}, \text{int}, \text{int}). \text{unit} \rightarrow ((y : \text{int}) \rightarrow ((s : \text{int}) \rightarrow \{z : \text{int} \mid X(z, s, y)\})) \\ & \rightarrow ((s : \text{int}) \rightarrow \{z : \text{int} \mid X(z, s, s)\}) \} \end{aligned}$$

Then, each sub-computation in the handled computation can be typed as follows:

$$\begin{aligned} & \vdash \text{set } 3 : \Sigma \triangleright \text{int} / (\forall_. (s : \text{int}) \rightarrow \{z : \text{int} \mid z = s + 5\}) \Rightarrow (s : \text{int}) \rightarrow \{z : \text{int} \mid z = 3 + 5\} \\ & \vdash \text{get } () : \Sigma \triangleright \text{int} / (\forall n. (s : \text{int}) \rightarrow \{z : \text{int} \mid z = n + 5\}) \Rightarrow (s : \text{int}) \rightarrow \{z : \text{int} \mid z = s + 5\} \\ n : \text{int} & \vdash \text{set } 5 : \Sigma \triangleright \text{int} / (\forall_. (s : \text{int}) \rightarrow \{z : \text{int} \mid z = n + s\}) \Rightarrow (s : \text{int}) \rightarrow \{z : \text{int} \mid z = n + 5\} \\ n : \text{int} & \vdash \text{get } () : \Sigma \triangleright \text{int} / (\forall m. (s : \text{int}) \rightarrow \{z : \text{int} \mid z = n + m\}) \Rightarrow (s : \text{int}) \rightarrow \{z : \text{int} \mid z = n + s\} \\ n : \text{int}, m : \text{int} & \vdash \mathbf{return } n + m : \\ \Sigma \triangleright \{x_r : \text{int} \mid x_r = n + m\} & / (\forall x_r. (s : \text{int}) \rightarrow \{z : \text{int} \mid z = x_r\}) \Rightarrow (s : \text{int}) \rightarrow \{z : \text{int} \mid z = n + m\} \end{aligned}$$

The first four judgments are derived by (T-OP) with appropriate instantiation of the type schemes of `set` and `get`. The last judgement is derived using (S-EMBED) as in the first example. The type of the handled computation is derived from these computation types, taking the following form:

$$\Sigma \triangleright \text{int} / (\forall x_r. (s : \text{int}) \rightarrow \{z : \text{int} \mid z = x_r\}) \Rightarrow (s : \text{int}) \rightarrow \{z : \text{int} \mid z = 8\}.$$

Therefore, by (T-HNDL), the type of the handling construct is $(s : \text{int}) \rightarrow \{z : \text{int} \mid z = 8\}$, and by (T-APP), the type of the whole program is $\{z : \text{int} \mid z = 8\}$ as promised.

⁴An alternative approach is to use intersection types (i.e., allow a set of types to be given to an operation). But, we find our approach more uniform and modular as it is able to give a single compact type scheme and enables operation signatures to be unaware of in which contexts operations are called.

3.3.3 *Example 3: File Manipulation.* Finally, we consider the last example in Section 2.1 that manipulates a specified file. Because the example uses nondeterministic while-loop constructs **while** (\star) $\{c\}$, we informally extend our language with them.⁵ The semantics of the while-loop constructs is given by the reduction rules **while** (\star) $\{c\} \longrightarrow c$; **while** (\star) $\{c\}$ and **while** (\star) $\{c\} \longrightarrow \text{return } ()$, and the typing rule is given as follows:

$$\frac{\Gamma \vdash c : \Sigma \triangleright \text{unit} / C \Rightarrow C}{\Gamma \vdash \text{while } (\star) \{c\} : \Sigma \triangleright \text{unit} / C \Rightarrow C} \text{ (T-LOOP)}$$


Note that it is easy to adapt the type safety to this extension.

Recall that the example for file manipulation is the following function:

$$v \triangleq \lambda x. \text{while } (\star) \{ \text{open } x; \text{while } (\star) \{ \text{let } y = \text{read}() \text{ in write } (y \wedge "X") \}; \text{close } () \} .$$

The regular scheme stipulating the valid use of the file operations is $(\text{open } (\text{read} \mid \text{write})^* \text{close})^*$, which is equivalent to the automaton to the right.

Our idea to verify the correctness of the file manipulation is to encode the automaton states as program states, simulate the state transitions in the automaton by state-passing, and check that the file operations are used only in appropriate states. Let $Q_0 \triangleq 0$ and $Q_1 \triangleq 1$; they represent the automaton states q_0 and q_1 , respectively. We suppose that an effect handler implements the file operations `open`, `close`, `read`, and `write` in a state-passing style for states Q_0 and Q_1 . Then, the type scheme of each file operation can be given as an instance of the following template:



$$F(T_{\text{in}}, T_{\text{out}}, Q_{\text{pre}}, Q_{\text{post}}) \triangleq T_{\text{in}} \rightarrow (T_{\text{out}} \rightarrow (\{x : \text{int} \mid x = Q_{\text{post}}\} \rightarrow C)) \rightarrow (\{x : \text{int} \mid x = Q_{\text{pre}}\} \rightarrow C)$$

where the parameters T_{in} and T_{out} are the input and output types, respectively, of the operation, and Q_{pre} and Q_{post} are the states before and after, respectively, performing the operation. We do not specify the final answer type C concretely here because it is not important. Using this template, an operation signature Σ of the file operations is given as

$$\{ \text{open} : F(\text{str}, \text{unit}, Q_0, Q_1), \text{close} : F(\text{unit}, \text{unit}, Q_1, Q_0), \\ \text{read} : F(\text{unit}, \text{str}, Q_1, Q_1), \text{write} : F(\text{str}, \text{unit}, Q_1, Q_1) \} .$$

Note that the state transitions represented in Σ are matched with those in the automaton. Let

$$S(Q_{\text{pre}}, Q_{\text{post}}) \triangleq (\{x : \text{int} \mid x = Q_{\text{post}}\} \rightarrow C) \Rightarrow (\{x : \text{int} \mid x = Q_{\text{pre}}\} \rightarrow C) .$$

Given an effect handler h conforming to Σ and a computation c with control effect $S(n_{\text{pre}}, n_{\text{post}})$ for some n_{pre} and n_{post} , if a handling construct **with** h **handle** c is well typed, the body of h 's return clause is typed at $\{x : \text{int} \mid x = n_{\text{post}}\} \rightarrow C$ —i.e., the computation c terminates at the state n_{post} —and the handling construct **with** h **handle** c is typed at $\{x : \text{int} \mid x = n_{\text{pre}}\} \rightarrow C$ —i.e., it requires n_{pre} as the initial state to start the computation c . Therefore, if $n_{\text{pre}} = n_{\text{post}} = Q_0$, then it is guaranteed that the file operations are used in a valid manner. Furthermore, even if c is non-terminating, our type system can ensure that it does not use the file operations in an invalid manner. For example, suppose that c is a computation `close ()`; Ω where Ω is a diverging computation. If it is well typed, its final answer type is $\{x : \text{int} \mid x = Q_1\} \rightarrow C$, which indicates that **with** h **handle** c requires Q_1 as the initial state. It is clearly inconsistent with the above automaton. As another instance, suppose that c involves a computation \dots ; `close ()`; `write "X"`; \dots . This is illegal because it tries to call `write` after `close` without `open`. Our type system rejects it because the initial answer type

⁵An alternative is to encode the while-loop constructs in our language by supposing that the termination of a while-loop construct is determined by some function parameter $f : \text{unit} \rightarrow \text{bool}$.

$\{x : \text{int} \mid x = Q0\} \rightarrow C$ of `close ()` is not matched with the final answer type $\{x : \text{int} \mid x = Q1\} \rightarrow C$ of `write "X"` while they must be matched for the computation to be well typed.

We end this section by showing that the example function v can be typed at $\text{str} \rightarrow \Sigma \triangleright \text{unit} / S(Q0, Q0)$, which means that v 's body uses the file operations appropriately. Note that, for any typing context Γ' and file operation `op`, if $\text{op} : F(T_1, T_2, n_{\text{pre}}, n_{\text{post}}) \in \Sigma$ and $\Gamma' \vdash v : T_1$, then $\Gamma' \vdash \text{op } v : \Sigma \triangleright T_2 / S(n_{\text{pre}}, n_{\text{post}})$ by (T-OP). Let $\Gamma \triangleq x : \text{str}$. For the inner while-loop construct, we have the following typing derivation:

$$\frac{\frac{\Gamma \vdash \text{read}() : \Sigma \triangleright \text{str} / S(Q1, Q1) \quad \Gamma, y : \text{str} \vdash \text{write } (y^{\wedge} \text{"X"}) : \Sigma \triangleright \text{unit} / S(Q1, Q1)}{\Gamma \vdash \text{let } y = \text{read}() \text{ in write } (y^{\wedge} \text{"X"}) : \Sigma \triangleright \text{unit} / S(Q1, Q1)} \text{(T-LETIP)}}{\Gamma \vdash \text{while } (\star) \{ \text{let } y = \text{read}() \text{ in write } (y^{\wedge} \text{"X"}) \} : \Sigma \triangleright \text{unit} / S(Q1, Q1)} \text{(T-LOOP)}$$

Thus, the sub-computations of the outer while-loop construct can be typed as follows:

$$\begin{array}{ll} \Gamma \vdash \text{open } x & : \Sigma \triangleright \text{unit} / S(Q0, Q1) \\ \Gamma \vdash \text{while } (\star) \{ \text{let } y = \text{read}() \text{ in write } (y^{\wedge} \text{"X"}) \} & : \Sigma \triangleright \text{unit} / S(Q1, Q1) \\ \Gamma \vdash \text{close } () & : \Sigma \triangleright \text{unit} / S(Q1, Q0) . \end{array}$$

where the control effects express how the state changes according to the operation calls. By (T-LETIP), (T-LOOP), and (T-FUN), they then imply that v is typed at $\text{str} \rightarrow \Sigma \triangleright \text{unit} / S(Q0, Q0)$ as desired.

3.4 Discussion

In this section, we discuss the current limitations and future extensions of our system.

3.4.1 Abstraction of Effects. Our type system has no mechanism for abstraction of effects. Therefore, if we cannot know possible effects of the handled computation in advance (e.g., as in $\lambda f. \text{with } h \text{ handle } (f \text{ } ())$, where the effects of the handled computation $f \text{ } ()$ are determined by function parameter f), we have to fix its effects (both the operation signature and the control effect). A possible way to address this issue is to incorporate some mechanism to abstract effects. For operation signatures, effect polymorphism as in the existing effect systems for algebraic effects and handlers [Leijen 2017; Lindley et al. 2017], is a promising solution. However, adapting it to our system is not trivial. Effect polymorphism enables specifying a part of an operation signature as a parameter, and handling constructs implicitly forward operations in the parameter. The problem is that *our type system modifies the type schemes of forwarded operations* (see the supplementary material for detail). Therefore, even though the type schemes are involved in an operation signature parameter, we need to track how they are modified. We leave addressing this challenge for future work. For control effects, we conjecture that bounded polymorphism can be used to abstract control effects while respecting the necessary sub-effecting constraints.

3.4.2 Combination with Other Computational Effects. Algebraic effects and handlers are sometimes used with other computational effects. For example, when implementing a scheduler with algebraic effects and handlers, an imperative queue is often used to keep suspended continuations, like in an example from Multicore OCaml [2022]. Even though some computational effects can be simulated by algebraic effects and handlers themselves, it is often convenient to address them as primitive operations for efficiency. Our system does not support such primitive computational effects. It is left for future work to combine these features in one system.

3.4.3 Shallow Handlers. The handlers we adopt in this work are called *deep handlers* [Kammar et al. 2013], which is the most widely used variant. Another variant of algebraic effect handlers is *shallow handlers* [Hillerström and Lindley 2018], which we do not address in the present work. Just as

deep handlers are related to `shift0/reset0`, shallow handlers are related to `control0/prompt0` [Piróg et al. 2019]. Therefore, the type system for `control0/prompt0` with ATM [Ishio and Asai 2022] may be adapted to develop a refinement type system for shallow handlers, as we have developed our refinement type system for deep handlers based on the type systems for `shift0/reset0` with ATM [Materzok and Biernacki 2011; Sekiyama and Unno 2023].

3.4.4 Recursive Computation Types. Some useful programs with algebraic effect handlers are ill typed in our system due to the lack of support for recursive computation types. For example, consider the following program:

```
rec(f, n).with h handle if n = 0 then Err "error" else f (n - 1)
```

where $h \triangleq \{\text{Err}(msg, k) \mapsto \text{Err}(\text{sprintf } "called\ at\ \%d.\ \%s" n\ msg)\}$. This recursive function handles the error in each function call, producing its own stack trace. It cannot be typed without recursive types because the type of the handled computations appears recursively as its answer type. To see this, assume that the type of the handled computation (i.e., the conditional branch) is assigned a type $\Sigma \triangleright T / C_1 \Rightarrow C_2$ (here we consider only simple types for simplicity). Then, the type of the handling construct (i.e., the body of the function) is C_2 , which implies that the overall function has type $\text{int} \rightarrow C_2$. And so, the recursive call to the function $f(n - 1)$ also has type C_2 . Then, the type C_2 should be a subtype of $\Sigma \triangleright T / C_1 \Rightarrow C_2$ since $f(n - 1)$ is the else-branch of the conditional branch. However, we cannot derive $\Gamma \vdash C_2 <: \Sigma \triangleright T / C_1 \Rightarrow C_2$ (for some Γ) in our system because while the type on the left-hand side is C_2 itself, C_2 appears as the answer type in the control effect of the type on the right-hand side. On the other hand, using recursive types, we can give this function the following type (again, we consider only simple types for simplicity): $\text{int} \rightarrow \mu\alpha.\Sigma_\alpha \triangleright T / T \Rightarrow \alpha$ where $\Sigma_\alpha \triangleq \{\text{Err} : \text{str} \rightarrow (T \rightarrow T) \rightarrow \alpha\}$ and T is an arbitrary value type. Type $\mu\alpha.C$ denotes a recursive computation type where the type variable α refers to the whole type itself. The control effect of this type is recursively nested, which reflects the fact that the handling construct is recursively nested due to the recursive call to the function.

3.4.5 Type Polymorphic Effect Operations. Consider the following program that evaluates to `[[21]]`:

```
with {xr ↦ xr, wrap((), k) ↦ [k ()]} handle (wrap (); wrap (); 21)
```

This does not type-check in our current system because a type polymorphic operation signature like $\Sigma \triangleq \{\text{wrap} : \forall\alpha.\text{unit} \rightarrow (\text{unit} \rightarrow \alpha / S) \rightarrow \alpha \text{ list} / S'\}$ is required. It is, however, easy to extend our type system to support type polymorphic operation signatures to handle such examples. Specifically, in the typing of operation clauses c_i in the (T-HNDL) rule, one would generalize type variables, and in the (T-OP) rule, one would instantiate type polymorphism.

3.4.6 Modularity (or Abstraction) versus Preciseness (or Concreteness). In our system, operation signatures are of the form $\text{op}_i : T_i \rightarrow (T'_i \rightarrow C_i) \rightarrow C'_i$ where the types C_i and C'_i represent behavior of the effect handler. In other words, the signature reveals specific implementation details regarding effect handlers. This design, from our perspective of precise specification and verification, is valuable. Indeed, our type system can formally specify and verify the assume-guarantee-like contracts between the handler and operation-call sides. However, from the perspective of modularity and abstraction, this design choice is not the optimal one. In fact, one of the purposes of effect handlers is to abstract away the specifics so that one could later choose a different implementation.

To ensure that the handler implementation details do not leak in the operation signatures, one can introduce computation type polymorphism: The types C_i and C'_i in the operation signature above will be replaced by computation type variables, thus hiding the details. However, completely hiding the information of handler implementations in this way implies that we are not providing and verifying a detailed specification requirement for the handler implementations.

Practically speaking, rather than the two extremes, we believe that it is engineering-wise desirable to allow for a gradient between modularity (abstraction) and preciseness (concreteness) and to describe and verify types at the appropriate level of detail depending on the use case. Introducing all the polymorphisms discussed in this section might achieve this goal, but we plan to investigate whether it is indeed the case by specifying and verifying various real-world programs. In our view, the issue of how to describe types at an appropriate level of abstraction, as discussed above, is an important open problem not just for algebraic effects but for general control operators and, more broadly, for effectful computation.

4 IMPLEMENTATION

4.1 Description of Our Implementation

In this section, we describe our prototype implementation of a refinement type checking and inference system, RCAML⁶. It takes a program written in a subset of the OCaml 5 language (including algebraic data types, pattern matching, recursive functions, exceptions, mutable references⁷, let-polymorphism, and effect handlers) and a refinement-type specification for the function of interest. It first (1) obtains an ML-typed AST of the program using OCaml's compiler library, (2) infers refinement-free operation signatures and control effects, (3) generates refinement constraints for the program and its specification as Constrained Horn Clauses (CHCs) (see e.g., the work of Bjørner et al. [2015]), and finally (4) solves these constraints to verify if the program satisfies the specification. The steps (3) and (4), where the refinement type checking is reduced to CHC solving, follow existing standard approaches such as those proposed by Rondon et al. [2008] and Unno and Kobayashi [2009]. The inference of (refinement-free) operation signatures is similar to that of record types using row variables, and is mutually recursive with the inference of control effects. It is based on the inference of control effects for `shift0/reset0` [Materzok and Biernacki 2011]. As we split the steps of CHC generation and solving, we can use different solvers as the backend CHC solver depending on benchmarks. In this experiment, we used two kinds of CHC solvers: SPACER [Komuravelli et al. 2013] that is based on Property Directed Reachability (PDR) [Bradley 2011; Een et al. 2011], and PCSAT [Unno et al. 2021] that is based on template-based CEGIS [Solar-Lezama et al. 2006; Unno et al. 2021] with Z3 [de Moura and Bjørner 2008] as an SMT solver.

Because inputs to the implementation are OCaml programs that are type-checked by OCaml's type checker which does not allow ATM, the underlying OCaml types corresponding to the answer types cannot be modified. However, as remarked before in Section 1, our aim is to verify *existing* programs with algebraic effects and handlers, and, as remarked before, our ARM, that allows only modification in the refinements, is useful for that purpose.

Our implementation supports several kinds of polymorphism. In addition to the standard let-polymorphism on types, it supports refinement predicate polymorphism. The implementation extends the formal system by allowing *bounded* predicate polymorphism in which abstracted predicates can be bounded by constraints on them, and further allows predicate-polymorphic types to be assigned to let-bound terms. However, because the implementation can infer predicate-polymorphic types only at let-bindings, we used a different approach, which we will discuss in Section 4.2, to simulate predicate polymorphism in operation signatures.

Another notable point is that our implementation deals with operations and exceptions uniformly. That is, exception raising is treated as an operation invocation and it can be handled by a certain kind of effect handlers which have clauses for exceptions (the exception clauses are included in the effect handlers of OCaml by default).

⁶available at <https://github.com/hiroshi-unno/coar>

⁷Strong updates [Foster et al. 2002] are not supported.

4.2 Evaluation

We performed a preliminary experiments to evaluate our method on some benchmark programs that use algebraic effect handlers. The benchmarks are based on example programs from [Bauer and Pretnar \[2015\]](#) and the repository of the Eff language [[Pretnar 2022](#)]. We gathered the effect handlers in those examples and created benchmark programs each of which uses one of the effect handlers. We also added a refinement type specification of the function of interest to each benchmark. (Other auxiliary functions are not given such extra information, and so their types are *inferred automatically* even for recursive functions.) Most benchmarks could be solved automatically without the annotations, but some need them as hints. We discuss the details at the end of this section. It is also notable that, although the examples presented in [Section 3.3](#) focus on the specifications specialized in concrete, constant values such as $\{z : \text{int} \mid z = 19\}$ for Example 1, the benchmarks include programs that demonstrate that our type system and implementation can address more general specifications. For instance, the specification for the benchmark `choose-max-SAT.ml`, which is a general version of Example 1 where the constants 10, 20, 1, and 2 are replaced by parameters u , v , x , and y , respectively, of a function `main` to be verified, is as follows:

$$\vdash \text{main} : (u : \text{int}) \rightarrow (v : \{z : \text{int} \mid z \geq u\}) \rightarrow (x : \text{int}) \rightarrow (y : \{z : \text{int} \mid z \geq x\}) \rightarrow \{z : \text{int} \mid z = v - x\}$$

We refer to the supplementary material for the source code and the specifications of our benchmarks. All the experiments were conducted on Intel Xeon Platinum 8360Y, 256 GB RAM.

[Table 1](#) shows the results of the evaluation. The files that are suffixed with `-SAT` are expected to result in “SAT”, that is, the programs are expected to be typed with the refinement types given as their specifications. The other files (suffixed with `-UNSAT`) are expected to result in “UNSAT”, that is, the programs are expected not to be typed with the given refinement types. For each program, we conducted verification in two configurations ((1) `SPACER`, and (2) `PCSAT`). The field “time” indicates the time spent in the whole process of the verification. We set the timeout to 600 seconds. Our implementation successfully answered correct result for most programs. For instance, we show the benchmark `io-write-2-SAT.ml` as an example (where `@annot_MB` is an effect annotation written in the underlying OCaml type, explained in the last paragraph of this section):

```
let[@annot_MB "(unit -> ({Write: s} |> unit / s3 => s3)) -> unit * int list"
  accumulate (body: unit -> unit) = match_with body () {
    retc = (fun v -> (v, [])); exnc = raise;
    effc = fun (type a) (e: a eff) -> match e with
      | Write x -> Some (fun (k: (a, _) continuation) ->
        let (v, xs) = continue k () in (v, x :: xs) ) }
let write_all l = accumulate (fun () ->
  let rec go li = match li with
    | [] -> () | s :: ss -> let _ = perform (Write s) in go ss
  in go l )
```

It iterates over a list l to pass its elements to the operation `Write`, and the handler for `Write` accumulates the passed elements into another list. It is checked against the following specification:

$$\vdash \text{write_all} : \{z : \text{int list} \mid z \neq []\} \rightarrow \{z : \text{unit} \times \text{int list} \mid \forall u, v. z = (u, v) \Rightarrow v \neq []\}$$

That is, if the iterated list is not empty, the accumulated list is not, either. Our implementation successfully answered that `write_all` satisfies the specification, with the following inferred type:

$$(l : \{z : \text{int list} \mid z \neq []\}) \rightarrow \{z : \text{unit} \times \{z' : \text{int} \mid l \neq []\} \text{ list} \mid \phi\}$$

where $\phi \triangleq \exists t : \text{int list}. (t = [] \vee z.2 \neq []) \wedge t \neq [] \wedge l \neq []$ and $z.2$ means the second element of the pair z . ARM is indispensable for this example because the initial answer type of the body of the

Table 1. Evaluation results

file name	SPACER		PCSAT	
	result correct?	time (sec.)	result correct?	time (sec.)
amb-1-SAT.ml	Yes	0.55	Yes	15.30
amb-1-UNSAT.ml	Yes	0.72	Yes	63.62
amb-2-SAT.ml	Yes	2.31	Yes	31.48
amb-2-UNSAT.ml	Yes	2.26	-	timeout [†]
amb-3-SAT.ml	Yes	3.20	Yes	182.41
amb-3-simpl-SAT.ml	Yes	1.71	Yes	16.79
bfs-SAT.ml	No ^{*1}	1.67	-	timeout ^{*1}
bfs-UNSAT.ml	Yes	2.00	-	timeout [‡]
bfs-simpl-SAT.ml	No ^{*1}	2.22	-	timeout ^{*1}
choose-all-SAT.ml	Yes	16.23	-	timeout [†]
choose-all-UNSAT.ml	Yes	12.56	-	timeout [†]
choose-max-SAT.ml	Yes	23.08	-	timeout [†]
choose-max-UNSAT.ml	Yes	15.97	-	timeout [†]
choose-sum-SAT.ml	Yes	1.54	-	timeout [†]
choose-sum-UNSAT.ml	Yes	7.99	Yes	15.00
deferred-1-SAT.ml	Yes	0.46	Yes	4.49
deferred-1-UNSAT.ml	Yes	0.27	Yes	4.09
deferred-2-SAT.ml	Yes	0.43	Yes	4.38
distribution-SAT.ml	Abort [‡]	-	-	timeout ^{*2}
distribution-UNSAT.ml	Abort [‡]	-	-	timeout ^{*2}
expectation-SAT.ml	Yes	0.51	Yes	7.25
expectation-UNSAT.ml	Yes	1.45	Yes	7.33
io-read-1-SAT.ml	Yes	0.43	Yes	13.90
io-read-1-UNSAT.ml	Yes	0.41	Yes	12.21
io-read-2-SAT.ml	Yes	0.56	Yes	21.10
io-read-3-SAT.ml	Yes	0.54	Yes	14.88
io-write-1-SAT.ml	Yes	0.32	Yes	8.48
io-write-1-UNSAT.ml	Yes	0.32	Yes	8.76
io-write-2-SAT.ml	Yes	0.46	Yes	11.33
io-write-2-UNSAT.ml	Yes	0.68	Yes	11.65
modulus-SAT.ml	Yes	14.23	Yes	11.89
modulus-UNSAT.ml	Yes	26.56	Yes	11.91
queue-1-SAT.ml	Yes	0.78	Yes	19.22
queue-1-UNSAT.ml	Yes	0.52	Yes	16.93
queue-2-SAT.ml	Yes	0.89	Yes	22.63
round-robin-SAT.ml	Yes	0.96	-	timeout [†]
round-robin-UNSAT.ml	Yes	0.73	-	timeout [†]
safe-div-1-SAT.ml	Abort [‡]	-	Yes	2.71
safe-div-1-UNSAT.ml	Abort [‡]	-	Yes	2.73
safe-div-2-SAT.ml	Abort [‡]	-	Yes	2.55
safe-div-2-UNSAT.ml	Abort [‡]	-	Yes	3.58
select-SAT.ml	-	timeout [‡]	Yes	13.28
select-UNSAT.ml	-	timeout [‡]	Yes	13.26
shift-SAT.ml	Yes	0.28	Yes	2.92
shift-UNSAT.ml	Yes	1.25	Yes	3.93
state-SAT.ml	-	timeout [‡]	Yes	33.69
state-UNSAT.ml	Yes	0.63	Yes	13.56
state-easy-SAT.ml	Yes	0.90	Yes	35.54
transaction-SAT.ml	-	timeout [‡]	Yes	15.36
transaction-UNSAT.ml	-	timeout [‡]	Yes	15.77
yield-SAT.ml	Yes	1.51	Yes	17.57
yield-UNSAT.ml	Yes	1.52	-	timeout [†]

function `go` should be $\{z : \text{unit} \times \text{int list} \mid z.2 = []\}$ (since it should be matched with the type of the return clause of the handler), while its final answer type should be $\{z : \text{unit} \times \text{int list} \mid z.2 \neq []\}$. We also present another interesting example (`queue-2-SAT.ml`) in detail in the supplementary material.

The benchmarks that were not verified correctly in both configurations are `bfs(-simpl)-SAT.ml` (marked with *1) and `distribution-(UN)SAT.ml` (marked with *2). They need some specific features which the implementation does not support. The formers need an invariant which states that there exists an element of a list that satisfies a certain property. The latter needs recursive predicates in the type of an integer list, which states a property about the sum of the elements of the list. These issues are orthogonal to the main contributions of this paper; they are about the expressiveness of the background theory used for refinement predicates, to which our novel refinement type system is agnostic. Also, `bfs(-simpl)-SAT.ml` uses mutable references which our implementation does not handle in a flow-sensitive manner (as mentioned in the footnote 7). One solution to this issue is to encode references with an effect handler as in Section 2.1, but our implementation does not do such encoding automatically. More advanced support for native effects including references is left for future work, as discussed in Section 3.4.

We discuss pros and cons between the two configurations. First, SPACER does not support division operator, and so it cannot verify some programs that use division (marked with \div , aborting with the message “Z3 Error: Uninterpreted ‘div’ in <null>”). Also, some programs can be solved in one configuration but not in the other. Among those solved by SPACER but not by PCSAT (marked with †), `round-robin-(UN)SAT.ml` timed out during the simplification of its constraints. For the remaining programs, their constraints tend to contain predicate variables that take a large number of arguments, which makes it hard for PCSAT to find solutions. Conversely, the programs solved by PCSAT but not by SPACER (marked with ‡) involve constraints where some predicate variables occur many times, which leads to complicated solutions that are difficult for SPACER to solve.

It is worth noting that our benchmarks do not rely on refinement type annotation in most places, even for recursive functions and recursive ADTs. However, a few kinds of annotations are still needed. First, as mentioned in Section 3.4, our type system does not support effect polymorphism. Therefore, we added effect annotations to function-type arguments which may perform operations when executed, as the one given to the benchmark `io-write-2-SAT.ml` using `@annot_MB`. These annotations are written in the underlying OCaml types, that is, we did not specify concrete refinements in the annotations. Second, we provided refinement type annotations for two small parts of `state-SAT.ml`, because otherwise it could not be verified within the timeout period in both configurations. Third, because our implementation infers predicate-polymorphic types only at let-bindings, we added *ghost parameters* to some operations and functions to infer precise refinement types of them which are not let-bound but need some abstraction of refinements. Ghost parameters are parameters which are used to express dependencies in dependent type checking, but have no impact on the dynamic execution of the program so they can be removed at runtime. In automated verification, completely inferring predicate variables requires higher-order predicate constraints, which are not expressible with CHC. Therefore, we provided ghost parameters to make it possible to reduce the verification to CHC solving. For example, the following is a part of `state-SAT.ml`:

```
let rec counter c =
  let i = perform (Lookup c) in
  if i = 0 then c else (perform (Update (c, i - 1))); counter (c + 1))
in counter 0
```

which is handled by a handler that simulates a mutable reference similar to that of Example 2 in Section 3.3.2. Here, we pass the variable `c` to the operation `Lookup` and `Update` as the ghost

Evaluation, Typing and subtyping rules $\boxed{c \longrightarrow c'}$ $\boxed{\Gamma \vdash c : \tau}$ $\boxed{\Gamma \vdash \tau_1 <: \tau_2}$

$$\begin{array}{c}
(c : \tau) \longrightarrow c \quad (\Lambda\alpha.c) \tau \longrightarrow c[\tau/\alpha] \quad (\widetilde{\Lambda X : \widetilde{B}.c}) \widetilde{A} \longrightarrow c[\widetilde{A}/\widetilde{X}] \quad \{(\text{op}_i = v_i)_i\} \# \text{op}_i \longrightarrow v_i \\
\frac{\Gamma \vdash c : \forall\alpha.\tau' \quad \Gamma \vdash \tau}{\Gamma \vdash c \tau : \tau'[\tau/\alpha]} \quad \frac{(\Gamma \vdash \tau_{1i} <: \tau_{2i})_i}{\Gamma \vdash \{(\text{op}_i : \tau_{1i})_i, (\text{op}'_i : \tau'_i)_i\} <: \{(\text{op}_i : \tau_{2i})_i\}} \\
\frac{\Gamma, \beta \vdash \tau_1[\tau/\alpha] <: \tau_2 \quad \Gamma, \beta \vdash \tau \quad \beta \notin \text{fv}(\forall\alpha.\tau_1)}{\Gamma \vdash \forall\alpha.\tau_1 <: \forall\beta.\tau_2}
\end{array}$$

Fig. 5. The operational semantics and the type system of the target language (excerpt).

parameter. In the formal system presented in Section 3.2 where predicate polymorphism is available in operation signatures, we can give Update the type

$$\forall X:(\text{int}, \text{int}). (x:\text{int}) \rightarrow (\text{unit} \rightarrow ((s:\text{int}) \rightarrow \{z:\text{int} \mid X(z, s)\})) \rightarrow ((s:\text{int}) \rightarrow \{z:\text{int} \mid X(z, x)\})$$

in the same way as Example 2 in Section 3.3.2, and instantiate the predicate variable X with $\lambda(z, s).z = c + 1 + s$ to correctly verify `state-SAT.ml`. On the other hand, in the implementation, since predicate polymorphism is not available in operation signatures, the handler needs to know the concrete predicate which replaces X . However, the predicate contains c , which the handler cannot know without receiving some additional information. Therefore, we need to add the ghost parameter c to Update (and the same for Lookup). This time we added them manually, but one possible approach for automating insertion of ghost parameters is to adopt the technique proposed by Unno et al. [2013]. We conjecture that a similar technique can be used for our purpose.

5 CPS TRANSFORMATION

5.1 Definitions and Properties

This section presents the crux of our CPS transformation that translate the language defined in Section 3 to a λ -calculus without effect handlers. Readers interested in the complete definitions of the target language and the CPS transformation are referred to the supplementary material.

The target language of the CPS transformation is a polymorphic λ -calculus with records and recursion. Its program and type syntax are defined as follows:

$$\begin{array}{l}
v ::= x \mid p \mid \mathbf{rec}(f : \tau_1, x : \tau_2).c \mid \widetilde{\Lambda X : \widetilde{B}.c} \mid \{(\text{op}_i = v_i)_i\} \mid \Lambda\alpha.c \\
c ::= v \mid c \ v \mid \mathbf{if} \ v \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2 \mid c \ \widetilde{A} \mid v \# \text{op} \mid c \ \tau \mid (c : \tau) \\
\tau ::= \{x : B \mid \phi\} \mid (x : \tau_1) \rightarrow \tau_2 \mid \widetilde{\forall X : \widetilde{B}. \tau} \mid \{(\text{op}_i : \tau_i)_i\} \mid \alpha \mid \forall\alpha.\tau
\end{array}$$

In the target language, values are not strictly separated from computations as those in the source language; for example, functions in function applications can be computations. The metavariables α and β range over type variables. Expressions $\Lambda\alpha.c$ and $c \ \tau$ are a type abstraction and application, respectively. Type polymorphism is introduced to express the pure control effect in the target language using *answer type polymorphism* [Thielecke 2003]. Expressions $\{(\text{op}_i = v_i)_i\}$ and $v \# \text{op}$ are a record literal and projection, respectively. We use operation names as record labels for the target language to encode handlers using records. Our CPS transformation produces programs with type annotations for proving bidirectional type-preservation. Recursive functions with type annotations and type ascriptions $(c : \tau)$ are used to annotate programs. We abbreviate $\mathbf{rec}(f : T_1, x : T_2).c$ to $\lambda x : T_2.c$ if f does not occur in c . Types are defined in a standard manner. Typing contexts Γ are extended to include type variables. The operational semantics is almost standard. Figure 5 shows

$$\begin{aligned}
\llbracket \{x : B \mid \phi\} \rrbracket &\triangleq \{x : B \mid \phi\} & \llbracket (x : T) \rightarrow C \rrbracket &\triangleq (x : \llbracket T \rrbracket) \rightarrow \llbracket C \rrbracket \\
\llbracket \Sigma \triangleright T / (\forall x. C_1) \Rightarrow C_2 \rrbracket &\triangleq \forall _ . \llbracket \Sigma \rrbracket \rightarrow ((x : \llbracket T \rrbracket) \rightarrow \llbracket C_1 \rrbracket) \rightarrow \llbracket C_2 \rrbracket \\
\llbracket \Sigma \triangleright T / \square \rrbracket &\triangleq \forall \alpha . \llbracket \Sigma \rrbracket \rightarrow (\llbracket T \rrbracket \rightarrow \alpha) \rightarrow \alpha \\
\llbracket \{(\text{op}_i : \forall X_i : \widetilde{B}_i . F_i)_i\} \rrbracket &\triangleq \{(\text{op}_i : \forall X_i : \widetilde{B}_i . \llbracket F_i \rrbracket^{\mathcal{F}})_i\} \\
\llbracket (x : T_1) \rightarrow ((y : T_2) \rightarrow C_1) \rightarrow C_2 \rrbracket^{\mathcal{F}} &\triangleq (x : \llbracket T_1 \rrbracket) \rightarrow \llbracket ((y : T_2) \rightarrow C_1) \rrbracket \rightarrow \llbracket C_2 \rrbracket \\
\llbracket (\text{op}^{\widetilde{A}} v)^{\Sigma \triangleright T / (\forall y. C_1) \Rightarrow C_2} \rrbracket &\triangleq \overline{\Lambda} \alpha . \overline{\lambda} h : \llbracket \Sigma \rrbracket . \overline{\lambda} k : (y : \llbracket T \rrbracket) \rightarrow \llbracket C_1 \rrbracket) . h \# \text{op} \widetilde{A} \llbracket v \rrbracket (\lambda y' : \llbracket T \rrbracket) . k y' \\
\llbracket (\text{with } h \text{ handle } c)^C \rrbracket &\triangleq \llbracket c \rrbracket \overline{\text{@}} \llbracket C \rrbracket \overline{\text{@}} \llbracket h^{\text{ops}} \rrbracket \overline{\text{@}} \llbracket h^{\text{ret}} \rrbracket \\
\text{where } \begin{cases} h = \{\text{return } x_r^{T_r} \mapsto c_r, (\text{op}^{\widetilde{X}_i \cdot \widetilde{B}_i} (x_i^{T_{x_i}}, k_i^{T_{k_i}}) \mapsto c_i)_i\} \\ \llbracket h^{\text{ops}} \rrbracket \triangleq \{(\text{op}_i = \overline{\Lambda} X_i : \widetilde{B}_i . \lambda x_i : \llbracket T_{x_i} \rrbracket . \lambda k_i : \llbracket T_{k_i} \rrbracket . \llbracket c_i \rrbracket)_i\} \\ \llbracket h^{\text{ret}} \rrbracket \triangleq \lambda x_r : \llbracket T_r \rrbracket . \llbracket c_r \rrbracket \end{cases}
\end{aligned}$$

Fig. 6. CPS transformation of types and expressions (excerpt).

four evaluation rules. Type ascriptions simply drop the ascribed type τ . Type applications substitute a given type τ for the bound type variable α . Predicate applications are similar. Record projections with op_i extract the associated field v_i . The type system is also standard, presented in Figure 5. We write $\Gamma \vdash \tau$ to state that all the free variables (including type and predicate ones) in the type τ are bound in the typing context Γ . The subtyping for record types allows supertypes to forget some fields in subtypes, and the types of each corresponding field in two record types to be in the subtyping relation (we deem record types, as well as records, to be equivalent up to permutation of fields). The subtyping rule for type polymorphism is a weaker variant of the containment rule for polymorphic types [Mitchell 1988]. It is introduced to emulate (S-EMBED) in the target language.

We show the key part of the CPS transformation in Figure 6. The upper half presents the transformation of types. The transformation of value types is straightforward. Operation signatures are transformed into record types, which means that operation clauses in a handler are transformed into a record. The transformations of computation types indicate that computations are transformed into functions that receive two value parameters: handlers and continuations. If the control effect is pure, the answer types of computations become polymorphic in CPS. This treatment of control effects is different from that of Materzok and Biernacki [2011], who define CPS transformation for control effects in the simply typed setting. Their CPS transformation transforms, in our notation, a computation type T / \square into the type $\llbracket T \rrbracket$, and a type $T / C_1 \Rightarrow C_2$ into the type $(\llbracket T \rrbracket \rightarrow \llbracket C_1 \rrbracket) \rightarrow \llbracket C_2 \rrbracket$ (note that they address neither operation signatures nor dependent typing). Because the latter takes continuations whereas the former does not, CPS transformation needs to know where pure computations are converted into impure ones (via subtyping). To address this issue, Materzok and Biernacki's CPS transformation focuses on typing derivations in the source language rather than expressions. However, because our aim is at reducing the typing of programs with algebraic effects and handlers to that of programs without them, we cannot assume typing derivations in the source language to be available. By treating two kinds of control effects uniformly using answer type polymorphism, our CPS transformation can focus only on expressions (with type annotations).

The lower half of Figure 6 shows the key cases of the transformation of expressions. We separate abstractions and applications in the target language into *static* and *dynamic* ones, as in the work of Hillerström et al. [2017], for proving the preservation of the operational semantics (Theorem 5.1). Redexes represented by static applications are known as *administrative redexes*, inserted and reduced

at compile (CPS-transformed) time. By contrast, redexes represented by dynamic applications are reduced at run time because they originate in the source program. Constructors for static expressions are denoted by the overline notation, like $\bar{\lambda}$, $\bar{\Lambda}$, and $\bar{\omega}$. We use the “at” symbol explicitly as an infix operator of static applications for clarification. Non-overlined abstractions and applications are dynamic ones, which are treated as ordinary expressions. Also, for backward type-preservation (Theorem 5.3), we extend the source language with type annotations. For example, in an operation call $(\text{op}^{\bar{A}} v)^{\Sigma \triangleright T / (\forall y. C_1) \Rightarrow C_2}$, \bar{A} are predicates used to instantiate the type scheme of the operation op , and $\Sigma \triangleright T / (\forall y. C_1) \Rightarrow C_2$ is the type of the operation call $\text{op } v$. Without type annotations, CPS-transformed expressions may have a type that cannot be transformed back to a type in the source language. An operation call $(\text{op}^{\bar{A}} v)^{\Sigma \triangleright T / (\forall y. C_1) \Rightarrow C_2}$ is transformed into a function that seeks the corresponding operation clause in a given handler and then applies it to a given sequence of predicates, argument, and continuation. Note that the continuation is in the η -expanded form because, for the preservation of the operational semantics, we need a dynamic lambda abstraction that corresponds to the continuation $\lambda y. \text{with } h \text{ handle } K[\text{return } y]$ introduced in the rule (E-HNDLOP) of the source language. An expression **with** h **handle** c is transformed into a function that applies the CPS-transformed handled computation to the record of the CPS-transformed operation clauses and the CPS-transformed return clause (because the return clause works as the continuation of c). The transformation preserves operational semantics bidirectionally in the following way:

THEOREM 5.1 (SIMULATION). *Let \equiv_{β} be the smallest congruence relation over expressions in the target language that satisfies $(\bar{\lambda}x : \tau.c) \bar{\omega} v \equiv_{\beta} c[v/x]$ and $(\bar{\Lambda}\alpha.c) \bar{\omega} \tau \equiv_{\beta} c[\tau/\alpha]$. If $c \longrightarrow^* \text{return } v$, then $\llbracket c \rrbracket_{\bar{\omega}\tau\bar{\omega}\{\}}(\lambda x : \tau.x) \longrightarrow^+ v'$ for some v' such that $\llbracket v \rrbracket \equiv_{\beta} v'$. Also, if $\llbracket c \rrbracket_{\bar{\omega}\tau\bar{\omega}\{\}}(\lambda x : \tau.x) \longrightarrow^+ v'$, then $c \longrightarrow^* \text{return } v$ and $\llbracket v \rrbracket \equiv_{\beta} v'$ for some v .*

(Note that τ can be any type since types are irrelevant to the operational semantics.) The first half states that if a computation c in the source language evaluates to a value-return of v , the transformed computation $\llbracket c \rrbracket$ applied to a type, an empty handler $\{\}$, and a trivial continuation $\lambda x : \tau.x$ evaluates to the transformed value $\llbracket v \rrbracket$. Similarly, the second half states the reverse direction.

Now, we state forward and backward type-preservation of the CPS transformation.

THEOREM 5.2 (FORWARD TYPE-PRESERVATION). *The following holds: (1) If $\Gamma \vdash v : T$ then $\llbracket \Gamma \rrbracket \vdash \llbracket v \rrbracket : \llbracket T \rrbracket$. (2) If $\Gamma \vdash c : C$ then $\llbracket \Gamma \rrbracket \vdash \llbracket c \rrbracket : \llbracket C \rrbracket$.*

THEOREM 5.3 (BACKWARD TYPE-PRESERVATION). *The following holds: (1) If $\emptyset \vdash \llbracket v \rrbracket : \tau$, then there exists some T such that $\emptyset \vdash v : T$ and $\emptyset \vdash \llbracket T \rrbracket <: \tau$. (2) If $\emptyset \vdash \llbracket c \rrbracket : \tau$, then there exists some C such that $\emptyset \vdash c : C$ and $\emptyset \vdash \llbracket C \rrbracket <: \tau$.*

Theorem 5.3 is implied immediately by backward type preservation of the CPS transformation for *open* expressions. See the supplementary material for the statement for open expressions. Theorem 5.3 indicates that it is possible to reduce typechecking in our source language to that in a language without effect handlers. That is, if ones want to verify whether an expression c has type C , they can obtain the same result as the direct verification by first applying CPS transformation to c and C , and then checking whether $\llbracket c \rrbracket$ has type $\llbracket C \rrbracket$ with a refinement type verification tool that does not support algebraic effect handlers.

Type annotations in the source language are necessary to restrict the image of the transformation. Without them, a CPS-transformed program may be of a type τ that cannot be transformed to a type in the source language inversely (i.e., there exists no type C in the source language satisfying $\llbracket C \rrbracket = \tau$). For example, consider $\bar{\Lambda}\alpha.\bar{\lambda}h.\bar{\lambda}k.k \ 0$, the CPS form (without annotations) of expression **return** 0. Without annotations, we can pick arbitrary types as the type of h . Therefore, it can have type $\forall\alpha.\text{bool} \rightarrow (\text{int} \rightarrow \alpha) \rightarrow \alpha$. However, there is no type C in the source language such that

$\llbracket C \rrbracket = \forall \alpha. \text{bool} \rightarrow (\text{int} \rightarrow \alpha) \rightarrow \alpha$. Even worse, the source language has no type that is a *subtype* of the type of the CPS form since `bool` and record types are incomparable with each other. Another example is $\lambda x. \bar{\lambda} \alpha. \bar{\lambda} h. \bar{\lambda} k. k \ x$, the CPS form (again, without annotations) of expression $\lambda x. \text{return } x$. Its type can be $(\text{int} \rightarrow \text{int}) \rightarrow \forall \alpha. \{ \} \rightarrow ((\text{int} \rightarrow \text{int}) \rightarrow \alpha) \rightarrow \alpha$, that is, x can be of type $\text{int} \rightarrow \text{int}$. However, there is no value type T in the source language such that $\llbracket T \rrbracket$ is a subtype of $\text{int} \rightarrow \text{int}$. Note that since a function type in the source language is in the form $(x : T_x) \rightarrow C$, the right hand side of the arrow in the CPS-transformed function type must be in the form $\forall \alpha. \{ \dots \} \rightarrow \dots$, which does not match with int . Therefore, without type annotations, Theorem 5.3 does not hold.

While our formalization requires concrete refinement type annotations in the source language, actually we can relax this restriction by using predicate variables as placeholders instead of concrete refinements in type annotations. This is because type annotations are only for prohibiting occurrences of types with unintended *structures*, not for restricting refinements. Those predicate variables are instantiated after CPS transformation with concrete predicates inferred by generating and solving CHC constraints that contain these predicate variables from the CPS-transformed expression. Formally, by allowing occurrences of predicate variables in type annotations of both the source and target language, and introducing predicate variable substitution σ , we can state that $\llbracket \sigma(c) \rrbracket = \sigma(\llbracket c \rrbracket)$. This means that, for an expression c that is annotated with types containing predicate variables, both of the followings result in the same expression: (1) first instantiating the predicate variables in c with concrete refinements, and then CPS-transforming it (i.e., CPS-transforming the concretely-annotated expression), and (2) first CPS-transforming c , and then instantiating the predicate variables in the CPS-transformed expression with the concrete refinements. In other words, concrete refinements are irrelevant to the CPS transformation. This irrelevance is ensured by the fact that refinements can depend only on first-order values because it means that handler variables h and continuation variables k , which occur only in CPS-transformed expressions, cannot be used in instantiated refinements. The reason why we have defined the CPS transformation with concrete refinements is just to state Theorem 5.2 and Theorem 5.3.

5.2 Comparison between the Direct Verification and the Indirect Verification

In this section, we compare the direct verification using our refinement system presented in Section 3 with the indirect verification via the CPS transformation presented above. One of the differences is that the direct verification requires special support of verification tools for algebraic effect handlers, while the indirect one can be done by existing tools without such support. On the other hand, the indirect verification has some disadvantages. First, in most cases, CPS-transformed programs tend to be complicated and be in the forms quite different from the source programs. This complexity incurred in the indirect typechecking may lead to confusing error messages when the typechecking fails. Transforming the inferred complex types back to the types of the source language would be helpful, but it is unclear whether we can do this because the inferred types of the CPS-transformed expressions do not necessarily correspond to the CPS-transformed types of the source expressions, as stated in Section 5.1. By contrast, because the direct typechecking deals with the structures of the source programs as they are, error messages can be made more user-friendly. Second, our CPS transformation needs a non-negligible amount of type annotations—type annotations are necessary in `let`-expressions, conditional branches, and recursive functions as well as operation calls and handling constructs. In practice, it is desired to infer as many types as possible. However, it seems quite challenging to define a CPS transformation that enjoys backward type-preservation and needs no, or few, type annotations. One of the possible approaches for addressing type annotations in more automated way is to use the underlying simple type system of our refinement type system for algebraic effect handlers. As mentioned in Section 5.1, concrete refinements are not necessary for

type annotations. Therefore, we can generate type annotations for an expression using its simple type inferred by the underlying type system.

We also compare these two approaches based on an experiment. We used some direct style (DS) programs (i.e., programs using algebraic effect handlers), and for each program, we applied our CPS transformation manually and ran the verification on both DS one and CPS one. Additionally, we also compared them with optimized CPS programs where administrative redexes were reduced. We used the same implementation as the one in Section 4 with the configuration of SPACER. We added annotations of source programs to only top-level, closed first-order expressions, but the correctness of the verification can be justified by the preservation of dynamic semantics.

Table 2 shows the results of the experiment. The columns “✓?” show whether the verification result is correct. The columns “time” are in seconds. Some programs have no big difference in verification time among the three variants, but there are two notable things. First, optimized CPS version of choose-sum took more time than the other versions. This seems because the size of the program became larger by the optimization.

Table 2. Evaluation results of CPS transformation

program	DS		CPS		CPS (opt)	
	✓?	time	✓?	time	✓?	time
amb-2	Yes	1.30	Yes	1.32	Yes	0.91
choose-easy	Yes	0.26	Yes	0.27	Yes	0.22
choose-sum	Yes	2.18	Yes	1.79	Yes	12.87
io-read-2	Yes	0.66	No	1.29	No	0.62
simple	Yes	0.11	Yes	0.16	Yes	0.14

The CPS choose-sum program contains some branching expressions and each branch uses variables representing its continuation and the outer handler. By reducing administrative redexes in the program, these variables are instantiated with a concrete continuation and handler, that is, the continuation and handler are copied to each branch, which results in larger size of the program and its constraints generated during the verification. Second, CPS version of io-read-2 could not be verified correctly. One possible reason is lack of support for higher-order predicate polymorphism. Since CPS programs explicitly pass around continuations, their types tend to be higher-order. Then, in some cases, higher-order predicate polymorphism becomes necessary by CPS transformation.

6 RELATED WORK

6.1 Algebraic Effects and Handlers

Algebraic effect handlers introduced by Plotkin and Pretnar [2013] turned out to be greatly expressive, which have inspired researchers and programming language designers and leads to a variety of implementations [Bauer and Pretnar 2015; Brady 2013; Kammar et al. 2013; Leijen 2017; Lindley et al. 2017; Sivaramakrishnan et al. 2021]. For advanced verification of algebraic effects and handlers, Ahman [2017] proposed a dependent type system for algebraic effects and handlers. Brady [2013] introduced algebraic effect handlers to Idris, a dependently typed programming language. In contrast to our system, these systems do not allow initial answer types to depend on values passed to continuations. Ahman and Plotkin [2015] investigated an algebraic treatment of computational effects with refinement types, but their language is not equipped with effect handlers. To our knowledge, there is no research focusing on refinement type systems with support for algebraic effect handlers and their implementations for automated verification.

Cong and Asai [2022] provided a type system with ATM for algebraic effect handlers in a simply typed setting. Compared with ours, their system is limited in a few points. First, it allows programs to use only one operation. Second, the operation can be invoked two or more times only when it is handled by an effect handler where the result types of the return and operation clause are the

same. This limitation is particularly critical for our aim, program verification, because it means that there is no way to track the state of continuations that changes with the execution of programs. For instance, the examples presented in Section 3.3 cannot be verified under such a restriction because they include multiple calls to an operation and each call changes the state of continuations. Our type system has none of these limitations—it supports multiple operations and an unlimited number of calls to operations even under a handler with clauses of different types. The key idea of our system to allow such a handler is to introduce the abstraction of operation clauses over predicates. By this abstraction, our type system can represent how the same operation clause behaves differently under different continuations.

Our CPS transformation is based on Hillerström et al. [2017]. They defined a CPS transformation from a language with effect handlers but without dependent/refinement types, and proved that it enjoys forward type-preservation, but they, and others, such as Cong and Asai [2022], who studied CPS transformation for effect handlers, did not consider the backward direction. Their transformation also assumes that programs are fully annotated with types.

6.2 Type Systems for Other Delimited Control Operators

ATM was proposed by Danvy and Filinski [1990] to type more expressions with the delimited control operators `shift/reset`. Cong and Asai [2018] proposed a dependent type system for `shift/reset`, where initial answer types cannot depend on values passed to continuations. A type system with ATM for another set of delimited control operators `shift0/reset0`, is developed by Materzok and Biernacki [2011]. They proposed a new subtyping relation that allow lifting pure expressions to impure ones. Based on their work, Sekiyama and Unno [2023] proposed a refinement type system for `shift0/reset0`. Their type system utilizes ATM for reasoning about traces (sequences of events) precisely. In their system, initial answer types *can* depend on values passed to continuations. Our control effects are inspired by their work, but they use the dependency of control effects mainly for reasoning about traces while we use it for refining properties of values. Their target operators `shift0/reset0` are closely related to our target operators, algebraic effect handlers [Forster et al. 2017; Piróg et al. 2019]. However, naively applying their approach to algebraic effect handlers does not enable precise verification. A critical difference between `shift0/reset0` and algebraic effect handlers is that, while `shift0/reset0` allows deciding the usage of captured delimited continuations per each call site of the continuation-capture operator `shift0`, algebraic effect handlers require all the calls to the same operation under a handler h to be interpreted by the same operation clause in h . This hinders precise verification of the use of continuations per each operation call. Our type system solves this problem by abstracting the type schemes of operations over predicates.

7 CONCLUSION

We developed a sound refinement type system for algebraic effects and handlers, which adopts the concept of ATM (especially, ARM) to capture how the use of effects and the handling of them influence the results of computations. This enables precise analysis of programs with algebraic effects and handlers. We also implemented the type checking and inference algorithm for a subset of OCaml 5 and demonstrated the usefulness of ARM. Additionally, we defined a bidirectionally-type-preserving CPS transformation from our language with effect handlers to the language without effect handlers. It enables the reuse of existing refinement type checkers to verify programs with effect handlers, but makes programs to be verified complicated and requires them to be fully annotated. One possible direction for future work is to incorporate temporal verification as in Sekiyama and Unno [2023] into algebraic effects and handlers. Also, it is interesting to apply ARM to other variants of effect handlers, such as lexically scoped effect handlers [Biernacki et al. 2020; Zhang and Myers 2019].

DATA-AVAILABILITY STATEMENT

Our artifact is available in the GitHub repository, at <https://github.com/hiroshi-unno/coar>. The experimental results shown in Table 1 and Table 2 can be reproduced by following the instructions in `popl24ae/README.md` of the repository.

ACKNOWLEDGMENTS

We are grateful to anonymous reviewers for their helpful and useful comments on the paper, especially regarding its presentation. We also thank Yiyang Guo and Kanaru Isoda for their contribution to our implementation. This work was supported by JSPS KAKENHI Grant Numbers JP19K20247, JP22K17875, JP20H00582, JP20H04162, JP22H03564, JP20H05703, JP20K20625, and JP22H03570 as well as JST CREST Grant Number JPMJCR21M3.

REFERENCES

- Danel Ahman. 2017. Handling Fibred Algebraic Effects. *Proc. ACM Program. Lang.* 2, POPL, Article 7 (dec 2017), 29 pages. <https://doi.org/10.1145/3158095>
- Danel Ahman and Gordon Plotkin. 2015. Refinement types for algebraic effects. In *Abstracts of the 21st Meeting ‘Types for Proofs and Programs’ (TYPES)*. Institute of Cybernetics, Tallinn University of Technology, 10–11.
- Andrew W. Appel. 1992. *Compiling with Continuations*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511609619>
- Kenichi Asai. 2009. On typing delimited continuations: three new solutions to the printf problem. *Higher-Order and Symbolic Computation* 22, 3 (01 Sep 2009), 275–291. <https://doi.org/10.1007/s10990-009-9049-5>
- Andrej Bauer. 2018. What is algebraic about algebraic effects and handlers? *CoRR* abs/1807.05923 (2018). arXiv:1807.05923 <http://arxiv.org/abs/1807.05923>
- Andrej Bauer and Matija Pretnar. 2013. An Effect System for Algebraic Effects and Handlers. In *Algebra and Coalgebra in Computer Science*, Reiko Heckel and Stefan Milius (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 1–16.
- Andrej Bauer and Matija Pretnar. 2015. Programming with algebraic effects and handlers. *Journal of Logical and Algebraic Methods in Programming* 84, 1 (2015), 108–123. <https://doi.org/10.1016/j.jlamp.2014.02.001> Special Issue: The 23rd Nordic Workshop on Programming Theory (NWPT 2011) Special Issue: Domains X, International workshop on Domain Theory and applications, Swansea, 5-7 September, 2011.
- Jesper Bengtson, Karthikeyan Bhargavan, Cédric Fournet, Andrew D. Gordon, and Sergio Maffei. 2011. Refinement Types for Secure Implementations. *ACM Trans. Program. Lang. Syst.* 33, 2, Article 8 (feb 2011), 45 pages. <https://doi.org/10.1145/1890028.1890031>
- Dariusz Biernacki, Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2020. Binders by day, labels by night: effect instances via lexically scoped handlers. *Proc. ACM Program. Lang.* 4, POPL (2020), 48:1–48:29. <https://doi.org/10.1145/3371116>
- Nikolaj Bjørner, Arie Gurfinkel, Kenneth L. McMillan, and Andrey Rybalchenko. 2015. Horn Clause Solvers for Program Verification. In *Fields of Logic and Computation II: Essays Dedicated to Yuri Gurevich on the Occasion of His 75th Birthday*, Vol. 9300. 24–51.
- Aaron R. Bradley. 2011. SAT-Based Model Checking without Unrolling. In *Proceedings of the 12th International Conference on Verification, Model Checking, and Abstract Interpretation (Austin, TX, USA) (VMCAI’11)*. Springer-Verlag, Berlin, Heidelberg, 70–87.
- Edwin Brady. 2013. Programming and Reasoning with Algebraic Effects and Dependent Types. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming (Boston, Massachusetts, USA) (ICFP ’13)*. Association for Computing Machinery, New York, NY, USA, 133–144. <https://doi.org/10.1145/2500365.2500581>
- Youyou Cong and Kenichi Asai. 2018. Handling Delimited Continuations with Dependent Types. *Proc. ACM Program. Lang.* 2, ICFP, Article 69 (jul 2018), 31 pages. <https://doi.org/10.1145/3236764>
- Youyou Cong and Kenichi Asai. 2022. Understanding Algebraic Effect Handlers via Delimited Control Operators. In *Trends in Functional Programming - 23rd International Symposium, TFP 2022, Virtual Event, March 17-18, 2022, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 13401)*, Wouter Swierstra and Nicolas Wu (Eds.). Springer, 59–79. https://doi.org/10.1007/978-3-031-21314-4_4
- Olivier Danvy and Andrzej Filinski. 1990. Abstracting Control. In *Proceedings of the 1990 ACM Conference on LISP and Functional Programming (Nice, France) (LFP ’90)*. Association for Computing Machinery, New York, NY, USA, 151–160. <https://doi.org/10.1145/91556.91622>
- Leonardo de Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *TACAS’08 (Budapest, Hungary, March 29 – April 6)*, Vol. 4963. Springer Berlin Heidelberg, 337–340.

- Niklas Een, Alan Mishchenko, and Robert Brayton. 2011. Efficient Implementation of Property Directed Reachability. In *Proceedings of the International Conference on Formal Methods in Computer-Aided Design (Austin, Texas) (FMCAD '11)*. FMCAD Inc, Austin, Texas, 125–134.
- Yannick Forster, Ohad Kammar, Sam Lindley, and Matija Pretnar. 2017. On the Expressive Power of User-Defined Effects: Effect Handlers, Monadic Reflection, Delimited Control. *Proc. ACM Program. Lang.* 1, ICFP, Article 13 (aug 2017), 29 pages. <https://doi.org/10.1145/3110257>
- Jeffrey S. Foster, Tachio Terauchi, and Alexander Aiken. 2002. Flow-Sensitive Type Qualifiers. In *Proceedings of the 2002 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), Berlin, Germany, June 17-19, 2002*, Jens Knoop and Laurie J. Hendren (Eds.). ACM, 1–12. <https://doi.org/10.1145/512529.512531>
- Timothy S. Freeman and Frank Pfenning. 1991. Refinement Types for ML. In *Proceedings of the ACM SIGPLAN'91 Conference on Programming Language Design and Implementation (PLDI), Toronto, Ontario, Canada, June 26-28, 1991*, David S. Wise (Ed.). ACM, 268–277. <https://doi.org/10.1145/113445.113468>
- Daniel Hillerström and Sam Lindley. 2018. Shallow Effect Handlers. In *Programming Languages and Systems*, Sukyoung Ryu (Ed.). Springer International Publishing, Cham, 415–435.
- Daniel Hillerström, Sam Lindley, Robert Atkey, and K. C. Sivaramakrishnan. 2017. Continuation Passing Style for Effect Handlers. In *2nd International Conference on Formal Structures for Computation and Deduction (FSCD 2017) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 84)*, Dale Miller (Ed.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 18:1–18:19. <https://doi.org/10.4230/LIPIcs.FSCD.2017.18>
- Chiaki Ishio and Kenichi Asai. 2022. Type System for Four Delimited Control Operators. In *Proceedings of the 21st ACM SIGPLAN International Conference on Generative Programming: Concepts and Experiences (Auckland, New Zealand) (GPCE 2022)*. Association for Computing Machinery, New York, NY, USA, 45–58. <https://doi.org/10.1145/3564719.3568691>
- Ohad Kammar, Sam Lindley, and Nicolas Oury. 2013. Handlers in Action. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional Programming (Boston, Massachusetts, USA) (ICFP '13)*. Association for Computing Machinery, New York, NY, USA, 145–158. <https://doi.org/10.1145/2500365.2500590>
- Ohad Kammar and Matija Pretnar. 2017. No value restriction is needed for algebraic effects and handlers. *Journal of Functional Programming* 27 (2017), e7. <https://doi.org/10.1017/S0956796816000320>
- Anvesh Komuravelli, Arie Gurfinkel, Sagar Chaki, and Edmund M. Clarke. 2013. Automatic Abstraction in SMT-Based Unbounded Software Model Checking. In *Proceedings of the 25th International Conference on Computer Aided Verification - Volume 8044 (Saint Petersburg, Russia) (CAV 2013)*. Springer-Verlag, Berlin, Heidelberg, 846–862.
- Daan Leijen. 2014. Koka: Programming with Row Polymorphic Effect Types. In *Proceedings 5th Workshop on Mathematically Structured Functional Programming, MSFP@ETAPS 2014, Grenoble, France, 12 April 2014 (EPTCS, Vol. 153)*, Paul Blain Levy and Neel Krishnaswami (Eds.). 100–126. <https://doi.org/10.4204/EPTCS.153.8>
- Daan Leijen. 2017. Type Directed Compilation of Row-Typed Algebraic Effects. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (Paris, France) (POPL '17)*. Association for Computing Machinery, New York, NY, USA, 486–499. <https://doi.org/10.1145/3009837.3009872>
- PaulBlain Levy, John Power, and Hayo Thielecke. 2003. Modelling environments in call-by-value programming languages. *Information and Computation* 185, 2 (2003), 182–210. [https://doi.org/10.1016/S0890-5401\(03\)00088-9](https://doi.org/10.1016/S0890-5401(03)00088-9)
- Sam Lindley, Conor McBride, and Craig McLaughlin. 2017. Do Be Do Be Do. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (Paris, France) (POPL '17)*. Association for Computing Machinery, New York, NY, USA, 500–514. <https://doi.org/10.1145/3009837.3009897>
- Marek Materzok and Dariusz Biernacki. 2011. Subtyping Delimited Continuations. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming (Tokyo, Japan) (ICFP '11)*. Association for Computing Machinery, New York, NY, USA, 81–93. <https://doi.org/10.1145/2034773.2034786>
- John C. Mitchell. 1988. Polymorphic type inference and containment. *Information and Computation* 76, 2 (1988), 211–249. [https://doi.org/10.1016/0890-5401\(88\)90009-0](https://doi.org/10.1016/0890-5401(88)90009-0)
- Multicore OCaml. 2022. *OCaml effects examples*. Retrieved November 5, 2022 from <https://github.com/ocaml-multicore/effects-examples>
- Yoji Nanjo, Hiroshi Unno, Eric Koskinen, and Tachio Terauchi. 2018. A Fixpoint Logic and Dependent Effects for Temporal Property Verification. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, Anuj Dawar and Erich Grädel (Eds.). ACM, 759–768. <https://doi.org/10.1145/3209108.3209204>
- Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2019. Typed Equivalence of Effect Handlers and Delimited Control. In *4th International Conference on Formal Structures for Computation and Deduction (FSCD 2019) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 131)*, Herman Geuvers (Ed.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 30:1–30:16. <https://doi.org/10.4230/LIPIcs.FSCD.2019.30>
- Gordon Plotkin. 1975. Call-by-name, call-by-value and the λ -calculus. *Theoretical Computer Science* 1, 2 (1975), 125–159. [https://doi.org/10.1016/0304-3975\(75\)90017-1](https://doi.org/10.1016/0304-3975(75)90017-1)

- Gordon Plotkin and John Power. 2003. Algebraic Operations and Generic Effects. *Applied Categorical Structures* 11, 1 (01 Feb 2003), 69–94. <https://doi.org/10.1023/A:1023064908962>
- Gordon Plotkin and Matija Pretnar. 2009. Handlers of Algebraic Effects. In *Programming Languages and Systems*, Giuseppe Castagna (Ed.), Springer Berlin Heidelberg, Berlin, Heidelberg, 80–94.
- Gordon D. Plotkin and John Power. 2001. Adequacy for Algebraic Effects. In *Foundations of Software Science and Computation Structures, 4th International Conference, FOSSACS 2001 Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2001 Genova, Italy, April 2-6, 2001, Proceedings (Lecture Notes in Computer Science, Vol. 2030)*, Furio Honsell and Marino Miculan (Eds.). Springer, 1–24. https://doi.org/10.1007/3-540-45315-6_1
- Gordon D Plotkin and Matija Pretnar. 2013. Handling Algebraic Effects. *Logical Methods in Computer Science* Volume 9, Issue 4 (Dec. 2013). [https://doi.org/10.2168/LMCS-9\(4:23\)2013](https://doi.org/10.2168/LMCS-9(4:23)2013)
- Matija Pretnar. 2015. An Introduction to Algebraic Effects and Handlers. Invited tutorial paper. *Electronic Notes in Theoretical Computer Science* 319 (2015), 19–35. <https://doi.org/10.1016/j.entcs.2015.12.003> The 31st Conference on the Mathematical Foundations of Programming Semantics (MFPS XXXI).
- Matija Pretnar. 2022. *Eff*. Retrieved November 5, 2022 from <https://github.com/matijapretnar/eff>
- Patrick M. Rondon, Ming Kawaguchi, and Ranjit Jhala. 2008. Liquid Types. In *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation (Tucson, AZ, USA) (PLDI '08)*. Association for Computing Machinery, New York, NY, USA, 159–169. <https://doi.org/10.1145/1375581.1375602>
- Taro Sekiyama and Hiroshi Unno. 2023. Temporal Verification with Answer-Effect Modification: Dependent Temporal Type-and-Effect System with Delimited Continuations. *Proc. ACM Program. Lang.* 7, POPL, Article 71 (jan 2023), 32 pages. <https://doi.org/10.1145/3571264>
- KC Sivaramakrishnan, Stephen Dolan, Leo White, Tom Kelly, Sadiq Jaffer, and Anil Madhavapeddy. 2021. Retrofitting Effect Handlers onto OCaml. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (Virtual, Canada) (PLDI 2021)*. Association for Computing Machinery, New York, NY, USA, 206–221. <https://doi.org/10.1145/3453483.3454039>
- Armando Solar-Lezama, Liviu Tancau, Rastislav Bodik, Sanjit Seshia, and Vijay Saraswat. 2006. Combinatorial Sketching for Finite Programs. In *Proceedings of the 12th International Conference on Architectural Support for Programming Languages and Operating Systems (San Jose, California, USA) (ASPLOS XII)*. Association for Computing Machinery, New York, NY, USA, 404–415. <https://doi.org/10.1145/1168857.1168907>
- Nikhil Swamy, Catalin Hritcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean Karim Zinzindohoue, and Santiago Zanella Béguelin. 2016. Dependent types and multi-monadic effects in F. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, Rastislav Bodik and Rupak Majumdar (Eds.). ACM, 256–270. <https://doi.org/10.1145/2837614.2837655>
- Tachio Terauchi. 2010. Dependent types from counterexamples. In *Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January 17-23, 2010*, Manuel V. Hermenegildo and Jens Palsberg (Eds.). ACM, 119–130. <https://doi.org/10.1145/1706299.1706315>
- Hayo Thielecke. 2003. From control effects to typed continuation passing. In *Conference Record of POPL 2003: The 30th SIGPLAN-SIGACT Symposium on Principles of Programming Languages, New Orleans, Louisiana, USA, January 15-17, 2003*, Alex Aiken and Greg Morrisett (Eds.). ACM, 139–149. <https://doi.org/10.1145/604131.604144>
- Hiroshi Unno and Naoki Kobayashi. 2009. Dependent Type Inference with Interpolants. In *Proceedings of the 11th ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (Coimbra, Portugal) (PPDP '09)*. Association for Computing Machinery, New York, NY, USA, 277–288. <https://doi.org/10.1145/1599410.1599445>
- Hiroshi Unno, Yuki Satake, and Tachio Terauchi. 2018. Relatively complete refinement type system for verification of higher-order non-deterministic programs. *Proc. ACM Program. Lang.* 2, POPL (2018), 12:1–12:29. <https://doi.org/10.1145/3158100>
- Hiroshi Unno, Tachio Terauchi, and Naoki Kobayashi. 2013. Automating Relatively Complete Verification of Higher-Order Functional Programs. *SIGPLAN Not.* 48, 1 (jan 2013), 75–86. <https://doi.org/10.1145/2480359.2429081>
- Hiroshi Unno, Tachio Terauchi, and Eric Koskinen. 2021. Constraint-Based Relational Verification. In *CAV'21*. Springer Berlin Heidelberg, 742–766.
- Niki Vazou, Eric L. Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and Simon L. Peyton Jones. 2014. Refinement types for Haskell. In *Proceedings of the 19th ACM SIGPLAN international conference on Functional programming, Gothenburg, Sweden, September 1-3, 2014*, Johan Jeuring and Manuel M. T. Chakravarty (Eds.). ACM, 269–282. <https://doi.org/10.1145/2628136.2628161>
- Panagiotis Vekris, Benjamin Cosman, and Ranjit Jhala. 2016. Refinement types for TypeScript. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2016, Santa Barbara, CA, USA, June 13-17, 2016*, Chandra Krintz and Emery D. Berger (Eds.). ACM, 310–325. <https://doi.org/10.1145/2908080.2908110>
- Hongwei Xi and Frank Pfenning. 1999. Dependent Types in Practical Programming. In *POPL '99, Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Antonio, TX, USA, January 20-22, 1999*, Andrew W. Appel and Alex Aiken (Eds.). ACM, 214–227. <https://doi.org/10.1145/292540.292560>

- Yizhou Zhang and Andrew C. Myers. 2019. Abstraction-safe effect handlers via tunneling. *Proc. ACM Program. Lang.* 3, POPL (2019), 5:1–5:29. <https://doi.org/10.1145/3290318>
- He Zhu and Suresh Jagannathan. 2013. Compositional and Lightweight Dependent Type Inference for ML. In *Verification, Model Checking, and Abstract Interpretation, 14th International Conference, VMCAI 2013, Rome, Italy, January 20-22, 2013. Proceedings (Lecture Notes in Computer Science, Vol. 7737)*, Roberto Giacobazzi, Josh Berdine, and Isabella Mastroeni (Eds.). Springer, 295–314. https://doi.org/10.1007/978-3-642-35873-9_19

Received 2023-07-11; accepted 2023-11-07