

Modular Denotational Semantics for Effects with Guarded Interaction Trees

DAN FRUMIN, University of Groningen, The Netherlands

AMIN TIMANY, Aarhus University, Denmark

LARS BIRKEDAL, Aarhus University, Denmark

We present *guarded interaction trees* – a structure and a fully formalized framework for representing higher-order computations with higher-order effects in Coq, inspired by domain theory and the recently proposed interaction trees. We also present an accompanying separation logic for reasoning about guarded interaction trees. To demonstrate that guarded interaction trees provide a convenient domain for interpreting higher-order languages with effects, we define an interpretation of a PCF-like language with effects and show that this interpretation is sound and computationally adequate; we prove the latter using a logical relation defined using the separation logic. Guarded interaction trees also allow us to combine different effects and reason about them modularly. To illustrate this point, we give a modular proof of type soundness of cross-language interactions for safe interoperability of different higher-order languages with different effects. All results in the paper are formalized in Coq using the Iris logic over guarded type theory.

CCS Concepts: • **Theory of computation** → **Program semantics**; *Logic and verification*; • **Software and its engineering** → *Software libraries and repositories*.

Additional Key Words and Phrases: Coq, Iris, denotational semantics, logical relations

ACM Reference Format:

Dan Frumin, Amin Timany, and Lars Birkedal. 2024. Modular Denotational Semantics for Effects with Guarded Interaction Trees. *Proc. ACM Program. Lang.* 8, POPL, Article 12 (January 2024), 30 pages. <https://doi.org/10.1145/3632854>

1 INTRODUCTION

Interaction trees [Xia et al. 2019] are a recently proposed formalism for representing and reasoning about (possibly) non-terminating programs with side effects in Coq (a terminating type theory without effects). Since its inception, interaction trees have been applied, including but not limited, to specifying and verifying network servers [Koh et al. 2019; Zhang et al. 2021], semantics of LLVM [Zakowski et al. 2021], semantics of a language for robotics [Ye et al. 2022], non-interference [Silver et al. 2023], and verification of concurrent objects with transactional memory [Lesani et al. 2022].

The introduction of interaction trees was motivated by a desire to simplify mechanized formalizations of interactive, effectful, non-terminating computations and the developers of the ITrees library argued that ITrees can represent computations in a way which is more *modular* than representations based on operational semantics and *executable* (in contrast to earlier representations based on traces represented as predicates on events). In particular, the idea is that interaction trees can be used to give *denotational semantics* to programming languages and thus allow one to abstract away from syntactic details and reuse meta-language features such as function composition so as

Authors' addresses: Dan Frumin, d.frumin@rug.nl, University of Groningen, The Netherlands; Amin Timany, Aarhus University, Denmark, timany@cs.au.dk; Lars Birkedal, Aarhus University, Denmark, birkedal@cs.au.dk.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

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

ACM 2475-1421/2024/1-ART12
<https://doi.org/10.1145/3632854>

to obtain more robust mechanizations. And, indeed, the applications mentioned above demonstrate that interaction trees work well for giving semantics to first-order programming languages with first-order effects.

The challenge we address in this paper is that interaction trees cannot easily be used as a model of *higher-order* programming languages with *higher-order effects*, which, of course, limits the applicability of interaction trees. Indeed, the ease of use of interaction trees is enabled, in part, by two restrictions imposed on the computations represented by the interaction trees: the computations must be first-order, and the effects that the computation performs must be first-order as well. With those restrictions, the type of interaction trees forms a monad, which allows one to compose the represented computations and reason about them modularly. (In principle, one could represent higher-order computations by means of closures in interaction trees, but that would defeat the purpose of interaction trees and force one to reason about syntactic representations, which interaction trees otherwise relieves one from.) To understand the limitations to first-order programs and first-order effects, we call to mind the definition of interaction trees.

Interaction trees are possibly infinite trees with two types of branching. The first type of branching represents a “delayed” computation (similar to that of the delay monad), or a computation performing a silent step. The second type of branching represents a computation that performs an effect; different results of the effect lead to different branches. Interaction trees are formalized as coinductive types in Coq, allowing one to leverage existing infrastructure for coinductive programs and proofs:

```
CoInductive itree (E : Type -> Type) (R : Type) :=
| Ret : R -> itree E R
| Tau : itree E R -> itree E R
| Vis {A : Type} : E A -> (A -> itree E R) -> itree E R
```

Now the point is that if we wish to represent higher-order computations, then we cannot simply add a constructor `Fun : (itree E R -> itree E R) -> itree E R`, as the resulting recursive type would have negative occurrences of the recursive variable (the `itree E R` on the left of the first arrow). Similarly, if we want to support computations with higher-order effects, i.e., the result of an effect is an interaction tree itself, we run into the same problems with positivity. For example, in the following potential signature for a higher-order effect, the parameter `test` occurs in a negative position:

```
Inductive test : Type -> Type :=
| T : nat -> test (itree test unit).
```

Guarded interaction trees: Iris and guarded type theory. Our goal is to address the challenge of extending interaction trees to allow for higher-order computations and higher-order effects, in such a way that we retain some of the advantages of interaction trees; in particular we wish to obtain a representation with which we can work efficiently in Coq. From the discussion above, it is clear that a way forward is to work in a setting that allows to solve mixed-variance recursive domain equations. There are several possible choices for such a setting, including classical Scott domain theory [Scott 1976; Smyth and Plotkin 1982] and guarded type theory [Birkedal et al. 2012, 2010]. We choose to use the latter since this choice allows us to leverage the Iris program logic framework in Coq and thence obtain an efficient environment in which we can work efficiently and formally in Coq.

Thus in this paper we introduce *guarded interaction trees*, which are formally defined in guarded type theory as a solution to a guarded recursive domain equation, and we show how guarded interaction trees can be used to represent higher-order computations and higher-order effects. Moreover, we demonstrate how we can retain some of the benefits of interaction trees, in particular

modularity with respect to effects and ease of use in Coq. The extension to higher-order computations and effects does come with a certain price, in that we need to reason about guardedness, but we believe the use of Iris alleviates this.

Our Coq formalization is available online at

<https://github.com/logsem/gitrees/tree/pop124>.

(tag pop124 in the Git repository)

Contributions. In this paper we present the following contributions, all formalized as part of our extensible and adaptable Coq formalization:

- (1) We present *guarded interaction trees*, describe the associated recursion principle, and demonstrate how to write combinators to program with guarded interaction trees (Section 3).
- (2) We describe a way of *reifying* effects in the guarded interaction trees, and the reduction semantics (Section 4).
- (3) We show how to give a model of a higher-order programming language with general recursion and effects in guarded interaction trees, and show that the model is sound (Section 5).
- (4) We build a separation logic (a program logic) on top of guarded interaction trees, allowing us to reason about their behavior (Section 6).
- (5) We use the separation logic to show that the model that we construct in Section 5 satisfies computational adequacy (Section 7).
- (6) We demonstrate how multiple different effects can be combined in guarded interaction trees, and how the separation logic is used to reason about the effects locally (Section 8).
- (7) Finally, we utilize the results above, and use guarded interaction trees to show type safety of cross-language interactions for safe interoperability of languages with different effects (Section 9).

We discuss related work in Section 10. Before we present our results, we briefly go over some preliminaries about the setting that we are working in.

2 IRIS LOGIC OVER GUARDED TYPE THEORY

In this section we describe the Iris logic, in which we shall define and work with guarded interaction trees. Our treatment is brief since Iris has been described in many other papers and we are just using a small extension of the usual presentation; we refer the reader to the literature on Iris [Jung et al. 2018] and guarded type theory [Birkedal et al. 2012] for more details.

Iris is usually presented as a separation logic over a simple type theory. The model of Iris, however, models a richer type theory and in this paper we are going to make use of that and consider Iris over a guarded type theory with (1) a modicum of dependent type theory, and (2) the ability to define guarded recursive types. Both of these features are supported by the existing Coq implementation of Iris and the associated Iris proof mode [Krebbers et al. 2017b].

Note that since we are working formally in Iris in Coq, there are two logical levels at play: the statements and proofs at the Coq level (which we refer to as the meta-logic level, or as the meta level), and the statements and proofs at the level of Iris (which we refer to as the logic level).

We recall the grammar of Iris in Figure 1; the syntax consists of types, terms, and propositions. Most of the grammar is standard for higher-order intuitionistic logic, with the parts related to guarded recursion highlighted in blue. As usual in higher-order logic, we have a well-typedness judgment $x_1 : \tau_1, \dots, x_n : \tau_n \vdash t : \tau$ stating that the term t has type τ , under the assumption that the variables x_i have types τ_i . In the grammar for types, I ranges over so-called discrete types, which are meta-level types embedded into the types of Iris. Note that types include dependent types over discrete types. While we have not shown it in the grammar, we can also form types as solutions to

$$\begin{aligned}
\tau ::= & \text{iProp} \mid 0 \mid \mathbf{1} \mid \mathbb{B} \mid \text{Nat} \mid \tau + \tau \mid \tau \times \tau \mid \tau \rightarrow \tau \mid \blacktriangleright \tau \mid I \mid \sum_{i \in I} \tau_i \mid \prod_{i \in I} \tau_i \mid \dots \\
t ::= & x \in \text{Var} \mid F(t_1, \dots, t_n) \mid \text{abort } t \mid () \mid (t, t) \mid \pi_i t \mid \lambda x : \tau. t \mid t(t) \mid \\
& \text{inj}_i t \mid \text{match } t \text{ with } \text{inj}_1 x. t \mid \text{inj}_2 x. t \text{ end} \mid \text{next}(t) \mid \text{fix}_\tau \\
P ::= & \text{False} \mid \text{True} \mid t =_\tau t \mid P \wedge P \mid P \vee P \mid P \rightarrow P \mid \exists x : \tau. P \mid \forall x : \tau. P \mid \blacktriangleright P \mid \mu x : \tau. P
\end{aligned}$$

Fig. 1. Grammar for the base logic.

guarded recursive domain equations, i.e., type equations where the recursive occurrence of the type being defined is guarded under the \blacktriangleright type modality. Such types are defined up to isomorphism; we will see an example shortly: the type of guarded interaction trees will be such a recursive type and will be introduced in the following section. A useful semantic intuition for the types of Iris is that they denote (certain kinds of) time-indexed sets, i.e., families of sets indexed over natural numbers. At time step $n > 0$, the later type $\blacktriangleright \tau$ consists of the elements of τ at a later time step, i.e., at $n - 1$. At time step $n = 0$, the type $\blacktriangleright \tau$ is a singleton set. Intuitively, guarded recursive types exist because to understand what a guarded recursive type is at time step n , one only needs to understand what it is at $n - 1$, since the recursion is guarded.

Elements of $\blacktriangleright \tau$ can be constructed from elements of τ , using the next constructor, and we can form fixed points for guarded endo-functions:

$$\frac{\Gamma \vdash t : \tau}{\Gamma \vdash \text{next}(t) : \blacktriangleright \tau} \qquad \Gamma \vdash \text{fix}_\tau : (\blacktriangleright \tau \rightarrow \tau) \rightarrow \tau$$

The \blacktriangleright type former is functorial and we write $\blacktriangleright f : \blacktriangleright \tau_1 \rightarrow \blacktriangleright \tau_2$ for its action on terms $f : \tau_1 \rightarrow \tau_2$. The fixed point satisfies the equation $\text{fix}_\tau(f) = f(\text{next}(\text{fix}_\tau(f)))$.

For propositions $P : \text{iProp}$ we also have the provability judgment $\Gamma \mid P \vdash Q$ stating that Q is derivable from P in the typing context $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n$. The rules corresponding to the intuitionistic fragment are standard. Here we present the rules concerning the guarded part of the logic.

On the level of propositions, we have a(nother) *later modality* \blacktriangleright . This is the later modality most users of Iris are already familiar with. It is related to the later modality on types in that $\blacktriangleright(t =_\tau t') \dashv\vdash \text{next}(t) =_{\blacktriangleright \tau} \text{next}(t')$. We recall that \blacktriangleright can be used to define guarded recursive predicates and that it supports reasoning via Löb induction:

$$\frac{\Gamma \vdash P : \text{iProp}}{\Gamma \vdash \blacktriangleright P : \text{iProp}} \qquad \frac{\Gamma, x : \tau \vdash P : \text{iProp} \quad x \text{ is guarded in } P}{\Gamma \vdash \mu x : \tau. P : \text{iProp}} \qquad \frac{\Gamma \mid \blacktriangleright P \vdash P}{\Gamma \mid \text{True} \vdash P}$$

Iris also includes separation logic connectives; we recall those later, when we need them, in [Section 6](#).

If P is a proposition that consists only of intuitionistic logic connectives without \blacktriangleright , then we can interpret it both as a meta-level proposition (i.e. a Coq proposition), and as an Iris proposition. For such propositions we have the following result, connecting Iris with the meta-level:

THEOREM 2.1 (IRIS ADEQUACY). *Let P be a proposition containing only intuitionistic connectives. Then, if $\text{True} \vdash P$ is derivable in Iris, then P also holds at the meta-level.*

3 GUARDED INTERACTION TREES

The type of *guarded interaction trees* (or *GITrees* for short) $\text{IT}_E(A)$ is defined for a ground type A and an effect signature E , as we explain below. In [Figure 2](#) $\text{IT}_E(A)$ is written down as a guarded

$$\begin{aligned}
\text{guarded type } \mathbf{IT}_E(A) = & \text{Ret} : A \rightarrow \mathbf{IT}_E \\
& | \text{Fun} : \blacktriangleright(\mathbf{IT}_E(A) \rightarrow \mathbf{IT}_E(A)) \rightarrow \mathbf{IT}_E(A) \\
& | \text{Err} : \text{Error} \rightarrow \mathbf{IT}_E(A) \\
& | \text{Tau} : \blacktriangleright\mathbf{IT}_E(A) \rightarrow \mathbf{IT}_E(A) \\
& | \text{Vis} : \prod_{i \in I} (\text{Ins}_i(\blacktriangleright\mathbf{IT}_E(A)) \times (\text{Outs}_i(\blacktriangleright\mathbf{IT}_E(A)) \rightarrow \blacktriangleright\mathbf{IT}_E(A))) \rightarrow \mathbf{IT}_E(A)
\end{aligned}$$

Fig. 2. Guarded datatype of interaction trees.

datatype.¹ The first constructor Ret says that any element a of the ground type A can be associated with a “terminated” guarded interaction tree $\text{Ret}(a)$. The second constructor Fun says that functions are also guarded interaction trees, and it is this constructor that allows us to model higher-order computations. Since the function constructor contains a negative occurrence of $\mathbf{IT}_E(A)$ in its argument, we must put it under a \blacktriangleright . The third constructor Err represents an error state, or a stuck computation, which we take from some predefined set Error of errors. We assume that it contains at least one element $\text{RunTime} \in \text{Error}$ representing a generic run-time error. The fourth constructor Tau denotes a delayed computation, or a computation that is available “later”. We also write $\text{Tick} : \mathbf{IT}_E(A) \rightarrow \mathbf{IT}_E(A)$ for the composition $\text{Tau} \circ \text{next}$. Then the term $\text{Tick}(\alpha)$ represents a guarded interaction tree that takes a silent step to α . It satisfies the following rule for equality: $\text{Tick}(\alpha) = \text{Tick}(\beta) \dashv\vdash \blacktriangleright(\alpha = \beta)$.

Finally, the last constructor Vis allows us to model effects. The possible effects are given by the signature $E = (I, \text{Ins}_-, \text{Outs}_-)$, where I is an indexing set on the meta-level (i.e. a set of operation names), and Ins and Outs are functors determining the arities of the operations. That is $\text{Ins}_i, \text{Outs}_i : \text{Type} \rightarrow \text{Type}$ for $i \in I$. The Type argument in Ins_i and Outs_i is instantiated with \mathbf{IT}_E itself, and is used for giving signatures to higher-order effects. With this, the first argument to Vis_i is then the input for the operation, and the second parameter is a continuation which, given an arbitrary output of the operations, produces the remainder of the computation. One way to visualize this is to think of Vis_i as a node in the tree, with the annotation Ins_i and having Outs_i many branches.

We refer to guarded interaction trees $\text{Ret}(a)$ and $\text{Fun}(f)$ as GITree values, and write $\mathbf{IT}_E^v(A) \subseteq \mathbf{IT}_E(A)$ for the set of values. When quantifying over an indexing set, we implicitly coerce E to I , i.e. we write $i \in E$ to mean $i \in E.I$. Similarly we write Ins for $E.\text{Ins}$ and Outs for $E.\text{Outs}$ when the signature E is clear from the context. When the signature E is obvious or unimportant we simply write $\mathbf{IT}(A)$ for $\mathbf{IT}_E(A)$ and $\mathbf{IT}^v(A)$ for $\mathbf{IT}_E^v(A)$.

Let us demonstrate the syntax of guarded interaction trees with some running examples of effects.

Example 3.1 (Input/output on a tape). Suppose we want to model two effectful operations, for reading a number from STDIN and for writing an output on STDOUT . We will model them as

¹Formally, the datatype is given by a solution to a recursive equation, which we examine in [Section 3.1](#). But it is convenient to think of $\mathbf{IT}_E(A)$ as a recursive datatype in which every recursive occurrence is behind a \blacktriangleright .

guarded interaction trees $\mathbf{IT}_{E_{io}}(\mathbf{1} + \mathbf{Nat})$, where $\mathbf{1} = \{\()\}$ is the unit type and

$$\begin{aligned} E_{io} &\triangleq \{\text{input, output}\} \\ \text{Ins}_{\text{input}}(X) &\triangleq \mathbf{1} & \text{Outs}_{\text{input}}(X) &\triangleq \mathbf{Nat} \\ \text{Ins}_{\text{output}}(X) &\triangleq \mathbf{Nat} & \text{Outs}_{\text{output}}(X) &\triangleq \mathbf{1} \end{aligned}$$

We write Input and $\text{Output}(n)$ for the GITrees

$$\text{Input} \triangleq \text{Vis}_{\text{input}}(\(), \lambda n. \text{next}(\text{Ret}(\text{inr}(n)))) \quad \text{Output}(n) \triangleq \text{Vis}_{\text{output}}(n, \lambda x. \text{next}(\text{Ret}(\text{inl}(\()))))$$

Here we use $\text{inl}(\()) : \mathbf{1} + \mathbf{Nat}$ as a “dummy” value, since we do not care about the return value of Output .

The operations Input and Output above are represented as GITrees $\mathbf{IT}_{E_{io}}(\mathbf{1} + \mathbf{Nat})$. However, the exact ground type is not important, as long as it contains the unit $\mathbf{1}$ and the natural numbers \mathbf{Nat} . As such, we assume that we can write down operations like Input and Output as GITrees $\mathbf{IT}_{E_{io}}(A)$ where $A \simeq \mathbf{1} + \mathbf{Nat} + B$ for some type B . We return again to this point in [Section 8](#), but for now we assume that we always pick a ground type A that is “large enough” to represent all the ground values that we need.

Example 3.2 (Higher-order store). We can model higher-order store with the following signature.

$$\begin{aligned} E_{\text{store}} &\triangleq \{\text{alloc, read, write, dealloc}\} \\ \text{Ins}_{\text{alloc}}(X) &\triangleq X & \text{Outs}_{\text{alloc}}(X) &\triangleq \text{Loc} \\ \text{Ins}_{\text{read}}(X) &\triangleq \text{Loc} & \text{Outs}_{\text{read}}(X) &\triangleq X \\ \text{Ins}_{\text{write}}(X) &\triangleq \text{Loc} \times X & \text{Outs}_{\text{write}}(X) &\triangleq \mathbf{1} \\ \text{Ins}_{\text{dealloc}}(X) &\triangleq \text{Loc} & \text{Outs}_{\text{dealloc}}(X) &\triangleq \mathbf{1} \end{aligned}$$

where Loc is a countable type of locations/pointers. We write Alloc , Read , Write , and Dealloc for the following GITrees:

$$\begin{aligned} \text{Alloc}(\alpha : \mathbf{IT}(A), k : \text{Loc} \rightarrow \mathbf{IT}(A)) &\triangleq \text{Vis}_{\text{alloc}}(\text{next}(\alpha), \text{next} \circ k) \\ \text{Read}(\ell : \text{Loc}) &\triangleq \text{Vis}_{\text{read}}(\ell, \lambda x. x) \\ \text{Write}(\ell : \text{Loc}, \alpha : \mathbf{IT}(A)) &\triangleq \text{Vis}_{\text{write}}((\ell, \text{next}(\alpha)), \lambda x. \text{next}(\text{Ret}(\text{inj}(\())))) \\ \text{Dealloc}(\ell : \text{Loc}) &\triangleq \text{Vis}_{\text{dealloc}}(\ell, \lambda x. \text{next}(\text{Ret}(\text{inj}(\())))) \end{aligned}$$

Here we assume that the ground type A is isomorphic to $\mathbf{1} + B$ for some B , with the injection $\text{inj} : \mathbf{1} \rightarrow A$.

3.1 Recursion Principle for Guarded Interaction Trees

In order to write programs that eliminate GITrees, i.e. programs of type $\mathbf{IT}_E(A) \rightarrow P$, we need to come up with a suitable recursion principle. Recursion principles for inductive datatypes usually follow from the initiality principles of the defined datatypes. However, the type of GITrees is not purely inductive, as it has mixed-variance recursive occurrences, and the corresponding recursion principle should reflect that. To understand the necessary recursion principle we need to understand first how the GITrees are defined. The definition at the beginning of this section presents GITrees as a guarded datatype, but how should such a datatype be constructed? In the type theory, the type $\mathbf{IT}_E(A)$ is given as the solution to the following guarded equation:

$$\begin{aligned} \mathbf{IT}_E(A) &\simeq A + \blacktriangleright[\mathbf{IT}_E(A) \rightarrow \mathbf{IT}_E(A)] + \text{Error} + \blacktriangleright\mathbf{IT}_E(A) + \\ &\quad \Sigma_{i \in E} (\text{Ins}_i(\blacktriangleright\mathbf{IT}_E(A)) \times (\text{Outs}_i(\blacktriangleright\mathbf{IT}_E(A)) \rightarrow \blacktriangleright\mathbf{IT}_E(A))) \end{aligned} \tag{1}$$

The isomorphism is witnessed by the pair of functions (*unfold*, *fold*), and the constructors we presented at the beginning of the section are compositions of injections and *fold*. Since Equation (1) contains recursive occurrences with mixed variance, we cannot use the usual recursion principle for inductive data type. Instead, we employ a mixed initial-algebra/final-coalgebra principle, following [Freyd 1991; Pitts 1996]. To understand it better, we first write out the bi-functor, where the fixed point corresponds to the type of GITrees:²

$$F(X, Y) \simeq A + \blacktriangleright[X \rightarrow Y] + \text{Error} + \blacktriangleright Y + \sum_{i \in E} (\text{Ins}_i(\blacktriangleright Y) \times (\text{Outs}_i(\blacktriangleright X) \rightarrow \blacktriangleright Y))$$

Here the bi-functor $F(-, -)$ is contravariant in the first argument and covariant in the second one. The bi-algebra corresponding to the type of GITrees is given by the (*fold*, *unfold*) pair:

$$F(\mathbf{IT}, \mathbf{IT}) \begin{array}{c} \xleftarrow{\text{unfold}} \\ \xrightarrow{\text{fold}} \end{array} \mathbf{IT}$$

where we write \mathbf{IT} as a shorthand for $\mathbf{IT}_E(A)$. The recursion principle that we are looking for then states that this bi-algebra is both initial and terminal. That is, for any other bi-algebra (P, f, g) we have unique maps h and k such that the following diagram commutes:

$$\begin{array}{ccc} F(P, P) & \begin{array}{c} \xleftarrow{F(k, h)} \\ \xrightarrow{F(h, k)} \end{array} & F(\mathbf{IT}, \mathbf{IT}) \\ f \downarrow \uparrow g & & \text{fold} \downarrow \uparrow \text{unfold} \\ P & \begin{array}{c} \xleftarrow{h} \\ \xrightarrow{k} \end{array} & \mathbf{IT} \end{array}$$

That is, in order to construct a function $k : \mathbf{IT} \rightarrow P$, one has to provide the “unfolding” $P \rightarrow F(P, P)$, as well as functions $A \rightarrow P$, $\text{Error} \rightarrow P$, $\blacktriangleright P \rightarrow P$, $\blacktriangleright(P \rightarrow P) \rightarrow P$, and $\prod_{i \in E} \text{Ins}_i(\blacktriangleright P) \rightarrow (\text{Outs}_i(\blacktriangleright P) \rightarrow \blacktriangleright P) \rightarrow P$. This alone would allow us to *iterate* over GITrees. However, we would run into trouble if we wish to write a primitive-recursive style function. For example, we might wish to write a destructor function k such that $k(\alpha)$ returns α if α itself is a function $\text{Fun}(f)$, and $\text{Err}(\text{RunTime})$ otherwise. We cannot do so with the scheme outlined above, since in the recursive call we don’t have access to the original argument, only to the result of applying recursion to the argument. This is similar to how the iteration scheme $B \rightarrow (B \rightarrow B) \rightarrow \mathbb{N} \rightarrow B$ for natural numbers does not allow us to (easily) write the predecessor function $p : \mathbb{N} \rightarrow \mathbb{N}$ satisfying $p(0) = 0$ and $p(n + 1) = n$ if we pick $B = \mathbb{N}$.

Recursion from iteration on inductive types. Let us then look at how to solve the issue of defining primitive recursive functions on inductive types using initiality. Suppose the function $p : \mathbb{N} \rightarrow B$ that we want to construct is defined by equations $p(0) = p_1$ and $p(n + 1) = p_2(n, p(n))$. Then we can obtain this function p using the following trick: instead of eliminating \mathbb{N} into B using initiality, we eliminate it into $\mathbb{N} \times B$, in such a way that the induced map $\mathbb{N} \rightarrow \mathbb{N} \times B$ is the identity on the first component. More concretely, suppose we have maps $p_1 : 1 \rightarrow B$ and $p_2 : \mathbb{N} \times B \rightarrow B$, forming together the equations for primitive recursion. Then we construct an \mathbb{N} -algebra over $\mathbb{N} \times B$ as

$$1 + (\mathbb{N} \times B) \xrightarrow{[(0, p_1), \langle S, p_2 \rangle]} \mathbb{N} \times B$$

where $S : \mathbb{N} \rightarrow \mathbb{N}$ is the successor function. The initiality of \mathbb{N} will then induce the unique map $p : \mathbb{N} \rightarrow \mathbb{N} \times B$, which, when composed with projection $\mathbb{N} \times B \rightarrow B$, determines the recursive function given by the clauses p_1 and p_2 .

²For a detailed category-theoretic treatment, see [Birkedal et al. 2010].

Recursion/corecursion from mixed-variance types. Dually, for coinductive datatypes we can obtain a form of primitive coinduction from coiteration by using coproducts. In our case we have a datatype with mixed variance, and we use coproducts for the negative occurrences and products for the positive ones. That is, in order to eliminate \mathbf{IT} into a type P we will assume an unfolding $P \rightarrow F(P, P)$, and a folding $F(\mathbf{IT} + P, \mathbf{IT} \times P) \rightarrow P$. More concretely:

Definition 3.3 (Recursion/corecursion principle). In order to define a pair of maps $P \xrightleftharpoons[k]{h} \mathbf{IT}$, one has to define maps

- $h_u : P \rightarrow A + \blacktriangleright[P \rightarrow P] + \text{Error} + \blacktriangleright P + \sum_{i \in E} (\text{Ins}_i(\blacktriangleright P) \times (\text{Outs}_i(\blacktriangleright P) \rightarrow \blacktriangleright P))$;
- $k_{\text{Ret}} : A \rightarrow P$;
- $k_{\text{Fun}} : \blacktriangleright((\mathbf{IT} + P) \rightarrow (\mathbf{IT} \times P)) \rightarrow P$;
- $k_{\text{Err}} : \text{Error} \rightarrow P$;
- $k_{\text{Tau}} : \blacktriangleright(\mathbf{IT} \times P) \rightarrow P$;
- $k_{\text{Vis}} : \prod_{i \in E} \text{Ins}_i(\blacktriangleright(\mathbf{IT} \times P)) \rightarrow (\text{Outs}_i(\blacktriangleright(\mathbf{IT} + P)) \rightarrow \blacktriangleright(\mathbf{IT} \times P)) \rightarrow P$.

The resulting maps (h, k) will then satisfy the following *computational rules*:

- $k(\text{Ret}(a)) = k_{\text{Ret}}(a)$;
- $k(\text{Fun}(f)) = k_{\text{Fun}}(\blacktriangleright s(f))$ where $s(f) = \langle \text{id}, k \rangle \circ f \circ [\text{id}, h]$;
- $k(\text{Err}(e)) = k_{\text{Err}}(e)$;
- $k(\text{Tau}(t)) = k_{\text{Tau}}(\blacktriangleright \langle \text{id}, k \rangle(t))$;
- $k(\text{Vis}_i(x, k)) = k_{\text{Vis}}(i, \text{Ins}_i(\blacktriangleright \langle \text{id}, k \rangle)(x), \blacktriangleright \langle \text{id}, k \rangle \circ k \circ \text{Outs}_i(\blacktriangleright [\text{id}, h]))$;
- plus equations for h .

(Recall that we write $\blacktriangleright s : \blacktriangleright A \rightarrow \blacktriangleright B$ for a function $s : A \rightarrow B$.)

This recursion/corecursion principle is constructed using guarded recursion, and can be used to define a large variety of combinators. For example, we can write a generalization of the aforementioned function k that returns its argument, if the argument is a function, and returns an error otherwise.

Using the recursion principle we can define a function $\text{get_fun}(\alpha : \mathbf{IT}, f : \blacktriangleright(\mathbf{IT} \rightarrow \mathbf{IT})) \rightarrow \mathbf{IT}$ satisfying the computational rules

- $\text{get_fun}(\text{Ret}(a), f) = \text{Err}(\text{RunTime})$;
- $\text{get_fun}(\text{Fun}(g), f) = f(g)$;
- $\text{get_fun}(\text{Err}(e), f) = \text{Err}(e)$;
- $\text{get_fun}(\text{Tau}(t), f) = \text{Tau}(\blacktriangleright \text{get_fun}(t, f))$ and $\text{get_fun}(\text{Tick}(\alpha), f) = \text{Tick}(\text{get_fun}(\alpha), f)$;
- $\text{get_fun}(\text{Vis}_i(x, k), f) = \text{Vis}_i(x, \blacktriangleright \text{get_fun}(-, f) \circ k)$.

In the next section we will see how to use get_fun to define an application function $\alpha \bullet \beta$ for applying a GITree function α to a GITree argument β .

In the rest of the paper we will define other operations on GITrees using just the computational rules, with the understanding that we can write down the explicit recursor for any such set of equations. Interested readers are referred to the Coq formalization for the details.

3.2 Programming with GITrees

Using the recursion principle we can define operations on GITrees that correspond to common programming constructs. For example, with get_fun we can write a function $\text{App}_I(\alpha, \beta)$, which applies α to β if α is a function, and returns $\text{Err}(\text{RunTime})$ otherwise.

$$\text{App}_I(\alpha, \beta) = \text{get_fun}(\alpha, \lambda g. \text{Tau}(\blacktriangleright g(\beta))).$$

This operation gives us “call-by-name” application, in the sense that it satisfies

$$\text{App}_I(\text{Fun}(\text{next}(g)), \beta) = \text{Tick}(g(\beta))$$

for $g : \mathbf{IT}_E(A) \rightarrow \mathbf{IT}_E(A)$ for any argument β . In particular, it invokes the underlying function g even if the argument β is a Tick or an effect, without evaluating the argument first. In order to define a “call-by-value” application, we compose App_I with the following operation.

The function $\text{get_val}(\alpha, f)$ recurses into its argument, looking under Ticks and Vis’s, until it reaches either a function or a ground type (i.e. a value from $\mathbf{IT}_E^v(A)$), after which it applies the function $f : \mathbf{IT}_E(A) \rightarrow \mathbf{IT}_E(A)$ to it:

$$\begin{aligned} \text{get_val}(\text{Ret}(a), f) &= f(\text{Ret}(a)) & \text{get_val}(\text{Fun}(g), f) &= f(\text{Fun}(g)) \\ \text{get_val}(\text{Err}(e), f) &= \text{Err}(e) & \text{get_val}(\text{Tick}(\alpha), f) &= \text{Tick}(\text{get_val}(\alpha, f)) \\ \text{get_val}(\text{Vis}_i(x, k), f) &= \text{Vis}_i(x, \blacktriangleright \text{get_val}(-, f) \circ k) \end{aligned}$$

As syntactic sugar, we write $\text{Let } x = \alpha \text{ in } \beta(x)$ for $\text{get_val}(\alpha, \lambda x. \beta(x))$.

Now we can define the “call-by-value” application $\alpha \bullet \beta$ as $\text{get_val}(\beta, \lambda \beta_v. \text{App}_I(\alpha, \beta_v))$. This strict application then satisfies the following computational rules

$$\begin{aligned} \alpha \bullet \text{Tick}(\beta) &= \text{Tick}(\alpha \bullet \beta) & \alpha \bullet \text{Vis}_i(x, k) &= \text{Vis}_i(x, \lambda y. \text{next}(\alpha) (\blacktriangleright \bullet) k y) \\ \text{Tick}(\alpha) \bullet \beta_v &= \text{Tick}(\alpha \bullet \beta_v) & \text{Vis}_i(x, k) \bullet \beta_v &= \text{Vis}_i(x, \lambda y. k y (\blacktriangleright \bullet) \text{next}(\beta_v)) \\ \text{Fun}(\text{next}(g)) \bullet \beta_v &= \text{Tick}(g(\beta_v)) & \alpha \bullet \beta &= \text{Err}(\text{RunTime}) \text{ in other cases} \end{aligned}$$

Where $-\ (\blacktriangleright \bullet) -$ is the lifting of $-\bullet-$ to $\blacktriangleright \mathbf{IT}_E(A) \rightarrow \blacktriangleright \mathbf{IT}_E(A) \rightarrow \blacktriangleright \mathbf{IT}_E(A)$, and $\beta_v \in \mathbf{IT}_E^v(A)$ is either $\text{Ret}(a)$ or $\text{Fun}(g)$. The application function $-\bullet-$ not only simulates strict application, but it also fixes a right-to-left evaluation order of effects and computation steps.

One can see that there are common properties for the computational rules between $\text{get_fun}(-, f)$, $\text{App}_I(-, \beta)$, $\alpha \bullet -$, and $-\bullet \beta_v$ (where $\beta_v \in \mathbf{IT}_E^v(A)$): they all preserve ticks, effects, and errors. Functions that have these preservation properties are called *homomorphisms* of GITrees and will play an important role in later sections.

Definition 3.4. A function $f : \mathbf{IT}_E(A) \rightarrow \mathbf{IT}_E(A)$ is a *homomorphism*, written as $f \in \text{Hom}$, if it satisfies the following equations:

- $f(\text{Err}(e)) = \text{Err}(e)$;
- $f(\text{Tick}(\alpha)) = \text{Tick}(f(\alpha))$;
- $f(\text{Vis}_i(x, k)) = \text{Vis}_i(x, \blacktriangleright f \circ k)$

As expected from the name, the identity function is a homomorphism and composition of two homomorphisms is a homomorphism. This notion of homomorphism is inspired by the one in [Hoshino 2012]. It follows from the definition, that in order to define a homomorphism it suffices to define its action on GITree values.

Programming with GITrees and natural numbers. In the remainder of this paper we work with a lot of examples involving programming with natural numbers (as an illustrative ground type). It is then useful to assume in the remainder of this paper that the ground type A in any type $\mathbf{IT}_E(A)$ of GITrees is “large enough” to contain natural numbers, and the unit type. That is, we assume that $A \simeq 1 + \text{Nat} + \dots$, and we simply write $\text{Ret}(n)$ for $\text{Ret}(\text{inj}(n))$ and $\text{Ret}()$ for $\text{Ret}(\text{inj}'(1))$ (for appropriate injections $\text{inj} : \text{Nat} \rightarrow A$ and $\text{inj}' : 1 \rightarrow A$). We will also abbreviate $\mathbf{IT}_E(A)$ as IT or IT_E when A is generic as above or is clear from the context.

In Figure 3 we summarize the operations on GITrees that we define using recursion and other functions. The computational rules described in Figure 3 are only for the base cases; the other computational rules follow from the fact that those operations are homomorphisms. Concretely, we

$$\begin{array}{c}
\frac{n \in \text{Nat}}{\text{get_nat}(\text{Ret}(n), f) = f(n)} \qquad \frac{b \notin \text{Nat}}{\text{get_nat}(\text{Ret}(b), f) = \text{Err}(\text{RunTime})} \\
\text{get_nat}(\text{Fun}(g), f) = \text{Err}(\text{RunTime}) \qquad \text{get_fun}(\text{Ret}(a), f) = \text{Err}(\text{RunTime}) \qquad \text{get_fun}(\text{Fun}(g), f) = f(g) \\
\text{If}(\text{Ret}(0), \alpha_1, \alpha_2) = \alpha_1 \qquad \frac{n > 0}{\text{If}(\text{Ret}(n), \alpha_1, \alpha_2) = \alpha_2} \qquad \text{If}(\text{Fun}(f), \alpha_1, \alpha_2) = \text{Err}(\text{RunTime}) \\
\frac{b \notin \text{Nat}}{\text{If}(\text{Ret}(b), \alpha_1, \alpha_2) = \text{Err}(\text{RunTime})} \qquad \frac{n_1, n_2 \in \text{Nat}}{\text{NatOp}_f(\text{Ret}(n_1), \text{Ret}(n_2)) = \text{Ret}(f(n_1, n_2))} \\
\frac{\alpha_v \text{ or } \beta_v \text{ are not Ret}(n)}{\text{NatOp}_f(\alpha_v, \beta_v) = \text{Err}(\text{RunTime})} \qquad \beta_v ; \alpha = \alpha \qquad \text{While } \alpha \text{ do } \beta = \text{If}(\alpha, (\beta ; \text{Tick}(\text{While } \alpha \text{ do } \beta)), \text{Ret}(()))
\end{array}$$

Fig. 3. Programming operations on GITrees.

have the following operations. The `get_nat` function extracts a natural number from a `GITree` and applies the function $f : \text{Nat} \rightarrow \text{IT}$ to it. It is a homomorphism in the first argument. If it encounters a function `Fun(g)` or a different ground value `Ret(b)`, then it returns an error. The `If` operations test whether the first argument is zero, and picks the appropriate branch. The function $\text{If}(-, \alpha_1, \alpha_2)$ is a homomorphism. Similarly, if the first argument is not a natural number then `If` returns an error. The `NatOpf` operation applies the binary function f to its integer arguments, returning an error on all the other values. The maps $\text{NatOp}_f(\alpha, -)$ and $\text{NatOp}_f(-, \beta_v)$ are homomorphisms for $\beta_v \in \text{IT}^v$. The $\alpha ; \beta$ is a sequencing operation: it puts all the effects and ticks in α before the effects and ticks in β . This is witnessed by the fact that $(-); \alpha$ is a homomorphism. The `While α do β` represents a while loop with the conditional α and the body β ; it is defined using guarded recursion, and is equal to its one-step unfolding using the `If` construct.

Let us look at some example programs that we can write using the operations we have defined.

Example 3.5 (Factorial). In the first example, we have a factorial function that we implement using the store operations (Example 3.2).

$$\begin{aligned}
\text{fact}(n) &\triangleq \text{Alloc}(\text{Ret}(1), \lambda \text{acc}. \text{Alloc}(\text{Ret}(n), \lambda \ell. \text{factBody}(\text{acc}, \ell) ; \text{Read}(\text{acc}))) \\
\text{factBody}(\text{acc}, \ell) &\triangleq \text{While } \text{Read}(\ell) \text{ do} \\
&\quad \text{Let } i = \text{Read}(\ell) \text{ in} \\
&\quad \text{Let } r = \text{NatOp}_\times(i, \text{Read}(\text{acc})) \text{ in} \\
&\quad \text{Let } i = \text{NatOp}_-(i, \text{Ret}(1)) \text{ in} \\
&\quad \text{Write}(\text{acc}, r) ; \text{Write}(\ell, i)
\end{aligned}$$

The program `factBody` computes the factorial of the number stored in the location ℓ using an intermediate location acc for the accumulated result. The complete program `fact` then allocates the needed references and runs `factBody` before reading off the result from the location acc .

Example 3.6 (Encoding of pairs). Our definition of `GITrees` does not include arbitrary algebraic datatypes, like pairs or sums. We can, however, encode them using a Church-style encoding. We write $(\alpha, \beta) : \text{IT}$ for the guarded interaction tree

$$\text{Let } y = \beta \text{ in Let } x = \alpha \text{ in Fun}(\text{next}(\lambda f. f \bullet x \bullet y)).$$

Note that (α_v, β_v) is a GITree value whenever α_v and β_v are. Furthermore, $(\alpha, -)$ and $(-, \beta_v)$ are homomorphisms. We then define the projection functions as

$$\pi_1(\alpha) = \alpha \bullet \text{Fun}(\text{next}(\lambda a. \text{Fun}(\text{next}(\lambda b. a)))) \quad \pi_2(\alpha) = \alpha \bullet \text{Fun}(\text{next}(\lambda a. \text{Fun}(\text{next}(\lambda b. b)))).$$

The projection functions then satisfy the following computational rules:

$$\pi_1(\alpha_v, \beta_v) = \text{Tick}^3(\alpha_v) \quad \pi_2(\alpha_v, \beta_v) = \text{Tick}^3(\beta_v).$$

We can use similar style encodings to represent other algebraic datatypes as guarded interaction trees.

4 REIFICATION OF EFFECTS AND REDUCTIONS OF GITREES

GITrees allow us to conveniently write down and combine various effects. But in order to reason about the effects we also need a way of giving them meaning. In this section we establish a way of *reifying* effects of GITrees and use reification to define *reductions* of GITrees, which explain how computations represented by GITrees reduce.

In order to interpret stateful effects we assume that we have a type *State*, and each effect is interpreted using the state monad with a function:

$$r : \prod_{i \in E} \text{Ins}_i(\blacktriangleright \text{IT}_E) \times \text{State} \rightarrow \text{option}(\text{Outs}_i(\blacktriangleright \text{IT}_E) \times \text{State}).$$

We call a tuple (E, State, r) a *reifier* for the effects E . Assuming we have such a reifier, we write a function $\text{reify} : \text{IT} \times \text{State} \rightarrow \text{IT} \times \text{State}$ that satisfies

$$\frac{r_i(x, \sigma) = \text{Some}(y, \sigma') \quad k \ y = \text{next}(\beta)}{\text{reify}(\text{Vis}_i(x, k), \sigma) = (\text{Tick}(\beta), \sigma')} \quad \frac{r_i(x, \sigma) = \text{None}}{\text{reify}(\text{Vis}_i(x, k), \sigma) = (\text{Err}(\text{RunTime}), \sigma)}$$

Example 4.1 (Reification for the input/output operations Example 3.1). We take the state *State* to be a pair of two lists of natural numbers, corresponding to input and output tapes. The reifier is defined as

$$\begin{aligned} r_{\text{input}}((), (\vec{n}, \vec{m})) &= \text{Some}(n, (\vec{n}, \vec{m})) & r_{\text{input}}((), (\epsilon, \vec{m})) &= \text{None} \\ r_{\text{output}}(x, (\vec{n}, \vec{m})) &= \text{Some}(), (\vec{n}, x\vec{m}) \end{aligned}$$

Example 4.2 (Reification for the higher-order store operations Example 3.2). For higher-order store we take *State* to be the type of finite partial maps $\text{Loc} \xrightarrow{\text{fin}} \blacktriangleright \text{IT}$. The reifier function is defined in the expected way:

$$\begin{aligned} r_{\text{alloc}}(\alpha, \sigma) &= \text{Some}(\ell, \sigma[\ell \mapsto \alpha]) \quad \text{where } \ell \text{ is the smallest location not present in } \sigma \\ r_{\text{read}}(\ell, \sigma) &= \begin{cases} \text{Some}(\alpha, \sigma) & \text{if } \sigma(\ell) = \alpha \\ \text{None} & \text{otherwise} \end{cases} \\ r_{\text{write}}((\ell, \beta), \sigma) &= \begin{cases} \text{Some}(), \sigma[\ell \mapsto \beta] & \text{if } \sigma(\ell) \text{ is defined} \\ \text{None} & \text{otherwise} \end{cases} \\ r_{\text{dealloc}}(\ell, \sigma) &= \begin{cases} \text{Some}(), \sigma \setminus \{\ell\} & \text{if } \sigma(\ell) \text{ is defined} \\ \text{None} & \text{otherwise} \end{cases} \end{aligned}$$

From reification to reductions. Using the reification function we can formulate the reduction relation on interaction trees. The *internal* reduction relation $\rightsquigarrow: (\mathbf{IT} \times \mathbf{State}) \rightarrow (\mathbf{IT} \times \mathbf{State}) \rightarrow \mathbf{iProp}$:

$$(\alpha, \sigma) \rightsquigarrow (\beta, \sigma') \triangleq (\alpha = \text{Tick}(\beta) \wedge \sigma = \sigma') \vee (\exists i x k. \alpha = \text{Vis}_i(x, k) \wedge \text{reify}(\alpha, \sigma) = (\text{Tick}(\beta), \sigma'))$$

Intuitively, a reduction of GITrees corresponds to either stripping away one computational step, or to reifying an effect. We consider an (annotated) transitive closure of the reduction relation:

$$\begin{aligned} (\alpha, \sigma) \rightsquigarrow^0 (\beta, \sigma') &\triangleq \alpha = \beta \wedge \sigma = \sigma' \\ (\alpha, \sigma) \rightsquigarrow^{n+1} (\beta, \sigma') &\triangleq \exists \alpha_0, \sigma_0. (\alpha, \sigma) \rightsquigarrow (\alpha_0, \sigma_0) \wedge (\alpha_0, \sigma_0) \rightsquigarrow^n (\beta, \sigma') \end{aligned}$$

We write \rightsquigarrow^* for the reflexive transitive closure of the reduction relation.

Reductions and homomorphisms. Homomorphisms (Definition 3.4) play an important role in the reduction relation, allowing us to compute the reductions more easily. Specifically, homomorphisms preserve and reflect reductions:

LEMMA 4.3. *Let f be a homomorphism. Then $(\alpha, \sigma) \rightsquigarrow (\beta, \sigma')$ implies $(f(\alpha), \sigma) \rightsquigarrow (f(\beta), \sigma')$.*

LEMMA 4.4. *Let f be a homomorphism. If $(f(\alpha), \sigma) \rightsquigarrow (\beta', \sigma')$ then either*

- α is a GITree-value, or;
- there exists β such that $(\alpha, \sigma) \rightsquigarrow (\beta, \sigma')$ and $\triangleright(f(\beta) = \beta')$.

These two lemmas suggest that homomorphisms play the role of evaluation contexts within the reduction relation \rightsquigarrow . For example, if $(\alpha ; \beta, \sigma) \rightsquigarrow (\delta, \sigma')$, then either α was a value, or $(\alpha, \sigma) \rightsquigarrow (\alpha', \sigma')$ and $\triangleright(\alpha' ; \beta = \delta)$.

Continuation-independent reifiers. The reifiers that we consider here produce an output based on the input, but do not have direct access to the continuation. The reification function `reify` just calls the continuation with the produced output. This *continuation-independence* is crucial for proving Lemma 4.4 (and the associated rule `WP-HOM` in separation logic in Section 6). Not all effects are continuation-independent, for example `call/cc` cannot be implemented this way. In this paper, just like in [Xia et al. 2019], we stick to working with continuation-independent effects, as it simplifies the separation logic and the reasoning principles, and we defer studying continuation-dependent effects to future work.

5 MODELING A HIGHER-ORDER EFFECTFUL PROGRAMMING LANGUAGE

In this section we show how guarded interaction trees provide a model for a programming language with recursion, higher-order functions, and effects. Specifically, we study a PCF-like higher-order programming language with input/output effects, give its interpretation into \mathbf{IT}_{io} (see Examples 3.1 and 4.1), and show its soundness, i.e., that the interpretation agrees with the operational semantics. The same approach applies to other classes of effects for which you can write operational semantics.

Syntax and operational semantics. The syntax for the programming language, which we dub $\lambda_{\text{rec},io}$, consists of values and expressions:

$$\begin{aligned} v \in \mathbf{Val} &::= n \mid \text{rec } f(x) = e \\ e \in \mathbf{Expr} &::= x \mid v \mid \text{if } e \text{ then } e_1 \text{ else } e_2 \mid e_1(e_2) \mid e_1 + e_2 \mid e_1 - e_2 \mid \text{input} \mid \text{output}(e) \end{aligned}$$

where n ranges over the set of natural numbers, and f, x range over the set \mathbf{Var} of variables.

The operational semantics for $\lambda_{\text{rec},io}$ is given in Figure 4 as a small-step reduction relation on the configurations $\mathbf{Expr} \times \mathbf{State}$, where \mathbf{State} is a pair of lists as in Example 4.1. The reductions are

$$\begin{array}{c}
\text{RED-BETA} \\
((\text{rec } f(x) = e) v, \sigma) \rightarrow (e[v/x][\text{rec } f(x) = e/f], \sigma) \\
\\
\text{RED-IF-FALSE} \\
(\text{if } 0 \text{ then } e_1 \text{ else } e_2, \sigma) \rightarrow (e_2, \sigma) \\
\\
\text{RED-OUTPUT} \\
(\text{output}(m), (\vec{n}, \vec{m})) \rightarrow (0, (\vec{n}, m' \vec{m})) \\
\\
\text{RED-NATOP} \\
\frac{n_1, n_2 \in \mathbb{N} \quad \oplus \in \{+, -, \times, \dots\} \quad n_1 \oplus n_2 = n}{(n_1 \oplus n_2, \sigma) \rightarrow (n, \sigma)} \\
\\
\text{RED-IF-TRUE} \\
\frac{n \in \mathbb{N} \quad n > 0}{(\text{if } n \text{ then } e_1 \text{ else } e_2, \sigma) \rightarrow (e_1, \sigma)} \\
\\
\text{RED-INPUT} \\
(\text{input}, (n' \vec{n}, \vec{m})) \rightarrow (n', (\vec{n}, \vec{m})) \\
\\
\text{RED-ECTX} \\
\frac{(e_1, \sigma_1) \rightarrow (e_2, \sigma_2)}{(K[e_1], \sigma_1) \rightarrow (K[e_2], \sigma_2)}
\end{array}$$

Fig. 4. Small-step operational semantics for $\lambda_{\text{rec}, \text{io}}$.

$$\begin{array}{c}
\llbracket x \rrbracket_\rho = \rho(x) \quad \llbracket n \rrbracket_\rho = \text{Ret}(n) \quad \llbracket \text{if } e \text{ then } e_1 \text{ else } e_2 \rrbracket_\rho = \text{If}(\llbracket e \rrbracket_\rho, \llbracket e_1 \rrbracket_\rho, \llbracket e_2 \rrbracket_\rho) \\
\\
\frac{\oplus \in \{+, -, \times, \dots\}}{\llbracket e_1 \oplus e_2 \rrbracket_\rho = \text{NatOp}_\oplus(\llbracket e_1 \rrbracket_\rho, \llbracket e_2 \rrbracket_\rho)} \quad \llbracket \text{input} \rrbracket_\rho = \text{Input} \quad \llbracket \text{output}(e) \rrbracket_\rho = \text{get_nat}(\llbracket e \rrbracket_\rho, \text{Output}) \\
\\
\llbracket e_1 e_2 \rrbracket_\rho = \llbracket e_1 \rrbracket_\rho \bullet \llbracket e_2 \rrbracket_\rho \quad \llbracket \text{rec } f(x) = e \rrbracket_\rho = \text{fix}_{\text{IT}}(\lambda(t : \blacktriangleright \text{IT}). \text{Fun}(\blacktriangleright(\lambda\alpha v. \llbracket e \rrbracket_{\rho[x \mapsto v][f \mapsto \alpha]}(t))))).
\end{array}$$

Fig. 5. Semantic interpretation for $\lambda_{\text{rec}, \text{io}}$.

defined, following [Felleisen and Hieb 1992], using *evaluation contexts* $K \in \text{Ectx}$, given as:

$$K \in \text{Ectx} ::= [\cdot] \mid \text{output}(K) \mid \text{if } K \text{ then } e_1 \text{ else } e_2 \mid e K \mid K v \mid e \oplus K \mid K \oplus v$$

By $K[e]$ we denote the result of replacing the hole $[\cdot]$ in the context K with the expression e . The evaluation contexts ensure the call-by-value right-to-left evaluation order of $\lambda_{\text{rec}, \text{io}}$, as having a predefined evaluation order is important in the presence of effects.

Interpretation in guarded interaction trees. We will interpret a closed program e as an interaction tree $\llbracket e \rrbracket : \text{IT}_{\text{io}}(A)$. The effects io are those of Examples 3.1 and 4.1, and we assume that the ground type A is “large enough” to have natural numbers. For convenience, we drop the ground type and write simply IT_{io} for $\text{IT}_{\text{io}}(A)$.

In order to provide a (compositional) denotational semantics we need to provide an interpretation not only for closed terms, but for open terms as well. Given a set $\text{fv}(e) = \{x_1, \dots, x_n\}$ of free variables of e , we define the interpretation $\llbracket e \rrbracket_\rho : \text{IT}_{\text{io}}$, where ρ maps the free variables of e to interaction trees. The interpretation function is defined in Figure 5. The definition follows the standard notion of semantics for (untyped) λ -calculus, adjusted for effects and explicit recursion. The interpretation of recursive functions $\text{rec } f(x) = e$ is defined using the guarded fixed pointed operation $\text{fix}_{\text{IT}} : (\blacktriangleright \text{IT} \rightarrow \text{IT}) \rightarrow \text{IT}$, and satisfies the following equality:

$$\llbracket \text{rec } f(x) = e \rrbracket_\rho = \text{Fun}(\text{next}(\lambda v. \llbracket e \rrbracket_{\rho[x \mapsto v][f \mapsto \llbracket \text{rec } f(x) = e \rrbracket_\rho]}))$$

We show that the interpretation is sound:

THEOREM 5.1 (SOUNDNESS). *If $(e_1, \sigma_1) \rightarrow (e_2, \sigma_2)$, then $(\llbracket e_1 \rrbracket, \sigma_1) \rightsquigarrow^* (\llbracket e_2 \rrbracket, \sigma_2)$.*

We prove Theorem 5.1 by induction on the \rightarrow -derivation. The most interesting cases are for the reductions RED-BETA and RED-ECTX, which we now sketch. For the former, we need a substitution lemma, and for the latter we need to extend the interpretation to evaluation contexts.

$$\begin{array}{c}
\text{WP-VAL} \\
\frac{\alpha \in \mathbf{IT}^v \quad \Phi(\alpha)}{\text{wp } \alpha \{ \Phi \}}
\end{array}
\qquad
\begin{array}{c}
\text{WP-TICK} \\
\frac{\triangleright \text{wp } \alpha \{ \Phi \}}{\text{wp Tick}(\alpha) \{ \Phi \}}
\end{array}
\qquad
\begin{array}{c}
\text{WP-HOM} \\
\frac{f \in \text{Hom} \quad \text{wp } \alpha \{ \beta_v. \text{wp } f(\beta_v) \{ \Phi \} \}}{\text{wp } f(\alpha) \{ \Phi \}}
\end{array}$$

$$\begin{array}{c}
\text{WP-REIFY} \\
\frac{\text{has_state}(\sigma) \quad \text{reify}(\text{Vis}_i(x, k), \sigma) = (\text{Tick}(\beta), \sigma') \quad \triangleright (\text{has_state}(\sigma') * \text{wp } \beta \{ \Phi \})}{\text{wp Vis}_i(x, k) \{ \Phi \}}
\end{array}$$

$$\begin{array}{c}
\text{WP-UPD} \\
\frac{\models \text{wp } \alpha \{ \alpha_v. \models \Phi(\alpha_v) \}}{\text{wp } \alpha \{ \Phi \}}
\end{array}
\qquad
\begin{array}{c}
\text{WP-MONO} \\
\frac{\text{wp } \alpha \{ \Psi \} \quad \forall \alpha_v. \Psi(\alpha_v) * \Phi(\alpha_v)}{\text{wp } \alpha \{ \Phi \}}
\end{array}
\qquad
\begin{array}{c}
\text{WP-LAM} \\
\frac{\text{wp } \beta \{ \beta_v. \triangleright \text{wp } f(\beta_v) \{ \Phi \} \}}{\text{wp Fun}(\text{next}(f)) \bullet \beta \{ \Phi \}}
\end{array}$$

Fig. 6. Selected weakest precondition calculus rules.

LEMMA 5.2 (SUBSTITUTION LEMMA). *For any expression e with a free variable x we have*

$$\llbracket e[e'/x] \rrbracket_\rho = \llbracket e \rrbracket_{\rho[x \mapsto \llbracket e' \rrbracket]}.$$

PROOF. By induction on e , using Löb induction in the case of recursive functions. \square

In order to handle **RED-ECTX** we provide the following auxiliary interpretation for evaluation contexts. Each evaluation context K is interpreted as a *homomorphism* $\llbracket K \rrbracket_\rho : \mathbf{IT} \rightarrow \mathbf{IT}$, such that $\llbracket K[e] \rrbracket_\rho = \llbracket K \rrbracket_\rho(\llbracket e \rrbracket_\rho)$, which together with Lemma 4.3 implies the soundness of the **RED-ECTX** reduction.

6 SEPARATION LOGIC OVER GITREES

In this section we define a separation logic as a program logic for guarded interaction trees. We define a proposition $\text{wp } \alpha \{ \Phi \}$ to denote that an interaction tree α is safe to reduce, and if α reduces to an interaction tree value β_v , then β_v satisfies the postcondition $\Phi : \mathbf{IT}^v \rightarrow \text{iProp}$.

In this section we make use of the separation logic connectives of Iris, which we recall here:

$$P ::= \dots \mid P * P \mid P * P \mid \models P \mid \Box P \mid \overline{P} \mid \dots$$

For brevity, we only briefly recall the intuitive reading of these propositions and refer to [Jung et al. 2018] for details. The proposition $P * Q$ says that the propositions P and Q hold over disjoint resources; the proposition $P * Q$ says that if we were to add any resources which satisfy P , then Q would be satisfied. The proposition $\models P$ says that the current resources can be updated to satisfy P . The proposition $\Box P$ states that P holds *persistently*, i.e., without asserting any resources. Crucially, such propositions can be duplicated: $\Box P \vdash \Box P * \Box P$. An example of a persistent proposition is the invariant proposition \overline{P} , which satisfies $\overline{P} \vdash \Box \overline{P}$.

Selected rules for the weakest precondition proposition $\text{wp } \alpha \{ \Phi \}$ are given in Figure 6. The rule **WP-VAL** states that to verify a value it suffices to check that the value satisfies the postcondition. The rule **WP-TICK** states that in order to verify $\text{Tick}(\alpha)$ it suffices to verify α , under a later \triangleright . The rule **WP-HOM** states that in order to verify $f(\alpha)$ for a homomorphism f , it suffices to reduce α to a value α_v , and then verify $f(\alpha_v)$.

The rule **WP-REIFY** tells us how to deal with effects. The rule uses the proposition $\text{has_state}(\sigma)$ which signifies the *exclusive ownership* of the current state σ . The use of separation logic is crucial in this case, as we do not want to allow duplicating that proposition. The rule then states that in order to verify an effect, one has to provide the current state σ' and the proof that the interaction tree with the effect reifies into some $\text{Tick}(\beta)$. Then, the user has to verify that the resulting β reduces to a value satisfying the postcondition, under the assumption that the state has been updated to σ' .

The rule **WP-UPD** states that one can update ghost resources before and after reducing α . The rule **WP-MONO** states that one can always weaken the postcondition in $\text{wp } \alpha \{ \Phi \}$. Finally, **WP-LAM** is an example of a derived rule. It combines the computational rule for function application of GITrees, and rules **WP-HOM** and **WP-TICK**. Let us look at an example derivation using these rules.

Example 6.1. Consider a $\lambda_{\text{rec,io}}$ expression $(\text{input} + 1)$. It is interpreted as the GITree $\llbracket \text{input} + 1 \rrbracket = \text{NatOp}_+(\text{Input}, \text{Ret}(1))$, for which we can prove the following specification:

$$\frac{\text{has_state}(\vec{n}, \vec{m}) \quad \triangleright (\text{has_state}(\vec{n}, \vec{m}) \multimap \Phi(\text{Ret}(n+1)))}{\text{wp NatOp}_+(\text{Input}, \text{Ret}(1)) \{ \Phi \}}$$

PROOF. Note that $\text{NatOp}_+(-, \text{Ret}(1))$ is a homomorphism. We apply **WP-HOM**, reducing our goal to:

$$\text{wp Input } \{ \beta_v. \text{wp NatOp}_+(\beta_v, \text{Ret}(1)) \{ \Phi \} \}.$$

At this point we can use the assumption $\text{has_state}(\vec{n}, \vec{m})$ and the rule **WP-REIFY**. By the reifier of input/output effects, $\text{reify}(\text{Input}, (\vec{n}, \vec{m})) = \text{Some}(\text{Tick}(\text{Ret}(n), (\vec{n}, \vec{m})))$, and we get the following goal:

$$\triangleright (\text{has_state}(\vec{n}, \vec{m}) \multimap \text{wp Ret}(n) \{ \beta_v. \text{wp NatOp}_+(\beta_v, \text{Ret}(1)) \{ \Phi \} \}).$$

Recall that we still have the assumption $\triangleright (\text{has_state}(\vec{n}, \vec{m}) \multimap \Phi(\text{Ret}(n+1)))$. By the monotonicity of \triangleright we can remove the \triangleright modality both from the goal and the assumption. Since $\text{Ret}(n)$ is a GITree value, we can use **WP-VAL** and reduce the goal to

$$\text{wp NatOp}_+(\text{Ret}(n), \text{Ret}(1)) \{ \Phi \}.$$

By calculation, $\text{NatOp}_+(\text{Ret}(n), \text{Ret}(1)) = \text{Ret}(n+1)$, which is also a GITree value. We can then apply **WP-VAL** again to reduce the goal to $\Phi(\text{Ret}(n+1))$, which follows from the assumption. \square

We define the weakest precondition as a guarded recursive predicate, as is standard in Iris. The weakest precondition then satisfies the following *adequacy* and *safety* theorem, the proof of which relies on the adequacy of Iris ([Theorem 2.1](#)).

THEOREM 6.2. *Let α be an interaction tree and σ be a state such that*

$$\text{has_state}(\sigma) \vdash \text{wp } \alpha \{ \Phi \}$$

is derivable for some meta-level predicate Φ (containing only intuitionistic logic connectives). Then for any β and σ' such that $(\alpha, \sigma) \rightsquigarrow^ (\beta, \sigma')$, one of the following two things hold:*

- (adequacy) *either $\beta \in \text{IT}^\circ$, and $\Phi(\beta)$ holds in the meta-logic;*
- (safety) *or there are β_1 and σ_1 such that $(\beta, \sigma') \rightsquigarrow (\beta_1, \sigma_1)$*

In particular, safety implies that $\beta \neq \text{Err}(e)$ for any error $e \in \text{Error}$.³

Finally, it is worth noting that separation logic/Iris is useful for reasoning about higher-order GITrees even in the absence of effects, as demonstrated by the following example.

Example 6.3. Using guarded recursion, we can write down a GITree lter that satisfies the equation:

$$\text{lter} \bullet f \bullet \alpha \bullet \beta = \text{lf}(\alpha, f \bullet (\text{lter} \bullet f \bullet \text{NatOp}_-(\alpha, \text{Ret}(1)) \bullet \beta), \beta).$$

That is, $\text{lter} \bullet f \bullet \text{Ret}(n) \bullet \beta$ computes the iterated application $f^n \bullet \beta$. We can give lter the following higher-order specification:

$$\frac{\text{wp } \beta \{ \Psi \} \quad \square \forall \beta_v. \Psi(\beta_v) \multimap \text{wp } f \bullet \beta_v \{ \Psi \}}{\text{wp } (\text{lter} \bullet f \bullet \text{Ret}(n) \bullet \beta) \{ \Psi \}}$$

³While the weakest precondition that we presented in this section disallow any errors in guarded interaction tree, we will consider in [Section 9.3](#) that allows *some* errors, at the user's discretion.

The specification says that if β initially satisfies Ψ and f preserves Ψ , then $\text{Iter} \bullet f \bullet \text{Ret}(n) \bullet \beta$ will also satisfy Ψ . The second premise is the specification of f , and it can be used multiple times in the proof. For that reason the that premise is behind the persistently modality \triangleright .

It is also worth noting that while Iter itself does not use state, the function f that we supply to it might as well use all sorts of effects internally, and our implementation and specification of Iter is oblivious to that.

6.1 Domain-Specific Logic for Higher-Order Store

Now we show how we can use the standard mechanisms in Iris to recover a fairly standard-looking separation logic for a programming language with references from the weakest precondition calculus presented above. We use Iris's notion of *higher-order ghost state* [Jung et al. 2016, 2018] to provide the following logical interface for the higher-order store operations:

$$\begin{array}{c}
 \text{WP-ALLOC} \\
 \text{heap_ctx} \quad \triangleright \forall \ell. \ell \mapsto \alpha \ast \text{wp } k \ell \{ \Phi \} \\
 \hline
 \text{wp Alloc}(\alpha, k) \{ \Phi \}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{WP-READ} \\
 \text{heap_ctx} \quad \triangleright \ell \mapsto \alpha \quad \triangleright (\ell \mapsto \alpha \ast \text{wp } \alpha \{ \Phi \}) \\
 \hline
 \text{wp Read}(\ell) \{ \Phi \}
 \end{array}$$

$$\begin{array}{c}
 \text{WP-WRITE} \\
 \text{heap_ctx} \quad \triangleright \ell \mapsto \alpha \quad \triangleright (\ell \mapsto \beta \ast \Phi(\text{Ret}(()))) \\
 \hline
 \text{wp Write}(\ell, \beta) \{ \Phi \}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{WP-DEALLOC} \\
 \text{heap_ctx} \quad \triangleright \ell \mapsto \alpha \quad \triangleright \Phi(\text{Ret}(())) \\
 \hline
 \text{wp Dealloc}(\ell) \{ \Phi \}
 \end{array}$$

$$\text{heap_ctx} \vdash \Box \text{heap_ctx}$$

Here heap_ctx is a persistent proposition, which is part of the logical interface.

Thus our goal is to provide definitions of heap_ctx and $\ell \mapsto \alpha$ that allow us to derive the rules above. The main challenge is that the resource $\text{has_state}(\sigma)$ provides a singular complete view of the state, without the ability to split it into local portions corresponding to individual locations. That has_state is not splittable by itself is not surprising – it is an abstract representation of an arbitrary state for arbitrary effects, and there is no a priori way of splitting it. However, for our specific state (a heap $\text{Loc} \xrightarrow{\text{fin}} \blacktriangleright \text{IT}$) we know how to do the splitting. What we need to do is to provide an alternative view of the state, amenable to splitting, and tie it together with the actual state of the has_state predicate.

Our first step is then to provide a *resource algebra* for this view of the state. Following the standard practice of Iris, we use an authoritative resource algebra of the heap. It contains two kinds of resources: the “full heap” $\bullet \sigma$ and the “fragmental heap” $\circ \sigma'$. The fragmental heap is guaranteed to be a subheap of the full one $\sigma' \subseteq \sigma$. Then we make the following definitions:

$$\text{heap_ctx} \triangleq \boxed{\exists \sigma. \text{has_state}(\sigma) \ast \boxed{\bullet \sigma}} \qquad \ell \mapsto \alpha \triangleq \boxed{\circ \boxed{\ell \mapsto \text{next}(\alpha)}}$$

The heap_ctx predicate is an invariant that says that the full view of the heap coincides with the actual state that we have as part of has_state , and the points-to predicate $\ell \mapsto \alpha$ states that $\boxed{\ell \mapsto \text{next}(\alpha)}$ is in the fragmental view of the heap. Together those predicates imply that the actual state σ maps ℓ to $\text{next}(\alpha)$, which is precisely what allows us to deduce the rules **WP-ALLOC**, **WP-READ** and **WP-WRITE** from the rule **WP-REIFY**.

As a simple example, we can use the rules for the store operations to verify the factorial program from [Example 3.5](#). We show the following specification: $\text{wp fact}(n) \{ \beta_\sigma. \beta_\sigma = \text{Ret}(!n) \}$. For this, we will use the intermediate lemma:

LEMMA 6.4. *Under the assumptions heap_ctx , $\text{acc} \mapsto \text{Ret}(m)$ and $\ell \mapsto \text{Ret}(n)$, we have*

$$\text{wp factBody}(\text{acc}, \ell) \{ _ . \text{acc} \mapsto \text{Ret}(m \times !n) \}.$$

$$\begin{array}{c}
x : \tau, \Gamma \vdash x : \tau \qquad \Gamma \vdash n : \text{Nat} \qquad \frac{\Gamma \vdash e_1 : \text{Nat} \quad \Gamma \vdash e_2 : \text{Nat} \quad \oplus \in \{-, +\}}{\Gamma \vdash e_1 \oplus e_2 : \text{Nat}} \\
\\
\frac{\Gamma \vdash e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash e_2 : \tau_1}{\Gamma \vdash e_1 e_2 : \tau_2} \qquad \frac{f : \tau_1 \rightarrow \tau_2, x : \tau_1, \Gamma \vdash e : \tau_1 \rightarrow \tau_2}{\Gamma \vdash \text{rec } f(x) = e : \tau_1 \rightarrow \tau_2} \\
\\
\frac{\Gamma \vdash e : \text{Nat} \quad \Gamma \vdash e_1 : \tau \quad \Gamma \vdash e_2 : \tau}{\Gamma \vdash \text{if } e \text{ then } e_1 \text{ else } e_2 : \tau} \qquad \Gamma \vdash \text{input} : \text{Nat} \qquad \frac{\Gamma \vdash e : \text{Nat}}{\Gamma \vdash \text{output}(e) : \text{Nat}}
\end{array}$$

Fig. 7. Typing rules for $\lambda_{\text{rec},\text{io}}$.

PROPOSITION 6.5. *Under the assumption heap_ctx we have*

$$\text{wp fact}(n) \{ \beta_v. \beta_v = \text{Ret}(!n) \}.$$

PROOF. We proceed by allocating the locations acc and ℓ symbolically using **WP-ALLOC**, and then appeal to [Lemma 6.4](#). \square

As one can see, the logic that we recovered for the higher-order store effects is very close to a normal separation logic one would normally see for a programming language with a heap [[Jung et al. 2018](#)]. Our logic, however, is amenable to extensions with other effects and programming language constructs. Indeed, we explain how to obtain a logic for reasoning about different combined effects in [Section 8](#). In the next section we show how to apply the separation logic to show computational adequacy of the model of $\lambda_{\text{rec},\text{io}}$.

7 COMPUTATIONAL ADEQUACY FOR $\lambda_{\text{rec},\text{io}}$

In [Section 5](#) we constructed a compositional model of $\lambda_{\text{rec},\text{io}}$ in guarded interaction trees and proved that it is sound: if a $\lambda_{\text{rec},\text{io}}$ program e terminates to a natural number n , then $\llbracket e \rrbracket$ terminates to $\text{Ret}(n)$. In this section we show the other direction, known as computational adequacy in domain theory [[Plotkin 1977](#)], for the well-typed fragment of $\lambda_{\text{rec},\text{io}}$; the typing relation $(\Gamma \vdash e : \tau)$ is given in [Figure 7](#). Computational adequacy is formally stated as the following theorem:

THEOREM 7.1 (ADEQUACY). *If $\vdash e : \text{Nat}$ and $(\llbracket e \rrbracket, \sigma) \rightsquigarrow^* (\text{Ret}(n), \sigma')$ then $(e, \sigma) \rightarrow^* (n, \sigma')$.*

Computational adequacy is usually proved using logical relations between the syntax (terms of $\lambda_{\text{rec},\text{io}}$ in our case) and semantics (guarded interaction trees in our case). Here we follow the recent practice [[Krebbbers et al. 2017b](#)] of using the separation logic (see [Section 6](#)) to define our logical relations model.

We define a logical relation $\Gamma \models \alpha \lesssim e : \tau$, relating a guarded interaction tree α and an expression e . Here, e is an open expression for which we have $\Gamma \vdash e : \tau$ while α is “an open interaction tree”, i.e., a function of type $(\text{fv}(\Gamma) \rightarrow \text{IT}) \rightarrow \text{IT}$. As usual, we first define the relation over closed GITrees and expressions, and then generalize it to the open case. The logical relation, given in [Figure 8](#), is, as usual for call-by-value languages, is decomposed into an expression relation $\mathcal{E}[\llbracket \tau \rrbracket]$ and a value relation $\mathcal{V}[\llbracket \tau \rrbracket]$. The expression relation simply states that related expressions should produce related values. The value relation is defined by induction on the type in the standard way: values of base types should be equal while functions take related values to related expressions. As values, once computed, can be used multiple times (cf. the logical relation in [Section 9.1](#)) the value relation is required to be persistent; hence the persistently modality \square in the value relation for functions. In order to define the relation on open terms we define a relation for typing contexts $\mathcal{V}^*[\llbracket \Gamma \rrbracket]$ which relates, point-wise, two substitutions respectively of the types $\text{fv}(\Gamma) \rightarrow \text{IT}^v$ and $\text{fv}(\Gamma) \rightarrow \text{Val}$.

$$\begin{aligned}
\mathcal{E}[\tau](\alpha, e) &\triangleq \forall \sigma. \text{has_state}(\sigma) \multimap \text{wp } \alpha \{ \beta_v. \exists v, \sigma'. (e, \sigma) \rightarrow^* (v, \sigma') \multimap \mathcal{V}[\tau](\beta_v, v) \multimap \text{has_state}(\sigma') \} \\
\mathcal{V}[\text{Nat}](\beta_v, v) &\triangleq \exists n \in \mathbb{N}. \beta_v = \text{Ret}(n) \wedge v = n \\
\mathcal{V}[\tau_1 \rightarrow \tau_2](\beta_v, v) &\triangleq \exists f. \beta_v = \text{Fun}(f) \wedge \square(\forall \alpha_w, w. \mathcal{V}[\tau_1](\alpha_w, w) \multimap \mathcal{E}[\tau_2](\beta_v \bullet \alpha_w, v w)). \\
\mathcal{V}^*[\Gamma](\rho_1, \rho_2) &\triangleq \forall (x : \tau) \in \Gamma. \mathcal{V}[\tau](\rho_1(x), \rho_2(x)) \\
\Gamma \models \alpha \lesssim e : \tau &\triangleq \forall \rho_1, \rho_2. \mathcal{V}^*[\Gamma](\rho_1, \rho_2) \implies \mathcal{E}[\tau](\alpha(\rho_1), e[\rho_2])
\end{aligned}$$

Fig. 8. Logical relation for $\lambda_{\text{rec}, \text{io}}$.

LEMMA 7.2 (FUNDAMENTAL PROPERTY). *For any $\Gamma \vdash e : \tau$, we have $\Gamma \models (\lambda \rho. \llbracket e \rrbracket_\rho) \lesssim e : \tau$.*

Computational adequacy follows from the fundamental property together with the following Lemma which itself is a consequence of the soundness of the weakest precondition calculus (Theorem 6.2):

LEMMA 7.3. *Suppose that $\models \alpha \lesssim e : \text{Nat}$. Then for any state σ ,*

- *if $(\alpha, \sigma) \rightsquigarrow^* (\text{Ret}(n), \sigma')$, then $(e, \sigma) \rightarrow^* (n, \sigma')$;*
- *if $(\alpha, \sigma) \rightsquigarrow^* (\beta, \sigma')$, then $\beta \neq \perp$.*

8 MODULAR REASONING ABOUT COMBINATIONS OF EFFECTS

Because (guarded) interaction trees define effects abstractly, one of the main advantages is the ability to combine programs with different effects modularly in the same setting. In this section we demonstrate how we achieve this for guarded interaction trees.

Given two signatures E and F , with indexing sets I and J , we say that E is a subsignature of F , written as $E \succrightarrow F$, if there is a mapping $\epsilon : I \rightarrow J$ such that $E.\text{Ins}_i(X) \simeq F.\text{Ins}_{\epsilon(i)}(X)$ and $E.\text{Outs}_i(X) \simeq F.\text{Outs}_{\epsilon(i)}(X)$ for any $i \in I$ and for any type X . Here, \simeq stands for isomorphism of types.

In regular interaction trees, a subsignature $E \succrightarrow F$ induces an embedding $\text{IT}_E \rightarrow \text{IT}_F$ of interaction trees. However, such an embedding is not possible for guarded interaction trees due to the mixed-variance definition: a function $\text{IT}_E \rightarrow \text{IT}_E$ cannot be converted to a function $\text{IT}_F \rightarrow \text{IT}_F$ which takes a guarded interaction tree with a larger set of effects.

To achieve modularity we will instead work with an open-ended collection of effects which is large enough to embed all the effects that we need. It is only at the “top-level”, e.g., when applying the adequacy theorem, that we pick a concrete signature of effects. For example, the precise type of the `Alloc` function from Example 3.2 is the following type, for any F such that $E_{\text{store}} \succrightarrow F$:

$$\text{Alloc} : \text{IT}_F \times (\text{Loc} \rightarrow \text{IT}_F) \rightarrow \text{IT}_F$$

(In practice, the function `Alloc` is polymorphic in F at the meta-level, i.e., in Coq.) With this, we can easily combine two programs with different collections of effects, assuming both of the programs are written in such an open ended manner; we just need to pick F to be large enough to embed the effects of both programs. For example, we can combine the factorial implementation from Example 3.5 with input/output effects, to write a program that takes a natural number from the input, computes its factorial, and prints the result to the output:

$$\text{fact_io} \triangleq \text{get_nat}(\text{get_nat}(\text{Input}, \text{fact}), \text{Output}).$$

The resulting program `fact_io` has type IT_F for any F such that $E_{\text{store}} \succrightarrow F$ and $E_{\text{io}} \succrightarrow F$.

Reifiers for modular effects. Writing down programs with modular combinations of effects is not enough by itself: we also want to reason about the reification of effects modularly. Suppose we write a program with effects E as an $\text{GITree } \mathbf{IT}_F$ with $E \mapsto F$, and suppose that we have a reifier for E . Recall that we defined a reifier for the effects E to be a tuple $(E, \text{State}, r : \prod_{i \in E} \text{Ins}_i(\blacktriangleright \mathbf{IT}_E) \times \text{State} \rightarrow \text{option}(\text{Outs}_i(\blacktriangleright \mathbf{IT}_E) \times \text{State}))$. However, if the state itself includes interaction trees, as in [Example 4.2](#), we need also to make the state and the reifier parametric in the effects. Therefore, instead of a fixed type State we consider a family of states $\text{State}(X)$, and instead a single reifier function r we consider a family of functions

$$\forall X. \prod_{i \in E} \text{Ins}_i(X) \times \text{State}(X) \rightarrow \text{option}(\text{Outs}_i(X) \times \text{State}(X)).$$

In practice, we assume that the global state is the product of states of reifiers for sub-effects, in which each sub-effect acts only on its own part of the state. Concretely, given a sequence $\vec{R} = (E_1, \text{State}_1, r_1), \dots, (E_m, \text{State}_m, r_m)$ of reifiers we define the *global reifier* $R_G = (G, \text{State}(-), r)$:

$$\begin{aligned} G.I &\triangleq \sum_{1 \leq i \leq m} E_i.I & G.\text{Ins}_{(i,j)}(X) &\triangleq E_i.\text{Ins}_j(X) & G.\text{Outs}_{(i,j)}(X) &\triangleq E_i.\text{Outs}_j(X) \\ \text{State}(X) &\triangleq \prod_{1 \leq i \leq m} \text{State}_i(X) \\ r_{X,(i,j)}(x, (\sigma_1, \dots, \sigma_i, \dots, \sigma_m)) &\triangleq \begin{cases} \text{Some}(y, (\sigma_1, \dots, \sigma'_i, \dots, \sigma_m)) & \text{if } r_i(x, \sigma_i) = \text{Some}(y, \sigma') \\ \text{None} & \text{otherwise} \end{cases} \end{aligned}$$

Turning to the separation logic, we specialize the rule [WP-REIFY](#) to the signature G and the reifier R_G , and simplify it to

$$\frac{\text{WP-REIFY-LOCAL} \quad \text{has_state}_i(\sigma_i) \quad r_i(x, \sigma_i) = \text{Some}(y, \sigma'_i) \quad k \ y = \text{next}(\beta) \quad \triangleright (\text{has_state}_i(\sigma'_i) * \text{wp } \beta \{ \Phi \})}{\text{wp Vis}_{(i,j)}(x, k) \{ \Phi \}}$$

where the predicate $\text{has_state}_i(\sigma)$ tracks the local component of the global state associated with the i^{th} reifier. The predicates are defined to validate the following rule, which allows us to split the global state into local subcomponents and combine them back together:

$$\text{has_state}(\vec{\sigma}) \dashv\vdash \text{has_state}_1(\sigma_1) * \dots * \text{has_state}_m(\sigma_m).$$

Then to write down the abstractions for the domain-specific logic in [Section 6.1](#) we change the heap_ctx definition to link together only the state corresponding to the specific effects:

$$\text{heap_ctx} \triangleq \boxed{\exists \sigma. \text{has_state}_j(\sigma) * \boxed{\sigma}}$$

where the higher-order store reifier is the i^{th} subreifier of \vec{R} .

Example 8.1. Recall the program fact_io from the beginning of this section. We use [Proposition 6.5](#) to show the following specification:

$$\text{heap_ctx} * \text{has_state}_j(k\vec{n}, \vec{m}) \vdash \text{wp fact_io} \{ _ . \text{has_state}_j(\vec{n}, (!k)\vec{m}) \}$$

where has_state_j tracks the state of the input/output effects. The specification tells us that if we run fact_io with the starting state $(k\vec{n}, \vec{m})$ for the input/output effects, then we end up with the state $(\vec{n}, (!k)\vec{m})$ for the input/output effects.

8.1 Modular reasoning with a generic ground type

As we have mentioned in [Section 3.2](#), we often would like to work with the GITrees $\text{IT}_E(A)$ for some generic ground type A that is “large enough” to contain ground values that we need to represent (e.g. the unit type, natural numbers, the type of locations, etc). That is, we assume that the ground type A is isomorphic to a sum $1 + \text{Nat} + \text{Loc} + \dots$, depending on ground values we need. We tackle this generic ground type in a similar way we deal with different effect signatures modularly.

Specifically, we write $B \rightarrow A$ if $A \simeq B + C$ for some type C . We then have the generic return constructor $\text{Ret} : B \rightarrow \text{IT}_E(A)$ for any $B \rightarrow A$. Similarly, we have a generic “destructor” $\text{get_ret} : \text{IT}_E(A) \times (B \rightarrow \text{IT}_E(A)) \rightarrow \text{IT}_E(A)$ which allows us to extract a ground value of type B , as long as we have $B \rightarrow A$. such that get_ret is a homomorphism in the first argument, which satisfies:

$$\text{get_ret}(\text{Ret}(b), g) = \begin{cases} g(b) & \text{if } b \in B, \\ \text{Err}(\text{RunTime}) & \text{otherwise.} \end{cases}$$

Then, the get_nat function from [Section 3.2](#) is just the specialization of get_ret to the situation $\text{Nat} \rightarrow A$.

The predicate $B \rightarrow A$ is formalized in Coq as a typeclass, making it easy to use the generic operations like Ret and get_ret . In the remainder of this paper we will stick to those generalized operations, and will assume that the ground type A contains all the ground values we need.

9 TYPE SAFETY FOR CROSS-LANGUAGE INTEROPERABILITY

One of the advantages of using GITrees for denotational semantics is that it provides a common setting for interpreting and reasoning about different languages with different effects, and then combining the results in a modular manner. In this section we demonstrate this point by verifying type safety of interoperability between two different languages. The interoperability is achieved by allowing embeddings from one language into another at a particular boundary [[Matthews and Findler 2007](#)]. We take inspiration for this case study from the approach of [Patterson et al. \[2022\]](#), who consider interoperability of different languages at a level of a common third language, which both the source languages are compiled down to. The communication between the source languages is done through *glue code* at the level of the target language, which converts types from one language to another. The type safety result then states that well-typed programs in a combined language can only go wrong due to conversion errors at the boundaries.

Specifically, we first consider an affine programming language $\lambda_{\rightarrow, \text{ref}}$ with linear references and strong updates, which we interpret in guarded interaction trees using the higher-order store effects ([Example 3.2](#)). We show the type safety of $\lambda_{\rightarrow, \text{ref}}$ by building a logical relations model.

We then consider cross-language interoperability, by allowing embedding of (higher-order) programs from the non-affine language $\lambda_{\text{rec}, \text{io}}$ ([Section 5](#)) into $\lambda_{\rightarrow, \text{ref}}$, thus combining two languages with different type systems and different effects. Following the approach outlined in [Section 8](#) we *reuse the standalone interpretations* of $\lambda_{\text{rec}, \text{io}}$ and $\lambda_{\rightarrow, \text{ref}}$ to interpret the combined language in $\text{IT}_{\text{store}, \text{io}}$. Finally, we show type safety of this combined language, by building the logical relations model out of the models for the individual languages.

We stress that our approach here allows us to prove type safety of $\lambda_{\text{rec}, \text{io}}$ and $\lambda_{\rightarrow, \text{ref}}$ *separately*, and then prove type safety of the combined language by *reusing* the logical relations for the individual languages, thus highlighting the modular nature of the GITrees framework.

9.1 An Affine Programming Language

First, we consider an affine programming language $\lambda_{\rightarrow, \text{ref}}$ with references with strong updates, and show how to interpret it in GITrees in a way that enforces linearity. We then consider the

$$\begin{array}{c}
\frac{b \in \mathbb{B}}{\Omega \vdash b : \text{Bool}} \quad \frac{n \in \mathbb{N}}{\Omega \vdash n : \text{Nat}} \quad \Omega \vdash () : \text{Unit} \quad \Omega_1, a : \tau, \Omega_2 \vdash a : \tau \quad \frac{a : \tau_1, \Omega \vdash e : \tau_2}{\Omega \vdash \lambda a. e : \tau_1 \multimap \tau_2} \\
\\
\frac{\Omega_1 \vdash e_1 : \tau_1 \multimap \tau_2 \quad \Omega_2 \vdash e_2 : \tau_1}{\Omega_1, \Omega_2 \vdash e_1 e_2 : \tau_2} \quad \frac{\Omega_1 \vdash e_1 : \tau_1 \quad \Omega_2 \vdash e_2 : \tau_2}{\Omega_1, \Omega_2 \vdash (e_1, e_2) : \tau_1 \otimes \tau_2} \\
\\
\frac{\Omega_1 \vdash e_1 : \tau_1 \otimes \tau_2 \quad a_1 : \tau_1, a_2 : \tau_2, \Omega_2 \vdash e_2 : \tau}{\Omega_1, \Omega_2 \vdash \text{let } (a_1, a_2) = e_1 \text{ in } e_2 : \tau} \quad \frac{\Omega \vdash e : \tau}{\Omega \vdash \text{alloc}(e) : \text{ref } \tau} \quad \frac{\Omega \vdash e : \text{ref } \tau}{\Omega \vdash \text{dealloc}(e) : \text{Unit}} \\
\\
\frac{\Omega_1 \vdash e_1 : \text{ref } \tau_1 \quad \Omega_2 \vdash e_2 : \tau_2}{\Omega_1, \Omega_2 \vdash \text{replace}(e_1, e_2) : \tau_1 \otimes \text{ref } \tau_2}
\end{array}$$

Fig. 9. Type system for $\lambda_{\multimap, \text{ref}}$.

combination $\lambda_{\multimap, \text{ref}} + \lambda_{\text{rec}, \text{io}}$, which allows us to embed $\lambda_{\text{rec}, \text{io}}$ programs, including functions, into $\lambda_{\multimap, \text{ref}}$. The syntax for the affine language $\lambda_{\multimap, \text{ref}}$ is as follows:

$$\begin{array}{l}
\tau \in \text{Type} ::= \text{Bool} \mid \text{Nat} \mid \text{Unit} \mid \tau_1 \otimes \tau_2 \mid \tau_1 \multimap \tau_2 \mid \text{ref } \tau \\
e \in \text{Expr} ::= n \mid b \mid () \mid a \mid \lambda a. e \mid e_1 e_2 \mid (e_1, e_2) \mid \text{let } (a_1, a_2) = e_1 \text{ in } e_2 \\
\quad \mid \text{alloc}(e) \mid \text{dealloc}(e_1) \mid \text{replace}(e_1, e_2)
\end{array}$$

To differentiate between the terms of $\lambda_{\multimap, \text{ref}}$ and the terms of $\lambda_{\text{rec}, \text{io}}$, we use the **orange** color to refer to types and programs of $\lambda_{\multimap, \text{ref}}$. The type system for $\lambda_{\multimap, \text{ref}}$ is given in Figure 9. Let us explain some of the details. The language $\lambda_{\multimap, \text{ref}}$ contains Booleans, natural numbers, and the unit type. It also features linear functions $\tau_1 \multimap \tau_2$. Note that in the typing rule for function application, the context is split between typing of the function and typing of the argument. This ensures that the function and its argument do not share any variables or resources in common.

The language also features linear pairs $\tau_1 \otimes \tau_2$. In the typing rule for pairs the typing environment has to be split between the two components. This ensures that we cannot have, e.g., pairs of the form (x, x) .

Finally, the language features references with strong updates, i.e., references that can store values of different types. The constructors **alloc** and **dealloc** are used to allocate and free the references, respectively. To ensure linearity, we have a single operation that combines reading from a reference and performing a strong update. The program **replace**(r, v) reads the value that is stored in the reference r and updates it to the value v . It then returns a linear pair consisting of the old value and the reference itself, allowing one to reuse the reference later on.

The meaning of $\lambda_{\multimap, \text{ref}}$ is given by the interpretation function $\llbracket e \rrbracket_\rho : \mathbf{IT}_F(A)$, where ρ is the environment mapping the free variables of e to GITrees , and where F is a signature which contains the higher-order store effects (Example 3.2). We assume that the ground type A contains, in addition to natural numbers and the unit type, the type Loc of locations. The semantic interpretation of $\lambda_{\multimap, \text{ref}}$ is given in Figure 10. In the interpretation of compound expression we split the environment ρ into the environments ρ_1 and ρ_2 , for the free variables of e_1 and e_2 respectively.

In order to ensure that the variables from the context Ω are used at most once, we wrap every variable in a thunk which can be evaluated at most once:

$$\begin{array}{l}
\text{Thunk}(\alpha) \triangleq \text{Alloc}((0), \lambda \ell. \text{Fun}(\text{next}(\lambda _ . \text{If}(\text{Read}(\ell), \text{Err}(\text{Lin}), \text{Write}(\ell, \text{Ret}(1)) ; \alpha)))) \\
\text{Force}(\alpha) \triangleq \alpha \bullet \text{Ret}(0)
\end{array}$$

$$\begin{aligned}
\llbracket b \rrbracket_\rho &\triangleq \begin{cases} \text{Ret}(1) & \text{if } b = \text{true} \\ \text{Ret}(0) & \text{otherwise} \end{cases} & \llbracket () \rrbracket_\rho &\triangleq \text{Ret}() \\
\llbracket n \rrbracket_\rho &\triangleq \text{Ret}(n) & \llbracket \lambda a. e \rrbracket_\rho &\triangleq \text{Fun}(\text{next}(\lambda \alpha. \llbracket e \rrbracket_{\rho[a \mapsto \alpha]})) \\
\llbracket a \rrbracket_\rho &\triangleq \text{Force}(\rho(a)) & \llbracket e_1 e_2 \rrbracket_\rho &\triangleq \text{Let } x = \llbracket e_2 \rrbracket_{\rho_2} \text{ in} \\
& & & \llbracket e_1 \rrbracket_{\rho_1} \bullet \text{Thunk}(x) \\
\llbracket \text{let } (a_1, a_2) = e_1 \text{ in } e_2 \rrbracket_\rho &\triangleq \text{Let } x = \llbracket e_1 \rrbracket_{\rho_1} \text{ in} & \llbracket (e_1, e_2) \rrbracket_\rho &\triangleq (\llbracket e_1 \rrbracket_{\rho_2}, \llbracket e_2 \rrbracket_{\rho_2}) \\
& \text{Let } y = \text{Thunk}(\pi_1(x)) \text{ in} \\
& \text{Let } z = \text{Thunk}(\pi_2(x)) \text{ in} \\
& \llbracket e_2 \rrbracket_{\rho_2[a_1 \mapsto y, a_2 \mapsto z]} \\
\llbracket \text{alloc}(e) \rrbracket_\rho &\triangleq \text{Let } x = \llbracket e \rrbracket_\rho \text{ in } \text{Alloc}(x, \text{Ret}) & \llbracket \text{dealloc}(e) \rrbracket_\rho &\triangleq \text{get_ret}(\llbracket e \rrbracket_\rho, \text{Dealloc}) \\
\llbracket \text{replace}(e_1, e_2) \rrbracket_\rho &\triangleq \text{Let } y = \llbracket e_2 \rrbracket_{\rho_2} \text{ in } \text{get_ret}(\llbracket e_1 \rrbracket_{\rho_1}, \lambda \ell. \text{Let } x = \text{Read}(\ell) \text{ in} \\
& \text{Write}(\ell, y) ; (x, \text{Ret}(n)))
\end{aligned}$$

Fig. 10. Interpretation of $\lambda_{\rightarrow, \text{ref}}$.

$$\begin{aligned}
\mathcal{V}[\text{Unit}](\beta_v) &\triangleq \beta_v = \text{Ret}() & \mathcal{V}[\text{Nat}](\beta_v) &\triangleq \exists n \in \mathbb{N}. \beta_v = \text{Ret}(n) \\
\mathcal{V}[\tau_1 \rightarrow \tau_2](\beta_v) &\triangleq \forall \alpha_w. \mathcal{V}[\tau_1](\alpha_w) * \mathcal{E}[\tau_2](\beta_v \bullet \alpha_w) & \mathcal{V}[\text{Bool}](\beta_v) &\triangleq \beta_v = \text{Ret}(0) \vee \beta_v = \text{Ret}(1) \\
\mathcal{V}[\text{ref } \tau](\beta_v) &\triangleq \exists \ell \in \text{Loc}. \alpha_v. (\beta_v = \text{Ret}(\ell)) * & \mathcal{V}[\tau_1 \otimes \tau_2](\beta_v) &\triangleq \exists \gamma_v, \delta_v. \beta_v = (\gamma_v, \delta_v) * \\
& \ell \mapsto \alpha_v * \mathcal{V}[\tau](\alpha_v) & & \mathcal{V}[\tau_1](\gamma_v) * \mathcal{V}[\tau_2](\delta_v) \\
\mathcal{E}[\tau](\alpha) &\triangleq \text{heap_ctx} * \text{wp } \alpha \{ \beta_v. \mathcal{V}[\tau](\beta_v) \} & \text{protected}(\Phi)(\beta_v) &\triangleq \text{wp } \text{Force}(\beta_v) \{ \Phi \} \\
\mathcal{V}^*[\Omega](\rho) &\triangleq \forall (a : \tau) \in \Omega. \text{protected}(\mathcal{V}[\tau])(\rho(a)) & \Omega \models \alpha : \tau &\triangleq \forall \rho. \mathcal{V}^*[\Omega](\rho) \implies \mathcal{E}[\tau](\alpha(\rho))
\end{aligned}$$

Fig. 11. Logical relation for $\lambda_{\rightarrow, \text{ref}}$.

When we called a thunked GITree for the second time, it will return the error $\text{Err}(Lin)$, signifying that we broke the linearity condition. Here we assume that we have a separate error state $Lin \in \text{Error}$, because we want to treat linearity condition errors separate from type errors or memory safety errors. As such, in the interpretation of a function application we put the argument in a `Thunk`, and whenever we use the argument (or any affine variable) we then have to `Force` it.

We can show that if we have a well-typed program, then it does not have any run-time errors, and that all the thunks are evaluated at most once:

PROPOSITION 9.1. *Suppose that $\vdash e : \tau$, and suppose that $(\sigma, \llbracket e \rrbracket) \rightsquigarrow^* (\sigma', \beta)$. Then $\beta \neq \text{Err}(err)$.*

To prove [Proposition 9.1](#) we use a logical relation, given in [Figure 11](#), defined similarly to the logical relation from [Section 7](#). The interpretation $\mathcal{V}[-]$ of the base types cover the appropriate subsets of natural numbers. Reference types are interpreted using the “pointsto” $\ell \mapsto \alpha_v$ proposition, and affine pairs are interpreted component-wise. The main differences to note here are: (1) variables in $\lambda_{\rightarrow, \text{ref}}$ are interpreted as thunks, and thus we adjust the interpretation $\mathcal{V}^*[\Omega]$ to account for that; (2) values in $\lambda_{\rightarrow, \text{ref}}$ can be used at most once; hence the value interpretation is not persistent, i.e., there is no persistently modality \square in the interpretation of function types.

LEMMA 9.2 (FUNDAMENTAL PROPERTY). *For any expression e , if $\Omega \vdash e : \tau$, then $\Omega \models \lambda \rho. \llbracket \alpha \rrbracket_\rho : \tau$.*

We prove the fundamental property by induction on the typing derivation. More specifically, for each typing rule we prove an associated *compatibility lemma*, by replacing expressions with interaction trees and the derivability \vdash with validity \models . For example, the compatibility lemma for `dealloc` looks as follows:

LEMMA 9.3. *Suppose that $\Omega \models \alpha : \text{ref } \tau$. Then $\Omega \models \lambda\rho. \text{get_ret}(\alpha(\rho), \text{Dealloc}) : \text{Unit}$.*

Proving all the compatibility lemmas is relatively straightforward using the separation logic rules. Having separate compatibility lemmas will be useful for us in the next section.

By combining the fundamental property with the safety theorem for the weakest precondition calculus (Theorem 6.2) we obtain a proof of Proposition 9.1.

Safety for $\lambda_{\text{rec},\text{io}}$. Similar to the logical relation for safety for $\lambda_{\rightarrow,\text{ref}}$, we define a logical relation for $\lambda_{\text{rec},\text{io}}$. It is simply an unary version of the logical relation from Section 7. For example, the expression relation is defined as

$$\mathcal{E}[\llbracket \tau' \rrbracket](\alpha) \triangleq \forall \sigma'. \text{has_state}_i(\sigma') * \text{wp } \alpha \{ \beta_v. \exists \sigma'_1. \mathcal{V}[\llbracket \tau' \rrbracket](\beta_v) * \text{has_state}_i(\sigma'_1) \}$$

where has_state_i tracks the state for the input/output effects. We omit the other details here and direct an interested reader to the Coq formalization. We only note that just like for $\lambda_{\rightarrow,\text{ref}}$, we split the proof of the fundamental property into compatibility lemmas, which we will use in the next section.

9.2 Interoperability Between Languages

Following the approach of [Patterson et al. 2022], we allow the interaction between the languages $\lambda_{\text{rec},\text{io}}$ and $\lambda_{\rightarrow,\text{ref}}$, using the guarded interaction trees as the “common ground”, combining the effects of the two languages. At the syntactic language level, this is done by allowing one to embed the expressions of $\lambda_{\text{rec},\text{io}}$ into the $\lambda_{\rightarrow,\text{ref}}$ programs. The embedding is given by the following typing rule:⁴

$$\frac{\text{TYPED-CONV} \quad \vdash e : \tau' \quad \tau' \sim \tau}{\Omega \vdash \langle e \rangle_{\tau' \sim \tau} : \tau}$$

The crucial ingredient for the interaction is a type conversion relation $\tau' \sim \tau$ stating that the $\lambda_{\text{rec},\text{io}}$ type τ' is convertible to the $\lambda_{\rightarrow,\text{ref}}$ type τ .

We have the following type conversions:

$$\text{Nat} \sim \text{Nat} \quad \text{Nat} \sim \text{Unit} \quad \text{Nat} \sim \text{Bool} \quad \frac{\tau'_1 \sim \tau_1 \quad \tau'_2 \sim \tau_2}{(\text{Nat} \rightarrow \tau'_1) \rightarrow \tau'_2 \sim \tau_1 \multimap \tau_2}$$

The first three type conversions say that we can freely convert between integers, Booleans, and the unit type (since all of them have similar internal representation). The last conversion is more interesting. It says that we can convert between affine functions and non-affine functions. The affine argument τ_1 is represented as a thunk $(\text{Nat} \rightarrow \tau'_1)$, which we will protect at runtime to ensure that it is not invoked more than once. The type Nat is used in lieu of the unit type which is absent from $\lambda_{\text{rec},\text{io}}$.

⁴For simplicity, we only consider one-way embeddings from $\lambda_{\text{rec},\text{io}}$ to $\lambda_{\rightarrow,\text{ref}}$, and we only consider embeddings of closed terms. This simplifies the type system, but does not lead to loss of expressivity, since we allow type conversions of functions.

Glue code for conversion functions. In order to (1) convert between different types, and (2) ensure the linearity of $\lambda_{\rightarrow, \text{ref}}$ programs that might cross the boundary to $\lambda_{\text{rec}, \text{io}}$, we need to interpret embedded terms with additional *glue code*. For every type conversion $\tau' \sim \tau$ we generate recursively a pair of conversion functions $C_{\tau', \tau}$ and $C_{\tau, \tau'}$ which convert the representations from τ' to τ and vice versa. The glue code for converting between the base types ensures that the underlying natural number representation stays within the range. For example, for $\text{Nat} \sim \mathbb{B}$ we have:

$$C_{\text{Nat}, \text{Bool}}(\alpha) \triangleq \text{If}(\alpha, \text{Ret}(1), \text{Ret}(0)) \qquad C_{\text{Bool}, \text{Nat}}(\alpha) \triangleq \alpha$$

The glue code for converting functions is a bit more involved:⁵

$$\begin{aligned} C_{(\text{Nat} \rightarrow \tau'_1) \rightarrow \tau'_2, \tau_1 \rightarrow \tau_2}(\alpha) &\triangleq \text{Let } p = \alpha \text{ in} \\ &\quad \text{Fun}(\text{next}(\lambda x. \text{Let } y = C_{\tau_1, \tau'_1}(\text{Force}(x)) \text{ in } C_{\tau'_2, \tau_2}(p \bullet \text{Thunk}(y)))) \\ C_{\tau_1 \rightarrow \tau_2, (\text{Nat} \rightarrow \tau'_1) \rightarrow \tau'_2}(\alpha) &\triangleq \text{Let } p = \alpha \text{ in Let } p' = \text{Thunk}(p) \text{ in} \\ &\quad \text{Fun}(\text{next}(\lambda x. \text{Let } f = \text{Force}(p) \text{ in Let } y = C_{\tau'_1, \tau_1}(\text{Force}(x)) \text{ in } C_{\tau'_2, \tau_2}(f \bullet \text{Thunk}(y)))) \end{aligned}$$

When we convert a function, we need to recursively convert the argument and the result; in addition the argument to the function needs to be protected with a `Thunk`. Furthermore, when affine functions are converted to non-affine ones, we need to protect the function itself with a `Thunk`, to ensure that the function is not invoked multiple times. Calling an affine function multiple times might be unsound, e.g., calling the following function twice will attempt to deallocate an already deallocated reference:

$$(\lambda \ell. \lambda _ . (\lambda _ . 7) \text{dealloc}(\ell)) (\text{alloc}(42)) : \text{Unit} \rightarrow \text{Nat}.$$

Partial safety for the combined language. We interpret the combined language with embeddings using the glue code. The embedding from $\lambda_{\text{rec}, \text{io}}$ to $\lambda_{\rightarrow, \text{ref}}$ is interpreted as follows:

$$\llbracket (e)_{\tau' \sim \tau} \rrbracket_{\rho} = C_{\tau', \tau}(\text{IO}\llbracket e \rrbracket),$$

where $\text{IO}\llbracket - \rrbracket$ is the interpretation function for $\lambda_{\text{rec}, \text{io}}$ expressions from Section 5. The interpretation for all the operations, except for the embedding, remains unchanged.

Using the modular approach we described in Section 8, we interpret the extended language $\lambda_{\text{rec}, \text{io}} + \lambda_{\rightarrow, \text{ref}}$ into the guarded interaction trees $\text{IT}_{\text{store}, \text{io}}$ with the signature that combines the input/output effects and the store effects. This ensures that the interpretation of the languages end up in the same domain, where they can interact. By combining the reifiers for the effects of the individual languages we also get the reduction relation $((\sigma_1, \sigma'_1), \alpha) \rightsquigarrow ((\sigma_2, \sigma'_2), \beta)$, where σ_1 and σ_2 are stores, and σ'_1 and σ'_2 are input/output tapes.

Of course, in the combined language with conversions, our programs can actually violate the linearity condition, since it is no longer enforced by the type system. However, we can prove that linearity violations at the boundary are the only errors that we will possibly get. Thus we will to prove the following safety theorem:

THEOREM 9.4. *Suppose that $\vdash e : \tau$ with the embedding rule, and suppose that $((\sigma_1, \sigma'_1), \llbracket e \rrbracket) \rightsquigarrow^* ((\sigma_2, \sigma'_2), \beta)$. Then either β is not an error, or $\beta = \text{Err}(\text{Lin})$.*

⁵The glue code for functions is a bit more complicated than in [Patterson et al. 2022]. Specifically, they do not protect the converted affine functions. This is fine in their setting, because their affine functions are pure, and invoking them multiple times does not lead to run-time errors. However, in the presence of references with strong updates this assumption is no longer true, and not protecting the converted functions will lead to unsound behavior!

9.3 Logical Relations for the Combined Safety

We will prove [Theorem 9.4](#) by constructing a logical relation similarly to what we did for the individual languages in [Section 9.1](#). Our goal is to do so modularly, by reusing as much material from [Section 9.1](#) as possible. In particular, we will reuse all the compatibility lemmas we used to prove [Lemma 9.2](#), and only prove one (!) new compatibility lemma for **TYPED-CONV**. Letting $\mathcal{E}[\tau']$ and $\mathcal{E}[\tau]$ be the expression relations for the logical relation for $\lambda_{\text{rec,io}}$ and $\lambda_{\text{-o,ref}}$ resp., this compatibility lemma is:

LEMMA 9.5. *Suppose that $\mathcal{E}[\tau'](\alpha)$ and $\tau' \sim \tau$. Then $\mathcal{E}[\tau](C_{\tau',\tau}(\alpha))$. Moreover, for the other direction, suppose that $\mathcal{E}[\tau](\alpha)$ and $\tau' \sim \tau$. Then $\mathcal{E}[\tau'](C_{\tau,\tau'}(\alpha))$.*

When we presented the separation logic and logical relations earlier, we presented a slightly simplified version which was sufficient for our purposes. However, in order to prove [Lemma 9.5](#) we need to make use of features that we have not yet presented. We describe those features now.

Separation logic for weak safety. The first feature is that our notion of weakest precondition is actually parameterized by a *stuckness predicate* $s : \text{Error} \rightarrow \text{iProp}$, and satisfies the following rule (in addition to the rules presented in [Section 6](#)):

$$s(e) \vdash \text{wp}_s \text{Err}(e) \{ \Phi \}$$

This means that if the stuckness predicate s holds for some error e , then the weakest precondition predicate holds for that error, irrespectively of the postcondition. The earlier presented weakest precondition, which did not allow for errors, is obtained by using $s(e) = \text{False}$. The general weakest precondition with the stuckness predicate satisfies a version of the adequacy/safety property ([Theorem 6.2](#)), in which the (*safety*) condition is replaced with the following condition:

- (*weak safety*) either there are β_1 and σ_1 such that $(\beta, \sigma') \rightsquigarrow (\beta_1, \sigma_1)$, or $\beta = \text{Err}(e)$ with e satisfying the predicate s .

All the logical relations presented earlier are actually parameterized by a predicate s and uses $\text{wp}_s \alpha \{ \Phi \}$ in the expression interpretation – to recover the earlier stated theorems for full safety we simply instantiate the logical relations with $s = \lambda e. \text{False}$.

Freely adjoined Kripke structure. In addition to the stuckness bit, we actually formulate our logical relations a bit more generally than what we have shown so far. This is because in the logical relations for individual languages require particular resources that we need to combine when constructing a logical relation for the combined language.

Indeed, to ensure that our logical relations are sufficiently modular, we parameterize the expression relation by an arbitrary predicate $P : A \rightarrow \text{iProp}$ of an arbitrary type A . We refer to this predicate P as *freely adjoined Kripke structure* (because, in the underlying model of Iris, it allows us to make arbitrary transitions between worlds, constrained by the predicate P). The general definition of the expression relation for all our logical relation is thus:

$$\mathcal{E}_s^P(\Phi)(\alpha) \triangleq \forall x : A. P(x) \text{ -* } \text{wp}_s \alpha \{ \beta_v. \exists y : A. \Phi(\beta_v) \text{ * } P(y) \}.$$

The idea is that the P parameter describes additional resources (for other effects) and ensures that ITrees in the expression relation preserve any such additional resources. The idea is akin to the “baking-in” of the frame rule in models of separation logic for higher-order languages [[Birkedal et al. 2008](#); [Birkedal and Yang 2008](#)].

The expression relations for individual languages $\lambda_{\text{rec,io}}$ and $\lambda_{\rightarrow,\text{ref}}$ are then both parameterized by predicates $P : A \rightarrow \text{iProp}$ and $s : \text{Error} \rightarrow \text{iProp}$ and defined as

$$\begin{aligned} \mathcal{E}[\tau'](\alpha) &\triangleq \mathcal{E}_s^{\lambda(\sigma',x).\text{has_state}_i(\sigma')*P(x)}(\mathcal{V}[\tau']) (\alpha) \\ \mathcal{E}[\tau](\alpha) &\triangleq \mathcal{E}_s^{\lambda x.\text{heap_ctx}*P(x)}(\mathcal{V}[\tau]) (\alpha). \end{aligned}$$

For $\lambda_{\text{rec,io}}$ and for $\lambda_{\rightarrow,\text{ref}}$ we prove the compatibility properties for arbitrary P and s (in the proofs of the compatibility lemmas, the resources described by P are just passed through). We obtain full safety for the individual languages by instantiating P with $\lambda x.\text{True}$ and s with $\lambda e.\text{False}$.

When we prove partial safety for the combined language, we reuse the same logical relations and the same compatibility lemmas for the individual languages, by instantiating P in $\mathcal{E}[\tau']$ with $P(x) = \text{heap_ctx}$, and by instantiating P in $\mathcal{E}[\tau]$ with $P(x) = \text{has_state}_i(x)$, and s with $\lambda e.\text{Lin}$ in both cases. Then, to prove the fundamental property of the combined language we can reuse the compatibility lemmas for individual languages, and it only remains to prove the compatibility [Lemma 9.5](#) for the type conversion.

The most interesting case of [Lemma 9.5](#) is the conversion between functions, which involves showing soundness of the glue code. The interpretation of non-affine functions is persistent, as it begins with \square , since it is expected that you can use non-affine functions multiple times. The interpretation of affine functions, however, is not persistent — functions can be invoked only once. Because of that, we cannot directly use the interpretation $\mathcal{V}[\tau_1 \rightarrow \tau_2]$ when proving $\mathcal{V}[(\text{Nat} \rightarrow \tau'_1) \rightarrow \tau'_2]$. Instead, we put the interpretation of the affine function in a persistent invariant, which states:

$$(\ell \mapsto \text{Ret}(0) * \mathcal{V}[\tau_1 \rightarrow \tau_2](\alpha_\nu)) \vee (\ell \mapsto \text{Ret}(1)).$$

This describes when the affine function α_ν is protected with the `Think` via a reference ℓ : either the function has not been invoked yet ($\ell \mapsto \text{Ret}(0)$) and it satisfies the value interpretation, or the function has already been invoked ($\ell \mapsto \text{Ret}(1)$) and forcing its `think` will result in $\text{Err}(\text{Lin})$, in which case we do not invoke the function. (As a side remark, we note that we here crucially rely on Iris's powerful notion of invariants — this is another example of why it is advantageous to use Iris as the basis for our separation logic.)

Having established the compatibility lemma for type conversions, and the compatibility lemmas for the operations for each individual language, we prove the fundamental property for the logical relation for the combined language. In particular, for a closed term $\vdash e : \tau$ we obtain $\mathcal{E}_{\lambda e. e=\text{Lin}}^{\lambda x.\text{heap_ctx}*has_state_i(x)}(\mathcal{V}[\tau])(\llbracket e \rrbracket)$. From the adequacy property of the weakest precondition we can then conclude [Theorem 9.4](#).

In summary, this approach to logical relations, with freely adjoined Kripke structure as a parameter, allows us to scale proofs to interoperability between multiple different languages in a modular way. For each individual language we can prove safety separately (without knowing in advance with which other languages we are going to interface). Then we can combine logical relations for individual languages together, by instantiating the freely adjoined Kripke structure with the shared resources or with resources needed to verify the glue code, and reusing the compatibility lemmas.

10 DISCUSSION AND RELATED WORK

We have already discussed a lot of related work throughout the paper; in this section we include some further discussion of related work.

Differences with interaction trees. Whilst our work takes direct inspiration from the interaction trees approach, there are some crucial differences. One of the main difference comes from the treatment of effect reification. The original type of interaction trees is a monad, and the effects in

an interaction tree can be reified “in one go”, for example, with a state monad transformer over ITrees. In our case, we cannot reify all the effects in a guarded interaction tree, due to higher-order functions and higher-order effects. For example, a guarded interaction tree can be a function that contains latent effects, but these effects can only be reified at the point when the function is invoked. Because regular interaction trees contain only first-order structures, it is possible to traverse them completely, reifying all the effects.

Another difference worth mentioning, is that regular interaction trees can be extracted and executed from Coq. Our formalization cannot be directly extracted, as it is built upon a layer of guarded type theory. One potential approach would be to erase the guarded parts from the formulation of GITrees and obtain that way a representation of GITrees in a functional language like OCaml or Haskell, which already supports mixed-variance datatypes. Then, the extraction can be set up in such a way as to use this representation. We have not researched this direction and leave it to future work.

Finally, an important difference between our work and that on ITrees, is that the ITrees development relies heavily on the weak bisimilarity theory of interaction trees, while we opt for developing separation logic and refinements instead. There are several things that complicate the study of bisimilarity for guarded interaction trees. Firstly, the higher-order nature of GITrees suggests that we need to study a more involved notion of behavioral equivalence, like applicative bisimilarity. Secondly, we believe that developing a theory of *weak* bisimilarity in the context of Iris and guarded type theory is still an open question, complicated by issues with step-indices. For these reasons we believe that developing weak applicative bisimilarity for guarded interaction trees will require new techniques and we leave it for future work.

Differences with the standard Iris approach to verification. The standard Iris-based approach to separation logic [Jung et al. 2016, 2015; Krebbers et al. 2017a] is based on operational semantics, and has been proven to scale well to complicated programming language features. The main difference with our work, is that we are the first to build an Iris-based separation logic over denotational semantics, in a way that is tightly integrated with the existing Iris ecosystem. In particular, we rely on the Iris ecosystem for various data structures, resource algebras, base logic (but not the program logic), and the Iris Proof Mode. As such, in terms of reasoning about specific concrete programs, the GITrees-based approach is not that different from what a normal Iris user expects, with the added advantage of using equational reasoning for many computation steps that usually requires some form of symbolic execution.

The main advantage of our approach comes into play when we want to consider new models of programming languages, or reasoning about programs with combinations of effects (as in Section 8), or reasoning about interoperability (as in Section 9).

Domain theory and guarded type theory. Guarded type theory [Birkedal and Møgelberg 2013; Bizjak et al. 2016] has been studied as a setting for domain theory before [Birkedal et al. 2012; Møgelberg and Paviotti 2019; Møgelberg and Vezzosi 2021; Paviotti et al. 2015], but previous works mostly focused on specific typed models and was not formalized. In contrast, here we work with guarded interaction trees as a “universal domain”, similar to domain theoretic models of untyped λ -calculus.

The previous work used (dependent) guarded type theory not just for modeling, but also for reasoning about programs. This required a more complicated type theory and precluded the work from being formalized in a traditional proof assistant like Coq. By contrast, our reasoning is done in the logic *over* a guarded type theory. This is arguably simpler, and allowed us to make use of Iris and formalize all of our results in Coq.

Logical relations and language interoperability. Our case study on language interoperability in Section 9 is inspired by the seminal work of Patterson et al. [2022]. We believe that the approach we take in our work is more modular. Firstly, our approach here is less syntactic, as we use the domain theory of guarded interaction trees as the common setting for interpreting different languages with different effects, and we do not need to come up with a target language for each pair of source languages for which we wish to set up interoperability. Secondly, the models that we construct for individual languages are “local”, and that is exactly what allows for true reuse of proofs and for constructing a common model for the combined language from the individual models. This opens up the possibility of a model for a single language to be reused for different cross-language interactions. In contrast, the type safety of individual source languages in [Patterson et al. 2022] requires to have a model for the combined language in advance. And finally, the treatment of effects and step-indices is more abstract and high-level in our work, since we construct our models using separation logic.

(Guarded) Interaction Trees and Algebraic Effects and Handlers. The treatment of effects in interaction trees is reminiscent of effects in programming languages based on algebraic effects and handlers [Bauer and Pretnar 2015; Plotkin and Pretnar 2013]. Algebraic effects and handlers have been extensively studied in various contexts, including separation logic [de Vilhena and Pottier 2021], and both higher-order effects [Bach Poulsen and van der Rest 2023; van den Berg et al. 2021; Wu et al. 2014] and reasoning about combinations of effects [Johann et al. 2010; Yang and Wu 2021] have been investigated. Despite the aesthetic and moral similarities to (guarded) interaction trees, there are some substantial differences between the two approaches. Under algebraic effects and handlers, both the representation and reification of effects is done *inside* the programming language. As such, a particular theory and implementation of algebraic effects is always tied to a specific programming language. Whereas in the interaction trees-based approach, the effects are handled in the ambient type theory, outside the type of the (guarded) interaction trees itself. See also the discussion in [Xia et al. 2019, Section 8.2].

To our knowledge, the two approaches have not been formally compared. It would be interesting to consider a denotational model of a programming language with algebraic effects inside guarded interaction trees, and to see what kind of properties can be proved using such a model.

Additionally, such a comparison might help us understand the exact class of effects that can be represented with the GITrees-based approach. As it currently stands, our approach to representing effects is “open-ended”, in the sense that we can consider different classes of effects by varying the reification procedure. Of course, different classes of effects allow for different reasoning principles. For example, as we mentioned in the end of Section 4, we consider context-independent effects, which preclude us from modeling call/cc, but allows us to use the bind rule for the weakest precondition calculus. We leave a formal comparison with algebraic effects and further investigations in that area to future work.

ACKNOWLEDGMENTS

This work was supported in part by a Villum Investigator grant (no. 25804), Center for Basic Research in Program Verification (CPV), from the VILLUM Foundation. The authors are grateful to the anonymous reviewers for their comments and suggestions.

DATA AVAILABILITY STATEMENT

The Coq formalization corresponding to this article is available as a Git repository at <https://github.com/logsem/gitrees/tree/pop124> (tag pop124), or under a permanent DOI at <https://doi.org/10.5281/zenodo.10124427>.

REFERENCES

- Casper Bach Poulsen and Cas van der Rest. 2023. Hefty Algebras: Modular Elaboration of Higher-Order Algebraic Effects. *Proceedings of the ACM on Programming Languages* 7, POPL (Jan. 2023), 62:1801–62:1831. <https://doi.org/10.1145/3571255>
- Andrej Bauer and Matija Pretnar. 2015. Programming with Algebraic Effects and Handlers. *Journal of Logical and Algebraic Methods in Programming* 84, 1 (Jan. 2015), 108–123. <https://doi.org/10.1016/j.jlamp.2014.02.001>
- Lars Birkedal and Rasmus Ejlers Møgelberg. 2013. Intensional Type Theory with Guarded Recursive Types qua Fixed Points on Universes. In *28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2013, New Orleans, LA, USA, June 25–28, 2013*. IEEE Computer Society, 213–222. <https://doi.org/10.1109/LICS.2013.27>
- Lars Birkedal, Rasmus Ejlers Møgelberg, Jan Schwinghammer, and Kristian Støvring. 2012. First steps in synthetic guarded domain theory: step-indexing in the topos of trees. *Log. Methods Comput. Sci.* 8, 4 (2012). [https://doi.org/10.2168/LMCS-8\(4:1\)2012](https://doi.org/10.2168/LMCS-8(4:1)2012)
- Lars Birkedal, Bernhard Reus, Jan Schwinghammer, and Hongseok Yang. 2008. A Simple Model of Separation Logic for Higher-Order Store. In *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7–11, 2008, Proceedings, Part II - Track B: Logic, Semantics, and Theory of Programming & Track C: Security and Cryptography Foundations (Lecture Notes in Computer Science, Vol. 5126)*, Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfssdóttir, and Igor Walukiewicz (Eds.). Springer, 348–360. https://doi.org/10.1007/978-3-540-70583-3_29
- Lars Birkedal, Kristian Støvring, and Jacob Thamsborg. 2010. The Category-Theoretic Solution of Recursive Metric-Space Equations. *Theoretical Computer Science* 411, 47 (Oct. 2010), 4102–4122. <https://doi.org/10.1016/j.tcs.2010.07.010>
- Lars Birkedal and Hongseok Yang. 2008. Relational Parametricity and Separation Logic. *Log. Methods Comput. Sci.* 4, 2 (2008). [https://doi.org/10.2168/LMCS-4\(2:6\)2008](https://doi.org/10.2168/LMCS-4(2:6)2008)
- Ales Bizjak, Hans Bugge Grathwohl, Ranald Clouston, Rasmus Ejlers Møgelberg, and Lars Birkedal. 2016. Guarded Dependent Type Theory with Coinductive Types. *CoRR abs/1601.01586* (2016). arXiv:1601.01586 <http://arxiv.org/abs/1601.01586>
- Paulo Emilio de Vilhena and François Pottier. 2021. A Separation Logic for Effect Handlers. *Proceedings of the ACM on Programming Languages* 5, POPL (Jan. 2021), 33:1–33:28. <https://doi.org/10.1145/3434314>
- Matthias Felleisen and Robert Hieb. 1992. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science* 103, 2 (1992), 235–271. [https://doi.org/10.1016/0304-3975\(92\)90014-7](https://doi.org/10.1016/0304-3975(92)90014-7)
- Peter Freyd. 1991. Algebraically Complete Categories. In *Category Theory (Lecture Notes in Mathematics)*, Aurelio Carboni, Maria Cristina Pedicchio, and Giuseppe Rosolini (Eds.). Springer, Berlin, Heidelberg, 95–104. <https://doi.org/10.1007/BFb0084215>
- Naohiko Hoshino. 2012. Step Indexed Realizability Semantics for a Call-by-Value Language Based on Basic Combinatorial Objects. In *2012 27th Annual IEEE Symposium on Logic in Computer Science*. 385–394. <https://doi.org/10.1109/LICS.2012.74>
- Patricia Johann, Alex Simpson, and Janis Voigtländer. 2010. A Generic Operational Metatheory for Algebraic Effects. In *Proceedings of the 25th Annual IEEE Symposium on Logic in Computer Science, LICS 2010, 11–14 July 2010, Edinburgh, United Kingdom*. IEEE Computer Society, 209–218. <https://doi.org/10.1109/LICS.2010.29>
- Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. 2016. Higher-order ghost state. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming, ICFP 2016, Nara, Japan, September 18–22, 2016*, Jacques Garrigue, Gabriele Keller, and Eijiro Sumii (Eds.). ACM, 256–269. <https://doi.org/10.1145/2951913.2951943>
- Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.* 28 (2018), e20. <https://doi.org/10.1017/S0956796818000151>
- Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15–17, 2015*, Sriram K. Rajamani and David Walker (Eds.). ACM, 637–650. <https://doi.org/10.1145/2676726.2676980>
- Nicolas Koh, Yao Li, Yishuai Li, Li-yao Xia, Lennart Beringer, Wolf Honoré, William Mansky, Benjamin C. Pierce, and Steve Zdancewic. 2019. From C to Interaction Trees: Specifying, Verifying, and Testing a Networked Server. In *Proceedings of the 8th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2019)*. Association for Computing Machinery, New York, NY, USA, 234–248. <https://doi.org/10.1145/3293880.3294106>
- Robbert Krebbers, Ralf Jung, Ales Bizjak, Jacques-Henri Jourdan, Derek Dreyer, and Lars Birkedal. 2017a. The Essence of Higher-Order Concurrent Separation Logic. In *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22–29, 2017, Proceedings (Lecture Notes in Computer Science, Vol. 10201)*, Hongseok Yang (Ed.). Springer, 696–723. https://doi.org/10.1007/978-3-662-54434-1_26
- Robbert Krebbers, Amin Timany, and Lars Birkedal. 2017b. Interactive proofs in higher-order concurrent separation logic. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18–20, 2017*, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM, 205–217. <https://doi.org/10.1145/3009837>

3009855

- Mohsen Lesani, Li-yao Xia, Anders Kaseorg, Christian J. Bell, Adam Chlipala, Benjamin C. Pierce, and Steve Zdancewic. 2022. C4: Verified Transactional Objects. *Proceedings of the ACM on Programming Languages* 6, OOPSLA1 (April 2022), 80:1–80:31. <https://doi.org/10.1145/3527324>
- Jacob Matthews and Robert Bruce Findler. 2007. Operational Semantics for Multi-Language Programs. *ACM SIGPLAN Notices* 42, 1 (Jan. 2007), 3–10. <https://doi.org/10.1145/1190215.1190220>
- Rasmus E. Møgelberg and Marco Paviotti. 2019. Denotational Semantics of Recursive Types in Synthetic Guarded Domain Theory. *Mathematical Structures in Computer Science* 29, 3 (March 2019), 465–510.
- Rasmus Ejlers Møgelberg and Andrea Vezzosi. 2021. Two Guarded Recursive Powerdomains for Applicative Simulation. *Electronic Proceedings in Theoretical Computer Science* 351 (Dec. 2021), 200–217. arXiv:2112.14056 [cs]
- Daniel Patterson, Noble Mushtak, Andrew Wagner, and Amal Ahmed. 2022. Semantic Soundness for Language Interoperability. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI 2022)*. Association for Computing Machinery, New York, NY, USA, 609–624. <https://doi.org/10.1145/3519939.3523703>
- Marco Paviotti, Rasmus Ejlers Møgelberg, and Lars Birkedal. 2015. A Model of PCF in Guarded Type Theory. *Electronic Notes in Theoretical Computer Science* 319 (Dec. 2015), 333–349.
- Andrew M. Pitts. 1996. Relational Properties of Domains. *Information and Computation* 127, 2 (June 1996), 66–90. <https://doi.org/10.1006/inco.1996.0052>
- G. D. Plotkin. 1977. LCF Considered as a Programming Language. *Theoretical Computer Science* 5, 3 (Dec. 1977), 223–255. [https://doi.org/10.1016/0304-3975\(77\)90044-5](https://doi.org/10.1016/0304-3975(77)90044-5)
- 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)
- Dana Scott. 1976. Data Types as Lattices. *SIAM J. Comput.* 5, 3 (Sept. 1976), 522–587. <https://doi.org/10.1137/0205037>
- Lucas Silver, Paul He, Ethan Cecchetti, Andrew K Hirsch, and Steve Zdancewic. 2023. Semantics for Noninterference with Interaction Trees. (2023).
- M. B. Smyth and G. D. Plotkin. 1982. The Category-Theoretic Solution of Recursive Domain Equations. *SIAM J. Comput.* 11, 4 (Nov. 1982), 761–783. <https://doi.org/10.1137/0211062>
- Birthe van den Berg, Tom Schrijvers, Casper Bach Poulsen, and Nicolas Wu. 2021. Latent Effects for Reusable Language Components. In *Programming Languages and Systems (Lecture Notes in Computer Science)*, Hakjoo Oh (Ed.). Springer International Publishing, Cham, 182–201. https://doi.org/10.1007/978-3-030-89051-3_11
- Nicolas Wu, Tom Schrijvers, and Ralf Hinze. 2014. Effect Handlers in Scope. In *Proceedings of the 2014 ACM SIGPLAN Symposium on Haskell (Haskell '14)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/2633357.2633358>
- Li-yao Xia, Yannick Zakowski, Paul He, Chung-Kil Hur, Gregory Malecha, Benjamin C. Pierce, and Steve Zdancewic. 2019. Interaction Trees: Representing Recursive and Impure Programs in Coq. *Proceedings of the ACM on Programming Languages* 4, POPL (Dec. 2019), 51:1–51:32. <https://doi.org/10.1145/3371119>
- Zhixuan Yang and Nicolas Wu. 2021. Reasoning about Effect Interaction by Fusion. *Proceedings of the ACM on Programming Languages* 5, ICFP (Aug. 2021), 73:1–73:29. <https://doi.org/10.1145/3473578>
- Kangfeng Ye, Simon Foster, and Jim Woodcock. 2022. Formally Verified Animation for RoboChart Using Interaction Trees. In *Formal Methods and Software Engineering (Lecture Notes in Computer Science)*, Adrian Riesco and Min Zhang (Eds.). Springer International Publishing, Cham, 404–420. https://doi.org/10.1007/978-3-031-17244-1_24
- Yannick Zakowski, Calvin Beck, Irene Yoon, Iliia Zaichuk, Vadim Zaliva, and Steve Zdancewic. 2021. Modular, Compositional, and Executable Formal Semantics for LLVM IR. *Proceedings of the ACM on Programming Languages* 5, ICFP (Aug. 2021), 67:1–67:30. <https://doi.org/10.1145/3473572>
- Hengchu Zhang, Wolf Honoré, Nicolas Koh, Yao Li, Yishuai Li, Li-Yao Xia, Lennart Beringer, William Mansky, Benjamin Pierce, and Steve Zdancewic. 2021. Verifying an HTTP Key-Value Server with Interaction Trees and VST. In *12th International Conference on Interactive Theorem Proving (ITP 2021) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 193)*, Liron Cohen and Cezary Kaliszyk (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 32:1–32:19. <https://doi.org/10.4230/LIPIcs.ITP.2021.32>

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