

Context-Dependent Effects in Guarded Interaction Trees

Sergei Stepanenko¹, Emma Nardino², Dan Frumin³, Amin Timany¹, and Lars Birkedal¹

¹ Aarhus University, Aarhus, Denmark

{sergei.stepanenko,timany,birkedal}@cs.au.dk

² Univ Lyon, ENS de Lyon, UCBL, CNRS, Inria, LIP, UMR 5668, F-69342, Lyon
CEDEX 07, France emma.nardino@ens-lyon.fr

³ University of Groningen, Groningen, The Netherlands
d.frumin@rug.nl

Abstract. Guarded Interaction Trees are a structure and a fully formalized framework for representing higher-order computations with higher-order effects in Coq. We present an extension of Guarded Interaction Trees to support formal reasoning about context-dependent effects. That is, effects whose behaviors depend on the evaluation context, e.g., `call/cc`, `shift` and `reset`. Using and reasoning about such effects is challenging since certain compositionality principles no longer hold in the presence of such effects. For example, the so-called “bind rule” in modern program logics (which allows one to reason modularly about a term inside a context) is no longer valid. The goal of our extension is to support representation and reasoning about context-dependent effects in the most painless way possible. To that end, our extension is conservative: the reasoning principles (and the Coq implementation) for context-independent effects remain the same. We show that our implementation of context-dependent effects is viable and powerful. We use it to give direct-style denotational semantics for higher-order programming languages with `call/cc` and with delimited continuations. We extend the program logic for Guarded Interaction Trees to account for context-dependent effects, and we use the program logic to prove that the denotational semantics is adequate with respect to the operational semantics. This is achieved by constructing logical relations between syntax and semantics inside the program logic. Additionally, we retain the ability to combine multiple effects in a modular way, which we demonstrate by showing type soundness for safe interoperability of a programming language with delimited continuations and a programming language with higher-order store.

Keywords: Coq · Iris · denotational semantics · logical relations · control flow operators · continuations · delimited continuations

1 Introduction

Despite a lot of recent progress, representing and reasoning about programming languages in proof assistants, such as Coq, is still considered a major challenge.

The design space is wide and many approaches have been considered. Recently, research on a novel point in the design space was initiated with the introduction of Interaction Trees [33], or ITrees for short. ITrees were introduced to simplify representation and reasoning about possibly non-terminating programs with side effects in Coq. In a sense, ITrees provide a target for denotational semantics of programming languages, which allows one to abstract from syntactic details often found in models based on operational semantics. ITrees specifically allow one to easily represent and reason about various effects and their combinations in a modular way. A wide range of subsequent applications of ITrees (see, e.g., [18,35,34,25], among others) show that they indeed work very well for representing and reasoning about first-order programs with first-order effects. As part of the trade-offs, ITrees could not support higher-order representations and higher-order effects. To address this challenge, Guarded Interaction Trees (or GITrees for short) were introduced [13].

While GITrees support *higher-order effects*, in particular, the challenging case of higher-order store, they are limited to effects that do not alter the control flow of the program (more specifically, the continuation). As a consequence, one cannot use GITrees to give direct-style denotational semantics of programming languages with context-dependent effects such as `call/cc`, exceptions, or delimited continuations. It is of course possible to give semantics to, e.g., `call/cc` using a CPS translation, but this would require a global transformation which complicates representation and reasoning, especially in combination with other effects present. In this paper we extend GITrees to support direct-style representation and reasoning about higher-order programs with higher-order context-dependent effects in Coq, and evaluate its modularity.

We want to stress that our extension to context-dependent effects is not only theoretically interesting, but also important for scalability, since many real mainstream programming languages include context-dependent effects. Indeed, exceptions are now a standard feature in many languages, and while other context-dependent effects such as delimited continuations are not as widespread in mainstream programming languages, they are present in the core calculi used for some such languages: for example, the Glasgow Haskell Compiler core language was recently extended with delimited continuations to support the introduction of effect systems, which can be efficiently developed on top of delimited continuations [16]. Moreover, effect handlers, which rely on control flow operators, have recently been introduced in OCaml 5.0, and as a design feature in Helium [7], Koka [21,20], and other languages. Note that these languages do not only include forms of delimited continuations, but also other effects, which underscores the importance of considering delimited control in combination with other effects.

Overview of technical development and key challenges. Similarly to how ITrees are defined as a coinductive type in Coq, GITrees are defined as a guarded recursive type (this is to support function spaces in GITrees; in the presence of function spaces there are negative occurrences of the recursive type and hence one cannot simply define the type of GITrees using coinduction). Coq does not

directly support guarded recursive types, so GITrees are defined using a fragment of guarded type theory implemented in Coq, as part of the Iris framework [14]. To work efficiently with GITrees we make use of other Iris features like separation logic and Iris Proof Mode [19], which we use to define custom program logics for different (combinations of) effects. This enables us to reason about GITrees smoothly in the Iris logic in Coq, in much the same way as one works directly in Coq. We recall the precise definition of GITrees in Section 2.

Broadly speaking, GITrees model effects in the following way. The type of GITrees is parameterized over a set of effectful operations. Each operation is given meaning by a *reifier* function, using a form of state monad. From this, we define the *reduction* relation of GITrees, which gives semantics to computations represented by GITrees.

To support context-dependent effects, we extend (Section 3) the notion of a reifier so that reification of effects can also depend on the context; technically, the reifier operation becomes parameterized by a suitable GITrees continuation. This extension allows us to give semantics to context-dependent effects, but it comes at a price. In particular, following the change in semantics, we need to reformulate the program logic for GITrees: in the presence of context-dependent effects (like `call/cc`), the so-called “bind”-rule becomes unsound. Of course, we do not want this reformulation to complicate reasoning about computations that do *not* include context-dependent effects. To that end, we parameterize the GITrees (and the program logic) by a flag, which allows us to recover the original proof rules and make sure that all of the original GITrees framework still works with our extension.

To motivate the extension to context-dependent effects, we give direct-style denotational semantics to a higher-order programming language $\lambda_{\text{call/cc}}$ with `call/cc` (Section 3). Furthermore, we use the derived program logic to construct a logical relation between the denotational and the operational semantics to prove computational adequacy of our model.

Our main interest, however, lies in the treatment of delimited continuations. In Section 4 we show how to represent delimited continuations as effects in GITrees, and we use them to define a novel denotational semantics for a programming language with `shift` and `reset` operators. We prove that our denotational semantics is sound with respect to the operational semantics (given by an extension of the CEK abstract machine). We additionally use the program logic to define a logical relation, and prove computational adequacy and *semantic* type soundness. We recall that semantic type soundness is interesting because it allows one to combine syntactically well-formed programs with syntactically ill-typed, but semantically well-behaved programs [29].

As we mentioned, it is important to consider delimited continuations not only on their own, but in combinations with other effects. And indeed, one of the key points of ITrees, and therefore also of GITrees, is that they support reasoning about effecting and language interoperability by establishing a common unifying semantic framework. In this paper, we consider (in Section 5) an example of such interaction: we show a type-safe embedding of λ_{delim} with delimited continuations

into a language λ_{embed} with higher-order store. We allow λ_{delim} expressions to be embedded into λ_{embed} by surrounding them by simple glue code, and use a type system to ensure type safety of the combined language. To define the semantics of the combined language we rely on the modularity of GITrees, and combine reifiers for delimited continuations with reifiers for higher-order store. We prove type safety of the combined language by constructing a logical relation and use the program logic both to define the logical relation and to verify the glue code between the two languages. The type system for the combined language naturally requires that the embedded code is well-typed according to the type system for λ_{delim} and thus we can rely on the type soundness of λ_{delim} (proved in the earlier Section 4) when proving type safety for the combined language. At the end of Section 5, we give an example of how to verify a more involved interaction of effects, albeit without the type system.

Summary of Contributions. In summary, we present:

- A conservative (with respect to the old results) extension to GITrees for representing and reasoning about context-dependent effects (Section 3).
- A sound and adequate model of a calculus with `call/cc` and `throw`, implemented in a direct style (Section 3.3).
- A sound and adequate model of a calculus with delimited continuations, with operations `shift` and `reset`, implemented in a direct style (Section 4).
- A type system for interoperability between a programming language with delimited continuations and a programming language with higher-order store, with a semantic type safety proof (Section 5).

All results in the paper have been formalized in Coq as a modification to the GITrees library and the previously proved results have been ported to our extension. We conclude and discuss related work in Section 6. Before we go on with the main part of the paper, we recall some background material on GITrees.

2 Guarded Interaction Trees

In this section we provide an introduction to guarded interaction trees. Our treatment is brief, and we refer the reader to the original paper for details [13].

Iris and Guarded Type Theory. Guarded Interaction trees (GITrees) are defined in Iris logic. Here we briefly touch Iris, and refer the reader to the literature on Iris [15] and guarded type theory [8] for more in-depth details. Iris is a separation logic framework built on top of a model of *guarded type theory*, the main use of which is to solve recursive equations and define guarded recursive types, such as the type of GITrees described below. Moreover, Iris has a specialized proof mode [19], implemented in Coq. This allows the users of Iris to carry out formal reasoning in separation logic as if they are proving things normally Coq, as we have done in the formalization of this work. For this reason, in the paper we will

$$\begin{aligned}
\tau ::= & \text{iProp} \mid 0 \mid \mathbf{1} \mid \mathbb{B} \mid \text{Nat} \mid \tau + \tau \mid \tau \times \tau \mid \tau \rightarrow \tau \mid \blacktriangleright \tau \mid I \mid \Sigma_{i \in \mathbb{I}} \tau_i \mid \Pi_{i \in \mathbb{I}} \tau_i \mid \dots \\
t ::= & x \mid F(t_1, \dots, t_n) \mid \text{abort } t \mid () \mid (t, t) \mid \pi_i t \mid \lambda x : \tau. t \mid \\
& \text{inj}_i t \mid \text{match } t \text{ with } \overline{\text{inj}_i x.t} \text{ end} \mid \text{next}(t) \mid \text{fix}_\tau \mid \dots \\
P ::= & \text{False} \mid t =_\tau t \mid P \vee P \mid P \rightarrow P \mid \forall x : \tau. P \mid P * P \mid P \multimap P \mid \square P \mid \boxed{P} \mid \triangleright P \mid \dots
\end{aligned}$$
Fig. 1. Grammar for the Iris base logic.
$$\begin{aligned}
& \text{guarded type } \mathbf{IT}_E(A) = \text{Ret} : A \rightarrow \mathbf{IT}_E(A) \\
& \mid \text{Fun} : \blacktriangleright(\mathbf{IT}_E(A) \rightarrow \mathbf{IT}_E(A)) \rightarrow \mathbf{IT}_E(A) \\
& \mid \text{Err} : \text{Error} \rightarrow \mathbf{IT}_E(A) \\
& \mid \text{Tau} : \blacktriangleright \mathbf{IT}_E(A) \rightarrow \mathbf{IT}_E(A) \\
& \mid \text{Vis} : \prod_{i \in \mathbb{I}} (\text{Ins}_i(\mathbf{IT}_E(A)) \times (\text{Outs}_i(\mathbf{IT}_E(A)) \rightarrow \blacktriangleright \mathbf{IT}_E(A))) \rightarrow \mathbf{IT}_E(A)
\end{aligned}$$
Fig. 2. Guarded datatype of interaction trees.

work with Iris and its type theory informally. Still, we need to say a few things about the foundations.

The syntax of Iris, shown in Figure 1, contains types, terms, and propositions. The grammar is standard for higher-order logic, with the exception of the guarded types fragment, and separation logic connectives. The type of propositions is denoted iProp . The *guarded* part of guarded type theory is the “later” modality \blacktriangleright . Intuitively, we view all types as indexed by a natural number, where τ_n contains elements of τ “at time” n . Then $\blacktriangleright \tau$ contains elements of τ at a later time; that is, $(\blacktriangleright \tau)_n = \tau_{n-1}$. There is an embedding $\text{next} : \tau \rightarrow \blacktriangleright \tau$, and there is a *guarded* fixed point combinator $\text{fix}_\tau : (\blacktriangleright \tau \rightarrow \tau) \rightarrow \tau$, similar to the unguarded version in PCF. We can also lift functions to \blacktriangleright : given $f : A \rightarrow B$, we have $\blacktriangleright f : \blacktriangleright A \rightarrow \blacktriangleright B$.

For the proposition, Iris contains the usual separation logic connectives, and the two modalities: “later” \triangleright and “persistently” \square . The propositional \triangleright modality reflects the type-level later modality \blacktriangleright on the level of propositions, as justified by the following rule: $\triangleright(\alpha =_\tau \beta) \dashv\vdash \text{next}(\alpha) =_{\blacktriangleright \tau} \text{next}(\beta)$. The persistence modality $\square P$ states that the proposition P is available without claiming any resources (as it normally is the case in separation logic); crucially it makes the proposition duplicable: $\square P \vdash (\square P) * (\square P)$. An example of a persistent proposition is the invariant proposition \boxed{P} , which satisfies $\boxed{P} \vdash \square \boxed{P}$.

Guarded Interaction Trees. Guarded recursive datatypes are datatypes obtained from recursive equations of the form $X = F(\blacktriangleright X)$. In other words, guarded recursive datatypes are similar to the regular datatypes you see in normal programming languages, but every recursive occurrence of the type must be

guarded by the \blacktriangleright modality. The datatype we are concerned with here is the type of GITrees, shown in Figure 2. It is parameterized over two types: the ground type A and the effect signature E (more on it below).

Guarded Interaction Trees represent computational trees in which the leaves are of the ground type ($\text{Ret}(a)$), error states ($\text{Err}(e)$), and functions ($\text{Fun}(f)$). The leaves $\text{Ret}(a)$ and $\text{Fun}(f)$ are also called *values*, and we write $\mathbf{IT}_E^v(A)$ for the type of $\mathbf{IT}_E(A)$ -values.

The nodes of the computation trees are of the two kinds. The first one is a “silent step” constructor $\text{Tau}(\alpha)$. It represents an unobservable internal step of the computation. For convenience, we use the function $\text{Tick} \triangleq \text{Tau} \circ \text{next} : \mathbf{IT}_E(A) \rightarrow \mathbf{IT}_E(A)$ that “delays” its argument. This function satisfies the following equation: $\text{Tick}(\alpha) = \text{Tick}(\beta) \dashv\vdash \triangleright(\alpha = \beta)$.

The second kind of nodes are effects given by $\text{Vis}_i(x, k)$. The parameters I , Ins and Outs are part of the effect signature E . The set I is the set of *names* of operations. The *arities* of an operation $i \in I$ are given by functors Ins_i and Outs_i . Let us give an example.

Consider the following signature for store effects. The signature E_{state} consists of effects $\{\text{write}, \text{read}, \text{alloc}\}$ with the following input/output arities:

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

For example, **write** expect a location and a new GITree as its input, and simply returns the unit value as an output. We write $\text{Vis}_{\text{write}}((\ell, \alpha), \lambda_. \beta)$ for the computation that invokes the **write** effect with arguments ℓ and α , waits for it to return, and proceeds as β . Thus, the first argument for Vis_i is the input, and the second one is the continuation dependent on the output. This continuation determines the branching in (G)Itrees.

For effects like above, it is usually convenient to provide wrappers:

$$\begin{aligned} \text{Alloc}(\alpha : \mathbf{IT}, k : \text{Loc} \rightarrow \mathbf{IT}) &\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}) &\triangleq \text{Vis}_{\text{write}}((\ell, \text{next}(\alpha)), \lambda x. \text{next}(\text{Ret}(\text{inj}(\)))) \end{aligned}$$

When the signature and the return type are clear from the context, we simply write \mathbf{IT} and \mathbf{IT}^v for the GITrees and GITree-values.

Equational theory. GITrees come with a number of operations (defined using the recursion principle) that are used for writing and composing computations. Here we list some of those operations which we will be using. The function $\text{get_val}(\alpha, f : \mathbf{IT}^v \rightarrow \mathbf{IT})$ are used for sequencing computations, and its corresponding equations are shown in Figure 3. Intuitively, $\text{get_val}(\alpha, f)$ first tries to compute α to a value (a $\text{Ret}(a)$ or a $\text{Fun}(g)$), and then calls f on that value. Similarly, $\text{get_fun}(\alpha : \mathbf{IT}, f : \blacktriangleright(\mathbf{IT} \rightarrow \mathbf{IT}) \rightarrow \mathbf{IT})$ and $\text{get_ret}(\alpha : \mathbf{IT}_E(A), f : A \rightarrow \mathbf{IT}_E(A))$ first compute α to a value; if that value is a function $\text{Fun}(g)$ (resp., $\text{Ret}(a)$), then it proceeds with $f(g)$ (resp., $f(a)$). Otherwise it results in an runtime error.

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

Fig. 3. Example function on Guarded Interaction Trees.

$$\begin{array}{c}
r : \prod_{i \in E} \text{Ins}_i(\mathbf{IT}_E) \times \text{State} \rightarrow \text{option}(\text{Outs}_i(\mathbf{IT}_E) \times \text{State}) \\
\frac{r_i(x, \sigma) = \text{Some}(y, \sigma') \quad k \ y = \text{next}(\beta)}{\text{reify}(\text{Vis}_i(x, k), \sigma) = (\text{Tick}(\beta), \sigma')} \quad \frac{r_i(x, \sigma) = \text{None}}{\text{reify}(\text{Vis}_i(x, k), \sigma) = (\text{Err}(\text{RunTime}), \sigma)}
\end{array}$$

Fig. 4. Signature of reifiers and the reification function

Crucially, to work with higher-order computations, GITrees provide the “call-by-value” application $\alpha \bullet \beta$ satisfying the following equations:

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

where $-\ (\blacktriangleright \bullet) -$ is defined as 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 $\alpha \bullet \beta$ simulates strict function application. It first tries to evaluate β to a value β_v . Then it tries to evaluate α to a function f . If it succeeds, then it invokes $f(\beta_v)$. If at any point it fails, application results in a runtime error.

For the often-used case of GITrees where the ground type includes natural numbers, we use the function $\text{NatOp} : (\mathbb{N} \rightarrow \mathbb{N} \rightarrow \mathbb{N}) \rightarrow \mathbf{IT} \rightarrow \mathbf{IT} \rightarrow \mathbf{IT}$ which lifts binary functions on natural numbers to binary functions on GITrees. That is, $\text{NatOp}_f(\alpha, \beta)$ first evaluate GITrees β and α to values. If those values are natural numbers, then it computes f of those numbers and returns the result as a GITree. Otherwise, it returns a runtime error $\text{Err}(\text{RunTime})$.

Reification and reduction relation. The semantics for effects are given in terms of *reifiers*. A reifier for the signature E is a tuple (State, r) , where State is a type representing the internal state needed to reify the effects, and r is a reifier function of the type given in Figure 4. The idea is that r_i uses the internal state State to compute the output of the effect i based on its input.

For example, for the store effects we take $State$ to be a map from locations to $\blacktriangleright\mathbf{IT}$ (representing the heap); and we define the following reifier functions:

$$\begin{aligned} r_{\text{write}}((\ell, \alpha), \sigma) &= \text{Some}(\ell, \sigma[\ell \mapsto \alpha]) && \text{(where } \ell \in \sigma, \text{ and None otherwise)} \\ r_{\text{read}}(\ell, \sigma) &= \text{Some}(\alpha, \sigma) && \text{(where } \sigma(\ell) = \alpha, \text{ and None otherwise)} \\ r_{\text{alloc}}(\alpha, \sigma) &= \text{Some}(\ell, \sigma[\ell \mapsto \alpha]) && \text{(where } \ell \notin \sigma) \end{aligned}$$

Given reifiers for all the effects, we define a function $\text{reify} : \mathbf{IT} \times State \rightarrow \mathbf{IT} \times State$ (as in Figure 4) that, given (α, σ) reifies the top-level effect in α using the state σ , and returns the reified GITree and the updated state.

The reify function is then used to give reduction semantics for GITrees. We write $(\alpha, \sigma) \rightsquigarrow (\beta, \sigma')$ for such a reduction step. The definition of \rightsquigarrow is given internally in the logic:

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

That is, either α is a “delayed” computation $\text{Tick}(\beta)$, which then reduces to β ; or it is an effect that can be reified. Recall that we write Tick for the composition $\text{Tau} \circ \text{next}$.

Note that the reify function operates on the top-level effect of the GITree. But what if the top-level constructor is not Vis , e.g. if we have an effect inside an “evaluation context”? The role of evaluation contexts in GITrees is played by *homomorphisms*, which also allow us to bubble up necessary effects to the top of the GITree.

Definition 1 (Homomorphism). *A map $f : \mathbf{IT} \rightarrow \mathbf{IT}$ is a homomorphism, written $f \in \text{Hom}$, if it satisfies:*

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

For example, $\lambda x. \alpha \bullet x$ is a homomorphism, and so is $\lambda x. \text{get_val}(x, f)$. On the other hand, $\lambda x. \text{Vis}_{\text{alloc}}(\text{next}(x), k)$ (for some fixed k) is *not* a homomorphism.

Program logic. In order to reason about GITrees, we employ the full power of the Iris separation logic framework. The program logic operates on the propositions of the form $\text{wp } \alpha \{ \Phi \}$. This weakest precondition proposition intuitively states that the GITree α is safe to reduce, and when it fully reduces, the resulting value satisfies the predicate Φ . Another important predicate is $\text{has_state}(\sigma)$, which signifies ownership of the current state σ .

In Figure 5 we show the rules, on which we focus in this work. Let us describe their meaning. The rule WP-REIFY allows us to symbolically execute effects in GITrees. It is given in a general form, and is used to derive domain-specific rules for concrete effects. Another important rule is WP-HOM which allows one to separate the reasoning about the computation from the reasoning about the context. The reason why WP-HOM is sound (this is going to be important in the next section when we make it unsound), is because the reduction \rightsquigarrow of GITrees satisfies the following properties which allow one disentangle a homomorphism from the GITree it’s applied to:

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

Fig. 5. Selected weakest precondition rules.

Lemma 1. *Let f be a homomorphism. Then,*

- $(\alpha, \sigma) \rightsquigarrow (\beta, \sigma')$ implies $(f(\alpha), \sigma) \rightsquigarrow (f(\beta), \sigma')$;
- 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 $\triangleright(f(\beta) = \beta')$.

Finally, as usual in Iris, the program logic satisfies an adequacy property, which allows one to relate propositions proved in the logic to the actual semantics:

Theorem 1. *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 $\beta \in \mathbf{IT}^v$, 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 $e \in \mathbf{Error}$.

The role of meta-logic is played by the Coq system; thus, the adequacy theorem allows us to relate proofs inside the program logic (Iris) to the proofs on the level of Coq. This aspect is important in Iris and GITrees in general, but it is orthogonal to the work that we present in this paper. See [13] for more details.

3 Context-Dependent Reification

In this section we extend reification to handle context-dependent effects, using a language $\lambda_{\text{call/cc}}$ with `call/cc` as a concrete example. In Section 3.1 we present $\lambda_{\text{call/cc}}$'s syntax and operational semantics (in the usual style with evaluation contexts). We then show why the current GITrees framework cannot be used as a denotational model for $\lambda_{\text{call/cc}}$ directly. In Section 3.2 we introduce our generalization of reification for context-dependent effects and corresponding extensions to the GITrees program logic. In Section 3.3 we demonstrate that our extension works as intended: we give a denotational semantics for $\lambda_{\text{call/cc}}$, and we show how the general program logic for GITrees specializes to a logic for reasoning about `call/cc`. We prove soundness and computational adequacy of denotational semantics using a logical relation defined within our program logic.

$$\begin{aligned}
\mathit{Ins}_{\text{callcc}}(X) &\triangleq ((\blacktriangleright X \rightarrow \blacktriangleright X) \rightarrow \blacktriangleright X) & \mathit{Outs}_{\text{callcc}}(X) &\triangleq \blacktriangleright X \\
\mathit{Ins}_{\text{throw}}(X) &\triangleq \blacktriangleright X \times \blacktriangleright (X \rightarrow X) & \mathit{Outs}_{\text{throw}}(X) &\triangleq 0 \\
\text{Callcc}(f) &\triangleq \text{Vis}_{\text{callcc}}(f, \text{id}) & \text{Throw}(e, f) &\triangleq \text{Vis}_{\text{throw}}(e, f, \lambda x. \text{abort } x)
\end{aligned}$$

Fig. 7. Signatures and operations on GITrees with `call/cc`.

$$\begin{aligned}
r &: \prod_{i \in E} \mathit{Ins}_i(\mathbf{IT}_E) \times \mathit{State} \times (\mathit{Outs}_i(\mathbf{IT}_E) \rightarrow \blacktriangleright \mathbf{IT}_E) \rightarrow \mathit{option}(\blacktriangleright \mathbf{IT}_E \times \mathit{State}) \\
\frac{r_i(x, \sigma, \kappa) = \text{Some}(\beta, \sigma')}{\text{reify}(\text{Vis}_i(x, \kappa), \sigma) = (\text{Tau}(\beta), \sigma')} & \qquad \frac{r_i(x, \sigma, \kappa) = \text{None}}{\text{reify}(\text{Vis}_i(x, \kappa), \sigma) = (\text{Err}(\text{RunTime}), \sigma)}
\end{aligned}$$

Fig. 8. Type of context-dependent reifiers and the context-dependent reify function

For convenience, we will use the abbreviations $\text{Callcc}(f)$ and $\text{Throw}(e)$, defined in Figure 7, for representing denotations of `throw` and `call/cc` as effects in GITrees.

To complete all the ingredients for the denotational semantics, we need reifiers for the `callcc` and `throw` effects. Given our operational understanding of continuations, the natural choice for the local state type State is $\mathbf{1}$ (since we do not have any state). However, the current reifier signature (Figure 4) poses a problem. Reifiers, as they are now, cannot access their current continuation, which is essential for both effects. $\text{Callcc}(f)$ needs to pass the current continuation to f , while Throw must redirect control to a provided continuation instead of returning normally. The current reifiers lacks this capability, and in the next subsection we show how to generalize the notion of reification to context-dependent effects.

3.2 Context-dependent Reifiers

This section presents our extension to context-dependent reification, and the limitations it imposes on the program logic. In order to allow reifiers to manage continuations, we change the type of reifiers to accept continuations as an extra parameter, as shown in Figure 8. Continuations for a given effect are functions from the effect's outputs to GITrees: $\mathit{Outs}_i(\mathbf{IT}_E) \rightarrow \blacktriangleright \mathbf{IT}_E$. Given a set of context-dependent reifiers, we define a context-dependent `reify` function, also shown in Figure 8. As before, `reify` dispatches to the correct individual reifier for the effect. Note that now it is the user's responsibility to pass the output of an effect to the given continuation if the control flow is not supposed to be interrupted. For example, since the evaluation of a `call/cc` ($x. e$) expression does not modify the control flow itself, but simply passes the current continuation to its body, the context-dependent reifier for `callcc` is simply $r_{\text{callcc}}(x, \sigma, \kappa) = \text{Some}(\kappa(x, \kappa), \sigma)$.

Before we move on to discussing the consequences of this for program logic, we would like to note that our treatment of continuation (with top-level reifiers

$$\begin{array}{c}
\text{WP-WRITE} \\
\frac{\text{heap_ctx} \triangleright \ell \mapsto \alpha \quad \triangleright(\ell \mapsto \beta \text{ -*}) \quad \text{wp Ret()} \{\Phi\}}{\text{wp Write}(\ell, \beta) \{\Phi\}}
\end{array}
\qquad
\begin{array}{c}
\text{WP-WRITE-CTX-DEP} \\
\frac{\kappa \in \text{Hom} \quad \text{heap_ctx} \triangleright \ell \mapsto \alpha \quad \triangleright(\ell \mapsto \beta \text{ -*}) \quad \text{wp } \kappa \text{ (Ret())} \{\Phi\}}{\text{wp } \kappa \text{ (Write}(\ell, \beta)) \{\Phi\}}
\end{array}
\qquad
\begin{array}{c}
\text{WP-REIFY-CTX-DEP} \\
\frac{\text{has_state}(\sigma) \quad r_i(x, \sigma, k) = \text{Some}(\text{next}(\beta), \sigma') \quad \triangleright(\text{has_state}(\sigma') \text{ -*}) \quad \text{wp } \beta \{\Phi\}}{\text{wp Vis}_i(x, k) \{\Phi\}}
\end{array}$$

Fig. 9. Program logic in the presence of context-dependent reifiers.

dispatching them) parallels Cartwright and Felleisen’s “extensible direct models” [9], which also aimed to support extensible denotational semantics in classical domain theory. We discuss this more in Section 6.

Program logic for GITrees in the presence of context-dependent reifiers. To reflect the generalization to context-dependent reifiers in the program logic, we replace the proof rule WP-REIFY by WP-REIFY-CTX-DEP, shown in Figure 9. This is, however, not the only change we need to make. In the presence of context-dependent effects, WP-HOM is not sound! (A similar observation was also made by [27] in their development of a program logic for `call/cc`.) The reason is that context-dependent reification invalidates Lemma 1. Now, since WP-HOM is not sound anymore, one might expect that we need to adapt all the other program logic rules to include a homomorphism similarly to how the rules of [27] were adapted to include an evaluation context. However, this is not necessary, because our program logic is defined on *denotations* on which we have a non-trivial equational theory, which can be used to reason about ‘pure’ GITrees. Only for effectful operations, the proof rules will now have to include a surrounding homomorphism. E.g., WP-WRITE from [13] is generalized to WP-WRITE-CTX-DEP, and considers ambient homomorphisms explicitly.

Our context-dependent reification extension, though simple, allows us to build sound and adequate denotational models for languages with control-flow operators, including λ_{callcc} (shown in the next subsection). Moreover, our extension is conservative, and we recover previous case studies (computational adequacy of $\lambda_{\text{rec,io}}$ and type safety for $\lambda_{\rightarrow, \text{ref}}$ [13]) with minimal modifications; see the accompanying Coq formalization.

3.3 Denotational Semantics of λ_{callcc}

In this section we show that context-dependent reifiers are sufficient for providing a sound and adequate semantic model of λ_{callcc} . We define context-dependent reifiers for `callcc` and `throw`, then prove that this gives a sound interpretation w.r.t. operational semantics. To show adequacy, we define a logical relation, which relates the denotational and operational semantics. The logical relation is defined in the (updated) program logic for GITrees (following the approach in [13]), and validates the utility of the program logic.

$$\begin{aligned}
\mathbf{E}\llbracket x \rrbracket_\rho &= \rho(x) \\
\mathbf{E}\llbracket \text{call/cc } (x. e) \rrbracket_\rho &= \text{Callcc}(\lambda(f : \blacktriangleright \mathbf{IT} \rightarrow \blacktriangleright \mathbf{IT}). \mathbf{E}\llbracket e \rrbracket_{\rho[x \mapsto \text{Fun}(\text{next}(\lambda y. \text{Tau}(f(\text{next}(y))))])}) \\
\mathbf{E}\llbracket \text{throw } e_1 \text{ to } e_2 \rrbracket_\rho &= \text{get_val}(\mathbf{E}\llbracket e_1 \rrbracket_\rho, \lambda x. \text{get_fun}(\mathbf{E}\llbracket e_2 \rrbracket_\rho, \lambda f. \text{Throw}(x, f))) \\
\mathbf{V}\llbracket \text{cont } K \rrbracket_\rho &= \text{Fun}(\text{next}(\lambda x. \text{Tau}(\mathbf{K}\llbracket K \rrbracket_\rho (\blacktriangleright \bullet) \text{next}(x)))) \\
\mathbf{K}\llbracket \text{throw } K \text{ to } e \rrbracket_\rho &= \lambda x. \text{get_val}(\mathbf{K}\llbracket K \rrbracket_\rho x, \lambda y. \text{get_fun}(\mathbf{E}\llbracket e \rrbracket_\rho, \lambda f. \text{Throw}(y, f))) \\
\mathbf{K}\llbracket \text{throw } v \text{ to } K \rrbracket_\rho &= \lambda x. \text{get_val}(\mathbf{V}\llbracket v \rrbracket_\rho, \lambda y. \text{get_fun}(\mathbf{K}\llbracket K \rrbracket_\rho x, \lambda f. \text{Throw}(y, f)))
\end{aligned}$$

Fig. 10. Denotational semantics of λ_{callcc} (selected clauses).

Interpretation of λ_{callcc} . The denotational semantics of λ_{callcc} is shown in Figure 10 (selected clauses only; see Coq formalization for the complete definition). The interpretation is split into three parts: $\mathbf{E}\llbracket - \rrbracket$ for expressions, $\mathbf{V}\llbracket - \rrbracket$ for values, and $\mathbf{K}\llbracket - \rrbracket$ for contexts. For the interpretation of `throw` e_1 `to` e_2 , the left-to-right evaluation order is enforced by the functions `get_val` and `get_fun`. They first evaluate their argument to a GITree value, and then pass it on (c.f. Figure 3).

The context-dependent reifiers for the effects `callcc` and `throw` are defined as follows:

$$r_{\text{callcc}}(f, (), \kappa) = \text{Some}(\kappa (f \kappa), ()) \quad r_{\text{throw}}((\alpha, f), (), \kappa) = \text{Some}(f \alpha, ())$$

To show that the denotational semantics is sound, we need the following lemma that shows that interpretations of expressions in evaluation contexts are decomposed into applications of homomorphisms.

Lemma 2. *For any context K and an environment ρ , we have $\mathbf{K}\llbracket K \rrbracket_\rho \in \text{Hom}$. For any context K , expression e , and an environment ρ , $\mathbf{E}\llbracket K[e] \rrbracket_\rho = \mathbf{K}\llbracket K \rrbracket_\rho(\mathbf{E}\llbracket e \rrbracket_\rho)$.*

With these results at hand, we can show soundness of our interpretation:

Lemma 3. *Soundness.* Suppose $e_1 \mapsto e_2$. Then $(\mathbf{E}\llbracket e_1 \rrbracket_\rho, ()) \rightsquigarrow^* (\mathbf{E}\llbracket e_2 \rrbracket_\rho, ())$, where $() : \mathbf{1}$ is the unique element of the unit type, representing the (lack of) state.

Program logic for λ_{callcc} . We now specialize the general program logic rule `WP-REIFY-CTX-DEP` using the reifiers for `callcc` and `throw` to obtain the following program logic rules:

$$\begin{array}{c}
\text{WP-THROW} \\
\frac{\kappa \in \text{Hom} \quad \text{has_state}(\sigma) \quad \triangleright(\text{has_state}(\sigma) \multimap \text{wp } f \ x \ \{\Phi\})}{\text{wp } \kappa \ (\text{Throw}(\text{next}(x), \text{next}(f))) \ \{\Phi\}}
\end{array}
\qquad
\begin{array}{c}
\text{WP-CALLCC} \\
\frac{\kappa \in \text{Hom} \quad \text{has_state}(\sigma) \quad \triangleright(\text{has_state}(\sigma) \multimap \text{wp } \kappa \ (f \ \kappa) \ \{\Phi\})}{\text{wp } \kappa \ (\text{Callcc}(\text{next} \circ f)) \ \{\Phi\}}
\end{array}$$

where κ is a homomorphism representing the current evaluation context on the level of GITrees. The reader may wonder why these rules include the `has_state`(σ) predicates, since it is just 'threaded around'. The reason is that these rules also apply when there are other effects around and the state is composed of different substates for different effects, cf. the discussion of modularity in Section 2.

$$\begin{array}{l}
\mathcal{O}(\alpha, e) \triangleq \text{has_state}(\alpha) \multimap \text{wp } \alpha \{ \beta. \exists v. (e \mapsto^* v) * \llbracket \mathbb{N} \rrbracket(\beta, v) * \text{has_state}(\alpha) \} \\
\mathcal{K}(R)(\kappa, K) \triangleq \square \forall (\beta, v). R(\beta, v) \multimap \mathcal{O}(\kappa \beta, K[v]) \\
\mathcal{E}(R)(\alpha, e) \triangleq \forall (\kappa, K). \mathcal{K}(R)(\kappa, K) \multimap \mathcal{O}(\kappa \alpha, K[e]) \\
\llbracket \mathbb{N} \rrbracket(\alpha, v) \triangleq \exists n : \mathbb{N}. \alpha = \text{Ret}(n) \wedge v = n \\
\llbracket \tau_1 \rightarrow \tau_2 \rrbracket(\beta, v) \triangleq \exists f. \beta = \text{Fun}(f) \wedge \square \forall (\alpha, v'). \llbracket \tau_1 \rrbracket(\alpha, v') \multimap \mathcal{E}(\llbracket \tau_2 \rrbracket)(\text{Fun}(f) \bullet \alpha, v v') \\
\llbracket \text{cont}(\tau) \rrbracket(\beta, v) \triangleq \exists \kappa K. \beta = \text{Fun}(\text{next}(\lambda x. \text{Tick}(\kappa x))) \wedge v = \text{cont } K \wedge \mathcal{K}(\llbracket \tau \rrbracket)(\kappa, K) \\
\llbracket \Gamma \rrbracket(\rho, \gamma) \triangleq \forall (x : \tau) \in \Gamma. \llbracket \tau \rrbracket(\rho x, \gamma x) \\
\Gamma \vDash e : \tau \triangleq \square \forall (\rho, \gamma). \llbracket \Gamma \rrbracket(\rho, \gamma) \multimap \mathcal{E}(\llbracket \tau \rrbracket)(\mathbf{E}\llbracket e \rrbracket_\rho, e[\gamma])
\end{array}$$

$\mathcal{O} : \text{ERel}$

$\mathcal{K} : \text{VRel} \rightarrow \text{CRel}$

$\mathcal{E} : \text{VRel} \rightarrow \text{ERel}$

$\llbracket \tau \rrbracket : \text{VRel}$

$\text{CRel} \triangleq \text{Hom} \times \text{Ectx} \rightarrow \text{iProp}$

$\text{VRel} \triangleq \mathbf{IT}^v \times \text{Val} \rightarrow \text{iProp}$

$\text{ERel} \triangleq \mathbf{IT} \times \text{Expr} \rightarrow \text{iProp}$

Fig. 11. Logical relation for λ_{callcc} .

Adequacy and logical relation. Having established soundness, we now turn our attention to *adequacy*, which is usually much more complicated to prove.

Lemma 4. *Adequacy.* Suppose that $\emptyset \vdash e : \mathbb{N}$ and $(\mathbf{E}\llbracket e \rrbracket_\emptyset, \sigma_1) \rightsquigarrow^* (\text{Ret}(n), \sigma_2)$, for a natural number n . Then $e \mapsto^* n$.

To prove Lemma 4, we define a logical relation between syntax (λ_{callcc} programs) and semantics (GITrees denotations) using the program logic from Figure 11. To handle control effects, we use a *biorthogonal* logical relation [23], adapted from [27] for adequacy, following the Iris approach [29].

The core observational refinement $\mathcal{O}(\alpha, e)$ ensures that if α reduces to a GITree value \mathbf{IT}^v , then this value is a natural number, and e also reduces to the same number. The evaluation context relation $\mathcal{K}(P)(\kappa, K)$ relates homomorphisms and evaluation contexts when they map related arguments to expressions satisfying \mathcal{O} . The expression relation $\mathcal{E}(P)(\alpha, e)$ connects related \mathbf{IT} 's and expressions in related evaluation contexts. Types are inductively interpreted: functions relate if they map related arguments to related results, and continuations relate via the context relation. For open terms, the validity judgment $\Gamma \vDash e : \tau$ uses closing substitutions, with $e[\gamma]$ denoting applying a substitution γ to e .

The proof of adequacy relies on the fact that the interpretation of evaluation contexts are homomorphisms, which allows us to use a limited version of the bind rule:

Lemma 5. *Limited bind rule.* If $\mathcal{E}(P)(\alpha, e)$ and $\mathcal{K}(P)(\kappa, K)$, then $\mathcal{O}(\kappa \alpha, K[e])$.

With this in mind we show the fundamental lemma, stating that every well-typed expression is related to its own interpretation:

Lemma 6. *Fundamental lemma.* Let $\Gamma \vdash e : \tau$ then $\Gamma \vDash e : \tau$.

		$\Gamma; \alpha \vdash e : \tau; \beta$ $\Gamma \vdash_{\text{pure}} e : \tau$		
types	$Ty \ni \tau, \sigma, \alpha, \beta, \delta, \gamma ::= \mathbb{N} \mid \tau/\alpha \rightarrow \sigma/\beta \mid \text{cont}(\tau, \alpha)$			
expressions	$Expr \ni e ::= v \mid x \mid e_1 e_2 \mid e_1 \oplus e_2$ $\mid \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \mid \mathcal{S} x. e \mid \mathcal{D} e \mid e_1 @ e_2$			
values	$Val \ni v ::= n \mid \text{rec } f(x) = e \mid \text{cont } K$			
eval. contexts	$Ectx \ni K ::= \square \mid K[\text{if } \square \text{ then } e_1 \text{ else } e_2] \mid K[v \square] \mid K[\square e]$ $\mid K[e \oplus \square] \mid K[\square \oplus v] \mid K[\square @ v] \mid K[e @ \square]$			
	$\frac{\Gamma \vdash_{\text{pure}} e : \tau}{\Gamma; \alpha \vdash e : \tau; \alpha}$	$\frac{\Gamma, x : \text{cont}(\tau, \alpha); \sigma \vdash e : \sigma; \beta}{\Gamma; \alpha \vdash \mathcal{S} x. e : \tau; \beta}$	$\frac{\Gamma; \tau \vdash e : \tau; \sigma}{\Gamma \vdash_{\text{pure}} \mathcal{D} e : \sigma}$	$\frac{x : \tau \in \Gamma}{\Gamma \vdash_{\text{pure}} x : \tau}$
	$\frac{\Gamma, f : \sigma/\alpha \rightarrow \tau/\beta, x : \sigma; \alpha \vdash e : \tau; \beta}{\Gamma \vdash_{\text{pure}} \text{rec } f(x) = e : \sigma/\alpha \rightarrow \tau/\beta}$	$\frac{\Gamma; \gamma \vdash e_1 : \sigma/\alpha \rightarrow \tau/\beta; \delta}{\Gamma; \beta \vdash e_2 : \sigma; \gamma}$ $\Gamma; \alpha \vdash e_1 e_2 : \tau; \delta$		
	$\frac{\Gamma; \beta \vdash e_1 : \mathbb{N}; \alpha \quad \Gamma; \sigma \vdash e_2 : \tau; \beta \quad \Gamma; \sigma \vdash e_3 : \tau; \beta}{\Gamma; \sigma \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \tau; \alpha}$	$\Gamma \vdash_{\text{pure}} n : \mathbb{N}$		
	$\frac{\Gamma; \alpha \vdash e_1 : \mathbb{N}; \beta \quad \Gamma; \beta \vdash e_2 : \mathbb{N}; \sigma}{\Gamma; \alpha \vdash e_1 \oplus e_2 : \mathbb{N}; \sigma}$	$\frac{\Gamma; \sigma \vdash e_1 : \text{cont}(\tau, \alpha); \delta \quad \Gamma; \delta \vdash e_2 : \tau; \beta}{\Gamma; \sigma \vdash e_1 @ e_2 : \alpha; \beta}$		

Fig. 12. Syntax and typing rules of λ_{delim} .

Computational adequacy now follows easily from the fundamental lemma.

Proof (of Lemma 4). By Lemma 6, we have that $\emptyset \vdash e : \mathbb{N}$ implies that $\emptyset \models e : \mathbb{N}$. Now, the statement follows from Theorem 1 and the assumption that $\mathbf{E}[e]_{\emptyset}, \sigma_1 \rightsquigarrow^* \text{Ret } n, \sigma_2$.

4 Modeling Delimited Continuations

In this section we scale our approach to delimited continuations, which is a challenging example of context-dependent effects. We provide a denotational semantics for a programming language λ_{delim} with **shift/reset**, and prove its soundness and adequacy relative to an abstract machine semantics [6]. The semantics and proofs are more complex than for λ_{callcc} due to the nature of delimited continuations and associated type system. To the best of our knowledge, this represents the first formalized sound and adequate direct-style denotational semantics for delimited continuations.

4.1 Syntax and Operational Semantics of λ_{delim}

The syntax and the type system of λ_{delim} is given in Figure 12. It is similar to λ_{callcc} , but instead of **call/cc** ($- . -$) there are operators $\mathcal{D} e$ (delimit the current evaluation context, also known as **reset**) and $\mathcal{S} x. e$ (grab the current delimited continuation, and bind it to x in e , also known as **shift**).

$$\begin{array}{l}
\langle e \rangle_{\text{term}} \mapsto \langle e, \square, [] \rangle_{\text{eval}} \\
\langle K :: mk, v \rangle_{\text{mcont}} \mapsto \langle K, v, mk \rangle_{\text{cont}} \\
\langle [], v \rangle_{\text{mcont}} \mapsto \langle v \rangle_{\text{ret}} \\
\langle \square, v, mk \rangle_{\text{cont}} \mapsto \langle mk, v \rangle_{\text{mcont}} \\
\langle K[\square @ v], \text{cont } K', mk \rangle_{\text{cont}} \mapsto \langle K', v, K :: mk \rangle_{\text{cont}} \\
\langle K[e @ \square], v, mk \rangle_{\text{cont}} \mapsto \langle e, K[\square @ v], mk \rangle_{\text{eval}} \\
\langle v, K, mk \rangle_{\text{eval}} \mapsto \langle K, v, mk \rangle_{\text{cont}} \\
\langle e_0 @ e_1, K, mk \rangle_{\text{eval}} \mapsto \langle e_1, K[e_0 @ \square], mk \rangle_{\text{eval}} \\
\langle \mathcal{D} e, K, mk \rangle_{\text{eval}} \mapsto \langle e, \square, K :: mk \rangle_{\text{eval}} \\
\langle \mathcal{S} k. e, K, mk \rangle_{\text{eval}} \mapsto \langle e[K/k], \square, mk \rangle_{\text{eval}}
\end{array}$$

metacontinuations:

 $M\text{cont} \ni mk ::= [] \mid K :: mk$

abstract machine config.:

 $Config ::= \langle e, K, mk \rangle_{\text{eval}} \mid$
 $\langle K, v, mk \rangle_{\text{cont}} \mid$
 $\langle mk, v \rangle_{\text{mcont}} \mid$
 $\langle e \rangle_{\text{term}} \mid$
 $\langle v \rangle_{\text{ret}}$

Fig. 13. Operational semantics of λ_{delim} (excerpt).

The type system follows Danvy and Filinski [10], extending simply-typed λ -calculus with answer types α, β . The main typing judgment $\Gamma; \alpha \vdash e : \tau; \beta$ means: under the typing context Γ , expression e can be plugged into a context expecting a value of type τ and producing a value of type α ; in that case the resulting program will have the type β . You can think of it a computation $\Gamma \rightarrow (\tau \rightarrow \alpha) \rightarrow \beta$ under the CPS translation. Thus, the type of the (delimited) continuation corresponds to $\tau \rightarrow \alpha$, while the type of the overall expression is β . The pure typing judgment $\Gamma \vdash_{\text{pure}} e : \tau$ indicates e does not depend on the surrounded context, and is context-independent for any answer types. Expressions can change their context's answer type, as seen in the $\mathcal{S} x. e$ typing rule.

For example, suppose we extend the type system with booleans, \mathbb{B} , and add a primitive function `isprime` that does not modify answer types. That is $\emptyset; \alpha \vdash \text{isprime} : \mathbb{N}/\beta \rightarrow \mathbb{B}/\beta; \alpha$. The expression $\mathcal{D} ((\text{rec } f(x) = \text{isprime } (\mathcal{S} k. x - 1)) 2)$ is well-typed in this type system as a \mathbb{N} , even though it changes the answer type from \mathbb{B} to \mathbb{N} .

Answer types appear in both judgments and type constructors. Continuation type `cont` (τ, α) represents contexts expecting something of the type τ and producing something of the type α . Function type $\sigma/\alpha \rightarrow \tau/\beta$, in addition to the input type σ and the output type τ , record the typing of the surrounding context at the point of the function call. See [10] for details.

The operational semantics for λ_{delim} uses a CEK machine, following [6,12,5]. Selected reduction rules appear in Figure 13 (see Coq formalization for the full set of rules). The abstract machine operates on various *configurations*, which can be of several forms. The first one is the initial configuration $\langle e \rangle_{\text{term}}$, which is just a starting state for evaluating expressions. Similarly, there is a terminal configuration $\langle v \rangle_{\text{ret}}$ signifying that the program has terminated with the value v .

From the initial configuration, we go on to $\langle e, K, mk \rangle_{\text{eval}}$, which signifies that we are evaluating an expression e inside the current delimited context K , with the metacontinuation mk (a stack of continuations based on different delimiters). It

$$\begin{array}{ll}
\mathit{Ins}_{\mathbf{reset}}(X) \triangleq \blacktriangleright X & \mathit{Ins}_{\mathbf{shift}}(X) \triangleq (\blacktriangleright X \rightarrow \blacktriangleright X) \rightarrow \blacktriangleright X \\
\mathit{Ins}_{\mathbf{pop}}(X) \triangleq \blacktriangleright X & \mathit{Ins}_{\mathbf{appcont}}(X) \triangleq \blacktriangleright X \times \blacktriangleright (X \rightarrow X) \\
\mathit{Outs}_{\mathbf{reset}}(X) \triangleq \blacktriangleright X & \mathit{Outs}_{\mathbf{shift}}(X) \triangleq \blacktriangleright X \\
\mathit{Outs}_{\mathbf{pop}}(X) \triangleq 0 & \mathit{Outs}_{\mathbf{appcont}}(X) \triangleq \blacktriangleright X \\
r_{\mathbf{reset}}(e, \sigma, \kappa) = \mathbf{Some}(e, \kappa :: \sigma) & r_{\mathbf{shift}}(f, \sigma, \kappa) = \mathbf{Some}(f \ \kappa, \sigma) \\
r_{\mathbf{pop}}(e, [], _) = \mathbf{Some}(e, []) & r_{\mathbf{pop}}(e, \kappa :: \sigma, _) = \mathbf{Some}(\kappa \ e, \sigma) \\
r_{\mathbf{appcont}}((e, \kappa), \sigma, \kappa') = \mathbf{Some}(\kappa \ e, \kappa' :: \sigma) & \mathbf{P}(\beta) \triangleq \mathbf{get_val}(\beta, \mathbf{Pop}) \\
\mathbf{Reset}(e) \triangleq \mathbf{Vis}_{\mathbf{reset}}(e, \mathbf{id}) & \mathbf{Shift}(f) \triangleq \mathbf{Vis}_{\mathbf{shift}}(f, \mathbf{id}) \\
\mathbf{Appcont}(e, f) \triangleq \mathbf{Vis}_{\mathbf{appcont}}((e, f), \mathbf{id}) & \mathbf{Pop}(e) \triangleq \mathbf{Vis}_{\mathbf{pop}}(e, \lambda x. \mathbf{abort} \ x)
\end{array}$$

Fig. 14. Effects for $\lambda_{\mathbf{delim}}$.

is this configuration type which takes care of delimited control operations. The \mathbf{D} operator saves the current continuation on top of the metacontinuation, limiting the scope of $\mathcal{S} \ x. \ e$. The $\mathcal{S} \ x. \ e$ operation behaves similarly to $\mathbf{call/cc} \ (x. \ e)$, except that it prevents later control operators from capturing its evaluation context.

The last two configuration types are for dealing with continuations and metacontinuations. A configuration $\langle K, v, mk \rangle_{\mathbf{cont}}$ signifies that we are trying to plug in the value v into the context K , with the metacontinuation mk . A configuration $\langle mk, v \rangle_{\mathbf{mcont}}$ signifies that we are done with the current continuation (ending with the value v), but we still have to unwind the continuation stack mk .

4.2 Denotational Semantics of $\lambda_{\mathbf{delim}}$

Our model represents delimited continuations with effects mimicking an abstract machine, operating on semantic rather than syntactic components. The effect signature and reifiers (Figure 14) define a state with a stack of continuations, manipulated explicitly. The effect signature $E_{\lambda_{\mathbf{delim}}}$ includes four operators: $\{\mathbf{reset}, \mathbf{shift}, \mathbf{pop}, \mathbf{appcont}\}$. The signature of \mathbf{reset} simply tells us that the corresponding effect does not directly modify its argument. The auxiliary effect \mathbf{pop} , which does not have an equivalent in the surface syntax, is used to enforce unwinding of the continuation stack. As the output arity of \mathbf{pop} signifies, it does not return. We describe the importance of that below. The rest of the signature is more straightforward: \mathbf{shift} and $\mathbf{appcont}$ are defined exactly as \mathbf{callcc} and \mathbf{throw} . The semantics of these effects, in terms of reification, is more intricate. As we mentioned, the state for reification is $State = List(\blacktriangleright \mathbf{IT} \rightarrow \blacktriangleright \mathbf{IT})$.

In comparison with $\mathbf{call/cc} \ (x. \ e)$, the control operator $\mathcal{S} \ x. \ e$ does not necessarily continue from the same continuation; hence, the corresponding reifier passes the current continuation to the body, but does not return control back. The reifier for \mathbf{reset} simply saves the current continuation κ onto the stack σ . It is then the job of the \mathbf{pop} operation to restore the continuation from the stack.

$$\begin{array}{c}
\text{WP-SHIFT} \\
\frac{\text{has_state}(\sigma) \quad \triangleright(\text{has_state}(\sigma) \multimap \text{wp } \beta \{ \Phi \}) \quad \blacktriangleright \mathbf{P}(f(\blacktriangleright \kappa)) = \text{next}(\beta)}{\text{wp } \kappa(\text{Shift}(f)) \{ \Phi \}}
\end{array}
\qquad
\begin{array}{c}
\text{WP-RESET} \\
\frac{\text{has_state}(\sigma) \quad \triangleright(\text{has_state}(\blacktriangleright \kappa :: \sigma) \multimap \text{wp } \mathbf{P}(e) \{ \Phi \})}{\text{wp } \kappa(\text{Reset}(\text{next}(e))) \{ \Phi \}}
\end{array}$$

$$\begin{array}{c}
\text{WP-POP} \\
\frac{\text{has_state}(\sigma) \quad \kappa' = \kappa \text{ if } \sigma = \kappa :: \sigma' \text{ and id otherwise} \quad \triangleright(\text{has_state}(\text{tail}(\sigma)) \multimap \text{wp } \kappa'(v) \{ \Phi \})}{\text{wp } \mathbf{P}(v) \{ \Phi \}}
\end{array}
\qquad
\begin{array}{c}
\text{WP-APPCONT} \\
\frac{\text{has_state}(\sigma) \quad \triangleright(\text{has_state}(\blacktriangleright \kappa :: \sigma) \multimap \text{wp } \beta \{ \Phi \}) \quad \blacktriangleright \kappa'(e) = \text{next}(\beta)}{\text{wp } \kappa(\text{Appcont}(e, \kappa')) \{ \Phi \}}
\end{array}$$

Fig. 15. Weakest precondition rules for delimited continuations.

$$\begin{aligned}
\mathbf{E}[\mathcal{D} \ e]_{\rho} &= \text{Reset}(\mathbf{P}(\mathbf{E}[e]_{\rho})) \\
\mathbf{E}[\mathcal{S} \ x. \ e]_{\rho} &= \text{Shift}(\mathbf{P} \circ (\lambda \kappa. \mathbf{E}[e]_{\rho, x \mapsto \text{Fun}(\text{next}(\lambda y. \text{Tau}(\kappa(\text{next}y)))}))) \\
\mathbf{E}[e_1 \ @ \ e_2]_{\rho} &= \text{get_val}(\mathbf{E}[e_2]_{\rho}, \lambda x. \text{get_fun}(\mathbf{E}[e_1]_{\rho}, \lambda y. \text{Appcont}(\text{next}(x), y))) \\
\mathbf{V}[\text{cont } K]_{\rho} &= \text{Fun}(\text{next}(\lambda x. \text{Tick}(\mathbf{P}(\mathbf{K}[K]_{\rho} \ x)))) \\
\mathbf{K}[K[\square \ @ \ v]]_{\rho} &= \lambda x. \mathbf{K}[K]_{\rho}(\mathbf{E}[x \ @ \ v]_{\rho}) \\
\mathbf{M}[mk]_{\rho} &= \text{map}(\lambda k. \mathbf{P} \circ \mathbf{K}[k]_{\rho})mk \\
\mathbf{S}[\langle e, K, mk \rangle_{\text{eval}}]_{\rho} &= (\mathbf{P}(\mathbf{E}[K[e]_{\rho}], \mathbf{M}[mk]_{\rho})) \\
\mathbf{S}[\langle K, v, mk \rangle_{\text{cont}}]_{\rho} &= (\mathbf{P}(\mathbf{E}[K[v]_{\rho}], \mathbf{M}[mk]_{\rho})) \\
\mathbf{S}[\langle mk, v \rangle_{\text{mcont}}]_{\rho} &= (\mathbf{P}(\mathbf{V}[v]_{\rho}), \mathbf{M}[mk]_{\rho}) \\
\mathbf{S}[\langle e \rangle_{\text{term}}]_{\rho} &= (\mathbf{P}(\mathbf{E}[e]_{\rho}), []) \\
\mathbf{S}[\langle v \rangle_{\text{ret}}]_{\rho} &= (\mathbf{V}[v]_{\rho}, [])
\end{aligned}$$

Fig. 16. Denotational semantics for a calculus with delimited control (selected clauses).

The reifier for `shift` is similar to that of `callcc`, except that it removes the current continuation entirely. The reifier for `appcont`, in comparison with `throw`, does not simply pass control, but also saves the current continuation on the stack. This corresponds to the fact that whenever a delimited continuation is invoked, the result is wrapped in a `reset`; that is done to prevent the continuation from escaping the delimiter. As part of instantiating `GITrees` with these effects, we obtain the specialized program logic rules shown in Figure 15. We will use those rules later for defining a logical relation between the syntax and the semantics of λ_{delim} . As mentioned above, we will use `Pop` to unwind the continuation stack and restore the continuation after finishing with a `reset`. This means that we will need to insert explicit calls to `Pop` in the interpretation of λ_{delim} . For these purposes, we use an abbreviation $\mathbf{P}(\beta)$, which first evaluates β to a value, and then executes the `pop` operation.

The interpretation of λ_{delim} uses this auxiliary function and is given in Figure 16. Similarly to the operational semantics, the interpretation is divided into five categories. First, we have $\mathbf{E}[-]$ and $\mathbf{V}[-]$ for the interpretation of expressions and values, which is what we need for the surface syntax. All of those interpretations return GITrees. Note that in the interpretation of \mathcal{D} – we insert explicit calls to \mathbf{P} , and similarly in the interpretation of continuations.

The other group of interpretations, $\mathbf{K}[-]$, $\mathbf{M}[-]$ and $\mathbf{S}[-]$, are for interpreting continuations, metacontinuations, and other configurations; these are used for showing soundness (preservation of operational semantics by the interpretation). The interpretation $\mathbf{K}[-]$ of continuations returns a semantic continuation (a function $\mathbf{IT} \rightarrow \mathbf{IT}$). Similarly, the interpretations $\mathbf{M}[-]$ (resp. $\mathbf{S}[-]$) of metacontinuations (resp. configurations) returns a stack of semantic continuations (resp. a semantic configuration).

We now show that our interpretation is sound w.r.t. the abstract machine semantics. For this we prove lemmas similar to Lemma 2, and put them to use in the soundness theorem:

Theorem 2. *Soundness.* Let $c_0, c_1 \in \text{Config}$ and suppose $c_0 \mapsto c_1$. Then $\mathbf{S}[c_0] \rightsquigarrow^* \mathbf{S}[c_1]$.

4.3 Logical Relation and Adequacy

We now show that our denotational semantics is adequate with regards to the abstract machine semantics. Specifically, we show the following result:

Theorem 3. *Adequacy.* Suppose $\emptyset; \mathbb{N} \vdash e : \mathbb{N}; \mathbb{N}$ is a well-typed term, and that $(\mathbf{P}(\mathbf{E}[e]_{\emptyset}), \square) \rightsquigarrow^* (\text{Ret}(n), \sigma)$ for a natural number n and a metacontinuation σ . Then $\langle e \rangle_{\text{term}} \mapsto^* \langle n \rangle_{\text{ret}}$.

We prove adequacy using a logical relation. It relates expressions to their interpretations and also connects syntactic and semantic configurations. The logical relation is shown in Figure 17. It is again a form of biorthogonal logical relation, with the main focus being the observational refinement \mathcal{O} : two configurations are related if they reduce to the same natural number. This coincides with what we want to show in Theorem 3. To facilitate this we then lift \mathcal{O} to the levels of metacontinuations, continuations and expressions. The relation $\mathcal{M}(P)$, where $P : \text{VRel}$ is a relation on values, states that two metacontinuations are related if, whenever we plug in P -related values, the resulting configurations become \mathcal{O} -related. Both \mathcal{M} and \mathcal{O} are then used to define the relation between semantic and syntactic continuations. The relation $\mathcal{K}(Q, P)$, where $P, Q : \text{VRel}$ are relations on values, states that two continuations are related if, whenever we plug them into Q -related metacontinuations with P -related values, the resulting configurations become \mathcal{O} -related. Finally, we use \mathcal{K} , \mathcal{M} , and \mathcal{O} to define the relation between GITrees and λ_{delim} terms. The relation $\mathcal{E}(P, Q, R)$, where $P, Q, R : \text{VRel}$ are relations on values, states that β is related to e if, whenever we plug them into (P, Q) -related continuations and an R -related metacontinuation, the resulting configurations become \mathcal{O} -related.

$$\begin{aligned}
\text{SynConf} &\triangleq \text{Expr} \times \text{Ectx} \times \text{Mcont} & \text{VRel} &\triangleq \mathbf{IT}^v \times \text{Val} \rightarrow i\text{Prop} \\
\text{SemConf} &\triangleq \mathbf{IT} \times \text{Hom} \times \text{List}(\text{Hom}) & \mathcal{O} &: \text{ConfRel} \\
\text{ConfRel} &\triangleq \text{SemConf} \times \text{SynConf} \rightarrow i\text{Prop} & \mathcal{M} &: \text{VRel} \rightarrow \text{MRel} \\
\text{MRel} &\triangleq \text{list Hom} \times \text{Mcont} \rightarrow i\text{Prop} & \mathcal{K} &: \text{VRel} \rightarrow \text{VRel} \rightarrow \text{CRel} \\
\text{CRel} &\triangleq \text{Hom} \times \text{Ectx} \rightarrow i\text{Prop} & \mathcal{E} &: \text{VRel} \rightarrow \text{VRel} \rightarrow \text{VRel} \rightarrow \text{ERel} \\
\text{ERel} &\triangleq \mathbf{IT} \times \text{Expr} \rightarrow i\text{Prop} & \llbracket \tau \rrbracket &: \text{VRel}
\end{aligned}$$

$$\begin{aligned}
\mathcal{O}((\alpha, \kappa, \sigma), (e, K, mk)) &\triangleq \text{has_state}(\sigma) \text{ -* } \\
&\quad \text{wp } \mathbf{P}(\kappa \alpha) \{ \beta. \exists v. \langle (e, K, mk) \rangle_{\text{eval}} \mapsto^* \langle v \rangle_{\text{ret}} \} \\
&\quad * (\beta, v) \in \llbracket \mathbb{N} \rrbracket * \text{has_state}(\beta) \} \\
\mathcal{M}(P)(\sigma, mk) &\triangleq \forall (\alpha, v). P(\alpha, v) \text{ -* } \mathcal{O}((\alpha, \iota, \sigma), (v, \square, mk)) \\
\mathcal{K}(Q, P)(\kappa, K) &\triangleq \square \forall (\alpha, v). Q(\alpha, v) \text{ -* } \forall (\sigma, mk). \mathcal{M}(P)(\sigma, mk) \text{ -* } \\
&\quad \mathcal{O}((\alpha, \kappa, \sigma), (v, K, mk)) \\
\mathcal{E}(P, Q, R)(\beta, e) &\triangleq \forall (\kappa, K). \mathcal{K}(P, Q)(\kappa, K) \text{ -* } \forall (\sigma, mk). \mathcal{M}(R)(\sigma, mk) \text{ -* } \\
&\quad \mathcal{O}((\beta, \kappa, \sigma), (e, K, mk)) \\
\llbracket \mathbb{N} \rrbracket(\beta, v) &\triangleq \exists n \in \mathbb{N}. \beta = \text{Ret}(n) \wedge v = n \\
\llbracket \tau / \alpha \rightarrow \sigma / \beta \rrbracket(\theta, v) &\triangleq \exists F. \theta = \text{Fun}(F) \wedge \square \forall (\eta, w). \llbracket \tau \rrbracket(\eta, w) \text{ -* } \\
&\quad \mathcal{E}(\llbracket \sigma \rrbracket, \llbracket \alpha \rrbracket, \llbracket \beta \rrbracket)(\theta \bullet \eta, v w) \\
\llbracket \text{cont}(\tau, \alpha) \rrbracket(\beta, v) &\triangleq \exists \kappa. K. \beta = \text{Fun}(\text{next}(\lambda x. \text{Tick}((\mathbf{P} \circ \kappa) x))) \wedge v = \text{cont } K \wedge \\
&\quad \mathcal{K}(\llbracket \tau \rrbracket, \llbracket \alpha \rrbracket)(\kappa, K) \\
\llbracket \Gamma \rrbracket(\rho, \gamma) &\triangleq \forall (x : \tau \in \Gamma). \square \forall \Phi. \mathcal{E}(\llbracket \tau \rrbracket, \Phi, \Phi)(\rho(x), \gamma(x)) \\
\Gamma \vDash_{\text{pure}} e : \tau &\triangleq \square \forall (\rho, \gamma). \llbracket \Gamma \rrbracket(\rho, \gamma) \text{ -* } \forall \Phi. \mathcal{E}(\llbracket \tau \rrbracket, \Phi, \Phi)(\mathbf{E}[e]_{\rho}, e[\gamma]) \\
\Gamma; \alpha \vDash e : \tau; \beta &\triangleq \square \forall (\rho, \gamma). \llbracket \Gamma \rrbracket(\rho, \gamma) \text{ -* } \mathcal{E}(\llbracket \tau \rrbracket, \llbracket \alpha \rrbracket, \llbracket \beta \rrbracket)(\mathbf{E}[e]_{\rho}, e[\gamma])
\end{aligned}$$

Fig. 17. Logical relation for λ_{delim} .

The relations \mathcal{E} and \mathcal{K} are used to give semantics $\llbracket \tau \rrbracket$ to types. The idea is that $\mathcal{E}(\llbracket \tau \rrbracket, \llbracket \alpha \rrbracket, \llbracket \beta \rrbracket)$ relates terms $\emptyset; \alpha \vdash e : \tau; \beta$ to their semantic counterparts. This is then used, as expected for logical relations, for defining the logical relation for function types and for open terms. The relation $\llbracket \Gamma \rrbracket(\rho, \gamma)$ relates the semantic environment $\rho : \text{Var} \rightarrow \mathbf{IT}$ to the syntactic substitution $\gamma : \text{Var} \rightarrow \text{Expr}$; they are related if they map the same variables to related GITrees/expressions. Then we say that an expression e is semantically valid, $\Gamma; \alpha \vdash e : \tau; \beta$, if its interpretation $\mathbf{E}[e]_{\rho}$ is related to $e[\gamma]$ under related substitutions ρ, γ . Note that if we ignore the answer types we can see that the logical relation exhibits a lot of similarities to the logical relation we gave in Section 3.3, and follows the same roadmap.

For this logical relation we obtain the fundamental property, which we will use for the proof of adequacy.

Lemma 7. *Fundamental lemma.* Let $\Gamma; \alpha \vdash e : \tau; \beta$ then $\Gamma; \alpha \vDash e : \tau; \beta$; and if $\Gamma \vdash_{\text{pure}} e : \tau$ then $\Gamma \vDash_{\text{pure}} e : \tau$.

$$\begin{array}{l}
 \text{types } Ty \ni \tau ::= \mathbb{N} \mid \mathbf{1} \mid \tau \rightarrow \sigma \mid \text{ref}(\tau) \\
 \text{expressions } Expr \ni e ::= x \mid () \mid e_1 e_2 \mid e_1 \oplus e_2 \mid n \mid \lambda x. e \\
 \phantom{\text{expressions }} \mid \ell \mid \text{ref}(e) \mid !e \mid e_1 \leftarrow e_2 \mid \text{embed } e
 \end{array}
 \quad
 \frac{\emptyset \vdash_{\text{pure}} e : \mathbb{N}}{\Gamma \vdash \text{embed } e : \mathbb{N}}$$

Fig. 18. Syntax and the new typing rule of the λ_{embed} .

Proof (of Theorem 3). Note that the empty (meta)continuation is related to its denotation: $\mathcal{K}(P, P)(\text{id}, \square)$ and $\mathcal{M}(P)([], [])$ hold for any relation P .

With this, we instantiate $\emptyset; \mathbb{N} \models e : \mathbb{N}; \mathbb{N}$ (that we get from Lemma 7) with the empty continuation/metacontinuation, and get the observational refinement between e and $\mathbf{E}[e]$. \square

This completes our treatment of denotational semantics of λ_{delim} . The next section examines interoperability of delimited continuations and other effects.

5 Modeling Interoperability Between Languages

A key advantage of using (G)Itrees for semantics is that they can provide a common framework for multi-language interaction. This section presents a case study on the interaction between the languages λ_{embed} (with higher-order store effects) and λ_{delim} (with delimited continuations). Specifically, we allow embedding closed λ_{delim} programs into λ_{embed} , and equip λ_{embed} with a type system that guarantees safe interoperability

The embedding we provide is restrictive, preventing programs with delimited continuations from accessing outer-language continuations. We leave developing a more permissive type system for future work. At the end of the section we give an example of how to verify a more involved interaction of effects, albeit without the type system.

In this section we reuse the semantics of λ_{delim} from the previous section and higher-order store effects from Section 2. For λ_{delim} , we reuse the semantics from the previous section and higher-order store reifiers from [13]. In this section, we use **magenta** to explicitly highlight programs written in λ_{delim} , and for the interpretation functions of the denotational model of λ_{delim} .

Language λ_{embed} . λ_{embed} is a λ -calculus with base types \mathbb{N} and $\mathbf{1}$, references types $\text{ref}(\tau)$ and function types $\tau \rightarrow \sigma$, with syntax given in Figure 18. Additionally, it includes a construct **embed** e for embedding λ_{delim} programs. The typing rules are all standard, except for the new typing rule for the embedding.

The idea behind the rule is that we can embed an expression from λ_{delim} if it is a “pure” expression that can evaluate to a natural number. The use of pure typing judgment for the embedded program ensures that it does not to alter the answer type. This means that we can treat an embedded expression as a “complete” program, that does not require outer continuation delimiters, even though it may rely on delimited continuations internally. Those restrictions are

$\mathbf{E}[-] : Expr \rightarrow (Var \rightarrow \mathbf{IT}) \rightarrow \mathbf{IT}$	$\mathbf{E}[\mathbf{ref}(e)]_\rho = \text{get_val}(\mathbf{E}[e]_\rho, \lambda x. \text{Alloc}(x, \text{Ret}))$
$\mathbf{E}[\mathbf{embed } e]_\rho = \text{Reset}(\text{next}(\mathbf{E}[e]_\emptyset))$	$\mathbf{E}[\mathbf{!} e]_\rho = \text{get_val}(\mathbf{E}[e]_\rho, \lambda x. \text{Read}(x))$
$\mathbf{E}[x]_\rho = \rho(x)$	$\mathbf{E}[e_1 \leftarrow e_2]_\rho = \text{get_val}(\mathbf{E}[e_2]_\rho, \lambda x.$
$\mathbf{E}[\ell]_\rho = \text{Ret}(\ell)$	$\text{get_ret}(\mathbf{E}[e_1]_\rho, \lambda y. \text{Write}(y, x)))$

Fig. 19. Denotational semantics of λ_{embed} (selected clauses).

$\text{VRel} \triangleq \mathbf{IT}^v \rightarrow iProp$	$\llbracket \tau \rrbracket : \text{VRel}$
$\text{ERel} \triangleq \mathbf{IT} \rightarrow iProp$	$\llbracket \mathbf{1} \rrbracket(\beta) \triangleq \beta = ()$
$\mathcal{O} : \text{VRel} \rightarrow \text{ERel}$	$\llbracket \mathbf{N} \rrbracket(\beta) \triangleq \exists n. \beta = \text{Ret}(n)$
$\mathcal{O}(P)(\beta) \triangleq \text{clwp } \beta \{x. P \ x * \text{has_state}(\llbracket \cdot \rrbracket)\}$	$\llbracket \tau \rightarrow \sigma \rrbracket(\beta) \triangleq \exists F. \beta = \text{Fun}(F) \wedge$
$\mathcal{E} : \text{VRel} \rightarrow \text{ERel}$	$\square \forall \beta. \llbracket \tau \rrbracket(\beta) \dashv\ast$
$\mathcal{E}(P)(\beta) \triangleq \text{heap_ctx} \dashv\ast \text{has_state}(\llbracket \cdot \rrbracket) \dashv\ast$	$\mathcal{E}(\llbracket \sigma \rrbracket)(\text{Fun}(F) \bullet \beta)$
$\mathcal{O}(P)(\beta)$	$\llbracket \mathbf{ref}(\tau) \rrbracket(\beta) \triangleq \exists \ell. \beta = \text{Ret}(\ell) \wedge$
	$\boxed{\exists \nu. \ell \mapsto \nu * \llbracket \tau \rrbracket(\nu)}$
$\llbracket \Gamma \rrbracket(\rho) \triangleq \forall (x : \tau \in \Gamma). \square \mathcal{E}(\llbracket \tau \rrbracket)(\gamma \ x)$	$\Gamma \vDash e : \tau \triangleq \forall \rho. \llbracket \Gamma \rrbracket(\rho) \dashv\ast \mathcal{E}(\llbracket \tau \rrbracket)(\mathbf{E}[e]_\rho)$

Fig. 20. Logical relation for λ_{embed} .

crucial for the type safety of the embedding. The typing guarantees that e does not expect any additional delimiters, but it does not, by itself guarantee that any continuations in e escape the embedding boundary. To prevent that we enforce the continuation delimiter along the embedding boundary in the interpretation of embedded expressions.

Denotational model of λ_{embed} . For denotational semantics of λ_{embed} , we start by defining reifiers for the effect signature, which includes higher-order store operations (allocating, reading, and storing references) as E_{state} , and effects related to delimited continuations ($E_{\lambda_{\text{delim}}}$). Then the combined effect signature is $E_{\lambda_{\text{delim}}} \times E_{\text{state}}$, and thus we also let $State \triangleq State_{\text{delim}} \times State_{\text{state}}$, and reifiers are defined component-wise. Figure 19 shows the key parts of the denotational semantics. For most of the syntactic constructs we give the standard interpretation. For $\mathbf{embed } e$ we use the interpretation $\mathbf{E}[-]$ for λ_{delim} from Section 4.2, and explicitly wrap the resulting GITree in a Reset . This continuation delimiter acts as a sort of *glue code* to protect the rest of the program from being captured by control operators from the embedded λ_{delim} program.

To show type safety of λ_{embed} , we construct a logical relation (shown in Figure 20), which is similar to the other logical relations we considered in this paper, mainly different in the observation relation \mathcal{O} . Given that the type system for λ_{embed} effectively prevents expressions of λ_{delim} to access contexts from λ_{embed} , we refine the observation relation to get access to a version of the WP-HOM rule

for expression interpretations of well-typed programs of λ_{embed} , which we do not have in general, as discussed in Section 3.

Instead of the standard weakest precondition wp , we utilize a *context-local weakest precondition* clwp which bakes-in the bind rules [27].

Definition 2. *Context-local weakest precondition (clwp) is defined as follows:*
 $\text{clwp } \alpha \{ \Phi \} \triangleq \forall \kappa (\Psi : \mathbf{IT}^v \rightarrow iProp). (\forall v. \Phi v \multimap \text{wp } (\kappa v) \{ \Psi \}) \rightarrow \text{wp } (\kappa \alpha) \{ \Psi \}$

Note that clwp always implies wp and validates a form of the bind rule:

$$\text{clwp } \alpha \{ \beta. \text{clwp } (\kappa \beta) \{ \Phi \} \} \vdash \text{clwp } (\kappa \alpha) \{ \Phi \}$$

for any homomorphism κ . We use the clwp in the definition of observational refinement in the model. Observational refinement asserts that if the evaluation of a top-level expression of λ_{embed} starts with an empty continuation stack, then the evaluation does not introduce new elements into the continuation stack. By using clwp we get a semantical bind lemma, which can be seen as a version of WP-HOM for semantically valid expressions of λ_{embed} .

Lemma 8. *Semantical bind.* $\mathcal{E}(\lambda x : \mathbf{IT}^v. \mathcal{E}(P)(\kappa x))(\beta)$ implies $\mathcal{E}(P)(\kappa \beta)$ for any homomorphism κ .

As before, we obtain fundamental lemma and denotational type soundness.

Lemma 9. *Fundamental lemma.* Let $\Gamma \vdash e : \tau$ then $\Gamma \vDash e : \tau$.

Lemma 10. *Denotational type soundness.* Let $\emptyset \vdash e : \tau$ and $(\mathbf{E}[e]_{\emptyset}, []) \rightsquigarrow^* (\alpha, \sigma)$, then $(\exists \beta \sigma'. (\alpha, \sigma) \rightsquigarrow (\beta, \sigma')) \vee (\alpha \in \mathbf{IT}^v)$.

Unrestricted interaction of delimited continuations and higher-order state. Even though the type system we considered here is restrictive, we can still reason about unrestricted interactions of events in the “untyped” setting. Here we show an example of such an unrestricted interaction, and demonstrate how to reason about context-dependent and context-independent effects at the same time. While this kind of interactions is forbidden by our type system, we can still write and prove meaningful specifications for such programs. Consider the program in Figure 21, written in GITrees directly. The function `prog` utilizes both delimited continuations and state. It takes a reference y as its argument and begins by allocating the value 1 in the store at reference x . Then it captures the continuation `get_ret(y, (λl . Let $p = \text{NatOp}_+(\text{Read}(l), \dots)$ in $\text{Write}(l, p)$))` as k . Invoking the continuation k with a number n increments the current value of y by n . The program then invokes this continuation twice. First with the original value of x . Then, with an incremented value of x . Since the starting value of x is 1, the reference y is incremented first by 1 and then by 2. We capture this behavior in the specification for `prog` stated in Figure 21.

It is important to note that this program features a bidirectional interaction between state and continuations. Specifically, the body of `Shift` involves state operations, while the result of `Shift` is subsequently used to increment a value in the heap. As we have seen, while this type of interaction is not allowed in the

<code>prog</code>	\triangleq	<code>Fun(next(λy.</code>	
<code>Let $x = \text{Alloc}(\text{Ret}(1))$ in</code>		Initial offset value	
<code>Let $n = \text{Shift}(\lambda k.\text{next}(\text{Appcont}'(\text{Read}(x), k);$</code>		Capture continuation as k	
<code>Let $m =$</code>		First call to k with the init. value of x	
<code>NatOp₊(Read(x), Ret(1)) in Write(x, m);</code>		$x := x + 1;$	
<code>Appcont'(Read(x), k))</code>		Second call to k with updated x	
<code>in get_ret($y, (\lambda l.\text{Let } p =$</code>			
<code>NatOp₊(Read(l), n) in Write(l, p))))</code>		$y := y + n;$	
where $\text{Appcont}'(x, y) \triangleq \text{Appcont}(\text{next}(x), \text{next}(\text{Tau} \circ y \circ \text{next}))$			
$\frac{\text{heap_ctx} \quad \text{has_state}(\sigma) \quad y \mapsto \text{Ret}(n)}{\text{wp}(\text{Reset}(\text{next}(\text{prog} \bullet \text{Ret}(y)))) \{y \mapsto \text{Ret}(n+3) * \text{has_state}(\sigma)\}}$			

Fig. 21. Example program with delimited continuations and state and its specification.

type system, we can still reason about them in program logic. We stipulate that our proposed type system could potentially be extended to support embedding “pure” functions, allowing for bi-directional interaction between the two languages. We believe that such an extension would require implementing answer-type polymorphism, following the approach of Asai and Kameyama [2].

6 Discussion and Related Work

We conclude the paper by discussing related work and future directions.

This paper extends GITrees to handle context-dependent effects, which allows us to model higher-order languages with control operators like `call/cc` ($x. e$) and `S` $x. e$. We showed this extension supports interoperability between languages with different context-dependent effects, while preserving reasoning about context-independent effects. Our approach leverages the native support for higher-order functions and effects in GITrees. This differs from the first-order effect representation of effects in ITrees [33], which would require explicitly first-order representation of functions and continuations, if we want to model first-class continuation. Such model would mix syntactic and semantic concerns, which is part of what we are trying to avoid by working with (G)ITrees.

Another difference with ITrees is our approach to reasoning. While ITrees use bisimulation-based equational theory, we follow GITrees in using tailored program logics and defining refinements. Our logic are expressive enough to define logical relations and carry out computational adequacy proofs. In future work, it would be interesting to develop techniques for reasoning about weaker notions of equality than the basic equational theory that GITrees comes equipped with, see the more extensive discussion of this point in [13].

The “classical” domain theory remains an important source of inspiration and ideas for our development, and we want to mention some of the related work along those lines. Cartwright and Felleisen [9] introduced a framework of

extensible direct models for constructing modular denotational semantics of programming languages. Their framework centers on an abstract notion of *resources* for representing effects (such as store or continuations) and a central `admin` function that manages these resources. Each language extension defines both the types (domains) of additional values and resources, and specifies the *actions* that the `admin` function can perform on these resources. Building on this framework, Sitaram and Felleisen [26] demonstrated that such models can provide direct-style fully abstract semantics for control operators. Their approach interprets effects, including continuations, by delegating them to a top-level handler. Our work adopts several key ideas from this line of research but reformulates them in the context of GITrees rather than classical domain theory. In our framework, effect signatures define resources, reifiers specify actions, and the reduction relation serves as the central authority dispatching the effects. The transition to GITrees enables us to formalize the extensibility of this approach in a practical manner, and it allows us to develop program logics where “resources” (as above) become resources in separation logic.

Compared to other programming languages paradigms and effects, type systems and logical relations for delimited continuations have not been studied as comprehensively. The original type system for `shift/reset` is due to Danvy and Filinski [10], where they employ *answer-type modification*. Materzok and Biernacki [22] generalized this type system to account for more involved control operators `shift0/reset0`; an alternative substructural type system for these operators was designed by Kiselyov and Shan [17]. Dyvbig, Peyton Jones, and Sabry [11] provide a typed monadic account of CPS for delimited continuations with dynamic prompt generation. Asai and Kameyama [2] present a polymorphic variant of the Danvy and Filinski’s type system.

Biernacka and Biernacki [5] prove termination for a language with $\mathcal{S} \times . e$ and $\mathcal{D} e$ (but without recursion) using logical relations based on abstract machine semantics. The shape of their logical relation is similar to our logical relation used for showing adequacy in that they also have relations for configurations, metacontinuations, etc. Asai [1] uses type-directed logical relations to verify a direct-style specializer (partial evaluator) for a language with $\mathcal{S} \times . e$ and $\mathcal{D} e$, proving correctness against evaluation-context based operational semantics. In contrast with those works, we define our logical relations on denotations, using the semantics of GITrees and the derived program logics.

In our interoperability example (Section 5) we showed type safety of a combined language, with respect to denotational semantics. In future work we would like to examine other properties: for example, that λ_{delim} expressions cannot disrupt λ_{embed} ’s control flow, perhaps establishing some form of well-bracketedness using ideas from [28].

We would also like to study other context-dependent effects like exceptions, handlers, and algebraic effects [24,4,32,30,3]. In particular, it would be interesting to give a denotational semantics to a language with handlers, derive program logic proof rules for it, and compare the resulting program logic to the one in [31].

References

1. Asai, K.: Logical relations for call-by-value delimited continuations. In: van Eekelen, M.C.J.D. (ed.) Revised Selected Papers from the Sixth Symposium on Trends in Functional Programming, TFP 2005, Tallinn, Estonia, 23-24 September 2005. Trends in Functional Programming, vol. 6, pp. 63–78. Intellect (2005)
2. Asai, K., Kameyama, Y.: Polymorphic Delimited Continuations. In: Shao, Z. (ed.) Programming Languages and Systems. pp. 239–254. Springer. https://doi.org/10.1007/978-3-540-76637-7_16
3. Bach Poulsen, C., van der Rest, C.: Hefty Algebras: Modular Elaboration of Higher-Order Algebraic Effects. Proceedings of the ACM on Programming Languages **7**(POPL), 62:1801–62:1831 (Jan 2023). <https://doi.org/10.1145/3571255>
4. Bauer, A., Pretnar, M.: Programming with algebraic effects and handlers. Journal of Logical and Algebraic Methods in Programming **84**(1), 108–123 (Jan 2015). <https://doi.org/10.1016/j.jlamp.2014.02.001>
5. Biernacka, M., Biernacki, D.: Context-based proofs of termination for typed delimited-control operators. In: Porto, A., López-Fraguas, F.J. (eds.) Proceedings of the 11th International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming, September 7-9, 2009, Coimbra, Portugal. pp. 289–300. ACM (2009). <https://doi.org/10.1145/1599410.1599446>, <https://doi.org/10.1145/1599410.1599446>
6. Biernacka, M., Biernacki, D., Danvy, O.: An operational foundation for delimited continuations in the CPS hierarchy. Log. Methods Comput. Sci. **1**(2) (2005). [https://doi.org/10.2168/LMCS-1\(2:5\)2005](https://doi.org/10.2168/LMCS-1(2:5)2005), [https://doi.org/10.2168/LMCS-1\(2:5\)2005](https://doi.org/10.2168/LMCS-1(2:5)2005)
7. Biernacki, D., Piróg, M., Polesiuk, P., Sieczkowski, F.: Abstracting algebraic effects. Proc. ACM Program. Lang. **3**(POPL) (jan 2019). <https://doi.org/10.1145/3290319>, <https://doi.org/10.1145/3290319>
8. Birkedal, L., Møgelberg, R.E., Schwinghammer, J., Støvring, K.: 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), [https://doi.org/10.2168/LMCS-8\(4:1\)2012](https://doi.org/10.2168/LMCS-8(4:1)2012)
9. Cartwright, R., Felleisen, M.: Extensible denotational language specifications. In: Hagiya, M., Mitchell, J.C. (eds.) Theoretical Aspects of Computer Software. pp. 244–272. Springer, Berlin, Heidelberg (1994). https://doi.org/10.1007/3-540-57887-0_99
10. Danvy, O., Filinski, A.: A functional abstraction of typed contexts. Tech. rep., DIKU – Computer Science Department, University of Copenhagen (1989)
11. Dyvbig, R.K., Peyton Jones, S., Sabry, A.: A monadic framework for delimited continuations **17**(6), 687–730. <https://doi.org/10.1017/S0956796807006259>, <https://www.cambridge.org/core/journals/journal-of-functional-programming/article/monadic-framework-for-delimited-continuations/D99D1394370DFA8EA8428D552B5D8E7E>
12. Felleisen, M., Friedman, D.P.: Control operators, the secd-machine, and the λ -calculus. In: Wirsing, M. (ed.) Formal Description of Programming Concepts - III: Proceedings of the IFIP TC 2/WG 2.2 Working Conference on Formal Description of Programming Concepts - III, Ebberup, Denmark, 25-28 August 1986. pp. 193–222. North-Holland (1987)
13. Frumin, D., Timany, A., Birkedal, L.: Modular denotational semantics for effects with guarded interaction trees. Proc. ACM Program. Lang. **8**(POPL), 332–361 (2024). <https://doi.org/10.1145/3632854>, <https://doi.org/10.1145/3632854>

14. Iris team: The Iris Project website and Coq development (2025), <https://iris-project.org/>
15. Jung, R., Krebbers, R., Jourdan, J., Bizjak, A., Birkedal, L., Dreyer, D.: Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.* **28**, e20 (2018). <https://doi.org/10.1017/S0956796818000151>, <https://doi.org/10.1017/S0956796818000151>
16. King, A.: Delimited continuation primops. <https://github.com/ghc-proposals/ghc-proposals/blob/master/proposals/0313-delimited-continuation-primops.rst> (2021), accessed: 2024-06-27
17. Kiselyov, O., Shan, C.c.: A Substructural Type System for Delimited Continuations. In: Della Rocca, S.R. (ed.) *Typed Lambda Calculi and Applications*. pp. 223–239. Springer, Berlin, Heidelberg (2007). https://doi.org/10.1007/978-3-540-73228-0_17
18. Koh, N., Li, Y., Li, Y., Xia, L.y., Beringer, L., Honoré, W., Mansky, W., Pierce, B.C., Zdancewic, S.: 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*. pp. 234–248. CPP 2019, Association for Computing Machinery, New York, NY, USA (Jan 2019). <https://doi.org/10.1145/3293880.3294106>
19. Krebbers, R., Timany, A., Birkedal, L.: Interactive proofs in higher-order concurrent separation logic. In: Castagna, G., Gordon, A.D. (eds.) *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*. pp. 205–217. ACM (2017). <https://doi.org/10.1145/3009837.3009855>, <https://doi.org/10.1145/3009837.3009855>
20. Leijen, D.: Koka: Programming with row polymorphic effect types. In: Levy, P.B., Krishnaswami, N. (eds.) *Proceedings 5th Workshop on Mathematically Structured Functional Programming, MSFP@ETAPS 2014, Grenoble, France, 12 April 2014*. EPTCS, vol. 153, pp. 100–126 (2014). <https://doi.org/10.4204/EPTCS.153.8>, <https://doi.org/10.4204/EPTCS.153.8>
21. Leijen, D.: Structured asynchrony with algebraic effects. In: Lindley, S., Yorgey, B.A. (eds.) *Proceedings of the 2nd ACM SIGPLAN International Workshop on Type-Driven Development, TyDe@ICFP 2017, Oxford, UK, September 3, 2017*. pp. 16–29. ACM (2017). <https://doi.org/10.1145/3122975.3122977>, <https://doi.org/10.1145/3122975.3122977>
22. Materzok, M., Biernacki, D.: Subtyping delimited continuations **46**(9), 81–93. <https://doi.org/10.1145/2034574.2034786>, <https://doi.org/10.1145/2034574.2034786>
23. Pitts, A.M.: Typed operational reasoning. In: Pierce, B.C. (ed.) *Advanced Topics In Types And Programming Languages*, chap. 7, pp. 245–289. The MIT Press (2004)
24. Plotkin, G.D., Pretnar, M.: 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)
25. Silver, L., He, P., Cecchetti, E., Hirsch, A.K., Zdancewic, S.: *Semantics for Noninterference with Interaction Trees* (2023)
26. Sitaram, D., Felleisen, M.: Models of continuations without continuations. In: *Proceedings of the 18th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. pp. 185–196. POPL '91, Association for Computing Machinery. <https://doi.org/10.1145/99583.99611>, <https://dl.acm.org/doi/10.1145/99583.99611>
27. Timany, A., Birkedal, L.: Mechanized relational verification of concurrent programs with continuations. *Proc. ACM Program. Lang.* **3**(ICFP), 105:1–105:28 (2019). <https://doi.org/10.1145/3341709>, <https://doi.org/10.1145/3341709>

28. Timany, A., Guéneau, A., Birkedal, L.: The logical essence of well-bracketed control flow. *Proc. ACM Program. Lang.* **8**(POPL), 575–603 (2024). <https://doi.org/10.1145/3632862>, <https://doi.org/10.1145/3632862>
29. Timany, A., Krebbers, R., Dreyer, D., Birkedal, L.: A logical approach to type soundness. *Journal of the ACM* **71** (2024)
30. van den Berg, B., Schrijvers, T., Poulsen, C.B., Wu, N.: Latent Effects for Reusable Language Components. In: Oh, H. (ed.) *Programming Languages and Systems*. pp. 182–201. *Lecture Notes in Computer Science*, Springer International Publishing, Cham (2021). https://doi.org/10.1007/978-3-030-89051-3_11
31. de Vilhena, P.E., Pottier, F.: A separation logic for effect handlers. *Proc. ACM Program. Lang.* **5**(POPL) (jan 2021). <https://doi.org/10.1145/3434314>, <https://doi.org/10.1145/3434314>
32. Wu, N., Schrijvers, T., Hinze, R.: Effect handlers in scope. In: *Proceedings of the 2014 ACM SIGPLAN Symposium on Haskell*. pp. 1–12. Haskell '14, Association for Computing Machinery, New York, NY, USA (Sep 2014). <https://doi.org/10.1145/2633357.2633358>
33. Xia, L., Zakowski, Y., He, P., Hur, C., Malecha, G., Pierce, B.C., Zdancewic, S.: Interaction trees: representing recursive and impure programs in coq. *Proc. ACM Program. Lang.* **4**(POPL), 51:1–51:32 (2020). <https://doi.org/10.1145/3371119>, <https://doi.org/10.1145/3371119>
34. Zakowski, Y., Beck, C., Yoon, I., Zaichuk, I., Zaliva, V., Zdancewic, S.: Modular, compositional, and executable formal semantics for LLVM IR. *Proceedings of the ACM on Programming Languages* **5**(ICFP), 67:1–67:30 (Aug 2021). <https://doi.org/10.1145/3473572>
35. Zhang, H., Honoré, W., Koh, N., Li, Y., Li, Y., Xia, L.Y., Beringer, L., Mansky, W., Pierce, B., Zdancewic, S.: Verifying an HTTP Key-Value Server with Interaction Trees and VST. In: Cohen, L., Kaliszyk, C. (eds.) *12th International Conference on Interactive Theorem Proving (ITP 2021)*. *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 193, pp. 32:1–32:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2021). <https://doi.org/10.4230/LIPIcs.ITP.2021.32>