Mechanized Relational Verification of Concurrent Programs with Continuations

AMIN TIMANY, imec-Distrinet, KU Leuven, Belgium

LARS BIRKEDAL, Aarhus University, Denmark

Concurrent higher-order imperative programming languages with continuations are very flexible and allow for the implementation of sophisticated programming patterns. For instance, it is well known that continuations can be used to implement cooperative concurrency. Continuations can also simplify web server implementations. This, in particular, helps simplify keeping track of the state of server's clients. However, such advanced programming languages are very challenging to reason about. One of the main challenges in reasoning about programs in the presence of continuations is due to the fact that the non-local flow of control breaks the *bind* rule, one of the important modular reasoning principles of Hoare logic.

In this paper we present the first completely formalized tool for interactive mechanized relational verification of programs written in a concurrent higher-order imperative programming language with continuations (*call/cc* and *throw*). We develop novel logical relations which can be used to give mechanized proofs of relational properties. In particular, we prove correctness of an implementation of cooperative concurrency with continuations. In addition, we show that that a rudimentary web server implemented using the continuation-based pattern is contextually equivalent to one implemented without the continuation-based pattern. We introduce context-local reasoning principles for our calculus which allows us to regain modular reasoning principles for the fragment of the language without non-local control flow. These novel reasoning principles can be used in tandem with our (non-context-local) Hoare logic for reasoning about programs that do feature non-local control flow. Indeed, we use the combination of context-local and non-context-local reasoning to simplify reasoning about the examples.

CCS Concepts: • **Theory of computation** → **Logic and verification; Hoare logic; Separation logic; Program specifications; Program verification; Invariants; Pre- and post-conditions**; • **Software and its engineering** → **Formal software verification; Semantics**.

Additional Key Words and Phrases: Logical relations, Continuations, Concurrency

ACM Reference Format:

Amin Timany and Lars Birkedal. 2019. Mechanized Relational Verification of Concurrent Programs with Continuations. *Proc. ACM Program. Lang.* 3, ICFP, Article 105 (August 2019), 28 pages. <https://doi.org/10.1145/3341709>

1 INTRODUCTION

In a programming language with continuations, a computation can be suspended into a continuation object which can be resumed later. Continuations enable interesting programming patterns. For instance, it is well-known that they can be used to implement cooperative concurrency [Haynes et al. 1984]: switching between threads can be implemented by suspending the running thread, storing the suspension and running another thread. Another notable application of continuations

Authors' addresses: Amin Timany, Department of Computer Science, imec-Distrinet, KU Leuven, Leuven, Belgium, amin.timany@cs.kuleuven.be; Lars Birkedal, Department of Computer Science, Aarhus University, Aarhus, Denmark, birkedal@cs.au.dk.

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

© 2019 Copyright held by the owner/author(s).

2475-1421/2019/8-ART105

<https://doi.org/10.1145/3341709>

is the implementation of continuation-based web servers [Flatt 2017; Krishnamurthi et al. 2007; Queinnec 2004]. Web servers store the state of their communication with each client in order to provide a coherent experience for returning clients. For this purpose continuation-based web servers store the continuation of the server-side program. This helps simplify the web server program because the only thing that the server needs to do in order to serve a returning client appropriately is to resume the continuation corresponding to the last communication with the client.

Both of the aforementioned programming patterns involve concurrency, higher-order and imperative aspects of the programming language in combination with continuations in sophisticated and interesting ways. Such expressive and advanced programming languages and programs written in them are known to be challenging to model and reason about. In this paper we present novel techniques for relational reasoning about such programming languages and programs written in them.

Specifically, we develop a new logical relations model for proving contextual refinement of programs written in $F_{conc, cc}^{\mu, ref}$, a call-by-value programming language featuring concurrency (*conc*), impredicative polymorphism (*F*), recursive types (μ), dynamically allocated higher-order store (*ref*) and first-class continuations (*cc*) with *call/cc* and *throw* primitives. We employ this logical relations model to prove (1) contextual equivalence of two simple web server implementations: a continuation-based one and a state-storing one; and (2) that one-shot continuations (continuations that can only be used once) can be used to simulate ordinary continuations [Friedman and Haynes 1985]. The latter is a well-known result for sequential programs, see, e.g., Dreyer et al. [2012]. Here we show that it also holds in $F_{conc, cc}^{\mu, ref}$, that is, in the presence of concurrency.

In addition, we develop a relational model for showing the correctness of a continuation-based implementation of cooperative concurrency. We consider two programming languages: a language, $F_{cc, coop}^{\mu, ref}$, with built-in cooperative concurrency (*coop*) and a sequential language, $F_{cc}^{\mu, ref}$, featuring continuations. We develop a cross-language logical relation between $F_{cc}^{\mu, ref}$ and $F_{cc, coop}^{\mu, ref}$ and use it to show correctness of a translation of cooperative concurrency into one based on continuations.

We define our logical relations models in a variant of the Iris program logic framework [Jung et al. 2016; Jung et al. 2018; Jung et al. 2015; Krebbers et al. 2017a]. Iris is a framework for state-of-the-art higher-order concurrent separation logics. We use Iris because (1) it allows us to define our logical relations and reason about them at a higher level of abstraction compared to an explicit model construction; (2) we side-step the well-known type-world-circularity problems [Ahmed 2004; Ahmed et al. 2002; Birkedal et al. 2011] involved in defining logical relations for programming languages with higher-order store (since that is already “taken care of” by the model of Iris); and (3) we can leverage the Coq implementation of the Iris base logic [Krebbers et al. 2017a] and the Iris Proof Mode [Krebbers et al. 2017b] when mechanizing our development in Coq. Indeed, accompanying this paper is a tool for mechanized relational verification of concurrent programs with continuations. The mechanization has been done in Coq and all the results in the paper have been formally verified.

1.1 Context-local Reasoning Principles

Some of the most important features of concurrent separation logics for *modular/local* reasoning about concurrent imperative programs, e.g. da Rocha Pinto et al. [2014]; Dinsdale-Young et al. [2013, 2010]; Jung et al. [2016, 2015]; Krebbers et al. [2017a,b]; Ley-Wild and Nanevski [2013]; Nanevski et al. [2014]; O’Hearn [2007]; Sergey et al. [2015]; Svendsen and Birkedal [2014]; Turon et al. [2013a] are *thread-local* reasoning and *context-local* reasoning. Thread-local reasoning means that we can reason about each thread in isolation: when we reason about a particular thread, we need not

explicitly consider interactions from other concurrently executing threads. Similarly, context-local means that when we reason about a particular expression, we need not consider under which evaluation context it is being evaluated. The latter is sometimes codified by the soundness of a proof rule such as the following:

$$\frac{\text{HOARE-BIND (INADMISSIBLE IN PRESENCE OF CONTINUATIONS)} \quad \{P\} e \{ \Psi \} \quad \forall w. \{ \Psi(w) \} K[w] \{ \Phi \}}{\{P\} K[e] \{ \Phi \}}$$

The Hoare-triple $\{P\} e \{ \Psi \}$ intuitively means that, given precondition P , expression e is *safe* and, whenever it reduces to a value v , we are guaranteed that $\Psi(v)$ holds. Intuitively, the above rule expresses that to prove a Hoare triple for an expression e in an evaluation context K , it suffices to prove a property for e in isolation from K , and then show that the desired postcondition Φ can be obtained when substituting a value w satisfying the postcondition Ψ for e into the evaluation context. In a programming language with non-local control operators, e.g. `call/cc` and `throw`, the context under which a program is being evaluated is of utmost importance, and thus the above proof rule is *not* sound in general for languages with non-local control operators. In spite of this, we ought be able to reason about those parts of the program that do not feature any non-local control flow operators in a context-local way. We explain this via an example which we use to demonstrate how non-local control flow operators can break context-local reasoning. We then discuss how our context-local program logic allows us to regain context-local reasoning in a sound way, and, how we can employ our context-local program logic to derive a context-local specification for the this example.

To illustrate how non-local control flow operators can break context-local reasoning, consider the following program:

$$\text{CallIncr} \triangleq \lambda f. \text{let } x = \text{ref}(0) \text{ in let } g = f() \text{ in } x \leftarrow !x + 2; g(); !x$$

(In Section 2 we present the syntax and semantics of $F_{\text{conc}, \text{cc}}^{\mu, \text{ref}}$ formally; here we simply explain informally what is intended by our ML-like syntax above.) The program *CallIncr* takes a function f as argument, allocates a local reference with value 0, calls f , and binds its result (again a function) to g . It then increments the internal reference and subsequently calls g before returning the value stored in the internal reference. Intuitively, in a programming language without non-local control flow this program should always return 2 (if it terminates, of course). Hence, using the bind rule we would be able to derive the following (incorrect) specification for *CallIncr*:

$$\{\text{true}\} h() \{x. \{\text{true}\} x() \{y. y = ()\}\} \text{ implies } \{\text{true}\} \text{CallIncr } h \{x. x = 2\} \quad (\text{CallIncr-incorrect-spec})$$

This specification says that, if h is a function that upon call returns a function x for which we know the Hoare-triple $\{\text{true}\} x() \{y. y = ()\}$ holds, then *CallIncr* applied to h should return 2. This specification *CallIncr-incorrect-spec* does not hold though. For a concrete counter-example, let h to be

$$h = \lambda_. \text{call/cc}(x. \lambda_. \text{throw}(\lambda. ()) \text{ to } x)$$

This function, when called, captures the current continuation x and then returns a function that returns to the point x , this time, with the result $(\lambda. ())$, a function that simply returns the unit value $()$. Hence, when *CallIncr* is called with h , the internal reference of *CallIncr* will be incremented twice, and thus return value be 4 and not 2. Notice that for h above we can prove the spec $\{\text{true}\} h() \{x. \{\text{true}\} x() \{y. y = ()\}\}$ since this spec considers h in isolation, i.e., the captured continuation will be the empty continuation.

Hence, as expected, the program *CallIncr* only behaves context-locally if the argument it is applied to does. This is indeed reflected in the context-local specification of *CallIncr* expressed in terms of *context-local Hoare triples* (Specification *CallIncr-context-local-spec* below). We introduce context-local Hoare triples in Section 4. A context-local Hoare triple $\{P\}^{\text{cl}} e \{x. Q\}$ not only implies safety of e (as ordinary Hoare triples do), but also implies that e behaves context-locally. Hence, they can be used to reason about programs context-locally, and thus the bind rule holds for them. In particular, we can prove the following context-local specification for the program *CallIncr* above:

$$\{\text{true}\}^{\text{cl}} h () \{x. \{\text{true}\}^{\text{cl}} x () \{y. y = ()\}\} \text{ implies } \{\text{true}\}^{\text{cl}} \text{CallIncr } h \{x. x = 2\} \\ (\text{CallIncr-context-local-spec})$$

This specification captures exactly the intuitive idea explained above: if we know that h , and the function it returns, behave context-locally, then so does *CallIncr* h .

Since context-local Hoare-triples cannot be used to reason about all programs, particularly about non-trivial uses of *call/cc* and *throw*, we have to express our logical relations model using ordinary (non-context-local) Hoare-triples, following ideas from earlier work [Krebbers et al. 2017b; Turon et al. 2013a]. In Section 4, we discuss how ordinary (non-context-local) Hoare-triples interact with their context-local counterparts. Using a combination of non-context-local and context-local triples, we can simplify reasoning about contextual equivalence of concurrent programs with continuations using our logical relations model. We achieve this by reasoning about parts of programs that do not use non-local control operators using context-local Hoare-triples.

Contributions. In this paper, we make the following contributions:

- We present a program logic (weakest preconditions and Hoare-triples) for reasoning about programs written in $F_{\text{conc}, \text{cc}}^{\mu, \text{ref}}$, a programming language with impredicative polymorphism, recursive types, higher-order functions, higher-order store, concurrency and first-class continuations.
- We present context-local weakest-preconditions and Hoare-triples which simplify reasoning about programs without non-local control flow.
- We present a novel logical relations model for $F_{\text{conc}, \text{cc}}^{\mu, \text{ref}}$.
- We use our logical relations model and context-local reasoning to prove equivalence of two simple web server implementations: a continuation-based one and a state-storing one.
- We further use our logical relations model to prove correctness of Friedman and Haynes [1985] encoding of continuations by means of one-shot continuations in a concurrent programming language.
- We develop a cross-language logical relations model between $F_{\text{cc}}^{\mu, \text{ref}}$ and $F_{\text{cc}, \text{coop}}^{\mu, \text{ref}}$ for proving program refinement.
- We use our cross-language logical relations model to prove correctness of a continuation-based implementation of cooperative concurrency.
- We have developed a fully formalized tool for mechanized interactive relational verification of concurrent programs with continuations. Our tool is developed on top of Iris, a state-of-the-art program logic framework, and we have used it to mechanize all of our contributions in the Coq proof assistant.

2 THE LANGUAGE: $F_{\text{conc}, \text{cc}}^{\mu, \text{ref}}$

The language that we consider in this paper, $F_{\text{conc}, \text{cc}}^{\mu, \text{ref}}$, is a typed lambda calculus with a standard call-by-value small-step operational semantics. It features impredicative polymorphism (F), recursive types (μ), higher-order mutable references (*ref*), fine-grained concurrency (*conc*) and first-class

continuations (cc). The types of $F_{conc,cc}^{\mu,ref}$ are as follows:

$$\tau ::= \alpha \mid \mathbf{1} \mid \mathbb{B} \mid \mathbb{N} \mid \tau \rightarrow \tau \mid \forall \alpha. \tau \mid \mu \alpha. \tau \mid \tau \times \tau \mid \tau + \tau \mid \text{ref}(\tau) \mid \text{cont}(\tau)$$

Here α ranges over type variables. The types $\mathbf{1}$, $\mathbb{B}$ and $\mathbb{N}$ are the unit type, the type of Boolean values and the type of natural numbers respectively. The type $\text{ref}(\tau)$ is the type of references with contents of type τ and $\text{cont}(\tau)$ is the type of continuations that can be resumed by throwing them a value of type τ . The symbols $\times$, $+$, μ , $\forall$ and $\rightarrow$ are the usual product, sum, recursive type, polymorphic type and function type formers, respectively.

The syntax for expressions and values is:

$$\begin{aligned} e ::= & x \mid () \mid \text{true} \mid \text{false} \mid n \mid e \odot e \mid \text{rec } f(x) = e \mid e \ e \mid \Lambda e \mid e _ \mid \text{fold } e \mid \text{unfold } e \mid (e, e) \mid \pi_i e \\ & \mid \text{inj}_i e \mid \text{match } e \text{ with } \text{inj}_i x \Rightarrow e_i \text{ end} \mid \ell \mid \text{ref}(e) \mid !e \mid e \leftarrow e \mid \text{cas}(e, e, e) \mid \text{fork } \{e\} \\ & \mid \text{cont}(K) \mid \text{call/cc}(x. e) \mid \text{throw } e \text{ to } e \\ v ::= & () \mid \text{true} \mid \text{false} \mid n \mid \text{rec } f(x) = e \mid \Lambda e \mid \text{fold } v \mid (v, v) \mid \text{inj}_i v \mid \ell \mid \text{cont}(K) \end{aligned}$$

We write n for natural numbers and the symbol $\odot$ stands for binary operations on natural numbers (both basic arithmetic operations and basic comparison operations). We consider both recursive functions $\text{rec } f(x) = e$ and polymorphic type abstractions Λe to be values. We write $e _$ for type level application (e is a polymorphic expression). We use **fold** and **unfold** to fold and unfold elements of recursive types. Memory locations ℓ are values of reference types. The expression $!e$ reads the memory location e evaluates to, and $e \leftarrow e'$ is an assignment of the value computed by e' to the memory location computed by e . The expression **fork** $\{e\}$ is for forking off a new thread to compute e and we write **cas**(e, e', e'') for the compare-and-set operation. A continuation, **cont**(K), is essentially a suspended evaluation context (see the operational semantics below).

Evaluation contexts of $F_{conc,cc}^{\mu,ref}$ are as follows:

$$\begin{aligned} K ::= & [] \mid K \odot e \mid v \odot K \mid K e \mid v K \mid K _ \mid \text{fold } K \mid \text{unfold } K \mid \text{if } K \text{ then } e \text{ else } e \mid (K, e) \mid (v, K) \\ & \mid \pi_i K \mid \text{inj}_i K \mid \text{match } K \text{ with } \text{inj}_i x \Rightarrow e_i \text{ end} \mid \text{ref}(K) \mid !K \mid K \leftarrow e \mid v \leftarrow K \\ & \mid \text{cas}(K, e, e) \mid \text{cas}(v, K, e) \mid \text{cas}(v, v, K) \mid \text{throw } K \text{ to } e \mid \text{throw } v \text{ to } K \end{aligned}$$

The evaluation context $[]$ is the empty evaluation context.

2.1 Typing

An excerpt of the typing rules is depicted in Figure 1. The context $\Xi = \alpha_1, \dots, \alpha_n$ is a list of distinct type variables and the context $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n$ assigns types to program variables.

2.2 Operational Semantics

We define the call-by-value small-step operational semantics of $F_{conc,cc}^{\mu,ref}$ in two stages. We first define a head-step relation $\rightarrow_K$. Here, K is the context under which the head step is being performed. Based on this, we define the operational semantics of programs by what we call *the thread-pool step relation* $\rightarrow$. A thread pool reduces by making a head reduction step in one of the threads, by forking off a new thread, or by resuming a captured continuation in one of the threads. These rules are depicted in Figure 2. In this figure, σ is the physical state of the program, i.e., the program heap, which is a finite partial map from memory locations to values. An excerpt of the head-step relation is given in Figure 3. Notice that the head-step for **call/cc** captures the continuation that is the index of the head-step relation.

$\boxed{\Xi \mid \Gamma \vdash e : \tau}$			
$\frac{\text{T-VAR} \quad x : \tau \in \Gamma}{\Xi \mid \Gamma \vdash x : \tau}$	$\text{T-UNIT} \quad \Xi \mid \Gamma \vdash () : 1$	$\text{T-NAT} \quad \Xi \mid \Gamma \vdash n : \mathbb{N}$	$\frac{\text{T-TLAM} \quad \Xi, \alpha \mid \Gamma \vdash e : \tau}{\Xi \mid \Gamma \vdash \Lambda e : \forall \alpha. \tau}$
$\frac{\text{T-REC} \quad \Xi \mid \Gamma, x : \tau, f : \tau \rightarrow \tau' \vdash e : \tau'}{\Xi \mid \Gamma \vdash \text{rec } f(x) = e : \tau \rightarrow \tau'}$	$\frac{\text{T-APP} \quad \Xi \mid \Gamma \vdash e : \tau \rightarrow \tau' \quad \Xi \mid \Gamma \vdash e' : \tau}{\Xi \mid \Gamma \vdash e e' : \tau'}$		$\frac{\text{T-TAPP} \quad \Xi \mid \Gamma \vdash e : \forall \alpha. \tau}{\Xi \mid \Gamma \vdash e _ : \tau[\tau'/\alpha]}$
$\frac{\text{T-FOLD} \quad \Xi \mid \Gamma \vdash e : \tau[\mu\alpha. \tau/\alpha]}{\Xi \mid \Gamma \vdash \text{fold } e : \mu\alpha. \tau}$	$\frac{\text{T-UNFOLD} \quad \Xi \mid \Gamma \vdash e : \mu\alpha. \tau}{\Xi \mid \Gamma \vdash \text{unfold } e : \tau[\mu\alpha. \tau/\alpha]}$	$\frac{\text{T-REF} \quad \Xi \mid \Gamma \vdash e : \tau}{\Xi \mid \Gamma \vdash \text{ref}(e) : \text{ref}(\tau)}$	
$\frac{\text{T-DeREF} \quad \Xi \mid \Gamma \vdash e : \text{ref}(\tau)}{\Xi \mid \Gamma \vdash !e : \tau}$	$\frac{\text{T-FORK} \quad \Xi \mid \Gamma \vdash e : \tau}{\Xi \mid \Gamma \vdash \text{fork } \{e\} : 1}$	$\frac{\text{T-ASSIGN} \quad \Xi \mid \Gamma \vdash e : \text{ref}(\tau) \quad \Xi \mid \Gamma \vdash e' : \tau}{\Xi \mid \Gamma \vdash e \leftarrow e' : 1}$	
$\frac{\text{T-CAS} \quad \Xi \mid \Gamma \vdash e_1 : \text{ref}(\tau) \quad \Xi \mid \Gamma \vdash e_2 : \tau \quad \Xi \mid \Gamma \vdash e_3 : \tau}{\Xi \mid \Gamma \vdash \text{cas}(e_1, e_2, e_3) : \mathbb{B}}$			$\frac{\text{T-CALL/CC} \quad \Xi \mid \Gamma, x : \text{cont}(\tau) \vdash e : \tau}{\Xi \mid \Gamma \vdash \text{call/cc}(x, e) : \tau}$
$\frac{\text{T-THROW} \quad \Xi \mid \Gamma \vdash e : \tau \quad \Xi \mid \Gamma \vdash e' : \text{cont}(\tau)}{\Xi \mid \Gamma \vdash \text{throw } e \text{ to } e' : \tau'}$			

Fig. 1. An excerpt of the typing rules

$\boxed{(\vec{e}; \sigma) \rightarrow (\vec{e}'; \sigma')}$	
$\frac{(e, \sigma) \rightarrow_K (e', \sigma')}{(\vec{e}_1, K[e], \vec{e}_2; \sigma) \rightarrow (\vec{e}_1, K[e'], \vec{e}_2; \sigma')}$	$(\vec{e}_1, K[\text{fork } \{e\}], \vec{e}_2; \sigma) \rightarrow (\vec{e}_1, K[], \vec{e}_2, e; \sigma)$
$(\vec{e}_1, K[\text{throw } v \text{ to cont}(K')], \vec{e}_2; \sigma) \rightarrow (\vec{e}_1, K'[v], \vec{e}_2; \sigma)$	

Fig. 2. The thread-pool step relation.

Contextual refinement/equivalence. A program e contextually refines a program e' if both programs have type τ and no *well-typed* context (a closed top-level program with a hole) can distinguish a situation where e' is replaced by e . We write $C : (\Xi \mid \Gamma; \tau) \rightsquigarrow (\Xi' \mid \Gamma'; \tau')$ for a context (a term with a hole) such that $\Xi' \mid \Gamma' \vdash C[e] : \tau'$ holds whenever $\Xi \mid \Gamma \vdash e : \tau$ does. We define contextual refinement of e' by e , written $\Xi \mid \Gamma \vdash e \leq_{\text{ctx}} e' : \tau$, as follows:

$$\Xi \mid \Gamma \vdash e \leq_{\text{ctx}} e' : \tau \triangleq \Xi \mid \Gamma \vdash e : \tau \text{ and } \Xi \mid \Gamma \vdash e' : \tau \text{ and}$$

$$\text{for any } C \text{ such that } C : (\Xi \mid \Gamma; \tau) \rightsquigarrow (\cdot \mid \cdot; 1) \text{ holds } C[e] \Downarrow \text{ implies } C[e'] \Downarrow$$

where $e \Downarrow$ stands for termination of e when run under the empty heap defined as follows:

$$e \Downarrow \triangleq \exists v, \vec{e}, \sigma. (e; \emptyset) \rightarrow^* (v, \vec{e}; \sigma)$$

$$\boxed{(e, \sigma) \rightarrow (e', \sigma')}$$

$$\begin{aligned}
& ((\text{rec } f(x) = e) v, \sigma) \rightarrow_K (e[v, (\text{rec } f(x) = e)/x, f], \sigma) & (\text{unfold } (\text{fold } v), \sigma) \rightarrow_K (v, \sigma) \\
& ((\Lambda e) _ , \sigma) \rightarrow_K (e, \sigma) & (\text{if true then } e_2 \text{ else } e_3, \sigma) \rightarrow_K (e_2, \sigma) \\
& (\text{if false then } e_2 \text{ else } e_3, \sigma) \rightarrow_K (e_3, \sigma) & (\pi_1(v_1, v_2), \sigma) \rightarrow_K (v_1, \sigma) & (\pi_2(v_1, v_2), \sigma) \rightarrow_K (v_2, \sigma) \\
& (\text{match inj}_1 v \text{ with inj}_1 x \Rightarrow e_1 \mid \text{inj}_2 x \Rightarrow e_2 \text{ end}, \sigma) \rightarrow_K (e_1[v/x], \sigma) \\
& (\text{match inj}_2 v \text{ with inj}_1 x \Rightarrow e_1 \mid \text{inj}_2 x \Rightarrow e_2 \text{ end}, \sigma) \rightarrow_K (e_2[v/x], \sigma) \\
& \frac{\ell \notin \text{dom}(\sigma)}{(\text{ref}(v), \sigma) \rightarrow_K (\ell, \sigma \uplus \{\ell \mapsto v\})} & \frac{\sigma = \sigma' \uplus \{\ell \mapsto v'\}}{(\ell \leftarrow v, \sigma) \rightarrow_K ((_, \sigma' \uplus \{\ell \mapsto v'\}))} & \frac{v = \sigma(\ell)}{(!\ell, \sigma) \rightarrow_K (v, \sigma)} \\
& \frac{\sigma = \sigma' \uplus \{\ell \mapsto v\}}{(\text{cas}(\ell, v, v'), \sigma) \rightarrow_K (\text{true}, \sigma' \uplus \{\ell \mapsto v'\})} & \frac{\sigma = \sigma' \uplus \{\ell \mapsto v''\} \quad v \neq v''}{(\text{cas}(\ell, v, v'), \sigma) \rightarrow_K (\text{false}, \sigma)} \\
& (\text{call/cc } (x. e), \sigma) \rightarrow_K (e[\text{cont}(K)/x], \sigma)
\end{aligned}$$

Fig. 3. An excerpt of the head-reduction rules

The intuitive explanation above for contextual refinement is the reason why in a contextual refinement $e \leq_{\text{ctx}} e'$ or in a logical relatedness relation $e \leq_{\text{log}} e'$, usually, the program on the left hand side, e , is referred to as the implementation side and the program on the right hand side, e' , is referred to as the specification side. Two programs are contextually equivalent, if each contextually refines the other:

$$\Xi \mid \Gamma \vdash e \approx_{\text{ctx}} e' : \tau \triangleq \Xi \mid \Gamma \vdash e \leq_{\text{ctx}} e' : \tau \wedge \Xi \mid \Gamma \vdash e' \leq_{\text{ctx}} e : \tau$$

3 LOGICAL RELATIONS

It is challenging to construct logical relations for languages with higher-order store because of the so-called type-world circularity [Ahmed 2004; Ahmed et al. 2002; Birkedal et al. 2011]. The logic of Iris is rich enough to allow for a direct inductive specification of the logical relations for programming languages with advanced features such as higher-order references, recursive types, and concurrency [Krebbers et al. 2017b; Krogh-Jespersen et al. 2017; Timany et al. 2018].

3.1 An Iris Primer

Iris [Jung et al. 2016; JUNG et al. 2018; Jung et al. 2015; Krebbers et al. 2017a] is a state-of-the-art higher-order concurrent separation logic designed for verification of programs. In Iris one can quantify over the Iris types κ :

$$\kappa ::= 1 \mid \kappa \times \kappa \mid \kappa \rightarrow \kappa \mid \text{Ectx} \mid \text{Var} \mid \text{Expr} \mid \text{Val} \mid \mathbb{N} \mid \mathbb{B} \mid \kappa \xrightarrow{\text{fin}} \kappa \mid \text{finset}(\kappa) \mid \text{Monoid} \mid \text{Names} \mid \text{iProp} \mid \dots$$

Here *Ectx*, *Var*, *Expr* and *Val* are Iris types for evaluation contexts, variables, expressions and values of $F_{\text{conc}, \text{cc}}^{\mu, \text{ref}}$. Natural numbers, $\mathbb{N}$, and Booleans $\mathbb{B}$ are also included among the base types of Iris. Iris also features partial maps with finite support, $\kappa \xrightarrow{\text{fin}} \kappa$, and finite sets, $\text{finset}(\kappa)$. Resources in Iris are represented using partial commutative monoids, *Monoid*, and instances of resources are

named using so-called ghost-names, *Names*. Finally, and most importantly, there is a type of Iris propositions *iProp*. The grammar for Iris propositions is as follows:

$$\begin{aligned} P ::= & \top \mid \perp \mid P * P \mid P \multimap P \mid P \wedge P \mid P \Rightarrow P \mid P \vee P \mid \forall x : \kappa. \Phi(x) \mid \exists x : \kappa. \Phi(x) \\ & \mid \triangleright P \mid \mu r. P \mid \Box P \mid \text{wp } e \{x. P\} \mid \{P\} e \{x. Q\} \mid \models P \mid \overline{P}^N \mid \dots \end{aligned}$$

Here, $\top$, $\perp$, $\wedge$, $\vee$, $\Rightarrow$, $\forall$, $\exists$ are the standard higher-order logic connectives. The predicates Φ are Iris predicates, *i.e.*, terms of type $\kappa \rightarrow iProp$.

The connective $*$ is the separating conjunction. Intuitively, $P * Q$ holds if resources can be split into two *disjoint* pieces such that one satisfies P and the other Q . The magic wand connective $P \multimap Q$ is satisfied by resources such that when these resources are combined with some resource satisfying P the resulting resources would satisfy Q .

The $\triangleright$ modality, pronounced “later” is a modality that intuitively corresponds to some abstract form of step-indexing [Appel and McAllester 2001; Appel et al. 2007; Dreyer et al. 2011]. Intuitively, $\triangleright P$ holds if P holds one step into the future. Iris has support for taking fixed points of *guarded* propositions, $\mu r. P$. This fixed point can only be defined if all occurrences of r in P are guarded, *i.e.*, appear under a $\triangleright$ modality. We use guarded fixed points for defining the interpretation of recursive types in $F_{\text{conc}, \text{cc}}^{\mu, \text{ref}}$. For any proposition P we have $P \vdash \triangleright P$.

When the modality $\Box$ is applied to a proposition P , the non-duplicable resources in P are forgotten, and thus $\Box P$ is “persistent.” In general, we say that a proposition P is *persistent* if $P \dashv\vdash \Box P$ (where $\dashv\vdash$ is the logical equivalence of Iris propositions). A key property of persistent propositions is that they are duplicable: $P \dashv\vdash P * P$. The type system of $F_{\text{conc}, \text{cc}}^{\mu, \text{ref}}$ is not a sub-structural type system and variables (in the typing environment) may be used multiple times. Therefore when we interpret types as logical relations in Iris, those relations should be duplicable. We use the persistence modality $\Box$ to ensure this.

The Iris program logic for $F_{\text{conc}, \text{cc}}^{\mu, \text{ref}}$. Iris facilitates specification and verification of programs by means of weakest-preconditions $\text{wp } e \{x. P\}$, which intuitively hold whenever e is *safe* and, moreover, whenever e terminates with a resulting value v , then $P[v/x]$ holds. When x does not appear in P we write $\text{wp } e \{x. P\}$ as $\text{wp } e \{P\}$. Also, we sometimes write $\text{wp } e \{\Phi\}$ for $\text{wp } e \{x. \Phi(x)\}$.

In Iris, Hoare triples are defined in terms of weakest preconditions:

$$\{P\} e \{x. Q\} \triangleq \Box (P \multimap \text{wp } e \{x. Q\})$$

Note that the $\Box$ modality ensures that the Hoare triples are persistent and hence duplicable (in separation logic jargon, Hoare triples should just express “knowledge” and not claim ownership of any resources).

A key feature of Iris (as for other concurrency logics) is that specification and verification is done *thread-locally*: the weakest precondition only describes properties of execution of a single thread. Concurrent interactions are abstracted and reasoned about in terms of resources (rather than by explicit reasoning about interleavings). For programming languages that do not include continuations or other forms of non-local control flow, the weakest precondition is not only thread-local, but also what we may call *context-local*. Context-local means that to reason about an expression in an evaluation context, it suffices to reason about the expression in isolation, and then separately about what the context does to the resulting value. This form of context-locality is formally expressed by the soundness of the following bind rule

$$\frac{\text{INADMISSIBLE-BIND} \quad \text{wp } e \{x. \text{wp } K[x] \{\Phi\}\}}{\text{wp } K[e] \{\Phi\}}$$

This rule is not sound when expressions include `call/cc` since `call/cc` captures the evaluation context and hence its behaviour depends on it. We discussed inadmissibility of this rule in [Introduction](#). Thus, for reasoning about $F_{conc,cc}^{\mu,ref}$ we cannot use the “standard” Iris rules [\[Jung et al. 2016; JUNG et al. 2018; Jung et al. 2015; Krebbers et al. 2017a,b\]](#) for weakest preconditions. Instead, we use new rules such as the following:

$$\begin{array}{c}
\text{FST-WP} \\
\frac{\triangleright \text{wp } K[v] \{ \Phi \}}{\text{wp } K[\pi_1(v, w)] \{ \Phi \}}
\end{array}
\quad
\begin{array}{c}
\text{IF-TRUE-WP} \\
\frac{\triangleright \text{wp } K[e] \{ \Phi \}}{\text{wp } K[\text{if true then } e \text{ else } e'] \{ \Phi \}}
\end{array}
\quad
\begin{array}{c}
\text{REC-WP} \\
\frac{\triangleright \text{wp } K[e[\text{rec } f(x) = e, v/f, x]] \{ \Phi \}}{\text{wp } K[(\text{rec } f(x) = e) v] \{ \Phi \}}
\end{array}$$

$$\begin{array}{c}
\text{CALLCC-WP} \\
\frac{\triangleright \text{wp } K[e[\text{cont}(K)/x]] \{ \Phi \}}{\text{wp } K[\text{call/cc}(x. e)] \{ \Phi \}}
\end{array}
\quad
\begin{array}{c}
\text{THROW-WP} \\
\frac{\triangleright \text{wp } K'[v] \{ \Phi \}}{\text{wp } K[\text{throw } v \text{ to cont}(K')] \{ \Phi \}}
\end{array}$$

The difference from the standard rules is that our new rules include an explicit context K . Earlier, such rules could be derived using the bind rule, but that is not sound in general in our settings. Note that the context is used in the rules `CALLCC-WP` and `THROW-WP` for `call/cc` and `throw`. These two rules directly reflect the operational semantics of `call/cc` and `throw`.

To demonstrate how the rules of our program logic are used to reason about a program we discuss the following program:

`call/cc (x. (throw 5 to x) + 4) + 2`

This program first captures the current continuation, $[] + 2$, and then continues with evaluating the `throw` operation. At that point, the `throw` operation causes the program to forget about the rest of computation, *i.e.*, adding 4 and 2, and jumps to the captured continuation, *i.e.*, $[] + 2$, with value 5. Hence, the overall result is 7. We prove this fact, *i.e.*, $\text{wp } \text{call/cc}(x. (\text{throw } 5 \text{ to } x) + 4) + 2 \{x. x = 7\}$, in our program logic as follows:

$$\frac{\frac{\frac{\text{wp } 5 + 2 \{x. x = 7\}}{\text{wp } ((\text{throw } 5 \text{ to cont}([] + 2)) + 4) + 2 \{x. x = 7\}} \text{THROW-WP with } K = ([] + 4) + 2, \text{ and, } K' = [] + 2}{\text{wp } \text{call/cc}(x. (\text{throw } 5 \text{ to } x) + 4) + 2 \{x. x = 7\}} \text{CALLCC-WP with } K = [] + 2$$

Notice the importance of the context K in applying the rule `CALLCC-WP` in this proof.

In summa, for $F_{conc,cc}^{\mu,ref}$ we use new non-context-local rules for reasoning about weakest preconditions, and the non-context-local rules allow us to reason about `call/cc` and `throw`.

Because of the explicit context K , the non-context-local rules for weakest preconditions are somewhat more elaborate to use than the corresponding context-local rules. However, that is the price we have to pay to be able to reason in general about non-local control flow. In [Section 4](#) we will see how we can still recover a form of context-local weakest precondition for reasoning about those parts of the program that do not use non-local control flow.

In the rules above for weakest preconditions, the antecedent is only required to hold a step of computation later ($\triangleright$) — that is because these rules correspond to expressions performing a reduction step.

The update modality and invariants. The update modality $\models$ accounts for updating (allocation, deallocation and mutation) of resources.¹ Intuitively, $\models P$ is satisfied by resources that can be updated to new resources for which P holds. For any proposition P , we have that $P \vdash \models P$. If P holds, then resources can be updated (trivially) so as to have that P holds. The update modality is

¹This modality is called the fancy update modality in [Krebbers et al. \[2017a\]](#).

idempotent, $\models \models P \dashv \vdash \models P$. We write $P \equiv Q$ as a shorthand for $P \multimap Q$. Crucially, resources can be updated throughout a proof of weakest preconditions:

$$\models \text{wp } e \{ \Phi \} \dashv \vdash \text{wp } e \{ \Phi \} \dashv \vdash \text{wp } e \{ x. \models \Phi(x) \}$$

Iris features invariants $\boxed{P}^N$ for enforcing concurrent protocols. Each invariant $\boxed{P}^N$ has a name, N , associated to it. Names are used to keep track of which invariants are open.² Intuitively, $\boxed{P}^N$ states that P always holds. The following rules govern invariants.

$$\begin{array}{c} \text{INV-ALLOC} \\ \frac{\triangleright P}{\models \boxed{P}^N} \end{array} \qquad \begin{array}{c} \text{INV-OPEN-WP} \\ \frac{\boxed{R}^N \quad (\triangleright R) \multimap \text{wp } e \{ y. (\triangleright R) * \text{wp } K[y] \{ x. Q \} \}}{\text{wp } K[e] \{ x. Q \}} \quad e \text{ is atomic} \end{array}$$

These rules say that invariants can always be allocated by giving up the resources being protected by the invariant and they can be kept opened only during the execution of *physically atomic* operations. Iris invariants are impredicative, *i.e.*, they can state P holds invariantly for any proposition P , including invariants. This is why the later operator is used as a guard to avoid self-referential paradoxes [Krebbers et al. 2017a]. Invariants essentially express the knowledge that some proposition holds invariantly. Hence, invariants are always persistent, *i.e.*, $\boxed{P}^N \dashv \vdash \Box \boxed{P}^N$.

3.2 Resources Used in Defining Logical Relations

We need some resources in order to define our logical relations in Iris. We need resources for representing memory locations of the implementation side, the memory locations of the specification side and the expression being evaluated on the specification side. These resources are written as follows:

- $\ell \mapsto_i v$: memory location ℓ contains value v on the implementation side.
- $\ell \mapsto_s v$: memory location ℓ contains value v on the specification side.
- $j \models e$: the thread j on the specification side is about to execute e .

These resources are defined using more primitive resources in Iris, but we omit such details here. What is important is that we can use these resources to reason about programs. In particular, we can derive the following rules (and similarly for other basic expressions) for weakest preconditions and for execution on the specification side.

$$\begin{array}{c} \frac{\forall \ell. \ell \mapsto_i v \multimap \text{wp } K[\ell] \{ \Phi \}}{\text{wp } K[\text{ref}(v)] \{ \Phi \}} \qquad \frac{\ell \mapsto_i v \multimap \text{wp } K[v] \{ \Phi \} \quad \triangleright \ell \mapsto_i v}{\text{wp } K[! \ell] \{ \Phi \}} \\ \\ \frac{\ell \mapsto_i w \multimap \text{wp } K[()] \{ \Phi \} \quad \triangleright \ell \mapsto_i v}{\text{wp } K[\ell \leftarrow w] \{ \Phi \}} \qquad \frac{j \models K[\text{ref}(v)]}{\models \exists \ell. \ell \mapsto_s v * j \models K[\ell]} \qquad \frac{\ell \mapsto_s v \quad j \models K[! \ell]}{\models \ell \mapsto_s v * j \models K[v]} \\ \\ \frac{\ell \mapsto_s v \quad j \models K[\ell \leftarrow w]}{\models \ell \mapsto_s w * j \models K[()] } \end{array}$$

These resources are all exclusive in the sense that:

$$\ell \mapsto_i v * \ell \mapsto_i v' \vdash \perp \qquad \ell \mapsto_s v * \ell \mapsto_s v' \vdash \perp \qquad j \models e * j \models e' \vdash \perp$$

²Officially in Iris, the update modality is in fact annotated with so-called masks (sets of invariant names), which are used to ensure that invariants are not re-opened. For simplicity, we do not include masks in this paper.

3.3 Logical Relations in Iris

Figure 4 presents our binary logical relation for $F_{conc,cc}^{\mu,ref}$. We define the logical relation in several stages. The first thing we define is the relation of observational refinement. Intuitively, an expression e observationally refines an expression e' if, whenever e reduces to a value so does e' . We define this in Iris using magic wand and weakest precondition. The whole formula reads as follows: assuming that there is some thread j on the specification side that is about to execute e' (represented in Iris by $j \Rightarrow e'$) then, after execution of e , we know that thread j on the specification side has also been executed to some value w .

We then use the notion of observational refinement defined above to define the value relation, the expression relation and the evaluation context relation for each type. In contrast to earlier definitions of logical relations in Iris [Krebbers et al. 2017b; Krogh-Jespersen et al. 2017; Timany et al. 2018], our logical relation is an example of so-called biorthogonal logical relations [Pitts 2005], also known as top-top closed logical relations. That is, we define two expressions to be related if plugging them into related evaluation contexts results in observationally related expressions. Two evaluation contexts are defined to be related if plugging related values into them results in observationally related expressions.

The value relation interpretation $\llbracket \Xi \vdash \tau \rrbracket_{\Delta}$ of a type τ in context Ξ is defined by induction on τ . Here Δ is an environment mapping type variables in Ξ to Iris relations. For all the non-continuation types, the definition is exactly as in for the language without `call/cc` [Krebbers et al. 2017b].

Pairs of values of base types, $\mathbf{1}$, $\mathbb{B}$, and $\mathbb{N}$, are related if they are equal values of their corresponding type. A pair of values are related at the sum type if they are both formed by applying the same injection to values that are in turn related at the appropriate type. A pair of values are related at the product type, $\tau \times \tau'$, if they are each a pair of values such that their first components are related at τ and their second components are related at τ' . The value interpretation of recursive types are defined using Iris's guarded fixpoint operator $\mu r.P$. It intuitively states that two values of are related at a recursive types if they are both folded values with their underlying values again related at recursive types. The value relation for functions types $\tau \rightarrow \tau'$ expresses that two values of the function type are related if whenever applied to related values of the domain type, τ , the resulting *expressions* are related at the codomain type, τ' . The use of the *persistently modality*, $\Box$, is to make sure that value interpretations of polymorphic programs are persistent. The value relation for polymorphic values requires two related values to produce two related expressions when instantiated, regardless of what predicate f we take as the interpretation of the type variable; provided that f is a persistent predicate. The use of the persistently modality here, as well as the side-condition on f being persistent, are required to ensure that the value interpretations are persistent. Two values of a reference type are related if they are both locations that always (expressed using Iris invariants) store related values. Finally, the relational interpretation of $\text{cont}(\tau)$ expresses that two continuations are related whenever their corresponding evaluation contexts are related at the evaluation context relation for the type in question.

The evaluation context relation $\mathcal{K} \llbracket \Xi \vdash \tau \rrbracket_{\Delta}$ relates evaluation contexts K and K' if plugging related values of type τ in them results in observationally related expressions.

The expression relation is the standard biorthogonal expression relation. It states that $\mathcal{E} \llbracket \Xi \vdash \tau \rrbracket_{\Delta}(e, e')$ holds whenever, for any two related evaluation contexts K and K' , the expressions $K[e]$ and $K'[e']$ are observationally related.

The notion of logical relatedness states, as usual for call-by-value languages, that two expressions e and e' are logically related if substituting related values for their free variables results in related expressions.

Observational refinement: $O : Expr \times Expr \rightarrow iProp$

$$O(e, e') \triangleq \forall j. j \models e' \ast \text{wp } e \{ \exists w. j \models w \}$$

Value interpretation of types: $\llbracket \Xi \vdash \tau \rrbracket_{\Delta} : Val \times Val \rightarrow iProp$

$$\llbracket \Xi \vdash \alpha \rrbracket_{\Delta} \triangleq \Delta(\alpha)$$

$$\llbracket \Xi \vdash \mathbf{1} \rrbracket_{\Delta}(v, v') \triangleq v = v' = ()$$

$$\llbracket \Xi \vdash \mathbb{B} \rrbracket_{\Delta}(v, v') \triangleq v = v' = \text{true} \vee v = v' = \text{false}$$

$$\llbracket \Xi \vdash \mathbb{N} \rrbracket_{\Delta}(v, v') \triangleq \exists n. v = v' = n$$

$$\llbracket \Xi \vdash \tau_1 + \tau_2 \rrbracket_{\Delta}(v, v') \triangleq \bigvee_{i \in \{1, 2\}} (\exists w, w'. v = \text{inj}_i w \wedge v' = \text{inj}_i w' \wedge \llbracket \Xi \vdash \tau_i \rrbracket_{\Delta}(w, w'))$$

$$\llbracket \Xi \vdash \tau \times \tau' \rrbracket_{\Delta}(v, v') \triangleq \exists w_1, w_2, w'_1, w'_2. v = (w_1, w_2) \wedge v' = (w'_1, w'_2) \wedge \llbracket \Xi \vdash \tau \rrbracket_{\Delta}(w_1, w'_1) \ast \llbracket \Xi \vdash \tau' \rrbracket_{\Delta}(w_2, w'_2)$$

$$\llbracket \Xi \vdash \mu\alpha. \tau \rrbracket_{\Delta} \triangleq \mu f : Val \times Val \rightarrow iProp.$$

$$\lambda(v, v'). \exists w, w'. v = \text{fold } w \wedge v' = \text{fold } w' \wedge$$

$$\triangleright \llbracket \alpha, \Xi \vdash \tau \rrbracket_{\Delta, \alpha \mapsto f}(w, w')$$

$$\llbracket \Xi \vdash \tau \rightarrow \tau' \rrbracket_{\Delta}(v, v') \triangleq \Box (\forall w, w'. \llbracket \Xi \vdash \tau \rrbracket_{\Delta}(w, w') \Rightarrow \mathcal{E} \llbracket \Xi \vdash \tau' \rrbracket_{\Delta}(w, v' w'))$$

$$\llbracket \Xi \vdash \forall \alpha. \tau \rrbracket_{\Delta}(v, v') \triangleq \forall f : Val \times Val \rightarrow iProp.$$

$$\text{persistent}(f) \Rightarrow \Box (\mathcal{E} \llbracket \alpha, \Xi \vdash \tau \rrbracket_{\Delta, \alpha \mapsto f}(v _, v' _))$$

$$\llbracket \Xi \vdash \text{ref}(\tau) \rrbracket_{\Delta}(v, v') \triangleq \exists \ell, \ell'. v = \ell \wedge v' = \ell' \wedge$$

$$\boxed{\exists w, w'. \ell \mapsto_i w \ast \ell' \mapsto_s w' \ast \llbracket \Xi \vdash \tau \rrbracket_{\Delta}(w, w')}^{\mathcal{N}. \ell. \ell'}$$

$$\llbracket \Xi \vdash \text{cont}(\tau) \rrbracket_{\Delta}(v, v') \triangleq \exists K, K'. v = \text{cont}(K) \wedge v' = \text{cont}(K') \wedge \mathcal{K} \llbracket \Xi \vdash \tau \rrbracket_{\Delta}(K, K')$$

Evaluation context interpretation of types: $\mathcal{K} \llbracket \Xi \vdash \tau \rrbracket_{\Delta} : Ectx \times Ectx \rightarrow iProp$

$$\mathcal{K} \llbracket \Xi \vdash \tau \rrbracket_{\Delta}(K, K') \triangleq \forall v, v'. \llbracket \Xi \vdash \tau \rrbracket_{\Delta}(v, v') \Rightarrow O(K[v], K'[v'])$$

Expression interpretation of types: $\mathcal{E} \llbracket \Xi \vdash \tau \rrbracket_{\Delta} : Expr \times Expr \rightarrow iProp$

$$\mathcal{E} \llbracket \Xi \vdash \tau \rrbracket_{\Delta}(e, e') \triangleq \forall K, K'. \mathcal{K} \llbracket \Xi \vdash \tau \rrbracket_{\Delta}(K, K') \Rightarrow O(K[e], K'[e'])$$

Logical relatedness: $\Xi \mid \Gamma \models e \leq_{\log} e' : \tau : iProp$ for $\Gamma = x_1 : \tau_1, \dots, x_n : \tau_n$

$$\Xi \mid \Gamma \models e \leq_{\log} e' : \tau \triangleq \forall \Delta, \vec{v}, \vec{v}'. \left(\bigstar_{x_i : \tau_i} \llbracket \Xi \vdash \tau_i \rrbracket_{\Delta}(v_i, v'_i) \right) \Rightarrow \mathcal{E} \llbracket \Xi \vdash \tau \rrbracket_{\Delta} \left(e[\vec{v}/\vec{x}], e'[\vec{v}'/\vec{x}] \right)$$

The map $\Delta : Var \rightarrow (Val \times Val) \rightarrow iProp$ maps free type variables to their value interpretations. The λ in the definition of $\llbracket \Xi \vdash \mu\alpha. \tau \rrbracket_{\Delta}$ is the meta-level λ used for forming an Iris predicate.

Fig. 4. An excerpt of the logical relations for $F_{conc, cc}^{\mu, ref}$

We can now state and prove the fundamental theorem of logical relations for $F_{conc,cc}^{\mu,ref}$. The theorem expresses that any well-typed expression is logically related to itself.

THEOREM 3.1 (FUNDAMENTAL THEOREM OF LOGICAL RELATIONS).

$$\Xi \mid \Gamma \vdash e : \tau \Rightarrow \Xi \mid \Gamma \models e \leq_{log} e : \tau$$

This theorem is proven by induction on the typing derivation using the basic rules for weakest-preconditions and executions on the specification side.

The above theorem, together with some basic properties of observational refinement, implies the soundness of our logical relations, *i.e.*, that logical relatedness implies contextual refinement:

THEOREM 3.2 (SOUNDNESS OF LOGICAL RELATIONS).

$$\Xi \mid \Gamma \models e \leq_{log} e' : \tau \Rightarrow \Xi \mid \Gamma \models e \leq_{ctx} e' : \tau$$

Our logical relation is expressed in terms of weakest preconditions and the proofs of the above theorems use the earlier presented proof rules for weakest preconditions. Before turning to applications, we pause to present *context-local weakest preconditions*, which we can use to simplify reasoning about program fragments, which do not use non-local control flow.

4 CONTEXT-LOCAL WEAKEST PRECONDITIONS

To make it simpler to reason about expressions that do not use non-local control flow, we define a new notion of *context-local weakest precondition* (CLWP). The definition is given in terms of the earlier weakest precondition, which, as we will explain below, means that we will be able to mix and match reasoning steps using (non-context local) weakest preconditions and context-local weakest preconditions.

Definition 4.1. The *context-local weakest precondition* of e with respect to Φ is defined as follows:

$$\text{clwp } e \{ \Phi \} \triangleq \forall K, \Psi. (\forall v. \Phi(v) \multimap \text{wp } K[v] \{ \Psi \}) \multimap \text{wp } K[e] \{ \Psi \}$$

Based on this, we also define context-local Hoare triples:

$$\{P\}^{cl} e \{x. Q\} \triangleq \Box (P \multimap \text{clwp } e \{x. Q\})$$

Note how the above definition essentially says that $\text{clwp } e \{ \Phi \}$ holds if the bind rule holds for e , which intuitively means that e behaves context-locally. Another way to look at the definition of context-local weakest preconditions above is the following. In order to prove $\text{clwp } e \{ \Phi \}$, we have to show that given any evaluation context K and predicate Ψ we have to show that running e under K can satisfy postcondition Ψ . Since K and Ψ are universally quantified and the *only thing* we know about them is that $\forall v. \Phi(v) \multimap \text{wp } K[v] \{ \Psi \}$, the *only way* for the program $K[e]$ to guarantee Ψ as its postcondition is that it would, in its execution, eventually reach a point $K[v]$ for a value v such that $\Phi(v)$. This means that intuitively e does not tamper with the evaluation context it is running under. In other words, for all intents and purposes e behaves context-locally. Therefore, the bind rule below is sound for context-local weakest preconditions.

$$\frac{\text{BIND} \quad \text{clwp } e \{x. \text{clwp } K[x] \{ \Phi \} \}}{\text{clwp } K[e] \{ \Phi \}}$$

Moreover, the “standard” rules for the basic language constructs (excluding `call/cc` and `throw`, of course) can also be derived for context-local weakest preconditions, as shown in Figure 5. We

$$\begin{array}{c}
\text{FST-CLWP} \\
\frac{\triangleright \text{clwp } v \{ \Phi \}}{\text{clwp } \pi_1(v, w) \{ \Phi \}} \\
\\
\text{IF-TRUE-CLWP} \\
\frac{\triangleright \text{clwp } e \{ \Phi \}}{\text{clwp } \text{if true then } e \text{ else } e' \{ \Phi \}} \\
\\
\text{REC-CLWP} \\
\frac{\triangleright \text{clwp } e[v/x][\text{rec } f(x) = e/f] \{ \Phi \}}{\text{clwp } (\text{rec } f(x) = e) v \{ \Phi \}} \\
\\
\text{ALLOC-CLWP} \\
\text{clwp } \text{ref}(v) \{ x. \exists \ell. x = \ell * \ell \mapsto_i v \} \\
\\
\text{LOAD-CLWP} \\
\frac{\triangleright \ell \mapsto_i v}{\text{clwp } !\ell \{ x. x = v * \ell \mapsto_i v \}}
\end{array}$$

Fig. 5. An excerpt of inference rules for CLWP's.

can also use invariants during atomic steps of computation while proving context-local weakest preconditions.

$$\begin{array}{c}
\text{INV-OPEN-CLWP} \\
\frac{\boxed{R}^N \quad (\triangleright R) * \text{clwp } e \{ x. (\triangleright R) * Q \} \quad e \text{ is atomic}}{\text{clwp } e \{ x. Q \}}
\end{array}$$

Now we have both (non-context-local) weakest preconditions and context-local weakest preconditions. What is the upshot of this? The key point is that when we prove correctness / relatedness of programs, we can use the simpler context-local weakest preconditions for reasoning about those parts of the program which are context local (do not use `call/cc` or `throw`) and only use the (non-context-local) weakest preconditions for reasoning about those parts that may involve non-local control flow. This fact is expressed formally by the rule **CLWP-WP** below, which establishes a connection between weakest-preconditions and context-local weakest preconditions.

$$\begin{array}{c}
\text{CLWP-WP} \\
\frac{\text{clwp } e \{ \Psi \} \quad \forall v. \Psi(v) * \text{wp } K[v] \{ \Phi \}}{\text{wp } K[e] \{ \Phi \}}
\end{array}$$

This rule basically says that if we know that e *context-locally* guarantees postcondition Ψ then we can prove $\text{wp } K[e] \{ \Phi \}$ by assuming that, *locally* under the context K , it will only evaluate to values that satisfy Ψ . Moreover, it guarantees that the evaluation of e does not tamper with the evaluation context that we are considering it under.

It is easy to show that $\text{clwp } e \{ \Phi \}$ implies $\text{wp } e \{ \Phi \}$; simply take K to be the empty evaluation context, $[\]$, and Ψ to be Φ . Hence, similarly to a (non-context-local) weakest precondition, the context-local weakest precondition $\text{clwp } e \{ \Phi \}$ also implies the safety of e , and, that whenever e terminates with a value v , $\Phi(v)$ holds.

The proposition $\text{clwp } e \{ \Phi \}$ only says that e behaves context-locally. Therefore, in principle, we should be able to give context-local specifications to programs that do use `call/cc` and `throw`, but still behave context-locally. For instance, the following example from Section 3.1:

`call/cc` ($x. (\text{throw } 5 \text{ to } x) + 4$) + 2

This program does indeed use `call/cc` and `throw`, but it does not tamper with the evaluation context that it runs under. That is to say that the effects of `call/cc` and `throw` are confined within the program itself. Hence, we can prove that the following *context-local* specification holds for this program.

$\text{clwp } \text{call/cc}(x. (\text{throw } 5 \text{ to } x) + 4) + 2 \{ x. x = 7 \}$

However, as there are no context-local reasoning rules for `call/cc` and `throw`, we have to unfold the definition of context-local weakest precondition proposition and prove this fact manually:

$$\begin{array}{c}
\frac{}{(\forall v. v = 7 \multimap \text{wp } K[v] \{ \Psi \}) \vdash \text{wp } K[5 + 2] \{ \Psi \}} \text{ trivial} \\
\frac{}{(\forall v. v = 7 \multimap \text{wp } K[v] \{ \Psi \}) \vdash \text{wp } K[(\text{throw } 5 \text{ to } \text{cont}(K[\] + 2)) + 4] + 2 \{ \Psi \}} \text{ THROW-WP} \\
\frac{}{(\forall v. v = 7 \multimap \text{wp } K[v] \{ \Psi \}) \vdash \text{wp } K[\text{call/cc}(x. (\text{throw } 5 \text{ to } x) + 4) + 2] \{ \Psi \}} \text{ CALLCC-WP} \\
\frac{}{\forall K, \Psi. (\forall v. v = 7 \multimap \text{wp } K[v] \{ \Psi \}) \multimap \text{wp } K[\text{call/cc}(x. (\text{throw } 5 \text{ to } x) + 4) + 2] \{ \Psi \}} \text{ introducing assumptions} \\
\frac{}{\text{clwp } \text{call/cc}(x. (\text{throw } 5 \text{ to } x) + 4) + 2 \{ x. x = 7 \}} \text{ unfolding clwp}
\end{array}$$

5 WEB SERVER REFINEMENT

The combination of continuations and concurrency allows for a simplified implementation of web servers. Such servers store explicitly captured (using `call/cc`) server-side continuations in order to track the state of communication with the client [Flatt 2017; Hendershott 2017; Krishnamurthi et al. 2007; Might 2017; Queinnec 2004]. To serve a returning user such web servers simply resume the stored continuation for that user. We refer to these servers as *continuation-based* as opposed to the more traditional *state-storing*, where server stores data that it later uses to reconstruct the state of the server for returning clients. In this section we show that two simple servers, one implemented in continuation-based style and one implemented in state-storing style, are equivalent.

Two Servers. The two servers that we consider take a number as a request. They reply to each user with the sum of the numbers that *that user* has submitted. Figure 6 shows implementations of the handler functions of the servers.³

The handlers take a connection (of type `serverConnT`) as input which consists of a pair of functions for reading the request and writing the response. The idea is that these functions are an abstraction of a TCP connection and thus the contextual equivalence can be understood as showing that clients cannot distinguish between the two implementations. These servers both internally use a table to associate a *resumption id* (a number to remember the client by) to each client. Here the functions `associate` and `get` are used for storing into the table and looking up resumption ids in the table, respectively. The state-storing implementation associates to a resumption id the sum so far. The continuation-based implementation associates a captured continuation to a resumption id. Each returning client sends its resumption id along with its request. The function `sumLoop` in the continuation based implementation is essentially a loop that sends the sum so far to the client and asks for the next number by calling `read_client`. The function `read_client` captures the current continuation and associates it with a resumption id. This resumption id is then sent to the client. After this, the server stops as the request in question is served.⁴ Notice that the `read_client` never returns. In practice the control is returned to the point after the call to this function when the client returns with the resumption id associated to the current continuation. In this case, as the connection is new, the server will supply the new connection (`reader`, `writer`) along with the new request.

Since the two servers, apart from their handlers, are identical (they just pass requests to their handlers) we here only discuss the contextual equivalence of the two handlers by showing that each handler refines the other. Here we show one of these refinements:

$$\exists \mid \Gamma \models \text{handler2} \leq_{\log} \text{handler1} : \text{ServerConnT} \rightarrow 1$$

The other is similar.

The two handlers internally use a concurrent (protected by a spin lock) table to store and look-up resumptions. The table and lock implementations are straightforward and thus omitted. Since these

³We use an ML-like syntax for the sake of brevity and legibility; our Coq formalization includes the $F_{\text{conc}, \text{cc}}^{\mu, \text{ref}}$ code for these handlers.

⁴The command `abort` is the command that ends the program (thread) and can be written in our programming language as `throw () to []` where `[]` is the empty evaluation context.


```

(* In the code below, the function `resumptionid` converts the index of a continuation in
   the table into a session id, i.e., a session cookie. Similarly, the function `result`
   formats the response appropriately, making it ready to be sent to the client. *)

1 let handler1 : ServerConnT -> 1 =
2   let tb = newTable () in
3   fun (cn : ServerConnT) ->
4     let (reader, writer) = cn in
5     match reader () with
6     (Some cid, n) ->
7       begin
8         match get tb cid with
9         None -> () (* unknown resumption id! *)
10        | Some sum -> writer (result (sum + n));
11                    writer (resumptionid (associate tb (sum + n)));
12          abort
13        end
14      | (None, n) ->
15        writer (result n);
16        let cid = associate tb n in
17        writer (resumptionid cid)

1 let read_client tb writer =
2   callcc (k. writer (resumptionid (associate tb k))); abort)
3
4 let rec sumloop m reader writer =
5   writer (result m);
6   let (v, reader, writer) = read_client tb writer in
7   sumloop (m + v) reader writer
8
9 let handler2 : ServerConnT -> 1 =
10  let tb = newTable () in
11  fun (cn : ServerConnT) ->
12    let (reader, writer) = cn in
13    match reader () with
14    (Some cid, n)->
15      begin
16        match get tb cid with
17        None -> () (*unknown resumption id!*)
18        | Some k -> throw (n,reader,writer) to k
19      end
20    | (None, n)-> sumloop n reader writer

```

Fig. 6. Two server handlers: one state-storing (top) and one continuation-based (bottom).

implementations do not use `call/cc` and `throw`, we give their relational specs using context-local weakest preconditions. Given the rule `CLWP-WP` we can use these specs during the proof of contextual refinement of handlers. We discuss the relational specs of the table before presenting the refinement of handlers.

5.1 Relational Spec for the Table and the Lock

The essence of relating the tables on both sides (specification side and implementation side) is simple. Two tables are related if their contents are. For this purpose we introduce the predicate $relTables(v, v', \gamma, \Phi)$ which states that the table v is related to the table v' where their contents are related by the binary predicate Φ . We ignore γ for now. It is only used for the internal lock protecting the table. With this definition, new tables are related (as they are both empty). The related specs for the `get` and `associate` operations require that only related values can be stored

into tables and guarantee that when looking the table up we are guaranteed to receive related values, if any.

What does it mean for the contents, *i.e.*, a number and a continuation to be related here, *i.e.*, what should we take for Φ ? The answer is simple. The number n is related to the continuation K if n is the same as the number that the suspended program in K considers to be the sum. Notice that this relation specifies some details about the program captured into K . This program uses the table itself! Therefore, it is not possible to write down the relation for the contents of the table before the table is created. As a result, the standard relational spec for the table, *e.g.*, the relational spec that we get *for free* from the fundamental theorem of logical relations, does not suffice for our application. Such standard specs require the relation on the contents to be given before the table is created:

$$\begin{aligned} & \forall \Phi. \{j \Rightarrow K[\text{newTable } ()]\}^{\text{cl}} && \text{(weaker standard spec)} \\ & \text{newTable } () \\ & \{x. \exists v'. j \Rightarrow K[v'] * \exists \gamma. \text{relTables}(x, v', \gamma, \Phi)\} \end{aligned}$$

Note the quantification over Φ outside the whole triple. Hence, we give the following stronger relational spec to our tables:

$$\begin{aligned} & \{j \Rightarrow K[\text{newTable } ()]\}^{\text{cl}} \\ & \text{newTable } () \\ & \{x. \exists v'. j \Rightarrow K[v'] * \forall \Phi. \models \exists \gamma. \text{relTables}(x, v', \gamma, \Phi)\} \\ & \{\text{relTables}(tb, tb', \gamma, \Phi) * j \Rightarrow K[\text{get } tb' \ n]\}^{\text{cl}} \\ & \text{get } tb \ n \\ & \left\{ \begin{array}{l} x. \exists v'. j \Rightarrow K[v'] * (x = v' = \text{None} \vee \\ (\exists w, w' \ x = \text{Some}(w) \wedge v' = \text{Some}(w') * \Phi(v, v')) \end{array} \right\} \\ & \{\text{relTables}(tb, tb', \gamma, \Phi) * \Phi(v, v') * j \Rightarrow K[\text{associate } tb' \ v]\}^{\text{cl}} \\ & \text{associate } tb \ v \\ & \{x. \exists n. x = n * j \Rightarrow K[n]\} \end{aligned}$$

Notice that with our stronger specification we can refer to the tables themselves in the predicate Φ that we pick for relating the contents, whereas in the (weaker standard spec) specification one has to pick this relation beforehand, and hence one cannot refer to the tables v and v' because they have not been created yet!

The predicate $\text{relTables}(tb, tb', \gamma, \Phi)$ is defined in terms of the relLocks predicate, which pertains to the relational specification of spin locks given below.

$$\text{relTables}(tb, tb', \gamma, \Phi) \triangleq \text{relLocks}(tb.\text{lock}, tb'.\text{lock}, \gamma, P_\Phi)$$

$$P_\Phi \triangleq \exists ls. \text{contents}(tb, \text{map } \pi_1 \ ls) * \text{contents}(tb', \text{map } \pi_2 \ ls) * \bigstar_{(x, x') \in ls} \Phi(x, x')$$

Here $tb.\text{lock}$ is the lock protecting the table tb . The proposition P_Φ above simply states that there is a list of pairs of values, which are pairwise related by Φ and, moreover, that the first projections of these pairs are stored in the implementation side table and the second projections of these pairs are stored in the specification side table. The contents predicate simply specifies that the index of an element in the table is its index in the list.

Relational spec for the spin lock. Similar to the table specification the lock specification also needs to be strengthened. The relational specs for locks is given below. In this specs the persistent proposition $relLocks(v, v', \gamma, P)$ states that v is a lock protecting two things: resources P and the fact that v' is not acquired; this is crucial as we discuss below. The proposition $locked(\gamma)$ states that both of the locks associated to γ are currently acquired.

$$\begin{aligned}
& \{j \Rightarrow K[\mathbf{newlock} ()]\}^{cl} \\
& \quad \mathbf{newlock} () \\
& \{x. \exists v'. j \Rightarrow K[v'] * \forall P. P \Rightarrow \exists \gamma. relLocks(x, v', \gamma, P)\} \\
& \{relLocks(v, v', \gamma, P) * j \Rightarrow K[\mathbf{acquire} v']\}^{cl} \\
& \quad \mathbf{acquire} v \\
& \{_ . j \Rightarrow K[] * locked(\gamma) * P\} \\
& \{relLocks(v, v', \gamma, P) * P * locked(\gamma) * j \Rightarrow K[\mathbf{release} v']\}^{cl} \\
& \quad \mathbf{release} v \\
& \{_ . j \Rightarrow K[]\}
\end{aligned}$$

The specification captures that whenever we acquire the lock on the implementation side, the lock on the specification side is free and can be acquired. This is necessary for showing contextual refinements because if the implementation side converges, then we need to show that so does the specification side and the acquire operation is potentially non-terminating. This also means that whenever we release the lock on the implementation side, the lock on the specification side is also released.

5.2 Proving Equivalence of Handlers

We prove only one direction here: $\Xi \mid \Gamma \models handler2 \leq_{log} handler1 : \mathbf{ServerConnT} \rightarrow \mathbf{1}$. We use the rules for weakest preconditions and executions on the specification side explained above and make use of the relational specification given above for tables, which is justified by the **CLWP-WP** rule. A key element of the proof is the choice of predicate for relating the contents of the two tables. We use the following predicate:

$$\begin{aligned}
\Phi_{handlers}(w, w') &= \exists sum \in \mathbb{N}. w' = sum \wedge \\
&\exists K. w = \mathbf{cont} \left(K \left[\begin{array}{l} \mathbf{let} (v, reader, writer) = [] \mathbf{in} \\ \mathbf{sumloop} (sum + v) \mathbf{reader} \mathbf{writer} \end{array} \right] \right)
\end{aligned}$$

The relation $\Phi_{handlers}$ essentially captures what was explained in prose earlier: the two sides consider the sum stored (as part of a continuation object in the continuation-based server) to be the same value. The relation $\Phi_{handlers}$ above is indeed capturing the essence of the intuitive reason why the two implementations of handlers have contextually equivalent behavior. According to the definition of our logical relations, to show logical relatedness we need to show that given any two related contexts the two programs behave in an observationally related way. Since at the time of picking the predicate above we do not know what contexts we will have to operate under, we have to consider that our code of interest is inside some arbitrary (hence existentially quantified in the continuation-based) evaluation context. The actual value of this evaluation context is not important as the control never reaches this evaluation context; the thread is ended with an **abort** before that. Note the use of the table itself (referenced inside **sumloop**) in the $\Phi_{handlers}$.

6 ONE-SHOT CALL/CC

In this section we consider a more technical verification challenge involving continuations, due to Friedman and Haynes [1985]. The challenge is to show that `call/cc` can be implemented using references and one-shot continuations, *i.e.*, continuations that can only be called once. This problem has been studied for *sequential* higher-order languages with references in Dreyer et al. [2012]; Støvring and Lassen [2007], with pen-and-paper proofs. Here we show that the equivalence also holds in our concurrent language (subtly so; because we are using *may* contextual equivalence) and we give a mechanized formal proof thereof.

For this verification we introduce two closed programs $\mathbb{C}\mathbb{C}$ (essentially a `call/cc`) and $\mathbb{C}\mathbb{C}_1$ (encoding a one shot continuation) both of type $\forall\alpha. (\text{cont}(\alpha) \rightarrow \alpha) \rightarrow \alpha$.

$$\begin{aligned}\mathbb{C}\mathbb{C} &\triangleq \Lambda \lambda f. \text{ call/cc } (x. f \ x) \\ \mathbb{C}\mathbb{C}_1 &\triangleq \Lambda \lambda f. \text{ let } b = \text{ref}(\text{false}) \text{ in} \\ &\quad \text{call/cc} \left(x. f \left(\text{cont} \left(\begin{array}{l} \text{let } y = [] \text{ in} \\ \text{if } !b \text{ then } \Omega \text{ else } b \leftarrow \text{true}; \\ \text{throw } y \text{ to } x \end{array} \right) \right) \right)\end{aligned}$$

Here Ω is the trivially diverging expression. When applied, the one-shot continuation, $\mathbb{C}\mathbb{C}_1$, first allocates a *one-shot bit* b and then calls the given function with a continuation that uses b to ensure that the continuation is only called once. Using one-shot continuations, we simulate normal continuations by defining $\mathbb{C}\mathbb{C}'$:

$$\begin{aligned}\mathbb{C}\mathbb{C}' &\triangleq \Lambda \lambda f. \text{ let } \ell = \text{ref}(\text{cont}([])) \text{ in } G \ f \\ G &\triangleq \text{rec } G(f) = \text{let } x = \mathbb{C}\mathbb{C}_1 _ \left(\begin{array}{l} \lambda y. \ell \leftarrow y; \\ f \ (\text{cont}(\text{throw } \text{cont}([]) \text{ to } !\ell)) \end{array} \right) \text{ in} \\ &\quad \mathbb{C}\mathbb{C}_1 _ (\lambda y. G \ (\lambda z. \text{throw } x \text{ to } y))\end{aligned}$$

The expression $\mathbb{C}\mathbb{C}'$ above has the same type as $\mathbb{C}\mathbb{C}$. $\mathbb{C}\mathbb{C}'$ perhaps looks fairly complex but the intuition is straightforward. It first allocates ℓ with the trivial continuation, then it takes a one-shot continuation and updates ℓ . When the one-shot continuation is used, it will first grab another *fresh* one-shot continuation and update ℓ with it before continuing. Hence, intuitively, every time the one-shot continuation stored in ℓ is used, it is immediately refreshed, thus mimicking the behavior of $\mathbb{C}\mathbb{C}$. We now prove that $\mathbb{C}\mathbb{C}$ is contextually equivalent to $\mathbb{C}\mathbb{C}'$:

THEOREM 6.1. $\cdot \mid \cdot \models \mathbb{C}\mathbb{C} \approx_{\text{ctx}} \mathbb{C}\mathbb{C}' : \forall\alpha. (\text{cont}(\alpha) \rightarrow \alpha) \rightarrow \alpha$

We only discuss one side of the refinement, namely, $\mathbb{C}\mathbb{C}' \leq_{\text{ctx}} \mathbb{C}\mathbb{C}$. The invariant that we need in order to prove this refinement is intuitively that the one-shot bit of the continuation stored in ℓ is always storing the value `false` indicating that the one-shot continuation is unused. This fact is expressed in terms of Iris invariants as follows:

$$\boxed{\exists b. b \mapsto_i \text{false} * \ell \mapsto_i \text{cont} \left(\begin{array}{l} \text{let } y = [] \text{ in} \\ \text{if } !b \text{ then } \Omega \text{ else } b \leftarrow \text{true}; \\ \text{throw } y \text{ to } \text{cont}(K[\text{restore}(\ell)]) \end{array} \right)}^{\mathcal{N}.\mathbb{C}\mathbb{C}}$$

where

$$\text{restore}(\ell) \triangleq \text{let } x = [] \text{ in } \mathbb{C}\mathbb{C}_1 _ (\lambda y. G \ (\lambda _ . \text{throw } x \text{ to } y))$$

Here K is the continuation that is captured by $\mathbb{C}\mathbb{C}$. The invariant above is exactly the invariant that Dreyer et al. [2012] use when translated to our system.⁵

This invariant suffices for a *sequential* programming language. However, in our concurrent setting, the “continuation” captured by $\mathbb{C}\mathbb{C}'$ may be shared among multiple threads and, if they use it concurrently, it may happen that a thread is using the continuation captured by $\mathbb{C}\mathbb{C}'$ and before this thread manages to capture another one-shot continuation and restore ℓ , another thread attempts to use the *then invalid* one-shot continuation, and hence it diverges.

We prove that the contextual refinement still holds (despite the possibility of divergence). However, because of the possible racing, we need to use a weaker invariant:

$$\boxed{\exists b, M. \text{OneShotBits}(M) * \text{isOneShotBit}(b) * \left(\bigstar_{r \in M} \exists v \in \{\text{true}, \text{false}\}. r \mapsto_i v \right) * \ell \mapsto_i \text{cont} \left(\begin{array}{l} \text{let } y = [] \text{ in} \\ \text{if } !b \text{ then } \Omega \text{ else } b \leftarrow \text{true}; \\ \text{throw } y \text{ to cont}(K[\text{restore}(\ell)]) \end{array} \right)} \quad \mathcal{N}, \mathbb{C}\mathbb{C}$$

This invariant says that ℓ stores a one-shot continuation with a one-shot bit b and that we have a set of bits that, intuitively, have been associated to one-shot continuations. We also know that b is one such one-shot bit, $\text{isOneShotBit}(b)$. The predicates $\text{OneShotBits}()$ and $\text{isOneShotBit}()$ are defined using Iris resources.⁶ Here, we only need to know two things about them: $\text{isOneShotBit}(b)$ is persistent and:

$$\text{OneShotBits}(M) * \text{isOneShotBit}(b) \vdash b \in M \quad (\text{in-bits})$$

Persistence allows us to retain the information $\text{isOneShotBit}(b)$ once we have opened the invariant and have read ℓ . Due to the race condition explained above, when we open the invariant we know, by (in-bits), that there is a value $v \in \{\text{true}, \text{false}\}$ stored in b , and this suffices for being able to complete the refinement proof.

7 CORRECTNESS OF CONTINUATION BASED COOPERATIVE CONCURRENCY

Cooperative concurrency, a.k.a. light-weight concurrency, is a form of concurrency where threads cooperate and use a *yield* command to relinquish control to other threads. This is in contrast to *preemptive* concurrency, where the operating system preempts and schedules threads. Cooperative concurrency is often implemented using continuations [Haynes et al. 1984]. Forking a new thread suspends the execution of the current thread, enqueues the suspension in a queue, and runs the forked thread. The yield command dequeues a previously enqueued suspension (thread) and resumes it, after enqueueing the current continuation.

In this section, we prove correctness of a continuation-based implementation of *cooperative concurrency*. It is not entirely obvious how to state the desired correctness property. Here we use a relational approach inspired by compiler correctness, and show that a language with built-in cooperative concurrency can be compiled into a continuation-based implementation of cooperative concurrency. We prove the correctness of this compilation, by showing that a compiled program refines its source program.

The programming language with built-in cooperative concurrency is called $F_{cc,coop}^{\mu,ref}$. This language serves as our specification of cooperative concurrency. Concretely, $F_{cc,coop}^{\mu,ref}$ provides two primitive commands **Cfork** and **yield**. The semantics of $F_{cc,coop}^{\mu,ref}$ keeps track of the currently running thread

⁵In the work of Dreyer et al. [2012], invariants were called *islands*.

⁶The one-shot bit predicates are defined using the authoritative resource algebra over the resource algebra of finite sets of locations (where the algebra operation is set union). See JUNG et al. [2018]

and keeps executing that thread until it reaches a fork or a yield command. In the former case, it starts a new thread and starts executing that. In the latter case, the semantics picks another thread and proceeds with executing that. The programming language that we consider as the target of translation of $F_{cc,coop}^{\mu,ref}$ is $F_{cc}^{\mu,ref}$, the sequential fragment of $F_{conc,cc}^{\mu,ref}$, i.e., $F_{conc,cc}^{\mu,ref}$ without **fork** and **cas**.

In the rest of this section, we define the precise syntax and semantics of the source and target languages, the translation of cooperative concurrency, and a cross-language logical relation between the source and target language, which we use to show the correctness of the translation.

7.1 Syntax and Semantics of $F_{cc}^{\mu,ref}$ and $F_{cc,coop}^{\mu,ref}$

The programming languages $F_{conc,cc}^{\mu,ref}$, $F_{cc}^{\mu,ref}$ and $F_{cc,coop}^{\mu,ref}$ have the same types and similar typing rules for shared parts; we use $\vdash_{coop}$ and $\vdash_{seq}$ to distinguish them. The typing rules for **Cfork** and **yield** in $F_{cc,coop}^{\mu,ref}$ are as follows:

$$\begin{array}{c} \text{T-CFORK} \\ \frac{\Xi \mid \Gamma \vdash_{coop} e : \tau}{\Xi \mid \Gamma \vdash_{coop} \text{Cfork } \{e\} : 1} \end{array} \qquad \begin{array}{c} \text{T-YIELD} \\ \Xi \mid \Gamma \vdash_{coop} \text{yield} : 1 \end{array}$$

The language $F_{cc}^{\mu,ref}$ is a fragment of $F_{conc,cc}^{\mu,ref}$ with the same operational semantics. Formally, the head steps for $F_{cc}^{\mu,ref}$ are identical to the corresponding fragment in $F_{conc,cc}^{\mu,ref}$. For the general execution, however, we define the *sequential step* relation, $\rightarrow_{seq}$, in place of the thread-pool step $\rightarrow$ for $F_{conc,cc}^{\mu,ref}$.

$$\frac{(e, \sigma) \rightarrow_K (e', \sigma')}{(K[e]; \sigma) \rightarrow_{seq} (K[e']; \sigma')} \qquad (K[\text{throw } v \text{ to cont}(K')]; \sigma) \rightarrow_{seq} (K'[v]; \sigma)$$

The head-step relation for the fragment of $F_{cc,coop}^{\mu,ref}$, apart from **Cfork** and **yield**, is defined identically to $F_{conc,cc}^{\mu,ref}$ and $F_{cc}^{\mu,ref}$. For general reduction, in place of a thread-pool step we define a *cooperative step*, $\rightarrow_{coop}$. The step $(\vec{e}; n; \sigma) \rightarrow_{coop} (\vec{e}'; n'; \sigma')$ is to be understood as a step transforming the thread pool $\vec{e}$ into the thread pool $\vec{e}'$ and state σ into σ' while changing the current thread being executed from thread number n to thread number n' . All head steps (and **throw**) do not change the current thread. For **Cfork** and **yield** we have:

$$\frac{\text{length}(e_1) = n \quad m = \text{length}(\vec{e}_1, K[\text{Cfork } \{e\}], \vec{e}_2)}{(\vec{e}_1, K[\text{Cfork } \{e\}], \vec{e}_2; n; \sigma) \rightarrow_{coop} (\vec{e}_1, K[], \vec{e}_2, e; m; \sigma)}$$

$$\frac{\text{length}(e_1) = n \quad 0 \leq m < \text{length}(\vec{e}_1, K[\text{yield}], \vec{e}_2)}{(\vec{e}_1, K[\text{yield}], \vec{e}_2; n; \sigma) \rightarrow_{coop} (\vec{e}_1, K[], \vec{e}_2; m; \sigma)}$$

A quick inspection of the cooperative step relation should make it clear that $F_{cc,coop}^{\mu,ref}$ does really capture cooperative concurrency. Hence, by showing that $F_{cc,coop}^{\mu,ref}$ can be *correctly* compiled to $F_{cc}^{\mu,ref}$ by a compiler that compiles **Cfork** and **yield** using continuations, we establish correctness of an implementation of cooperative concurrency using continuations.

7.2 Translation and its Correctness

The translation from $F_{cc,coop}^{\mu,ref}$ into $F_{cc}^{\mu,ref}$ is very straightforward. It simply translates **Cfork** and **yield** to programs that use the the light-weight thread library LiThr given in Figure 7.⁷ This simple and

⁷Note that the specification side, $F_{cc,coop}^{\mu,ref}$, does not restrict scheduling, i.e., **yield** non-deterministically chooses another thread. Here, on the implementation side we have chosen to use a queue for scheduling of threads.

$$\begin{aligned}
\text{LiThr} &\triangleq \text{let } Q = \text{newQueue in} \\
&\quad \text{let } Frk = \lambda x. \text{call/cc} \left(\begin{array}{l} y.\text{enqueue } Q \ y; \\ \text{throw } x \text{ to cont}([\] ()) \end{array} \right) \text{ in} \\
&\quad \text{let } Yld = \lambda _ . \text{call/cc} \left(\begin{array}{l} x. \text{enqueue } Q \ x; \\ \text{let } y = \text{dequeue } Q \text{ in} \\ \text{match } y \text{ with} \\ \quad \text{Some}(z) \Rightarrow \text{throw } () \text{ to } z \\ \mid \text{None} \quad \Rightarrow () \\ \text{end} \end{array} \right) \text{ in} \\
&\quad (Frk, Yld)
\end{aligned}$$

Fig. 7. The light-weight thread library LiThr.

minimalistic light-weight thread library provides two functions: one for forking a new thread, *Frk*, and one for relinquishing control to other threads, *Yld*. Notice that the *dequeue* operation can return *None* if the queue empty. However, since in *Yld* the *dequeue* operation is immediately preceded by an *enqueue* operation, this will never happen.

We translate programs in $F_{cc,coop}^{\mu,ref}$ into programs of $F_{cc}^{\mu,ref}$ as follows:

$$\mathbb{C}omp(e) \triangleq \text{let } (\mathbb{F}, \mathbb{Y}) = \text{LiThr in } \langle\langle e \rangle\rangle$$

Where $\langle\langle e \rangle\rangle$ translates $F_{cc,coop}^{\mu,ref}$ programs into programs that use special free variables $\mathbb{F}$ and $\mathbb{Y}$ as functions for forking and yielding respectively. The translation $\langle\langle e \rangle\rangle$ is very simple. It only changes *Cfork* and *yield* and leaves the rest of the program untouched (similar to *throw* below):

$$\begin{aligned}
\langle\langle \text{throw } e \text{ to } e' \rangle\rangle &\triangleq \text{throw } \langle\langle e \rangle\rangle \text{ to } \langle\langle e' \rangle\rangle \\
\langle\langle \text{yield} \rangle\rangle &\triangleq \mathbb{Y} () \quad \quad \quad \langle\langle \text{Cfork } \{e\} \rangle\rangle \triangleq \mathbb{F} (\lambda _ . \langle\langle e \rangle\rangle)
\end{aligned}$$

Notice that for *Cfork* we use a λ to turn e into a thunk.

LEMMA 7.1 (TYPING OF TRANSLATION). *Let e be a program in $F_{cc,coop}^{\mu,ref}$ such that $\Xi \mid \Gamma \vdash_{coop} e : \tau$. The following typing judgement holds for the translation of e .*

$$\Xi \mid \Gamma, \mathbb{Y} : 1 \rightarrow 1, \mathbb{F} : (1 \rightarrow 1) \rightarrow 1 \vdash_{seq} \langle\langle e \rangle\rangle : \tau$$

Correctness of translation. We will show correctness of our translation by showing observational refinement. To this end, we define propositions $e \Downarrow_{seq}$ and $e \Downarrow_{coop}$, which state when programs of $F_{cc}^{\mu,ref}$ and $F_{cc,coop}^{\mu,ref}$ terminate:

$$\begin{aligned}
e \Downarrow_{seq} &\triangleq \exists v, \sigma. (e; \emptyset) \rightarrow_{seq}^* (v; \sigma) \\
e \Downarrow_{coop} &\triangleq \exists \vec{e}, \sigma. (e; 0; \emptyset) \rightarrow_{coop}^* (\vec{e}; n; \sigma) \wedge e_n \text{ is a value}
\end{aligned}$$

Notice that programs of $F_{cc}^{\mu,ref}$ terminate whenever the current thread at the time has terminated.

THEOREM 7.2 (CORRECTNESS OF TRANSLATION). *Let e be a closed program of type τ , $\cdot \mid \cdot \vdash_{coop} e : \tau$. The following holds*

$$\text{if } \mathbb{C}omp(e) \Downarrow_{seq} \text{ then } e \Downarrow_{coop}$$

Intuitively, this theorem expresses that if the compiled program $\mathbb{C}omp(e)$ produces a result so would the original program (standard observational refinement). We prove this theorem by setting up a cross-language logical relation between $F_{cc}^{\mu, ref}$ and $F_{cc, coop}^{\mu, ref}$. We show that the translation $\langle\langle e \rangle\rangle$ is suitably related to e . Theorem 7.2 then follows essentially because $\mathbb{C}omp(e)$ reduces to $\langle\langle e \rangle\rangle[Frk/\mathbb{F}][Yld/\mathbb{Y}]$.

7.3 Cross-language Logical Relation

The types and basic terms of both $F_{cc, coop}^{\mu, ref}$ and $F_{cc}^{\mu, ref}$ are the same as $F_{conc, cc}^{\mu, ref}$. Hence, our cross-language logical relation is very similar to the logical relation presented in Section 3, except that it is defined for pairs of expressions, values and evaluation contexts where the first component pertains to $F_{cc}^{\mu, ref}$ and the second to $F_{cc, coop}^{\mu, ref}$. In fact, the only part of the logical relation in Section 3 that we need to change is the definition of observational refinement. To define the observational refinement relation between $F_{cc}^{\mu, ref}$ and $F_{cc, coop}^{\mu, ref}$ we need to introduce resources that we may use for keeping track of the execution on the specification side, i.e., the $F_{cc, coop}^{\mu, ref}$ side. For this purpose, in addition to propositions $\ell \mapsto_s v$ and $j \models e$ for keeping track of the heap and the (light-weight) threads on the specification side, we need a proposition to keep track of the current thread on the specification side. Hence, we introduce the proposition $CurTh(j)$, for asserting that the current thread being run is thread j . For all the basic terms of $F_{cc, coop}^{\mu, ref}$, the rules for execution on the specification side remain similar to the ones in $F_{conc, cc}^{\mu, ref}$, except that they require the thread being evaluated to be the current thread. As an example, the rule for storing a value in a reference on the specification side is given below. The only rules that change the current thread are those pertaining to **Cfork** and **yield**.

$$\frac{\ell \mapsto_s v \quad CurTh(j) \quad j \models K[\ell \leftarrow w]}{\models CurTh(j) * \ell \mapsto_s w * j \models K[()]} \quad \frac{CurTh(j) \quad j \models K[\mathbf{Cfork} \{e\}]}{\models \exists j'. CurTh(j') * j \models K[()] * j' \models e}$$

$$\frac{CurTh(j) \quad j \models K[\mathbf{yield}]}{\models CurTh(j) * j \models K[()]} \quad \frac{CurTh(j) \quad j \models K[\mathbf{yield}] \quad j' \models e'}{\models CurTh(j') * j \models K[()] * j' \models e'}$$

Note that there are two rules pertaining to the execution of **yield** as it may result in continuing the execution of the same thread.

For our cross-language logical relations we define the observational refinement relation, O^{cross} as follows:

$$O^{\text{cross}}(e, e') \triangleq \forall j. CurTh(j) * j \models e' * \text{wp } e \{ \exists j', w. j' \models w * CurTh(j') \}$$

Note how the specification side is expected to be in a thread that has reached a value in its execution, similarly to how we defined $\Downarrow_{coop}$ above.

7.4 Proof of Correctness of Translation

In order to prove correctness of the translation we need to reason about a relation between the internal state of the LiThr library and threads on the specification side. In particular, for each continuation K stored in the internal queue of LiThr there must be a thread $j' \models e'$ on the specification side such that $K[()]$ observationally refines e' . We use the proposition $LiThrInv$ for

this purpose.⁸

$$LiThrInv \triangleq \boxed{\exists l. isQueue(Q, l) * \bigstar_{K \in l} \exists j', e', j' \Rightarrow e' * \square(O^{\text{cross}}(K[], e'))}^{\mathcal{N}.LiThr}$$

The proposition $isQueue(Q, l)$ asserts that Q is a queue whose contents are the list l . The two operations of the queue, *enqueue* and *dequeue*, have the following specs:

$$\begin{aligned} & \{isQueue(Q, l)\}^{cl} \\ & \text{enqueue } Q \ v \\ & \{x. x = () * isQueue(Q, v :: l)\} \\ & \{isQueue(Q, l)\}^{cl} \\ & \text{dequeue } Q \\ & \{x. (\exists l'. l = l' \# [v] * x = Some(v) * isQueue(Q, l')) \vee (x = None * l = [] * isQueue(Q, l))\} \end{aligned}$$

where $\#$, $::$ and $[\cdot]$ are the usual list operations and $[]$ is the empty list.

To prove the desired correctness theorem (Theorem 7.2), we now prove the fundamental theorem for our cross-language logical relation. It says that, under the assumption that the internal state of the library *LiThr* is appropriate, any well-typed $F_{cc,coop}^{\mu,ref}$ program is refined by its translation when linked with the *LightTh* library.

THEOREM 7.3 (FUNDAMENTAL THEOREM OF CROSS-LANGUAGE LOGICAL RELATIONS). *Let e be a program in $F_{cc,coop}^{\mu,ref}$ such that $\Xi \mid \Gamma \vdash_{coop} e : \tau$. Then we have*

$$LiThrInv \Rightarrow \Xi \mid \Gamma \models \llbracket e \rrbracket [Frk, Yld/\mathbb{F}, \mathbb{Y}] \leq_{log} e : \tau$$

PROOF. By induction on the derivation of $\Xi \mid \Gamma \vdash_{coop} e : \tau$. All cases, except for **Cfork** and **yield** follow similarly to their counterpart in the proof of Theorem 3.1. Cases **Cfork** and **yield** follow by the fact that their translations are applications to $\mathbb{F}$ and $\mathbb{Y}$ respectively under the assumption that the invariant *LiThrInv* holds for the internal queue of *LiThr*. $\square$

Theorem 7.2 now follows from Theorem 7.3 (in the same way that Theorem 3.2 followed from Theorem 3.1), using and the fact that the program $\text{Comp}(e)$ reduces to $\llbracket e \rrbracket [Frk, Yld/\mathbb{F}, \mathbb{Y}]$.

8 MECHANIZATION IN COQ

Taking advantage of the Coq formalization of Iris and Iris Proof Mode (IPM) [Krebbers et al. 2017b], we have mechanized all the technical development and results presented in this paper in Coq. This includes mechanizing the small-step operational semantics of $F_{conc,cc}^{\mu,ref}$, $F_{cc}^{\mu,ref}$ and $F_{cc,coop}^{\mu,ref}$, and instantiating Iris with them. Our Coq development is about 15800 lines and includes proofs of contextual refinements for pairs of fine-grained/coarse-grained stacks and counters which we omitted discussion of for reasons of space.

For binders, we use the Autosubst library [Schäfer et al. 2015] which facilitates the use of de Bruijn indices by providing support for simplification of substitutions. In $F_{conc,cc}^{\mu,ref}$, $F_{cc}^{\mu,ref}$ and

⁸This is a slight simplification. The proposition *LiThrInv* is defined using Iris's non-atomic invariants. These are invariants that can be kept open for multiple steps of computation. They are admissible in our system because we have no real concurrency causing racy behavior. It is crucial to be able to keep the *LiThrInv* open for multiple steps so as to prove that the *dequeue* operation in *Yld* never returns a *None* value. (If the languages had included concurrent racy behaviour, then we could have used a lock in the library and then we would have been able to use standard Iris atomic invariants.) The definition of O^{cross} is also slightly simplified here. It should be slightly adjusted to allow for the use of non-atomic invariants. Iris's weakest preconditions by default only allow atomic invariants.

$F_{cc,coop}^{\mu,ref}$, evaluation contexts are also values and hence also expressions. This forces us to define these mutually inductively. This means that we need to derive the induction principle for these inductive types in Coq by hand. Furthermore, we have to help Autosubst in deriving substitution and simplification lemmas for $F_{conc,cc}^{\mu,ref}$ that it should otherwise automatically infer. This is mainly why the definition of $F_{conc,cc}^{\mu,ref}$, $F_{cc}^{\mu,ref}$ and $F_{cc,coop}^{\mu,ref}$ combined takes up about 20% of the whole Coq development.

9 RELATED WORK

There has been a considerable body of work on (delimited) continuations, but, we are not aware of any logics or relational models for reasoning about concurrent programs with continuations, let alone a mechanized framework for relational verification of concurrent programs with continuations.

Program logics for reasoning about continuations. Delbianco and Nanevski [2013] present a type theory for Hoare-style reasoning about an imperative higher-order programming language with (algebraic) continuations, but without concurrency. The system of Delbianco and Nanevski [2013] does not allow higher-order code (including continuations) to be stored in the heap. Note that storing higher-order code in the heap is essential for all our case studies: implementing cooperative concurrency with continuations, implementing the continuation-based web servers and implementing continuations in terms of one-shot continuations. Crolard and Polonowski [2012] develop a program logic for reasoning about jumps but their sequential programming language features no heap or recursive types. Berger [2010] presents a program logic for reasoning about programs in a programming language which is essentially an extension of PCF [Plotkin 1977] with continuations.

Relational reasoning about continuations. The work most closely related to ours is that of Dreyer et al. [2012] who consider a variety of different stateful programming languages and investigate the impact of the higher-order state and control effects (including `call/cc` and `throw`). In contrast to our work, they do not consider concurrency. Moreover, they reason directly in a model, whereas we define our logical relation using a program logic (Iris), which means that we can reason more compositionally and at a higher level of abstraction. Another advantage of using Iris, is that we have been able to leverage its Coq formalization and thus to mechanize all of our development. As mentioned in Section 6, our proof that continuations can be expressed in terms of one-shot continuations is inspired by *loc. cit.*

There are several other works on relational reasoning for sequential programming languages with continuations, e.g., Felleisen and Hieb [1992]; Laird [1997]; Støvring and Lassen [2007]. These differ from our work at least in that they do not consider concurrency.

Relational reasoning about concurrency. There has been much work on relational reasoning about concurrent higher-order imperative programs, without continuations. The work most closely related to ours also is that of Krebbers et al. [2017b], who develop mechanized logical relations (in Iris) for reasoning about contextual equivalence of programs in $F_{conc}^{\mu,ref}$, a language similar to the one we consider but without `call/cc` and `throw`. The approach in *loc. cit.* is based on earlier, non-mechanized logical relations for fine-grained concurrent programs [Birkedal et al. 2012; Turon et al. 2013a,b]. These relational models give an alternative method to linearizability [Herlihy and Wing 1990] for reasoning about contextual refinement for fine-grained concurrent programs. The logical relations method also works in the presence of higher-order programs, which linearizability traditionally struggles with, although there has been some recent promising developments [Cerone et al. 2014; Murawski and Tzevelekos 2017]. In this paper, we have extended the method of logical relations

for reasoning about contextual refinement for higher-order fine-grained concurrent programs to work for programs that also use continuations.

Correctness of compilation of cooperative concurrency. The only other work that we are aware of which proves correctness of compilation of cooperative concurrency is by Nakata and Saar [2013]. Nakata and Saar [2013] compile a simple while language extended with procedures, first-order store (only integers can be stored) and cooperative concurrency. Their target language is again a simple while language extended with procedures, delimited continuations and a store that can store in each cell either an integer value or a captured continuation. They use a syntactic proof of correctness as opposed to our semantic (using a logical relations model) proof. Such syntactic proof methods do not scale to higher-order programming languages with advanced features like higher-order store and impredicative polymorphism, e.g., $F_{conc,cc}^{\mu,ref}$ in our case.

10 CONCLUSION AND FUTURE WORK

We have developed a logical relation for $F_{conc,cc}^{\mu,ref}$, a programming language with advanced features such as impredicative polymorphism à la system F, higher-order mutable references, recursive types, concurrency and most notably continuations. We have devised new non-context-local proof rules for reasoning about weakest preconditions in Iris in the presence of continuations and also introduced context-local weakest preconditions for regaining context-local reasoning about expressions that do not involve non-local control flow. We have defined our relational model and proved properties thereof in the Iris program logic framework. This has greatly simplified the definition of our relational model, the existence of which is non-trivial because of the type-world circularity [Ahmed 2004; Ahmed et al. 2002; Birkedal et al. 2011]. Furthermore, working inside Iris has enabled us to mechanize the entire development presented in this paper on top of the Coq proof assistant.

We have demonstrated how our logical relation can be used to establish contextual equivalence for a pair of simplified web-server implementations: one storing the state explicitly and one storing the current continuation. The application of context local reasoning in the middle of our logical relatedness proofs demonstrates the usefulness and versatility of context-local weakest preconditions. Finally, we have also given the first (mechanized) proof of the correctness of Friedman and Haynes [1985] encoding of continuations by means of one-shot continuations in a concurrent programming language.

We developed a cross-language logical relation between $F_{cc}^{\mu,ref}$ and $F_{cc,coop}^{\mu,ref}$. We used this logical relation to give a compiler-correctness-inspired proof of correctness of the continuation-based implementation of cooperative concurrency. This is to the best of our knowledge the first formal proof of correctness of continuation-based cooperative concurrency for a programming language with a rich advanced features and types.

In the future, we wish to extend our mechanization to reason about delimited continuations [Danvy and Filinski 1990; Felleisen 1988]. Currently our mechanized reasoning is done interactively, in the same style as one reasons in Coq. In the future, we would also like to complement that with more automated reasoning methods.

ACKNOWLEDGMENTS

The first author is a postdoctoral fellow of the Flemish research fund (FWO). This project was supported in part by the FWO grant (grant no. G.0962.17N), the FWO travel grant (V435817N), the EU Types (CA15123) short scientific mission (STSM) grant (reference: 40667) and by the ModuRes Sapere Aude Advanced Grant from The Danish Council for Independent Research for the Natural Sciences (FNU).

REFERENCES

- Amal Ahmed. 2004. *Semantics of Types for Mutable State*. Ph.D. Dissertation. Princeton University.
- Amal J. Ahmed, Andrew W. Appel, and Roberto Virga. 2002. A Stratified Semantics of General References Embeddable in Higher-Order Logic. In *Proceedings of 17th Annual IEEE Symposium Logic in Computer Science*. IEEE Computer Society Press, 75–86.
- Andrew Appel and David McAllester. 2001. An Indexed Model of Recursive Types for Foundational Proof-Carrying Code. *TOPLAS* 23, 5 (2001), 657–683.
- Andrew Appel, Paul-André Mellès, Christopher Richards, and Jérôme Vouillon. 2007. A Very Modal Model of a Modern, Major, General Type System. In *POPL*.
- Martin Berger. 2010. *Program Logics for Sequential Higher-Order Control*. Springer Berlin Heidelberg, Berlin, Heidelberg, 194–211.
- Lars Birkedal, Bernhard Reus, Jan Schwinghammer, Kristian Støvring, Jacob Thamsborg, and Hongseok Yang. 2011. Step-Indexed Kripke Models over Recursive Worlds. In *POPL*.
- Lars Birkedal, Filip Sieczkowski, and Jacob Thamsborg. 2012. A Concurrent Logical Relation. In *CSL*.
- Andrea Cerone, Alexey Gotsman, and Hongseok Yang. 2014. Parameterised Linearisability. In *ICALP*.
- T. Crolard and E. Polonowski. 2012. Deriving a Floyd-Hoare logic for non-local jumps from a formulæ-as-types notion of control. *The Journal of Logic and Algebraic Programming* 81, 3 (2012), 181 – 208. The 22nd Nordic Workshop on Programming Theory (NWPT 2010).
- Pedro da Rocha Pinto, Thomas Dinsdale-Young, and Philippa Gardner. 2014. TaDA: A Logic for Time and Data Abstraction. In *ECOOP*. 207–231.
- Olivier Danvy and Andrzej Filinski. 1990. Abstracting Control. In *Proceedings of the 1990 ACM Conference on LISP and Functional Programming*.
- Germán Andrés Delbianco and Aleksandar Nanevski. 2013. Hoare-style reasoning with (algebraic) continuations. *ACM SIGPLAN Notices* 48, 9 (2013), 363–376.
- Thomas Dinsdale-Young, Lars Birkedal, Philippa Gardner, Matthew Parkinson, and Hongseok Yang. 2013. Views: Compositional Reasoning for Concurrent Programs. In *POPL*.
- T. Dinsdale-Young, M. Dodds, P. Gardner, M. Parkinson, and V. Vafeiadis. 2010. Concurrent abstract predicates. In *ECOOP*. 504–528.
- D. Dreyer, A. Ahmed, and L. Birkedal. 2011. Logical Step-Indexed Logical Relations. *LMCS* 7, 2:16 (2011).
- Derek Dreyer, Georg Neis, and Lars Birkedal. 2012. The impact of higher-order state and control effects on local relational reasoning. *Journal of Functional Programming* 22, 4-5 (2012), 477–528.
- Mattias Felleisen. 1988. The Theory and Practice of First-class Prompts. In *Proceedings of the 15th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*.
- Matthias Felleisen and Robert Hieb. 1992. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science* 103, 2 (1992), 235 – 271.
- Matthew Flatt. 2017. More: Systems Programming with Racket. <https://docs.racket-lang.org/more/index.html>.
- Daniel P. Friedman and Christopher T. Haynes. 1985. Constraining Control. In *Proceedings of the 12th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages (POPL '85)*. ACM, New York, NY, USA, 245–254.
- Christopher T. Haynes, Daniel P. Friedman, and Mitchell Wand. 1984. Continuations and Coroutines (*LFP '84*).
- Greg Hendershott. 2017. <http://www.greghendershott.com/2014/09/written-in-racket.html>.
- Maurice P. Herlihy and Jeannette M. Wing. 1990. Linearizability: a correctness condition for concurrent objects. *TOPLAS* 12, 3 (1990), 463–492.
- Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. 2016. Higher-order ghost state. In *ICFP*. 256–269.
- RALF JUNG, ROBBERT KREBBERS, JACQUES-HENRI JOURDAN, ALEŠ BIZJAK, LARS BIRKEDAL, and DEREK DREYER. 2018. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *Journal of Functional Programming* 28 (2018), e20. <https://doi.org/10.1017/S0956796818000151>
- Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *POPL*. 637–650.
- Robbert Krebbers, Ralf Jung, Aleš Bizjak, Jacques-Henri Jourdan, Derek Dreyer, and Lars Birkedal. 2017a. The essence of higher-order concurrent separation logic. In *European Symposium on Programming (ESOP)*.
- Robbert Krebbers, Amin Timany, and Lars Birkedal. 2017b. Interactive Proofs in Higher-Order Concurrent Separation Logic. In *POPL*.
- Shriram Krishnamurthi, Peter Walton Hopkins, Jay McCarthy, Paul T Graunke, Greg Pettyjohn, and Matthias Felleisen. 2007. Implementation and use of the PLT Scheme web server. *Higher-Order and Symbolic Computation* 20, 4 (2007), 431–460.
- Morten Krogh-Jespersen, Kasper Svendsen, and Lars Birkedal. 2017. A Logical Account of a Type-and-Effect System. In *POPL*.

- James Laird. 1997. Full Abstraction for Functional Languages with Control. In *Proceedings of the 12th Annual IEEE Symposium on Logic in Computer Science (LICS '97)*. IEEE Computer Society, Washington, DC, USA, 58–. <http://dl.acm.org/citation.cfm?id=788019.788859>
- Ruy Ley-Wild and Aleksandar Nanevski. 2013. Subjective Auxiliary State for Coarse-Grained Concurrency. In *POPL*.
- Matt Might. 2017. <http://matt.might.net/articles/low-level-web-in-racket/>.
- Andrzej S. Murawski and Nikos Tzevelekos. 2017. Higher-Order Linearisability. In *CONCUR 2017*.
- Keiko Nakata and Andri Saar. 2013. Compiling Cooperative Task Management to Continuations. In *Fundamentals of Software Engineering*, Farhad Arbab and Marjan Sirjani (Eds.).
- Aleksandar Nanevski, Ruy Ley-Wild, Ilya Sergey, and Germán Andrés Delbianco. 2014. Communicating State Transition Systems for Fine-Grained Concurrent Resources. In *ESOP*.
- Peter W. O’Hearn. 2007. Resources, Concurrency and Local Reasoning. *Theor. Comput. Sci.* 375, 1-3 (2007), 271–307.
- Andrew M. Pitts. 2005. Typed Operational Reasoning. In *Advanced Topics in Types and Programming Languages*, B. C. Pierce (Ed.). The MIT Press, Chapter 7, 245–289.
- Gordon D. Plotkin. 1977. LCF considered as a programming language. *Theoretical computer science* 5, 3 (1977), 223–255.
- Christian Queinnec. 2004. Continuations and web servers. *Higher-Order and Symbolic Computation* 17, 4 (2004), 277–295.
- Steven Schäfer, Tobias Tebbi, and Gert Smolka. 2015. Autosubst: Reasoning with de Bruijn Terms and Parallel Substitutions. In *ITP (LNCS)*, Vol. 9236. 359–374.
- Ilya Sergey, Aleksandar Nanevski, and Anindya Banerjee. 2015. Mechanized verification of fine-grained concurrent programs. In *PLDI*. 77–87.
- Kristian Støvring and Søren Lassen. 2007. A Complete, Co-Inductive Syntactic Theory of Sequential Control and State. In *POPL*.
- Kasper Svendsen and Lars Birkedal. 2014. Impredicative Concurrent Abstract Predicates. In *ESOP*. 149–168.
- Amin Timany, Léo Stefanescu, Morten Krogh-Jespersen, and Lars Birkedal. 2018. A Logical Relation for Monadic Encapsulation of State: Proving contextual equivalences in the presence of runST. *Proc. ACM Program. Lang.* 2, POPL (Jan. 2018), to appear.
- Aaron Turon, Derek Dreyer, and Lars Birkedal. 2013a. Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency. In *ICFP*.
- Aaron Turon, Jacob Thamsborg, Amal Ahmed, Lars Birkedal, and Derek Dreyer. 2013b. Logical relations for fine-grained concurrency. In *POPL*.