

Logical Relations in Iris

Amin Timany

imec-Distrinet, KU Leuven, Belgium
amin.timany@cs.kuleuven.be

Robbert Krebbers

Aarhus University, Denmark
mail@robbertkrebbers.nl

Lars Birkedal

Aarhus University, Denmark
birkedal@cs.au.dk

Abstract

We present a formalization of logical relations for the language $F_{\mu, \text{ref}, \text{conc}}$: a call-by-value higher-order language with impredicative polymorphism, recursive types, general references, and concurrency. The logical relation interpretation is defined in Iris, a state-of-the-art higher-order concurrent separation logic, which in turn is formalized in Coq. The proof effort is made simpler by the use of the novel interactive proof mode for Iris, called IPM [?].

1. Introduction

It is well-known that it is challenging to define logical relations for higher-order programming languages with general references and concurrency [?, ?]. One of the main challenges is the so-called type-world circularity [?]. Here we side-step that challenge because we define the logical relation in Iris, a state-of-the-art higher-order concurrent separation logic. Iris is a logic of resources, with built-in support for defining guarded recursive predicates and invariants. We present both a unary and a binary logical relation interpretation of the types of $F_{\mu, \text{ref}, \text{conc}}$.

The Coq source code for our development can be found at <https://bitbucket.org/logsem/iris-logrel>. The material presented in this abstract is briefly described in [?] as a case study.

2. The language $F_{\mu, \text{ref}, \text{conc}}$

Iris is not tied to a fixed programming language, but can be instantiated with different programming languages. Here we instantiate it to $F_{\mu, \text{ref}, \text{conc}}$, which we formalize in Coq using the Autosubst library for De Bruijn terms [?]. De Bruijn terms make it easier to deal with substitution on open terms, which is needed for our proofs.

The terms and types of $F_{\mu, \text{ref}, \text{conc}}$ can be found in Figure ?? . Terms are untyped, so type-level abstraction is written as Λe and type application as $e _$. The operational semantics is split into two parts: thread-local head reduction \rightarrow_h and thread-pool reduction \rightarrow_{tp} , see Figure ?? . Both are defined using standard call-by-value evaluation contexts K , whose definition is omitted. Thread-pool reduction is defined on configurations $\equiv; (\vec{e}, \sigma)$ consisting of a state σ (a finite partial map from locations to values) and a thread-pool \vec{e} (a list of expressions corresponding to the threads). The thread-pool reduction is defined by interleaving, *i.e.*, by picking a thread and executing it, thread-locally, for one step. The only special case is `fork {e}`, which spawns a new thread e , and reduces itself to the unit value $()$.

Typing judgments take the form $\Xi \mid \Gamma \vdash e : \tau$, where Ξ is a context of type variables, and Γ is a context assigning types to program variables. The inference rules for the typing judgment are mostly standard and hence omitted.

3. Unary logical relation

The unary logical relation for $F_{\mu, \text{ref}, \text{conc}}$ is presented in Figure ?? . The logical relation is defined by two relations, indexed over types τ ,

$$\begin{aligned} e &::= x \mid \ell \mid \text{rec } f(x) = e \mid \Lambda e \mid \text{fold } e \mid \text{unfold } e \mid e \ e \\ &\quad \mid e _ \mid \text{fork } \{e\} \mid \text{ref}(e) \mid !e \mid e \leftarrow e \mid \text{CAS}(e, e, e) \\ v &::= n \mid \ell \mid \text{rec } f(x) = e \mid \Lambda e \mid \text{fold } v \\ \tau &::= X \mid \mathbb{N} \mid \tau \rightarrow \tau \mid \forall X. \tau \mid \mu X. \tau \mid \text{ref}(\tau) \end{aligned}$$

Figure 1. The syntax of $F_{\mu, \text{ref}, \text{conc}}$ (sums and products omitted).

Thread-local CBV head-reduction (omitted): $(e, \sigma) \rightarrow_h (e', \sigma')$

Thread-pool reduction: $(\vec{e}, \sigma) \rightarrow_{\text{tp}} (\vec{e}', \sigma')$

$$\frac{(e, \sigma) \rightarrow_h (e', \sigma')}{(\vec{e}_1 \ K[e] \ \vec{e}_2, \sigma) \rightarrow_{\text{tp}} (\vec{e}_1 \ K[e'] \ \vec{e}_2, \sigma')}$$

$$(\vec{e}_1 \ K[\text{fork } \{e\}] \ \vec{e}_2, \sigma) \rightarrow_{\text{tp}} (\vec{e}_1 \ K[()] \ \vec{e}_2 \ e, \sigma)$$

Figure 2. Operational semantics of $F_{\mu, \text{ref}, \text{conc}}$.

$$\llbracket X \rrbracket_{\Delta}(v) \triangleq \Delta(X)(v)$$

$$\llbracket \tau_1 \rightarrow \tau_2 \rrbracket_{\Delta}(v) \triangleq \square(\forall w. \llbracket \tau_1 \rrbracket_{\Delta}(w) \rightarrow \llbracket \tau_2 \rrbracket_{\Delta}^{\mathcal{E}}(w))$$

$$\llbracket \forall X. \tau \rrbracket_{\Delta}(v) \triangleq \forall f. \square \llbracket \tau \rrbracket_{\Delta[X \mapsto f]}^{\mathcal{E}}(v _)$$

$$\llbracket \mu X. \tau \rrbracket_{\Delta}(v) \triangleq \mu f. \exists w. v = \text{fold } w \wedge \triangleright \llbracket \tau \rrbracket_{\Delta[X \mapsto f]}(w)$$

$$\llbracket \text{ref}(\tau) \rrbracket_{\Delta}(v) \triangleq \exists \ell. v = \ell \wedge \boxed{\exists w. \ell \mapsto w * \llbracket \tau \rrbracket_{\Delta}(w)}^{\mathcal{N}. \ell}$$

$$\llbracket \tau \rrbracket_{\Delta}^{\mathcal{E}}(e) \triangleq \text{wp } e \{v. \llbracket \tau \rrbracket_{\Delta}(v)\}$$

Figure 3. The unary logical relation for $F_{\mu, \text{ref}, \text{conc}}$.

namely a value interpretation $\llbracket \tau \rrbracket_{\Delta} : \text{Val} \rightarrow i\text{Prop}$ and an expression interpretation $\llbracket \tau \rrbracket_{\Delta}^{\mathcal{E}} : \text{Expr} \rightarrow i\text{Prop}$. Note that these relations are Iris relations; $i\text{Prop}$ is the type of Iris propositions. Formally, these logical relations are defined on types τ in context Ξ , but we omit the Ξ here for notational simplicity. We furthermore omit product, sum and base types, but these are present in the Coq formalization. The Δ is a mapping from type variables to value interpretations, *i.e.*, $\Delta : \text{Tvar} \rightarrow \text{Val} \rightarrow i\text{Prop}$.

Experts on logical relations will recognize that this definition is extremely compact; this is because we express the relations in Iris. The case for function types expresses the usual requirement that a value is in the interpretation if it maps a value in the argument type to an expression in the result type. The \square modality (here and elsewhere) is used to ensure that the relation is persistent, which means that it does not depend on any resources. We use this modality, because Iris is a logic of resources, which means that its propositions ($i\text{Prop}$) generally express ownership of (ghost) resources. The logical relation must be persistent since typing is

intuitionistic (consider for example that the context Γ is copied in the usual typing rule for products). The definition for recursive types is given using a recursively defined predicate (the second μ is a fixpoint combinator); this is well-defined in Iris since the recursion variable occurs under the \triangleright guard¹.

We use the Iris invariant $\boxed{\exists w. \ell \mapsto w * \llbracket \tau \rrbracket_{\Delta}(w)}$ ^{$\mathcal{N}.\ell$} to express that a value is in the interpretation of a reference type $\text{ref}(\tau)$. The box is (on paper) Iris notation for invariants and the superscript $(\mathcal{N}.\ell)$ is the name of the invariant. Iris uses names to make sure invariants are not used more than once (which would be unsound). This semantics of the type $\text{ref}(\tau)$ states that a value of this type is a location ℓ and, invariantly, the location ℓ contains a value w in memory that is in the interpretation of τ . This use of Iris invariants dispels the need for explicit possible worlds and explicit treatment of the type-world circularity, which is otherwise typical for logical relation for reference types [?, ?].

Finally, the expression relation $\llbracket \tau \rrbracket_{\Delta}^{\mathcal{E}}$ uses Iris weakest preconditions to say that e is in the semantic interpretation of τ , if it is a computation whose possibly resulting value v is in the semantic interpretation of τ . Using the expression relation, we can define the semantic interpretation of types as:

$$\exists | \Gamma \models e : \tau \triangleq \forall \Delta \vec{v}. \left(\bigwedge_i \llbracket \sigma_i \rrbracket_{\Delta}(v_i) \right) \vdash \llbracket \tau \rrbracket_{\Delta}^{\mathcal{E}}(e[\vec{v}/\vec{x}])$$

where $\Gamma = x_1 : \sigma_1, \dots, x_n : \sigma_n$ and the environments $\Delta : Tvar \rightarrow Val \rightarrow iProp$ map into persistent interpretations.

For this logical relation, we can now prove:

1. The fundamental theorem of logical relation²:

Theorem `fundamental` $\Gamma \models e : \tau \vdash_t e : \tau \rightarrow \Gamma \models e : \tau$.

2. Type soundness, *i.e.*, that reduction of any well-typed expression can never get stuck:

Corollary `type_soundness` $e \tau e' \text{ thp } \sigma \sigma' :$
 $\square \vdash_t e : \tau \rightarrow \text{rtc step } ([e], \sigma) (e' :: \text{thp}, \sigma') \rightarrow$
 $\text{is_Some } (\text{to_val } e') \vee \text{reducible } e' \sigma'.$

The first result is proven in Iris using IPM. The latter result is formalized in plain Coq and relies on the fundamental theorem and the adequacy result for Iris, which formalizes that the weakest precondition predicate of Iris really is connected to the operational semantics of $F_{\mu, \text{ref}, \text{conc}}$ in the way you would expect.

Note that the corollary `type_soundness` shows the true power of using a proof assistant instead of a standalone tool: we can compose a proof in Iris with the adequacy result of Iris into a corollary that only mentions the typing judgment and the operational semantics. So, one no longer has to trust Iris or IPM!

4. Binary logical relation

We have also defined a binary logical relation for $F_{\mu, \text{ref}, \text{conc}}$ and proven that logical relatedness implies contextual approximation. It is not too hard to generalize the unary logical *value* interpretation to a binary relation, but to generalize the *expression* interpretation from the unary logical relation to the binary logical relation, one needs to find some way of expressing a *relation* between two expressions e

and e' using weakest precondition predicates, which are unary. This can be done as follows:

$$\llbracket \tau \rrbracket_{\Delta}^{\mathcal{E}}(e, e') \triangleq \forall j K. j \Vdash K[e'] \rightarrow \text{wp } e \{v. \exists w. j \Vdash K[w] * \llbracket \tau \rrbracket_{\Delta}(v, w)\}$$

Here $j \Vdash K[e']$ is a predicate on ghost state, which expresses that the specification side e' is in some evaluation context K for some thread j before we run e . In the post-condition, we have $j \Vdash K[w]$. Together with an appropriate invariant on ghost state, these predicates ensure that we really are relating the execution of e , which results in a value v to an execution of e' , which results in a value w , and that those values are related at the type τ .

Logical relatedness of e and e' is then defined as:

$$\exists | \Gamma \vdash e \leq_{\log} e' : \tau \triangleq \forall \vec{v} \vec{v}' \Delta . ;$$

$$\boxed{\text{heap}}^{\mathcal{N}} * \boxed{\Psi();}^{\mathcal{N}'} * \left(\bigwedge_i \llbracket \sigma_i \rrbracket_{\Delta}(v_i, v'_i) \right) \vdash \llbracket \tau \rrbracket_{\Delta}^{\mathcal{E}}(e[\vec{v}/\vec{x}], e[\vec{v}'/\vec{x}])$$

where $\Gamma = x_1 : \sigma_1, \dots, x_n : \sigma_n$ and the environments $\Delta : Tvar \rightarrow Val \rightarrow iProp$ map into persistent interpretations. Here $\boxed{\Psi();}^{\mathcal{N}'}$ is the invariant on ghost states mentioned above.

The fact that logical relatedness implies contextual approximation is shown by a series of congruence lemmas, corresponding to each of the typing rules (each on average about 10 lines of code). These lemmas are proved in Iris using IPM.

5. Proving logical refinements

We have used the binary logical relation to prove that two fine-grained concurrent implementations of modules contextually refines their coarse-grained counterparts. The first example is a counter module and the second is a stack module. The fine-grained implementations use optimistic concurrency and no locks, whereas the coarse-grained implementations use a spin lock (implemented using a CAS loop) to lock the data structure of the module before and after an operation is performed on the data structure.

For the counter, we use an Iris invariant relating the reference cells in the two implementations and the lock used in the coarse-grained implementation. The stack example comes from [?], where it was proved on paper using an invariant formulated using state transition system. State transition systems can be encoded in Iris [?], but in our experience, it is often easier to use direct monoid constructions when working in Coq.

¹ Formally, one has to show that the value interpretation of types is suitably *non-expansive*, which relies on all the Iris connectives being non-expansive. In the Iris formalization, non-expansiveness is encoded using Coq's setoid machinery [?], which makes it possible to hide most details about non-expansiveness from the end-user. Using this machinery, non-expansiveness of the logical relation is proved automatically.

² Contexts Ξ do not appear in the Coq code since we use De Bruijn indices.