A Relational Model of Types-and-Effects in Higher-Order Concurrent Separation Logic

Morten Krogh-Jespersen

Aarhus University, Denmark

mkj@cs.au.dk

Kasper Svendsen

University of Cambridge, UK

ks775@cl.cam.ac.uk

Lars Birkedal

Aarhus University, Denmark

birkedal@cs.au.dk

Abstract

Recently we have seen a renewed interest in programming languages that tame the complexity of state and concurrency through refined type systems with more fine-grained control over effects. In addition to simplifying reasoning and eliminating whole classes of bugs, statically tracking effects opens the door to advanced compiler optimizations.

In this paper we present a relational model of a type-and-effect system for a higher-order, concurrent programming language. The model precisely captures the semantic invariants expressed by the effect annotations. We demonstrate that these invariants are strong enough to prove advanced program transformations, including automatic parallelization of expressions with suitably disjoint effects. The model also supports refinement proofs between abstract data type implementations with different internal data representations, including proofs that fine-grained concurrent algorithms refine their coarse-grained counterparts. This is the first model for such an expressive language that supports both effect-based optimizations and data abstraction.

The logical relation is defined in Iris, a state-of-the-art higher-order concurrent separation logic. This greatly simplifies proving well-definedness of the logical relation and provides us with a powerful logic for reasoning in the model.

1. Introduction

Programming with and reasoning about effects in higher-order programs is well-known to be very challenging. Over the years, there have therefore been many proposals of refined type systems for taming and simplifying reasoning about effectful programs. Examples include alias types [28], capability type systems [23], linear type systems [14, 17, 20] Hoare type theory [21], permissions-based type systems [24], type-and-effect systems [5, 6, 15, 19], etc. Lately, we have also witnessed some larger-scale implementation efforts on higher-order programming languages, e.g., the Mezzo programming language [24] and the Rust programming language [27], which employ refined type systems to control the use of state in the presence of concurrency.

In this paper, we provide a logical account of an expressive region-based type-and-effect system for a higher-order concurrent programming language $\lambda_{ref,conc}$ with general references (higher-order store). The type-and-effect system is taken from [11]; it is inspired by Lucassen and Gifford's seminal work [15, 19], but also features a notion of public and private regions, which can be used to limit interference from threads running in parallel. Hence it can be used to express effect-based optimizations, as emphasized for type-and-effect systems for sequential languages by Benton *et al.*, see, *e.g.,* [5, 6]. Effect-based optimizations are examples of so-called "free theorems", *i.e.,* they just depend on the types and effects of the involved expressions, not on the particular expressions involved. The most interesting effect-based optimization is a parallelization theorem expressing the equivalence of running expressions e_1 and e_2 in parallel and running them sequentially, assuming their effects are suitably disjoint. Note that this is a relational property, i.e., the intended invariants of the type-and-effect system are relational in nature. Our logical account of the type-and-effect system thus consists of a logical relations interpretation of the types in a program logic, and we prove that logical relatedness implies contextual equivalence. We show that our logical relations interpretation is strong enough to prove the soundness of effect-based optimizations, in particular the challenging parallelization theorem.

Since the programming language $\lambda_{ref,conc}$ includes higher-order store, it is non-trivial to define a logical relations interpretation of the types, as one is faced with the well-known type-world circularity [1] (see [10] for an overview). Here we factor out this challenge, by using a state-of-the-art program logic, Iris [16], as the logic in which we express the logical relations. Iris has direct support for impredicative invariants, as needed for defining logical relations for general references. Iris also supports reasoning about concurrency; in particular, it supports a form of rely-guarantee reasoning about shared state. We use this facility to capture invariants of private and public regions. Moreover, we show, using simple synthetic examples, how we can also use the logic to prove that syntactically ill-typed programs obey the semantic invariants enforced by the type system. This is important in practice: both Mezzo and Rust contain facilities for programming with statically ill-typed expressions (Mezzo uses dynamic type checks [25] and Rust allows for including unsafe code in statically typed programs [27]) thus models of type-and-effect systems should preferably support reasoning about combinations of statically ill-typed and statically well-typed programs.

1.1 Overview of Challenges and Contributions

The typing judgments of our type-and-effect system take the form

$$\Pi \mid \Lambda \mid \Gamma \vdash e : \tau, \varepsilon$$

and express that the term e is of type τ and has effect ε in the typing context Γ mapping variables to types. The additional contexts Π and Λ consist of region variables ρ denoting, respectively, the public regions and the private regions that e may use. Intuitively, public regions are those that other threads may also use, whereas private regions are not subject to interference from other threads. Thus, from a thread-local perspective, the segregation describes an expression's expectations of interference from the environment. The effect ε is a finite set of read rd_ρ, write wr_ρ, or allocation effects, al_ρ, the intuition being that if, *e.g.*, $rd_\rho \in \varepsilon$, then e may read a reference belonging to region ρ.

Effect-based optimizations. Using effect annotations we can express the idea of parallelization mentioned above formally as follows (where rds ε is the set of regions with read effects in ε and likewise for wrs ε and als ε):

Theorem 1 (Parallelization). *If* $\Lambda = \Lambda_1, \Lambda_2, \Lambda_3$ *and*

1. $\Lambda_3 \mid \Lambda_1 \mid \Gamma_1 \vdash e_1 : \tau_1, \varepsilon_1$ *and* $\Lambda_3 \mid \Lambda_2 \mid \Gamma_2 \vdash e_2 : \tau_2, \varepsilon_2$
2. *als* $\varepsilon_1 \cup$ *wrs* $\varepsilon_1 \subseteq \Lambda_1$, *als* $\varepsilon_2 \cup$ *wrs* $\varepsilon_2 \subseteq \Lambda_2$
3. *rds* $\varepsilon_1 \subseteq \Lambda_1 \cup \Lambda_3$ *and* *rds* $\varepsilon_2 \subseteq \Lambda_2 \cup \Lambda_3$

then $\cdot \mid \Lambda \mid \Gamma_1, \Gamma_2 \vdash e_1 \,\|\, e_2 \cong_{ctx} (e_1, e_2) : \tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2$.

Here Λ_1 are the private regions of e_1, Λ_2 are the private regions of e_2, and Λ_3 are regions that can be used by both e_1 and e_2. The theorem then says, that if the expressions e_i only write and allocate in their private regions (item 2) and only read in private or shared regions (item 3), then, running e_1 in parallel with e_2 is contextually equivalent with running e_1 and e_2 sequentially, *if* the context is not allowed to access any locations used by the two expressions (expressed by the fact that $\Lambda_1, \Lambda_2, \Lambda_3$ are all private in the conclusion).

Intuitively, this theorem sounds very plausible, perhaps even quite obvious, but proving it formally was an open problem for more than 25 years [11] and it is still very difficult to prove for higher-order languages with general references, such as ours. Indeed, one of our key contributions is a novel proof technique for proving parallelization. To outline our approach, consider proving the left-to-right approximation of the parallelization theorem (Theorem 1). Then we, in particular, have to show that any reduction step taken by $e_1 \,\|\, e_2$ can also be taken by (e_1, e_2). In the case where the expression $e_1 \,\|\, e_2$ takes a step in the right branch, we cannot yet take the corresponding step in (e_1, e_2), unless e_1 has already reduced to a value. Previous methods for proving parallelization therefore relied on reordering steps taken in e_2 with steps from e_1, while preserving the semantic invariants - resulting in very difficult proofs [11] or trace-based arguments [9], which are not known to scale to programming languages with general references and dynamic allocation.

Our new technique is instead based on framing. We suspend the reduction on the right hand side temporarily, and first disentangle the reduction of $e_1 \,\|\, e_2$ into two semi-independent ("semi" because they can read from shared regions) reductions for e_1 and e_2 respectively, which can then be reassembled into a reduction for (e_1, e_2) using framing. The disentanglement and the reassembling qua framing, of course, depends on the effect annotations, and our formal argument leverages Iris' facility for capturing sophisticated ownership disciplines. We present a more detailed description and the formal argument in Section 3.4.

Data abstraction and local state. The $\lambda_{ref,conc}$ language supports hiding of local state using closures. Hiding can be used to imple-

ment abstract data types (ADTs) that manipulate an internal data representation, which can only be accessed through the provided operations. Relating ADT implementations that use different internal data representations is well-studied in the setting of ML-like type systems (see, *e.g.*, [2, 32] and the references therein); effect tracking adds several interesting dimensions.

In the ML setting the type system imposes no constraints on local state when relating ADT implementations. This is not the case in our setting. To illustrate, consider the following counter implemented using local state:

$$e_{\text{count}} \triangleq \textbf{let } x = \textbf{new } 0 \textbf{ in rec } inc().\textbf{let } y = {!x} \textbf{ in}$$
$$\textbf{if } \textbf{CAS}(x, y, y + 1) \textbf{ then } y \textbf{ else } inc()$$

e_{count} allocates a local reference x and returns a function that try to increment x inside a loop, until it succeeds, and returns the old value. To allow for concurrent access, the function uses a compare-and-set operation (**CAS**), which atomically sets the value of x to $y + 1$ if the value of x is equal to y and returns **true** or **false** depending on the result. The counter has the following type:

$$\rho \mid - \mid - \vdash e_{\text{count}} : 1 \to_{\{rd_\rho, wr_\rho\}}^{\rho, -} \textbf{int}, \{al_\rho\}$$

The type is a function type, which is annotated with a latent effect, expressing that the returned function may read and write in the public region ρ. To prove soundness of effect-based transformations, it is, of course, crucial that the semantic model also enforces the semantic invariants expressed by the effect annotations on local state. Otherwise, if our semantic model would allow us to forget about the effects on the local reference x, then we would be able to show, using a semantic version of Theorem 1, that **let** $g = e_{\text{count}}$ **in** $g() \,\|\, g()$ is contextually equivalent to **let** $g = e_{\text{count}}$ **in** $(g(), g())$, which is not the case (the first expression may evaluate to $(1, 0)$, while the second always evaluates to $(0, 1)$).

We can use the type-and-effect system to limit interference from the environment on the internal state of ADTs, when relating ADTs. For example, consider the two stack modules listed in $Figure$ 1. The left stack module, STACK_1, uses a single reference to a pure functional list whereas the right module, STACK_2, uses a linked list representation. Both stack implementations use a **CAS** operation to ensure that they function correctly in the presence of concurrent interference. The implementations (*i.e.*, $stack_1$ and $stack_2$) can be given the following type τ_{STACK}:

$$\textbf{1} \to_{\{al_\rho\}}^{\rho, -} (\textbf{int} \to_{\{wr_\rho, rd_\rho, al_\rho\}}^{\rho, -} \textbf{1}) \times (\textbf{1} \to_{\{wr_\rho, rd_\rho\}}^{\rho, -} \textbf{1} + \textbf{int})$$

This type expresses that each module will allocate in region ρ and return two functions $push$ and pop. The type further expresses that $push$ is allowed to have read, write and allocate effects on the local state described by ρ and that pop is allowed to read and write.

Intuitively, the two implementations are equivalent at this type, because their internal data representations are purely local and hidden from clients of the modules. Indeed, we can use our logical relation to prove:

Theorem 2. $\rho \mid - \mid - \vdash stack_1 \cong_{ctx} stack_2 : \tau_{\text{STACK}}, \{al_\rho\}$

Now, if we restrict possible interference from the environment by making the ρ region private, as expressed by the type τ'_{STACK} (ρ is now private on the latent effects, since it comes *after* the comma):

$$1 \to_{al_\rho}^{-, \rho} (\textbf{int} \to_{\{wr_\rho, rd_\rho, al_\rho\}}^{-, \rho} \textbf{1}) \times (\textbf{1} \to_{\{wr_\rho, rd_\rho\}}^{-, \rho} \textbf{1} + \textbf{int})$$

then the two implementations are still contextually equivalent at this type.

Moreover, for this type, we can also prove that we can safely omit the **CAS** operation from the stack implementations (intuitively, because there is no possible concurrent interference). Thus, writing $stack\_nc_i$ for the implementation of $stack_i$ without a **CAS**

$$stack_1() = \textbf{let } h = \textbf{new inj}_1 \; () \textbf{ in } (push_1(h), pop_1(h))$$
$$push_1(h) = \textbf{rec } loop(n).\textbf{let } v = !h \textbf{ in}$$
$$\textbf{if CAS}(h, v, \textbf{inj}_2 \; (n, v)) \textbf{ then } () \textbf{ else } loop(n)$$
$$pop_1(h) = \textbf{rec } loop(\_).\textbf{let } v = !h \textbf{ in case}(v, \textbf{inj}_1 \; () \Rightarrow \textbf{inj}_1 \; (),$$
$$\textbf{inj}_2 \; (n, v') \Rightarrow \textbf{if CAS}(h, v, v') \textbf{ then inj}_2 \; n \textbf{ else } loop())$$

$$stack_2() = \textbf{let } h = \textbf{new } (\textbf{new inj}_1 \; ()) \textbf{ in } (push_2(h), pop_2(h))$$
$$push_2(h) = \textbf{rec } loop(n).\textbf{let } v = !h \textbf{ in}$$
$$\textbf{if CAS}(h, v, \textbf{new inj}_2 \; (n, v)) \textbf{ then } () \textbf{ else } loop(n)$$
$$pop_2(h) = \textbf{rec } loop(\_).\textbf{let } v = !h \textbf{ in case}(!v, \textbf{inj}_1 \; () \Rightarrow \textbf{inj}_1 \; (),$$
$$\textbf{inj}_2 \; (n, v'') \Rightarrow \textbf{if CAS}(h, v, v'') \textbf{ then inj}_2 \; n \textbf{ else } loop())$$

Figure 1. The left stack module, Stack_1, has a single reference to a pure list where Stack_2 uses a reference to a linked list.

loop, we can use our logical relation to prove the following equivalences.

Theorem 3.
$$- \mid \rho \mid - \vdash stack\_nc_1 \cong_{ctx} stack_1 : \tau'_{\text{STACK}}, \{al_\rho\}$$
$$and \; - \mid \rho \mid - \vdash stack\_nc_2 \cong_{ctx} stack_2 : \tau'_{\text{STACK}}, \{al_\rho\}$$

Our proofs of data abstraction, detailed in 3.4 and [18], leverage Iris's facility for expressing invariants on local state. As pointed out in [11] the logical relation in *loc.cit.* could not be used to prove equivalences such as this one, since the logical relation there only allowed for much more restricted invariants.

Ill-typed terms. Here is a simple example of a statically ill-typed expression which nevertheless satisfies the semantic invariants enforced by the type system:

$$e \triangleq x := (); x := \textbf{true}$$

This expression first assigns the unit value to a boolean reference, and then assigns **true** to it.

This expression is not statically typable, due to the assignment of unit to a boolean reference. However, if the boolean reference is *private* then the rest of the program is not allowed to observe the ill-typed intermediate value and will thus never observe that the typing discipline has been broken. It is thus perfectly safe to use the untypable term e *as if* it had the following type:

$$- \mid \rho \mid x : \textbf{ref}_\rho \; \textbf{B} \vdash e : \textbf{1}, \{rd_\rho, wr_\rho\}$$

Our logical relations model allows us also to reason about such statically ill-typed terms and, *e.g.,* prove that e is equivalent to a statically well-typed expression which only assigns **true** to x.

Summary of Contributions In summary, the contributions of this paper are:

- We show how to interpret types of a region-based type-and-effect system for a concurrent higher-order imperative programming language with higher-order store as logical relations in the state-of-the-art program logic Iris.

- We use the interpretation to prove soundness of effect-based optimizations. In particular, we prove the soundness of the parallelization theorem. Our parallelization theorem is a strengthening of the one in [11] and for our proof we use a novel proof technique, based on framing. The resulting proof is arguably a lot clearer and more abstract than the one in [11], thanks to the use of the logical features of Iris.

- We use the interpretation to prove contextual equivalence of fine-grained concurrent data structures that use local state to hide internal data representations. Our examples could not be proved with the logical relation in [11].

- We show how the logic may be used to prove that syntactically ill-typed expressions obey the semantic properties enforced by the type system.

- We demonstrate that the logic allows us to give a modular definition of the logical relation and explain the relation by breaking it down into more manageable parts.

Outline We begin by formally defining the syntax and semantics of $\lambda_{ref,conc}$, the type-and-effect system, and contextual equivalence in Section 2. In Section 3 we turn our attention to logical relations for $\lambda_{ref,conc}$. We present our logical relation in four stages, starting from a unary relation that characterizes type inhabitance and ending with a binary relation for reasoning about contextual equivalence that supports advanced effect-based optimizations, each building on the previous relation. We conclude and discuss related and future work in Section 4.

2. $\lambda_{ref,conc}$ with Types, Regions and Effects

In this section we present the operational semantics and the type-and-effect system for $\lambda_{ref,conc}$, a call-by-value language with general references and concurrency primitives $\|$ and **CAS** (compare-and-set).

2.1 Syntax and Operational Semantics of $\lambda_{ref,conc}$

The syntax of $\lambda_{ref,conc}$ is shown in Figure 2 and the operational semantics is summarized in Figure 3. We assume given denumerably infinite sets of variables VAR, ranged over by x, y, f, and locations LOC, ranged over by l. We use v to range over the set of values, VAL, and e to range over the set of expressions, EXP. Note that expressions do not include types. We use **B**, **true**, **false** and **if** e **then** e_1 **else** e_2 as shorthands for booleans and branching encoded using sums.

$$v ::= () \mid n \mid (v, v) \mid \textbf{inj}_i \; v \mid \textbf{rec } f(x).e \mid x \mid l$$
$$e ::= v \mid e = e \mid e \; e \mid (e, e) \mid \textbf{prj}_i \; e \mid \textbf{inj}_i \; e \mid e + e \mid \textbf{new } e \mid !e$$
$$\mid e := e \mid \textbf{CAS}(e, e, e) \mid e \| e \mid \textbf{case}(e, \textbf{inj}_1 \; x \Rightarrow e, \textbf{inj}_2 \; y \Rightarrow e)$$

Figure 2. Syntax of $\lambda_{ref,conc}$.

The operational semantics is defined by a small-step relation between configurations consisting of a heap and an expression. Heaps h are finite partial maps from locations to values. The semantics is defined in terms of evaluation contexts, $K \in \text{ECTX}$. We use $K[e]$ to denote the expression obtained by plugging e into the context K and $e[v/x]$ to denote capture-avoiding substitution of value v for variable x in expression e.

2.2 Types and Effects for $\lambda_{ref,conc}$

The set of types is defined by the following grammar:

$$\text{TYPE } \tau ::= \textbf{1} \mid \textbf{int} \mid \textbf{ref}_\rho \; \tau \mid \tau \times \tau \mid \tau + \tau \mid \tau \to_\varepsilon^{\Pi, \Lambda} \tau$$

Π and Λ are finite sets of region variables, taken from a denumerably infinite set REGVAR ranged over by ρ. We use comma to denote disjoint union of sets of region variables. An atomic effect on a region ρ is either a read effect, rd_ρ, a write effect, wr_ρ, or an allocation effect, al_ρ. An effect ε is a finite set of atomic effects. Typing judgments take the form $\Pi \mid \Lambda \mid \Gamma \vdash e : \tau, \varepsilon$. An excerpt of the typing rules are shown in Figure 4. All typing rules can be found in [18].

$K ::= [\,] \mid K = e \mid v = K \mid K\ e \mid v\ K \mid (K, e) \mid (v, K) \mid \mathbf{prj}_i\ K$
$\quad \mid\ \mathbf{inj}_i\ K \mid \mathbf{case}(K, \mathbf{inj}_1\ x \Rightarrow e, \mathbf{inj}_2\ y \Rightarrow e) \mid \mathbf{new}\ K$
$\quad \mid\ !K \mid K := e \mid v := K \mid K\,||\,e \mid e\,||\,K \mid \mathbf{CAS}(K, e, e)$
$\quad \mid\ \mathbf{CAS}(v, K, e) \mid \mathbf{CAS}(v, v, K) \mid K + e \mid v + K$

Pure reduction $\qquad\qquad\qquad\qquad\qquad \boxed{e \overset{\text{pure}}{\rightarrow} e'}$

$$v_1\,||\,v_2 \overset{\text{pure}}{\rightarrow} (v_1, v_2)$$

Reduction $\qquad\qquad\qquad\qquad\qquad\qquad \boxed{h; e \rightarrow h'; e'}$

$$h; e \rightarrow h; e' \qquad\qquad \text{if } e \overset{\text{pure}}{\rightarrow} e'$$
$$h; \mathbf{new}\ v \rightarrow h \uplus [l \mapsto v]; l$$
$$h; !l \rightarrow h; v \qquad\qquad \text{if } h(l) = v$$
$$h[l \mapsto -]; l := v \rightarrow h[l \mapsto v]; ()$$
$$h; \mathbf{CAS}(l, v_o, v_n) \rightarrow h; \mathbf{false} \qquad \text{if } h(l) \neq v_o$$
$$h[l \mapsto v_o]; \mathbf{CAS}(l, v_o, v_n) \rightarrow h[l \mapsto v_n]; \mathbf{true}$$
$$h; K[e] \rightarrow h'; K[e'] \qquad \text{if } h; e \rightarrow h'; e'$$

Figure 3. Operational semantics of $\lambda_{ref,conc}$. Remaining pure reductions are standard (see [18]).

Regions can be introduced by the masking rule (TMASK). The masking rule expresses when we can introduce a new private region ρ for the evaluation of an expression e and hide all of e's effects on region ρ. The condition $\rho \notin FRV(\Gamma, \tau)$ ensures that we do not leak any locations of ρ and hence, from the perspective of e, region ρ is private. The masking rule has been used to do memory-management [30] and to hide local effects to enable more program-transformations [4, 29].

Since the masking rule allows us to hide local state effects, a pure operation is not necessarily deterministic in our setting. For instance, the following code-snippet which non-deterministically returns **true** or **false** can be typed as a pure expression:

$$- \vdash \mathbf{let}\ x = \mathbf{new}\ \mathbf{true}\ \mathbf{in}\ (x := \mathbf{true}\,||\,x := \mathbf{false}); !x : \mathbf{B}, \emptyset$$

2.3 Contextual Equivalence for $\lambda_{ref,conc}$

We take contextual equivalence as our basic notion of equivalence. Contextual equivalence relates two expressions if no suitably typed context can distinguish them. For a concurrent language such as $\lambda_{ref,conc}$ we have to choose whether there simply has to exist an indistinguishable reduction (may-equivalence) or whether all possible reductions must be indistinguishable (must-equivalence). In this paper we study may-equivalence and may-approximation, as defined below.

Definition 1. $\Pi \mid \Lambda \mid \Gamma \vdash e_1 \leq_{ctx} e_2 : \tau, \varepsilon$ *iff for all contexts C, values v, and heaps h_1 such that $C : (\Pi \mid \Lambda \mid \Gamma \vdash \tau, \varepsilon) \rightsquigarrow (- \mid - \mid - \vdash \mathbf{B}, \emptyset)$ and $[]; C[e_1] \rightarrow^* h_1; v$ there exists a heap h_2 such that $[]; C[e_2] \rightarrow^* h_2; v$.*

The $C : (\Pi \mid \Lambda \mid \Gamma \vdash \tau, \varepsilon) \rightsquigarrow (\Pi' \mid \Lambda' \mid \Gamma' \vdash \tau', \varepsilon')$ relation expresses that the context C takes a term e of the former type to a term of the latter type; the definition is standard and relegated to [18]. Note that e_1 and e_2 are not required to be well-typed in the definition above. Contextual equivalence $\Pi \mid \Lambda \mid \Gamma \vdash e_1 \cong_{ctx} e_2 : \tau, \varepsilon$ is then defined as contextual approximation in both directions.

3. A Logical Relation for $\lambda_{ref,conc}$

In this section we present a logical relation for $\lambda_{ref,conc}$. To aid exposition we present the logical relation in four steps. We start by defining a unary logical relation for a simplified type system without regions and effects. This allows us to focus on the use of

Iris as a meta-language for logical relations and provide a gentle introduction to Iris. We then extend the unary relation to the full type system with regions and effects, focusing on how effects are translated into abstract descriptions of possible interference. These unary logical relations characterize type inhabitance, which suffices for establishing type soundness, but not for proving equivalences. As the third step we naively extend the unary relation for the full type system to a binary relation, focusing on how to express a binary relation in a unary program logic. This yields a logical relation that is sound with respect to contextual approximation and suffices for proving equivalences of many concrete examples, but not advanced effect-based equivalences such as parallelization. For the fourth and final relation, we refine the third relation further, to support reasoning using multiple simulations. This final relation validates parallelization and is also sound with respect to contextual approximation.

This staged presentation also highlights the modularity of using Iris as a meta-language for logical relations. In particular, each step builds naturally on the previous, only requiring small changes or additions between each relation.

3.1 Unary Relation for $\lambda_{ref,conc}$ Without Effects

We begin by defining a unary logical relation for $\lambda_{ref,conc}$ with a standard ML-like type system without regions and effects. The goal is to define a unary relation, $\mathrm{LR}_{\mathrm{ML}}$ that characterizes type inhabitance semantically and is sound with respect to the syntactic typing rules. More precisely, we wish to define two unary relations, a value relation, $[\![\tau]\!]$, that characterizes values of type τ and an expression relation, $\mathcal{E}[\![\tau]\!]$, that characterizes expressions that either diverge or evaluate to values of type τ.

For ground types the definition of $[\![\tau]\!]$ is obvious: it is the values of the given type τ. The main difficulty arises when defining the interpretation of reference types. The idea is to take a location l to be an inhabitant of type **ref** τ if location l contains a value of type τ in the current heap. $\lambda_{ref,conc}$ is a concurrent language and the context is free to update the heap as it sees fit. However, the context must preserve typing and we can thus think of $[\![\mathbf{ref}\ \tau]\!](l)$ as expressing an *invariant* that l must always contain a value v of the semantic type $[\![\tau]\!]$ in the heap. To formalize this we introduce our meta-language, Iris.

Iris and invariants. Iris is a generic framework for constructing higher-order separation logics. For the purposes of this paper we present one particular instance of this framework for the $\lambda_{ref,conc}$ language and we refer to this instance simply as Iris.

Figure 5 contains an excerpt of the Iris syntax. Iris is a higher-order logic over a simply-typed term language. The set of Iris types, ranged over by κ, includes a type of $\lambda_{ref,conc}$ expressions **Exp** and values **Val**, a type of propositions, **Prop**, and is closed under products and function spaces. Iris includes the usual connectives $(\bot, \top, \wedge, \vee, \Rightarrow, \forall, \exists, *, -\!\!*, =_\kappa)$ and proof rules of higher-order separation logic. Iris extends this with a few new primitives, which we explain below.

Iris makes no distinction between assertions and specifications. Specifications are simply treated as special assertions that do not express ownership of any state. This is captured by the always modality, $\Box P$, which expresses that P holds and does not assert ownership of any state. Since $\Box P$ does not assert any ownership, it can be freely duplicated $(\Box P \implies \Box P * \Box P)$. We therefore call assertions of the form $\Box P$ *pure*.

One of the main features of Iris is *invariants* for reasoning about shared state. The pure assertion $\boxed{P}^\iota$ asserts the existence of an invariant with the name ι that owns a resource satisfying the assertion P. Resources owned by invariants are shared by every thread and can be accessed freely by atomic operations, provided the invariant is preserved. For atomic operations we can thus *open* an invariant

$$\frac{(x:\tau)\in\Gamma}{\Pi\mid\Lambda\mid\Gamma\vdash x:\tau,\emptyset}\qquad\overline{\Xi\vdash():\mathbf{1},\emptyset}\qquad\frac{\Xi\vdash e_1:\tau,\varepsilon_1\quad\Xi\vdash e_2:\tau,\varepsilon_2\quad eq_{type}(\tau)}{\Xi\vdash e_1=e_2:\mathbf{B},\varepsilon_1\cup\varepsilon_2}\qquad\frac{\Pi\mid\Lambda\mid\Gamma,f:\tau_1\rightarrow_\varepsilon^{\Pi,\Lambda}\tau_2,x:\tau_1\vdash e:\tau_2,\varepsilon}{\Pi\mid\Lambda\mid\Gamma\vdash\mathbf{rec}\ f(x).e:\tau_1\rightarrow_\varepsilon^{\Pi,\Lambda}\tau_2,\emptyset}$$

$$\frac{\Pi\mid\Lambda\mid\Gamma\vdash e_1:\tau_1\rightarrow_\varepsilon^{\Pi,\Lambda}\tau_2,\varepsilon_1\quad\Pi\mid\Lambda\mid\Gamma\vdash e_2:\tau_1,\varepsilon_2}{\Pi\mid\Lambda\mid\Gamma\vdash e_1\ e_2:\tau_2,\varepsilon\cup\varepsilon_1\cup\varepsilon_2}\qquad\frac{\Xi\vdash e:\mathbf{ref}_\rho\ \tau,\varepsilon}{\Xi\vdash!e:\tau,\varepsilon\cup\{rd_\rho\}}\qquad\frac{\Pi\mid\Lambda\mid\Gamma\vdash e:\tau,\varepsilon\quad\rho\in\Pi,\Lambda}{\Pi\mid\Lambda\mid\Gamma\vdash\mathbf{new}\ e:\mathbf{ref}_\rho\ \tau,\varepsilon\cup\{al_\rho\}}$$

$$\frac{\Xi\vdash e_1:\mathbf{ref}_\rho\ \tau,\varepsilon_1\quad\Xi\vdash e_2:\tau,\varepsilon_2}{\Xi\vdash e_1:=e_2:\mathbf{1},\varepsilon_1\cup\varepsilon_2\cup\{wr_\rho\}}\qquad\frac{\Pi\mid\Lambda,\rho\mid\Gamma\vdash e:\tau,\varepsilon\quad\rho\notin FRV(\Gamma,\tau)}{\Pi\mid\Lambda\mid\Gamma\vdash e:\tau,\varepsilon-\rho}\ \text{TMask}\qquad\overline{eq_{type}(\mathbf{1})}$$

$$\frac{\Pi,\Lambda_3\mid\Lambda_1\mid\Gamma_1\vdash e_1:\tau_1,\varepsilon_1\quad\Pi,\Lambda_3\mid\Lambda_2\mid\Gamma_2\vdash e_2:\tau_2,\varepsilon_2}{\Pi\mid\Lambda_1,\Lambda_2,\Lambda_3\mid\Gamma_1,\Gamma_2\vdash e_1\|e_2:\tau_1\times\tau_2,\varepsilon_1\cup\varepsilon_2}\qquad\frac{\Xi\vdash e_1:\mathbf{ref}_\rho\ \tau,\varepsilon_1\quad\Xi\vdash e_2:\tau,\varepsilon_2\quad\Xi\vdash e_3:\tau,\varepsilon_3\quad eq_{type}(\tau)}{\Xi\vdash\mathbf{CAS}(e_1,e_2,e_3):\mathbf{B},\varepsilon_1\cup\varepsilon_2\cup\varepsilon_3\cup\{wr_\rho,rd_\rho\}}$$

$$\frac{\Pi\mid\Lambda\mid\Gamma\vdash e:\tau_1,\varepsilon_1\quad\Pi,\Lambda\vdash\tau_1\le\tau_2\quad\varepsilon_1\subseteq\varepsilon_2}{\Pi\mid\Lambda\mid\Gamma\vdash e:\tau_2,\varepsilon_2}\qquad\frac{eq_{type}(\tau)\quad eq_{type}(\sigma)\quad op\in\{+,\times\}}{eq_{type}(\tau\ op\ \sigma)}$$

$$\frac{FRV(\tau)\in\Pi}{\Pi\vdash\tau\le\tau}\qquad\frac{\Pi\vdash\tau_1\le\tau_1'\quad\Pi\vdash\tau_2\le\tau_2'}{\Pi\vdash\tau_1\times\tau_2\le\tau_1'\times\tau_2'}\qquad\frac{\Pi\vdash\tau_1'\le\tau_1\quad\Pi\vdash\tau_2\le\tau_2'\quad\varepsilon_1\subseteq\varepsilon_2\quad\Pi_1\subseteq\Pi_2\quad\Lambda_1\subseteq\Lambda_2}{\Pi\vdash\tau_1\rightarrow_{\varepsilon_1}^{\Pi_1,\Lambda_1}\tau_2\le\tau_1'\rightarrow_{\varepsilon_2}^{\Pi_2,\Lambda_2}\tau_2'}$$

Figure 4. Excerpt of typing and sub-typing inference rules. We write $FV(e)$ and $FRV(e)$ for the sets of free program variables and region variables respectively. For all typing judgments $\Pi\mid\Lambda\mid\Gamma\vdash e:\tau,\varepsilon$ we implicitly assume that $FRV(\Gamma,\tau,\varepsilon)\in\Pi\cup\Lambda$. The equality type predicate, eq_{type}, defines the types we may test for equality. We use Ξ as shorthand for $\Pi\mid\Lambda\mid\Gamma$.

$$\kappa::=1\mid\mathbf{Exp}\mid\mathbf{Val}\mid\mathbf{Name}\mid\mathbf{Prop}\mid\mathbf{Monoid}\mid\kappa\times\kappa\mid\cdots$$
$$t,\varphi,\iota::=x\mid\lambda x:\kappa.\,t\mid\varphi\,t\mid(t,t)\mid\pi_1(t)\mid\pi_2(t)\mid t=_\kappa t$$
$$P,\mathcal{M}\quad\mid\quad t\mapsto t\mid\bot\mid\top\mid P\wedge P\mid P\vee P\mid P\Rightarrow P\mid$$
$$\mid\quad\forall x:\kappa.\,P\mid\exists x:\kappa.\,P\mid P\ast P\mid P\mathbin{-\!\ast}P\mid\Box P$$
$$\mid\quad\triangleright P\mid\boxed{P}^\iota\mid\{P\}\,e\,\{v.\,Q\}_\mathcal{M}\mid P\,{}^\mathcal{M}\!\!\Rrightarrow^\mathcal{M}Q\mid\cdots$$

Figure 5. Excerpt of Iris syntax.

and take local ownership of the resource owned by the invariant for the duration of the operation, provided we transfer back a resource that satisfies the invariant assertion after the operation. In Iris this is captured formally by *view-shifts*. A view-shift, written $P\Rrightarrow Q$ expresses that it is possible to transform a resource satisfying P into a resource satisfying Q, without changing the underlying physical state. To reason about opening of invariants, view-shifts are further annotated with *invariant masks* indicating which invariants are required to hold before and after the view-shift. In the view-shift $P\,{}^{\mathcal{M}_1}\!\!\Rrightarrow^{\mathcal{M}_2}Q$, $\mathcal{M}_1$ and $\mathcal{M}_2$ are invariant masks (sets of invariant names) required to hold before and after the view-shift respectively. The invariant masks ensure that we do not open an invariant twice (which would not be sound in general). Opening and closing of invariants is captured by the two following view-shift axioms:

$$\overline{\boxed{P}^\iota\,{}^{\{\iota\}}\!\!\Rrightarrow^\emptyset\triangleright P}\ \text{InvOpen}\qquad\overline{\boxed{P}^\iota\ast\triangleright P\,{}^\emptyset\!\!\Rrightarrow^{\{\iota\}}\top}\ \text{InvClose}$$

The INVOPEN rule allows us to take ownership of P upon opening the invariant ι, while the INVCLOSE rule requires us to relinquish ownership of P to close the invariant ι. In both rules, the resource P is guarded by a modality, $\triangleright$, which we explain shortly.

To apply these view-shifts to open an invariant for the duration of an atomic operation, such as reading ($!e$), writing ($x:=e$) or allocating (**new** e), Iris features the following atomic rule-of-consequence.

$$\frac{e\ \text{atomic}\quad P_1\,{}^{\mathcal{M}\uplus\mathcal{M}'}\!\!\Rrightarrow^\mathcal{M}P_2}{\{P_2\}\,e\,\{Q_2\}_\mathcal{M}\quad\forall v.\,Q_2(v)\,{}^\mathcal{M}\!\!\Rrightarrow^{\mathcal{M}\uplus\mathcal{M}'}Q_1(v)}{\{P_1\}\,e\,\{Q_1\}_{\mathcal{M}\uplus\mathcal{M}'}}\ \text{ACSQ}$$

Iris triples, $\{P\}\,e\,\{Q\}_\mathcal{M}$, are also annotated with an invariant mask, $\mathcal{M}$, indicating which invariants are required to hold be-

fore, during and after the execution of e. The atomic rule-of-consequence allows us to change this mask to open invariants for the duration of an atomic expression e. View-shifts include implication ($\Box(p\Rightarrow q)\vdash p\Rrightarrow q$) and we can thus recover the usual rule-of-consequence from ACSQ.

The "later" modality, $\triangleright$, is used to express that a property is only required to hold after one step of execution. It is used in connection with invariants because an Iris invariant may contain *any* predicate P, including one referring to the invariant itself. To ensure this is well-defined, Iris uses a form of guarded recursion, an abstract version of step-indexing — where $\triangleright$ is used to *guard* the resource in the invariant. Since the later modality expresses that a predicate holds after one step of execution, we can remove a $\triangleright$ modality from a precondition whenever our program makes an operational step. This is captured by the frame rule for atomic expressions:

$$\frac{\{P\}\,e\,\{Q\}\quad e\ \text{atomic}}{\{P\ast\triangleright R\}\,e\,\{v.\,Q(v)\ast R\}}\ \text{AFrame}$$

For many assertions it is also possible to remove laters without an operational step. We call these assertions *timeless* as they are independent of the number of steps left. For timeless assertions P we can view-shift away laters: $\triangleright P\Rrightarrow P$. Timeless assertions are closed under the connectives and quantifiers of first-order separation logic, but crucially does not include invariant assertions $\boxed{P}^\iota$, as the steps are precisely needed to model potentially self-referential invariants. With the exception of the reference invariants we use, all the invariants used throughout this article are timeless.

While the invariant names ι on invariant assertions, view shifts and Hoare triples are important for soundness, they are not important for understanding our encodings of logical relations in Iris. We have therefore chosen to elide all invariant names in the article, and refer interested readers to [18], where everything is fully annotated.
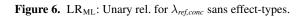
Logical relation. We now have enough logical machinery to define the first unary logical relation in Iris. The full definition of LR_{ML} is given in Figure 6.

The value relation, $[\![\tau]\!]$, is defined by induction on τ and defines an Iris assertion of type $\mathbf{Val}\to\mathbf{Prop}$. The expression relation, $\mathcal{E}$, is defined independently of the value relation and takes an arbitrary

$$\llbracket 1 \rrbracket \triangleq \lambda x.\ x = () \qquad\qquad \llbracket \mathbf{int} \rrbracket \triangleq \lambda x.\ x \in \mathbb{N}$$

$$\mathrm{REF}(x, \phi) \triangleq \exists v.\ x \mapsto v * \phi(v) \qquad \llbracket \mathbf{ref}\ \tau \rrbracket \triangleq \lambda x.\ \boxed{\mathrm{REF}(x, \llbracket \tau \rrbracket)}$$

$$\llbracket \tau_1 \rightarrow \tau_2 \rrbracket \triangleq \lambda x.\ \Box \forall y.\ (\triangleright \llbracket \tau_1 \rrbracket (y)) \Rightarrow \mathcal{E}(\llbracket \tau_2 \rrbracket)(x\ y)$$

$$\mathcal{E}(\phi)(e) \triangleq \{\top\}\ e\ \{v.\ \phi(v)\}_\top$$

Logical relatedness

$$\overline{x : \tau} \models_{\mathrm{ML}} e : \tau \triangleq\ \vdash_{\mathrm{IRIS}} \forall \overline{x'}.\ \overline{\llbracket \tau \rrbracket}(\overline{x'}) \implies \mathcal{E}(\llbracket \tau \rrbracket)(e[\overline{x'}/\overline{x}])$$

Figure 6. $\mathrm{LR}_{\mathrm{ML}}$: Unary rel. for $\lambda_{ref,conc}$ sans effect-types.

value predicate and extends it to expressions. It has the following type in Iris:

$$\mathcal{E} : (\mathbf{Val} \rightarrow \mathbf{Prop}) \rightarrow (\mathbf{Exp} \rightarrow \mathbf{Prop})$$

As already mentioned, for ground types, $\llbracket \tau \rrbracket$ is simply the set of values of the given type τ. The definition for arrow-types follows the usual idea of related arguments to related values, with the added wrinkle that we only require the argument to be related later. This suffices since applying a function takes a step in the operational semantics. The always modality in the value relation for arrow types is there to ensure the value relation is pure, which allows us to duplicate the resource that witnesses that a value is well-typed. It is needed in the arrow-case as implication does not preserve purity in general. For space reasons we omit the cases for products and sum types and refer the reader to [18].

Finally, for reference types, **ref** τ, the value relation is the set of locations l such that there exists an invariant that owns the location l and contains a value v in $\llbracket \tau \rrbracket$. Resources owned by invariants are shared, which allow all concurrently executing threads to freely update references, provided they respect the typing of the reference. This type of invariant can be seen as a particularly simple instance of rely / guarantee reasoning, where the rely and the guarantee are the same: namely, to preserve the invariant. A large part of the challenge throughout the rest of this article boils down to refining this invariant to limit possible interference from the environment, based on the region and effect system.

The expression relation $\mathcal{E}(\phi)$ extends a value predicate ϕ to expressions e by requiring that, if e terminates, then it terminates with a value satisfying ϕ. Finally, $\Gamma \models_{\mathrm{ML}} e : \tau$ extends this to open expressions, by closing under all substitutions. This semantic typing judgment is sound with respect to the usual typing rules, in the sense that for any well-typed term $\Gamma \vdash e : \tau$, the Iris assertion $\Gamma \models_{\mathrm{ML}} e : \tau$ is provable in Iris.

Lemma 1 (Soundness). *If* $\Gamma \vdash e : \tau$ *then* $\vdash_{\mathrm{IRIS}} \Gamma \models_{\mathrm{ML}} e : \tau$.

We note in passing that this logical relation shows the power of using Iris as a meta-language for defining logical relations: Usually, to define logical relations for a language with general references, one would need to index semantic types by worlds containing semantic types for allocated locations and the worlds and semantic types would have to be recursively defined [1, 10]; here this is all taken care of by Iris' built-in general logical facility for defining and working with invariants.

3.2 Unary Relation for $\lambda_{ref,conc}$ with Effects

In this section we extend the unary relation from the previous section to the full type system with regions and effects. Note that simply extending the relation from the previous section to the full type system by ignoring all region and effect annotations already yields a relation that is sound with respect to the full type system. However, this is needlessly conservative and by interpreting region

and effect annotations as restricting interference, we obtain a more precise semantic typing relation that is also sound.

The idea is to use the distinction between public and private regions to limit interference from the context, and the effect annotations to limit the effects of the given expression. We can encode this in Iris using tokens indexed by a region r corresponding to each type of effect: $[\mathrm{RD}]_r^\pi$, $[\mathrm{WR}]_r^\pi$, $[\mathrm{AL}]_r^\pi$. Each token is intended to grant permission to perform the corresponding effect on region r and, depending on the fractional permission $\pi \in \{\pi \in \mathbb{Q} \mid 0 < \pi \leq 1\}$, prevent the context from performing the given effect. These tokens must satisfy the following properties (and likewise for WR, AL):

$$[\mathrm{RD}]_r^1 * [\mathrm{RD}]_r^\pi \Rightarrow \bot \qquad [\mathrm{RD}]_r^{\pi_1 + \pi_2} \Leftrightarrow [\mathrm{RD}]_r^{\pi_1} * [\mathrm{RD}]_r^{\pi_2} \quad (1)$$

expressing that the full permission ($\pi = 1$) really means exclusive ownership of the token and that these tokens can be split and recombined arbitrarily. In Iris we can define such tokens using *ghost state*. Below we give a brief introduction to ghost state in Iris; for a more thorough treatment, we refer the reader to [16].

Iris and ghost state. Ghost resources provide a modular way of reasoning about *knowledge* and *rights* to modify some shared state. Ghost state is modeled using partial commutative monoids in Iris. Formally, these partial commutative monoids are presented as total commutative monoids with a distinguished zero element, $\bot$. The assertion $\overline{m : M}^\gamma$ asserts ownership of a ghost resource $m \in |M|$ of the monoid instance γ. Separating conjunction on ghost state is simply the lifting of the underlying monoid composition:

$$\overline{m_1 : M}^\gamma * \overline{m_2 : M}^\gamma \Leftrightarrow \overline{m_1 \cdot_M m_2}^\gamma \qquad (2)$$

The zero-element represents an ill-defined resource and thus cannot be owned: $\overline{\bot : M}^\gamma \Leftrightarrow \mathbf{false}$.

Tokens are a degenerate form of ghost state, consisting only of rights. The FRAC monoid, defined below, allows us to define an effect token for a single region. The carrier is rationals between 0 and 1, with addition as composition and 0 as the unit (we typically omit the explicit zero element from the definition of the monoid carrier and composition):

$$\mathrm{FRAC} = [0, 1] \cap \mathbb{Q} \qquad q \cdot q' = q + q', \quad \text{if } q + q' \leq 1$$

The idea is that 0 represents no ownership, 1 exclusive ownership and anything rational in the interval $(0, 1)$ non-exclusive ownership.

To define effect tokens for arbitrary regions, we also need the partial finite function monoid, $\mathrm{FPFUN}(X, M)$, with unit ϵ being the empty map and whose carrier is functions f from a set X into the non-zero elements of a monoid M, such that the set $\{x \in X \mid f(x) \neq \epsilon\}$ is finite. Composition on $\mathrm{FPFUN}(X, M)$ is defined point-wise, but is only defined if all point-wise compositions are well-defined:

$$(f \cdot g)(x) \triangleq f(x) \cdot g(x) \quad \text{if } f(x) \cdot g(x) \neq \bot \text{ for all } x \in X$$

Effect tokens can now be defined as follows and proven to satisfy property (1). The proof is an easy consequence of (2) and the definition of the monoid.

$$[X]_r^\pi \triangleq \overline{[r \mapsto \pi] : \mathrm{FPFUN}(\mathcal{RN}, \mathrm{FRAC})}^X, \ X \in \{\mathrm{RD}, \mathrm{WR}, \mathrm{AL}\}$$

Ghost state is a purely logical construct and is updated using view-shifts rather than assignments. To update a ghost resource we must ensure that our update is consistent with all ghost resources potentially owned by the environment. This is captured by the GHOSTUPD rule given below:

$$\frac{\text{GHOSTUPD}}{\forall m_f.\ m \cdot m_f \neq \bot \Rightarrow \exists m' \in M'.\ m' \cdot m_f \neq \bot}{\overline{m : M}^\gamma \Rightarrow \exists m' \in M'.\overline{m' : M}^\gamma}$$

To update a ghost resource m to some element $m' \in M'$, we have to show that doing so preserves all possible frames m_f composable with the resource m.

We can instantiate the finite partial functions monoid with locations and values to obtain the standard monoid of heaps used in separation logic. To define a monoid on values, we extend it with a unit element and a composition operator that is only defined if one of the two elements is unit.

$$\text{HEAP} \triangleq \text{FPFUN}(\text{LOC}, \text{VAL} + \{\epsilon\})$$

The ghost resource $\overline{[l \mapsto v] : \text{HEAP}}^{\gamma}$ asserts the exclusive right to modify location l in ghost heap γ and that location l currently contains the value v (here we use $[l \mapsto v]$ for the function that maps l to v and every other argument to ϵ). Using the GHOSTUPD rule, we can update ghost locations we own and allocate new ghost locations:

$$\overline{[l \mapsto v] : \text{HEAP}}^{\gamma} \Rrightarrow \overline{[l \mapsto v'] : \text{HEAP}}^{\gamma} \tag{3}$$

$$\overline{[] : \text{HEAP}}^{\gamma} \Rrightarrow \exists l. \overline{[l \mapsto v] : \text{HEAP}}^{\gamma} \tag{4}$$

Throughout the rest of the article we will need many ghost resources, including the HEAP monoid. We will introduce them by explaining the properties we expect them to satisfy. Naturally, we must define monoids for all of these resources and prove that the desired properties hold. All of these definitions and proofs can be found in [18].

Encoding effects using ghost state. Now that we have seen how to define and work with ghost state in Iris, we proceed with how to encode effects using ghost state.

A read-effect on a private region translates into exclusive ownership of the corresponding read token, while a read-effect on a public region translates into ownership of the corresponding read token with an arbitrary fractional permission π (and likewise for write and allocation effects).

The intended meaning of these tokens is enforced through the interplay between two invariants: a new *region* invariant, $\text{REG}(r)$, linking references with their corresponding region, and an updated reference invariant, $\text{REF}(r, \phi, x)$, indexed by a region identifier r and the reference's semantic type ϕ. Before we define these formally, we review some properties that should hold. If we own part of the read token for a region r then the context knows we might read references belonging to this region and must ensure that their values are well-typed. This is captured by the following property (where all free variables are universally quantified):

$$\boxed{\text{REG}(r)}, \boxed{\text{REF}(r, \phi, x)} \vdash \{[\text{RD}]_r^{\pi}\} \, !x \, \{y. \, [\text{RD}]_r^{\pi} * \phi(y)\} \tag{5}$$

Preservation of well-typedness is expressed by $\phi(y)$ in the postcondition. If we own part of the write token for a region r then we should be allowed to write any well-typed value to a reference belonging to region r:

$$\boxed{\text{REG}(r)}, \boxed{\text{REF}(r, \phi, x)}, \phi(v) \vdash \{[\text{WR}]_r^{\pi}\} x := v \, \{y. \, [\text{WR}]_r^{\pi}\} \tag{6}$$

Likewise, if we own any part of the allocation token for a region r we should be allowed to allocate a new reference and associate it with region r:

$$\boxed{\text{REG}(r)}, \phi(v) \vdash \{[\text{AL}]_r^{\pi}\} \, \textbf{new} \, v \, \{y. \, [\text{AL}]_r^{\pi} * \boxed{\text{REF}(r, \phi, y)}\}$$

Those three properties were fairly uneventful; the interesting properties deal with exclusive ownership of effect tokens.

Exclusive read effect. If we own the read token for region r exclusively, then the context cannot rely on references in region r containing well-typed values. If we additionally own a write token for region r, then we should be allowed to assign arbitrary values to references belonging to region r, provided we restore them with

well-typed values before returning the exclusive read token. To capture this formally, we introduce two new tokens, $[\text{RD}(x)]_r$ and $[\text{NORD}(x)]_r$, which express that if location x belongs to region r then it contains a well-typed value and may contain a value that is not well-typed, respectively. If we own the read token on a region r exclusively, then the following property allows us to exchange it for tokens that force all locations belonging to region r to contain well-typed values.

$$\boxed{\text{REG}(r)} \vdash [\text{RD}]_r^1 \iff \circledast_x [\text{RD}(x)]_r \tag{7}$$

By giving up the token that expresses that a location contains a well-typed value, we can assign an arbitrary value to the location. If we later assign a well-typed value, we can recover the token witnessing the well-typedness of the location. This is captured by the following two properties.

$$\boxed{\text{REG}(r)}, \boxed{\text{REF}(r, \phi, x)} \tag{8}$$
$$\vdash \{[\text{WR}]_r^{\pi} * [\text{RD}(x)]_r\} \, x := v \, \{y. \, [\text{WR}]_r^{\pi} * [\text{NORD}(x)]_r\}$$

$$\boxed{\text{REG}(r)}, \boxed{\text{REF}(r, \phi, x)}, \phi(v) \tag{9}$$
$$\vdash \{[\text{WR}]_r^{\pi} * [\text{NORD}(x)]_r\} \, x := v \, \{y. \, [\text{WR}]_r^{\pi} * [\text{RD}(x)]_r\}$$

Exclusive write effect. If we own the full write token, $[\text{WR}]_r^1$, then the context should not be allowed to modify references belonging to region r. Again, we capture this property by introducing new tokens $[\text{WR}(x)]_r$ and $x \xrightarrow{\pi}_r v$. Both tokens express that if location x belongs to region r, then we own the exclusive right to update it; the latter token further asserts that the current value is v. As before, we can trade ownership of a per-region write token for region r for all per-location write tokens for region r:

$$\boxed{\text{REG}(r)} \vdash [\text{WR}]_r^1 \iff \circledast_x [\text{WR}(x)]_r \tag{10}$$

Given ownership of a per-location write token for a location x belonging to region r, we can trade the token for a *points-to proxy* for x with fractional permission $\frac{1}{2}$:

$$\boxed{\text{REF}(r, \phi, x)} \vdash [\text{WR}(x)]_r \iff \exists v. \, x \xrightarrow{\frac{1}{2}}_r v \tag{11}$$

This points-to proxy satisfies similar properties as the standard separation logic points-to: If we own the points-to proxy $x \xrightarrow{\pi}_r v$ and read location x, we will read the value v:

$$\boxed{\text{REG}(r)} \vdash \{x \xrightarrow{\pi}_r v\} \, !x \, \{y. \, y = v * x \xrightarrow{\pi}_r v\} \tag{12}$$

If we own half of the points-to proxy for a location x we can also use it to assign a well-typed value to x:

$$\boxed{\text{REG}(r)}, \boxed{\text{REF}(r, \phi, x)}, \tag{13}$$
$$\phi(v_2) \vdash \{x \xrightarrow{\frac{1}{2}}_r v_1\} \, x := v_2 \, \{y. \, x \xrightarrow{\frac{1}{2}}_r v_2\}$$

Exclusive ownership of the per-region write token thus allows us to reason about the exact value of all references belonging to the region.

Exclusive allocation effect. Exclusive ownership of a per-region allocation token allows us to lock the domain of the heap associated with the given region. By trading our exclusive per-region allocation token, we can take ownership of a new token, $[\text{AL}(h)]_r^{\pi}$, that witnesses the domain of the heap associated with the given region:

$$\boxed{\text{REG}(r)} \vdash [\text{AL}]_r^1 \iff \exists h. \, [\text{AL}(h)]_r^{\frac{1}{2}} \tag{14}$$

As usual, we use fractional permissions to share the $[\text{AL}(h)]_r^{\pi}$ token. Given fractional ownership of two parts of the $[\text{AL}(h)]_r^{\pi}$ token, the domains of the two heaps must agree:

$$[\text{AL}(h_1)]_r^{\pi_1} * [\text{AL}(h_2)]_r^{\pi_2}$$
$$\Rightarrow [\text{AL}(h_1)]_r^{\pi_1} * [\text{AL}(h_2)]_r^{\pi_2} * \text{dom}(h_1) = \text{dom}(h_2)$$

New predicates

$$effs(r, \phi, x, v) \triangleq ([\text{Wr}(x)]_r \vee x \stackrel{\frac{1}{2}}{\hookrightarrow}_r v) *$$
$$([\text{Rd}(x)]_r \vee (\phi(v) * [\text{NoRd}(x)]_r))$$

$$\text{Reg}(r) \triangleq locs(r) * toks(r)$$

$$locs(r) \triangleq \exists h.\, rheap(h, r) * alloc(h, r) * \circledast_{(l,v) \in h} l \mapsto v$$
$$\circledast_{\{x | x \in (Loc \setminus \text{dom}(h))\}} [\text{NoRd}(x)]_r$$

$$toks(r) \triangleq ([\text{Rd}]_r^1 \vee \circledast_{x \in Loc} [\text{Rd}(x)]_r) *$$
$$([\text{Wr}]_r^1 \vee \circledast_{x \in Loc} [\text{Wr}(x)]_r)$$

$$alloc(h, r) \triangleq ([\text{Al}]_r^1 * [\text{Al}(h)]_r^{\frac{1}{2}}) \vee [\text{Al}(h)]_r^1$$

$$P_{toks}(\rho, r, \pi, \varepsilon) \triangleq (\rho \notin \text{rds }\varepsilon \vee [\text{Rd}]_r^\pi) * (\rho \notin \text{wrs }\varepsilon \vee [\text{Wr}]_r^\pi) *$$
$$(\rho \notin \text{als }\varepsilon \vee [\text{Al}]_r^\pi)$$

$$P_{reg}(R, g, \varepsilon, M) \triangleq \underset{\rho \in R}{\circledast}\, P_{toks}(\rho, M(\rho), g(\rho), \varepsilon) * \boxed{\text{Reg}(M(\rho))}$$

$$P^{\Pi, \Lambda}(g, \varepsilon, M) \triangleq P_{reg}(\Lambda, \mathbf{1}, \varepsilon, M) * P_{reg}(\Pi, g, \varepsilon, M)$$

Changes to previous definitions

$$\text{Ref}(r, \phi, x) \triangleq \exists v.\, x \stackrel{\frac{1}{2}}{\hookrightarrow}_r v * effs(r, \phi, x, v)$$

$$[\![\tau_1 \to_\varepsilon^{\Pi, \Lambda} \tau_2]\!]^M \triangleq \lambda x.\, \Box \forall y.\, (\triangleright y \in [\![\tau_1]\!]^M) \Rightarrow \mathcal{E}_{\varepsilon, M}^{\Pi, \Lambda}([\![\tau_2]\!]^M)(x\, y)$$

$$[\![\mathbf{ref}_\rho\, \tau]\!]^M \triangleq \lambda x.\, \boxed{\text{Ref}(M(\rho), [\![\tau]\!]^M, x)}$$

$$\mathcal{E}_{\varepsilon, M}^{\Pi, \Lambda}(\phi)(e) \triangleq \forall g.$$
$$\{P^{\Pi, \Lambda}(g, \varepsilon, M)\}\, e\, \{v.\, \phi(v) * P^{\Pi, \Lambda}(g, \varepsilon, M)\}_\top$$

Logical relatedness

$$\Pi \mid \Lambda \mid \overline{x : \tau} \models_{\text{Eff}} e : \tau, \varepsilon \triangleq$$
$$\vdash_{\text{Iris}} \forall M. \forall \overline{x'}.\, \overline{[\![\tau]\!]^M(x')} \implies \mathcal{E}_{\varepsilon, M}^{\Pi, \Lambda}([\![\tau]\!]^M)(e[\overline{x'}/\overline{x}])$$

Figure 7. LR_{Eff}: Unary rel. for $\lambda_{ref,conc}$ with effect-types.

Exclusive ownership ($\pi = 1$) is required to update the domain of the heap. It is possible to update the heap without exclusive access, as long as the domain is preserved:

$$[\text{Al}(h)]_r^1 \Rightarrow [\text{Al}(h')]_r^1$$
$$[\text{Al}(h)]_r^\pi * \text{dom}(h) = \text{dom}(h') \Rightarrow [\text{Al}(h')]_r^\pi$$

Logical relation. The LR_{Eff} logical relation, including the new region invariant and updated reference invariant is defined in Figure 7. Changes to existing predicates are in green. The value relation, $[\![\tau]\!]^M$, is now indexed by an injective mapping M from region variables to Iris invariant names. This mapping allows us to model that the same region variable might refer to different regions in the case where a region ρ is created after hiding a region with the same name ρ. Likewise, the expression relation, $\mathcal{E}_{\varepsilon, M}^{\Pi, \Lambda}(\phi)$, is also indexed by the region mapping M, in addition to the region contexts Π, Λ and the effect typing ε.

To explain the relation, let us start with the reference invariant $\text{Ref}(r, \phi, x)$. Note first that the reference invariant no longer owns the underlying physical location (i.e., $x \mapsto v$). Instead it owns a proxy $x \stackrel{\frac{1}{2}}{\hookrightarrow}_r v$. The *effs* predicate encodes the meaning of the per-location read and write tokens. It allows us to exchange a write token $[\text{Wr}(x)]_r$ for a proxy that describes the current value of the location (property (11)) and track when the location contains a well-typed value (properties (8) and (9)).

The region invariant $\text{Reg}(r)$ consists of two resources, a token resource $toks(r)$ that ties all the per-region tokens together with the per-location tokens, and the $locs(r)$ resource that ties together the points-to proxies with the physical state. The $toks(r)$ resource allows us to exchange an exclusive per-region read or write token for all the corresponding per-location read or write tokens (properties (7) and (10)). It also enforces that if we only own a fraction of the per-region read or write token then the region invariant must own all per-location read and write tokens for the given region. This ensures that the location must contain a well-typed value and that we are allowed to update it, respectively (properties (5) and (6)).

The local points-to proxies for a region r are tied to the physical state using the global counter-part $rheap(h, r)$ resource in $locs(r)$. The local points-to proxy always agrees with the global heap proxy:

$$rheap(h, r) * x \stackrel{\pi}{\hookrightarrow}_r v \Rightarrow rheap(h, r) * x \stackrel{\pi}{\hookrightarrow}_r v * h(x) = v$$

To update a points-to proxy thus requires ownership of both the corresponding global heap proxy and exclusive ownership of the local points-to proxy:

$$rheap(h, r) * x \stackrel{1}{\hookrightarrow}_r v \Rightarrow rheap(h[x \mapsto v'], r) * x \stackrel{1}{\hookrightarrow}_r v'$$

The $locs(r)$ resource asserts ownership of physical points-to resources for each location and value in the global points-to proxy for region r. Since these are tied together with the local points-to proxies, the local points-to proxies must agree with the underlying physical state (properties (12) and (13)). $locs(r)$ also asserts ownership of all the $[\text{NoRd}(x)]_r$ resources for locations x not belonging to the region.

The reason for introducing the indirection of proxies, is to allow reasoning about the set of locations belonging to a region, to interpret allocation effects. This is captured by the $alloc(h, r)$ resource, which allows clients to trade the exclusive allocation token for a lock on the set of locations belonging to the region (property (14)).

Finally, P_{reg} specifies how the effect annotation translates into ownership of the corresponding effect tokens. Effects in private regions yield exclusive ownership, while effects in public regions yield non-exclusive ownership.

Example: type violating update. To illustrate how we can use this stronger semantic typing judgment to semantically type-check code that is not syntactically well-typed, recall the previous mentioned type-violating example.

$$- \mid \rho \mid x : \mathbf{ref}_\rho\, \mathbf{B} \models_{\text{Eff}} x := (); x := \mathbf{true} : \mathbf{1}, \{rd_\rho, wr_\rho\}$$

The read and write effect on the private region ρ translates into exclusive ownership of the read and write token. Using properties (7), (8) and (9) we can thus easily verify that the example is semantically well-typed.

Context: $\boxed{\text{Reg}(r)}$, $\boxed{\text{Ref}(r, [\![\mathbf{B}]\!]^M, x)}$
$$\{[\text{Wr}]_r^1 * [\text{Rd}]_r^1\} \Rrightarrow \{[\text{Wr}]_r^1 * \circledast_{y \in Loc} [\text{Rd}(y)]_r\}$$
$$x := ()$$
$$\{[\text{Wr}]_r^1 * [\text{NoRd}(x)]_r * \circledast_{y \in Loc \setminus \{x\}} [\text{Rd}(y)]_r\}$$
$$x := \mathbf{true}$$
$$\{[\text{Wr}]_r^1 * \circledast_{y \in Loc} [\text{Rd}(y)]_r\} \Rrightarrow \{[\text{Wr}]_r^1 * [\text{Rd}]_r^1\}$$

3.3 Binary Relation for $\lambda_{ref,conc}$ with Effects

Previously we looked at unary relations for semantically characterizing type inhabitance. Now, we switch to binary relations intended to imply contextual approximation.

We define two families of binary relations, $[\![\tau]\!]^M$ and $\mathcal{E}([\![\tau]\!]^M)$, that characterize contextual approximation on values and expressions of type τ, respectively. Generalizing the value relation to contextual approximation is fairly straightforward: on ground types it

New predicates

$$\text{SPEC}(h_0, e_0) \triangleq \exists h, e.\, heap_S(h) * mctx(e) * (h_0; e_0 \to^* h; e)$$

Changes to previous definitions

$$\text{REF}(r, \phi, x) \triangleq \exists v.\, x_I \overset{\frac{1}{2}}{\hookrightarrow}_{I,r} v_I * x_S \overset{\frac{1}{2}}{\hookrightarrow}_{S,r} v_S * \textit{effs}(r, \phi, x, v)$$

$$\textit{effs}(r, \phi, x, v) \triangleq ([\text{WR}(x)]_r \vee (x_I \overset{\frac{1}{2}}{\hookrightarrow}_{I,r} v_I * x_S \overset{\frac{1}{2}}{\hookrightarrow}_{S,r} v_S)) *$$
$$([\text{RD}(x)]_r \vee (\phi(v_I, v_S) * [\text{NORD}(x)]_r))$$

$$\textit{locs}(r) \triangleq \exists h.\, rheap_I(h_I, r) * rheap_S(h_S, r) * alloc(h, r)$$
$$\circledast_{(l,v)\in h_I}\, l \mapsto_I v * \circledast_{(l,v)\in h_S} l \mapsto_S v$$
$$\circledast_{x\in Loc\backslash\text{dom}(h_I) \times (Loc\backslash\text{dom}(h_S))} [\text{NORD}(x)]_r$$

$$\textit{tokens}(r) \triangleq ([\text{WR}]^1_r \vee \circledast_{x\in Loc^2}[\text{WR}(x)]_r) *$$
$$([\text{RD}]^1_r \vee \circledast_{x\in Loc^2}[\text{RD}(x)]_r)$$

$$\textit{alloc}(h, r) \triangleq ([\text{AL}]^1_r * [\text{AL}(h_I, h_S)]^{\frac{1}{2}}_r) \vee [\text{AL}(h_I, h_S)]^1_r$$

$$[\![\mathbf{1}]\!]^M \triangleq \lambda x.\, x_I = x_S = ()$$

$$[\![\mathbf{int}]\!]^M \triangleq \lambda x.\, x_I, x_S \in \mathbb{N} \wedge x_I = x_S$$

$$[\![\tau_1 \to^{\Pi,\Lambda}_\varepsilon \tau_2]\!]^M \triangleq \lambda x.\, \Box \forall y.\, (\triangleright[\![\tau_1]\!]^M)(y_I, y_S) \Rightarrow$$
$$\mathcal{E}^{\Pi,\Lambda}_{\varepsilon,M}([\![\tau_2]\!]^M)(x_I\, y_I, x_S\, y_S)$$

$$[\![\mathbf{ref}_\rho\, \tau]\!]^M \triangleq \lambda x.\, \boxed{\text{REF}(M(\rho), [\![\tau]\!]^M, x)}$$

$$\mathcal{E}^{\Pi,\Lambda}_{\varepsilon,M}(\phi)(e_I, e_S) \triangleq \forall g, j, h_0, e_0.\, \boxed{\text{SPEC}(h_0, e_0)} \vdash$$
$$\{j \Rrightarrow_S e_S * P^{\Pi,\Lambda}(g, \varepsilon, M)\}\, e_I$$
$$\{v_I.\, \exists v_S.\, j \Rrightarrow_S v_S * \phi(v_I, v_S) * P^{\Pi,\Lambda}(g, \varepsilon, M)\}_\top$$

Logical relatedness

$$\Pi \mid \Lambda \mid \overline{x:\tau} \models_{\text{BIN}} e_1 \leq_{log} e_2 : \tau, \varepsilon \triangleq$$
$$\vdash_{\text{IRIS}} \forall M.\, \forall \overline{x}.\, \overline{[\![\tau]\!]^M(\overline{x})} \Rightarrow \mathcal{E}^{\Pi,\Lambda}_{\varepsilon,M}([\![\tau]\!]^M)(e_1[\overline{x_I/x}], e_2[\overline{x_S/x}])$$

Figure 8. LR$_{\text{BIN}}$: Binary rel. for $\lambda_{ref,conc}$ with effect-types.

---

is simply the identity relation on the values of the given type; for arrow types it relates functions that map related arguments to related expressions, and for reference types it relates two locations if they contain related values.

The expression relation is more interesting: intuitively it should express that e_I approximates e_S, if any step that e_I can make can be simulated by zero or more steps of e_S. We think of e_I as an "implementation" and of e_S as a "specification". We follow the approach of Turon *et al.* [31] and capture this relational property as a unary Hoare triple on e_I by requiring the triple to update ghost resources that force the execution of e_S. The idea is to introduce a ghost resource $j \Rrightarrow_S e$ that expresses that the expression e is in an evaluation context on the "specification" side and the exclusive right to reduce this expression. With this ghost resource we can express a simulation between an implementation e_I and a specification e_S as follows:

$$e_I \leq e_S \approx \{j \Rrightarrow_S e_S\}\, e_I\, \{v_I.\, \exists v_S.\, j \Rrightarrow_S v_S * \phi(v_I, v_S)\}$$

By requiring e_I to update the ghost resource from e_S to a value v_S, we are forced to show that we can reduce e_S, which appears in an evaluation context on the specification side, to the value v_S. We refer to $j \Rrightarrow_S e$ as a local expression resource, as it allows us to reason locally about reductions on sub-expressions of the full specification program.

The generalization to expressions in evaluation contexts is necessary to prove that the relation is a congruence. In particular, to prove the following congruence property:

$$e_{1I} \leq e_{1S} \wedge e_{2I} \leq e_{2S} \Rightarrow e_{1I}||e_{2I} \leq e_{1S}||e_{2S}$$

We need to be able to split the local expression resource $j \Rrightarrow_S e_{1S}||e_{2S}$ into two separate resources, one for e_{1S} and another for e_{2S}. Then we can pass one to e_{1I} and the other to e_{2I} and they can each reduce their corresponding expression on the right, independently of the other. This is possible because e_{1S} and e_{2S} both occur in evaluation contexts. This is also the reason why local expression resources $j \Rrightarrow_S e$ are indexed. The j serves as a logical "thread identifier", allowing us to distinguish different local expression resources.

Specification resources. To formalize this idea, we need a number of ghost resources. In addition to the local expression resource, $j \Rrightarrow_S e$ we also need a global expression resource, $mctx(e)$, for reasoning about the full specification program. Naturally, the global expression resource and all local expression resources must agree on the specification program, so splitting a local expression resource requires ownership of both. The following lemma allows us to introduce and eliminate a local expression resource for an expression that occurs in an evaluation context inside another local expression resources:

$$j \Rrightarrow_S \kappa[e_1] * mctx(e) \iff \exists i.\, j \Rrightarrow_S \kappa[i] * i \Rrightarrow_S e_1 * mctx(e) \tag{15}$$

This is achieved by introducing a new logical thread identifier i for the new local expression resource for e_1 and replacing e_1 with i in the original local expression resource. Here κ is an evaluation context extended to expressions that may contain logical thread identifiers. By applying the above property twice, we can split a local expression resource for a parallel composition into two:

$$j \Rrightarrow_S e_1||e_2 * mctx(e) \iff \exists i_1, i_2.$$
$$j \Rrightarrow_S i_1||i_2 * i_1 \Rrightarrow_S e_1 * i_2 \Rrightarrow_S e_2 * mctx(e)$$

Since all local specification expressions are in an evaluation context of the global specification expression, any reduction of a local specification expression can be extended to the global specification expression:

$$j \Rrightarrow_S e_1 * mctx(e) * (h; e_1 \to h'; e_1') \Rightarrow \tag{16}$$
$$\exists e'.\, j \Rrightarrow_S e_1' * mctx(e') * (h; e \to h'; e')$$

In the case where there exists just one local expression resource that contains no free logical thread identifiers, then the local expression should agree with the global expression. To formalize this, we treat the thread identifier 0 as the "root" local expression:

$$0 \Rrightarrow_S e_1 * mctx(e_2) * FA(e_1) = \emptyset \Rightarrow \tag{17}$$
$$0 \Rrightarrow_S e_1 * mctx(e_2) * e_1 = e_2$$

where $FA(e)$ is the set of free logical thread identifiers in e. These local and global expression resources are definable in Iris and we refer the reader to [18] for detailed definitions.

We need another two ghost resources, $heap_S(h)$ and $l \mapsto_S v$, for reasoning about specification heaps. This is in fact the HEAP monoid we have seen before, with some additional structure. The $heap_S(h)$ resource asserts global ownership of the full specification heap h, while $l \mapsto_S v$ asserts local ownership of a single location l, respectively. We require that the global heap agrees with the local heap resources:

$$heap_S(h) * l \mapsto_S v \Rightarrow heap_S(h) * l \mapsto_S v * h(l) = v \tag{18}$$

Updating a location l requires both local ownership of l and the global heap resource and allocation requires ownership of the

global heap resource, both lifted from updating and allocating ghost locations seen before in (3) and (4).

$$heap_S(h) * l \mapsto_S v \Rightarrow heap_S(h[l \mapsto v']) * l \mapsto_S v' \quad (19)$$

$$heap_S(h) * l \notin \mathrm{dom}(h) \Rightarrow heap_S(h[l \mapsto v]) * l \mapsto_S v \quad (20)$$

With these resources in hand, we can now formally define a simulation as a Hoare triple. We define a specification invariant that asserts ownership of the global specification heap and expression. Additionally, it also requires that there exists a reduction from some initial configuration $h_0; e_0$ to the current global specification heap and expression:

$$\mathrm{SPEC}(h_0, e_0) \triangleq \exists h, e.\, heap_S(h) * mctx(e) * (h_0; e_0 \to^* h; e)$$

By requiring the Hoare triple to update the local expression e_S of a to a value v_S, we thus force it to show the existence of a reduction:

$$\boxed{\mathrm{SPEC}(h_0, e_0)} \vdash \{j \Rightarrow_S e_S\}\, e_I\, \{v_I.\, \exists v_S.\, j \Rightarrow_S v_S\}$$

The only way to update the local expression resource $j \Rightarrow_S e$ is through property (16) which also requires opening and reestablishing the specification invariant to gain access to the global expression resource.

The logical relation. Now that we have seen how we can express relational properties as unary Hoare triples, we just need to integrate this idea with the techniques from the previous section for translating region and effect annotations into specifications of abstract interference.

Consider the reference invariant, $\mathrm{REF}(r, \phi, x)$. In the unary setting it asserts ownership of a proxy for the underlying heap location that, depending on ownership of the per-location read and write tokens, contains a well-typed value and may be updated. In the binary setting, x is now a pair of locations (x_I, x_S) and the invariant asserts ownership of proxies for both the implementation and specification side heaps, but otherwise the structure of the definition remains the same. The binary reference invariant is defined in Figure 8. We use x_I and x_S as shorthand for the first and second projection of a pair x. Note that, in the binary setting, per-location read, write and no-read tokens are now indexed by a pair of locations, rather than just a single location. The $[\mathrm{RD}(x_I, x_S)]_r$ token now expresses that if locations x_I and x_S are related and belong to region r, then they contain related values, and likewise for the other tokens.

The $\mathrm{LR_{BIN}}$ logical relation satisfies the fundamental theorem of logical relations (Theorem 4), which expresses that all well-typed terms are related to themselves. It is also sound with respect to contextual approximation (Theorem 5).

Theorem 4 (Fundamental Theorem). *If* $\Pi \mid \Delta \mid \Gamma \vdash e : \tau, \varepsilon$ *then* $\Pi \mid \Delta \mid \Gamma \models_{\mathrm{BIN}} e \leq_{log} e : \tau, \varepsilon$

Theorem 5 (Soundness). *If* $\Pi \mid \Delta \mid \Gamma \models_{\mathrm{BIN}} e_I \leq_{log} e_S : \tau, \varepsilon$ *then* $\Pi \mid \Delta \mid \Gamma \vdash e_I \leq_{ctx} e_S : \tau, \varepsilon.$

3.4 Binary Relation for $\lambda_{ref,conc}$ with Effects Using Multiple Simulations

The $\mathrm{LR_{BIN}}$ relation supports proofs of contextual approximations by showing that each step on the left can be simulated on the right. However, it requires that each thread on the left owns the local expression resource of the thread on the right that simulates the thread on the left. This is too restrictive in cases where multiple threads on the left are simulated by a single thread on the right, such as the case of parallelization. In this section we introduce our final logical relation, $\mathrm{LR_{PAR}}$, that removes this restriction.

The idea is simple: The $\mathrm{LR_{BIN}}$ relation allowed us to reason about a single simulation; now, we generalize the relation to allow reasoning about multiple simulations, such that multiple threads on the left can be given ownership of an expression resource for the same thread on the right, in different simulations.

Multiple simulations. To make this precise, we generalize the existing specification ghost resources, so that we can have multiple independent copies, by indexing the global and local expression resources ($mctx(e, \zeta)$ and $j \stackrel{\zeta}{\Rightarrow}_S e$) and heap resources ($heap_S(h, \zeta)$ and $l \mapsto_S^\zeta v$) with a simulation identifier ζ. For each simulation identifier ζ, the resources $mctx(e, \zeta)$ and $j \stackrel{\zeta}{\Rightarrow}_S e$ satisfy the same properties as before (properties (15) to (17)) and likewise for the heap resources (properties (18) and (19)). We can allocate new expression and heap resources initialized with an arbitrary expression e and an empty heap:

$$\top \Rightarrow \exists \zeta.\, mctx(e, \zeta) * heap_S([], \zeta) * 0 \stackrel{\zeta}{\Rightarrow}_S e \quad (21)$$

The idea is to relate e_I and e_S if e_S can simulate any step performed by e_I in an arbitrary simulation ζ in which e_S appears in an evaluation context:

$$\forall \zeta, i, h_0, e_0.\, \boxed{\exists h, e.\, heap_S(h, \zeta) * mctx(e, \zeta) * (h_0; e_0 \to^* h; e)}$$
$$\vdash \{i \stackrel{\zeta}{\Rightarrow}_S e_S\}\, e_I\, \{v_I.\, \exists v_S.\, i \stackrel{\zeta}{\Rightarrow}_S v_S\}$$

This allows the caller of e_I to choose in which simulation ζ the specification e_S must simulate e_I. It also allows e_I to simulate sub-expressions of e_I in different simulations than ζ, provided it can still prove a simulation in ζ at the end.

We can allocate a new simulation with an arbitrary initial configuration $h; e$ and take ownership of the local heap and expression resources for this simulation, using (21) and (20).

This ability to simulate sub-expressions in different simulations is exactly what we need to disentangle an execution of $e_1 \,\|\, e_2$ into two independent executions of e_1 and e_2, when proving parallelization. To show that $e_1 \,\|\, e_2$ is related to (e_1, e_2) in the expression relation, we (roughly) prove the following triple:

$$\boxed{\exists h, e.\, heap_S(h, \zeta) * mctx(e, \zeta) * (h_0; e_0 \to^* h; e)} \vdash$$
$$\{i \stackrel{\zeta}{\Rightarrow}_S (e_1, e_2) * \cdots\}\, e_1 \,\|\, e_2\, \{v_I.\, \exists v_S.\, i \stackrel{\zeta}{\Rightarrow}_S v_S * \cdots\}$$

Recall from the Introduction that the idea is to use the effect annotations to prove that an execution of $e_1 \,\|\, e_2$ can be disentangled into semi-independent executions of e_1 and e_2.

Since e_1 and e_2 are well-typed, it follows by the fundamental theorem of logical relations that they are related to themselves. To use these assumptions we must pass ownership of a local expression resource to each of e_1 and e_2 with e_1 and e_2 in an evaluation context, respectively. We could use the ζ simulation with e_1 since e_1 is already in an evaluation context in the ζ simulation. However, this leaves us without an expression resource for e_2.

Instead, the idea is to suspend the ζ simulation and create two new simulations ζ_1 and ζ_2 with e_1 as the full specification of the ζ_1 simulation and e_2 as the full specification of the ζ_2 simulation. Then we can appeal to relatedness of e_1 and e_2 to themselves with ζ_1 and ζ_2 as the respective simulations, which will show the existence of the two independent executions of e_1 and e_2. Once $e_1 \,\|\, e_2$ has terminated on the left, we can resume the ζ simulation and use the two independent executions of e_1 and e_2 to take the appropriate steps in the ζ simulation.

The reason this works, is the effect annotations, which ensure that e_1 and e_2 are semi-independent. In particular, all locations accessed by both e_1 and e_2 are read-only and can therefore soundly be shared between the ζ_1 and ζ_2 simulations.

Relating heaps in multiple simulations. In previous relations, the region invariant ensured that all specification heap proxies ($l \stackrel{\pi}{\hookrightarrow}_{S,r} v$) matched the contents of the actual specification heap. With multiple simulations, we have multiple specification heaps. The idea is to allow proxies to be tied to multiple specification heaps, provided we can guarantee that the given references are immutable.

In cases where we cannot guarantee immutability, we still only allow proxies to be tied to a single simulation, to ensure we can reason locally about reductions in simulations.

To capture this formally, we introduce a new ghost resource, to specify whether a region is immutable or not, and which simulations the region proxies are tied to. The $[\text{IM}(r, S, h)]^\pi$ resource asserts that region r is immutable and the current specification heap of the region is h, while $[\text{MU}(r, S)]^\pi$ asserts that it is mutable. In both cases the set S specifies which simulations the proxies of region r are tied to. In the mutable state, we require that the set S is a singleton. We call this ghost resource the *specification link* resource.

The fractional permission is used to track whether we are allowed to change the state of a region. If we own a specification link resource exclusively (i.e., $\pi = 1$), then we can change its state between mutable and immutable and which simulations the proxies of the region are tied to.

$$[\text{MU}(r, S)]^1 \Longleftrightarrow [\text{IM}(r, S', h)]^1 \qquad (22)$$

Both tokens can be split arbitrarily using the fraction. Any two fractional immutable tokens must agree on the current heap and which simulations the region is tied to:

$$[\text{IM}(r, S_1, h_1)]^{\pi_1} * [\text{IM}(r, S_2, h_2)]^{\pi_2} \Longrightarrow \qquad (23)$$
$$[\text{IM}(r, S_1, h_1)]^{\pi_1} * [\text{IM}(r, S_2, h_2)]^{\pi_2} * (h_1 = h_2) * (S_1 = S_2)$$

Disjointness of allocations. We need two final bits of ghost state before we can define the full logical relation. Namely, we need a way to control which locations simulations use when allocating new locations.

To facilitate this level of control over locations, we introduce a ghost resource, $[X]$, for asserting ownership of a set of locations X. These can be split and recombined and ensure that disjoint resources refer to disjoint sets of locations:

$$[X_1 \uplus X_2] \Longleftrightarrow [X_1] * [X_2] \qquad (24)$$
$$[X_1] * [X_2] \Longrightarrow [X_1] * [X_2] * X_1 \cap X_2 = \emptyset \qquad (25)$$

The idea is to give each specification invariant ownership of a countably infinite set of locations that only that simulation may use for future allocations.

To allow simulations to replay reductions from other simulations, we also need a way of deactivating a simulation, such that we can take back ownership of that simulation's locations. To achieve this we introduce a ghost resource $[\text{SR}]^\pi_\zeta$, which we refer to as a *specification runner resource*, to track whether a simulation is active. Ownership of any fraction of this token witnesses that the simulation is active.

The LR_{PAR} logical relation is defined in Figure 9. The most important difference compared to the LR_{BIN} relation, is in the $locs(r)$ predicate contained in the region invariant REG. REG now asserts fractional ownership of a specification link resource for the given region through the *slink* predicate. In case the region is immutable, the pair of heaps given by the specification link resource must match the actual implementation and specification heap for the references belonging to the given region. The region invariant further asserts ownership of the local specification heap resource $l \mapsto^\zeta_S v$ for every simulation $\zeta \in S$ tied to the given region through the specification link resource.

The specification invariant, SPEC, has been extended to support global freshness when allocating, as explained above. Either the specification invariant owns half of the specification runner resource, in which case it also asserts ownership of countably infinite sets of fresh locations through the $disj$ predicate. Otherwise, the specification invariant is inactive and asserts exclusive ownership of the specification runner resource.

New predicates

$$P_{par}(R, g, \varepsilon, M, \zeta) \triangleq \circledast_{\rho \in mutable(R,g,\varepsilon)} [\text{MU}(M(\rho), \{\zeta\})]^{g(\rho)} *$$
$$\circledast_{\rho \in R \setminus mutable(R,g,\varepsilon)} \exists h, S. slink(M(\rho), \{\zeta\} \uplus S, h, g(\rho), g(\rho))$$

$$slink(r, S, h, \pi, \pi') \triangleq [\text{MU}(r, S)]^\pi \vee [\text{IM}(r, S, h)]^{\pi'}$$

$$disj(X_0, X) \triangleq \exists Y. [Y] \wedge \text{dom}(X_0) \cap Y = \emptyset \wedge$$
$$(\text{dom}(X) \setminus \text{dom}(X_0)) \subset Y$$

$$mutable(R, g, \varepsilon) \triangleq \text{wrs } \varepsilon \cup \text{als } \varepsilon \cup \left\{\rho \mid \rho \in R \wedge g(\rho) = \tfrac{1}{2}\right\}$$

Changes to previous definitions

$$locs(r) \triangleq \exists h, S. \, slink(r, S, h_S, \tfrac{1}{2}, \tfrac{1}{4}) * rheap_I(h_I, r) *$$
$$rheap_S(h_S, r) * alloc(h, r) *$$
$$\circledast_{(l,v)\in h_I} l \mapsto_I v * \circledast_{\zeta \in S} \circledast_{(l,v)\in h_S} l \mapsto^\zeta_S v$$

$$\text{SPEC}(h_0, e_0, \zeta) \triangleq \exists h, e, \pi. \, heap_S(h, \zeta) * mctx(e, \zeta) *$$
$$(h_0, e_0) \rightarrow^* (h, e) * ([\text{SR}]^1_\zeta \vee ([\text{SR}]^{\frac{1}{2}}_\zeta * disj(h_0, h_S)))$$

$$P_{reg}(\cdots, \zeta) \triangleq \cdots * P_{par}(R, \tfrac{1}{2} \circ g, \varepsilon, M, \zeta)$$

$$\mathcal{E}^{\Pi, \Lambda}_{\varepsilon, M}(\phi)(e_I, e_S) \triangleq \forall g, j, h_0, e_0, \pi, \zeta. \boxed{\text{SPEC}(h_0, e_0, \zeta)} \vdash$$
$$\left\{ j \xRightarrow{\zeta}_S e_S * [\text{SR}]^\pi_\zeta * P_{reg}(\Lambda, \mathbf{1}, \varepsilon, M, \zeta) * P_{reg}(\Pi, g, \varepsilon, M, \zeta) \right\}$$
$$e_I$$
$$\left\{ \begin{array}{l} v_I. \exists v_S. \, j \xRightarrow{\zeta}_S v_S * [\text{SR}]^\pi_\zeta * \phi(v_I, v_S) * \\ P_{reg}(\Lambda, \mathbf{1}, \varepsilon, M, \zeta) * P_{reg}(\Pi, g, \varepsilon, M, \zeta) \end{array} \right\}_\top$$

Figure 9. LR_{PAR}: Binary rel. for $\lambda_{ref,conc}$ with effect-types and effect-based simulations.

Finally, the expression relation now asserts fractional ownership of the specification runner resource and fractional ownership of specification link resources, for all regions in the context. The specification runner resource ensures that the ζ simulation is active. In case a region is private or the effect mask contains a write or allocation effect for the given region, then the region must be in the mutable state and tied only to the simulation ζ. Otherwise, the region may be in either the mutable or the immutable state, as long as it is tied to the ζ simulation.

The LR_{PAR} relation is sound with respect to contextual approximation and supports parallelization.

Theorem 6 (Soundness). *If* $\Pi \mid \Delta \mid \Gamma \models_{\text{PAR}} e_I \leq_{log} e_S : \tau, \varepsilon$ *then* $\Pi \mid \Delta \mid \Gamma \vdash e_I \leq_{ctx} e_S : \tau, \varepsilon$.

Theorem 7 (Parallelization (semantically)). *If*

1. $\mathcal{E}^{\Lambda_3, \Lambda_1}_{\varepsilon_1, M}([\![\tau_1]\!]^M)(e_{1I}, e_{1S})$ *and* $\mathcal{E}^{\Lambda_3, \Lambda_2}_{\varepsilon_2, M}([\![\tau_2]\!]^M)(e_{2I}, e_{2S})$
2. *als* $\varepsilon_1 \cup$ *wrs* $\varepsilon_1 \subseteq \Lambda_1$, *als* $\varepsilon_2 \cup$ *wrs* $\varepsilon_2 \subseteq \Lambda_2$
3. *rds* $\varepsilon_1 \subseteq \Lambda_1 \cup \Lambda_3$ *and* *rds* $\varepsilon_2 \subseteq \Lambda_2 \cup \Lambda_3$

then $\mathcal{E}^{-, (\Lambda_1, \Lambda_2, \Lambda_3)}_{\varepsilon_1 \cup \varepsilon_2, M}([\![\tau_1 \times \tau_2]\!]^M)(e_{1I} \| e_{2I}, (e_{1S}, e_{2S}))$.

To illustrate how we can use the LR_{PAR} relation to prove contextual refinements that depend on the effect annotations, we give a proof sketch of Theorem 3 below. The full proof can be found in [18].

Recall that Theorem 3 states that each of the two stack modules is contextually equivalent to their counterpart without a **CAS** loop, at an effect type where the local state of the stack module belongs to a private region:

$$- \mid \rho \mid - \vdash stack\_nc_1 \cong_{ctx} stack_1 : \tau'_{\text{STACK}}, \{al_\rho\} \text{ and}$$

$$- \mid \rho \mid - \vdash stack\_nc_2 \cong_{ctx} stack_2 : \tau'_{\text{STACK}}, \{al_\rho\} \text{ with}$$

$$\tau'_{\text{STACK}} = \mathbf{1} \rightarrow^{-, \rho}_{\{al_\rho\}} (\text{int} \rightarrow^{-, \rho}_{\{wr_\rho, rd_\rho\}} \mathbf{1}) \times (\mathbf{1} \rightarrow^{-, \rho}_{\{wr_\rho, rd_\rho\}} \mathbf{1} + \text{int})$$

$stack_{nc_1}() = \textbf{let } h = \textbf{new inj}_1 \ () \ \textbf{in} \ (push_{nc_1}(h), pop_{nc_1}(h))$

$push_{nc_1}(h) = \textbf{rec } p(n).h := \textbf{inj}_2 \ (n, !h)$

$pop_{nc_1}(h) = \textbf{rec } p(\_).\textbf{case}(!h, \textbf{inj}_1 \ () \Rightarrow \textbf{inj}_1 \ (),$

$\quad \textbf{inj}_2 \ (n, v') \Rightarrow h := v'; \textbf{inj}_2 \ n)$

Figure 10. Stack module without **CAS**.

We will focus on the first contextual equivalence where the $stack_{nc_1}$ module is implemented using a single reference to a pure functional list as shown in Figure 10.

To show logical relatedness between $stack_{nc_1}$ and $stack_1$ we will have to assert a relation between the state maintained by the modules. Since the state is local to each module we are not required to use the $[\![\textbf{ref}_\rho \ \tau]\!]^M$ interpretation and are free to pick any invariant to relate the state of the two modules.

A suitable relation would assert ownership of each head-pointer for the region ρ and would state that each pair-wise entry in the stacks are related. STACK is a promising candidate:

$$STACK(h, r, l, v) \triangleq h_I \overset{1}{\hookrightarrow}_{I,r} v_I * h_S \overset{1}{\hookrightarrow}_{S,r} v_S * vals(l, v)$$

$$vals(\text{nil}, v) \triangleq v_I = \textbf{inj}_1 \ () \land v_S = \textbf{inj}_1 \ ()$$

$$vals(x :: xs, v) \triangleq [\![\textbf{int}]\!](x) * \exists v'. \ v_I = \textbf{inj}_2 \ (x_I, v_I') \land$$
$$v_S = \textbf{inj}_2 \ (x_S, v_S') \land vals(xs, v')$$

Note that while we are free to pick an invariant to relate the internal state of the modules, we still use the points-to proxy resources to ensure that the state is tied to simulations correctly.

The STACK relation allows us to read from the head-pointer using a simple extension of property 12:

$$\boxed{REG(r)} \vdash \{STACK(h, r, l, v)\} \ !h_I \qquad (26)$$
$$\{v_I. \exists v_S. \ STACK(h, r, l, (v_I, v_S))\}$$

Similarly, we can use a variant of property 13 to do assignment to the head-pointer location:

$$\boxed{REG(r)} \vdash \{STACK(h, r, l, v)\} \ h_I := v' \qquad (27)$$
$$\{h_I \overset{1}{\hookrightarrow}_{I,r} v' * h_S \overset{1}{\hookrightarrow}_{S,r} v_S * vals(l, v)\}$$

Putting STACK into an invariant is not sufficient, however, for showing the direction $stack_{nc_1} \leq_{log} stack_1$. The reason is that both $push_1(h_S)(n_S)$ and $pop_1(h_S)()$ has a **CAS** operation, that we must guarantee succeeds in a compatible state with $push_{nc_1}(h_I)(n_I)$ and $pop_{nc_1}(h_I)()$, for related n_I and n_S. This relies on the fact that the ρ region is private which ensures that the environment cannot access the local state during the stack operations. We can capture this by defining a REL predicate with a property that allow us to exchange the exclusive write permission $[WR]_r^1$ for ownership of the stack module's points-to proxies:

$$\boxed{REL(h, r)} \vdash [WR]_r^1 \Longleftrightarrow \exists l, v. \ STACK(h, r, l, v) \qquad (28)$$

We will need to establish the invariant when the data structures in the implementation and specification side are both empty:

$$\vdash h_I \overset{1}{\hookrightarrow}_{I,r} \textbf{inj}_1 \ () * h_S \overset{1}{\hookrightarrow}_{S,r} \textbf{inj}_1 \ () \Rightarrow \boxed{REL(h, r)} \qquad (29)$$

REL as defined below allows for the above view-shifts:

$$REL(h, r) \triangleq \exists l, v. \ STACK(h, r, l, v) \lor [WR]_r^1$$

For this particular example we have no need for interpreting read effects on the local state. The REL invariant therefore makes no mention of the $[RD]_r^1$ token.

We show logical equivalence by showing logical approximation in both directions. Here we present a proof outline of the direction

$stack_{nc_1} \leq_{ctx} stack_1$, the full proof of both directions can be found in [18]. Since $stack_{nc_1}$ and $stack_1$ are already values, it suffices to show they are related in the value relation for τ'_{STACK}, which reduces to showing that:

$$\mathcal{E}^{-,\rho}_{\{al_\rho\},M}(\phi)(stack_{nc_1}(), stack_1())$$

where $\phi = [\![(\textbf{int} \to^{-,\rho}_{\{wr_\rho, rd_\rho\}} \textbf{1}) \times (\textbf{1} \to^{-,\rho}_{\{wr_\rho, rd_\rho\}} \textbf{1} + \textbf{int})]\!]^M$.

Thus we first show that we can establish the REL invariant using the local state allocated by the two modules. Next, we show that this invariant is preserved by the push and pop operations and that they are pairwise related assuming this invariant. The proof outline is given below and uses the following two properties to allocate points-to proxies on the implementation and specification side, respectively:

$$\boxed{REG(r)} \vdash \{[AL]_r^\pi\} \ \textbf{new} \ v \ \{y. \ [AL]_r^\pi * h_I \overset{1}{\hookrightarrow}_{I,r} v\} \qquad (30)$$

$$\boxed{REG(r)} \vdash [AL]_r^\pi * [MU(r, \{\zeta\})]^\pi * i \overset{\zeta}{\Rightarrow}_S \textbf{new} \ v \qquad (31)$$
$$\Rightarrow \exists l. \ [AL]_r^\pi * [MU(r, \{\zeta\})]^\pi * i \overset{\zeta}{\Rightarrow}_S l * l \overset{1}{\hookrightarrow}_{S,r} v$$

Context: $\boxed{SPEC(h_0, e_0, \zeta)}, \boxed{REG(r)}, [\![\textbf{1}]\!]^M(y)$

$\left\{ \begin{array}{l} j \overset{\zeta}{\Rightarrow}_S \textbf{let } h_S = \textbf{new inj}_1 \ () \ \textbf{in} \ (push_1(h_S), pop_1(h_S)) * \\ [SR]_\zeta^\pi * [AL]_r^1 * [MU(r, \{\zeta\})]^{\frac{1}{2}} \end{array} \right\}$

$\left\{ \begin{array}{l} \exists i. \ j \overset{\zeta}{\Rightarrow}_S \textbf{let } h_S = i \ \textbf{in} \ (push_1(h_S), pop_1(h_S)) * \\ i \overset{\zeta}{\Rightarrow}_S \textbf{new inj}_1 \ () * [SR]_\zeta^\pi * [AL]_r^1 * [MU(r, \{\zeta\})]^{\frac{1}{2}} \end{array} \right\}$

$\quad$ **new inj**$_1$ ()

`// Follows from Properties 30 and 31`

$\left\{ \begin{array}{l} h_I. \ \exists h'_S, i. \ j \overset{\zeta}{\Rightarrow}_S \textbf{let } h_S = i \ \textbf{in} \ (push_1(h_S), pop_1(h_S)) * \\ i \overset{\zeta}{\Rightarrow}_S h'_S * [SR]_\zeta^\pi * [AL]_r^1 * [MU(r, \{\zeta\})]^{\frac{1}{2}} * \\ h_I \overset{1}{\hookrightarrow}_{I,r} \textbf{inj}_1 \ () * h_S \overset{1}{\hookrightarrow}_{S,r} \textbf{inj}_1 \ () \end{array} \right\}$

`// Follows from Property 29`

$\left\{ \begin{array}{l} \exists h_I, h'_S. \ j \overset{\zeta}{\Rightarrow}_S \textbf{let } h_S = h'_S \ \textbf{in} \ (push_1(h_S), pop_1(h_S)) * \\ [SR]_\zeta^\pi * [AL]_r^1 * [MU(r, \{\zeta\})]^{\frac{1}{2}} * \boxed{REL((h_I, h'_S), r)} \end{array} \right\}$

$\left\{ \begin{array}{l} \exists h. \ j \overset{\zeta}{\Rightarrow}_S (push_1(h_S), pop_1(h_S)) * [SR]_\zeta^\pi * [AL]_r^1 * \\ [MU(r, \{\zeta\})]^{\frac{1}{2}} * \boxed{REL(h, r)} \end{array} \right\}$

$\quad (push_{nc_1}(h_I), pop_{nc_1}(h_I))$

$\left\{ \begin{array}{l} v_I. \ \exists v_S. \ j \overset{\zeta}{\Rightarrow}_S v_S * [SR]_\zeta^\pi * [AL]_r^1 * [MU(r, \{\zeta\})]^{\frac{1}{2}} * \\ [\![(\textbf{int} \to^{-,\rho}_{\{wr_\rho, rd_\rho\}} \textbf{1}) \times (\textbf{1} \to^{-,\rho}_{\{wr_\rho, rd_\rho\}} \textbf{1} + \textbf{int})]\!]^M(v_I, v_S) \end{array} \right\}$

For the last step we need to show the following two refinements:

$\boxed{REL(h, r)}, [\![\textbf{int}]\!]^M(n) \vdash$
$$\mathcal{E}^{-,\rho}_{\{wr_\rho, rd_\rho\},M}([\![\textbf{1}]\!]^M)(push_{nc_1}(h_I)(n_I), push_1(h_S)(n_S))$$

$\boxed{REL(h, r)} \vdash$
$$\mathcal{E}^{-,\rho}_{\{wr_\rho, rd_\rho\},M}([\![1 + \textbf{int}]\!]^M)(pop_{nc_1}(h_I)(), pop_1(h_S)())$$

We sketch a proof of the first refinement below. The proof of the second refinement can be found in [18].

The interpretation of the region and effect annotations for the first refinement is as follows. Since we do not interpret read effects on the local state, $[RD]_r^1$ is framed off in the proof outline below.

$$P_{reg}(\{\rho\}, \textbf{1}, \{wr_\rho, rd_\rho\}, M[\rho \mapsto r], \zeta)$$
$$= [WR]_r^1 * [RD]_r^1 * [MU(r, \{\zeta\})]^{\frac{1}{2}} * \boxed{REG(r)}$$

The proof starts by trading the exclusive write token for ownership of the local state (28). The we use (27) to push n_I onto the implementation-side stack.

Context: $\boxed{\mathrm{SPEC}(h_0, e_0, \zeta)}$, $\boxed{\mathrm{REL}(h, r)}$, $\boxed{\mathrm{REG}(r)}$, $[\![\mathbf{int}]\!](n)$

$\left\{ j \overset{\zeta}{\Rightarrow}_S push_1(h_S)(n_S) * [\mathrm{SR}]_\zeta^\pi * [\mathrm{WR}]_r^1 * [\mathrm{MU}(r, \{\zeta\})]^{\frac{1}{2}} \right\}$

$\left\{ \begin{array}{l} \exists l, v.\ j \overset{\zeta}{\Rightarrow}_S push_1(h_S)(n_S) * [\mathrm{SR}]_\zeta^\pi * [\mathrm{MU}(r, \{\zeta\})]^{\frac{1}{2}} * \\ \mathrm{STACK}(h, r, l, v) \end{array} \right\}$

$\quad !h_I$

$\left\{ \begin{array}{l} v_I.\ \exists l, v_S.\ j \overset{\zeta}{\Rightarrow}_S push_1(h_S)(n_S) * [\mathrm{SR}]_\zeta^\pi * [\mathrm{MU}(r, \{\zeta\})]^{\frac{1}{2}} * \\ \mathrm{STACK}(h, r, l, (v_I, v_S)) \end{array} \right\}$

$\quad h_I := \mathbf{inj}_2\ (n_I, v_I)$

$\left\{ \begin{array}{l} v_I'.\ \exists l, v_S.\ j \overset{\zeta}{\Rightarrow}_S push_1(h_S)(n_S) * [\mathrm{SR}]_\zeta^\pi * [\mathrm{MU}(r, \{\zeta\})]^{\frac{1}{2}} * \\ h_I \overset{1}{\hookrightarrow}_{I,r} \mathbf{inj}_2\ (n_I, v_I) * h_S \overset{1}{\hookrightarrow}_{S,r} v_S * vals(l, v) * v_I' = () \end{array} \right\}$

After pushing n_I on the stack on the implementation side, we simulate pushing n_S on the specification side. Let:

$$K_1 \triangleq \mathbf{let}\ v = [] \ \mathbf{in}$$
$$\qquad \mathbf{if}\ \mathbf{CAS}(h_S, v, \mathbf{inj}_2\ (n_S, v))\ \mathbf{then}\ ()\ \mathbf{else}\ loop(n_S)$$
$$K_2 \triangleq \mathbf{if}\ []\ \mathbf{then}\ ()\ \mathbf{else}\ loop(n_S)$$

be the evaluation contexts which require a non-trivial reduction. Notice that $push_1(h_S)(n_S) = K_1[!h_S]$. We can now perform the simulation:

$\boxed{\mathrm{SPEC}(h_0, e_0, \zeta)}$, $\boxed{\mathrm{REG}(r)}$ $\vdash$

$\quad j \overset{\zeta}{\Rightarrow}_S push_1(h_S)(n_S) * [\mathrm{SR}]_\zeta^\pi * [\mathrm{MU}(r, \{\zeta\})]^{\frac{1}{2}} * h_S \overset{1}{\hookrightarrow}_{S,r} v_S$

$\Rightarrow \exists i.\ j \overset{\zeta}{\Rightarrow}_S K_1[i] * i \overset{\zeta}{\Rightarrow}_S !h_S * [\mathrm{SR}]_\zeta^\pi * [\mathrm{MU}(r, \{\zeta\})]^{\frac{1}{2}} *$

$\quad h_S \overset{1}{\hookrightarrow}_{S,r} v_S$

$\Rightarrow \exists i.\ j \overset{\zeta}{\Rightarrow}_S K_2[i] * i \overset{\zeta}{\Rightarrow}_S \mathbf{CAS}(h_S, v_S, \mathbf{inj}_2\ (n_S, v_S)) * [\mathrm{SR}]_\zeta^\pi *$

$\quad [\mathrm{MU}(r, \{\zeta\})]^{\frac{1}{2}} * h_S \overset{1}{\hookrightarrow}_{S,r} v_S$

$\Rightarrow j \overset{\zeta}{\Rightarrow}_S () * [\mathrm{SR}]_\zeta^\pi * [\mathrm{MU}(r, \{\zeta\})]^{\frac{1}{2}} * h_S \overset{1}{\hookrightarrow}_{S,r} \mathbf{inj}_2\ (n_S, v_S)$

The simulation follows from repeatedly stepping by using property 15 and 16. Observe that $\mathbf{CAS}(h_S, v_S, \mathbf{inj}_2\ (n_S, v_S))$ always succeeds since we have ownership of $h_S \overset{1}{\hookrightarrow}_{S,r} v_S$. We can now finish the proof by reestablishing the relation between the local state of the modules and trading it for the exclusive write permission:

$\left\{ \begin{array}{l} \exists l, v_S.\ j \overset{\zeta}{\Rightarrow}_S () * [\mathrm{SR}]_\zeta^\pi * [\mathrm{MU}(r, \{\zeta\})]^{\frac{1}{2}} * \\ h_I \overset{1}{\hookrightarrow}_{I,r} \mathbf{inj}_2\ (n_I, v_I) * h_S \overset{1}{\hookrightarrow}_{S,r} \mathbf{inj}_2\ (n_S, v_S) * vals(l, v) \end{array} \right\}$

$\left\{ \begin{array}{l} \exists l, v'.\ j \overset{\zeta}{\Rightarrow}_S () * [\mathrm{SR}]_\zeta^\pi * [\mathrm{MU}(r, \{\zeta\})]^{\frac{1}{2}} * h_I \overset{1}{\hookrightarrow}_{I,r} v_I' * \\ h_S \overset{1}{\hookrightarrow}_{S,r} v_S' * vals(n :: l, v') \end{array} \right\}$

$\left\{ \begin{array}{l} \exists l, v'.\ j \overset{\zeta}{\Rightarrow}_S () * [\mathrm{SR}]_\zeta^\pi * [\mathrm{MU}(r, \{\zeta\})]^{\frac{1}{2}} * \\ \mathrm{STACK}(h, r, n :: l, v') \end{array} \right\}$

$\left\{ \exists l, v'.\ j \overset{\zeta}{\Rightarrow}_S () * [\mathrm{SR}]_\zeta^\pi * [\mathrm{WR}]_r^1 * [\mathrm{MU}(r, \{\zeta\})]^{\frac{1}{2}} \right\}$

4. Discussion

We have already mentioned some related work along the way; here we discuss some other related work.

Benton *et al.* initiated a line of work on relational models of type-and-effect systems to formally justify effect-based program transformations for increasingly sophisticated sequential programming languages and increasingly expressive effect systems [4–8, 29]. Birkedal *et al.* showed how to extend this approach to a concurrent language [11]. The effect system we use here is from *loc. cit.* Birkedal *et al.*'s relational interpretation is defined by a concrete step-indexed Kripke logical relation. They used the model to prove a parallelization theorem similar to ours, but the proof was very technical and consisted of manual disentangling and re-ordering of computation steps. Part of the reason for this was that support for parallelization was not built into their logical relation and had to be proven separately. In contrast, we build in support for parallelization in the LR_{PAR} relation through its support for multiple simulations. This allows us to reduce the proof of the parallelization theorem to the essence of why it holds: framing. Moreover, as mentioned in the Introduction, it makes it possible to use the program logic to show that an expression satisfies the semantic invariants imposed by the type system even if the expression is not statically well-typed and to reason about refinements.

In recent work, Benton *et al.* [9] have also considered a concurrent language, which in contrast to the language considered here only includes first-order store. Technically, this makes the construction of a logical relations model simpler, since one avoids having to deal with the type-world circularity mentioned in the Introduction. Their type-and-effect system does not support dynamic allocation of abstract locations (which correspond to regions in our setup), requiring all abstract locations to be given up front. Our type-and-effect system supports dynamic allocation and hiding of regions, through the masking rule. On the other hand, their effect system supports a notion of abstract effects, which means, *e.g.,* that an operation in a data structure module can be considered pure even if it uses effects internally, as long as those effects are not observable outside the module boundary. Benton *et al.* use this facility for treating refinement of fine-grained concurrent data structures, illustrated using an idealized Michael-Scott queue. Our semantics also supports refinements between fine-grained concurrent data structures, using Iris' support for general invariants. In this paper we have focused on an example of a refinement that only holds by restricting interference through the type-and-effect system. Our method also scales to fine-grained concurrent data structures that use helping, thanks to Iris [16].

Raza *et al.* [26] and Botincan *et al.* [12], both explore automatic parallelization of sequential programs verified in separation logic. Raza *et al.* rely on specifications inferred from a shape analysis. Botincan *et al.* explore the idea of using the proof to insert synchronization that ensures the dependencies of the original program are preserved. These analyses focus on first-order programs, whereas our type-and-effect system applies to higher-order programs.

The idea of defining logical relations in a program logic goes back at least to Plotkin and Abadi, who used a second-order logic to define logical relations for a second-order lambda calculus [22]. Dreyer *et al.* used a second-order logic with a Löb modality, inspired by [3], to give a logical relations interpretation of a programming language with recursive types [13]. The logic used by Dreyer *et al.* did not support invariants and hence it did not support the interpretation of reference types. Turon *et al.* showed how to use a variant of second-order concurrent separation logic with invariants for giving a logical relations interpretation of an ML-like type system for a language similar to the one considered in this paper [31]. To define logical relations in the unary separation logic, their logic had a *built-in* notion of specification resources and a single specification invariant. In contrast, here we use a higher-order concurrent separation logic, Iris, which is flexible enough that one can define specification resources and invariants in it. We rely crucially on this flexibility for the LR_{PAR} relation to support multiple simulations, as discussed in Section 3.4.

References

[1] A. Ahmed, A. Appel, and R. Virga. A Stratified Semantics of General References. In *LICS*, 2002.

[2] A. Ahmed, D. Dreyer, and A. Rossberg. State-dependent representation independence. In *POPL*, 2009.

[3] A. Appel, P.-A. Melliès, C. Richards, and J. Vouillon. A very modal model of a modern, major, general type system. In *POPL*, 2007.

[4] N. Benton and P. Buchlovsky. Semantics of an effect analysis for exceptions. In *TLDI*, 2007.

[5] N. Benton, A. Kennedy, M. Hofmann, and L. Beringer. Reading, writing and relations. In *PLAS*. Springer, 2006.

[6] N. Benton, A. Kennedy, L. Beringer, and M. Hofmann. Relational semantics for effect-based program transformations with dynamic allocation. In *PPDP*, 2007.

[7] N. Benton, A. Kennedy, L. Beringer, and M. Hofmann. Relational semantics for effect-based program transformations: higher-order store. In *PPDP*, 2009.

[8] N. Benton, M. Hofmann, and V. Nigam. Abstract effects and proof-relevant logical relations. In *POPL*, 2014.

[9] N. Benton, M. Hofmann, and V. Nigam. Effect-dependent transformations for concurrent programs. In *PPDP*, 2016.

[10] L. Birkedal, B. Reus, J. Schwinghammer, K. Støvring, J. Thamsborg, and H. Yang. Step-Indexed Kripke models over recursive worlds. In *POPL*, 2011.

[11] L. Birkedal, F. Sieczkowski, and J. Thamsborg. A concurrent logical relation. In *CSL*, 2012.

[12] M. Botincan, M. Dodds, and S. Jagannathan. Proof-Directed Parallelization Synthesis by Separation Logic. *TOPLAS*, 35(2), 2013.

[13] D. Dreyer, A. Ahmed, and L. Birkedal. Logical step-indexed logical relations. *Logical Methods in Computer Science*, 7(2), 2011.

[14] M. Fähndrich and R. DeLine. Adoption and focus: practical linear types for imperative programming. In *PLDI*, 2002.

[15] D. K. Gifford and J. M. Lucassen. Integrating functional and imperative programming. In *LISP*, 1986.

[16] R. Jung, D. Swasey, F. Sieczkowski, K. Svendsen, A. Turon, L. Birkedal, and D. Dreyer. Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. In *POPL*, 2015.

[17] N. Krishnaswami, P. Pradic, and N. Benton. Integrating linear and dependent types. In *POPL*, 2015.

[18] M. Krogh-Jespersen, K. Svendsen, and L. Birkedal. A Relational Model of Types-and-Effects in Higher-Order Concurrent Separation Logic: Technical Appendix. `http://www.kasv.dk/irisrelmodel-tr.pdf`.

[19] J. M. Lucassen and D. K. Gifford. Polymorphic effect systems. In *POPL*, 1988.

[20] G. Morrisett, A. Ahmed, and M. Fluet. L^3: A linear language with locations. In *TLCA*, 2005.

[21] A. Nanevski, G. Morrisett, and L. Birkedal. Polymorphism and separation in hoare type theory. In *ICFP*, 2006.

[22] G. Plotkin and M. Abadi. A logic for parametric polymorphism. In *TLCA*, 1993.

[23] F. Pottier. Hiding local state in direct style: a higher-order anti-frame rule. In *LICS*, 2008.

[24] F. Pottier and J. Protzenko. Programming with permissions in Mezzo. In *ICFP*, 2013.

[25] F. Pottier and J. Protzenko. A few lessons from the mezzo project. In *SNAPL*, 2015.

[26] M. Raza, C. Calcagno, and P. Gardner. Automatic Parallelization with Separation Logic. In *ESOP*, 2009.

[27] Rust Language. `https://doc.rust-lang.org`, 2016.

[28] F. Smith, D. Walker, and G. Morrisett. Alias types. In *ESOP*, 2000.

[29] J. Thamsborg and L. Birkedal. A kripke logical relation for effect-based program transformations. In *ICFP*, 2011.

[30] M. Tofte and J.-P. Talpin. Implementation of the typed call-by-value λ-calculus using a stack of regions. In *POPL*, 1994.

[31] A. Turon, D. Dreyer, and L. Birkedal. Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency. In *ICFP*, 2013.

[32] A. Turon, J. Thamsborg, A. Ahmed, L. Birkedal, and D. Dreyer. Logical relations for fine-grained concurrency. In *POPL*, 2013.