

---

# Towards Modular Reasoning for Realistic Programming Languages

Filip Sieczkowski

---

PhD Dissertation  
IT University of Copenhagen

November 2013



# Abstract

Computer programs, even the ones whose correctness is vital, are growing ever larger. Thus, if we are to be capable of verifying software, we need to do so *modularly*, i.e., reason separately about separate parts of the big programs and conclude that the whole program is correct based on correctness of its subcomponents. In recent years, we have seen a steady advance in modular techniques for software verification. However, these new developments are still far from encompassing all the features that modern programming languages utilize.

In this thesis, we study several approaches to extend modular reasoning techniques to more realistic programming languages. We consider three distinct directions, which deal with both traditional program logic style verification, as well as relational reasoning.

In **Verifying Object-Oriented Programs with Higher-Order Separation Logic in Coq** we construct a logic for giving modular separation logic specifications to object-oriented code that uses dynamic dispatch through interfaces. This powerful programming technique is commonly used in virtually all modern object-oriented languages, and most previous work only targeted the most common design patterns. We present a logic that is powerful enough to express the rarer, complex patterns of usage, while allowing for simple specification of simpler, more concrete interfaces. We follow this direction with **Formalized Verification of Snapshotable Trees: Separation and Sharing**, where we show that our approach scales to non-trivial programs. In this article we consider a family of sophisticated data structures that use complex sharing patterns, notoriously difficult to reason about with separation logic, and show how to verify one of them using our approach.

**A Separation Logic for Fictional Sequential Consistency** takes our work on separation logic in a different direction: based on a cutting-edge logic for reasoning about concurrent programs, it tackles the problem of real-world memory models, focusing on the common TSO model. The paper provides two interconnected logics: the TSO logic, in which one can reason about the additional observable behaviours, which the user can observe on the weak memory model, and the SC logic, in which reasoning is standard. Since most well-behaved code (e.g., well synchronized programs) cannot observe any additional behaviours admitted by the memory model, we thus provide the much

needed abstraction from the machine model for a large class of higher-level code.

Finally, in **A Concurrent Logical Relation** we investigate a different, *relational*, approach to reasoning about programs. We treat a concurrent extension of an ML-like language, build the first logical relation for a concurrent programming language, and prove a parallelization theorem, a long standing conjecture that describes, in terms of types and effects, when it is safe to run two expressions in parallel.

# Acknowledgments

First and foremost, I would like to thank my advisors, Lars Birkedal and Peter Sestoft, for their continued guidance through these three years, and for outstanding patience and support. I could not imagine better supervisors.

A big thank-you to my collaborators, for the long and interesting discussions, and the ideas we shared. In particular, to my colleagues from the ToMeSo project: Jesper Bengtson, Jonas B. Jensen, Hannes Mehnert and Jacob Thamsborg, and to Kasper Svendsen; they have all been great company, and I will fondly remember the times spent together, both in and out of work.

Finally, to my the friends in Copenhagen, Aarhus, Wrocław and everywhere else, for sharing both in the good times, and the slightly worse. To the wonderful PLS group at the IT University for making it an exceptionally fun environment to work in. And to my family, for always being there when I needed them. Thank you all.

*Filip Sieczkowski,  
Copenhagen–Aarhus, 29th November 2013.*



# Contents

<b>Abstract</b>	<b>i</b>
<b>Acknowledgments</b>	<b>iii</b>
<b>Contents</b>	<b>v</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 Verifying Object-Oriented Programs in Coq</b>	<b>15</b>
<b>3 Formalized Verification of Snapshotable Trees</b>	<b>31</b>
<b>4 A Separation Logic for Fictional Sequential Consistency</b>	<b>49</b>
<b>5 A Concurrent Logical Relation</b>	<b>61</b>
<b>A iCAP-TSO: Technical Appendix</b>	<b>93</b>





# Introduction

The recent spread of computers and computer programs to virtually all areas of human life has increased the need for understanding whether these programs are correct. Especially in the case of safety critical systems, which are more and more computerized, we would like to be able to conclude whether the systems will run correctly. However, at the same time the systems have grown in size and complexity, to the extent where analyzing a complete program is virtually impossible. This underlines the need for *modular* techniques of reasoning about software. By modularity we mean, in this context, that we should be able to reason independently about separate components of the system, and deduce the properties of the whole system from the properties of these parts.

In the last decade, we have seen a steady advance in modular techniques for software verification. However, these new developments are still far from encompassing all the features that modern programming languages utilize. In this thesis, we concentrate on two such techniques: separation logic and Kripke logical relations. We consider three distinct directions in which we can extend modular reasoning techniques to more realistic programming languages. First, however, we present a short overview of the key problem.

## 1 The problems of shared mutable state.

One of the biggest problems in verification of real-life software has, for a long time, been the treatment of shared mutable state. This is a feature pervasive in almost all modern programming languages. In simple imperative languages the pointers allow the programmer to define complex data structures; in object-oriented programs the objects are stored on the heap and can be referenced from multiple different points of the program; even in most functional languages imperative, higher-order state is used to provide the means for more efficient implementations. Also in the concurrent programming languages, most communication primitives are usually developed in terms of shared memory.

To understand how shared mutable state constitutes a problem, consider the following code snippet:

```
x := 1;
y := 2
```

Let us try to specify what value can be stored in the cell  $x$  after the snippet executes. Clearly, it can be 1, as long as  $y$  denotes a different memory cell than  $x$ . However, what if both  $x$  and  $y$  denote the same location? In this case the value of  $x$  should be 2, since this is the final value written to that memory cell. This is an instance of the problem of *pointer aliasing*, and it is one of the problems underlying the considerations in this thesis.

From the program verification perspective, this problem basically reduces to being able to describe pointer aliasing — or, rather, lack of aliasing — in a

concise and modular way. However, reasoning about pointer aliasing directly, using explicit equality of the references, is problematic. In the case of our example, this is still relatively simple: we could informally specify our snippet by saying, for example “if  $x = y$ , then the value of  $x$  is 2, otherwise the value of  $x$  is 1”. However, this approach does not readily scale even to the simplest data structures, especially if we want to keep the specification abstract. For example, if we want to append two singly-linked lists in constant time, using a pointer to the end of the list, we need to know that the two lists are not overlaid: the structures cannot share any pointers. If the list were overlaid, the list resulting from appending one to the other would become a cyclic structure, a result we might assume the programmer did not expect of an append function.

The problem runs even deeper, actually. Even if we manage to express the disjointness of the two lists in question, we do not know how the append action would affect any of the other lists we might know about. This makes specifying functions in a concise way virtually impossible. We discuss one of the ways to solve this problem in the following section.

## 2 Program logics: the case for separation.

We begin this section with a short overview of Hoare logic. Then, we proceed to discuss the basics of separation logic, motivating its design by how it helps with the problem of pointer aliasing.

Hoare logic [16] was one of the first attempts at reasoning about programs in a modular way. It was originally developed for a simple imperative language, but since its inception a lot of work was put into extending it to more realistic languages, a process that still continues, and of which this thesis is a small part. Its key characteristic is the use of *Hoare triples* of the form

$$\{P\} c \{Q\}$$

where  $P$  and  $Q$  are *assertions* and  $c$  is the command that we are reasoning about. The assertions are usually interpreted as predicates on the program states, and the interpretation of a triple  $\{P\} c \{Q\}$  states that if we evaluate the program from a state  $s$  that satisfies  $P$  and the evaluation terminates in a state  $s'$ , then  $s'$  satisfies  $Q$ . Assertions themselves form a logic, and we can use the reasoning within this logic via a rule of consequence:

$$\frac{P_1 \vdash P_2 \quad \{P_2\} c \{Q_2\} \quad Q_2 \vdash Q_1}{\{P_1\} c \{Q_1\}}.$$

Hoare triples can also be easily adapted to a language of expressions, by taking the postcondition  $Q$  to be parameterized on the final value of the expression.

The notion of Hoare triples allows us to reason about simple imperative code, but in fact we need a lot more power if we are to handle realistic programming languages. As Reynolds puts it in [25], where he introduces the notion of *specification* logic,

[T]o obtain universal specifications about commands in Algol-like languages, it is necessary to extend the language of specifications well beyond the Hoare triples.

In fact, we can build a general higher-order logic of specifications, with Hoare triples being just one of the basic formulae. Another kind of basic formula might be, for example, a specification of a procedure. Adding this kind of specification would allow us to verify programs that use procedure calls, as shown in [25]. Sufficiently strong specification logics can also allow us to reason about correctness of recursive or mutually recursive functions in a general way, for example using Löb induction [7], or, as we show further on, specify object-oriented code that uses interfaces.

These developments, however, do not readily scale to shared mutable state. In fact, the way to reason in a modular way in the classical Hoare logic was through an admissible conjunction proof rule,

$$\frac{\{P\} c \{Q\}}{\{P \wedge R\} c \{Q \wedge R\}},$$

which usually requires side conditions dependent on the programming language for which the logic is designed. This rule, however, is not admissible if the state of the program includes shared mutable state: the changes to the heap that  $c$  makes might well invalidate  $R$ !

## Separation Logic.

This brings us to separation logic. Introduced independently by Reynolds in [26], and Ishtiaq and O’Hearn in [18], separation logic extends the classic assertion logic of Hoare’s in a way that makes a variant of the proof rule discussed above its cornerstone. Specifically, let us consider *when* would the conjunction rule hold: clearly, if the assertion  $R$  were made about a part of state that was *different*, disjoint from the part of the state that  $P$  was made about, such a property should hold. This observation leads us to a new assertion logic connective, called *separating conjunction* and written  $*$ . Intuitively,  $P * Q$  holds in any state that can be split in two *disjoint* sub-states, one of which satisfies  $P$ , and the other  $Q$ . We also need some primitive way of specifying what values are stored in the heap; for this, we use a *points-to* predicate, written  $l \mapsto v$ , which asserts that the location  $l$  is allocated in the heap, and the value stored at  $l$  is indeed  $v$ .

Coming back to the example from Section 1, we can equip the snippet with a precondition

$$x \mapsto u * y \mapsto v.$$

This precondition ensures that  $x$  and  $y$  *can* actually be dereferenced, since we have a points-to assertion for both of these locations. But equally importantly, it states that  $x \neq y$  — since if they were equal, we would not be able to split the state into two disjoint parts such that the same location were in both parts.

As long as our assertion logic also allows for existential quantification, these new connectives are a powerful tool for describing the shape and contents of disjoint data structures, such as lists or trees. Consider, for example, the following, classic specification of a singly-linked list:

$$\begin{aligned} \text{list}(x, \text{vs}) &\stackrel{\text{def}}{=} (x = \mathbf{null} \wedge \text{vs} = \varepsilon) \\ &\vee \exists y, v, \text{vs}' . x \mapsto (v, y) * \text{list}(y, \text{vs}') \wedge \text{vs} = v \cdot \text{vs}' \end{aligned}$$

This specification, defined by induction on the mathematical representation of the list,  $\text{vs}$ , ensures that no parts of the list are shared within the list. If we did not require this, our definition would admit circular lists, which would make reasoning about updates we might make to the list much more complicated. Moreover, if we have an assertion  $\text{list}(x, \text{us}) * \text{list}(y, \text{vs})$ , we are sure that no pointers except for the final null pointer are shared between the two lists. In contrast, a separate assertion that would describe that the lists are well-formed and not overlaid would have to be at least as complex as the definition itself — and it would not give us any knowledge about other lists that we might know about.

These definitions lead us to rephrase the Hoare logic rule for conjoining an assertion to a triple in a way that works in the presence of shared mutable state: the *frame rule*. As with its predecessor, it may need programming language-dependent side conditions, but its general form is

$$\frac{\{P\} c \{Q\}}{\{P * R\} c \{Q * R\}}.$$

This proof rule states that if we can split our state into two parts, one satisfying  $P$  and the other  $R$ , and if the part that satisfies  $P$  suffices to run the program  $c$  and end in a state that satisfies  $Q$ , then running the program in the state consisting of both parts will give us a state that again can be split in two parts: one that satisfies  $Q$ , and one that still satisfies  $R$ . Intuitively, this rule should hold, since the program does not “need” the part of the state described by  $R$ , so it should not modify it, but in some settings more complex arguments are needed to show that the rule is sound.

This ability to succinctly and accurately describe the shapes of data structures ties in well with some of the useful features of specification logics. For example, if the specification logic admits second-order quantification, existential quantifiers can be used to make the predicates that encode the shape of the data structure abstract [10]. This is similar in spirit to using existential quantification for data abstraction. Specific proof systems sometimes use more syntactic variants with a particular mode of use, an example developed for the object-oriented setting are *abstract predicates*, introduced by Parkinson and Bierman in [23].

This new expressive power does not come without a cost, however. One can easily notice that the separating conjunction does not admit all the usual structural rules that we would expect of a conjunction. In particular, contraction

rule does not hold in general, since we have

$$l \mapsto v \not\vdash l \mapsto v * l \mapsto v.$$

Admissibility of the weakening rule depends on the design and precise semantics of the logic. A good source for an overview of design considerations and techniques is, for example, a recent review by Jensen [19].

### 3 Relational reasoning and higher-order store.

In this section, we briefly discuss the development of the logical relation techniques to the point where they are able to handle languages with higher-order store.

The logical relations technique was developed in the 1970s by, among others, Tait [27] and Plotkin [24]. The technique was first developed for showing unary properties of languages, for instance the strong normalization of simply-typed lambda calculus, and later also for binary properties, most importantly contextual equivalence. The essence of the technique is interpreting syntactic types as relations that are determined by the structure of the type: for example,  $\llbracket \sigma \rightarrow \tau \rrbracket$  should be determined from the relations  $\llbracket \sigma \rrbracket$  and  $\llbracket \tau \rrbracket$  in a way that guarantees closure under lambda abstraction and application; see [21, Chapter 8] for a good introduction to the technique, as well as more historical references.

In the most basic setting, the types can be interpreted simply as sets of well-typed values or expressions. However, this simple mode of interpretation does not readily scale to richer type systems, including polymorphic calculi. The crucial idea that allows us to interpret the universal quantification, is to have a space of *semantic* types, closed under intersections. Once we have such a space, we can give an interpretation to an open type that is well-formed in a type variable context,  $\Delta \vdash \tau$ . We interpret  $\Delta$  as a sequence of semantic types, one for each variables, and the interpretation of the judgment is a map from the interpretation of  $\Delta$  into semantic types. This treatment admits interpretations of universal and existential quantifiers via intersections and unions respectively, as long as the space of semantic types is closed under these operations. This idea, though stated differently, dates back to Girard’s original paper on System F [15]. A good reference for the treatment of polymorphism is [21, Chapter 9].

Still more involved is the treatment of the recursive types. In contrast to universal and existential quantifiers, we cannot directly express the interpretation of the judgment  $\Delta \vdash \mu \alpha. \tau$  in terms of the interpretation of  $\Delta, \alpha \vdash \tau$  — at least not without enriching the space of semantic types with some extra structure. One of the possibilities is to ensure that the space of semantic types is non-empty and has a metric structure. Then, as long as the interpretation of  $\Delta, \alpha \vdash \tau$  that we want to use to interpret the recursive type is a contractive map, we would be able to use Banach’s fixed point theorem to give an interpretation of the recursive type. This approach was first explored by Amadio [5], and Abadi and Plotkin [1]; see for example Birkedal et al. [12] and the references therein for more extensive treatment.

For all their usefulness, the domain-theoretic models mentioned above are complex, and require solid mathematical grounding. Thus, simpler, more operational treatment of recursive types would also be useful. A line of research that provided such a technique is step-indexing. In this setting, one explicitly stratifies the space of semantic types by using a natural number that intuitively corresponds to the number of evaluation steps for which the type provides information. At index 0, the type does not provide any information, while at each higher index the information is more refined. Since the stratified relation is tied to the evaluation steps, this setup provides a relatively simple way to interpret iso-recursive types: unfolding a recursive value requires an evaluation step, so the interpretation of the recursive type only needs to depend on information from the previous stratum, and thus is well-defined. Step-indexing has been introduced by Appel and McAllester [6], and a good reference for this treatment of recursive types is Ahmed’s work [3]. Note that the syntactic treatment has been since reconciled with the semantic one by Birkedal et al. in [11], by showing that the space of semantic types in the step-indexed models can also be seen as an ultrametric space.

The higher-order store, the final language extension we consider in this section, presents us with a different challenge still. This is due to the fact that we can no longer consider the semantic type in isolation: after all, we have to consider references, and whether they are well-typed should depend on the store. More explicitly, assume we have a space of semantic types  $T^\circ$ . To be able to provide an interpretation of reference types, we need a space of semantic types that would know, at the least, what locations are defined. This gives us a domain equation of the following shape:

$$T = W \rightarrow T^\circ,$$

where  $W$  denotes the space of *worlds*, the semantic descriptions of heaps. In the case of a ground store, it is enough for the world to describe which locations are allocated, and so we could take  $W = \mathcal{P}_{fin}(L)$ , a finite set of locations. However, if we have higher-order store, we need to know what type is associated with any given location. This leaves us, in the simplest case, with an equation of the shape

$$W = L \multimap_{fin} T.$$

Solving this sort of recursive domain equations is the crucial challenge in modeling the higher-order store, and since the recursion occurs in a negative position, the solution is not trivial. That said, in the recent years different solutions for these types of domain equations have been provided, including [2, 4, 17, 12, 11].

Note that the notion of worlds given by the equation above is a very simplistic one. It can be extended to much more complex structures, see for example [4] or some of the recent work on models of concurrent calculi with higher-order state, e.g., [30].

## Relational models and aliasing

The problem of aliasing is just as pressing in the relational setting, as it is when reasoning in unary models. Coming back to the code snippet of Section 1, let us consider the following equation:

$$e_1; e_2 \simeq e_2; e_1.$$

Should we be able to reorder the assignments in the example with this equality? Clearly we should only be allowed to do that if we know that  $x$  and  $y$  are not aliased, but this means that the equation does not hold in general. Thus, if we are to provide useful equivalences for programs that use higher-order store, we need to be able to express what aliasing patterns are permitted for the program equivalence to hold.

In the logical relations setting, one way to express such restrictions is to strengthen the underlying type system. To this end, we can use a *type-and-effect* system, which, in addition to the information on the type of an expression, also provides the information on side effects that it may perform. However, since the type and effect information should be determined statically, the precise locations that may or may not be aliased cannot be determined in general. Thus, the type-and-effect systems usually partition the heap into disjoint *regions*, with each reference being allocated in one of the regions. Thus, they can statically infer the *lack* of aliasing — as long as the references belong to different regions of the memory. The effects usually include read, write and allocate actions in a region, but it is possible to extend this set. Type-and-effect systems were introduced by Gifford and Lucassen in [14, 20], and have since been used, for example, to define static analyses [22] or even as a basis for a runtime system of SML [29, 13].

What interests us in the context of this thesis is the use of a type-and-effect system to enable relational reasoning, and in particular to justify effect-based program transformations. This possibility has been conjectured in the original papers of Gifford and Lucassen, but the relational models of higher-order programming languages were not strong enough to handle this kind of reasoning. However, with the development of rich relational models by Benton et al. [8, 9], this line of research has opened to new possibilities, and the work of Thamsborg and Birkedal [28] has extended and adapted it to the operational model, on which we base some of the work presented in this thesis.

## 4 The contents of this thesis.

This thesis aims to extend some of the modular reasoning techniques towards more realistic programming language features. In the following, we present a short overview of the subsequent chapters.

In Chapter 2, **Verifying Object-Oriented Programs with Higher-Order Separation Logic in Coq**, we define a higher-order separation logic

that allows us to give modular specification to object-oriented code that uses dynamic dispatch. We believe this an important development, since dynamic dispatch, typically using *interfaces* is a widespread and powerful feature of modern object-oriented programming languages. Most previous work targeted specific, widespread design patterns, but didn't readily scale to other modes of use that are relatively common when using interfaces. A particular example that we use in Section 2 of that chapter is reasoning about methods that take *any* instance of a given interface as an argument. This technique can be used, for example, to parameterize an algorithm on an underlying collection data structure. We show how, using higher-order separation logic, one can give an abstract specifications to an interface, and a concrete instantiation thereof to verify a class that implements the interface. Thus, specifications of methods that *use* the interface can also be expressed abstractly. Additionally, if these methods in turn are called by a client that knows which particular implementation is being used, the concrete instantiation of the interface can be used to easily regain the information about the state of the object.

The other aspect of our work in Chapter 2 is the *Coq* formalization. The framework that we lay down in Section 3 is, to a large extent, language independent, and can be used to experiment with different programming language features in a sound manner. The framework itself has been extended with tactic support, and eventually became the *Charge!* platform.

In Chapter 3, **Formalized Verification of Snapshotable Trees: Separation and Sharing**, we show how the techniques shown in Chapter 2 scale to realistic, sophisticated data structures. We begin by presenting a family of challenging case studies, the snapshotable trees, derived from a library of collections for *C#*. This data structure is a usual binary search tree in all but one aspects: it also provides a method to take a persistent *snapshot* of the collection in constant time. Since this time bound means we cannot simply copy the tree to build a snapshot, the memory has to be *shared* between the tree and its snapshots. This makes reasoning about the data structure challenging, particularly so if one wants to give it modular specifications. We also provide four implementations of the data structure, varying in terms of complexity of sharing involved — and, thus, the space usage.

We show how using higher-order separation logic one can achieve relatively modular specifications and use our formalization to prove one of the provided implementations, thus showing that our approach to verification of object-oriented programs does indeed scale to complex, realistic data structure.

We tackle an altogether different problem in Chapter 4, **A Separation Logic for Fictional Sequential Consistency**: instead of considering a sophisticated programming language, we turn to underlying memory models. In particular, we consider a simple, class-based, concurrent language running on a total store order (TSO) memory model. In TSO, each thread has an associated *store buffer* — a FIFO queue that buffers writes. Any time a value should be written into the memory, it is instead added to the buffer, and the writes from the buffer are flushed to the memory nondeterministically. This setup admits additional states that would not be observable with “standard” memory mod-



els, and is commonly used to model some real-life processors, notably of the x86 family.

We provide a separation logic for the TSO model, including some novel connectives that allow us to reason about the new states of the memory, but the real strength of our approach is that it allows us to abstract from the complex memory model, and use a much simpler, standard separation logic for a large class of programs. We achieve this by giving programs that utilize the weak behaviors in a restricted way specifications that provide their clients with a *fiction of sequential consistency*, thus bringing standard modular reasoning to much more realistic memory models.

Finally, in Chapter 5, **A Concurrent Logical Relation**, we again consider a different problem: giving a relational account of concurrency in the presence of higher-order store. We build on a logical relation for an ML-like language with higher-order store developed by Thamsborg and Birkedal, and extend it to support concurrency. To the best of our knowledge, this is the first such logical relation published. Since we are using a type-and-effect system to qualify not just the type of the expression, but also the kinds of side effects the program performs, this allows us to reason about correctness of concurrent program transformations — like the ones one might find in an optimizing compiler. One of such program transformations, proposed by Gifford and Lucassen as one of the arguments for introducing the type-and-effect systems, is automatic, effect-guided parallelization. We prove that, under certain conditions, this is indeed a valid optimization, thus showing that our logical relation is both strong and useful.

Additionally, we show how one can use relational reasoning about concurrent programs to show that certain low-level, nonblocking algorithms fulfill their specification, in this case, given by a high-level programming language construct instead of a logical formula.

## 5 Overview of the publications.

Listed below are the papers I have contributed to during the PhD studies. Included is a brief account of the role I have had in the research and writing phase.

1. **Verifying Object-Oriented Programs with Higher-Order Separation Logic in Coq.** Jesper Bengtson, Jonas Braband Jensen, Filip Sieczkowski and Lars Birkedal. In *Proceedings of ITP*, 2011. Forms Chapter 2 of this thesis.

This paper formed the basis for what later became *Charge!* — a sound framework for interactive verification of object-oriented programs, build entirely within the *Coq* proof assistant. My share of the work on this theory was proportional; in particular, the *Coq* formalization was developed jointly by Jesper Bengtson, Jonas B. Jensen and myself, with all of us collaborating closely on all of its features.

In the writing of this paper, my input was concentrated on the sections that concerned formalization of abstract separation logic and its instantiation to the chosen subset of Java. The technical development formalized in *Coq* can be found at [http://itu.dk/people/birkedal/papers/hosl\\_coq-201105.tar.gz](http://itu.dk/people/birkedal/papers/hosl_coq-201105.tar.gz).

2. **Formalized Verification of Snapshotable Trees: Separation and Sharing.** Hannes Mehnert, Filip Sieczkowski, Lars Birkedal and Peter Sestoft. In *Proceedings of VSTTE 2012*. Forms Chapter 3 of this thesis.

This article was the first big case study conducted within the framework presented above, and informed some of the decisions taken with respect to proof automation in *Charge!* Hannes and me contributed the majority of this theory, and most of the formalization work fell to me.

I have taken a smaller part in writing the article, focusing mostly on the treatment of iterators. The technical development formalized in *Coq* can be found at <http://itu.dk/people/hame/snapshots.tar.gz>.

3. **A Concurrent Logical Relation.** Lars Birkedal, Filip Sieczkowski and Jacob Thamsborg. In *Proceedings of CSL 2012*. Forms Chapter 5 of this thesis.

This paper extends the theory developed earlier by Jacob and Lars to a concurrent setting. My contribution here was proportional; in particular I was responsible for proving the parallelization theorem, the main application we provide of our logical relation. Some of the formal proofs are included in the appendix attached to the paper.

4. **A Separation Logic for Fictional Sequential Consistency.** Filip Sieczkowski, Kasper Svendsen and Lars Birkedal. Under submission. Forms Chapter 4 of this thesis.

I have developed this theory with help from Kasper and Lars. My contribution has been major, although Kasper and Lars made sure it evolved towards the principled, two-level structure the logic has in its current form.

The writing has been a joint work, in which I have been responsible for explaining the programming language and the complex TSO logic (Sections 2 and 4). I have also written most of the technical appendix included in this thesis.

5. **A Kripke Logical Relation for Effect-based Program Transformations.** Lars Birkedal, Guilhem Jaber, Filip Sieczkowski and Jacob Thamsborg. Under submission. Not included in this thesis.

This is an extended version of Jacob and Lars's ICFP 2011 article. My contributions to it were centered around extending the theory presented there to support region polymorphism, and preparing the text for the journal submission.

## References

- [1] M. Abadi and G. D. Plotkin. A per model of polymorphism and recursive types. In *Logic in Computer Science, 1990. LICS'90, Proceedings., Fifth Annual IEEE Symposium on*, pages 355–365. IEEE, 1990. 5
- [2] A. Ahmed. *Semantics of types for mutable state*. PhD thesis, Princeton University, 2004. 6
- [3] A. Ahmed. Step-indexed syntactic logical relations for recursive and quantified types. In *Programming Languages and Systems*, pages 69–83. Springer, 2006. 6
- [4] A. Ahmed, D. Dreyer, and A. Rossberg. State-dependent representation independence. In *ACM SIGPLAN Notices*, volume 44, pages 340–353. ACM, 2009. 6
- [5] R. M. Amadio. Recursion over realizability structures. *Information and Computation*, 91(1):55–85, 1991. 5
- [6] A. W. Appel and D. McAllester. An indexed model of recursive types for foundational proof-carrying code. *ACM Trans. Program. Lang. Syst.*, 23(5):657–683, Sept. 2001. 6
- [7] A. W. Appel, P.-A. Melliès, C. D. Richards, and J. Vouillon. A very modal model of a modern, major, general type system. In *Proceedings of the 34th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '07, pages 109–122, New York, NY, USA, 2007. ACM. 3
- [8] N. Benton, A. Kennedy, L. Beringer, and M. Hofmann. Relational semantics for effect-based program transformations with dynamic allocation. In *Proceedings of the 9th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*, PPDP '07, pages 87–96, New York, NY, USA, 2007. ACM. 7
- [9] N. Benton, A. Kennedy, L. Beringer, and M. Hofmann. Relational semantics for effect-based program transformations: Higher-order store. In *Proceedings of the 11th ACM SIGPLAN Conference on Principles and Practice of Declarative Programming*, PPDP '09, pages 301–312, New York, NY, USA, 2009. ACM. 7
- [10] B. Biering, L. Birkedal, and N. Torp-Smith. Bi-hyperdoctrines, higher-order separation logic, and abstraction. *ACM Trans. Program. Lang. Syst.*, 29(5), Aug. 2007. 4
- [11] L. Birkedal, B. Reus, J. Schwinghammer, K. Støvring, J. Thamsborg, and H. Yang. Step-indexed Kripke models over recursive worlds. In *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of*

- Programming Languages*, POPL '11, pages 119–132, New York, NY, USA, 2011. ACM. 6
- [12] L. Birkedal, K. Støvring, and J. Thamsborg. Realizability semantics of parametric polymorphism, general references, and recursive types. In *Foundations of Software Science and Computational Structures*, pages 456–470. Springer, 2009. 5, 6
- [13] L. Birkedal, M. Tofte, and M. Vejlstrup. From region inference to von neumann machines via region representation inference. In *Proceedings of the 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '96, pages 171–183, New York, NY, USA, 1996. ACM. 7
- [14] D. K. Gifford and J. M. Lucassen. Integrating functional and imperative programming. In *Proceedings of the 1986 ACM Conference on LISP and Functional Programming*, LFP '86, pages 28–38, New York, NY, USA, 1986. ACM. 7
- [15] J.-Y. Girard. Une extension de l'interprétation de Gödel à l'analyse, et son application à l'élimination des coupures dans l'analyse et la théorie des types. 2nd Scandinavian Logic Symposium, pages 63–92, Amsterdam, Netherlands, 1971. 5
- [16] C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):576–580, Oct. 1969. 2
- [17] A. Hobor, R. Dockins, and A. W. Appel. A theory of indirection via approximation. In *Proceedings of the 37th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '10, pages 171–184, New York, NY, USA, 2010. ACM. 6
- [18] S. S. Ishtiaq and P. W. O'Hearn. BI as an assertion language for mutable data structures. In *ACM SIGPLAN Notices*, volume 36, pages 14–26. ACM, 2001. 3
- [19] J. B. Jensen. Techniques for model construction in separation logic. 2013. 5
- [20] J. M. Lucassen and D. K. Gifford. Polymorphic effect systems. In *Proceedings of the 15th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '88, pages 47–57, New York, NY, USA, 1988. ACM. 7
- [21] J. C. Mitchell. *Foundations for programming languages*, volume 1. MIT press Cambridge, 1996. 5
- [22] F. Nielson and H. R. Nielson. Type and effect systems. In *Correct System Design*, pages 114–136. Springer, 1999. 7

- [23] M. Parkinson and G. Bierman. Separation logic and abstraction. In *Proceedings of the 32Nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '05, pages 247–258, New York, NY, USA, 2005. ACM. 4
- [24] G. D. Plotkin. *Lambda-definability and logical relations*. School of Artificial Intelligence, University of Edinburgh, 1973. 5
- [25] J. C. Reynolds. Idealized algol and its specification logic. *Tools and notions for program construction*, pages 121–161, 1982. 2, 3
- [26] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Logic in Computer Science, 2002. Proceedings. 17th Annual IEEE Symposium on*, pages 55–74. IEEE, 2002. 3
- [27] W. W. Tait. Intensional interpretations of functionals of finite type i. *The Journal of Symbolic Logic*, 32(2):198–212, 1967. 5
- [28] J. Thamsborg and L. Birkedal. A Kripke logical relation for effect-based program transformations. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming, ICFP '11*, pages 445–456, New York, NY, USA, 2011. ACM. 7
- [29] M. Tofte and J.-P. Talpin. Region-based memory management. *Information and Computation*, 132(2):109–176, 1997. 7
- [30] A. J. Turon, J. Thamsborg, A. Ahmed, L. Birkedal, and D. Dreyer. Logical relations for fine-grained concurrency. In *Proceedings of the 40th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 343–356. ACM, 2013. 6



# Verifying object-oriented programs with higher-order separation logic in Coq

Jesper Bengtson, Jonas Braband Jensen, Filip Sieczkowski, and Lars Birkedal

IT University of Copenhagen

**Abstract.** We present a shallow Coq embedding of a higher-order separation logic with nested triples for an object-oriented programming language. Moreover, we develop novel specification and proof patterns for reasoning in higher-order separation logic with nested triples about programs that use interfaces and interface inheritance. In particular, we show how to use the higher-order features of the Coq formalisation to specify and reason modularly about programs that (1) depend on some unknown code satisfying a specification or that (2) return objects conforming to a certain specification. All of our results have been formally verified in the interactive theorem prover Coq.

## 1 Introduction

Separation Logic [12,16] is a Hoare-style program logic for modular reasoning about programs that use shared mutable data structures. *Higher-order* separation logic [3] (HOSL) is an extension of separation logic that allows for quantification over predicates in both the assertion logic (the logic of pre- and post-conditions) and the specification logic (the logic of Hoare triples). HOSL was proposed with the purposes of (1) reasoning about data abstraction via quantification over resource invariants, and (2) making formalisations of separation logic easier by having one general expressive logic in which it is possible to define predicates, etc., needed for applications. In this article we explore these two purposes further; we discuss each in turn.

The first purpose (data abstraction) has been explored for a first-order language [4], for higher-order languages [9,11], and for reasoning about generics and delegates in object-oriented languages (without interfaces and without inheritance) [18]. In this article we show how HOSL can be used for modular reasoning about interfaces and interface-based inheritance in an object-oriented language like Java or  $C\sharp$ . Our current work is part of a research project in which we aim to formally specify and verify the C5 generic collection library [8], which is an extensive collection library that is used widely in practice and whose implementation makes extensive use of shared mutable data structures. A first case-study of one of the C5 data structures is described in [7]. C5 is written in  $C\sharp$  and is designed mainly using interface inheritance, rather than class-to-class inheritance; different collection modules are related via an inheritance hierarchy among interfaces. For this reason we focus on verifying object-oriented programs that use interfaces and interface-based inheritance.

We explore the second purpose (formalisation) by developing a Coq formalisation of HOSL for an object-oriented class-based language and show through verified examples how it can be used to reason about interfaces and inheritance.

Our formalisation makes use of ideas from abstract separation logic [6] and thus consists of a general treatment of the assertion logic that works for many models and for a general operationally-inspired notion of semantic command. Our general treatment of the logic is also rich enough to cover so-called nested triples [17], which are useful for reasoning about unknown code, either in the form of closures or delegates [18] or, as we show here, in the form of code matching an interface. To reason about object-oriented programs, we instantiate the general development with the heap model for our object-oriented language and derive suitable proof rules for the language. This approach makes it easier in the future to experiment with other storage models and languages, e.g., variants of separation logic with fractional permissions.

*Summary of contributions.* We formalize a shallow Coq embedding of a higher-order separation logic for an object-oriented programming language. We have designed a system that allows us to write programs together with their specifications, and then prove that each program conforms to its specification. All meta-theoretical results have been verified in Coq<sup>1</sup>.

We introduce a pattern for interface specifications that allows for a modular design. An interface specification is parametrised in such a way that any class implementing the interface can be given a suitably expressive specification by a simple instantiation of the interface specification. Moreover, we show how to use nested triples to, e.g., write postconditions in the assertion logic that require a returned object to match a certain specification. Our approach enables us to verify dynamically dispatched method calls, where the dynamic types of the objects are unknown.

*Outline.* The rest of this article is structured as follows. In Section 2 we demonstrate the patterns we use for writing interfaces by providing a small example program that uses interface inheritance and proving that it conforms to its specification. In Section 3 we cover the language and memory-model independent kernel of our Coq formalisation. In Section 4 we specialise our system to handle Java-like programs by providing constructs and a suitable memory model for a subset of Java. Section 5 covers related work, and Section 6 concludes.

## 2 Reasoning with interfaces

To demonstrate how our logic is applied, we will use the example of a class `Cell` that stores a single value and which is extended by a subclass `Recell` that maintains a backup of the last overwritten value and has an `undo` operation. This example is originally due to Abadi and Cardelli [1]; a variant of it was also used

<sup>1</sup> The Coq development accompanying this article can be found at [http://itu.dk/people/birkedal/papers/hosl\\_coq-201105.tar.gz](http://itu.dk/people/birkedal/papers/hosl_coq-201105.tar.gz)



<pre> <b>interface</b> ICell {   <b>int</b> get();   <b>void</b> set(<b>int</b> v); }  <b>class</b> ProxySet {   <b>static void</b> proxySet(ICell c, <b>int</b> v) {     c.set(v);   } }  <b>class</b> Cell <b>implements</b> ICell {   <b>int</b> value;    Cell() {}   <b>int</b> get() {     <b>return</b> <b>this</b>.value;   }   <b>void</b> set(<b>int</b> v) {     <b>this</b>.value = v;   } } </pre>	<pre> <b>interface</b> IRecell <b>extends</b> ICell {   <b>void</b> undo(); }  <b>class</b> Recell <b>implements</b> IRecell {   Cell cell;   <b>int</b> bak;    Recell() {     <b>this</b>.cell = <b>new</b> Cell();   }   <b>int</b> get() {     <b>return</b> <b>this</b>.cell.get();   }   <b>void</b> set(<b>int</b> v) {     <b>this</b>.bak = <b>this</b>.cell.get();     <b>this</b>.cell.set(v);   }   <b>void</b> undo() {     <b>this</b>.cell.set(<b>this</b>.bak);   } } </pre>
---	--

**Fig. 1.** Java code for the Cell-Recell example with interface inheritance.

by Parkinson and Bierman [14] to show how their logic deals with class-to-class inheritance.

We add to this example a method `proxySet`, which calls the `set` method of a given object reference. It is a challenge to give a single specification to this method that is powerful enough to expose any additional side effects the `set` method might have in arbitrary subclasses. We will see in this section how our specification style achieves this, and it is sketched in Section 5 how this compares to related work.

Our model programming language is a subset of both Java and  $C\sharp$ . It leaves out class-to-class inheritance and focuses on interface inheritance. This mode of inheritance captures the essential object-oriented aspect of dynamic dispatch, while the code-reuse aspect has to be explicitly encoded with class composition. A Java implementation of the Cell-Recell example can be found in in Figure 1.

## 2.1 Interface ICell

Interface `ICell` from Figure 1 is modelled as a parametrised specification that states conditions for whether a class  $C$  behaves “Cell-like”. In the following,  $val$  denotes the type of program values, in our case the union of integers, Booleans and object references. Also,  $UPred(heap)$  is the type of logical propositions over heaps, i.e., the spatial component of the assertion logic (see Section 3.1 for the

precise definition).

$$\begin{aligned}
ICell &\triangleq \lambda C : \text{classname}. \quad \lambda T : \text{Type}. \quad \lambda R : \text{val} \rightarrow T \rightarrow \text{UPred}(\text{heap}). \\
&\quad \lambda g : T \rightarrow \text{val}. \quad \lambda s : T \rightarrow \text{val} \rightarrow T. \\
&\quad (\forall t : T. C::\text{get}(\text{this}) \mapsto \{\widehat{R} \text{ this } t\}_{r. \widehat{R} \text{ this } t \wedge r = g \ t}) \wedge & (1) \\
&\quad (\forall t : T. C::\text{set}(\text{this}, x) \mapsto \{\widehat{R} \text{ this } t\}_{\widehat{R} \text{ this } (\widehat{s} \ t \ x)}) \wedge & (2) \\
&\quad (\forall t, v. g \ (s \ t \ v) = v) & (3)
\end{aligned}$$

There is some notation to explain here.  $ICell$  is a function that takes five arguments and returns a result of type  $spec$ , which is the type of specifications. The logical connectives at the outer level ( $\wedge$  and  $\forall$ ) thus belong to the specification logic. The parameter  $R$  is the representation predicate of class  $C$ , so  $R \ c \ t$  intuitively means that  $c$  is a reference to an object that is mathematically modelled by the value  $t$  of type  $T$ . The parameters  $g$  and  $s$  are functions that describe how `get` and `set` inspect and transform this mathematical value. They are constrained by (3) to ensure that `get` will actually return the value set with `set`.

The notation  $C::m(\bar{p}) \mapsto \{P\}_{r. Q}$  from (1) and (2) specifies that method  $m$  of class  $C$  has precondition  $P$  and postcondition  $Q$ . The arguments in a call will be bound to the names  $\bar{p}$  in  $P$  and  $Q$ , and the return value will be bound to  $r$  in  $Q$ . We support both static and dynamic methods, where dynamic methods have an additional first argument, as seen in (1) and (2). The precise definition is given in Section 4.2.

The notation  $f$  from (1) and (2) lifts a function  $f$  such that it operates on expressions, including program variables, rather than operating directly on  $val$ . It is a technical point that can be ignored for a first understanding of this example, but it is crucial for making HOSL work in a stack-based language. Details are in Section 3.2.

The type of  $T$  refers to the  $Type$  universe hierarchy in Coq.

## 2.2 Method proxySet

Consider method `proxySet` from Figure 1. Operationally, calling `proxySet(c, v)` does the same as calling `c.set(v)`, and we seek a specification that reflects this. It is crucial for modularity that `proxySet` can be specified and verified only once and then used with any implementation of `ICell` that may be defined later. We give it the following specification.

$$\begin{aligned}
ProxySet\_spec &\triangleq \forall C, T, R, g, s. \quad ICell \ C \ T \ R \ g \ s \rightarrow \\
&\quad \forall t : T. \quad ProxySet::\text{proxySet}(c, x) \mapsto \{c : C \wedge \widehat{R} \ c \ t\}_{\widehat{R} \ c \ (\widehat{s} \ t \ x)}
\end{aligned}$$

The assertion  $c : C$  means that the object referenced by  $c$  is of class  $C$ . Thus, the caller of `proxySet` can pass in an object reference of any class  $C$  as long as  $C$  can be shown to satisfy  $ICell$ .

This specification is as powerful as that of `set` in  $ICell$  since it essentially forwards it. Any class that behaves Cell-like should be able to encode the behaviour of its `set` method by a suitable choice of  $R$  and  $s$ . We will see in Section 2.6 that it, for instance, is possible to pass in a `Recell` and deduce how `proxySet` affects its backup value.

### 2.3 Class Cell

A Java implementation of `Cell` can be found in Figure 1. We model constructors as static methods that allocate the object before running the initialisation code and return the allocated object, which is what happens in the absence of class-to-class inheritance.

We give class `Cell` the following specification, which is a conjunction of what we will call an *interface specification* and a *class specification*. These correspond respectively to the *dynamic* and *static* specifications in [14].

$Cell\_spec \triangleq \exists R_{Cell}. ICell\ Cell\ val\ R_{Cell}\ (\lambda v. v)\ (\lambda -, v. v) \wedge Cell\_class\ R_{Cell}$   
where

$$\begin{aligned} Cell\_class &\triangleq \lambda R_{Cell} : val \rightarrow val \rightarrow UPred(heap). \\ &Cell::new() \mapsto \{true\} \_ \{\exists v. \widehat{R}_{Cell}\ this\ v\} \wedge \\ &(\forall v. Cell::get(this) \mapsto \{\widehat{R}_{Cell}\ this\ v\} \_ \{r. \widehat{R}_{Cell}\ this\ v \wedge r = v\}) \wedge \\ &(\forall v. Cell::set(this, x) \mapsto \{\widehat{R}_{Cell}\ this\ v\} \_ \{\widehat{R}_{Cell}\ this\ x\}) \end{aligned}$$

The representation predicate  $R_{Cell}$  is quantified such that its definition is visible only while proving the specifications of `Cell`, thus hiding the internal representation of the class from clients [4,13].

It is crucial that  $R_{Cell}$  is quantified outside both the class and the interface specification such that the representation predicate is the same in the two. In practice, a client will allocate a `Cell` by calling `new`, which establishes  $R_{Cell}$ ; later, to model casting the object reference to its interface type, the client knows that `ICell` holds for this same  $R_{Cell}$ .

The specifications of `get` and `set` in `Cell_class` are identical to their counterparts in `ICell` when  $C, T, R, g$ , and  $s$ , are instantiated as in `Cell_spec`. In general, the class specification can be more precise than the interface specification, similarly to the dynamic and static specifications of [14].

To prove `Cell_spec`, the existential  $R_{Cell}$  is chosen as  $\lambda c, v. c.value \mapsto v$ . We can then show that `Cell_class`  $R_{Cell}$  holds by verifying the method bodies of `get`, `set` and `init`, and the correctness of `get` and `set` can be used as a lemma in proving the interface specification. In this way, each method body is verified only once.

### 2.4 Interface IRecell

To show the analogy to interface inheritance at the specification level, we examine an interface for classes that behave `Recell`-like. The Java code for that is `IRecell` in Figure 1. The specification corresponding to this interface follows the same pattern as `ICell`:

$$\begin{aligned} IRecell &\triangleq \lambda C : classname. \quad \lambda T : Type. \quad \lambda R : val \rightarrow T \rightarrow UPred(heap). \\ &\lambda g : T \rightarrow val. \quad \lambda s : T \rightarrow val \rightarrow T. \quad \lambda u : T \rightarrow T. \\ &ICell\ C\ T\ R\ g\ s \wedge & (4) \\ &(\forall t : T. C::undo(this) \mapsto \{\widehat{R}\ this\ t\} \_ \{\widehat{R}\ this\ (u\ t)\}) \wedge & (5) \\ &(\forall t, v. g\ (u\ (s\ t\ v)) = g\ t) & (6) \end{aligned}$$

Notice that interface extension is modelled by referring to *ICell* in (4). We do not have to respecify `get` and `set` since they were already general enough in *ICell* due to it being parametric in  $g$  and  $s$ . Note how equation (6) specifies the abstract behaviour of `undo` via  $g$  and  $s$ .

There is a pattern to how we construct a specification-logic interface predicate from a Java interface declaration. For each method  $m(x_1, \dots, x_n)$ , we add a parameter  $f_m : T \rightarrow val^n \rightarrow (val \times T)$ . The product  $(val \times T)$  can be replaced with just  $val$  or  $T$  if the method should have no side effects or no return value, respectively. We then add a method specification of the form:

$$\forall t : T. C::m(\bar{p}) \mapsto \{\widehat{R} \text{ this } t\} \cdot \{r. \widehat{R} \text{ this } (\pi_2 (\widehat{f}_m \bar{p} t)) \wedge r = \pi_1 (\widehat{f}_m \bar{p} t)\}.$$

## 2.5 Class Recell

The specification of class `Recell` follows the same pattern as with `Cell`:

$$\begin{aligned} \text{Recell\_spec} &\triangleq \exists R_{\text{Recell}} : val \rightarrow val \rightarrow val \rightarrow UPred(\text{heap}). \\ &\quad I\text{Recell } \text{Recell } (val \times val) R g s u \wedge \text{Recell\_class } R_{\text{Recell}} \\ \text{where} \quad R &= \lambda \text{this}, (v, b). R_{\text{Recell}} \text{ this } v b, & g &= \lambda(v, b). v, \\ s &= \lambda(v, b), v'. (v', v), & u &= \lambda(v, b). (b, b), \end{aligned}$$

and *Recell\_class* is defined analogously to *Cell\_class*.

## 2.6 Class World

The correctness of the above specifications only matters if it enables client code to instantiate and use the classes. The client code in `World` demonstrates this:

```
class World {
  static ICell make() {
    Recell r = new Recell();
    r.set(5);
    ProxySet::proxySet(r, 3);
    r.undo();
    return r;
  }
  static void main() {
    ICell c = World::make();
    assert c.get() == 5;
  }
}
```

The body of `make` demonstrates the use of `proxySet`. Operationally, it should be clear that `r` has the value 3 and the backup value 5 after the call to `proxySet`. This can also be proved in our logic despite using a specification of `proxySet` that was verified without knowledge of `Recell` and its backup field.

Upon returning from `make`, we choose to forget that the returned object is really a `Recell`, upcasting it to `ICell`. Its precise class is not needed by the caller, `main`, which only needs to know that the returned object will return 5 from `get`.

We capture the interaction between these two methods with the following specification, in which  $FunI : spec \rightarrow UPred(\text{heap})$  injects the specification logic

into the logic of propositions over heaps, thus generalising the concept of nested triples. Section 3.5 describes *FunI* in more detail.

$$\begin{aligned} \text{World\_spec} &\triangleq \text{World}::\text{main}() \mapsto \{true\}\_-\{true\} \wedge \\ \text{World}::\text{make}() &\mapsto \{true\}\_-\left\{ \begin{array}{l} r. \exists C, T, R, g, s. \widehat{\text{FunI}} (\text{ICell } C \ T \ R \ g \ s) \wedge \\ \exists t. \widehat{R} \ r \ t \wedge g \ t = 5 \wedge r : C \end{array} \right\} \end{aligned}$$

The `make` method is specified to return an object whose class  $C$  is unknown, but we know that  $C$  satisfies *ICell*.

This pattern of returning an object of an unknown type that satisfies a particular specification often comes up in object-oriented programming: think of the method on a collection that returns an iterator, for example. The essence of this pattern is to have a parametrised specification  $S : \text{classname} \rightarrow \text{spec}$  and a method specified as  $D::m() \mapsto \{true\}\_-\{r. \exists C. r : C \wedge \widehat{\text{FunI}} (S \ C)\}$ . A more straightforward alternative to such a specification – one that does not require an embedding of the specification logic in the assertion logic – would be  $\exists C. S \ C \wedge D::m() \mapsto \{true\}\_-\{r. r : C\}$ . However, this restricts the body of  $m$  to only being able to return objects of one class. The method body cannot, for example, choose at run time to return either a  $C_1$  or a  $C_2$ , where both  $C_1$  and  $C_2$  satisfy  $S$ . We find that the most elegant way to allow the method body to make such a choice is to embed the specification in the postcondition.

Using the notion of validity from Definition 5 in Section 3.4 we can now prove that the whole program will behave according to specification:

**Theorem 1.** (*ProxySet\_spec*  $\wedge$  *Cell\_spec*  $\wedge$  *Recell\_spec*  $\wedge$  *World\_spec*) is valid.

### 3 Abstract representation

The core of our system is designed to be language independent. To allow for different memory models, we adopt the notion of separation algebras from Calcagno et al. [6]; we can then instantiate an assertion logic with any separation algebra suitable for the problem at hand. Commands are modelled as relations on the program state, which in turn consists of a mutable stack and a heap. Finally, we define an expressive specification logic that can be used to reason about semantic commands.

We use set-theoretic notation to describe our formalisation as this makes the theories easier to read; in Coq we model these sets as functions into *Prop*, which is the sort of propositions in Coq.

#### 3.1 Uniform predicates

**Definition 1 (Separation algebra).** A separation algebra is a partial, cancellative, commutative monoid  $(\Sigma, \circ, \mathbf{1})$  where  $\Sigma$  is the carrier,  $\circ$  is the monoid operator, and  $\mathbf{1}$  is the unit element.

Intuitively,  $\Sigma$  can be thought of as a type of heaps, and the  $\circ$ -operator as composition of disjoint heaps. Hence we refer to the elements of  $\Sigma$  as heaps. Two heaps are compatible, written  $h_1 \# h_2$  if  $h_1 \circ h_2$  is defined. A heap  $h_1$  is a subheap of a  $h_2$ , written  $h_1 \sqsubseteq h_2$ , if there exists an  $h_3$  such that  $h_2 = h_1 \circ h_3$ . We will commonly refer to a separation algebra by its carrier  $\Sigma$ .

A uniform predicate [5] over a separation algebra is a predicate on heaps and natural numbers; it is upwards closed in the heaps and downwards closed in the natural numbers.

$$UPred(\Sigma) \triangleq \{p \subseteq \Sigma \times \mathbb{N} \mid \forall g, m. \forall h \sqsupseteq g. \forall n \leq m. (g, m) \in p \rightarrow (h, n) \in p\}$$

The upward closure in heaps ensures that we have an intuitionistic separation logic as is desirable for garbage-collected languages.

The natural numbers are used to connect the uniform predicates with the step-indexed specification logic – this connection will be covered in Section 3.5.

We define the standard connectives for the uniform predicates as in [5]:

$$\begin{aligned} true &\triangleq \Sigma \times \mathbb{N} & false &\triangleq \emptyset \\ p \wedge q &\triangleq p \cap q & p \vee q &\triangleq p \cup q \\ \forall x : U. f &\triangleq \bigcap_{x:U} f x & \exists x : U. f &\triangleq \bigcup_{x:U} f x \\ p \rightarrow q &\triangleq \{(h, n) \mid \forall g \sqsupseteq h. \forall m \leq n. (g, m) \in p \rightarrow (g, m) \in q\} \\ p * q &\triangleq \{(h_1 \circ h_2, n) \mid h_1 \# h_2 \wedge (h_1, n) \in p \wedge (h_2, n) \in q\} \\ p \multimap q &\triangleq \{(h, n) \mid \forall m \leq n. \forall h_1 \# h. (h_1, m) \in p \rightarrow (h \circ h_1, m) \in q\} \end{aligned}$$

For the quantifiers,  $U$  is of type *Type*, i.e. the sort of types in Coq, and  $f$  is any Coq function from  $U$  to  $UPred(\Sigma)$ . This allows us to quantify over *any* member of *Type* in Coq.

### 3.2 Stacks

Stacks are functions from variable names to values:  $stack \triangleq var \rightarrow val$ .

Two stacks are said to agree on a set  $V$  of variables if they assign the same value to all members of  $V$ :  $s \simeq_V s' \triangleq \forall x \in V. s x = s' x$ . In order to define operators that take values from the stack as arguments we introduce the notion of a *stack monad*. This approach is similar to that of Varming and Birkedal [20].

$$sm T \triangleq \{(f : stack \rightarrow T, V : \mathcal{P}(var)) \mid \forall s, s'. s \simeq_V s' \rightarrow f s = f s'\}$$

Intuitively,  $V$  is an over-approximation of the free program variables in  $f$ . For any  $m = (f, V) \in sm T$ , we write  $m s$  to mean  $f s$  and  $fv m$  to mean  $V$ .

**Theorem 2.** *sm is a monad with return operation  $\lambda x : T. ((\lambda \_ . x), \emptyset)$  and bind operation  $\lambda m : sm T. \lambda f : T \rightarrow sm U. ((\lambda s. f (m s) s), fv m \cup \bigcup_{t \in T} fv (f t))$ .*

We use the stack monad to model expressions (which can be evaluated to values using data from the stack), pure assertions (that represent logical propositions that are evaluated without using the heap), and assertions (that represent logical propositions that are evaluated using both the heap and the stack).

$$expr \triangleq sm val \quad pure \triangleq sm Prop \quad asn(\Sigma) \triangleq sm UPred(\Sigma)$$

We create an assertion logic by lifting all connectives from  $UPred(\Sigma)$  into  $asn(\Sigma)$ . The definitions and properties of the liftings follow from the fact that  $sm$  is a monad (Theorem 2). We prove that both the uniform predicates and the assertions model separation logic [3].

**Theorem 3.** *For any separation algebra  $\Sigma$ ,  $UPred(\Sigma)$  and  $asn(\Sigma)$  are complete BI-algebras.*

The stack monad is also used for the lifting operator  $\widehat{f}$  that was introduced in Section 2.1. The operator takes a function  $f$ , and returns a function  $\widehat{f}$  where any argument type  $T$  that is passed to  $f$  is replaced with  $sm T$ , and any return type  $U$  with  $sm U$ . As an example, the representation predicate  $R$  in the specification for  $ICell$ , which has type  $val \rightarrow T \rightarrow UPred(heap)$ , is lifted to  $\widehat{R}$  in the assertion-logic formulas of the specification. The resulting type for  $\widehat{R}$  is  $sm val \rightarrow sm T \rightarrow sm UPred(heap)$ , i.e.  $expr \rightarrow sm T \rightarrow asn(heap)$ .

We have to make this lifting explicit in specifications because it restricts how program variables behave under substitution. We have that  $(\widehat{f} e)[e'/x] = \widehat{f}(e[e'/x])$  for any  $f : val \rightarrow UPred(\Sigma)$ , but it is not the case that  $(g e)[e'/x] = g(e[e'/x])$  for any  $g : expr \rightarrow asn(\Sigma)$  because  $g e$  may have more free program variables than those appearing in  $e$ , whereas  $\widehat{f} e$  cannot, by construction. To make HOSL useful in a stack-based language, where such substitutions are commonplace, we therefore typically quantify over functions into  $UPred(\Sigma)$  that we then lift to  $asn(\Sigma)$  where needed.

### 3.3 Semantic commands

To obtain a language-independent core, we model commands as indexed relations on program states (each consisting of a stack and a heap) – a semantic command will relate, in a certain number of steps, a state either to another state or to an error. The only requirements we impose on these commands are that they do not relate to anything in zero steps, and that they satisfy a frame property that will allow us to infer a frame-rule for all semantic commands. Intuitively, the semantic commands can be seen as abstractions of rules of a step-indexed big-step operational semantics. More formally, we have the following definitions.

**Definition 2 (pre-command).** *A pre-command  $\tilde{c}$  relates an initial state to either a terminal state or the special **err** state:*

$$precmd \triangleq \mathcal{P}(stack \times \Sigma \times ((stack \times \Sigma) \uplus \{\mathbf{err}\}) \times \mathbb{N})$$

We write  $(s, h, \tilde{c}) \rightsquigarrow^n x$  to mean that  $(s, h, x, n) \in \tilde{c}$ .

**Definition 3 (Frame property).** *A pre-command  $\tilde{c}$  has the frame property in case the following holds. If  $(s, h_1, \tilde{c}) \not\rightsquigarrow^n \mathbf{err}$  and  $(s, h_1 \circ h_2, \tilde{c}) \rightsquigarrow^n (s', h')$  then there exists  $h'_1$  such that  $h' = h'_1 \circ h_2$  and  $(s, h_1, \tilde{c}) \rightsquigarrow^n (s', h'_1)$ .*

**Definition 4 (Semantic command).** *A semantic command satisfies the frame property and does not evaluate to anything in zero steps.*

$$\text{semcmd} \triangleq \{\hat{c} \in \text{precmd} \mid \hat{c} \text{ has the frame property} \wedge \forall s, h, x. (s, h, \hat{c}) \not\rightsquigarrow^0 x\}$$

To facilitate the encoding of imperative programming languages in our framework, we create the following semantic commands that can be used as building blocks for that purpose. These commands are similar to the ones found in [6].

$$\mathbf{id} \quad \mathbf{seq} \hat{c}_1 \hat{c}_2 \quad \hat{c}_1 + \hat{c}_2 \quad \hat{c}^* \quad \mathbf{assume} P \quad \mathbf{check} P$$

Intuitively, these semantic commands are defined as follows: The **id**-command is the identity command – it does nothing; the **seq**-command executes two commands in sequence; the **+**-operator nondeterministically executes one of two commands; the **\***-command executes a command an arbitrary amount of times; the **assume**-command assumes a pure assertion that can be used to prove correctness of future commands; the **check**-command works like the **id**-command as long as a pure assertion can be inferred. Recall that pure assertions are logical formulas that are evaluated without using the heap.

**Theorem 4.** *id, seq, +, \*, assume, and check are semantic commands.*

### 3.4 Specification logic

With the assertion logic and the semantic commands in place, we can define the specification logic. Semantically, a specification is a downwards-closed set of natural numbers; this allows us to reason about (mutually) recursive programs via step-indexing.

$$\text{spec} \triangleq \{S \subseteq \mathbb{N} \mid \forall m, n. m \leq n \wedge n \in S \rightarrow m \in S\}$$

The set *spec* is a complete Heyting algebra under the subset ordering, i.e., logical entailment ( $\models$ ) is modelled as subset inclusion. Hence a specification  $S$  is *valid* if  $S = \mathbb{N}$ .

Given assertions  $P$  and  $Q$ , and semantic command  $\hat{c}$ , we define a Hoare triple specification:

$$\{P\}\hat{c}\{Q\} \triangleq \{n \mid \forall m \leq n. \forall k \leq m. \forall s, h. (h, m) \in P \rightarrow (s, h, \hat{c}) \not\rightsquigarrow^k \mathbf{err} \wedge \forall h', s'. (s, h, \hat{c}) \rightsquigarrow^k (s', h') \rightarrow (h', m - k) \in Q\}$$

A program is proved correct by proving that its specification is valid:

**Definition 5.** *A specification is valid, written  $\models S$ , when  $\text{true} \models S$ .*

### 3.5 Connecting the assertion logic with the specification logic

We define an embedding of the specification logic into the assertion logic as follows:

$$\text{FunI} : \text{spec} \rightarrow \text{UPred}(\Sigma) \triangleq \lambda S. \Sigma \times S.$$



**Lemma 1.** *FunI is monotone, preserves implication, and has a left and a right adjoint, when spec and UPred( $\Sigma$ ) are treated as poset categories.*

From the second part of this lemma it follows that *FunI* preserves both finite and infinite conjunctions and disjunctions, which entails that all specification logic connectives are preserved by the translation.

### 3.6 Recursion

The specification connectives defined in the previous section are not enough for our purposes. When proving a program correct (by proving a formula of the form  $\models S$ ), it is commonplace that the proof of one part of specification in  $S$  requires other parts of  $S$  – a typical example is recursive method calls, where the specification of the method called must be available in the context during its own verification. To accomplish this, we borrow the *later* operator ( $\triangleright$ ) from Gödel-Löb logic (see [2]).

$$\triangleright S \triangleq \{n + 1 \mid n \in S\} \cup \{0\}$$

This operator can be used via the Löb rule, which allows us to do induction on the step-indexes of the semantic commands.

$$\frac{\Gamma \wedge \triangleright S \models S \quad 0 \in \Gamma \rightarrow 0 \in S}{\Gamma \models S} \text{LÖB}$$

In the inductive case  $\triangleright S$  is found on the left hand side of the turnstile and can hence be used to prove  $S$ .

## 4 Instantiation to an object-oriented language

We define a Java-like language with syntax of programs  $\mathcal{P}$  shown below. The language is untyped and does not need syntax for interfaces; these exist in the specification logic only.

We use a shallow embedding for expressions, which we denote with  $e$ , as shown in Section 3.2.

$$\begin{aligned} \mathcal{P} &::= \mathcal{C}^* && f \in (\text{field names}) \\ \mathcal{C} &::= \text{class } C \ f^* (m(\bar{x})\{c; \text{return } e\})^* \\ c &::= x := \text{alloc } C \mid x := e \mid x := y.f \mid x.f := e \mid x := y.m(\bar{e}) \\ &\mid x := C::m(\bar{e}) \mid \text{skip} \mid c_1; c_2 \mid \text{if } e \text{ then } c_1 \text{ else } c_2 \\ &\mid \text{while } e \text{ do } c \mid \text{assert } e \end{aligned}$$

In order to provide a concrete instance of the assertion logic, we construct a separation algebra of concrete heaps. The carrier set is  $\text{heap} \triangleq (\text{ptr} \times \text{field}) \xrightarrow{\text{fin}} \text{val}$ , with the values defined as the union of integers, Booleans and object references. The partial composition  $h_1 \circ h_2$  is defined as  $h_1 \cup h_2$  if  $\text{dom } h_1 \cap \text{dom } h_2 = \emptyset$ ; otherwise the result is undefined. The unit of the algebra is the empty map,  $\text{emp}$ . We denote this separation algebra  $(\text{heap}, \circ, \text{emp})$  with  $\text{heap}$ . The points-to predicate is defined as  $v.f \mapsto v' \triangleq \{(h, n) \mid h \sqsupseteq [(v, f) \mapsto v']\}$ .

$$\begin{array}{c}
\frac{}{\mathbf{skip} \sim_{\text{sem}} \mathbf{id}} \text{SKIP-SEM} \qquad \frac{c_1 \sim_{\text{sem}} \hat{c}_1 \quad c_2 \sim_{\text{sem}} \hat{c}_2}{c_1; c_2 \sim_{\text{sem}} \mathbf{seq} \hat{c}_1 \hat{c}_2} \text{SEQ-SEM} \\
\\
\frac{c \sim_{\text{sem}} \hat{c}}{\mathbf{while} \ e \ \mathbf{do} \ c \sim_{\text{sem}} \mathbf{seq} (\mathbf{seq} (\mathbf{assume} \ e) \hat{c})^* (\mathbf{assume} \ \neg e)} \text{WHILE-SEM} \\
\\
\frac{c_1 \sim_{\text{sem}} \hat{c}_1 \quad c_2 \sim_{\text{sem}} \hat{c}_2}{\mathbf{if} \ e \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2 \sim_{\text{sem}} (\mathbf{seq} (\mathbf{assume} \ e) \hat{c}_1) + (\mathbf{seq} (\mathbf{assume} \ \neg e) \hat{c}_2)} \text{IF-SEM}
\end{array}$$

**Fig. 2.** The skip, sequential composition, conditional and loop cases of the semantics relation

#### 4.1 Semantics of the programming language

We define the semantics of the programming language commands by relating them to semantic commands instantiated with *heap* as the separation algebra. We write  $c \sim_{\text{sem}} \hat{c}$  to denote that the syntactic command  $c$  is related to the semantic command  $\hat{c}$ . The  $\sim_{\text{sem}}$  relation can be thought of as a function; it is defined as a relation only because this was more straightforward in Coq.

The commands **skip**, **;**, **if**, and **while** can be related directly to composites of the general semantic commands, defined in Section 3.3. The definition of  $\sim_{\text{sem}}$  for these commands can be found in Figure 2. For the remaining commands, new semantic commands must be created.

In particular, for method calls, we define a semantic command

$$\mathbf{call} \ x \ C::m(\bar{e}) \ \mathbf{with} \ c \ \hat{c}$$

that, intuitively, calls method  $m$  of class  $C$  with arguments  $\bar{e}$  and assigns the return value to  $x$ ; the command  $c$  is the method body, and  $\hat{c}$  is its corresponding semantic command. This semantic command works uniformly for both static and dynamic methods, since in the dynamic case we can pass the object reference as an additional argument. The definition of this semantic command is shown in Figure 3. The definition makes use of a predicate

$$C::m(\bar{p})\{c; \mathbf{return} \ r\} \in \mathcal{P}$$

which holds in case method  $m$  in class  $C$  has parameters  $\bar{p}$  and method body  $c$  in program  $\mathcal{P}$ . The program parameter  $\mathcal{P}$  has been left implicit in the other rules. The notation  $[\bar{p} \mapsto (\bar{e} \ s)]$  denotes a finite map that associates each  $p$  in  $\bar{p}$  with the  $e$  at the corresponding position in  $\bar{e}$  evaluated in stack  $s$ .

The requirement that the method body is related to the semantic command is not enforced by the construction of the semantic command, but rather by the definition of  $\sim_{\text{sem}}$  for respectively static and dynamic method calls:

$$\frac{c \sim_{\text{sem}} \hat{c}}{x := C::m(\bar{e}) \sim_{\text{sem}} \mathbf{call} \ x \ C::m(\bar{e}) \ \mathbf{with} \ c \ \hat{c}} \text{SCALL-SEM}$$

$$\begin{array}{c}
 \frac{([\bar{p} \mapsto (\bar{e} s)], h, \hat{c}) \rightsquigarrow^n (s', h') \quad C::m(\bar{p})\{c; \mathbf{return} r\} \in \mathcal{P} \quad |\bar{p}| = |\bar{e}|}{(s, h, \mathbf{call} x C::m(\bar{e}) \mathbf{with} c \hat{c}) \rightsquigarrow^{n+1} (s[x \mapsto (r s')], h')} \text{CALL} \\
 \\
 \frac{C::m(\bar{p})\{c; \mathbf{return} r\} \notin \mathcal{P}}{(s, h, \mathbf{call} x C::m(\bar{e}) \mathbf{with} c \hat{c}) \rightsquigarrow^1 \mathbf{err}} \text{CALL-FAIL1} \\
 \\
 \frac{C::m(\bar{p})\{c; \mathbf{return} r\} \in \mathcal{P} \quad |\bar{p}| \neq |\bar{e}|}{(s, h, \mathbf{call} x C::m(\bar{e}) \mathbf{with} c \hat{c}) \rightsquigarrow^1 \mathbf{err}} \text{CALL-FAIL2} \\
 \\
 \frac{([\bar{p} \mapsto (\bar{e} s)], h, \hat{c}) \rightsquigarrow^n \mathbf{err} \quad C::m(\bar{p})\{c; \mathbf{return} r\} \in \mathcal{P} \quad |\bar{p}| = |\bar{e}|}{(s, h, \mathbf{call} x C::m(\bar{e}) \mathbf{with} c \hat{c}) \rightsquigarrow^{n+1} \mathbf{err}} \text{CALL-FAIL3}
 \end{array}$$

**Fig. 3.** Semantic call commands.

$$\frac{c \sim_{\text{sem}} \hat{c} \quad y : C}{x := y.m(\bar{e}) \sim_{\text{sem}} \mathbf{call} x C::m(y, \bar{e}) \mathbf{with} c \hat{c}} \text{DCALL-SEM}$$

## 4.2 Syntactic Hoare triples and the concrete assertion logic

Hoare triples for syntactic commands are defined in the following manner:

$$\{P\}c\{Q\} \triangleq \forall \hat{c}. c \sim_{\text{sem}} \hat{c} \rightarrow \{P\}\hat{c}\{Q\}.$$

From this definition we infer and prove sound Hoare rules for all commands of our language. To define the rule for method calls we first define the predicate that asserts the specification of methods, introduced in Section 2.1.

$$\begin{aligned}
 C::m(\bar{p}) \mapsto \{P\}_- \{r. Q\} &\triangleq \exists c, e. wf(\bar{p}, r, P, Q, c) \wedge C::m(\bar{p})\{c; \mathbf{return} e\} \in \mathcal{P} \\
 &\wedge \{P\}c\{Q[e/r]\},
 \end{aligned}$$

where  $wf$  is a predicate to assert the following static properties: the method parameter names do not clash; the pre- and postcondition do not use any stack variables other than the method parameters and **this** (the postcondition may also use the return variable); the method body does not modify the values of the method parameters or **this**.

Selected proof rules for syntactic commands are shown in Figure 4. Note the use of the *later* operator ( $\triangleright$ ) in the method call rule; this means that this method call rule will often be used in connection with the Löb rule.

**Theorem 5.** *The rules in Figure 4 are sound with respect to the operational semantics.*

$$\begin{array}{c}
\frac{}{\models \{P\}\mathbf{skip}\{P\}} \text{SKIP} \qquad \frac{}{\models \{P\}c_1\{Q\} \wedge \{Q\}c_2\{R\} \models \{P\}c_1; c_2\{R\}} \text{SEQ} \\
\frac{}{\models \{P \wedge e\}c_1\{Q\} \wedge \{P \wedge \neg e\}c_2\{Q\} \models \{P\}\mathbf{if } e \mathbf{ then } c_1 \mathbf{ else } c_2\{Q\}} \text{IF} \\
\frac{}{\models \{P \wedge e\}c\{P\} \models \{P\}\mathbf{while } e \mathbf{ do } c\{P \wedge \neg e\}} \text{WHILE} \qquad \frac{P \vdash e}{\models \{P\}\mathbf{assert } e\{P\}} \text{ASSERT} \\
\frac{}{\models \{true\}x := \mathbf{alloc } C\{\forall^* f \in \mathit{fields}(C). x.f \mapsto null\}} \text{ALLOC} \\
\frac{}{\models \{P\}x := e\{\exists v. P[v/x] \wedge x = e[v/x]\}} \text{ASSIGN} \qquad \frac{}{\models \{x.f \mapsto \_ \}x.f := e\{x.f \mapsto e\}} \text{WRITE} \\
\frac{P \vdash y.f \mapsto e}{\models \{P\}x := y.f\{\exists v. P[v/x] \wedge x = e[v/x]\}} \text{READ} \\
\frac{\Gamma \models \triangleright C::m(\bar{p}) \mapsto \{P\}\_ \{r. Q\} \quad |\bar{p}| = |y, \bar{e}|}{\Gamma \models \{y : C \wedge P[y, \bar{e}/\bar{p}]\}x := y.m(\bar{e})\{\exists v. Q[x, y[v/x], \bar{e}[v/x]/r, \bar{p}]\}} \text{DCALL} \\
\frac{\Gamma \models \triangleright C::m(\bar{p}) \mapsto \{P\}\_ \{r. Q\} \quad |\bar{p}| = |\bar{e}|}{\Gamma \models \{P[\bar{e}/\bar{p}]\}x := C::m(\bar{e})\{\exists v. Q[x, \bar{e}[v/x]/r, \bar{p}]\}} \text{SCALL} \\
\frac{P \vdash P' \quad Q' \vdash Q}{\models \{P'\}c\{Q'\} \models \{P\}c\{Q\}} \text{CONSEQUENCE} \qquad \frac{\forall x \in \mathit{fv } R. c \text{ does not modify } x}{\models \{P\}c\{Q\} \models \{P * R\}c\{Q * R\}} \text{FRAME} \\
\frac{P \vdash P' \quad Q' \vdash Q \quad \mathit{fv } P \subseteq \mathit{fv } P' \quad \mathit{fv } Q \subseteq \mathit{fv } Q'}{C::m(\bar{p}) \mapsto \{P'\}\_ \{r. Q'\} \models C::m(\bar{p}) \mapsto \{P\}\_ \{r. Q\}} \text{CONSEQUENCE-MSPEC} \\
\frac{\mathit{fv } R \subseteq \{\mathbf{this}\} \cup \bar{p}}{C::m(\bar{p}) \mapsto \{P\}\_ \{r. Q\} \models C::m(\bar{p}) \mapsto \{P * R\}\_ \{r. Q * R\}} \text{FRAME-MSPEC}
\end{array}$$

Fig. 4. Specification logic rules for syntactic Hoare triples

## 5 Related work

Formalisations of higher-order separation logic have been proposed before, e.g. by Varming and Birkedal [20], who developed an Isabelle/HOL formalisation of HOSL for partial correctness for a simple imperative language with first-order mutually recursive procedures, using a denotational semantics of the programming language, and by Preoteasa [15], who developed a PVS formalisation for total correctness using a predicate-transformer semantics for a similar programming language.

Parkinson and Bierman treated an extended version of the Cell-Recell example in [14], improving upon their earlier work in [13]. Their approach is to tailor the specification logic to build in a form of quantification over families of rep-

representation predicates following a fixed pattern determined by the inheritance tree of the program. This construction is known as *abstract predicate families* (APFs).

Where our logic allows quantification over a representation type  $T$ , as used in Section 2.1, APFs have a built-in notion of variable-arity predicates to achieve same effect: representation predicates of a subclass can add parameters to the representation predicate they inherit. Class `Cell` defines a two-parameter representation predicate family `Val`, which is extended to three arguments in `Recell`. A `Recell`  $r$  having value 2 and backup field 1 would be asserted as  $Val(r, 2, 1)$ . This assertion implies  $Val(r, 2)$ , which in turn implies  $\exists b. Val(r, 2, b)$  if it is known that  $r$  is a `Recell`. Thus, casting to the two-argument representation predicate that would be necessary for calling  $\{\exists v. Val(c, v)\} \text{proxySet}(c, x) \{Val(c, x)\}$  will lose any information about the backup field.

The logic of Parkinson and Bierman was extended by van Staden and Calcagno [19] to handle multiple inheritance, abstract classes and controlled leaking of facts about the abstract representation of either a single class or a class hierarchy. Using the latter feature, we observe that their logic can also be used to reason about the example in Section 2, by using parameters  $g$  and  $s$  to give a precise specification of `proxySet`. Instead of being functions,  $g$  and  $s$  would be abstract predicate families whose first argument would be an object reference used only for selecting the correct member of the APF.

Compared to the logics based on abstract predicate families, our logic allows families of not just predicates but also types, functions, class names or any other type that can be quantified over in Coq. This gives us strong typing of logical variables, and all this works without building it into the logic and requiring that quantifications and proofs follow the shape of the inheritance tree.

## 6 Conclusion and Future Work

We have presented a Coq implementation of a generic framework for higher-order separation logic. In this framework, instantiated with a simple object-oriented language, we have shown how HOSL can be used to reason about interfaces and interface inheritance.

Future work includes developing better support for automation via better use of tactics. Our Coq proofs of example programs are cluttered with manual reordering of the context because we do not yet have tactics to automate this. We also plan to integrate the current tool with an Eclipse front-end that is currently being researched within our project [10]. Moreover, we plan to use the tool for formal verification of interesting data structures from the C5 collection library.

Although it is not necessary for the code we mostly want to verify, proper support for class-to-class inheritance in both the logic and the design pattern would enable more direct comparison with related work. It would also make our Java subset more similar to actual Java.

*Acknowledgements* We are grateful for useful discussions with Hannes Mehnert, Matthew Parkinson, Peter Sestoft, Kasper Svendsen, and Jacob Thamsborg.

This work was supported in part by the ToMeSo project, funded by the Danish Research Council.

## References

1. M. Abadi and L. Cardelli. *A Theory of Objects*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1st edition, 1996.
2. A. W. Appel, P.-A. Melliès, C. D. Richards, and J. Vouillon. A very modal model of a modern, major, general type system. In *Proceedings of POPL*, 2007.
3. B. Biering, L. Birkedal, and N. Torp-Smith. BI hyperdoctrines and higher-order separation logic. In *Proceedings of ESOP*, pages 233–247, 2005.
4. B. Biering, L. Birkedal, and N. Torp-Smith. BI-hyperdoctrines, higher-order separation logic, and abstraction. *ACM Trans. Program. Lang. Syst.*, 29(5), 2007.
5. L. Birkedal, B. Reus, J. Schwinghammer, K. Støvring, J. Thamsborg, and H. Yang. Step-indexed kripke models over recursive worlds. In *Proceedings of POPL*, 2011.
6. C. Calcagno, P. W. O’Hearn, and H. Yang. Local action and abstract separation logic. In *Proceedings of LICS*, pages 366–378, 2007.
7. J. Jensen, L. Birkedal, and P. Sestoft. Modular verification of linked lists with views via separation logic. *Journal of Object Technology*, 2011. To Appear. Preliminary version in FTfJP’10, available at [www.itu.dk/people/birkedal/papers/views.pdf](http://www.itu.dk/people/birkedal/papers/views.pdf).
8. N. Kokholm and P. Sestoft. The C5 generic collection library for C# and CLI. Technical Report ITU-TR-2006-76, IT University of Copenhagen, 2006.
9. N. R. Krishnaswami, J. Aldrich, L. Birkedal, K. Svendsen, and A. Buisse. Design patterns in separation logic. In *In Proceedings of TLDI*, pages 105–116, 2009.
10. H. Mehnert. Kopitiam: modular incremental interactive full functional static verification of java code. In M. Bobaru, K. Havelund, G. Holzmann, and R. Joshi, editors, *Proceedings of the Third NASA Formal Methods Symposium (NFM 2011)*. NASA, April 2011.
11. A. Nanevski, A. Ahmed, G. Morrisett, and L. Birkedal. Abstract predicates and mutable ADTs in hoare type theory. In *In Proc. of ESOP*, pages 189–204, 2007.
12. P. W. O’Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *Proceedings of CSL*, pages 1–19, 2001.
13. M. J. Parkinson and G. M. Bierman. Separation logic and abstraction. In *Proceedings of POPL*, pages 247–258, 2005.
14. M. J. Parkinson and G. M. Bierman. Separation logic, abstraction and inheritance. In *Proceedings of POPL*, pages 75–86, 2008.
15. V. Preoteasa. Frame rules for mutually recursive procedures manipulating pointers. *Theoretical Computer Science*, 2009.
16. J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proceedings of LICS*, pages 55–74, 2002.
17. J. Schwinghammer, L. Birkedal, B. Reus, and H. Yang. Nested Hoare triples and frame rules for higher-order store. In *Proceedings of CSL*, 2009.
18. K. Svendsen, L. Birkedal, and M. Parkinson. Verifying generics and delegates. In *Proceedings of ECOOP*, pages 175–199, 2010.
19. S. van Staden and C. Calcagno. Reasoning about multiple related abstractions with multistar. In *Proceedings of OOPSLA*, pages 504–519, 2010.
20. C. Varming and L. Birkedal. Higher-order separation logic in Isabelle/HOLCF. *Electr. Notes Theor. Comput. Sci.*, 218:371–389, 2008.

# Formalized Verification of Snapshotable Trees: Separation and Sharing

Hannes Mehnert, Filip Sieczkowski, Lars Birkedal, and Peter Sestoft

IT University of Copenhagen  
{hame, fisi, birkedal, sestoft}@itu.dk

**Abstract.** We use separation logic to specify and verify a Java program that implements snapshotable search trees, fully formalizing the specification and verification in the Coq proof assistant. We achieve *local* and *modular* reasoning about a tree and its snapshots and their iterators, although the implementation involves shared mutable heap data structures with no separation or ownership relation between the various data. The paper also introduces a series of four increasingly sophisticated implementations and verifies the first one. The others are included as future work and as a set of challenge problems for full functional specification and verification, whether by separation logic or by other formalisms.

## 1 Introduction

This paper presents a family of realistic but compact challenge case studies for modular software verification. We fully specified and verified the first case study in Coq, using a domain-specific separation logic [10] and building upon our higher-order separation logic [2]. As future work we plan to verify the other implementations with the presented abstract interface specification. We believe this is the first mechanical formalization of this approach to modular reasoning about implementations that use shared heap data with no separation or ownership relation between the various data.

The family of case studies consists of a single interface specification for snapshotable trees, and four different implementations. A *snapshotable tree* is an ordered binary tree that represents a set of items and supports taking readonly *snapshots* of the set, in constant time, at the expense of slightly slower subsequent updates to the tree. A snapshotable tree also supports iteration (enumeration) over its items as do, e.g., the Java collection classes. The four implementations of the snapshotable tree interface all involve shared heap data as well as increasingly subtle uses of destructive heap update.

For practical purposes it is important that the same interface specification can support verification of multiple implementations with varying degrees of internal sharing and destructive update. Moreover, the specification must accommodate any number of data structure (tree) instances, each having any number of iterators and snapshots, each of which in turn can have any number of iterators. Most importantly, we show how we can have local reasoning (a frame rule) even though the tree and its snapshots share mutable heap data.

We welcome other solutions to the specification and verification of this case study; indeed R. Leino has already made one (unpublished) using Dafny [11].

The Java source code of the case studies of all four implementations and the Coq source is available at <http://www.itu.dk/people/hame/snapshots.tgz>.

Section 2 presents the interface of the case study data structure, shows an example use, and outlines four implementations. Section 3 gives a formal specification of the interface using separation logic and verifies the example code. Sections 4 and 5 verify the first implementation.

## 2 Case Study: Snapshotable Trees

The case study is a simplified version of snapshotable treesets from the C5 collection library [8].

### 2.1 Interface: Operations on Snapshotable Trees

Conceptually, a snapshot of a treeset is a readonly copy of the treeset. Subsequent updates to the tree do not affect any of its snapshots, so one can update the tree while iterating over a snapshot. Taking a snapshot must be a constant time operation, but subsequent updates to the tree may be slower after a snapshot has been taken. Implementations (Section 2.3) typically achieve this by making the tree and its snapshots share parts of their representation, gradually unsharing it as the tree gets updated, in a manner somewhat analogous to copy-on-write memory management schemes in operating systems.

All tree and snapshot implementations implement the same `ITree` interface:

```
public interface ITree extends Iterable<Integer> {
    public boolean contains(int x);
    public boolean add(int x);
    public ITree snapshot();
    public Iterator<Integer> iterator();
}
```

These operations have the following effect:

- `tree.contains(x)` returns true if the item is in the tree, otherwise false.
- `tree.add(x)` adds the item to the tree and returns true if the item was not already in the tree; otherwise does nothing and returns false.
- `tree.snapshot()` returns a readonly snapshot of the given tree. Updates to the given tree will not affect the snapshot. A snapshot cannot be made from a snapshot.
- `tree.iterator()` returns an iterator (also called enumerator, or stream) of the tree's items. Any number of iterators on a tree or snapshot may exist at the same time. Modifying a tree will invalidate all iterators on that tree (but not on its snapshots), so that the next operation on such an iterator will throw `ConcurrentModificationException`.



We include the somewhat complicated `iterator()` operation because it makes the distinction between a tree and its snapshots completely clear: While it is illegal to modify a tree while iterating over it, it is perfectly legal to modify the tree while iterating over one of its snapshots. Also, this poses an additional verification challenge when considering implementations with rebalancing (cases A2B1 and A2B2 in Section 2.3) because `tree.add(item)` may rebalance the tree in the middle of an iteration over a snapshot of the tree, and that should be legal and not affect the iteration.

Note that for simplicity, items are here taken to be integers; using techniques from [20] it is straightforward to extend our formal specification and verification to handle a generic version of snapshotable trees.

## 2.2 Example Client Code

To show what can be done with snapshots and iterators (and not without), consider this piece of client code. It creates a treeset `t`, adds three items to it, creates a snapshot `s` of the tree, and then iterates over the snapshot’s three items while adding new items (6 and 9) to the tree:

```
ITree t = new Tree();
t.add(2); t.add(1); t.add(3);
ITree s = t.snapshot();
Iterator<Integer> it = s.iterator();
boolean lc = it.hasNext();
while (lc) {
    int x = it.next();
    t.add(x * 3);
    lc = it.hasNext();
}
```

## 2.3 Implementations of Snapshotable Trees

One may consider four implementations of treesets, spanned by two orthogonal implementation features. First, the tree may be unbalanced (A1) or it may be actively rebalanced (A2) to keep depth  $O(\log n)$ . Second, snapshots may be kept persistent, that is, unaffected by tree updates, either by path copy persistence (B1) or by node copy persistence (B2):

	Without rebalancing	With rebalancing
Path copy persistence	A1B1	A2B1
Node copy persistence	A1B2	A2B2

The implementation closest to that of the C5 library [8, section 13.10] is A2B2, which is still somewhat simplified: only integer items, no comparer argument, no update events, and so on. In this paper we formalize and verify only implementation A1B1; the verification of the more sophisticated implementations A1B2, A2B1 and A2B2 will be addressed in future work.

Nevertheless, for completeness and in the hope that others may consider this verification challenge, we briefly discuss all four implementations and the expected verification challenges here.

With *path copy persistence* (cases AxB1), adding an item to a tree will duplicate the path from the root to the added node, if this is necessary to avoid modifying any snapshot of the tree. Thus an update will create  $O(d)$  new nodes where  $d$  is the depth of the tree.

With *node copy persistence* (cases AxB2), each tree node has a spare child reference. The first update to a node uses this spare reference, does not copy the node and does not update its parent; the node remains shared between the tree and its snapshots. Only the second update to a node copies it and updates its parent. Thus an update does not replicate the entire path to the tree root; the number of new nodes per update is amortized  $O(1)$ . See Driscoll [6] or [8].

To implement ordered trees without rebalancing (cases A1By), we use a Node class containing an item (here an integer) and left and right children; `null` is used to indicate the absence of a child. A tree or snapshot contains a stamp (indicating the “time” of the most recent update) and a reference to the root Node object; `null` if the tree is empty.

To implement rebalancing of trees (cases A2By), we use left-leaning red-black trees (LLRB) which encode 2-3 trees [1, 19], instead of general red-black trees [7] as in the C5 library. This reduces the number of rebalancing cases.

To implement iterators on a tree or snapshot we use a class `TreeIterator` that holds a reference to the underlying tree, a stamp (the creation “time” of the iterator) and a stack of nodes. The stamp is used to detect subsequent updates to the underlying tree, which will invalidate the iterator. Since snapshots cannot be updated, their iterators are never invalidated. The iterator’s stack holds its current state: for each node in the stack, the node’s own item and all items in the right subtree have yet to be output by the iterator.

*Case A1B1 = no rebalancing, path copy persistence* In this implementation there is shared data between a tree and its snapshots, but the shared data is not being mutated because the entire path from the root to an added node gets replicated. Hence no node reachable from the root of a snapshot, or from nodes in its iterators’ stacks, can be affected by an update to the live tree; therefore no operation on a snapshot can be affected by operations on the live tree. Although this case is therefore the simplest case, it already contains many challenges in finding a suitable specification for trees, snapshots and iterators, and in proving the stack-based iterator implementation correct.

*Case A2B1 = rebalancing, path copy persistence* In this case there is potential mutation of shared data, because the rebalancing rotations seem to be able to affect nodes just off the fresh unshared path from a newly added node to the root. This could adversely affect an iterator of a snapshot because a reference from the iterator’s node stack might have its right child updated (by a rotation), thus wrongly outputting the items of its right subtree twice or not at all. However, this does not happen because the receiver of a rotation (to be moved down) is

always a fresh node (we’re in case B1 = path copy persistence) and moreover we consider only `add` operations (not `remove`), so the child being rotated (moved up) is also a fresh node and thus not on the stack of any iterator – the rebalancing was caused by this child being “too deep” in the tree. Hence if we were to support `remove` as well, then perhaps the implementation of rotations needs to be refined.

*Case A1B2 = no rebalancing, node copy persistence* In this case, there is mutation of shared data not observable by the client. For example, a left-child update to a tree node that is also part of a snapshot will move the snapshot’s left-child value to the node’s extra reference field, and destructively update the left child as required for the live tree. There should be no observable change to the snapshot, despite the change to the data representing it. The basic reason for correctness is that any snapshot involving an updated node will use the extra reference and hence not see the update; this is true for nodes reachable from the root of a snapshot as well as for nodes reachable from the stack of an iterator. When we need to update a node whose extra reference is already in use, we leave the old node alone and create a fresh copy of the node for use in the live tree; again, existing snapshots and their iterators do not see the update.

*Case A2B2 = rebalancing, node copy persistence* In this case there is mutation of shared data (due both to moving child fields to the extra reference in nodes, and due to rotations), not observable for the client. Since the updates caused by rotations are handled exactly like other updates, the correctness of rebalancing with respect to iterators seems to be more straightforward than in case A2B1.

### 3 Abstract Specification and Client Code Verification

We use higher-order separation logic [18, 3] to specify and verify the snapshotable tree data structure. We build on top of our intuitionistic formalization of HOSL in Coq [2] with semantics for an untyped Java-like language.

To allow implementations to share data between a tree, its snapshots, and iterators and still make it possible for clients to reason locally (to focus only on a single tree / snapshot / iterator), we will use an idea from [10] (see also the verification of Union-Find in [9]). The idea is to introduce an abstract predicate, here named  $H$ , global to each tree data structure consisting of a single tree, multiple snapshots, and multiple iterators. This abstract predicate  $H$  is parameterized by a finite set of disjoint *abstract structures*. We have three kinds of abstract structures: Tree, Snap, and Iter. The use of  $H$  enables a client of our specification to consider each abstract structure to be separate or disjoint from the rest of the abstract structures and thus the client can reason modularly about client code using only those abstract structures she needs; the rest can be framed out. Since the abstract predicate  $H$  is existentially quantified, the client has no knowledge of how an implementation defines  $H$  (see [3, 16] for more on abstract predicates in higher-order separation logic). The implementor

of the tree data structure has a global view on the tree with its snapshots and iterators, and is able to define which parts of the abstract structures are shared in the concrete heap. Section 4 defines  $H$  for the A1B1 case from Section 2.3.

The Tree abstract structure consists of a handle (reference) to the tree and a model, which is an ordered finite set, containing the elements of the tree. The Snap structure is similar to Tree. The Iter structure consists of a handle to the iterator and a model, which is a list containing the remaining elements for iteration. Because  $H$  is tree-global, exactly one Tree structure must be present (“the tree”), while the number of Snap and Iter structures is not constrained.

### 3.1 Specification of the ITree Interface

We now present the formal abstract specification of the ITree interface informally described in Section 2.1. The specification also contains five axioms, which are useful for a client and obligations to an implementor of the interface. The specification is parametrized over an implementation class  $C$  and the above-mentioned predicate  $H$ , and each method specification is universally quantified over the model  $\tau$ , a finite set of integers and a finite set of abstract structures  $\phi$ .

```
interface ITree {
  {H({Tree(this,  $\tau$ )}  $\uplus$   $\phi$ )} contains(x) {ret = x  $\in$   $\tau$   $\wedge$  H({Tree(this,  $\tau$ )}  $\uplus$   $\phi$ )}
  {H({Snap(this,  $\tau$ )}  $\uplus$   $\phi$ )} contains(x) {ret = x  $\in$   $\tau$   $\wedge$  H({Snap(this,  $\tau$ )}  $\uplus$   $\phi$ )}
  {H({Tree(this,  $\tau$ )}  $\uplus$   $\phi$ )} add(x)      {ret = x  $\notin$   $\tau$   $\wedge$  H({Tree(this, {x}  $\cup$   $\tau$ )}  $\uplus$   $\phi$ )}
  {H({Tree(this,  $\tau$ )}  $\uplus$   $\phi$ )} snapshot() {H({Snap(ret,  $\tau$ )}  $\uplus$  {Tree(this,  $\tau$ )}  $\uplus$   $\phi$ )}
  {H({Snap(this,  $\tau$ )}  $\uplus$   $\phi$ )} iterator() {H({Iter(ret, [ $\tau$ ])}  $\uplus$  {Snap(this,  $\tau$ )}  $\uplus$   $\phi$ )  $\wedge$ 
                                         ret <: Iterator}
}
(a) H({Tree(t,  $\tau$ )}  $\uplus$   $\phi$ )  $\vdash$  t : C
(b) H({Snap(s,  $\tau$ )}  $\uplus$   $\phi$ )  $\vdash$  s : C
(c)  $\tau = \tau' \wedge H({Tree(t, \tau)} \uplus \phi) \vdash H({Tree(t, \tau')} \uplus \phi)$ 
(d) H({Snap(s,  $\tau$ )}  $\uplus$   $\phi$ )  $\vdash$  H( $\phi$ )
(e) H({Iter(it,  $\alpha$ )}  $\uplus$   $\phi$ )  $\vdash$  H( $\phi$ )
}
```

These specifications can be read as follows:

- **contains** requires either a Snap or Tree structure (written as separate specifications) for the **this** handle and some set  $\tau$ . The structure is unmodified in the postcondition, and the return value **ret** is true if the item **x** is in the set  $\tau$ , otherwise false.
- **add** requires a Tree structure for the **this** handle and some set  $\tau$ . The postcondition states that the given item **x** is added to the set  $\tau$ . The return value indicates whether the tree was modified, which is the case if the item was not already present in the set  $\tau$ .
- **snapshot** requires a Tree structure for the **this** handle and some set  $\tau$ . The postcondition constructs a Snap structure for the returned handle **ret** and the set  $\tau$ . So the Tree and the Snap structure contain the same elements.
- **iterator** requires a Snap structure for the **this** handle and some set  $\tau$ . The postcondition constructs an Iter structure with the return handle and the set  $\tau$  converted to an ordered list, written  $[\tau]$ . The returned handle conforms (written  $<:$ ) to the Iterator specification shown in Section 3.2.

The five axioms state that (a) the static type of the tree is the given class  $C$ ; (b) the static type of a snapshot is  $C$ ; (c) the model  $\tau$  of the tree can be replaced by an equal model  $\tau'^{-1}$ ; and we can forget about snapshots (d) and iterators (e).

In contrast to the description in Section 2.1 we leave iterators over the tree for future work. We could use the ramification operator [10] to express that any iterators over the tree become invalid when the tree is modified.

The abstract separation can be observed, e.g., in the specification of `add`: it only modifies the model of the Tree structure and does not affect the rest of the abstract structures ( $\phi$  is preserved in the postcondition). Hence the client can reason about calls to `add` locally, independently of how many snapshots and iterators there are.

In our Coq formalization we do not have any syntax for interfaces at the specification logic level [2], but represent interfaces using Coq-level definitions. Appendix A contains the formal representations (ITree, Iterator, Stack).

### 3.2 Iterator Specification

Our iterator specification is also parametrized over a class  $IC$  and a predicate  $H$ , and each method specification is universally quantified over a list of integers  $\alpha$  and a finite set of abstract structures  $\phi$ .

```
interface Iterator<Integer> {
  {H({Iter(this,  $\alpha$ )}  $\uplus$   $\phi$ )}   hasNext() {ret = ( $|\alpha| \neq 0$ )  $\wedge$  H({Iter(this,  $\alpha$ )}  $\uplus$   $\phi$ )}
  {H({Iter(this,  $x :: \alpha$ )}  $\uplus$   $\phi$ )} next()   {ret =  $x \wedge$  H({Iter(this,  $\alpha$ )}  $\uplus$   $\phi$ )}
}
```

The specification of the Iterator interface requires an `Iter` structure with the `this` handle and some list  $\alpha$ . The return value of the method `hasNext` captures whether the list  $\alpha$  is non-empty. The `Iter` structure in the postcondition is not modified. The method `next` requires an `Iter` structure with a non-empty list ( $x :: \alpha$ ). The list head is returned and the model of the `Iter` structure is updated to the remainder of the list.

### 3.3 Client Code Verification

To verify the client code from Section 2.2 we assume we are given a class  $C$  such that `ITree C H` holds for some  $H$  and then verify the client code under the precondition  $\{H(\{Tree(t, \{\})\})\}$ .

Figure 1 gives a step-by-step proof of the client code from Section 2.2, with client code lines to the left and their postconditions to the right.

After inserting some items (line 1) to the tree, the model contains these items,  $\{1, 2, 3\}$ . In line 2, a snapshot `s` of the tree `t` is created. The invariant  $H$  now consists of the Tree structure and a Snap structure containing the same elements. For the client the abstract structures are disjoint, but in an implementation, they will be realized using sharing. Indeed, for the A1B1 implementation, the concrete

<sup>1</sup> This is explicit for technical reasons: in our implementation  $H$  is defined inside a monad [2], and the client should not have to discharge obligations inside the monad.

```

1: t.add(2);t.add(1);t.add(3); {H({Tree(t, {1, 2, 3})})}
2: ITree s = t.snapshot(); {H({Tree(t, {1, 2, 3})} ⊔ {Snap(s, {1, 2, 3})})}
3: Iterator<Integer> it = {H({Tree(t, {1, 2, 3})} ⊔ {Snap(s, {1, 2, 3})} ⊔
  s.iterator());
  {Iter(it, [1, 2, 3])}}
4: boolean lc = {lc = true ∧ H({Tree(t, {1, 2, 3})} ⊔
  it.hasNext()); {Snap(s, {1, 2, 3})} ⊔ {Iter(it, [1, 2, 3])}}
5: while (lc) { invariant: ∃α, β. α@β = [1, 2, 3] ∧ lc = (|β| ≠
  0) ∧ H({Tree(t, {1, 2, 3} ∪ {3z|z ∈ α})}) ⊔
  {Snap(s, {1, 2, 3})} ⊔ {Iter(it, β)}
6: int x = it.next(); {α@β = [1, 2, 3] ∧ lc = (|β| ≠ 0) ∧ β =
  x :: β' ∧ H({Tree(t, {1, 2, 3} ∪ {3z|z ∈ α})}) ⊔
  {Snap(s, {1, 2, 3})} ⊔ {Iter(it, β')}}
7: t.add(x * 3); {α@β = [1, 2, 3] ∧ lc = (|β| ≠ 0) ∧ β = x ::
  β' ∧ H({Tree(t, {1, 2, 3} ∪ {3z|z ∈ α} ∪ {3x})}) ⊔
  {Snap(s, {1, 2, 3})} ⊔ {Iter(it, β')}}
8: lc = it.hasNext(); {α@β = [1, 2, 3] ∧ lc = (|β'| ≠ 0) ∧ β = x ::
  β' ∧ H({Tree(t, {1, 2, 3} ∪ {3z|z ∈ α} ∪ {3x})}) ⊔
  {Snap(s, {1, 2, 3})} ⊔ {Iter(it, β')}}
9: } {H({Tree(t, {1, 2, 3, 6, 9})} ⊔ {Snap(s, {1, 2, 3})})}

```

Fig. 1. Client code verification

heap will be as shown in Figure 2, where all the nodes are shared between the tree and the snapshot.

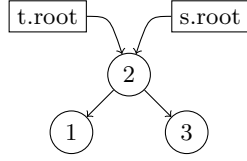
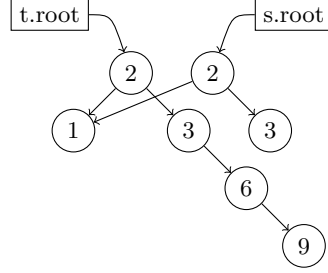
In line 3 an iterator `it` over the snapshot `s` is created. To apply the call rule of the `iterator` method, only the `Snap` structure is taken into account, the rest (the `Tree` structure) is framed out inside of  $H$  (via appropriate instantiation of  $\phi$  in the `iterator` specification). The result is that an `Iter` structure is constructed, whose model contains the same values as the model of the snapshot, but converted to an ordered list. We introduce the loop condition `lc` in line 4, and again use abstract framing to call `hasNext`.

Lines 5–9 contain a while loop with loop condition `lc`. The loop invariant splits the iteration list  $[1, 2, 3]$  into the list  $\alpha$  containing the elements already iterated over and the list  $\beta$  containing the remainder. The loop variable `lc` is false iff  $\beta$  is the empty list. The invariant  $H$  contains the `Tree` structure whose model is the initial set  $\{1, 2, 3\}$  joined with the set of the elements of  $\alpha$ , each multiplied by 3.  $H$  also contains the `Iter` and the `Snap` structures.

We omit detailed explanation of the remaining lines of verification.

Note that in the final postcondition, the client sees two disjoint structures (axiom (e) is used to forget the empty iterator), but in the A1B1 implementation, the concrete heap will involve sharing, as shown in Figure 3. Only the left subtree is shared by the tree and the snapshot; the root and right subtree were unshared by the first call to `add` in the loop.

In summary, we have shown the following theorem, which says that given any  $H$  and any classes  $C$  and  $IC$  satisfying the `ITree` and `Iterator` interface


**Fig. 2.** Heap after snapshot construction

**Fig. 3.** Live heap after loop

specifications, the client code satisfies its specification. The postcondition states that snapshot  $\mathbf{s}$  contains 1, 2 and 3, and tree  $\mathbf{t}$  contains additionally 6 and 9.

**Theorem 1.**  $\forall H.\forall C.\forall IC.ITree\ C\ H \wedge Iterator\ IC\ H$   
 $\vdash \{H(\{Tree(\mathbf{t}, \{\})\})\} client\_code \{H(\{Tree(\mathbf{t}, \{1, 2, 3, 6, 9\})\}) \uplus \{Snap(\mathbf{s}, \{1, 2, 3\})\}\}$

## 4 Implementation A1B1

In this section we show the partial correctness verification of the A1B1 implementation with respect to the abstract specification from the previous section. This involves defining a concrete  $H$  and showing that the methods satisfy the required specifications for this concrete  $H$ . The development has been formally verified in Coq (as has the client program verification above).

The Coq formalization uses a shallow embedding of higher-order separation logic, developed for verification of OO programs using interfaces. See [2].

Invariant  $H$  is radically different depending on whether snapshots of the tree are present or not. The reason is that method `add` mutates the existing tree if there are no snapshots present, see Section 5 for details. Here we focus on the case where snapshots are present.

The A1B1Tree class stores its data in three fields: the `root` node, a boolean field `isSnapshot`, indicating whether it is a snapshot, and a field `hasSnapshot`, indicating whether it has snapshots. The stamp field mentioned in Section 2.3 is only required for iterators over the tree and so not further discussed here.

The Node class is a nested class of the A1B1Tree with three fields, `item` containing its value, and a handle to the right (`right`) and left (`left`) subtree.

In the following we use standard separation logic connectives, in particular the separating conjunction  $*$  and the points to predicate  $\mapsto$ .

We now define our concrete  $H$  and also the realization of the abstract structures. We first explain the realization of Tree and Snap; the Iter structure is described in Section 4.1. Recall that  $\phi$  ranges over finite sets of abstract structures (Tree, Snap, Iter), with exactly one Tree structure, and recall that  $H$ , given a  $\phi$ , returns a separation logic predicate. The definition of  $H$  is:

$$H(\phi) \triangleq \exists \sigma. wf(\sigma) \wedge heap(\sigma) * \sigma \models \phi$$

Here  $\sigma$  is a finite map of type  $\text{ptr} \rightarrow \text{ptr} \times \mathbb{Z} \times \text{ptr}$ , with  $\text{ptr}$  being the type of Java pointers (handles), corresponding to the `Node` class. The map  $\sigma$  must be well-formed (pure predicate  $wf(\sigma)$ ), which simply means that all pointers in the codomain of  $\sigma$  are either `null` or in the domain of  $\sigma$ .

The *heap* function maps  $\sigma$  to a separation logic predicate, which describes the realization of  $\sigma$  as a linked structure in the concrete heap, starting with  $\top$ :

$$\text{heap}(\sigma) \triangleq \text{fold} (\lambda p n Q. \text{match } n \text{ with } (\text{pl}, v, \text{pr}) \Rightarrow \\ \text{p.left} \mapsto \text{pl} * \text{p.item} \mapsto v * \text{p.right} \mapsto \text{pr} * Q) \top \sigma$$

Finally, we present the definition of  $\sigma \models \phi$  (we defer the definition of  $\sigma \models \{ \text{Iter}(-, -) \}$  to the following subsection):

$$\begin{aligned} \sigma \models \phi \uplus \psi &\triangleq \sigma \models \phi * \sigma \models \psi \\ \sigma \models \{ \text{Tree}(\text{ptr}, \tau) \} &\triangleq \exists p. \text{Node}(\sigma, p, \tau) \wedge \text{ptr.root} \mapsto p * \\ &\quad \text{ptr.isSnapshot} \mapsto \text{false} * \text{ptr.hasSnapshot} \mapsto \text{true} \\ \sigma \models \{ \text{Snap}(\text{ptr}, \tau) \} &\triangleq \exists p. \text{Node}(\sigma, p, \tau) \wedge \text{ptr.root} \mapsto p * \\ &\quad \text{ptr.isSnapshot} \mapsto \text{true} * \text{ptr.hasSnapshot} \mapsto \text{false} \end{aligned}$$

The spatial structure of all the nodes is covered by *heap*( $\sigma$ ) so  $\sigma \models \phi$  just needs to describe the additional heap taken up by *Tree*, *Snap*, and *Iter* structures.

The pure *Node* predicate is defined inductively on  $\tau$  below. It is used to express that  $\tau$  is the finite set of items reachable from  $p$  in  $\sigma$ .

$$\begin{aligned} \text{Node}(\sigma, p, \tau) &\triangleq (p = \text{null} \wedge \tau = \{\}) \vee \\ &\quad (p \in \text{dom}(\sigma) \wedge \exists \text{pl}, v, \text{pr}. \sigma[p] = (\text{pl}, v, \text{pr}) \wedge \\ &\quad \exists \tau_l, \tau_r. \tau = \tau_l \cup \{v\} \cup \tau_r \wedge \\ &\quad (\forall x \in \tau_l. x < v) \wedge (\forall x \in \tau_r. x > v) \wedge \\ &\quad \text{Node}(\sigma, \text{pl}, \tau_l) \wedge \text{Node}(\sigma, \text{pr}, \tau_r)) \end{aligned}$$

The sortedness constraint (a strict total order) in the *Node* predicate enforces implicitly that  $\sigma$  has the right shape:  $\sigma$  cannot contain cycles and the left and right subtrees must be disjoint. The set  $\tau$  is split into three sets, one with strictly smaller elements ( $\tau_l$ ), the singleton  $v$  and with strictly bigger elements ( $\tau_r$ ).

#### 4.1 Iterator

The *TreeIterator* class implements the *Iterator* interface. It contains a single field, `context`, which is a stack of *Node* objects.

The constructor of the *TreeIterator* pushes all nodes on the leftmost path of the tree onto the stack. The method `next` pops the top node from the stack and returns the value held in that node. Before returning, it pushes the leftmost path of the node's right subtree (if any) onto the stack. The method `hasNext` returns true if and only if the stack is empty.

The verification of the iterator depends on the following specification of a stack class, generic in  $C$ . The specification is parametrized over a representation type  $T$  and existentially over a representation predicate  $SR$  (of type  $\text{classname} \rightarrow$



$(\text{val} \rightarrow T \rightarrow \text{HeapAsn}) \rightarrow \text{val} \rightarrow T^* \rightarrow \text{HeapAsn}$ ). The second argument is the predicate  $P$  (of type  $\text{val} \rightarrow T \rightarrow \text{HeapAsn}$ ), which holds for every stack element. This specification is kept in the style of [17], although we use a different logic.

```

class Stack<C> {
  ⊤
  SR C P this α          new() SR C P ret nil
  SR C P this α * P x t ∧ x : C empty() ret = (α = nil) ∧ SR C P this α
  SR C P this (t :: α)    push(x) SR C P this (t :: α)
  SR C P this (t :: α)    pop() P ret t * SR C P this α
  SR C P this (t :: α)    peek() P ret t *(∀u.P ret u -* SR C P this (u :: α))
(a) P v t ⊢ P' v t ⇒ SR C P v α ⊢ SR C P' v α
}
    
```

For the purpose of specifying the iterators over snapshotable trees, we instantiate the type  $T$  with  $\mathbb{Z}^*$ ; the model of a node on the stack is a list of integers. Intuitively, this list corresponds to the node value and the element list of its right subtree. The iterator is modelled as a list that is equal to the concatenation of the elements of the stack. We also require that the topmost element of the stack is nonempty (if present). This intuition is formalized in the interpretation of the *Iter* structure, where  $SR$  is a representation predicate of a stack:

$$\sigma \models \{Iter(\mathbf{p}, \alpha)\} \triangleq \exists \text{st. } \mathbf{p}.\text{context} \mapsto \text{st} * \exists \beta. \text{stack\_inv}(\beta, \alpha) \wedge SR \text{ Node } (NS \sigma) \text{ st } \beta.$$

To make this definition complete, we provide the definitions of  $stack\_inv$ , which connects the representation of the stack with the representation of the iterator, and the definition of the  $NS$  predicate.

$$stack\_inv(xss, ys) \triangleq ys = \text{concat}(xss) \wedge \begin{cases} \top & \text{iff } xss = \text{nil} \\ xs \neq \text{nil} & \text{iff } xss = xs :: xss' \end{cases}$$

$$NS \sigma \text{ node } \alpha \triangleq Node(\sigma, \text{node}, \tau) \wedge \alpha = [\{x \in \tau \mid x \geq \text{node.item}\}]$$

These definitions, along with an assumption that  $SR$  is the representation predicate of *Stack* (i.e., fulfills all the method specifications and axioms of *Stack\_spec*) suffice to show the correctness of *Iter*-dependent methods. The axiom present in *Stack\_spec* is needed to preserve iterators if some new memory is added to  $\sigma$ : it allows us to replace  $(NS \sigma)$  with  $(NS \sigma')$  as a representation predicate of stack objects under certain side conditions.

## 5 On the Verification of Implemented Code

We now give an intuitive description of how the A1B1 implementation was verified, given the concrete  $H$  defined above. We verified the complete implementation in Coq but only discuss the `add` method here. We used Kopitiam [13] to transform the Java code into SimpleJava, the fragment represented in Coq.

Method `add` calls method `addRecursive` with the root node to insert the item into the binary tree, respecting the ordering. Method `addRecursive`, shown below, must handle several cases:

- if there are no snapshots present, then
  - if the item  $x$  is already in the tree, then the heap is not modified.
  - if the item  $x$  is not in the tree, then a new node is allocated and destructively inserted into the tree.
- if there are snapshots present, then
  - if the item  $x$  is already in the tree, then the heap is not modified.
  - if the item  $x$  is not in the tree, then a new node is allocated and every node on the path from the root to the added node is replicated, so that the snapshots are unimpaired.

The implementation of `addRecursive` walks down the tree until a node with the same value, or a leaf, is reached. It uses the call stack to remember the path in the tree. If a node was added, either the entire path from the root to the added node is duplicated (if snapshots are present) or the handles to the left or right subtree are updated (happens destructively exactly once, the parent of the added node updates its left or right handle, previously pointing to `null`):

```
Node addRecursive (Node node, int item, RefBool updated) {
  Node res = node;
  if (node == null) {
    updated.value = true;
    res = new Node(item);
  } else {
    if (item < node.item) {
      Node newLeft = addRecursive(node.left, item, updated);
      if (updated.value && this.hasSnapshot)
        res = new Node(newLeft, node.item, node.rght);
      else
        node.left = newLeft;
    } else if (node.item < item) {
      Node newRght = addRecursive(node.rght, item, updated);
      if (updated.value && this.hasSnapshot)
        res = new Node(node.left, node.item, newRght);
      else
        node.rght = newRght;
    } //else item == node.item so no update
  }
  return res;
}
```

We now show the pre- and postcondition of `addRecursive` for the two cases where snapshots are present. If the item is already present in the tree, the pre- and postcondition are equal:

$$\begin{aligned} & \{\text{updated.value} \mapsto \text{false} * \text{this.hasSnapshot} \mapsto \text{true} * \\ & \quad \text{heap}(\sigma) * \text{wf}(\sigma) \wedge \text{Node}(\sigma, \text{node}, \tau) \wedge \text{item} \in \tau\} \\ & \quad \text{addRecursive}(\text{node}, \text{item}, \text{updated}) \\ & \{\text{updated.value} \mapsto \text{false} * \text{this.hasSnapshot} \mapsto \text{true} * \\ & \quad \text{heap}(\sigma) * \text{ret} = \text{node}\} \end{aligned}$$

The postcondition in the case that the item is added to the tree extends the map  $\sigma$  to  $\sigma'$ , for which the heap layout and the well-formedness condition must hold. The Node predicate uses  $\sigma'$  and the finite set is extended with `item`:

$$\begin{aligned} & \{\text{updated.value} \mapsto \text{false} * \text{this.hasSnapshot} \mapsto \text{true} * \\ & \quad \text{heap}(\sigma) * \text{wf}(\sigma) \wedge \text{Node}(\sigma, \text{node}, \tau) \wedge \text{item} \notin \tau\} \\ & \quad \text{addRecursive}(\text{node}, \text{item}, \text{updated}) \\ & \{\text{updated.value} \mapsto \text{true} * \text{this.hasSnapshot} \mapsto \text{true} * \\ & \quad \exists \sigma'. \sigma \subseteq \sigma' \wedge \text{heap}(\sigma') * \text{wf}(\sigma') \wedge \text{Node}(\sigma', \text{ret}, \{\text{item}\} \cup \tau)\} \end{aligned}$$

The call to `addRecursive` inside of `add` is verified for each specification of `addRecursive` independently.

To summarize Sections 4 and 5, we state the following theorem, which says there exists an  $H$  that given a stack fulfilling the stack specification, the `TreeIterator` class meets the `Iterator` specification and the `A1B1` class meets the `ITree` specification, and the constructor for the `A1B1Tree` establishes the  $H$  predicate.

**Theorem 2.**  $\exists H. \text{Stack\_spec} \vdash \text{Iterator } \text{TreeIterator } H \wedge \text{ITree } \text{A1B1 } H \wedge \{\top\} \text{A1B1Tree}() \{H(\{\text{Tree}(\text{ret}, \{\})\})\}$

The client code, independently verified, can be safely linked with the `A1B1` implementation!

## 6 Related Work

Malecha and Morrisett [12] presented a formalization of a Ynot implementation of B-trees with an iterator method. In their case, the iterator and the tree also share data in the concrete heap. However, they can only reason about “single-threaded” uses of trees and iterators: their specification of the iterator method transforms the abstract tree predicate into an abstract iterator predicate, which prohibits calling tree methods until the client turns the iterator back into a tree. In our setup, we have one tree, but multiple snapshots and iterators, and the tree can be updated after an iterator has been created. To permit sharing between a tree and an iterator, Malecha and Morrisett use fractional permissions, where we use the  $H$  predicate. They work in an axiomatic extension of Coq, whereas our proofs are done in a shallowly embedded program logic, since our programs are written in an existing programming language (Java).

Dinsdale-Young et al. [5] present another approach to reasoning about shared data structures, which gives the client a fiction of disjointness. Roughly speaking, they define a new abstract program logic for each module (they can be combined) for abstract client reasoning. Their approach allows one to give a client specification similar to ours, but without using the  $H$  and with the abstract structures (`Tree` / `Snap` / `Iter`) being predicates in the (abstract) program logic. This has the advantage that one can use ordinary framing for local reasoning.

Dinsdale-Young et al. [4] also presented an approach to reasoning about sharing. Sharing can happen in certain regions, and the module implementor

has to define a protocol that describes how data in the shared region can evolve. What corresponds to our abstract structures can now be seen as separation logic predicates and thus one can use ordinary framing for local reasoning.

In both approaches [5] and [4] the module implementor has more proof obligations than in our approach: In [5] he must show that the abstract operations satisfy some conditions related to the realization of the abstract structures in the concrete heap. In [4] she must show related properties phrased in terms of certain stability conditions.

Compared to the work of Dinsdale-Young et al., our approach has the advantage that it is arguably simpler, with no need to introduce new separation (or context) algebras for the modules. That is why we could build our formalization on an implementation of standard separation logic in Coq.

Rustan Leino made a solution for a custom implementation of this data structure (A1B1) using Dafny [11]. Dafny verifies that if a snapshot is present, the nodes are shared and not mutated by the tree operations. His solution does not (yet) verify the content of the tree, snapshots or iterators. Our verification specifies the concrete heap layout. Dafny does not support abstract specification due to the lack of inheritance. The trusted code base is different: Dafny relies on Boogie, Z3 and the CLR, whereas our proof trusts Coq.

## 7 Conclusion and Future Work

We have presented snapshotable trees as a challenge for formalized modular reasoning about mutable data structures that use sharing extensively, and given an abstract specification of the ITree interface. Moreover, we have presented a formalization of the A1B1 implementation of snapshotable trees.

The overall size of the formalization effort is roughly 5000 lines of Coq code and it takes 2 hours to `Qed` the proofs. This is quite big compared to other formalization efforts of imperative programs in Coq, such as Hoare Type Theory / Ynot [14, 15]. The main reason is that we are working in a shallowly embedded program logic for a Java-like language, whereas Hoare Type Theory / Ynot is an axiomatic extension of Coq. Thus our formalization includes both the operational semantics of the Java subset and the soundness theorems for the program logic; also, Java program variables cannot simply be represented by Coq variables.

We also plan to verify the even subtler implementations A1B2, A2B1 and A2B2, which are expected to provide further insight into the challenges of dealing with shared mutable data and unobservable state changes. Through those more complex applications of separation logic we hope to learn more about desirable tool support, including how to automate the “obvious” reasoning that currently requires much thought and excessive amounts of proof code. Although we have not formally verified these implementations yet, we are fairly certain they would match the interface specification presented in Section 3. In all four implementations the tree is conceptually separate from its snapshots, which is the property required by the interface, and the invariant  $H$  allows us to describe the heap layout very precisely, using techniques shown in Section 4.

Finally, we would like to explore how to combine the advantages of our approach and those of Dinsdale-Young’s approach discussed above.

## References

1. A. Andersson. Balanced search trees made simple. In F. Dehne et al., editors, *Algorithms and Data Structures. LNCS 709*, pages 60–71. Springer-Verlag, 1993.
2. J. Bengtson, J. B. Jensen, F. Sieczkowski, and L. Birkedal. Verifying object-oriented programs with higher-order separation logic in Coq. In *ITP 2011*, 2011.
3. B. Biering, L. Birkedal, and N. Torp-Smith. BI-hyperdoctrines, higher-order separation logic, and abstraction. *ACM Trans. Program. Lang. Syst.*, 29(5), 2007.
4. T. Dinsdale-Young, M. Dodds, P. Gardner, M. Parkinson, and V. Vafeiadis. Concurrent abstract predicates. In T. DHondt, editor, *ECOOP 2010*, volume 6183 of *LNCS*, pages 504–528. Springer Berlin / Heidelberg, 2010.
5. T. Dinsdale-Young, P. Gardner, and M. Wheelhouse. Abstraction and refinement for local reasoning. In *VSTTE*, pages 199–215, Berlin, Heidelberg, 2010. Springer.
6. J. Driscoll, N. Sarnak, D. Sleator, and R. Tarjan. Making data structures persistent. *Journal of Computer and Systems Sciences*, 38(1):86–124, 1989.
7. L. Guibas and R. Sedgwick. A dichromatic framework for balanced trees. In *19th FCS, Ann Arbor, Michigan*, pages 8–21, 1978.
8. N. Kokholm and P. Sestoft. The C5 Generic Collection Library for C# and CLI. Technical Report ITU-TR-2006-76, IT University of Copenhagen, January 2006.
9. N. Krishnaswami. *Verifying Higher-Order Imperative Programs with Higher-Order Separation Logic*. PhD thesis, Carnegie Mellon University, 2011. Forthcoming.
10. N. R. Krishnaswami, L. Birkedal, and J. Aldrich. Verifying event-driven programs using ramified frame properties. In *TLDI*, pages 63–76. ACM, 2010.
11. K. R. M. Leino. Dafny: An automatic program verifier for functional correctness. In E. M. Clarke and A. Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning. LNCS 6355*, pages 348–370, 2010.
12. G. Malecha and G. Morrisett. Mechanized verification with sharing. In *7th International Colloquium on Theoretical Aspects of Computing*, Sept. 2010.
13. H. Mehnert. Kopitiam: Modular incremental interactive full functional static verification of java code. In M. Bobaru, K. Havelund, G. Holzmann, and R. Joshi, editors, *NASA Formal Methods*, volume 6617, pages 518–524. Springer, 2011.
14. A. Nanevski, G. Morrisett, A. Shinnar, P. Govereau, and L. Birkedal. Ynot: dependent types for imperative programs. In J. Hook and P. Thiemann, editors, *Proc. of 13th ACM ICFP 2008*, pages 229–240. ACM, 2008.
15. A. Nanevski, V. Vafeiadis, and J. Berdine. Structuring the verification of heap-manipulating programs. In *Proceedings of POPL*, 2010.
16. M. J. Parkinson and G. M. Bierman. Separation logic and abstraction. In *Proceedings of POPL*, pages 247–258, 2005.
17. R. Petersen, L. Birkedal, A. Nanevski, and G. Morrisett. A realizability model for impredicative hoare type theory. In S. Drossopoulou, editor, *ESOP 2008*, volume 4960 of *Lecture Notes in Computer Science*, pages 337–352. Springer, 2008.
18. J. C. Reynolds. Separation logic: A logic for shared mutable data structures. *IEEE Proc. of 17th Symp. on Logic in CS*, Nov 2002.
19. R. Sedgwick. Left-leaning red-black trees. At <http://www.cs.princeton.edu/~rs/talks/LLRB/LLRB.pdf>.
20. K. Svendsen, L. Birkedal, and M. Parkinson. Verifying generics and delegates. In *ECOOP’10*, pages 175–199. Springer-Verlag, 2010.

## A Appendix

We define here the *ITree* and the *Iterator* interface specification as Coq definitions, as well as the *Stack* class. We use the name *SPred* for the finite set of abstract structures containing exactly one *Tree* structure and any number of *Snap* and *Iter* structures.

The notation  $\widehat{f}$  lifts the function  $f$  such that it operates on expressions rather than values.

A detailed explanation of the notation and of lifting can be found in [2].

$$\begin{aligned}
ITree &\triangleq \lambda C : \text{classname}. \lambda H : \mathcal{P}_{fin}(SPred) \rightarrow \text{HeapAsn}. \\
&(\forall \tau : \mathcal{P}_{fin}(\mathbb{Z}). \forall \phi : \mathcal{P}_{fin}(SPred). C::\text{contains}(\text{this}, \mathbf{x}) \mapsto \\
&\quad \{\widehat{H}(\{\widehat{\text{Tree}}(\text{this}, \tau)\} \uplus \phi)\}_{-}\{\mathbf{r}. \widehat{H}(\{\widehat{\text{Tree}}(\text{this}, \tau)\} \uplus \phi) \wedge \mathbf{r} = (\mathbf{x} \in \tau)\}) \\
&\wedge (\forall \tau : \mathcal{P}_{fin}(\mathbb{Z}). \forall \phi : \mathcal{P}_{fin}(SPred). C::\text{contains}(\text{this}, \mathbf{x}) \mapsto \\
&\quad \{\widehat{H}(\{\widehat{\text{Snap}}(\text{this}, \tau)\} \uplus \phi)\}_{-}\{\mathbf{r}. \widehat{H}(\{\widehat{\text{Snap}}(\text{this}, \tau)\} \uplus \phi) \wedge \mathbf{r} = (\mathbf{x} \in \tau)\}) \\
&\wedge (\forall \tau : \mathcal{P}_{fin}(\mathbb{Z}). \forall \phi : \mathcal{P}_{fin}(SPred). C::\text{add}(\text{this}, \mathbf{x}) \mapsto \\
&\quad \{\widehat{H}(\{\widehat{\text{Tree}}(\text{this}, \tau)\} \uplus \phi)\}_{-}\{\mathbf{r}. \widehat{H}(\{\widehat{\text{Tree}}(\text{this}, \{\mathbf{x}\} \cup \tau)\} \uplus \phi) \wedge \mathbf{r} = (\mathbf{x} \notin \tau)\}) \\
&\wedge (\forall \tau : \mathcal{P}_{fin}(\mathbb{Z}). \forall \phi : \mathcal{P}_{fin}(SPred). C::\text{snapshot}(\text{this}) \mapsto \\
&\quad \{\widehat{H}(\{\widehat{\text{Tree}}(\text{this}, \tau)\} \uplus \phi)\}_{-}\{\mathbf{r}. \widehat{H}(\{\widehat{\text{Tree}}(\text{this}, \tau), \widehat{\text{Snap}}(\mathbf{r}, \tau)\} \uplus \phi)\}) \\
&\wedge (\forall \tau : \mathcal{P}_{fin}(\mathbb{Z}). \forall \phi : \mathcal{P}_{fin}(SPred). C::\text{iterator}(\text{this}) \mapsto \\
&\quad \{\widehat{H}(\{\widehat{\text{Snap}}(\text{this}, \tau)\} \uplus \phi)\}_{-}\{\mathbf{r}. \exists IC : \text{classname}. \text{Iterator } IC \ H \wedge \mathbf{r} : IC \wedge \\
&\quad \widehat{H}(\{\widehat{\text{Snap}}(\text{this}, \tau), \widehat{\text{Iter}}(\mathbf{r}, [\tau])\} \uplus \phi)\}) \\
&\wedge (\forall v : \text{val}. \forall \tau : \mathcal{P}_{fin}(\mathbb{Z}). \forall \phi : \mathcal{P}_{fin}(SPred). \\
&\quad (H(\{\text{Tree}(v, \tau)\} \uplus \phi) \implies v : C) \wedge (H(\{\text{Snap}(v, \tau)\} \uplus \phi) \implies v : C)) \\
&\wedge (\forall v : \text{val}. \forall \tau : \mathcal{P}_{fin}(\mathbb{Z}). \forall \phi : \mathcal{P}_{fin}(SPred). \\
&\quad (H(\{\text{Snap}(v, \tau)\} \uplus \phi) \vdash H(\phi))) \\
&\wedge (\forall v : \text{val}. \forall \alpha : \mathbb{Z}^*. \forall \phi : \mathcal{P}_{fin}(SPred). \\
&\quad (H(\{\text{Iter}(v, \alpha)\} \uplus \phi) \vdash H(\phi))) \\
&\wedge (\forall v : \text{val}. \forall \tau, \tau' : \mathcal{P}_{fin}(\mathbb{Z}). \forall \phi : \mathcal{P}_{fin}(SPred). \\
&\quad \tau = \tau' \implies (H(\{\text{Tree}(v, \tau)\} \uplus \phi) \vdash H(\{\text{Tree}(v, \tau')\} \uplus \phi)))
\end{aligned}$$

$$\begin{aligned}
Iterator &\triangleq \lambda C : \text{classname}. \lambda H : \mathcal{P}_{fin}(SPred) \rightarrow \text{HeapAsn}. \\
&(\forall \alpha : \mathbb{Z}^*. \forall \phi : \mathcal{P}_{fin}(SPred). C::\text{hasNext}(\text{this}) \mapsto \\
&\quad \{\widehat{H}(\{\widehat{\text{Iter}}(\text{this}, \alpha)\} \uplus \phi)\}_{-}\{\mathbf{r}. \widehat{H}(\{\widehat{\text{Iter}}(\text{this}, \alpha)\} \uplus \phi) \wedge \mathbf{r} = (\alpha \neq \text{nil})\}) \\
&\wedge (\forall x : \mathbb{Z}. \forall \alpha : \mathbb{Z}^*. \forall \phi : \mathcal{P}_{fin}(SPred). C::\text{next}(\text{this}) \mapsto \\
&\quad \{\widehat{H}(\{\widehat{\text{Iter}}(\text{this}, x::\alpha)\} \uplus \phi)\}_{-}\{\mathbf{r}. \widehat{H}(\{\widehat{\text{Iter}}(\text{this}, \alpha)\} \uplus \phi) \wedge \mathbf{r} = x\})
\end{aligned}$$

$Stack\_spec \triangleq \forall T : Type.$

$\exists SR : classname \rightarrow (val \rightarrow T \rightarrow HeapAsn) \rightarrow val \rightarrow T^* \rightarrow HeapAsn.$

$(\forall C : classname. \forall P : val \rightarrow T \rightarrow HeapAsn.$

$\quad Stack::new() \mapsto \{\top\}_{-}\{r. \widehat{SR} C P r \text{ nil}\}$

$\wedge (\forall \alpha : T^*. Stack::empty(\text{this}) \mapsto$

$\quad \{\widehat{SR} C P \text{ this } \alpha\}_{-}\{r. \widehat{SR} C P \text{ this } \alpha \wedge r = (\alpha = \text{nil})\})$

$\wedge (\forall \alpha : T^*. \forall t : T. Stack::push(\text{this}, x) \mapsto$

$\quad \{\widehat{SR} C P \text{ this } \alpha * \widehat{P} x t \wedge x : C\}_{-}\{\widehat{SR} C P \text{ this } (t :: \alpha)\})$

$\wedge (\forall \alpha : T^*. \forall t : T. Stack::pop(\text{this}, x) \mapsto$

$\quad \{\widehat{SR} C P \text{ this } (t :: \alpha)\}_{-}\{r. \widehat{P} r t * \widehat{SR} C P \text{ this } \alpha\})$

$\wedge (\forall \alpha : T^*. \forall t : T. Stack::peek(\text{this}, x) \mapsto$

$\quad \{\widehat{SR} C P \text{ this } (t :: \alpha)\}_{-}\{r. \widehat{P} r t^*$

$\quad (\forall u : T. \widehat{P} r u \rightarrow \widehat{SR} C P \text{ this } (u :: \alpha))\})$

$\wedge (\forall C : classname. \forall P, P' : val \rightarrow T \rightarrow HeapAsn.$

$\quad (\forall v : val. \forall t : T. (P v t \vdash P' v t)) \implies$

$\quad \forall v : val. \forall \alpha : T^*. (SR C P v \alpha \vdash SR C P' v \alpha))$





# A Separation Logic for Fictional Sequential Consistency

Filip Sieczkowski  
IT University of Copenhagen  
fisi@itu.dk

Kasper Svendsen  
Aarhus University  
ksvendsen@cs.au.dk

Lars Birkedal  
Aarhus University  
birkedal@cs.au.dk

## Abstract

Modern multiprocessors and mainstream concurrent programming languages do not provide the strong *sequentially consistent* shared memory that has been assumed by most work on program logics for concurrent programs. Rather, they typically implement *weak* memory models and the programmer can observe reorderings of memory operations caused by the weak behavior. In particular, programmers have to reason about weak behavior when implementing low-level efficient algorithms, e.g., synchronization primitives. However, in practice, low-level algorithms are often implemented in such a way that clients do not have to reason about weak behavior; instead clients can pretend that they are working with a sequentially consistent model.

In this paper we present a new program logic, iCAP-TSO, for a TSO memory model. The logic supports formal reasoning at both levels of abstraction that programmers use informally. Our logic can be used to verify efficient low-level implementations of libraries against specifications that provide clients with a *fiction of sequential consistency*. Hence client-side reasoning can be done as in earlier program logics that assumed a sequentially consistent memory model.

## 1. Introduction

**What?** Modern multiprocessors and mainstream concurrent programming languages do not provide the strong *sequentially consistent* (SC) shared memory that has been assumed by most work on program logics for concurrent programs. Rather, they typically implement *weak* memory models with observable buffering and reorderings of memory operations. In the case of the Total Store Order (TSO) memory model, each thread is connected to main memory via a FIFO store buffer that buffers writes before committing them to main memory.

It is possible to extend existing program logics to a setting with a TSO memory model by reasoning explicitly about these store buffers. Ridge [14] defines a Rely-Guarantee-based program logic for reasoning about x86-TSO assembly code and Wehrman [18] defines a separation logic for a low-level language with a TSO memory model. Both of these logics can be used to reason about TSO programs that exhibit *racy* behavior. Importantly, large classes of programs (e.g., well-synchronized programs) only exhibit sequentially consistent behavior. Intuitively, for most programs, explicit reasoning about store buffers is thus unnecessary and if we are

forced to reason about store buffers, we will never be able to *scale* our techniques to *realistic* systems. Ridge’s and Wehrman’s logics both suffer from this scalability issue, by requiring explicit reasoning about store buffers even for well-behaved clients. The main challenge is thus to develop a program logic that supports *simple* high-level reasoning for well-behaved clients *and* low-level reasoning about *racy* library implementations.

In this paper we present a new program logic, iCAP-TSO, for a TSO memory model. This logic features *two* higher-order separation logics, a TSO logic for reasoning about *racy* low-level code and an SC logic with *standard* separation logic assertions and proof rules for reasoning about well-behaved client code. Crucially, our logic supports *simple* reasoning about programs that only exhibit sequentially consistent behaviors despite using libraries with *racy* implementations. Intuitively, these libraries provide clients with a *fiction of sequential consistency* and our logic exploits this to simplify reasoning about clients.

**How?** In the TSO memory model, each thread is connected to main memory via a FIFO store buffer, modeled as a sequence of (address, value) pairs, see, e.g., [15]. When a value is written to an address, the write is recorded in the writing thread’s store buffer. Threads can commit these buffered writes to main memory at any point in time. When reading from a location, threads first consult their own store buffer and only consult main memory if their own store buffer does not contain a buffered write for the given location. If it does, the thread reads the value of the last buffered write in its own store buffer. Each thread thus has its own *subjective* view of the current state of memory, which might differ from other threads.

In contrast, in a sequentially consistent memory model, threads read and write directly to main memory and thus share an *objective* view of the current state of the memory. In separation logics for languages with sequentially consistent memory models we thus use assertions such as  $x \mapsto v$ , which express an *objective* property about the value of location  $x$ . Since in the TSO setting each thread has a subjective view of the state, in order to preserve the standard proof rules for reading and writing, we need a *subjective* interpretation of pre- and postconditions.

In our SC logic we thus have specifications of the form

$$[P] e [r.Q],$$

which express that if  $e$  is executed by thread  $t$  from an initial state that satisfies  $P$  from the point of view of  $t$  and  $e$  terminates with value  $v$ , then the terminal state satisfies  $Q[v/r]$  from the point of view of thread  $t$ . Informally, an assertion  $P$  holds from point of view of a thread  $t$  if  $P$ ’s assertion about the heap holds from the point of view of main memory and  $t$ ’s store buffer and no other thread’s store buffer contains pending writes to these parts of the heap. In particular,  $x \mapsto v$  holds from the point of view of thread  $t$ , if the value of  $x$  that  $t$  can observe is  $v$ . We shall see that this interpretation justifies the standard separation logic read and write rules.

What about transfer of resources? In separation logics for sequentially consistent memory models, assertions about resources are objective and can thus be transferred freely. However, since assertions in the SC logic are interpreted subjectively, they may not hold from the point of view of other threads. Clearly, to transfer resources between threads, their views of the resources must match. In general, we thus cannot reason directly about code that facilitates resource transfer in the SC logic. To reason about such code, we use the TSO logic, which allows us to reason about the complete TSO machine state, including store buffers. Importantly, in cases where the implementation that facilitates the resource transfer between two threads *does* ensure that their views match, we can verify the implementation against an SC specification! This gives us a *fiction of sequential consistency* and allows us to reason about clients using the SC logic.

**Example.** To illustrate, consider a simple spin-lock library with acquire and release methods. We can specify the lock in the SC logic as follows.

$$\begin{aligned} & \exists \text{isLock, locked} : \text{Prop}_{\text{SC}} \times \text{Val} \rightarrow \text{Prop}_{\text{SC}}. \\ & \forall R : \text{Prop}_{\text{SC}}. \text{stable}(R) \Rightarrow \\ & \quad [\text{R}] \text{Lock}() [r. \text{isLock}(R, r)] \\ & \quad \wedge [\text{isLock}(R, \text{this})] \text{Lock.acquire}() [\text{locked}(R, \text{this}) * R] \\ & \quad \wedge [\text{locked}(R, \text{this}) * R] \text{Lock.release}() [\top] \\ & \quad \wedge \text{valid}(\forall x : \text{Val}. \text{isLock}(R, x) \Leftrightarrow \text{isLock}(R, x) * \text{isLock}(R, x)) \\ & \quad \wedge \forall x : \text{Val}. \text{stable}(\text{isLock}(R, x)) \wedge \text{stable}(\text{locked}(R, x)) \end{aligned}$$

Here  $\text{Prop}_{\text{SC}}$  is the type of propositions of the SC logic, and  $\text{isLock}$  and  $\text{locked}$  are thus abstract representation predicates.  $\text{isLock}(R, x)$  expresses that  $x$  is a lock protecting the resource invariant  $R$  and  $\text{locked}(R, x)$  expresses that the lock  $x$  is indeed locked. Acquiring the lock grants ownership of  $R$ , while releasing the lock requires the client to relinquish ownership of  $R$ . Since the resource invariant  $R$  is universally quantified, this is a very strong specification; in particular, the client is free to instantiate  $R$  with any SC proposition. This specification requires the resource invariant to be stable,  $\text{stable}(R)$ . The reason is that  $R$  could in general refer to shared resources and to reason about shared resources we need to ensure we only use assertions that are closed under interference from the environment. This is what stability expresses.

Note that this specification is expressed in the SC logic and the specification of the acquire method thus grants ownership of the resource  $R$  from the caller's point of view. Likewise, the release method only requires that  $R$  holds from the caller's point of view. This specification thus provides a fiction of sequential consistency, by allowing transfer of SC resources. Crucially, since the lock specification is an SC specification, we can reason about well-synchronized clients *entirely* using the standard proof rules of the SC logic. We illustrate this by verifying a shared bag in Section 3.

Using the TSO logic we can verify that an *efficient* spin-lock implementation satisfies this specification. The spin-lock that we verify is inspired by the Linux spin-lock implementation [1] which allows the release to be buffered. To verify the implementation we must prove that between releasing and acquiring the lock, the releasing and acquiring threads' views of the resource  $R$  match. Intuitively, this is the case because if  $R$  holds from the point of view of the releasing thread, once the buffered release makes it to main memory,  $R$  holds objectively. This style of reasoning relies on the ordering of buffered writes. To capture this style of reasoning, we introduce a new operator in the TSO logic for expressing such ordering dependencies. In Section 5 we illustrate how to use the TSO logic to verify the spin-lock.

$$\begin{aligned} \text{Val } \ni v ::= & x \mid \text{null} \mid \text{this} \mid o \mid n \mid b \mid () \\ \text{Exp } \ni e ::= & v \mid \text{let } x = e_1 \text{ in } e_2 \mid \text{if } v \text{ then } e_1 \text{ else } e_2 \mid \text{new } C(\bar{v}) \\ & \mid v.f \mid v_1.f := v_2 \mid v.m(\bar{v}) \mid \text{CAS}(v_1.f, v_2, v_3) \mid \text{fork}(v.m) \end{aligned}$$

**Figure 1.** Syntax of the programming language. In the definition of values,  $n$  ranges over machine integers,  $b$  over booleans, and  $o$  over object references. In the definition of expressions,  $f$  ranges over the field names, and  $m$  over the method names.

**Technical Overview.** While our proof system consists of two logics, SC triples are in fact interpreted as TSO triples, through an embedding that formalizes our subjective interpretation of SC assertions. Soundness of the standard read and write rules in the SC logic follow from this subjective interpretation. Soundness of the standard structural rules in the SC logic follow from the fact that this embedding is well-behaved with respect to the connectives and quantifiers of the TSO logic (Lemma 2).

iCAP-TSO builds on iCAP [16], a recent extension of higher-order separation logic [3] for modular reasoning about concurrent higher-order programs with shared mutable state.

**Summary of Contributions.** We provide a new program logic, iCAP-TSO, for a TSO memory model, which features:

- a novel logic for reasoning about low-level *racy* code on a TSO memory model, called TSO logic; this logic features new connectives for expressing ordering dependencies introduced by store buffers
- a *standard* separation logic, called SC logic, that allows *simple* reasoning from the perspective of a single thread
- and, importantly, a *fiction of sequential consistency* which allows us to reason about ownership transfer within the SC logic

Moreover, we prove soundness of iCAP-TSO with respect to a model. Using this logic we verify an *efficient* spin-lock implementation against an SC lock specification. Crucially, this means that we can reason about well-synchronized clients *entirely* using standard separation logic proof rules!

**Outline.** In Section 2 we introduce the programming language that we reason about and its operational semantics. Section 3 illustrates how the fiction of sequential consistency allows us to reason about shared resources using standard separation logic. Section 4 introduces the TSO logic and connectives introduced to reason about store buffers. In Section 5 we illustrate the use of the TSO logic to verify an efficient spin-lock. In Section 6 we discuss the iCAP-TSO soundness theorem. Finally, in Sections 7 and 8 we discuss related work and future work and conclude. Details and proofs can be found in the accompanying technical report [12].

## 2. Language

We build our logic for a simple, class-based programming language. For simplicity of both semantics and the logic, the language uses let-bindings and expressions, but it stays relatively low-level by ensuring that all the values are machine-word size. The values include variables, natural numbers, booleans, unit, object references (pointers), the null pointer and the special variable **this**. The expressions include values, let bindings, conditionals, constructor and method calls, field reads and writes, atomic compare-and-swap expression and a fork call. The syntax of values and expressions is shown in Figure 1. The class and method definitions are standard and so omitted here; they can be found in the accompanying technical report.

Following the Views framework [7], the operational semantics is split into two components: a thread-local small-step semantics labeled with actions that occur during the step, and action semantics that defines the effect of the action on the machine state — in our case, the heap and the store buffer pool. In the thread-local semantics, a thread, consisting of a thread identifier and expression takes a single step of evaluation to a finite map of threads that contains besides the original thread also the threads spawned by this step. It also emits the action that describes the interaction with the memory state. For instance, the WRITE rule in Figure 3 applies when the expression associated with thread  $t$  is an assignment (possibly in some evaluation context). It reduces by replacing the assignment with a unit value, and emits a write action that states that thread  $t$  wrote the value  $v$  to the field  $f$  of object  $o$ .

The non-fault memory state consists of a heap — a finite map from pairs of an object reference and a field to semantic values (i.e., all the values that are not variables) — and a store buffer pool, which contains a sequence of buffered updates for each of the finitely many thread identifiers. The memory can also be in a fault state (written  $\zeta$ ), which means that an error in the execution of the program has occurred. The action semantics interprets the actions as functions from memory states to sets of memory states: if it is impossible for the action to occur in the given state, the result is an empty set; if, however, the action may occur in the given state but it would be erroneous, the result is the fault state. Consider the write action emitted by reducing an assignment. In Figure 4 we can see the interpretation of the action: there are three distinct cases. The action is successful if there is a store buffer associated with the thread that emitted the action and the object is allocated in the heap, and has the appropriate field. In this case, the write gets added to the end of the buffer. However, the write action can have two additional outcomes: if there is no store buffer associated with the thread in the store buffer pool, the initial state had to be ill-formed, and so the interpretation of the action is an empty set; however, if the thread is defined, but the reference to the field  $o.f$  is not found in the heap, the program is erroneous, and the interpretation of the action is a fault state.

The state of a complete program consists of the thread pool and a memory state, and is consistent if the memory state is a fault, or the domain of the store buffer pool equals the domain of the thread pool. The complete semantics proceeds by reducing one of the threads using the thread-local semantics, then interpreting the resulting action with the action semantics, and reducing to a memory state in the resultant set, as in Figure 2. Note how in some cases, notably read, this might require “guessing” the return value, and checking that the guess was right using the action semantics. Some of the cases of the semantics are written out in Figures 3 and 4. In particular, we show the reduction and action semantics that correspond to the (nondeterministic) flushing of a store buffer: a flush action can be emitted by a thread at any time, and the action is interpreted by flushing the oldest buffered write to the memory. Note also the rules for the compare-and-swap expression: similarly to reading, the return value has to be guessed by the thread local semantics. In the case when the guess does not match the state of the memory, the result of the action semantics is empty set, much like when reading. However, if the compare-and-swap succeeds and the return value matches the outcome, the new value is not just written to the store buffer: instead, the *whole* content of the buffer, including the update resulting from compare-and-swap, is written to main memory. Thus, this expression can serve as a synchronization primitive.

Note that our operational semantics is relatively high-level with respect to the canonical TSO semantics [15]. However, it exhibits all the necessary weak behaviors of the TSO model, and as such can serve as a vehicle for studying the problem.

$$\frac{t \in \text{dom } T \quad (t, T(t)) \xrightarrow{a} T' \quad \mu' \in \llbracket a \rrbracket(\mu)}{(\mu, T) \rightarrow (\mu', (T - t) \uplus T')}$$

**Figure 2.** Single step evaluation of a thread pool.

$$\frac{}{(t, E[o.f := v]) \xrightarrow{\text{write}(t, o, f, v)} \{(t, E[\mathbf{0}])\}} \text{WRITE}$$

$$\frac{}{(t, E[o.f]) \xrightarrow{\text{read}(t, o, f, v)} \{(t, E[v])\}} \text{READ}$$

$$\frac{}{(t, E[\text{CAS}(o.f, v_o, v_n)]) \xrightarrow{\text{cas}(t, o, f, v_o, v_n, r)} \{(t, E[r])\}} \text{CAS}$$

$$\frac{}{(t, e) \xrightarrow{\text{flush}(t)} \{(t, e)\}} \text{FLUSH}$$

**Figure 3.** Selected cases of the thread-local operational semantics.

### 3. Reasoning in the SC logic

The SC logic of iCAP-TSO allows us to reason about well-behaved code, *using standard separation logic*, without having to reason about store buffers. Naturally, this class of well-behaved code includes standard mutable data structures without any sharing. We can thus easily verify a list library in the SC logic against the standard separation logic specification as it enforces a *unique* owner. Crucially, this class of well-behaved code also includes the large class of well-synchronized programs *with sharing*. Using the lock specification from the Introduction and the fiction of sequential consistency that it provides, we illustrate how to verify a well-synchronized program with sharing in the SC logic. In particular, we verify a shared bag library implemented using a list protected by a lock.

**The SC logic.** The SC logic is an intuitionistic higher-order separation logic. Recall that the SC logic features Hoare triples of the form  $\llbracket P \rrbracket e \llbracket r, Q \rrbracket$ , where  $P$  and  $Q$  are SC assertions. Formally, SC assertions are terms of type  $\text{Prop}_{\text{SC}}$ . SC assertions include the usual connectives and quantifiers of higher-order separation logic and language specific assertions such as points-to,  $x.f \mapsto v$ , for asserting the value of field  $f$  of object  $x$ .

Recall that SC triples employ a subjective interpretation of the pre- and postcondition:  $\llbracket P \rrbracket e \llbracket r, Q \rrbracket$  expresses that if thread  $t$  executes the expression  $e$  from an initial state where  $P$  holds from the point of view of thread  $t$  and  $e$  terminates with value  $v$  then  $Q[v/r]$  holds for the terminal state from the point of view of thread  $t$ . An assertion  $P$  holds from the point of view of a thread  $t$  if  $P$ 's assertions about the heap hold from the point of view of  $t$ 's store buffer and main memory *and* no other thread's store buffer contains a buffered write to these parts of the heap. The assertion  $x.f \mapsto v$  thus holds from the point of view of thread  $t$  if

- the value of the most recently buffered write to  $x.f$  in  $t$ 's store buffer is  $v$
- or  $t$ 's store buffer does not contain any buffered writes to  $x.f$  and the value of  $x.f$  in main memory is  $v$

and no other threads store buffer contains a buffered write to  $x.f$ . The condition that no other thread's store buffer can contain a buffered write to  $x.f$  ensures that flushing of store buffers cannot invalidate  $x.f \mapsto v$  from the point of view of a given thread.

If  $x.f \mapsto v$  holds from the point of view of thread  $t$  and thread  $t$  attempts to read  $x.f$  it will thus read the value  $v$  either from main memory or its own store buffer. Likewise, if  $x.f \mapsto v_1$  holds from the point of view of thread  $t$  and thread  $t$  writes  $v_2$  to  $x.f$ , afterwards

$$\begin{aligned}
\llbracket \text{read}(t, o, f, v) \rrbracket(h, U) &= \begin{cases} \{(h, U)\} & \text{if } (o, f) \in \text{dom } h \text{ and } \text{lookup}(o, f, U(t), h) = v \\ \emptyset & \text{if } t \notin \text{dom } U \text{ or } (o, f) \in \text{dom } h \text{ and } \text{lookup}(o, f, U(t), h) \neq v \\ \{\zeta\} & \text{if } (o, f) \notin \text{dom } h \end{cases} \\
\llbracket \text{write}(t, o, f, v) \rrbracket(h, U) &= \begin{cases} \{(h, U[t \mapsto U(t) \cdot (o, f, v)])\} & \text{if } (o, f) \in \text{dom } h \text{ and } t \in \text{dom } U \\ \{\zeta\} & \text{if } (o, f) \notin \text{dom } h \\ \emptyset & \text{if } t \notin \text{dom } U \end{cases} \\
\llbracket \text{cas}(t, o, f, v_o, v_n, r) \rrbracket(h, U) &= \begin{cases} \{(\text{flush}(h, U(t) \cdot (o, f, v_n)), U[t \mapsto \varepsilon])\} & \text{if } (o, f) \in \text{dom } h, r = \mathbf{true} \text{ and } \text{lookup}(o, f, U(t), h) = v_o \\ \{(h, U)\} & \text{if } (o, f) \in \text{dom } h, r = \mathbf{false} \text{ and } \text{lookup}(o, f, U(t), h) \neq v_o \\ \{\zeta\} & \text{if } (o, f) \notin \text{dom } h \\ \emptyset & \text{otherwise} \end{cases} \\
\llbracket \text{flush}(t) \rrbracket(h, U) &= \begin{cases} \{(h[(o, f) \mapsto v], U[t \mapsto \alpha])\} & \text{if } U(t) = (o, f, v) \cdot \alpha \text{ and } (o, f) \in \text{dom } h \\ \emptyset & \text{if } t \notin \text{dom}(U), U(t) = \varepsilon, \text{ or } (o, f) \notin \text{dom } h \end{cases}
\end{aligned}$$

**Figure 4.** Selected cases of the action semantics. The lookup function finds the newest value associated with the field, including the store buffer, while the flush function applies all the updates from the store buffer to the heap in order.

$$\begin{array}{c}
\frac{}{[x.f \mapsto v] \ x.f \ [r. x.f \mapsto v * r = v]} \text{S-READ} \quad \frac{}{[x.f \mapsto v_1] \ x.f := v_2 \ [r. x.f \mapsto v_2]} \text{S-WRITE} \quad \frac{\triangleright([P] \ C(\bar{x}) \ [r. Q])}{[P[\bar{v}/\bar{x}]] \ \mathbf{new} \ C(\bar{v}) \ [r. Q[\bar{v}/\bar{x}]]} \text{S-NEW} \\
\frac{\triangleright([P] \ C.m(\bar{y}) \ [r. Q])}{[P[\bar{v}/\bar{y}, x/\mathbf{this}] * x : C] \ x.m(\bar{v}) \ [r. Q[\bar{v}/\bar{y}, x/\mathbf{this}]]} \text{S-CALL} \quad \frac{[P] \ e_1 \ [r. Q] \quad [Q[x/r]] \ e_2 \ [r. R]}{[P] \ \mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2 \ [r. R]} \text{S-BIND} \quad \frac{}{[\top] \ v \ [r. r = v]} \text{S-VAL} \\
\frac{P_1 \vdash P_2 \quad [P_2] \ e \ [r. Q_2] \quad Q_2 \vdash Q_1}{[P_1] \ e \ [r. Q_1]} \text{S-CONS} \quad \frac{[P] \ e \ [r. Q] \quad \mathbf{stable}(R)}{[P * R] \ e \ [r. Q * R]} \text{S-FRAME}
\end{array}$$

**Figure 5.** Selected proof rules for the sequentially consistent logic.

$x.f \mapsto v_2$  holds from the point of view of thread  $t$ . We thus get the standard rules for reading and writing to a field in our SC logic (rules S-READ and S-WRITE in Figure 5.)

The S-NEW and S-CALL rules are the object allocation and method call rules of the SC logic. The assertion  $x : C$  expresses that  $x$  refers to an object of class  $C$ . The method call rule (S-CALL) thus allows us to reason about a call to a method  $m$  of a class  $C$ . The later operator, written  $\triangleright$  in the premise of the S-CALL rule expresses that the given specification must hold *one step later*. Among other things, these later operators internalizes a notion of steps of evaluation. This allows us to reason about mutually recursive methods using Löb induction. The later operator goes back to Gödel-Löb logic and has been used in several models and logics for reasoning about higher-order code, e.g., [2, 11, 17]. Note the side condition that  $R$  is stable in the S-FRAME rule: this is a requirement similar to the one we have seen in the specification of the lock in the Introduction, and asserts that the interference of the environment cannot invalidate the frame.

**A list library.** The standard separation logic specification for a list library enforces a unique owner and thus ensures that the underlying heap representation is not shared. We can express the standard

list specification in our SC logic as follows.

$$\begin{aligned}
&\exists \text{lst} : \text{Val} \times \text{seq Val} \rightarrow \text{Prop}_{\text{SC}}. \\
&[\top] \text{List}() \ [r. \text{lst}(r, \varepsilon)] \wedge \\
&\forall l. [\text{lst}(\mathbf{this}, l)] \ \text{List.push}(x) \ [\text{lst}(\mathbf{this}, x :: l)] \wedge \\
&\forall l. [\text{lst}(\mathbf{this}, l)] \\
&\quad \text{List.pop}() \\
&\quad \left[ \begin{array}{l} r. (r = \mathbf{null} * l = \varepsilon * \text{lst}(\mathbf{this}, l)) \\ \vee (r = \text{hd}(l) * \text{lst}(\mathbf{this}, \text{tail}(l))) \end{array} \right] \wedge \\
&\forall x : \text{Val}. \forall l_1, l_2 \in \text{seq Val}. \text{valid}(\text{lst}(x, l_1) * \text{lst}(x, l_2) \Rightarrow \perp)
\end{aligned}$$

Here  $\text{Val}$  is the type of values of our programming language and  $\text{seq Val}$  the type of finite sequences with elements of type  $\text{Val}$ . The  $\text{lst}$  predicate is a representation predicate relating a mathematical model of a list with its concrete heap representation. We use higher-order quantification to existentially quantify the  $\text{lst}$  predicate and thus hide the internal data representation of implementations from clients.

Consider the specification of the list constructor. Due to the subjective interpretation of SC triples the constructor specification expresses that if thread  $t$  creates a new list, then there exists a new list from the point of view of thread  $t$ . It says nothing about the list from the point of view of other threads (other than that other thread's store buffers cannot contain buffered writes to the list). Likewise, for the other methods: if the list contains the values  $l$  from the point of view of thread  $t$  and thread  $t$  pushes an element  $x$ , afterwards the list contains the values  $x :: l$  from the point of view of  $t$ .

This specification also explicitly rules our sharing, by asserting that  $\text{lst}$  is non-duplicable. Due to this subjective interpretation and

```

Bag {
  Lock lock; List elms;

  pop() {
    let x = this.lock in let y = this.elms in
      x.acquire(); let z = y.pop() in x.release(); z
  }
  ...
}

```

**Figure 6.** Sketch of a shared bag implementation.

the lack of sharing, the verification of a linked-list implementation against this specification is as in separation logic for sequentially consistent memory models. We can instantiate the `lst` representation predicate with the usual definition and verify the implementation using the usual proof rules.

We can thus reason about mutable data structures without any sharing in the SC logic as if we were in standard separation logic. However, the real power of our proof system comes from the fact that we can still reason as in standard separation logic about sharing of such lists through libraries that provide a fiction of sequential consistency.

**A shared bag.** To illustrate, consider a shared bag implemented using a list. Each shared bag maintains a list of elements and a lock to ensure exclusive access to the list of elements. Each bag method acquires the lock and calls the corresponding method of the list library before releasing the lock. In Figure 6 we sketch an implementation of a shared bag.

We take the following specification, which allows unrestricted sharing of the bag, to be our specification of a shared bag. This is not the most general specification we can express, but it suffices to illustrate that verification of the shared bag against such specifications is standard. Since the specification allows unrestricted sharing (the bag predicate is duplicable), no client can know the contents of the bag; instead, the specification allows clients to associate ownership of additional resources (expressed using the predicate  $P$ ) with each element in the bag.

$$\begin{aligned}
& \exists \text{bag} : \text{Val} \times (\text{Val} \rightarrow \text{Prop}_{\text{SC}}) \rightarrow \text{Prop}_{\text{SC}}. \forall P : \text{Val} \rightarrow \text{Prop}_{\text{SC}}. \\
& (\forall x : \text{Val}. \text{stable}(P(x))) \Rightarrow \\
& \quad [\top] \text{Bag}() [r. \text{bag}(r, P)] \wedge \\
& \quad [\text{bag}(\mathbf{this}, P) * P(x)] \text{Bag.push}(x) [\top] \wedge \\
& \quad [\text{bag}(\mathbf{this}, P)] \text{Bag.pop}() [r. (P(r) \vee r = \mathbf{null})] \wedge \\
& \quad \forall x : \text{Val}. \text{valid}(\text{bag}(x, P) \Leftrightarrow \text{bag}(x, P) * \text{bag}(x, P))
\end{aligned}$$

Pushing an element  $x$  thus requires the client to transfer ownership of  $P(x)$  to the bag. Likewise, either `pop` returns `null` or the client receives ownership of the resources associated with the returned element.

To verify the implementation against this specification, we first have to define the abstract bag representation predicate. To define bag we first need to define the resource invariant of the lock. Intuitively, the lock owns the list of elements and the resources associated with the elements currently in the list. This is expressed by the following resource invariant  $R_{\text{bag}}(x, P)$ , where  $x$  refers to the list of elements.

$$R_{\text{bag}}(x, P) \stackrel{\text{def}}{=} \exists l : \text{seq Val}. \text{lst}(x, l) * \otimes_{y \in \text{mem}(l)} P(y)$$

The bag predicate asserts *read-only* ownership of the lock and elms fields, and that the lock field refers to a lock with the above resource

```

pop() {
  [bag(this, P)]
  let x = this.lock in
  let y = this.elms in
    x.acquire();
  [this.elms ↦ y * locked(x, R_bag(P)) * R_bag(P)]
  let z = y.pop() in
  [locked(x, R_bag(P)) * R_bag(P) * (z = null ∨ P(z))]
  x.release();
  [isLock(x, R_bag(P)) * (z = null ∨ P(z))]
  z
  [r. r = null ∨ P(r)]
}

```

**Figure 7.** Proof outline of the bag `Pop` method.

invariant.

$$\text{bag}(x, P) \stackrel{\text{def}}{=} \exists y, z : \text{Val}. x.\text{lock} \mapsto y * x.\text{elms} \mapsto z * \text{isLock}(y, R_{\text{bag}}(z, P))$$

Now, we are ready to verify the bag methods. The most interesting method is `Pop`, as it actually returns the resources associated with the elements it returns. Figure 7 contains a proof outline of `Pop`. The crucial thing to note is that from the specification of the lock, once a thread  $t$  acquires the lock, it receives ownership of the resource invariant  $R_{\text{bag}}(P)$  from the point of view of  $t$ . This is exactly what we need to use the SC specification of the `List.pop` method with the `S-CALL` rule!

In general, we can thus reason about well-synchronized code with sharing in the SC logic as in separation logics for sequentially consistent memory models.

## 4. TSO logic and connectives

In this section we describe the TSO logic and introduce our new TSO connectives that allow us to reason about the kinds of weak behaviors that occur in low-level concurrency libraries.

We can express the configurations that lead to these weak behaviors by extending the space of states over which the assertions are built — in the case of our TSO model, we include the store buffer pool as an additional component of the memory state. However, reasoning about the buffers directly would be extremely complicated and contrary to the spirit of program logics. Hence, we introduce new logical connectives that help express these new configurations in a useful, abstract manner, and provide appropriate reasoning rules.

**The triples and assertions of the TSO logic** First, however, we need to consider how the TSO logic is built. As mentioned in the Introduction, its propositions extend the propositions of SC logic by adding the store buffer pool component. Just like SC assertions, this space forms a higher-order intuitionistic separation logic, with the usual rules for reasoning about assertion entailment. However, we are still reasoning about the code running in some thread and we often need to state properties that hold of its — or other thread's — store buffer. Thus, formally, the typing rule for the TSO logic triples is as follows:

$$\frac{P : \text{TId} \rightarrow \text{Prop}_{\text{TSO}} \quad Q : \text{TId} \rightarrow \text{Val} \rightarrow \text{Prop}_{\text{TSO}}}{\{P\} e \{Q\} : \text{Spec}}$$

where `TId` is the type of thread identifiers and `Spec` is the type of specifications. We usually keep this quantification over thread identifiers implicit, by introducing appropriate syntactic sugar for the TSO-level connectives. The logic also include another family

$$\begin{array}{c}
\frac{}{\triangleright x.f \mapsto v \vdash_w x.f \mapsto v, \top} \text{W-AX} \quad \frac{P \vdash_w x.f \mapsto v, R}{P * \triangleright \bar{Q} \vdash_w x.f \mapsto v, Q * R} \text{W-*} \quad \frac{Q \vdash_w x.f \mapsto v, R}{P \mathcal{U} Q \vdash_w x.f \mapsto v, R} \text{W-}\mathcal{U} \quad \frac{P \vdash_r x.f \mapsto v}{\langle P \rangle x.f \langle r. P * r = v \rangle^C} \text{A-READ} \\
\frac{P \vdash_w x.f \mapsto v', Q}{\langle P \rangle x.f := v \langle \dots P \mathcal{U} (\ulcorner Q * x.f \mapsto v \urcorner) \rangle^C} \text{A-WRITE} \quad \frac{P \vdash_w x.f \mapsto v, Q}{\langle P \rangle \text{CAS}(x.f, v', v) \langle r. r = \text{true} * \ulcorner Q * x.f \mapsto v \urcorner \rangle^C} \text{A-CAS-TRUE} \\
\frac{P \vdash_r x.f \mapsto v \quad v \neq v_2}{\langle P \rangle \text{CAS}(x.f, v_1, v_2) \langle r. r = \text{false} * P \rangle^C} \text{A-CAS-FALSE} \quad \frac{\langle P \rangle e \langle Q \rangle^C \quad \text{atomic}(e)}{\{P\} e \{Q\}} \text{A-START} \\
\frac{\{P\} e_1 \{r. Q(r)\} \quad \{Q(x)\} e_2 \{r. R(r)\}}{\{P\} \text{let } x = e_1 \text{ in } e_2 \{r. R(r)\}} \text{BIND} \quad \frac{\{P\} e \{Q\} \quad \text{stable}(R)}{\{P * R\} e \{Q * R\}} \text{FRAME} \quad \frac{\{\bar{P}\} e \{\bar{Q}\}}{\{P\} e \{Q\}} \text{S-SHIFT}
\end{array}$$

Figure 8. Select rules of the TSO logic.

of Hoare triples, the atomic triples, with the following typing rule:

$$\frac{\text{atomic}(e) \quad P : \text{TId} \rightarrow \text{Prop}_{\text{TSO}} \quad Q : \text{TId} \rightarrow \text{Val} \rightarrow \text{Prop}_{\text{TSO}}}{\langle P \rangle e \langle Q \rangle^C : \text{Spec}}$$

As the rule states, these triples can only be used to reason about atomic expressions — read, write and compare-and-swap. This feature is inherited from iCAP, as a means of reasoning about the way the shared state changes at the atomic updates. We give an example of such reasoning in Section 5.

**Two spaces of assertions and the connection between them.** As mentioned in the Introduction, our proof system consists of two logics, each with a corresponding space of assertions. To provide the fiction of sequential consistency and show SC specifications for implementations whose correctness involves reasoning about buffered updates, we need to use of both of these spaces in the TSO logic. To this end we use an embedding of  $\text{Prop}_{\text{SC}}$  into  $\text{Prop}_{\text{TSO}}$ .

The canonical embedding is denoted  $\ulcorner - \urcorner : \text{Prop}_{\text{SC}} \rightarrow \text{Prop}_{\text{TSO}}$ . The idea is that  $\ulcorner P \urcorner$  holds in a state that does include store buffers if  $P$  holds in the state where we ignore the buffers *and* none of the buffers contain buffered updates to the locations mentioned by  $P$ . The intuition behind this embedding is that  $P$  should hold *in main memory*. The condition that no buffered updates whatsoever to the state in  $P$  are allowed may at first seem overly restrictive. However, it is crucial to our setup, in particular to allow resource transfer and the fiction of sequential consistency. This is due to the nondeterministic nature of the store buffers: if more than one thread contains a buffered update to the same location, no thread can be certain what values of that location it can observe. Thus, we can maintain a reasonable level of complexity only as long as a thread can only write to a location for which no *other* thread has buffered updates. However, *transferring* a resource must mean passing the right to update it to a different thread. Thus, the thread that gives up a resource also can have no buffered writes at the point when the transfer occurs, so no updates of the resource can be buffered at all when the transfer occurs.

For a concrete example of what this embedding means, consider an assertion  $x.f \mapsto v : \text{Prop}_{\text{SC}}$ . Clearly, we can use our embedding to get  $\ulcorner x.f \mapsto v \urcorner$  — an assertion that means precisely that the reference  $x.f$  is defined, its associated value in the heap is  $v$ , and there are no buffered updates in any of the store buffers to the field  $x.f$ .

**Reasoning about buffered updates.** Even though there must be no buffered updates of the resource when we transfer resources between threads, most of the time we reason about situations where buffered updates *are* possible. Therefore we define two additional connectives that allow us to reason about the presence of buffered updates. The first of these is a binary function  $\ulcorner - \text{ in } - \urcorner : \text{Prop}_{\text{SC}} \times$

$\text{TId} \rightarrow \text{Prop}_{\text{TSO}}$ . Intuitively,  $\ulcorner P \text{ in } t \urcorner$  means that  $P$  holds from the perspective of thread  $t$  — including the possible buffered updates in the store buffer of  $t$ . However, for the reasons discussed prior, we have to disallow buffered updates to the locations described by  $P$  in any of the other store buffers.

The second of this pair of connectives is the one that allows us to express how the state *changes* due to an update. Because of its role, it has a certain temporal feel: in fact, it behaves in a way that is somewhat similar to the classic “until” operator. Intuitively,  $P \mathcal{U}_t Q$  means that there exists a buffered update in the store buffer of thread  $t$ , such that until this update is flushed the assertion  $P$  holds, while after the update gets written to memory, the assertion  $Q$  holds. Thus, it can be used to describe the ordering dependencies introduced by the presence of store buffers. This intuition should become clearer by observing the proof rules in Figure 8 (explained in the following).

Coming back to our example, consider first the assertion  $\ulcorner x.f \mapsto v \text{ in } t \urcorner$ . This proposition holds in any state where all the buffered updates to  $x.f$  are in the store buffer associated with thread  $t$  — including the boundary case when there are no buffered updates to  $x.f$  at all. Thus, if we are reasoning from the perspective of thread  $t$ , we know precisely the value we can read from the field  $x.f$ , but no other thread has any knowledge about the state, save that  $x.f$  is defined. On the other hand, in the state described by  $\ulcorner x.f \mapsto 1 \urcorner \mathcal{U}_t \ulcorner x.f \mapsto 2 \urcorner$ , we have information that any thread can use: we know that the value of  $x.f$  in the heap is 1, and that there exists a buffered update in thread  $t$ , such that before that update there is no updates to  $x.f$ , i.e., it is the first update to  $x.f$  in the store buffer of  $t$ , and after it gets flushed  $\ulcorner x.f \mapsto 2 \urcorner$  holds — so the update must set  $x.f$  to 2. Additionally, we know that there are no other buffered updates to  $x.f$ . This means that the thread  $t$  can observe the value of  $x.f$  to be 2, while all of the other threads can observe it to be 1. Note that, since  $\mathcal{U}$  is a binary operator on  $\text{Prop}_{\text{TSO}}$ , it is possible to use it to express multiple buffered updates.

Since most of the time we are reasoning from the perspective of a particular thread, we also include some syntactic sugar:  $\mathcal{U}$  is a shorthand for an update in the current thread, while  $\bar{\mathcal{U}}$  is a shorthand for an update in some thread *other than the current one*. We also use  $\bar{p}$  as a shorthand for  $\ulcorner p \text{ in } t \urcorner$ , where  $t$  is the current thread.

**Reading and writing state.** The presence of additional connectives that mention the state makes reading fields of an object and writing to them more involved than in standard separation logic. We deal with this by introducing additional judgments that specify when we can read or write a value. Intuitively,  $P \vdash_r x.f \mapsto v$  specifies that, from the perspective of the current thread  $x.f$  has the value  $v$  — and so the read expression can proceed. Similarly,  $P \vdash_w x.f \mapsto v, Q$  means that in the state specified by  $P$  the current thread has permission to write to  $x.f$ , the value of  $x.f$  from its perspective is  $v$ , and that once an update to  $x.f$  reaches main memory,

$Q$  will additionally hold *in main memory*. In this sense,  $Q$  behaves like a frame; however, since the rule is used to reason about writes, which affect the store buffers, we can learn something about the order of writes in the buffer, which we would not be able to do using the frame rule. Note that the reading of the write judgment corresponds to the intuitive understanding of the postcondition of the A-WRITE rule in Figure 8: the write we are performing introduces a buffered update, before which the precondition still holds, and after which  $\ulcorner Q * x.f \mapsto v \urcorner$  holds. This setup allows us to retain some knowledge of the order in which updates reach the main memory in an abstract manner; a property that is crucial to allow resource transfer in the presence of the store buffers. Note also that the A-CAS-TRUE rule can be viewed as a combination of writing and flushing. As mentioned above, a write adds a buffered write to the appropriate store buffer, which is expressed using  $\mathcal{U}$  in A-WRITE. However, if at that moment we *flush* the store buffer, we get a condition under which the right-hand side of  $\mathcal{U}$  is defined to hold: thus, we get to know that  $\ulcorner Q * x.f \mapsto v \urcorner$  holds in the postcondition of A-CAS-TRUE.

Also of interest are some of the proof rules for the write judgment (the read judgment behaves in a similar manner, with some complications when reading from  $\bar{\mathcal{U}}$ ). Note how in rules W-AX and W-\* the later operators ( $\triangleright$ ) appear. These arise from the fact that the model is defined using guarded recursion to break circularities and, since they play the role of guards, they can be removed at the atomic steps of the proof, as expressed by the rules. Moreover, the rules also match the intuition we gave about the store buffer related connectives. First, the judgment means that it is enough that we know the value of  $x.f$  from the thread’s perspective: hence, we are interested in the most recent updates, and look to the right of  $\mathcal{U}$  in W- $\mathcal{U}$ . For the same reason, it is enough to use  $\bar{\phantom{x}}$  in W-AX. Second, due to the way we interpret the “frame” in the write judgment, we only need  $Q$  to hold from the thread’s perspective in W-\*. This is also why we can use the frame taken from the right-hand-side of  $\mathcal{U}$  in W- $\mathcal{U}$ : the update that we are reasoning about in the write judgment will happen *after* the buffered update asserted in the definition of  $\mathcal{U}$ , so the frame taken from after the first update will hold in the memory by the time the later one reaches it.

**Stability and stabilization.** There is one potentially worrying issue with the explanation of the  $\mathcal{U}_t$  operator given in this section: since at any point in the program a *flush* action can occur nondeterministically, how can we know that there still exists a buffered update as asserted by  $\mathcal{U}_t$ ? After all, it might have been flushed to the memory. This is the question of *stability*<sup>1</sup> of the until operator — and the answer is that it is unstable by design. The rationale behind this choice is simple: Suppose we had made it stable by allowing the possibility that the buffered update has already been flushed. Then, if we were to read a field that had a buffered write to it in a different thread, we would not know whether the write was still buffered or had been flushed, and so we would not know what value we read. With the current definition, when we read, we know that the update is still buffered and so the result of the read is known. However, we only allow reasoning with unstable assertions in the *atomic* triples, i.e., when reasoning about a single read, write or compare-and-swap expression. Hence, we need a way to make  $\mathcal{U}$  stable. For this reason, we define an explicit *stabilization* operator,  $(\bar{\phantom{x}})$ . It is a closure operator, which means we have  $P \vdash (\bar{P})$ . Moreover, for stable assertions, the other direction  $(\bar{P}) \vdash P$  also holds. The important part, however, is how stabilization behaves with respect to  $\mathcal{U}$ : provided  $P$  and  $Q$  are stable, we have  $(\bar{P} \mathcal{U}_t Q) \dashv\vdash (P \mathcal{U}_t Q) \vee Q$ . This does indeed correspond with our intuition — even for stable assertions  $P$  and  $Q$ , the interference can flush the buffered update that is asserted in the definition

<sup>1</sup> Recall an assertion is stable, if it cannot be invalidated by the environment.

of  $\mathcal{U}$ , which would transition to a state in which  $Q$  holds. However, since  $P$  and  $Q$  are stable, this is also the *only* problem that the interference could cause. Explicit stabilization is not a new idea: it has been used recently in the context of program logics, most commonly in connection with rely-guarantee reasoning. In particular, Wickerson studies explicit stabilization in RGSep in his PhD thesis [19, Chapter 3], and Ridge [14] uses it to reason about the x86-TSO.

Considering stability of the other TSO-level connectives exposes a different issue: the fact that there are two distinct parts to the interference a thread is subject to — the effect of store buffer actions, and the effect other threads can have on shared state. We follow the CAP approach [8] of reasoning about shared state using shared regions with protocols. See Section 5 for an example. We thus have two notions of stability: *r*-stable and *b*-stable. The first of these considers the shared regions, and is the only kind of stability that concerns the SC-logic; it corresponds to stability in iCAP [16]. The second only considers the additional guarantees required by the TSO logic. Stability is defined as the conjunction of the two. Thanks to the restrictions of the contents of the store buffers imposed by the embeddings, for any  $P : \text{Prop}_{\text{SC}}$  we have  $\text{b-stable}(\ulcorner P \urcorner)$  and  $\text{b-stable}(\ulcorner P \text{ in } t \urcorner)$ . These two facts in conjunction with the explicit stabilization operator suffice for reasoning about low level implementations. Although important for reasoning about shared regions, stability under region interference is much as in iCAP, so we elide the proof rules for inferring that a predicate is *r*-stable and refer the reader to the technical report for details [12].

**Interpretation of the SC logic.** Note that the same intuition that lies behind the SC logic, discussed in the previous section, is expressed by the  $\ulcorner \phantom{x} \text{ in } t \urcorner$  embedding. This is more formally expressed by the rule S-SHIFT in Figure 8 (recall  $\bar{P}$  is syntactic sugar for  $\lambda t. \ulcorner P \text{ in } t \urcorner$ ), which states that the two ways of expressing that a triple holds from the perspective of the current thread are equivalent. In fact, we take this rule as the *definition* of the SC triples, and so we can prove that the SC triples actually form a standard separation logic by proving that the proof rules of SC logic correspond to admissible rules in the TSO logic. This is expressed by the following theorem:

**Theorem 1** (Soundness of SC logic). *The SC logic is sound wrt. its interpretation within TSO logic, i.e., the proof rules (in particular those shown in Figure 5) composed with the rule S-SHIFT are admissible rules of the TSO logic.*

For most of the proof rules, the soundness follows directly; the only ones that require additional properties to be proved are the frame, consequence, and standard quantifier rules, which additionally require the following property:

**Lemma 2.** *The embeddings  $\ulcorner \phantom{x} \urcorner$  and  $\ulcorner \phantom{x} \text{ in } t \urcorner$  distribute over quantifiers and separating conjunction, and preserve entailment.*

The formal statement of this property, along with the proof, can be found in the accompanying technical report.

## 5. Reasoning in the TSO logic

In Section 3 we illustrated that the fiction of sequential consistency provided by the lock specification allows us to reason about shared mutable data structures shared through locks, without explicitly reasoning about the underlying weak memory model. Of course, to verify a lock implementation against this lock specification, we *do* have to reason about the weak memory model. In this section we illustrate how to achieve this using our TSO logic. We focus on the use of the TSO-connectives introduced in Section 4 to describe

```

Lock {
  bool locked;

  Lock() {
    CAS(this.locked, false, null)
  }

  acquire() {
    let x = CAS(this.locked, true, false) in
    if x then () else acquire()
  }

  release() {
    this.locked := false
  }
}

```

Figure 9. Spin-lock implementation.

the machine states of the spin-lock and elide the details related to the use of concurrent abstract predicates.

The spin-lock implementation that we wish to verify is given in Figure 9. It uses a compare-and-swap (CAS) instruction to attempt to acquire the lock, but only a primitive write instruction to release the lock. While CAS flushes the store buffer of the thread that executes the CAS, a primitive write does not. To verify this implementation, we thus have to explicitly reason about the possibility of buffered releases in store buffers.

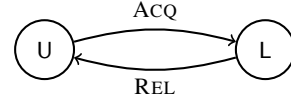
**Specification.** In the Introduction we introduced a lock specification expressed in our SC logic. When verifying the spin-lock implementation, we actually verify the implementation against the following slightly stronger specification, from which we can easily derive the SC specification.

$$\begin{aligned}
& \exists \text{isLock, locked} : \text{Prop}_{\text{SC}} \times \text{Val} \rightarrow \text{Prop}_{\text{SC}} \\
& \forall R : \text{Prop}_{\text{SC}}. \text{stable}(R) \Rightarrow \\
& \quad \{\overline{R}\} \text{Lock}() \{r. \overline{\text{isLock}(R, r)}\} \\
& \quad \wedge \{\overline{\text{isLock}(R, x)}\} \text{Lock.acquire}() \{\overline{\text{locked}(R, x)} * \ulcorner R \urcorner\} \\
& \quad \wedge \{\overline{\text{locked}(R, x)} * R\} \text{Lock.release}() \{\top\} \\
& \quad \wedge \text{valid}(\forall x : \text{Val}. \text{isLock}(R, x) \Leftrightarrow \text{isLock}(R, x) * \text{isLock}(R, x)) \\
& \quad \wedge \forall x : \text{Val}. \text{stable}(\text{isLock}(R, x)) \wedge \text{stable}(\text{locked}(R, x))
\end{aligned}$$

Note that this stronger specification is expressed using TSO triples. This specification of the Acquire method is slightly stronger: this specification asserts that upon termination of Acquire, the resource invariant  $R$  holds in main memory and there are no buffered writes affecting  $R$  in *any* store buffer ( $\ulcorner R \urcorner$ ). The weaker SC specification only asserts that the resource invariant  $R$  holds from the point of view of the acquiring thread and that there are no buffered writes affecting  $R$  in *any of the other threads'* store buffers ( $\overline{R}$ ).

**Lock protocol.** To verify the spin-lock implementation against the above specification, we first need to define the abstract representation predicates  $\text{isLock}$  and  $\text{locked}$ . Following CAP [8] and iCAP [16], to reason about sharing iCAP-TSO extends separation logic with shared regions, with protocols governing the resources owned by each shared region. In the case of the spin-lock, upon allocation of a new spin-lock the idea is to allocate a new shared region governing the state of the spin-lock and ownership of the resource invariant.

Conceptually, a spin-lock can be in one of two states: locked and unlocked. In iCAP-TSO, we express this formally using the following labeled transition system.



The above labeled transition system specifies an abstract model of the lock; to relate it to the concrete implementation, for each abstract state ( $L$  and  $U$ ), we choose an assertion that describes the resources the spin-lock owns in the given abstract state.

Since acquiring the lock flushes the store buffer of the acquiring thread, the locked state is fairly simple. In the locked state the spin-lock owns the `locked` field, which contains the value `true` in main memory and there are no buffered writes to `locked` in any store buffer. The spin-lock  $x$  with resource invariant  $R$  thus owns the resources described by  $I_L(x, R, n)$  in the abstract locked state.

$$I_L(x, R, n) = \ulcorner x.\text{locked} \mapsto \text{true} \urcorner$$

Due to the possibility of buffered releases in store buffers, the unlocked state is more complicated. In the unlocked state,

- either `locked` is `false` in main memory and there are no buffered writes to `locked` in any store buffer
- or `locked` is `true` in main memory, and there is *exactly one* store buffer with a buffered write to `locked`, and the value of this buffered write is `false`

Furthermore, in case there is a buffered write to `locked` that changes its value from `true` to `false`, then, once the buffered write reaches main memory, the resource invariant holds in main memory. Intuitively, the resource invariant holds from the point of view of the releasing thread before the lock is released; hence any buffered writes affecting the resource invariant must reach main memory before the buffered release. We can express this ordering dependency using the until operator as follows.

$$\begin{aligned}
I_U(x, R, n) = \exists t : \text{TId}. (\ulcorner x.\text{locked} \mapsto \text{true} \urcorner \mathcal{U}_t \\
\ulcorner x.\text{locked} \mapsto \text{false} * R * [\text{REL}]_1^{\ulcorner \urcorner} \urcorner)
\end{aligned}$$

Here  $[\text{REL}]_1^{\ulcorner \urcorner}$  is a CAP action permission used to ensure that only the current holder of the lock can release the lock. Since this is orthogonal to the underlying memory model, we refer the interested reader to the technical report [12] for details.

Since both arguments of  $\mathcal{U}_t$  are stable, as explained in Section 4,  $I_U(x, R, n)$  is equivalent to the following assertion.

$$\begin{aligned}
& \exists t : \text{TId}. \\
& \quad \ulcorner x.\text{locked} \mapsto \text{false} * R * [\text{REL}]_1^{\ulcorner \urcorner} \urcorner \vee \\
& \quad (\ulcorner x.\text{locked} \mapsto \text{true} \urcorner \mathcal{U}_t \ulcorner x.\text{locked} \mapsto \text{false} * R * [\text{REL}]_1^{\ulcorner \urcorner} \urcorner)
\end{aligned}$$

The first disjunct corresponds to the case where the release has already made its way to main memory and the second disjunct to the case where it is still buffered.

The definition of  $\text{isLock}$  in terms of  $I_L$  and  $I_U$  now follows standard iCAP.<sup>2</sup> The  $\text{isLock}$  predicate asserts the existence of a shared region governed by the above labeled transition system, where the resources owned by the shared region in the two abstract states are given by  $I_L$  and  $I_U$ . It further asserts that the abstract state of the shared region is either locked or unlocked and also a non-exclusive right to acquire the lock.

**Proof outline.** To verify the spin-lock implementation, it remains to verify each method against the specification instantiated with the concrete  $\text{isLock}$  and  $\text{locked}$  predicates. To illustrate the reasoning related to the weak memory model, we focus on the verification of the acquire method and the compare-and-swap instruction in par-

<sup>2</sup> See the accompanying technical report for a formal definition of  $\text{isLock}$ .



ticular. The full proof outline is given the accompanying technical report.

As the name suggests, the resources owned by a shared region are shared between all threads. Atomic instructions are allowed to access and modify resources owned by shared regions, provided they follow the protocol imposed by the region. In the case of the spin-lock, the spin-lock region owns the shared locked field and we thus need to follow the spin-lock protocol to access and modify the locked field. Since the precondition of acquire asserts that the lock is either in the locked or unlocked state, we need to consider two cases.

If the spin-lock region is already locked, then the compare-and-swap fails and we remain in the locked state. This results in the following proof obligation:

$$\langle \triangleright I_L(\mathbf{this}, R, n) \rangle$$

$$\text{CAS}(\mathbf{this}.locked, \mathbf{true}, \mathbf{false})$$

$$\langle r. \triangleright I_L(\mathbf{this}, R, n) * r = \mathbf{false} \rangle$$

That is, if locked contains the value true from the point of view of a thread  $t$ , then CAS'ing from false to true in thread  $t$  will fail. This is easily shown to hold by rule A-CAS-FALSE.

If the spin-lock region is unlocked, then the compare-and-swap may or may not succeed, depending on whether the buffered release has made it to main memory and which thread performed the buffered release. If it succeeds, the acquiring thread transitions the shared region to the locked state and takes ownership of the resource invariant; otherwise, the shared region remains in the unlocked state. This results in the following proof obligation:

$$\langle \triangleright I_U(\mathbf{this}, R, n) \rangle$$

$$\text{CAS}(\mathbf{this}.locked, \mathbf{true}, \mathbf{false})$$

$$\langle r. \exists y \in \{U, L\}. \triangleright I_y(\mathbf{this}, R, n) * Q(y, r, n) \rangle$$

where

$$Q(y, r, n) = (y = U * r = \mathbf{false}) \vee$$

$$(y = L * [\text{REL}]_1^n * \ulcorner R \urcorner * r = \mathbf{true})$$

Rewriting the explicit stabilization to a disjunction and commuting in  $\triangleright$ , this reduces to the following proof obligation:

$$\langle \exists t : \text{Tid}. \triangleright \ulcorner x.locked \urcorner \mapsto \mathbf{false} * R * [\text{REL}]_1^n \urcorner \vee$$

$$\langle \triangleright \ulcorner x.locked \urcorner \mapsto \mathbf{true} \urcorner \mathcal{U}_t \triangleright \ulcorner x.locked \urcorner \mapsto \mathbf{false} * R * [\text{REL}]_1^n \urcorner \rangle$$

$$\text{CAS}(\mathbf{this}.locked, \mathbf{true}, \mathbf{false})$$

$$\langle r. \exists y \in \{U, L\}. \triangleright I_y(\mathbf{this}, R, n) * Q(y, r, n) \rangle$$

In case the second disjunct holds and there exists buffered releases in store buffer  $t$ , the CAS will succeed if executed by thread  $t$  and fail if executed by any other thread. To prove this obligation, we thus do case analysis on whether  $t$  is our thread or not. This leaves us with three proof obligations (after strengthening the post-condition):

- either the buffered release is in our store buffer

$$\langle \ulcorner \mathbf{this}.locked \urcorner \mapsto \mathbf{true} \urcorner \mathcal{U} \triangleright \ulcorner x.locked \urcorner \mapsto \mathbf{false} * R * [\text{REL}]_1^n \urcorner \rangle$$

$$\text{CAS}(\mathbf{this}.locked, \mathbf{true}, \mathbf{false})$$

$$\langle r. \triangleright I_L(\mathbf{this}, R, n) * [\text{REL}]_1^n * \ulcorner R \urcorner * r = \mathbf{true} \rangle$$

- or in some other thread's store buffer

$$\langle \ulcorner \mathbf{this}.locked \urcorner \mapsto \mathbf{true} \urcorner \overline{\mathcal{U}} \triangleright \ulcorner x.locked \urcorner \mapsto \mathbf{false} * R * [\text{REL}]_1^n \urcorner \rangle$$

$$\text{CAS}(\mathbf{this}.locked, \mathbf{true}, \mathbf{false})$$

$$\langle r. \triangleright I_U(\mathbf{this}, R, n) * r = \mathbf{false} \rangle$$

- or it has already been flushed

$$\langle \ulcorner x.locked \urcorner \mapsto \mathbf{false} * R * [\text{REL}]_1^n \urcorner \rangle$$

$$\text{CAS}(\mathbf{this}.locked, \mathbf{true}, \mathbf{false})$$

$$\langle r. \triangleright I_L(\mathbf{this}, R, n) * [\text{REL}]_1^n * \ulcorner R \urcorner * r = \mathbf{true} \rangle$$

These three proof obligations are easily discharged using rules A-CAS-TRUE and A-CAS-FALSE.

Note that our TSO logic forces us to consider exactly those four cases that intuitively one has to consider when considering a TSO weak memory model.

## 6. Soundness

We prove soundness of iCAP-TSO with respect to a kind of Kripke model in which worlds consist of allocated regions and their associated protocols. Since iCAP-TSO inherits iCAP's impredicative protocols [16], worlds need to be recursively defined. Hence we use a meta-theory that supports the definition of sets by guarded recursion, namely the so-called internal language of the topos of trees [4]. As the overall structure of the model construction remains unchanged from iCAP, we elide the details. We refer the interested reader to the accompanying technical report [12] for details and proofs.

Following the Views framework [7], TSO assertions (terms of type  $\text{Prop}_{\text{TSO}}$ ) are modeled as predicates over *instrumented states*. In addition to the underlying machine state, instrumented states contain shared regions and protocols. Soundness is proven by relating the machine semantics with an instrumented semantics that, for instance, enforces that clients obey the chosen protocols when accessing shared state. This relation is expressed through an erasure function,  $[-]$ , that erases an instrumented state to a set of machine states.

The soundness theorem is stated in terms of the following  $eval(\mu, T, q)$  predicate, which asserts that for any terminating execution of the thread pool  $T$  from initial state  $\mu$ , the predicate  $q$  must hold for the terminal state and thread pool. The  $eval$  predicate is defined as a guarded recursive predicate (the recursive occurrence of  $eval$  is guarded by  $\triangleright$ ), to express that each step of evaluation in the machine semantics corresponds to a step in the topos of trees.

$$eval(\mu, T, q) \stackrel{\text{def}}{=} (\text{irr}(\mu, T) \wedge (\mu, T) \in q) \vee$$

$$(\forall T', \mu'. (\mu, T) \rightarrow (\mu', T') \Rightarrow \triangleright eval(\mu', T', q))$$

Here  $\text{irr}(\mu, T)$  means that  $(\mu, T)$  is irreducible. We can now state the soundness of iCAP-TSO.

**Theorem 3 (Soundness).** *If  $\{P\}e\{r.Q\}$  and  $\mu \in \llbracket [P] \rrbracket(t)$  then*

$$eval(\mu, [t \mapsto e], \lambda(\mu', T). \mu' \in \llbracket [Q] \rrbracket(t)(T(t)))$$

This theorem expresses that if a specification  $\{P\}e\{r.Q\}$  holds and the execution of the thread pool  $[t \mapsto e]$  with a single thread  $t$  from an initial state  $\mu$  in the erasure of  $P$  terminates (including threads spawned by  $t$ ), then the terminal state is in the erasure of  $Q$  instantiated with the return value  $T(t)$  of thread  $t$ .

## 7. Related Work

Our work builds directly on iCAP [16], which is an extension of separation logic for modular reasoning about concurrent higher-order programs with shared mutable state. Our work extends the model of iCAP with store buffers to implement a TSO memory model, extends the iCAP logic with TSO-connectives for reasoning about these store buffers and, importantly, defines a new logic for reasoning about sequentially consistent clients.

There has been a lot of previous work on reasoning about TSO-like and even weaker memory models. Owens [13] defines a property on the set of sequentially consistent traces of an x86 program, which ensures that every TSO trace of the program has a "memory equivalent" sequentially consistent trace. Owens shows how this property allows clients of synchronization primitives to reason as if the underlying memory model was sequentially consistent, despite

racy implementations of these synchronization primitives. In particular, he proves that all traces of spin-lock well-synchronized x86 programs satisfy this property. Unfortunately, Owens' approach is non-compositional: while Owens proves similar results for multiple synchronization primitives *in isolation*, these results do not apply to clients that *combine* two or more of these synchronization primitives. Our approach is compositional in this sense and we can reason about clients that combine multiple libraries with racy implementations that have been verified *independently*, without having to re-verify the library implementations.

Ridge [14] proposes a Rely-Guarantee-based proof system for reasoning about assembly code on the x86-TSO memory model. To ensure sound reasoning in the presence of a TSO memory model, the proof system enforces that the rely relation includes possible interference from write buffers. Ridge uses this proof system to verify an implementation of Simpson's four slot algorithm. Ridge's proof system does not provide anything like the fiction of sequential consistency provided by our logic or Owens' approach; in particular, to reason about clients of Simpson's four slot algorithm in Ridge's system, we would still have to explicitly reason about possible interference from store buffers.

In his thesis, Wehrman [18] proposes a program logic for reasoning about low-level code in a language with built-in locks and a TSO memory model. Wehrman's logic is based on separation logic, which he also extends with temporal operators to reason about the orderings introduced by the weak memory model. Wehrman's temporal operators differ from ours; in particular, Wehrman's temporal operators concern all store buffers, whereas our temporal operator refers to a particular store buffer. Wehrman's objective is to allow local reasoning about racy *implementations* and his logic does not provide any fiction of sequential consistency to simplify reasoning about *clients*. In his conclusion, Wehrman points out that it would be beneficial with a program logic for a weak memory model that reduces to a standard logic for non-racy clients. This is exactly what our SC logic and fiction of sequential consistency provides!

Another verification technique for programs in the TSO memory model, which does not involve explicit reasoning about store buffers, is to restrict programs to follow a certain programming discipline that ensures that all program behaviors can be simulated by a sequentially consistent machine. Cohen and Schirmer [6] propose such a programming discipline based on ownership. With this approach one can, of course, only reason about programs that follow the discipline. For instance, as Cohen and Schirmer remark, an efficient spin-lock implementation with a buffered release, like the one we verify in Section 5, does not obey their programming discipline. Essentially, the problem is one of compositionality: The programming discipline is formalized as a property of the set of reachable states of a *complete* program. This is a non-compositional property and the authors do not provide a compositional method for reasoning about it.

Gotsman et al. [10] propose another approach for providing clients with a fiction of sequential consistency based on linearizability. By relating racy library implementations on a TSO architecture with abstract specifications on an SC architecture, they can reason about data-race free clients that call racy libraries using an SC memory model. Their approach is only compositional for *non-interacting* libraries: assuming two libraries are linearizable with respect to SC specifications and the set of locations accessed by the libraries are disjoint, then the composition of the two libraries is linearizable with respect to the composition of the specifications. Their approach also requires the client and any libraries to be non-interacting. Our work does not suffer from either of these restrictions. Furthermore, these restrictions make it very unclear how to extend their approach to higher-order languages, as higher-order languages blur the lines between library and client. In con-

trast, our work extends easily to a higher-order programming language, by following the iCAP methodology [16].

## 8. Conclusion and Future Work

We have presented a new proof system, iCAP-TSO, to support modular and scalable reasoning for a language with a TSO memory model. The proof system consists of two logics, a TSO logic for reasoning about racy low-level code and an SC logic for reasoning about well-behaved code. This class of well-behaved code includes standard mutable data structures without any sharing but also well-synchronized code *with sharing*. For this class of well-behaved code the SC logic allows us to reason much as in standard separation logic, without having to worry about the underlying weak memory model.

In particular, we use the TSO logic to verify an efficient spin-lock implementation against an SC specification. This allows us to reason about well-synchronized clients with sharing entirely from within the SC logic! We illustrate this by verifying a standard list library and using it to verify a shared bag library. Both of these libraries can be verified entirely in the SC logic and their proofs are as in separation logics for sequentially consistent memory models.

We think of iCAP-TSO as a first steps towards more automated/interactive tools for reasoning about the TSO memory model. In this paper we have focused on the foundational issues of constructing a logic that allows simple reasoning for well-behaved code. As future work it would be interesting to try to extend tools like [5, 9] to support mostly automated verification in the SC logic, while requiring more interactive verification of low-level racy code in the TSO logic.

## Acknowledgements

We thank Mark Batty, Aleš Bizjak, Susmit Sarkar, and Peter Sewell for helpful discussions on this work. This research was supported in part by the ModuRes Sapere Aude Advanced Grant from The Danish Council for Independent Research for the Natural Sciences (FNU).

## References

- [1] Linux kernel mailing list, Nov. 1999. spin\_unlock optimization(i386).
- [2] A. W. Appel, P.-A. Melliès, C. D. Richards, and J. Vouillon. A Very Modal Model of a Modern, Major, General Type System. In *Proceedings of POPL*, 2007.
- [3] B. Biering, L. Birkedal, and N. Torp-Smith. BI-Hyperdoctrines, Higher-order Separation Logic, and Abstraction. *ACM TOPLAS*, 2007.
- [4] L. Birkedal, R. Møgelberg, J. Schwinghammer, and K. Støvring. First Steps in Synthetic Guarded Domain Theory: Step-Indexing in the Topos of Trees. In *Proceedings of LICS*, 2011.
- [5] A. Chlipala. Mostly-automated Verification of Low-level Programs in Computational Separation Logic. In *Proceedings of PLDI*, 2011.
- [6] E. Cohen and B. Schirmer. From Total Store Order to Sequential Consistency: A Practical Reduction Theorem. In *Proceedings of ITP*, 2010.
- [7] T. Dinsdale-Young, L. Birkedal, P. Gardner, M. Parkinson, and H. Yang. Views: Compositional Reasoning for Concurrent Programs. In *Proceedings of POPL*, 2013.
- [8] T. Dinsdale-Young, M. Dodds, P. Gardner, M. J. Parkinson, and V. Vafeiadis. Concurrent Abstract Predicates. In *Proceedings of ECOOP*, 2010.
- [9] A. Gotsman, J. Berdine, B. Cook, and M. Sagiv. Thread-modular Shape Analysis. In *Proceedings of PLDI*, 2007.
- [10] A. Gotsman, M. Musuvathi, and H. Yang. Show No Weakness: Sequentially Consistent Specifications of TSO Libraries. In *Proceedings of DISC*, 2012.

- [11] A. Hobor, R. Dockins, and A. Appel. A theory of indirection via approximation. In *Proceedings of POPL*, 2010.
- [12] NN. A Separation Logic for Fictional Sequential Consistency. Technical report, 2013.
- [13] S. Owens. Reasoning about the Implementation of Concurrency Abstractions on x86-TSO. In *Proceedings of ECOOP*, 2010.
- [14] T. Ridge. A Rely-Guarantee proof system for x86-TSO. In *Proceedings of VSTTE*, 2010.
- [15] P. Sewell, S. Sarkar, S. Owens, F. Zappa Nardelli, and M. O. Myreen. x86-TSO: A Rigorous and Usable Programmers Model for x86 Multiprocessors. In *Communications of the ACM*, 2010.
- [16] K. Svendsen and L. Birkedal. Impredicative Concurrent Abstract Predicates. 2013. Under Submission.
- [17] A. Turon, D. Dreyer, and L. Birkedal. Unifying Refinement and Hoare-Style Reasoning in a Logic for Higher-Order Concurrency. In *Proceedings of ICFP*, 2013.
- [18] I. Wehrman. *Weak-Memory Local Reasoning*. PhD thesis, University of Texas, 2012. Dissertation draft.
- [19] J. Wickerson. *Concurrent verification for sequential programs*. PhD thesis, University of Cambridge, 2012.



# A Concurrent Logical Relation

Lars Birkedal, Filip Sieczkowski, and Jacob Thamsborg

IT University of Copenhagen, Denmark  
{birkedal, fisi, thamsborg}@itu.dk

---

## Abstract

We present a logical relation for showing the correctness of program transformations based on a new type-and-effect system for a concurrent extension of an ML-like language with higher-order functions, higher-order store and dynamic memory allocation.

We show how to use our model to verify a number of interesting program transformations that rely on effect annotations. In particular, we prove a Parallelization Theorem, which expresses when it is sound to run two expressions in parallel instead of sequentially. The conditions are expressed solely in terms of the types and effects of the expressions. To the best of our knowledge, this is the first such result for a concurrent higher-order language with higher-order store and dynamic memory allocation.

**1998 ACM Subject Classification** F.3.1 Specifying and Verifying and Reasoning about Programs

**Keywords and phrases** logics of programs, verification, semantics

**Digital Object Identifier** 10.4230/LIPIcs.xxx.yyy.p

## 1 Introduction

Relational reasoning about program equivalence is useful for reasoning about the correctness of program transformations, data abstraction (representation independence), compiler correctness, etc. The standard notion of program equivalence is contextual equivalence and in recent years, there have been many improvements in reasoning methods for higher-order ML-like languages with general references, based on bisimulations, e.g., [1, 2, 3], traces [4], game semantics [5], and Kripke logical relations, e.g., [6, 7, 8, 9].

In this paper we present the first Kripke logical relation for reasoning about equivalence of a *concurrent* higher-order ML-like language with higher-order store and dynamic memory allocation.

To state and prove useful equivalences about concurrent programs, it is necessary to have some way of restricting the contexts under which one proves equivalences. This point was made convincingly in the recent paper by Liang et. al. [10], who presented a rely-guarantee-based simulation for verifying concurrent program transformations for a first-order imperative language (with first-order store). Here is a very simple example illustrating the point. Consider two expressions

$$e_1 \equiv x := 1; y := 1 \quad \text{and} \quad e_2 \equiv y := 1; x := 1.$$

Here  $x$  and  $y$  are variables of type `refint`. The expressions  $e_1$  and  $e_2$  are not contextually equivalent. (To see why, consider expression  $e_3 \equiv x := 0; y := 0$ , and note that running  $e_1$  in parallel with  $e_3$  may result in a state with  $!x = 0$  and  $!y = 1$ , but that cannot be the case when we run  $e_2$  in parallel with  $e_3$ .) The issue is, of course, that the context may also modify the references  $x$  and  $y$ . On the other hand, if we know that no other threads have access to  $x$  or  $y$ , then it should be the case that  $e_1$  and  $e_2$  are equivalent. We can express this restriction on the contexts using a refined region-based type-and-effect system.



© Lars Birkedal, Filip Sieczkowski and Jacob Thamsborg;  
licensed under Creative Commons License NC-ND

Conference title on which this volume is based on.

Editors: Billy Editor, Bill Editors; pp. 61–91



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

We first recall that a type-and-effect system is a type system that classifies programs according to which side effects the programs may have. A variety of effect systems have been proposed for higher-order programming languages, e.g., [11, 12, 13], see [14] for a recent overview. Effect systems can often be understood as specifying the results of a static analysis, in the sense that it is possible to automatically infer types and effects. Effect systems can be used for different purposes: they were originally proposed by Lucassen and Gifford [12] for parallelization purposes but they have also, e.g., been used as the basis for implementing ML using a stack of regions for memory management [13, 15]. In a recent series of papers, Benton et. al. have argued that another important point of effect systems is that they can be used as the basis for effect-based program transformations, e.g., compiler optimizations, [16, 17, 18, 19], see also [20]. The idea is that certain program transformations are only sound under additional assumptions about which effects program phrases may, or rather may not, have.

Now, returning to our example, we refine the types of  $x$  and  $y$  to be  $\text{ref}_\rho \text{int}$  and  $\text{ref}_\sigma \text{int}$ , respectively. Intuitively, this expresses that  $x$  and  $y$  are references in different regions, but it does not put any restrictions on whether other threads may access  $x$  or  $y$ . Thus, when we type  $e_1$  and  $e_2$  we will use two contexts of region variables, one for public regions that can be used by other concurrently running threads, and one for private regions that are under the control of the present thread. This idea is inspired by recent work on concurrent separation logic, e.g., [21, 22, 23, 24]. We use a vertical bar to separate public and private regions: the typing context

$$\rho, \sigma \mid \emptyset \mid x : \text{ref}_\rho \text{int}, y : \text{ref}_\sigma \text{int}$$

expresses that  $\rho$  and  $\sigma$  are public regions, whereas the typing context

$$\emptyset \mid \rho, \sigma \mid x : \text{ref}_\rho \text{int}, y : \text{ref}_\sigma \text{int}$$

expresses that  $\rho$  and  $\sigma$  are private regions. The expressions  $e_1$  and  $e_2$  are well-typed in the latter context and, with this refined typing, they are indeed contextually equivalent, because our type-and-effect system guarantees that no well-typed context can access regions  $\rho$  or  $\sigma$ . (The expressions are also well-typed in the former context, but not contextually equivalent with that refined typing.)

In this paper we present a step-indexed Kripke logical relations model of a type-and-effect system with public and private regions for a concurrent higher-order language with general references. Our model is constructed over the operational semantics of the programming language, and builds on recent work by Thamsborg and Birkedal on logical relations for the sequential sub-language [20]. Note that the type-and-effect annotations are just annotations; the operational semantics of the language is standard and regions only exist in our semantic model, not in the operational semantics.

As an important application of our model we prove a Parallelization Theorem, which expresses when it is sound to run two expressions in parallel instead of sequentially. To the best of our knowledge, this is the first such result for a higher-order language with higher-order store and dynamic memory allocation. Here is a very simple instance of the theorem. Consider two expressions

$$e_1 \equiv y := !x + !y \quad \text{and} \quad e_2 \equiv z := !x + !z,$$

each well-typed in a context

$$\emptyset \mid \rho_x, \rho_y, \rho_z \mid x : \text{ref}_{\rho_x} \text{int}, y : \text{ref}_{\rho_y} \text{int}, z : \text{ref}_{\rho_z} \text{int},$$

i.e., where  $x$ ,  $y$ , and  $z$  are references in distinct private regions. In this context, running  $e_1$  and  $e_2$  sequentially is contextually equivalent to running  $e_1$  and  $e_2$  in parallel. Intuitively, this also makes sense:  $e_1$  and  $e_2$  update references in distinct regions, and it is unproblematic that they both read (but not write) from the same region.

As mentioned, this was a simple instance of the Parallelization Theorem. We stress that the theorem is expressed solely in terms of the type and effects of the expressions  $e_1$  and  $e_2$ , so a compiler may automatically infer that it is safe to parallelize two expressions by looking at the inferred effect types, and without reasoning about all interleavings. Moreover, the theorem applies to contexts and expressions with general higher types (not just with references to integers and unit types). Note that the distinction between private and public regions is also crucial here (parallelization would not be sound if the effects of the expressions were on public regions).

Our type-and-effect system crucially also includes a region-masking rule. Traditionally, this rule has been used to hide local effects on regions, which makes it possible to view a computation as pure even if it uses effects locally and makes the effect system stronger, in the sense that it can justify more program transformations. Here we also observe that the masking rule can be used for introducing private regions, since the masking rule intuitively guarantees that effects on a region are not leaked to the context. It is well-known that region-masking makes the model construction for a sequential language technically challenging, see the extensive discussion in [20]. Here it is yet more challenging because of concurrency; we explain how our model ensures soundness of the masking rule in Section 3.

The extension with concurrency also means that when we define the logical relation for contextual approximation and relate two computations  $e_1$  and  $e_2$ , we cannot simply require relatedness after  $e_1$  has completed evaluation (as in the sequential case), since other threads should be allowed to execute as well. We explain our approach to relating concurrent computations in Section 3; it is informed by recent soundness proofs of unary models of concurrent separation logic [25, 26].

Another challenge arises from the fact that since our language includes dynamically allocated general references, the existence of the logical relation is non-trivial; in particular, the set of Kripke worlds must be recursively defined. Here we build on our earlier work [27] and define the worlds as a solution to a recursive metric-space equation. Indeed, to focus on the essential new aspects due to the extension with concurrency, we deliberately choose to use the exact same notion of worlds as we used for the sequential sub-language in [20]. In the same vein, we here consider a monomorphically typed higher-order programming language with general references, but leave out universal and existential types as well as recursive types. However, we want to stress that since our semantic techniques (step-indexed Kripke logical relations over recursively defined worlds) do indeed scale well to universal, existential, and recursive types, e.g. [27, 9], it is possible to extend our model to a language with such types. We conjecture that it is also possible to extend our model to richer effect systems involving region and effect polymorphism, but we have not done so yet.

## 2 Language and Typing

We consider a standard call-by-value lambda calculus with general references, and extended with parallel composition and an atomic construct. We assume countably infinite, pairwise disjoint sets of *region variables*  $\mathcal{RV}$  (ranged over by  $\rho$ ), *locations*  $\mathcal{L}$  (ranged over by  $l$ ) and program variables (ranged over by  $x, y, f$ ). As usual, the reduction relation is between configurations,  $(e | h) \mapsto (e' | h')$  where heaps  $\mathcal{H}$  are finite maps from locations to

$\begin{aligned} \pi &::= rd_\rho \mid wr_\rho \mid al_\rho \\ \varepsilon &::= \pi_1, \dots, \pi_n \\ \tau &::= 1 \mid \text{int} \mid \tau_1 \times \tau_2 \mid \text{ref}_\rho \tau \\ &\quad \mid \tau_1 \xrightarrow{\Pi, \Lambda} \tau_2 \\ v &::= x \mid \langle \rangle \mid \langle v_1, v_2 \rangle \\ &\quad \mid \text{fun } f(x).e \mid l \\ e &::= v \mid \text{proj}_i v \mid v e \mid \text{ref } v \mid !v \\ &\quad \mid v_1 := v_2 \mid \text{par } e_1 \text{ and } e_2 \\ &\quad \mid \text{cas } (v_1, v_2, v_3) \mid \text{atomic } e \\ E &::= [] \mid v E \mid \text{par } E \text{ and } e_2 \\ &\quad \mid \text{par } e_1 \text{ and } E \end{aligned}$	$\begin{aligned} (E[\text{proj}_i \langle v_1, v_2 \rangle] \mid h) &\mapsto (E[v_i] \mid h) \\ (E[(\text{fun } f(x).e) v] \mid h) &\mapsto (E[e[\text{fun } f(x).e/f, v/x]] \mid h) \\ (E[\text{ref } v] \mid h) &\mapsto (E[l \mid h[l \mapsto v]]) \text{ if } l \notin \text{dom}(h) \\ (E[l := v] \mid h) &\mapsto (E[\langle \rangle \mid h[l := v]]) \text{ if } l \in \text{dom}(h) \\ (E[!l] \mid h) &\mapsto (E[h(l)] \mid h) \text{ if } l \in \text{dom}(h) \\ (E[\text{par } v_1 \text{ and } v_2] \mid h) &\mapsto (E[\langle v_1, v_2 \rangle] \mid h) \\ (E[\text{cas } (l, n_1, n_2)] \mid h) &\mapsto (E[1] \mid h[l := n_2]) \\ &\quad \text{if } l \in \text{dom}(h) \text{ and } h(l) = n_1 \\ (E[\text{cas } (l, n_1, n_2)] \mid h) &\mapsto (E[0] \mid h) \\ &\quad \text{if } l \in \text{dom}(h) \text{ and } h(l) \neq n_1 \\ (E[\text{atomic } e] \mid h) &\mapsto (E[v] \mid h') \\ &\quad \text{if } (e \mid h) \mapsto^* (v \mid h') \\ (E[\text{atomic } e] \mid h) &\mapsto (E[\text{atomic } e] \mid h) \end{aligned}$
---	---

■ **Figure 1** Syntax

■ **Figure 2** Operational semantics

values. Figures 1 and 2 give the syntax and operational semantics; we denote the set of expressions  $\mathcal{E}$  and the set of values  $\mathcal{V}$ . The evaluation contexts allow parallel evaluation inside `par` expressions, and there is a new primitive reduction covering the case when the two subcomputations have terminated. For technical simplicity, we allow an `atomic`  $e$  expression to reduce to itself, possibly introducing more divergence than the diverging behaviours of  $e$ . The syntax is kept minimal; in examples we may use additional syntactic sugar, e.g., writing `let  $x = e_1$  in  $e_2$`  for `(fun  $f(x).e_2$ )  $e_1$`  for some fresh  $f$ . For  $e \in \mathcal{E}$ , we write  $\text{FV}(e)$  and  $\text{FRV}(e)$  for the sets of free program variables and region variables, respectively; also we define  $\text{rds } \varepsilon = \{\rho \in \mathcal{RV} \mid rd_\rho \in \varepsilon\}$  and similarly for writes and allocation.

The form of the judgments of our type-and-effect system is standard with one important refinement: regions are partitioned into *public* and *private* regions, with the purpose of restricting interference from the environment. In greater detail, a typing judgement looks like this:

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

The  $\Gamma$ ,  $e$  and  $\tau$  are the usual: the *variable context*  $\Gamma$  assigns types to program variables in the expression  $e$ , with the resulting type of  $\tau$ . To get an idea of — or rather an upper bound of — the side-effects of  $e$ , we split the heap into *regions*; these are listed in  $\Pi$  and  $\Lambda$ . We track memory accesses by adding a set  $\varepsilon$  of *effects* of the form  $rd_\rho$ ,  $wr_\rho$  and  $al_\rho$ , where  $\rho$  is a region. Roughly, a computation with effect  $rd_\rho$  may read one or more locations in region  $\rho$ , and similarly for writes and allocation. This setup goes back to Lucassen and Gifford [12].

The novelty, as mentioned in the Introduction, is our partition of regions into the *public* ones  $\Pi$  and the *private* ones  $\Lambda$ . As opposed to the rest of the judgment, this public-private division does not make promises about the behavior of  $e$ . Instead, it states the expectations that  $e$  has of the environment: threads running in parallel with  $e$  may — in a well-typed manner — read, write and allocate in the public regions but must leave the private regions untouched.

When running parallel threads, the private regions of the parent are shared between the children, and so are public from their point of view; this is reflected in the typing



$$\begin{array}{c}
\frac{}{\Pi | \Lambda | \Gamma, x : \tau \vdash x : \tau, \emptyset} \quad \frac{}{\Pi | \Lambda | \Gamma \vdash \langle \rangle : \mathbf{1}, \emptyset} \quad \frac{}{\Pi | \Lambda | \Gamma \vdash v : \tau_1 \times \tau_2, \varepsilon} \\
\frac{\Pi | \Lambda | \Gamma \vdash v_1 : \tau_1, \varepsilon_1 \quad \Pi | \Lambda | \Gamma \vdash v_2 : \tau_2, \varepsilon_2}{\Pi | \Lambda | \Gamma \vdash \langle v_1, v_2 \rangle : \tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2} \quad \frac{\Pi | \Lambda | \Gamma, f : \tau_1 \xrightarrow{\Pi, \Lambda} \tau_2, x : \tau_1 \vdash e : \tau_2, \varepsilon}{\Pi | \Lambda | \Gamma \vdash \text{fun } f(x).e : \tau_1 \xrightarrow{\Pi, \Lambda} \tau_2, \emptyset} \\
\frac{\Pi | \Lambda | \Gamma \vdash v : \tau_1 \xrightarrow{\Pi, \Lambda} \tau_2, \varepsilon_1 \quad \Pi | \Lambda | \Gamma \vdash e : \tau_1, \varepsilon_2}{\Pi | \Lambda | \Gamma \vdash v e : \tau_2, \varepsilon_1 \cup \varepsilon_2 \cup \varepsilon} \quad \frac{\Pi | \Lambda | \Gamma \vdash v : \tau, \varepsilon \quad \rho \in \Pi, \Lambda}{\Pi | \Lambda | \Gamma \vdash \text{ref } v : \text{ref}_\rho \tau, \varepsilon \cup \{al_\rho\}} \\
\frac{\Pi | \Lambda | \Gamma \vdash v_1 : \text{ref}_\rho \tau, \varepsilon_1 \quad \Pi | \Lambda | \Gamma \vdash v_2 : \tau, \varepsilon_2}{\Pi | \Lambda | \Gamma \vdash v_1 := v_2 : \mathbf{1}, \varepsilon_1 \cup \varepsilon_2 \cup \{wr_\rho\}} \quad \frac{\Pi | \Lambda | \Gamma \vdash v : \text{ref}_\rho \tau, \varepsilon}{\Pi | \Lambda | \Gamma \vdash !v : \tau, \varepsilon \cup \{rd_\rho\}} \\
\frac{\Pi | \Lambda, \rho | \Gamma \vdash e : \tau, \varepsilon}{\Pi | \Lambda | \Gamma \vdash e : \tau, \varepsilon - \rho} \quad (\rho \notin \text{FRV}(\Gamma, \tau)) \quad \frac{\cdot | \Pi, \Lambda | \Gamma \vdash e : \tau, \varepsilon}{\Pi | \Lambda | \Gamma \vdash \text{atomic } e : \tau, \varepsilon} \quad (\text{als } \varepsilon \subseteq \text{rds } \varepsilon \cap \text{wrs } \varepsilon) \\
\frac{\Pi, \Lambda | \cdot | \Gamma \vdash e_1 : \tau_1, \varepsilon_1 \quad \Pi, \Lambda | \cdot | \Gamma \vdash e_2 : \tau_2, \varepsilon_2}{\Pi | \Lambda | \Gamma \vdash \text{par } e_1 \text{ and } e_2 : \tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2} \\
\frac{\Pi | \Lambda | \Gamma \vdash v_1 : \text{ref}_\rho \text{int}, \varepsilon_1 \quad \Pi | \Lambda | \Gamma \vdash v_2 : \text{int}, \varepsilon_2 \quad \Pi | \Lambda | \Gamma \vdash v_3 : \text{int}, \varepsilon_3}{\Pi | \Lambda | \Gamma \vdash \text{cas } (v_1, v_2, v_3) : \text{int}, \{wr_\rho, rd_\rho\} \cup \varepsilon_1 \cup \varepsilon_2 \cup \varepsilon_3} \\
\frac{\Pi | \Lambda | \Gamma \vdash e : \tau_1, \varepsilon_1 \quad \Pi, \Lambda \vdash \tau_1 \leq \tau_2 \quad \varepsilon_1 \subseteq \varepsilon_2}{\Pi | \Lambda | \Gamma \vdash e : \tau_2, \varepsilon_2} \quad (\text{FRV}(\varepsilon_2) \subseteq \Pi, \Lambda) \\
\frac{}{\Theta \vdash \tau \leq \tau} \quad (\text{FRV}(\tau) \subseteq \Theta) \quad \frac{\Theta \vdash \tau_1 \leq \tau'_1 \quad \Theta \vdash \tau_2 \leq \tau'_2}{\Theta \vdash \tau_1 \times \tau_2 \leq \tau'_1 \times \tau'_2} \\
\frac{\Theta \vdash \tau'_1 \leq \tau_1 \quad \Theta \vdash \tau_2 \leq \tau'_2 \quad \varepsilon_1 \subseteq \varepsilon_2 \quad \Pi_1 \subseteq \Pi_2 \quad \Lambda_1 \subseteq \Lambda_2}{\Theta \vdash \tau_1 \xrightarrow{\Pi_1, \Lambda_1} \tau_2 \leq \tau'_1 \xrightarrow{\Pi_2, \Lambda_2} \tau'_2} \quad (\text{FRV}(\varepsilon_2), \Pi_2, \Lambda_2 \subseteq \Theta)
\end{array}$$

■ **Figure 3** Typing and subtyping relations. Notice that for a typing judgement  $\Pi | \Lambda | \Gamma \vdash e : \tau, \varepsilon$  we always have  $\text{FRV}(\Gamma, \tau, \varepsilon) \subseteq \Pi \cup \Gamma$ .

rule for parallel composition, c.f. Figure 3. Note that the parent thread only continues once both children have terminated; as a consequence, the parent regains ownership of its private regions before it goes on. Running an expression atomically temporarily makes all regions private. The side condition is a technical necessity. Finally, new, private regions are introduced by the so-called masking rule:

$$\frac{\Pi | \Lambda, \rho | \Gamma \vdash e : \tau, \varepsilon}{\Pi | \Lambda | \Gamma \vdash e : \tau, \varepsilon - \rho} \quad (\rho \notin \text{FRV}(\Gamma, \tau))$$

The subtraction of  $\rho$  in the conclusion removes any read, write or allocation effects tagged with  $\rho$ . The reading of the masking rule is that we make a brand new, empty region  $\rho$  for  $e$  to use, but once  $e$  has terminated we forget about  $\rho$  again; this works out since the side condition prevents  $e$  from leaking locations from  $\rho$ . Traditionally, the masking rule has been used to do memory-management [13] as well as a means of hiding local effects to facilitate effect-based program transformations [17, 20]. Here we make another use of the rule: we observe that, moreover,  $e$  cannot leak locations from  $\rho$  while running and so  $\rho$  is a *private* region for the duration of  $e$ . After all, the only means of inter-thread communication is shared memory. Note that from the perspective of the context, this rule allows to *remove* a private region, and prepare a setup for application of the parallel composition.

All the typing rules are in Figure 3. Note how reference types are tagged with the region

where the location resides and that function arrows are tagged with the latent effects as well as with the public and private regions that the function expects; the latter is natural once we remember that a function is basically just a suspended, well-typed expression.

Because of the nondeterminism arising from `par` and shared references, the definition of contextual equivalence could take into account both may- and must-convergence. In this paper we only consider may-equivalence and formally we define (may-) contextual approximation by:

► **Definition 1.**  $\Pi | \Lambda | \Gamma \vdash e \lesssim_{\downarrow} e' : \tau, \varepsilon$  if and only if for all  $h$  and  $C$  typed such that  $C : (\Pi | \Lambda | \Gamma \vdash \tau, \varepsilon) \rightsquigarrow (\cdot | \cdot | \cdot \vdash \text{int}, \emptyset)$ , whenever  $(C[e] | h) \downarrow$  then  $(C[e'] | h) \downarrow$ .

Here, as usual,  $(e | h) \downarrow$  means that  $(e | h) \mapsto^* (v | h')$  for some value  $v$  and some  $h'$ . Typing for contexts is defined in a standard way; selected rules are presented in Appendix A.1.

Contextual equivalence,  $\Pi | \Lambda | \Gamma \vdash e \approx e' : \tau, \varepsilon$ , is then defined as  $\Pi | \Lambda | \Gamma \vdash e \lesssim_{\downarrow} e' : \tau, \varepsilon$  and  $\Pi | \Lambda | \Gamma \vdash e' \lesssim_{\downarrow} e : \tau, \varepsilon$ . Note that the diverging behaviours introduced by our operational semantics of `atomic`  $e$  do not influence may-contextual equivalence.

### 3 Definition of the logical relation

**Semantic Types and Worlds** We give a Kripke or world-indexed logical relation. This is a fairly standard approach to modeling dynamic allocation; in combination with higher-order store, however, it comes with a fairly standard problem: the *type-world circularity*. Roughly, semantic types are indexed over worlds and worlds contain semantic types, so both need to be defined before the other. A specific instance of this circularity was solved recently by Thamsborg and Birkedal [20] based on metric-space theory developed by Birkedal et. al. [27]; we re-use that solution here. Semantic types (and worlds) are constructed as a fixed-point of an endo-functor on a certain category of metric-spaces. We do not care about that, though; we just give the result of the construction. In addition, we largely ignore the fact that we actually deal in metric spaces and not just plain sets; the little metric machinery we need is deferred to Appendix A.2.

There is a set  $\mathbf{T}$  of *semantic types* and a set  $\mathbf{W}$  of *worlds*; types are world-indexed relations on values and worlds describe the regions and type-layouts of heaps, roughly speaking. Take a type  $\mu \in \mathbf{T}$  and apply it to a world  $w \in \mathbf{W}$  and you get an indexed relation on values, i.e.,  $\mu(w) \subseteq \mathbb{N} \times \mathcal{V} \times \mathcal{V}$ . These relations are downwards closed in the first coordinate; we read  $(k, v_1, v_2) \in \mu(w)$  as saying that  $v_1$  and  $v_2$  are related at type  $\mu$  up to approximation  $k$  assuming world  $w$ .

We assume a countably infinite set of region names  $\mathcal{RN}$ ; a world  $w \in \mathbf{W}$  contains finitely many such  $|w| \subseteq_{\text{fin}} \mathcal{RN}$ . Some of these  $\text{dom}(w) \subseteq |w|$  are *live* and the rest are *dead*. To each live region  $r \in \text{dom}(w)$  we associate a finite partial bijection  $w(r)$  on locations decorated with types, i.e.,  $w(r) \subseteq_{\text{fin}} \mathcal{L} \times \mathcal{L} \times \widehat{\mathbf{T}}$  such that for  $(l_1, l_2, \mu), (m_1, m_2, \nu) \in w(r)$  we have that both  $l_1 = m_1$  and  $l_2 = m_2$  imply  $l_1 = m_1$ ,  $l_2 = m_2$  and  $\mu = \nu$ . We write  $\text{dom}_1(w(r))$  for the set of left hand side locations in the bijection and  $\text{dom}_2(w(r))$  for the right hand side ones; different regions must have disjoint left and right hand side locations. For convenience, we set  $\text{dom}_1^{\mathbf{A}}(w) = \bigcup_{r \in A \cap \text{dom}(w)} \text{dom}_1(w(r))$  whenever  $A \subseteq |w|$ , and we write  $\text{dom}_1(w)$  for  $\text{dom}_1^{|w|}(w)$ , i.e., the set of all left hand side locations. Similarly for the right hand side.

Worlds evolve and types adapt. Triples of two locations and a type can be added to a live region, as long as different regions remain disjoint. Orthogonal to this, one can add a fresh, i.e., neither live nor dead, region name with an associated empty partial bijection. And one can kill any live region, rendering it dead and losing the associated the partial bijection in

the process. The reflexive, transitive closure of all three combined is a preorder  $\sqsubseteq$  on worlds; it is a crucial property of types that they respect this, i.e., that  $w \sqsubseteq w' \implies \mu(w) \subseteq \mu(w')$  for any two  $w, w' \in \mathbf{W}$  and any  $\mu \in \mathbf{T}$ . This is *type monotonicity* and it prevents values from fleeing types over time.

Finally, to tie the knot, there is an isomorphism  $\iota : \widehat{\mathbf{T}} \rightarrow \mathbf{T}$  from the odd types stored in worlds to proper types. Whenever a type is extracted from a world it needs to be coerced by this isomorphism before it can be applied to some world.

**The Logical Relation and Interpretation of Types** Often, a logical relation goes like this: two computations are related if they (from related heaps) reduce to related values (and heaps); this is the *extensional* view: we do not care about the intermediate states. As we consider concurrency, however, a computation can be interrupted and so we need to start caring. In our setup, public regions are accessible from the environment. To address this, we assume that before each reduction step, the public regions hold related values; in return, we promise related values after the step. In other words, the *granularity of extensionality* is just one step for the public regions. For private regions, however, there is no interference and the granularity is an entire computation as usual. This is the fundamental idea; it is how we propose to stay extensional in the face of concurrency.

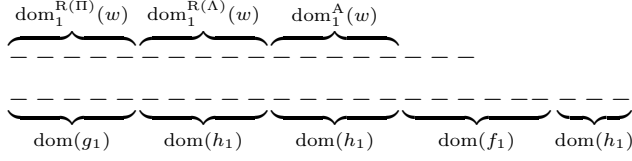
Without further ado, let us look into the cornerstone of our model: the safety relation defined in Figure 6; auxiliary relations are defined in Figure 8. What does it mean to have

$$(k, h_1^\circ, h_2^\circ, e_1, e_2, h_1, h_2) \in \mathbf{safe}_{\tau, \varepsilon}^{\Pi, \Lambda, A, R} w^\circ, w?$$

Overall, it says that after environment interference, we can match the behavior of  $e_1$ , i.e., termination or any one-step reduction, by zero or more steps of  $e_2$ ; match in the sense of (re-)establishing certain relations, including safety itself. Safety is a *local* property of a pair of computations, this is crucial: it has no knowledge of computations running concurrently and  $h_1$  and  $h_2$  are the *local heaps*, i.e., the parts of the global heaps that the  $e_1$  respectively  $e_2$  control exclusively. The computations consider  $R(\Pi)$  to be their public,  $R(\Lambda)$  to be their private and  $A$  to be their *anonymous* regions. The latter intuitively are private regions that have been masked out: they exist only for the duration of these computations, but we have to track them to deny the environment access; this is another difficulty imposed by concurrency. Safety is indexed by a world  $w$  as well; note that worlds are *global* things: all concurrent threads share one world, i.e., they agree about the division of the heap into regions and the types associated to locations. Finally  $k$  is intuitively the number of steps we are safe for,  $h_1^\circ$  and  $h_2^\circ$  are the (private parts of) the initial local heaps,  $\tau$  is the expected return type,  $\varepsilon$  the effects and  $w^\circ$  the initial world.

We unroll the definition in writing. The first pair of big square brackets — the *prerequisites* — translates to ‘the environment interferes’. This yields a new world  $w'$  subject to the constraints of the environment transition relation: no public, private or anonymous regions are killed, and the latter two see no allocation either. The actual contents of the public regions are unknown, but we are free to assume that they hold related values of the proper type, at least where we have read effects; this is the *public heaps*  $g_1$  and  $g_2$  in the precondition relation. In addition we have *frames*  $f_1$  and  $f_2$  that cover the remainder of the world and a triple-split relation that ensures coherence between the domains of corresponding parts of the world and the heaps, see Figures 4 and 8.

The left hand side is irreducible in the *termination branch* and takes one step in the *progress branch*. In either case, we must match this in zero or more steps on the right hand side, not touching the frame; this means finding a future world  $w''$  and relating a number of things. The choice of future world is restricted by the self transition relation: we must not



■ **Figure 4** The left hand side of the triple-split relation. The top dashed line is  $\text{dom}_1(w)$ , the bottom dashed line  $\text{dom}(g_1 \cdot h_1 \cdot f_1)$ . The local heap  $h_1$  has a *private* part matching the private regions, an *anonymous* part matching the anonymous regions and an *off-world* part outside the domain of the world. The frame  $f_1$  must cover regions that are neither public, private nor anonymous.

kill private or public regions, but we can allocate in them, and regions that we know nothing about must be left untouched; this is our promise to the environment. In the termination branch, we are furthermore required to kill off all anonymous regions as the computation is done; any new regions added in the progress branch go to the set of anonymous regions. In both branches, the changes made to the public heap must be well-typed and permitted by the effects and, if we are done, we check the changes made to (the private part of) the local heaps as well; the fact that the public heaps are compared across a single stage and the (private parts of) the local heaps are compared across the entire computations is the crux of the idea of having different granularities of extensionality.

In addition to performing actual allocation, we have the possibility of moving existing locations from, say, the off-world part of the local heap into the public heap or the private part of the local heap; this is a subtle point that permits the actual allocation of new locations and the corresponding extension of the world to be temporarily out of sync.

We have glossed over one aspect of safety: the right hand side takes steps in the ordinary operational semantics, but the left hand side works in the *instrumented operational semantics*. A reduction  $(e | h) \rightarrow_\mu^n (e' | h')$  in the latter implies a similar reduction in former; in addition it counts the steps of a reduction with all atomic commands ‘unfolded’ (with unfolding itself counting one step) and it records all heap accesses; the formal definition is deferred to the Appendix. We need the former for compatibility of the atomic typing rule below: atomic commands really unfold as they execute, hence we must count the number of ‘unfolded’ steps. It is less immediate that we must test the actual reads, writes and allocations, recorded by  $\mu$ , against the effects described by  $\varepsilon$ , as done in the progress branch of safety. But if omitted, our present proof of the Parallelization Theorem falls short, since it relies on the following simple, but crucial commutation property:

► **Lemma 2.** *If we have  $l \notin \mu$  and  $(e | h) \rightarrow_\mu^n (e' | h')$ , then  $(e | h[l \mapsto v]) \rightarrow_\mu^n (e' | h'[l \mapsto v])$ .*

The actual logical relation is given in Figure 7. The existentially quantified  $a \in \mathbb{N}$  is the minimal number of anonymous regions required to run; apart from that it uses safety in a straightforward way. There is some asymmetry to these definitions: the anonymous regions  $A$  are required to exist (and be empty) in the world beforehand, but are killed off in the termination branch; also the precondition on the (private parts of) the initial local heaps is in the logical relation whereas the postcondition lives in the termination branch. The interpretation of types is in Figure 5. Interpreting the function type looks daunting, but a function is just a suspended expression with a single free variable, hence we have to restate most of the logical relation in the definition. Apart from that, we just remark that the  $R(\rho) \not\subseteq \text{dom}(w)$  case of reference interpretation is part of an approach to handling dangling pointers (due to region masking) proposed recently in [20]; similarly for the  $R(\text{FRV}(\varepsilon)) \not\subseteq \text{dom}(w)$ .

$$\begin{aligned}
\llbracket 1 \rrbracket^R w &= \{(k, (), ()) \mid k \in \mathbb{N}\} & \llbracket \text{int} \rrbracket^R w &= \{(k, n, n) \mid k \in \mathbb{N} \wedge n \in \mathbb{Z}\} \\
\llbracket \tau_1 \times \tau_2 \rrbracket^R w &= \{(k, (v_{11}, v_{21}), (v_{12}, v_{22})) \mid (k, v_{11}, v_{12}) \in \llbracket \tau_1 \rrbracket^R w \wedge (k, v_{21}, v_{22}) \in \llbracket \tau_2 \rrbracket^R w\} \\
\llbracket \text{ref}_{\rho\tau} \rrbracket^R w &= \begin{cases} \left\{ \left\{ \begin{array}{l} (k, l_1, l_2) \mid \exists \mu \in \widehat{\mathbf{T}}. (l_1, l_2, \mu) \in w(R(\rho)) \wedge \\ \forall w' \sqsupseteq w. \llbracket \tau \rrbracket^R w' \stackrel{k}{=} (\iota \mu)(w') \end{array} \right\} \right\} & R(\rho) \in \text{dom}(w) \\ \{(k, v_1, v_2) \mid k \in \mathbb{N} \wedge v_1, v_2 \in \mathcal{V}\} & R(\rho) \notin \text{dom}(w) \end{cases} \\
\llbracket \tau_1 \xrightarrow{\Pi, \Lambda} \tau_2 \rrbracket^R w &= \left\{ \left( \begin{array}{l} (k, \text{fun } f(x).e_1, \text{fun } f(x).e_2) \mid \exists a \in \mathbb{N}. \forall j < k. \forall w' \sqsupseteq w. \\ \forall A \subseteq \text{dom}(w'). \forall v_1, v_2 \in \mathcal{V}. \forall h_1, h_2, h'_1, h'_2 \in \mathcal{H}. \\ \left[ \begin{array}{l} R(\text{FRV}(\varepsilon)) \subseteq \text{dom}(w') \wedge A \# R(\Pi \cup \Lambda) \wedge |A| \geq a \wedge w'(A) = \emptyset \wedge \\ (j, v_1, v_2) \in \llbracket \tau_1 \rrbracket^R w' \wedge h'_1 \subseteq h_1 \wedge h'_2 \subseteq h_2 \wedge (j, h'_1, h'_2) \in \mathbf{P}_{\varepsilon}^{\Lambda, R} w' \end{array} \right] \Rightarrow \\ (j, h'_1, h'_2, (\text{fun } f(x).e_1) v_1, (\text{fun } f(x).e_2) v_2, h_1, h_2) \in \mathbf{safe}_{\tau_2, \varepsilon}^{\Pi, \Lambda, A, R} w', w' \end{array} \right) \right\} \\
& \left\{ \begin{array}{l} \{(k, v_1, v_2) \mid k \in \mathbb{N} \wedge v_1, v_2 \in \mathcal{V}\} \\ R(\text{FRV}(\varepsilon)) \subseteq \text{dom}(w) \\ R(\text{FRV}(\varepsilon)) \not\subseteq \text{dom}(w) \end{array} \right.
\end{aligned}$$

■ **Figure 5** Interpretation of types. We require  $R : \mathcal{RV} \rightarrow_{\text{fin}} \mathcal{RN}$  injective with  $\text{FRV}(\tau) \subseteq \text{dom}(R)$ . We assume  $R(\text{FRV}(\tau)) \subseteq |w|$  above, otherwise we define  $\llbracket \tau \rrbracket^R w$  to be the empty set. We get that  $\llbracket \tau \rrbracket^R \in \mathbf{T}$ .

To conclude this subsection we give a theorem that, combined with the upcoming compatibility, means that logical relatedness implies contextual may-approximation. The proof is in Appendix A.4 and it is not hard, but it is worth noting that we need a proof at all: with sequential languages, this is a result one reads off the definition of the logical relation.

► **Theorem 3 (May-Equivalence).** *Assume that  $\cdot \mid \cdot \mid \cdot \models e_1 \preceq e_2 : \text{int}, \emptyset$  holds. Take any  $h_1, h_2 \in \mathcal{H}$ . If there are  $e'_1, h'_1$  with  $(e_1 \mid h_1) \xrightarrow{*} (e'_1 \mid h'_1)$  such that  $\text{irr}(e'_1 \mid h'_1)$  holds, then there is  $n \in \mathbb{Z}$  such that  $e'_1 = \underline{n}$  and  $h'_2$  such that  $(e_2 \mid h_2) \xrightarrow{*} (\underline{n}, h'_2)$ .*

**Compatibility of the Logical Relation** The logical relation is compatible, i.e., respects all typing rules. This is a *sine qua non* of logical relations; it implies the fundamental lemma stating that every well-typed expression is related to itself. And, as discussed just above, it makes the logical relation approximate contextual may-approximation:

► **Theorem 4.**  $\Pi \mid \Lambda \mid \Gamma \models e_1 \preceq e_2 : \tau, \varepsilon$  implies  $\Pi \mid \Lambda \mid \Gamma \vdash e_1 \lesssim_{\downarrow} e_2 : \tau, \varepsilon$ .

Compatibility means that each typing rule induces a lemma by reading the (unary) typing judgments as the corresponding (binary) logical relations. The three most interesting of these have to do with concurrency and the divide between public and private regions; they are listed here and proofs are given in Appendix A.6:

► **Lemma 5.**  $\Pi \mid \Lambda, \rho \mid \Gamma \models e_1 \preceq e_2 : \tau, \varepsilon$  implies  $\Pi \mid \Lambda \mid \Gamma \models e_1 \preceq e_2 : \tau, \varepsilon - \rho$  provided that  $\rho \notin \text{FRV}(\Gamma, \tau)$ .

► **Lemma 6.**  $\cdot \mid \Pi, \Lambda \mid \Gamma \models e_1 \preceq e_2 : \tau, \varepsilon$  implies  $\Pi \mid \Lambda \mid \Gamma \models \text{atomic } e_1 \preceq \text{atomic } e_2 : \tau, \varepsilon$  if also  $\varepsilon \subseteq \text{rds } \varepsilon \cap \text{wrs } \varepsilon$ .

► **Lemma 7.**  $\Pi, \Lambda \mid \cdot \mid \Gamma \models e_1 \preceq e_2 : \tau, \varepsilon$  and  $\Pi, \Lambda \mid \cdot \mid \Gamma \models e_1^{\dagger} \preceq e_2^{\dagger} : \tau^{\dagger}, \varepsilon^{\dagger}$  together imply  $\Pi \mid \Lambda \mid \Gamma \models \text{par } e_1$  and  $e_1^{\dagger} \preceq \text{par } e_2$  and  $e_2^{\dagger} : \tau \times \tau^{\dagger}, \varepsilon \cup \varepsilon^{\dagger}$ .

$$\begin{aligned}
& (k, h_1^\circ, h_2^\circ, e_1, e_2, h_1, h_2) \in \mathbf{safe}_{\tau, \varepsilon}^{\Pi, \Lambda, A, R} w^\circ, w \\
& \iff \\
& \forall j \leq k. \forall w', g_1, g_2, f_1, f_2. \\
& \left[ \mathbf{envtran}^{\Pi, \Lambda, A, R} w, w' \wedge (j, g_1, g_2) \in \mathbf{P}_\varepsilon^{\Pi, R} w' \wedge \right. \\
& \quad \left. (g_1, h_1, f_1, g_2, h_2, f_2) \in \mathbf{splits}^{\Pi, \Lambda, A, R} w' \right] \Rightarrow \\
& \left[ \mathbf{irr}(e_1 | g_1 \cdot h_1 \cdot f_1) \Rightarrow \right. \\
& \quad \exists e_2', w'', h_1', h_2', g_1', g_2'. \\
& \quad (e_2 | g_2 \cdot h_2 \cdot f_2) \mapsto^* (e_2' | g_2' \cdot h_2' \cdot f_2) \wedge \mathbf{selftran}^{\Pi, \Lambda, A, R} w', w'' \wedge \\
& \quad \emptyset = (A \cap \text{dom}(w'')) \cup (\text{dom}(w'') \setminus \text{dom}(w')) \wedge g_1 \cdot h_1 = g_1' \cdot h_1' \wedge \\
& \quad (g_1', h_1', f_1, g_2', h_2', f_2) \in \mathbf{splits}^{\Pi, \Lambda, \emptyset, R} w'' \wedge (j, g_1, g_2, g_1', g_2') \in \mathbf{Q}_\varepsilon^{\Pi, R} w', w'' \wedge \\
& \quad \left. (j, e_1, e_2') \in \llbracket \tau \rrbracket^R(w'') \wedge \exists h_1'' \subseteq h_1', h_2'' \subseteq h_2'. (j, h_1^\circ, h_2^\circ, h_1'', h_2'') \in \mathbf{Q}_\varepsilon^{\Lambda, R} w^\circ, w'' \right] \wedge \\
& \left[ \forall e_1', h_1^\dagger, \mu, n \leq j. (e_1 | g_1 \cdot h_1 \cdot f_1) \rightarrow_\mu^n (e_1' | h_1^\dagger) \Rightarrow \right. \\
& \quad \exists e_2', w'', A', h_1', h_2', g_1', g_2'. \\
& \quad (e_2 | g_2 \cdot h_2 \cdot f_2) \mapsto^* (e_2' | g_2' \cdot h_2' \cdot f_2) \wedge \mathbf{selftran}^{\Pi, \Lambda, A, R} w', w'' \wedge \\
& \quad A' = (A \cap \text{dom}(w'')) \cup (\text{dom}(w'') \setminus \text{dom}(w')) \wedge h_1^\dagger = g_1' \cdot h_1' \cdot f_1 \wedge \\
& \quad (g_1', h_1', f_1, g_2', h_2', f_2) \in \mathbf{splits}^{\Pi, \Lambda, A', R} w'' \wedge (j - n, g_1, g_2, g_1', g_2') \in \mathbf{Q}_\varepsilon^{\Pi, R} w', w'' \wedge \\
& \quad \left. \mu \in \mathbf{effs}_{\varepsilon, h_1'}^{A', R} w'' \wedge (j - n, h_1^\circ, h_2^\circ, e_1', e_2', h_1', h_2') \in \mathbf{safe}_{\tau, \varepsilon}^{\Pi, \Lambda, A', R} w^\circ, w'' \right]
\end{aligned}$$

■ **Figure 6** Safety. The predicate is defined by well-founded induction. Nontrivial requirements are:  $\Pi \# \Lambda$ ,  $\text{FRV}(\tau, \varepsilon) \subseteq \Pi \cup \Lambda$ ,  $\text{FV}(e_1, e_2) = \emptyset$ ,  $R : \Pi \cup \Lambda \leftrightarrow |w^\circ|$ ,  $R(\text{FRV}(\varepsilon)) \subseteq \text{dom}(w^\circ)$  and  $w \sqsupseteq w^\circ$  with  $\text{dom}(w^\circ) \cap R(\Pi \cup \Lambda) \subseteq \text{dom}(w)$ ,  $A \subseteq \text{dom}(w)$  and  $A \# R(\Pi \cup \Lambda)$ . See Figure 8 for auxiliary definitions. We refer to the contents of the big square brackets as the *prerequisites*, the *termination branch* and the *progress branch*, respectively.

$$\begin{aligned}
& \Pi \mid \Lambda \mid \Gamma \models e_1 \preceq e_2 : \tau, \varepsilon \\
& \iff \\
& \exists a \in \mathbb{N}. \forall k \in \mathbb{N}. \forall w \in \mathbf{W}. \forall R : \Pi \cup \Lambda \leftrightarrow |w|. \forall A \subseteq \text{dom}(w). \\
& \forall \gamma_1, \gamma_2 \in \mathcal{V}^{|\Gamma|}. \forall h_1, h_2, h_1', h_2' \in \mathcal{H}. \\
& \left[ R(\text{FRV}(\varepsilon)) \subseteq \text{dom}(w) \wedge A \# R(\Pi \cup \Lambda) \wedge |A| \geq a \wedge \forall r \in A. w(r) = \emptyset \wedge \right. \\
& \quad \left. (k, \gamma_1, \gamma_2) \in \llbracket \Gamma \rrbracket^R w \wedge h_1' \subseteq h_1 \wedge h_2' \subseteq h_2 \wedge (k, h_1', h_2') \in \mathbf{P}_\varepsilon^{\Lambda, R} w \right] \Rightarrow \\
& \quad (k, h_1', h_2', e_1[\gamma_1/\Gamma], e_2[\gamma_2/\Gamma], h_1, h_2) \in \mathbf{safe}_{\tau, \varepsilon}^{\Pi, \Lambda, A, R} w, w.
\end{aligned}$$

■ **Figure 7** The logical relation with anonymous regions. We require that  $\Pi \# \Lambda$ ,  $\text{FRV}(\Gamma, \tau, \varepsilon) \subseteq \Pi \cup \Lambda$  and, as always, that  $\text{FV}(e_1, e_2) \in |\Gamma|$ .

$$\begin{aligned} \mathbf{envtran}^{\Pi, \Lambda, A, R} w, w' &\iff w \sqsubseteq w' \wedge \forall r \in \text{dom}(w) \cap (R(\Pi \cup \Lambda) \cup A). r \in \text{dom}(w') \\ &\wedge \forall r \in \text{dom}(w) \cap (R(\Lambda) \cup A). w(r) = w'(r). \end{aligned}$$

$$\begin{aligned} \mathbf{selftran}^{\Pi, \Lambda, A, R} w, w' &\iff w \sqsubseteq w' \wedge \forall r \in \text{dom}(w) \setminus A. r \in \text{dom}(w') \\ &\wedge \forall r \in \text{dom}(w) \setminus (R(\Pi \cup \Lambda) \cup A). w(r) = w'(r). \end{aligned}$$

$$\begin{aligned} (g_1, h_1, f_1, g_2, h_2, f_2) \in \mathbf{splits}^{\Pi, \Lambda, A, R} w &\iff \\ \text{dom}(h_1) \# \text{dom}(g_1) \# \text{dom}(f_1) \wedge \text{dom}(h_2) \# \text{dom}(g_2) \# \text{dom}(f_2) \wedge \\ \text{dom}_1^{\text{R}(\Pi)}(w) = \text{dom}(g_1) \wedge \text{dom}_1^{\text{R}(\Lambda) \cup A}(w) \subseteq \text{dom}(h_1) \wedge \\ \text{dom}_1^{\text{dom}(w) \setminus (R(\Pi \cup \Lambda) \cup A)}(w) \subseteq \text{dom}(f_1) \wedge \\ \text{dom}_2^{\text{R}(\Pi)}(w) = \text{dom}(g_2) \wedge \text{dom}_2^{\text{R}(\Lambda) \cup A}(w) \subseteq \text{dom}(h_2) \wedge \\ \text{dom}_2^{\text{dom}(w) \setminus (R(\Pi \cup \Lambda) \cup A)}(w) \subseteq \text{dom}(f_2). \end{aligned}$$

$$\begin{aligned} (k, h_1, h_2) \in \mathbf{P}_\varepsilon^{\Theta, R} w &\iff \text{dom}(h_1) = \text{dom}_1^{\text{R}(\Theta)}(w) \wedge \text{dom}(h_2) = \text{dom}_2^{\text{R}(\Theta)}(w) \wedge \\ &\forall r \in R(\Theta) \cap \text{dom}(w). \forall (l_1, l_2, \mu) \in w(r). \\ &r \in R(\text{rds } \varepsilon) \Rightarrow k > 0 \Rightarrow (k-1, h_1(l_1), h_2(l_2)) \in (\iota \mu)(w). \end{aligned}$$

$$\begin{aligned} (k, h_1, h_2, h'_1, h'_2) \in \mathbf{Q}_\varepsilon^{\Theta, R} w, w' &\iff \\ \text{dom}(h_1) = \text{dom}_1^{\text{R}(\Theta)}(w) \wedge \text{dom}(h_2) = \text{dom}_2^{\text{R}(\Theta)}(w) \wedge \\ \text{dom}(h'_1) = \text{dom}_1^{\text{R}(\Theta)}(w') \wedge \text{dom}(h'_2) = \text{dom}_2^{\text{R}(\Theta)}(w') \wedge \\ (\forall r \in R(\Theta) \cap \text{dom}(w). \forall (l_1, l_2, \mu) \in w(r). \\ [h_1(l_1) = h'_1(l_1) \wedge h_2(l_2) = h'_2(l_2)] \vee [r \in R(\text{wrs } \varepsilon) \wedge \\ k > 0 \Rightarrow (k-1, h'_1(l_1), h'_2(l_2)) \in (\iota \mu)(w')]) \wedge \\ (\forall r \in R(\Theta) \cap \text{dom}(w). \\ \forall (l_1, l_2, \mu) \in w'(r) \setminus w(r). r \in R(\text{als } \varepsilon) \wedge \\ k > 0 \Rightarrow (k-1, h'_1(l_1), h'_2(l_2)) \in (\iota \mu)(w')). \end{aligned}$$

$$\begin{aligned} \mu \in \mathbf{effs}_{\varepsilon, h}^{A, R} w &\iff \{l \mid rd_l \in \mu\} \cap \text{dom}_1(w) \subseteq \text{dom}_1^{\text{R}(\text{rds } \varepsilon) \cup A}(w) \wedge \\ &\{l \mid wr_l \in \mu\} \cap \text{dom}_1(w) \subseteq \text{dom}_1^{\text{R}(\text{wrs } \varepsilon) \cup A}(w) \wedge \\ &\{l \mid al_l \in \mu\} \cap \text{dom}_1(w) \subseteq \text{dom}_1^{\text{R}(\text{als } \varepsilon) \cup A}(w) \wedge \\ &\{l \mid rd_l \in \mu \vee wr_l \in \mu \vee al_l \in \mu\} \setminus \text{dom}_1(w) \subseteq \text{dom}(h). \end{aligned}$$

■ **Figure 8** Six auxiliary definitions. The *environment transition* and *self transition* relations are defined for  $\Pi \# \Lambda$ ,  $R : \Pi \cup \Lambda \hookrightarrow |w|$ ,  $A \subseteq \text{dom}(w)$  and  $R(\Pi \cup \Lambda) \# A$ . The *triple-split* relation has the same prerequisites. The *precondition* relation is defined for  $R : \mathcal{RV} \xrightarrow{f_n} |w|$  injective with  $\Theta \cup \text{FRV}(\varepsilon) \subseteq \text{dom}(R)$ . The *postcondition* relation additionally requires  $w' \sqsupseteq w$  such that  $\text{dom}(w) \cap R(\Theta) \subseteq \text{dom}(w')$ . Finally the *actual-effects* relation expects  $R : \mathcal{RV} \xrightarrow{f_n} |w|$  injective with  $\text{FRV}(\varepsilon) \subseteq \text{dom}(R)$  and  $A \subseteq \text{dom}(w)$ .

## 4 Applications

### 4.1 Parallelization Theorem: Disjoint Concurrency

We now explain our Parallelization Theorem, which gives us an easy way to prove properties about the common case of disjoint concurrency, where disjointness is captured using private regions and effect annotations.

► **Theorem 8 (Parallelization).** *Assuming that*

1.  $\Pi, \Lambda \mid \cdot \mid \Gamma \vdash e_1 : \tau_1, \varepsilon_1,$
2.  $\Pi, \Lambda \mid \cdot \mid \Gamma \vdash e_2 : \tau_2, \varepsilon_2,$
3.  $\text{rds } \varepsilon_1 \cup \text{wrs } \varepsilon_1 \cup \text{rds } \varepsilon_2 \cup \text{wrs } \varepsilon_2 \subseteq \Lambda,$
4.  $\text{rds } \varepsilon_1 \cap \text{wrs } \varepsilon_2 = \text{rds } \varepsilon_2 \cap (\text{wrs } \varepsilon_1 \cup \text{als } \varepsilon_1) = \text{wrs } \varepsilon_1 \cap \text{wrs } \varepsilon_2 = \emptyset,$

*the following property holds:*

$$\Pi \mid \Lambda \mid \Gamma \models \langle e_1, e_2 \rangle \cong \text{par } e_1 \text{ and } e_2 : \tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2.$$

Intuitively, item 3 keeps the environment from detecting anything, and item 4 prevents the two computations from talking among themselves, thereby making them independent; the  $\text{als } \varepsilon_1$  in item 4 is a technicality that we cannot do without. We showed a concrete simple application of this theorem in the Introduction. More generally, example usage includes situations where we operate on two imperative data structures (say linked lists or graphs); if we only mutate parts of the data structures that are in different regions, then we may safely parallelize operations on the data structures.

The masking rule makes it possible to do more optimizations via the Parallelization Theorem: Consider, for simplicity, the familiar example of an efficient implementation *fib* of the Fibonacci function using two local references. We can use the masking rule to give it type and effect  $\text{int} \rightarrow_{\emptyset}^{\rho, \emptyset} \text{int}, \emptyset$ . This allows us to view the imperative implementation as pure, and thus by Theorem 8 we find that it is sound to optimize two sequential calls to *fib* to two parallel calls. This may sound like a simple optimization, but the point is that a compiler can perform it automatically, just based on the effect types. It also underlines how we are able to reason about more involved behaviors of concurrent threads, even though the type system provides only rough bounds on interference through the private-public distinction.

The proof of the Parallelization Theorem is quite tricky. Please see the appendix for an informal overview of the proof and the technical details.

### 4.2 Non-disjoint Concurrency

We now exemplify how our logical relations model can also be used to reason compositionally about equivalences of fine-grained concurrent programs operating on public regions.

Consider the following type

$$\tau \equiv \text{ref}_{\rho} \text{int} \rightarrow_{\{\text{rd}_{\rho}, \text{wr}_{\rho}\}}^{\rho, \emptyset} 1$$

of functions that take an integer reference in a public region, possibly read and write from the reference, and return unit. The following two functions

$$\text{fun inc}_1(x). \text{ let } y = !x \text{ in let } z = y + 1 \text{ in } \quad \text{and } \text{fun inc}_2(x). \text{ atomic } (x := !x + 1) \\ \text{if cas } (x, y, z) \text{ then } \langle \rangle \text{ else inc}_1(x)$$

both have type  $\tau$ . (We have allowed ourselves to use a standard conditional expression; 1 corresponds to true and 0 to false.) Both functions increment the integer given in their reference arguments;  $\text{inc}_1$  uses the fine-grained compare-and-swap to do it atomically, whereas



$\text{inc}_2$  uses the brute-force atomic operation. Using our logical relations model, we can prove that  $\text{inc}_1$  and  $\text{inc}_2$  are contextually equivalent:

$$\rho \mid \cdot \mid \cdot \vdash \text{inc}_1 \approx \text{inc}_2 : \tau, \emptyset. \quad (1)$$

Hence, replacing  $\text{inc}_2$  with  $\text{inc}_1$  in any well-typed client gives two contextually equivalent expressions. Thus our logical relation models a form of *data abstraction for concurrency* (where we abstract over the granularity of concurrency in the module).

We now show how to use the equivalence of  $\text{inc}_1$  and  $\text{inc}_2$  to derive equivalences of two different clients using the fine-grained concurrency implementation  $\text{inc}_1$ .

To this end, consider the following two client programs of type

$$\sigma \equiv \tau \xrightarrow{\rho, \emptyset} \text{ref}_{\rho} \text{int} \xrightarrow{\rho, \emptyset}_{\{\text{rd}_{\rho}, \text{wr}_{\rho}\}} \text{int},$$

$$\text{fun } c_1(\text{inc}). \lambda n. \text{inc } n; \text{inc } n; !n \quad \text{and} \quad \text{fun } c_2(\text{inc}). \lambda n. (\text{par } \text{inc } n \text{ and } \text{inc } n); !n$$

Note that  $c_1$  makes two sequential calls to  $\text{inc}$ , whereas  $c_2$  runs the two calls in parallel. Because of the use of compare-and-swap in  $\text{inc}_1$ , we would hope that the  $c_1 \text{inc}_1$  and  $c_2 \text{inc}_1$  are contextually equivalent (in typing context  $\rho \mid \emptyset \mid \emptyset$ ). We can prove that this is indeed the case using compositional reasoning as follows. Using our logical relation, we prove that  $c_1 \text{inc}_2$  is contextually equivalent to  $c_2 \text{inc}_2$ , i.e.,

$$\rho \mid \cdot \mid \cdot \vdash c_1 \text{inc}_2 \approx c_2 \text{inc}_2 : \text{ref}_{\rho} \text{int} \xrightarrow{\rho, \emptyset}_{\{\text{rd}_{\rho}, \text{wr}_{\rho}\}} \text{int}, \emptyset. \quad (2)$$

Finally, we conclude that  $c_1 \text{inc}_1$  is contextually equivalent to  $c_2 \text{inc}_1$  by transitivity of contextual equivalence (using (1), (2) and (1) again for the respective steps):

$$c_1 \text{inc}_1 \approx c_1 \text{inc}_2 \approx c_2 \text{inc}_2 \approx c_2 \text{inc}_1$$

This proof illustrates an important point: to show equivalence of two clients of a module implemented using fine-grained concurrency, it suffices to show that the clients are equivalent wrt. a coarse-grained implementation, and that the coarse-grained implementation is equivalent to the fine-grained implementation. This is often a lot simpler than trying to show the equivalence of the clients wrt. the fine-grained implementation directly. We can think of the coarse-grained implementation of the module (here  $\text{inc}_2$ ) as the *specification* of the module and the fine-grained implementation (here  $\text{inc}_1$ ) as its *implementation*.

The formal proofs of (1) and (2) follow by straightforward induction.

## 5 Discussion

Gifford and Lucassen [11, 12] originally proposed type-and-effect systems as a static analysis for determining which parts of a higher-order imperative program could be implemented using parallelism. Here we are able to express the formal correctness of these ideas in a succinct way by having a parallel construct in our programming language and establishing the Parallelization Theorem.

In Section 4.2 we showed how contextual equivalence can be used to *state* that compare-and-swap can be used to implement a simple form of locking, and how our logical relations model could be used to *prove* this statement. We believe that it should be possible to give similar succinct statements and proofs of other implementations of synchronization. For instance, we are currently working on a similar relational specification and correctness proof of Peterson's mutual exclusion algorithm, which involves (for succinctness of specification) extending the language with a primitive notion of critical section.

As mentioned earlier, we have deliberately used the same definition of worlds here as in [20]. As discussed there [20, Section 8.2], this notion of world has somewhat limited expressiveness: the only heap invariants we can state are those that relate values at two locations by a semantic type. To increase expressiveness, it would thus be interesting to extend our model using ideas from [9], where worlds are defined using state-transition systems, and then investigate more examples of equivalences.

Recently, Liang et. al. [10] have proposed RGSim, a simulation based on rely-guarantee, to verify program transformations in a concurrent setting. Their actual definition [10, Definition 4] bears some resemblance to our safety relation; indeed, an early draft of *loc.cit.* was a source of inspiration. They have no division of the heap into public and private parts, instead they give a pair of rely and guarantee that, respectively, constrain the interference of the environment and the actions of the computation. Their approach is essentially untyped; one point of view is that we ‘auto-instantiate’ the many parameters of their simulation based on our typing information. They consider first-order languages with ground store; this obviously keeps life simple, but the example equivalences they give are not.

Our simple example of data abstraction for concurrency in Section 4.2 suggests that there could be a relationship to linearizability. In [28], Filipović et. al. show a formal connection between linearizability and simulation relations, for a simple first-order imperative programming language. We intend to explore whether a similar kind of formal relationship can be established in our higher-order setting.

## 6 Conclusion and Future Work

We have presented a logical relations model of a new type-and-effect system for a concurrent higher-order ML-like language with general references. We have shown how to use the model for reasoning about both disjoint and non-disjoint concurrency. In particular, we have proved the first automatic Parallelization Theorem for such a rich language.

In this paper, we have focused on may contextual equivalence. Future work includes investigating models for must contextual equivalence. Since our language allows the encoding of countable nondeterminism, must equivalence is non-trivial, and will probably involve indexing over  $\omega_1$  rather than  $\omega$  [29]. Future work also includes extending the model to region and effect polymorphism, as well as the extension to more expressive worlds, and to other concurrency constructs such as fork-join.

In this paper we have used logical relations for reasoning about contextual equivalence for a concurrent higher-order imperative language with a type-and-effect system. In the future, it would be interesting to explore also the application of other methods, such as bisimulations and game semantics.

The authors would like to thank Jan Schwinghammer and Xinyu Feng for discussions of aspects of this work.

---

## References

- 1 V. Koutavas and M. Wand, “Small bisimulations for reasoning about higher-order imperative programs,” in *POPL*, 2006.
- 2 D. Sangiorgi, N. Kobayashi, and E. Sumii, “Environmental bisimulations for higher-order languages,” *TOPLAS*, 2011.
- 3 E. Sumii, “A complete characterization of observational equivalence in polymorphic  $\lambda$ -calculus with general references,” in *CSL*, 2009.
- 4 J. Laird, “A fully abstract trace semantics for general references,” in *ICALP*, 2007.

- 5 A. Murawski and N. Tzevelekos, “Game semantics for good general references,” in *LICS*, 2011.
- 6 A. Ahmed, “Semantics of types for mutable state,” Ph.D. dissertation, Princeton University, 2004.
- 7 A. Ahmed, D. Dreyer, and A. Rossberg, “State-dependent representation independence,” in *POPL*, 2009.
- 8 L. Birkedal, J. Thamsborg, and K. Støvring, “Realizability semantics of parametric polymorphism, general references, and recursive types,” in *FOSSACS*, 2009.
- 9 D. Dreyer, G. Neis, and L. Birkedal, “The impact of higher-order state and control effects on local relational reasoning,” in *ICFP 2010*. ACM, 2010, pp. 143–156.
- 10 H. Liang, X. Feng, and M. Fu, “A rely-guarantee-based simulation for verifying concurrent program transformations,” in *POPL*, 2012.
- 11 D. Gifford and J. Lucassen, “Integrating functional and imperative programming,” in *LISP and Functional Programming*, 1986.
- 12 J. Lucassen and D. Gifford, “Polymorphic effect systems,” in *POPL*, 1988.
- 13 M. Tofte and J.-P. Talpin, “Implementation of the typed call-by-value  $\lambda$ -calculus using a stack of regions,” in *Proceedings of POPL*, 1994.
- 14 F. Henglein, H. Makhholm, and H. Niss, “Effect types and region-based memory management,” in *Advanced Topics in Types and Programming Languages*, B. Pierce, Ed. MIT Press, 2005.
- 15 L. Birkedal, M. Tofte, and M. Vejlstrup, “From region inference to von Neumann machines via region representation inference,” in *POPL*, 1996.
- 16 N. Benton, A. Kenney, M. Hofmann, and L. Beringer, “Reading, writing and relations: Towards extensional semantics for effect analyses,” in *APLAS*, 2006.
- 17 N. Benton and P. Buchlovsky, “Semantics of an effect analysis for exceptions,” in *TLDI*, 2007.
- 18 N. Benton, L. Beringer, M. Hofmann, and A. Kennedy, “Relational semantics for effect-based program transformations with dynamic allocation,” in *PPDP*. ACM, 2007.
- 19 —, “Relational semantics for effect-based program transformations: Higher-order store,” in *PPDP*. ACM, 2009.
- 20 J. Thamsborg and L. Birkedal, “A Kripke logical relation for effect-based program transformations,” in *ICFP*, 2011.
- 21 P. W. O’Hearn, “Resources, concurrency, and local reasoning,” *TCS*, 2007.
- 22 M. Dodds, X. Feng, M. Parkinson, and V. Vafeiadis, “Deny-guarantee reasoning,” in *ESOP*, 2009.
- 23 V. Vafeiadis and M. Parkinson, “A marriage of rely/guarantee and separation logic,” in *CONCUR*, 2007.
- 24 X. Feng, R. Ferreira, and Z. Shao, “On the relationship between concurrent separation logic and assume-guarantee reasoning,” in *ESOP*, 2007.
- 25 V. Vafeiadis, “Concurrent separation logic and operational semantics,” in *MFPS*, 2011.
- 26 A. Buisse, L. Birkedal, and K. Støvring, “A step-indexed Kripke model of separation logic for storable locks,” in *MFPS*, 2011.
- 27 L. Birkedal, B. Reus, J. Schwinghammer, K. Støvring, J. Thamsborg, and H. Yang, “Step-indexed Kripke models over recursive worlds,” in *POPL*, 2011.
- 28 I. Filipovic, P. W. O’Hearn, N. Rinetzky, and H. Yang, “Abstraction for concurrent objects,” *TCS*, 2010.
- 29 J. Schwinghammer and L. Birkedal, “Step-indexed relational reasoning for countable non-determinism,” in *CSL*, 2011.

## A Appendix

### A.1 Typing of Contexts

As usually, typing rules for contexts are simple extensions of the type basic type system. It is, however, interesting to observe the region dynamics in the analogues of masking and parallel composition rules, since these allow contexts (when read top-to-bottom) to “kill off” private regions and change public regions into private ones. Selected rules are presented in Figure 9.

$$\begin{array}{c}
\overline{[] : (\Pi | \Lambda | \Gamma \vdash \tau, \varepsilon) \rightsquigarrow (\Pi | \Lambda | \Gamma \vdash \tau, \varepsilon)} \\
\frac{C : (\Pi | \Lambda | \Gamma \vdash \tau, \varepsilon) \rightsquigarrow (\Pi' | \Lambda', \rho | \Gamma' \vdash \tau', \varepsilon')}{C : (\Pi | \Lambda | \Gamma \vdash \tau, \varepsilon) \rightsquigarrow (\Pi' | \Lambda' | \Gamma' \vdash \tau', \varepsilon' - \rho)} \quad (\rho \notin \text{FRV}(\Gamma', \tau')) \\
\frac{C : (\Pi | \Lambda | \Gamma \vdash \tau, \varepsilon) \rightsquigarrow (\Pi', \Lambda' | \cdot | \Gamma' \vdash \tau_1, \varepsilon_1) \quad \Pi', \Lambda' | \cdot | \Gamma' \vdash e : \tau_2, \varepsilon_2}{\text{par } C \text{ and } e : (\Pi | \Lambda | \Gamma \vdash \tau, \varepsilon) \rightsquigarrow (\Pi' | \Lambda' | \Gamma' \vdash \tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2)} \\
\frac{C : (\Pi | \Lambda | \Gamma \vdash \tau, \varepsilon) \rightsquigarrow (\Pi' | \Lambda' | \Gamma' \vdash \tau_1 \xrightarrow{\Pi', \Lambda'}_{\varepsilon'} \tau_2, \varepsilon_1) \quad \Pi' | \Lambda' | \Gamma' \vdash e : \tau_1, \varepsilon_2}{C e : (\Pi | \Lambda | \Gamma \vdash \tau, \varepsilon) \rightsquigarrow (\Pi' | \Lambda' | \Gamma' \vdash \tau_2, \varepsilon_1 \cup \varepsilon_2 \cup \varepsilon')} \\
\frac{C : (\Pi | \Lambda | \Gamma \vdash \tau, \varepsilon) \rightsquigarrow (\Pi' | \Lambda' | \Gamma' \vdash \tau_1, \varepsilon_2) \quad \Pi' | \Lambda' | \Gamma' \vdash e : \tau_1 \xrightarrow{\Pi', \Lambda'}_{\varepsilon'} \tau_2, \varepsilon_1}{e C : (\Pi | \Lambda | \Gamma \vdash \tau, \varepsilon) \rightsquigarrow (\Pi' | \Lambda' | \Gamma' \vdash \tau_2, \varepsilon_1 \cup \varepsilon_2 \cup \varepsilon')} \\
\frac{C : (\Pi | \Lambda | \Gamma \vdash \tau, \varepsilon) \rightsquigarrow (\Pi' | \Lambda' | \Gamma', f : \tau_1 \xrightarrow{\Pi, \Lambda}_{\varepsilon'} \tau_2, x : \tau_1 \vdash \tau_2, \varepsilon')}{\text{fun } f(x).C : (\Pi | \Lambda | \Gamma \vdash \tau, \varepsilon) \rightsquigarrow (\Pi' | \Lambda' | \Gamma' \vdash \tau_1 \xrightarrow{\Pi, \Lambda}_{\varepsilon'} \tau_2, \emptyset)}
\end{array}$$

■ **Figure 9** Selected typing rules for contexts

### A.2 Well-definedness of Interpretation of Types

As explained in the main text, we import the worlds and types of Thamsborg and Birkedal [20] wholesale; hence their non-trivial construction is for free. The details of the construction, including metric prerequisites, can be found in Appendices A.1 and A.2 of the long version of *loc. cit.*; it is available online.<sup>1</sup> Thus it only remains to verify that interpreting a syntactic type actually gives a semantic ditto.

In metric terms, this means that for any syntactic type  $\tau$  and any  $R : \mathcal{RV} \rightarrow_{\text{fin}} \mathcal{RN}$  with  $\text{FRV}(\tau) \subseteq \text{dom}(R)$  we must have

$$[[\tau]]^R \in \mathbf{T} = \mathbf{W} \rightarrow_{\text{mon}} \text{URel}(\mathcal{V}),$$

i.e., it should map worlds to indexed, downwards closed relations on values in a non-expansive and monotone manner. To prove this, we need two lemmas on the pre- and postcondition relations:

► **Lemma 9.** *We have  $\mathbf{P}_\varepsilon^{\Theta, R} w \in \text{URel}(\mathcal{H})$ . And for  $w_1, w_2 \in \mathbf{W}$  with  $w_1 \stackrel{n}{=} w_2$  we have  $\mathbf{P}_\varepsilon^{\Theta, R} w_1 \stackrel{n}{=} \mathbf{P}_\varepsilon^{\Theta, R} w_2$  as well.*

<sup>1</sup> [www.itu.dk/people/thamsborg/longcarnival.pdf](http://www.itu.dk/people/thamsborg/longcarnival.pdf)

► **Lemma 10.** We have  $\mathbf{Q}_\varepsilon^{\ominus,R} w, w' \in UR\text{el}(\mathcal{H} \times \mathcal{H})$ . And for  $w_1, w'_1, w_2, w'_2 \in \mathbf{W}$  with  $w_1 \stackrel{n}{=} w_2$  and  $w'_1 \stackrel{n}{=} w'_2$  we have  $\mathbf{Q}_\varepsilon^{\ominus,R} w_1, w'_1 \stackrel{n}{=} \mathbf{Q}_\varepsilon^R w_2, w'_2$  too.

The proofs are quite straightforward. Then, by simultaneous induction, we can prove the following two propositions and we are done:

► **Proposition 11.** We have  $\llbracket \tau \rrbracket^R \in \mathbf{T}$ .

► **Proposition 12.** For  $w_1^\circ, w_1, w_2^\circ, w_2 \in \mathbf{W}$  we have that  $w_1^\circ \stackrel{n}{=} w_2^\circ$  together with  $w_1 \stackrel{n}{=} w_2$  implies  $\mathbf{safe}_{\tau, \varepsilon}^{\Pi, \Lambda, A, R} w_1^\circ, w_1 \stackrel{n}{=} \mathbf{safe}_{\tau, \varepsilon}^{\Pi, \Lambda, A, R} w_2^\circ, w_2$ .

### A.3 Instrumented Operational Semantics

$$\begin{aligned}
(E[\text{proj}_i \langle v_1, v_2 \rangle] | h) &\rightarrow_\emptyset^1 (E[v_i] | h) \\
(E[(\text{fun } f(x).e) v] | h) &\rightarrow_\emptyset^1 (E[e[f := \text{fun } f(x).e, x := v]] | h) \\
(E[\text{ref } v] | h) &\rightarrow_{\{al_i\}}^1 (E[l] | h[l \mapsto v]) \text{ if } l \notin \text{dom}(h) \\
(E[l := v] | h) &\rightarrow_{\{wr_l\}}^1 (E[\langle \rangle] | h[l := v]) \text{ if } l \in \text{dom}(h) \\
(E[!l] | h) &\rightarrow_{\{rd_l\}}^1 (E[h(l)] | h) \text{ if } l \in \text{dom}(h) \\
(E[\text{par } v_1 \text{ and } v_2] | h) &\rightarrow_\emptyset^1 (E[\langle v_1, v_2 \rangle] | h) \\
(E[\text{cas } (l, n_1, n_2)] | h) &\rightarrow_{\{rd_l, wr_l\}}^1 (E[1] | h[l := n_2]) \\
&\text{if } l \in \text{dom}(h) \text{ and } h(l) = n_1 \\
(E[\text{cas } (l, n_1, n_2)] | h) &\rightarrow_{\{rd_l\}}^1 (E[0] | h) \\
&\text{if } l \in \text{dom}(h) \text{ and } h(l) \neq n_1 \\
(E[\text{atomic } e] | h) &\rightarrow_\mu^{n+1} (E[v] | h') \text{ if } (e | h) \xrightarrow{\mu}^n (v | h') \\
(E[\text{atomic } e] | h) &\rightarrow_\emptyset^1 (E[\text{atomic } e] | h)
\end{aligned}$$

$$\begin{aligned}
(e | h) \xrightarrow{\emptyset}^0 (e' | h') &\iff e = e' \wedge h = h' \\
(e | h) \xrightarrow{\mu}^n (e'' | h'') &\iff (e | h) \xrightarrow{\mu_1}^{n_1} (e' | h') \wedge \\
&(e' | h') \xrightarrow{\mu_2}^{n_2} (e'' | h'') \wedge \\
&n = n_1 + n_2 \wedge \mu = \mu_1 \cup \mu_2.
\end{aligned}$$

Instrumented operational semantics is presented above.  $(e | h) \xrightarrow{\mu}^n (e' | h')$  is defined for  $e, e' \in \mathcal{E}$  with  $\text{FV}(e, e') = \emptyset$ ,  $h, h' \in \mathcal{H}$ ,  $n \geq 1$  and  $\mu \subseteq \{al_i \mid l \in \mathcal{L}\} \cup \{wr_l \mid l \in \mathcal{L}\} \cup \{rd_l \mid l \in \mathcal{L}\}$ ; the starred version additionally permits a zero as superscript. Note that  $(e | h) \mapsto (e' | h')$  if and only if there are  $n$  and  $\mu$  such that  $(e | h) \xrightarrow{\mu}^n (e' | h')$ .

### A.4 Proof of May-Equivalence Theorem

**Theorem 3.** Assume that  $\cdot | \cdot | \cdot \models e_1 \preceq e_2 : \text{int}, \emptyset$  holds. Take any  $h_1, h_2 \in \mathcal{H}$ . If there are  $e'_1, h'_1$  with  $(e_1 | h_1) \mapsto^* (e'_1 | h'_1)$  such that  $\text{irr}(e'_1 | h'_1)$  holds, then there is  $n \in \mathbb{Z}$  such that  $e'_1 = \underline{n}$  and  $h'_2$  such that  $(e_2 | h_2) \mapsto^* (\underline{n}, h'_2)$ .

**Proof.** Pick, by assumption,  $n \in \mathbb{N}$  such that  $(e_1 | h_1) \mapsto^n (e'_1 | h'_1)$ . We name the configurations in this multi-step reduction  $(e_1^0 | h_1^0), (e_1^1 | h_1^1), \dots, (e_1^n | h_1^n)$  in order, i.e., such that

$$\begin{aligned}
(e_1 | h_1) &= (e_1^0 | h_1^0) \mapsto (e_1^1 | h_1^1) \mapsto \dots \\
&\mapsto (e_1^n | h_1^n) = (e'_1 | h'_1).
\end{aligned}$$

For each of these standard reductions, there is a corresponding instrumented reduction. This means, that we can pick  $n_1, n_2, \dots, n_n \in \mathbb{N}$  non-zero and  $\mu_1, \mu_2, \dots, \mu_n$  such that

$$(e_1^0 | h_1^0) \rightarrow_{\mu_1}^{n_1} (e_1^1 | h_1^1) \rightarrow_{\mu_2}^{n_2} \dots \rightarrow_{\mu_n}^{n_n} (e_1^n | h_1^n).$$

For convenience, we furthermore write  $N_m = \sum_{i=m+1}^n n_i$  for each  $0 \leq m \leq n$ .

Now let  $a \in \mathbb{N}$  be the required number of anonymous regions according to the definition of the logical relation, and let  $w \in \mathbf{W}$  be any world with  $a$  empty, live regions and nothing else. We now prove by induction that for all  $0 \leq m \leq n$  the following holds:

$$\begin{aligned} & \exists e'_2, h'_2, w'. (e_2 | h_2) \xrightarrow{*} (e'_2 | h'_2) \wedge \\ & \text{dom}(h_1^m) \supseteq \text{dom}_1(w') \wedge \text{dom}(h'_2) \supseteq \text{dom}_2(w') \wedge \\ & (N_m, [], [], e_1^m, e'_2, h_1^m, h'_2) \in \mathbf{safe}_{\text{int}, \emptyset}^{\emptyset, \emptyset, \text{dom}(w'), \emptyset} w, w'. \end{aligned}$$

The base case is easy by the definition of the logical relation: Pick  $e'_2 = e_2$ ,  $h'_2 = h_2$ , and  $w' = w$  and apply the assumption that  $\cdot | \cdot | \cdot \models e_1 \preceq e_2 : \text{int}, \emptyset$ .

For the inductive case, let  $0 < m \leq n$  and unroll the assumptions for  $m-1$ : we have  $e'_2, h'_2, w'$  with  $(e_2 | h_2) \xrightarrow{*} (e'_2 | h'_2)$  and  $\text{dom}(h_1^{m-1}) \supseteq \text{dom}_1(w')$ ,  $\text{dom}(h'_2) \supseteq \text{dom}_2(w')$ , and

$$(N_{m-1}, [], [], e_1^{m-1}, e'_2, h_1^{m-1}, h'_2) \in \mathbf{safe}_{\text{int}, \emptyset}^{\emptyset, \emptyset, \text{dom}(w'), \emptyset} w, w'.$$

By the overall setup, we know that  $(e_1^{m-1} | h_1^{m-1}) \rightarrow_{\mu_m}^{n_m} (e_1^m | h_1^m)$  and the progress branch of safety provides for us: There are  $e''_2$  and  $h''_2$  with  $(e'_2 | h'_2) \xrightarrow{*} (e''_2 | h''_2)$  and hence  $(e_2 | h_2) \xrightarrow{*} (e''_2 | h''_2)$ . There is  $w''$  with  $\text{dom}(h_1^m) \supseteq \text{dom}_1(w'')$  and  $\text{dom}(h''_2) \supseteq \text{dom}_2(w')$  and, finally,

$$(N_m, [], [], e_1^m, e''_2, h_1^m, h''_2) \in \mathbf{safe}_{\text{int}, \emptyset}^{\emptyset, \emptyset, \text{dom}(w''), \emptyset} w, w''.$$

We have now finished the induction proof; all that remains is to observe that the property proved implies the overall goal in the case  $m = n$  by the termination branch of safety.  $\blacktriangleleft$

## A.5 Properties of the Pre- and Postcondition Relations

► **Lemma 13** (Precondition Separation). *Assume that we have  $\Pi \# \Lambda$ ,  $h_1 = f_1 \upharpoonright_{\text{dom}_1^{\text{R}(\Lambda)}(w)}$ ,  $g_1 = f_1 \upharpoonright_{\text{dom}_1^{\text{R}(\Pi)}(w)}$ ,  $h_2 = f_2 \upharpoonright_{\text{dom}_2^{\text{R}(\Lambda)}(w)}$  and  $g_2 = f_2 \upharpoonright_{\text{dom}_2^{\text{R}(\Pi)}(w)}$ . Then it holds that*

$$\begin{aligned} & (k, h_1, h_2) \in \mathbf{P}_\varepsilon^{\Lambda, R} w \wedge (k, g_1, g_2) \in \mathbf{P}_\varepsilon^{\Pi, R} w \\ & \iff \\ & (k, f_1, f_2) \in \mathbf{P}_\varepsilon^{\Pi \cup \Lambda, R} w \end{aligned}$$

► **Lemma 14** (Postcondition Separation). *Assume that we have  $\Pi \# \Lambda$ ,  $h_1 = f_1 \upharpoonright_{\text{dom}_1^{\text{R}(\Lambda)}(w)}$ ,  $g_1 = f_1 \upharpoonright_{\text{dom}_1^{\text{R}(\Pi)}(w)}$ ,  $h_2 = f_2 \upharpoonright_{\text{dom}_2^{\text{R}(\Lambda)}(w)}$ ,  $g_2 = f_2 \upharpoonright_{\text{dom}_2^{\text{R}(\Pi)}(w)}$ ,  $h'_1 = f'_1 \upharpoonright_{\text{dom}_1^{\text{R}(\Lambda)}(w')}$ ,  $g'_1 = f'_1 \upharpoonright_{\text{dom}_1^{\text{R}(\Pi)}(w')}$ ,  $h'_2 = f'_2 \upharpoonright_{\text{dom}_2^{\text{R}(\Lambda)}(w')}$  and  $g'_2 = f'_2 \upharpoonright_{\text{dom}_2^{\text{R}(\Pi)}(w')}$ . Then*

$$\begin{aligned} & (k, h_1, h_2, h'_1, h'_2) \in \mathbf{Q}_\varepsilon^{\Lambda, R} w, w' \wedge \\ & (k, g_1, g_2, g'_1, g'_2) \in \mathbf{Q}_\varepsilon^{\Pi, R} w, w' \\ & \iff \\ & (k, f_1, f_2, f'_1, f'_2) \in \mathbf{Q}_\varepsilon^{\Pi \cup \Lambda, R} w, w'. \end{aligned}$$

► **Lemma 15** (Precondition Composition).

$$\begin{aligned} (k, h_1, h_2) \in \mathbf{P}_\varepsilon^{\Theta, R} w \wedge (k, h_1, h_2, h'_1, h'_2) \in \mathbf{Q}_\varepsilon^{\Theta, R} w, w' \\ \implies \\ (k, h'_1, h'_2) \in \mathbf{P}_\varepsilon^{\Theta, R} w'. \end{aligned}$$

► **Lemma 16** (Postcondition Composition).

$$\begin{aligned} (k, h_1, h_2, h'_1, h'_2) \in \mathbf{Q}_\varepsilon^{\Theta, R} w, w' \wedge \\ (k, h'_1, h'_2, h''_1, h''_2) \in \mathbf{Q}_\varepsilon^{\Theta, R} w', w'' \\ \implies \\ (k, h_1, h_2, h''_1, h''_2) \in \mathbf{Q}_\varepsilon^{\Theta, R} w, w''. \end{aligned}$$

## A.6 Some Cases of Proof of Compatibility

**Lemma 5.**  $\Pi \mid \Lambda, \rho \mid \Gamma \models e_1 \preceq e_2 : \tau, \varepsilon$  implies  $\Pi \mid \Lambda \mid \Gamma \models e_1 \preceq e_2 : \tau, \varepsilon - \rho$  provided that  $\rho \notin \text{FRV}(\Gamma, \tau)$ .

**Proof.** Let  $a \in \mathbb{N}$  be the required number of anonymous regions from the assumption. To prove the desired, we unroll the definition of the logical relation: choose  $a + 1$ , pick arbitrary  $k^\circ \in \mathbb{N}$ ,  $w^\circ \in \mathbf{W}$ ,  $R : \Pi \cup \Lambda \hookrightarrow |w^\circ|$ ,  $A^\circ \subseteq \text{dom}(w^\circ)$ ,  $\gamma_1, \gamma_2 \in \mathcal{V}^{|\Gamma|}$  and  $h_1^\circ, h_2^\circ, h'_1, h'_2 \in \mathcal{H}$ . Assume that  $R(\text{FRV}(\varepsilon - \rho)) \subseteq \text{dom}(w^\circ)$ ,  $A^\circ \# R(\Pi \cup \Lambda)$ ,  $|A^\circ| \geq a$ ,  $\forall r \in A^\circ. w^\circ(r) = \emptyset$ ,  $(k^\circ, \gamma_1, \gamma_2) \in \llbracket \Gamma \rrbracket^{R} w^\circ$ ,  $h_1^\circ \subseteq h_1, h_2^\circ \subseteq h_2$ , and  $(k^\circ, h'_1, h'_2) \in \mathbf{P}_{\varepsilon - \rho}^{\Lambda, R} w^\circ$ . We must show that  $(k^\circ, h'_1, h'_2, e_1[\gamma_1/\Gamma], e_2[\gamma_2/\Gamma], h_1^\circ, h_2^\circ) \in \mathbf{safe}_{\tau, \varepsilon - \rho}^{\Pi, \Lambda, A^\circ, R} w^\circ, w^\circ$ .

We need, obviously, to make use of the assumption of the lemma. To do so, we claim that to have

$$(k, h'_1, h'_2, e_1, e_2, h_1, h_2) \in \mathbf{safe}_{\tau, \varepsilon - \rho}^{\Pi, \Lambda, A, R} w^\circ, w \quad (3)$$

it suffices to know that

$$(k, h'_1, h'_2, e_1, e_2, h_1, h_2) \in \mathbf{safe}_{\tau, \varepsilon}^{\Pi, \Lambda, \rho, A \setminus \{r\}, R[\rho \mapsto r]} w^\circ, w \quad (4)$$

whenever  $r \in A$  holds. Observe first, that if we can prove this, then we are done by an easy application of the assumption of the lemma. So all that remains is to prove the claim; this we do by well-founded induction on  $k \in \mathbb{N}$ .

Assume that the claim holds for all naturals strictly less than  $k \in \mathbb{N}$ ; we will try to prove it for  $k$ . To prove (3) we proceed to unroll the definition of safety. Pick arbitrary  $j \leq k$ ,  $w' \in \mathbf{W}$  and  $g_1, g_2, f_1, f_2 \in \mathcal{H}$ . Assume that the prerequisites hold, i.e., that we have  $\mathbf{envtran}^{\Pi, \Lambda, A, R} w, w'$ ,  $(j, g_1, g_2) \in \mathbf{P}_{\varepsilon - \rho}^{\Pi, R} w'$  and  $(g_1, h_1, f_1, g_2, h_2, f_2) \in \mathbf{splits}^{\Pi, \Lambda, A, R} w'$ . Before we look into the termination and progress branches, observe that the  $w'$  and the  $g_1, g_2, f_1, f_2$  match the prerequisites of safety for (4) as well, since  $R(\Lambda) \cup A = R[\rho \mapsto r](\Lambda, \rho) \cup A \setminus \{r\}$ .

So we follow the termination branch and assume that  $\text{irr}(e_1 | g_1 \cdot h_1 \cdot f_1)$  holds. From (4) we get  $e'_2, w'', h'_1, h'_2, g'_1$  and  $g'_2$  with properties aplenty:

- $(e_2 | g_2 \cdot h_2 \cdot f_2) \mapsto^* (e'_2 | g'_2 \cdot h'_2 \cdot f_2)$ .
- $\mathbf{selftran}^{\Pi, \Lambda, \rho, A \setminus \{r\}, R[\rho \mapsto r]} w', w''$ .
- $\emptyset = ((A \setminus \{r\}) \cap \text{dom}(w'')) \cup (\text{dom}(w'') \setminus \text{dom}(w'))$ .
- $g_1 \cdot h_1 = g'_1 \cdot h'_1$ .
- $(g'_1, h'_1, f_1, g'_2, h'_2, f_2) \in \mathbf{splits}^{\Pi, \Lambda, \rho, \emptyset, R[\rho \mapsto r]} w''$ .

- $(j, e_1, e_2) \in \llbracket \tau \rrbracket^{R[\rho \mapsto r]}(w'')$ .
- $(j, g_1, g_2, g'_1, g'_2) \in \mathbf{Q}_\varepsilon^{\Pi, R[\rho \mapsto r]} w', w''$ .
- $\exists h''_1 \subseteq h'_1, h''_2 \subseteq h'_2. (j, h''_1, h''_2, h'_1, h'_2) \in \mathbf{Q}_\varepsilon^{\Lambda, \rho, R[\rho \mapsto r]} w^\circ, w''$ .

Now note that we must have  $r \in \text{dom}(w'')$  and so we define the world  $w'''$  by  $w'' \rightarrow_{\text{mask}(r)} w'''$ , i.e., we kill off region  $r$ . So we need to discharge a number of obligations; most are straightforward, but the type of the resulting expressions as well as the postcondition on the local heaps take a bit of scrutiny: Note first that we have

$$\llbracket \tau \rrbracket^{R[\rho \mapsto r]}(w'') \subseteq \llbracket \tau \rrbracket^{R[\rho \mapsto r]}(w''') = \llbracket \tau \rrbracket^R(w'')$$

by type monotonicity and since  $\rho \notin \text{FRV}(\tau)$ ; this means that the resulting expressions are indeed well-typed. Now define heaps  $h''_1 \subseteq h'_1$  and  $h''_2 \subseteq h'_2$  by demanding that  $\text{dom}(h''_1) = \text{dom}_1^{R(\Lambda)}(w''')$  and  $\text{dom}(h''_2) = \text{dom}_2^{R(\Lambda)}(w''')$ , i.e., we restrict to the private parts of the local heaps, not including the locations (formerly) in region  $r$ . One easily verifies that this gives

$$(j, h''_1, h''_2, h'_1, h'_2) \in \mathbf{Q}_{\varepsilon-\rho}^{\Lambda, R} w^\circ, w'''$$

and the termination branch is done.

Finally we get to the progress branch: we assume that there are  $e'_1, h_1^\dagger, \mu$  and  $n \leq j$  such that  $(e_1 | g_1 \cdot h_1 \cdot f_1) \rightarrow_\mu^n (e'_1 | h_1^\dagger)$ . From (4) we get  $e'_2, w'', A', h'_1, h'_2, g'_1, g'_2$  with a range of properties:

- $(e_2 | g_2 \cdot h_2 \cdot f_2) \xrightarrow{*} (e'_2 | g'_2 \cdot h'_2 \cdot f_2)$ .
- **selftran** $^{\Pi, \Lambda, \rho, A \setminus \{r\}, R[\rho \mapsto r]} w', w''$ .
- $A' = (A \setminus \{r\} \cap \text{dom}(w'')) \cup (\text{dom}(w'') \setminus \text{dom}(w'))$ .
- $h_1^\dagger = g'_1 \cdot h'_1 \cdot f_1$ .
- $\mu \in \mathbf{effs}_{\varepsilon, h_1^\dagger}^{A', R[\rho \mapsto r]} w''$ .
- $(g'_1, h'_1, f_1, g'_2, h'_2, f_2) \in \mathbf{splits}^{\Pi, \Lambda, \rho, A', R[\rho \mapsto r]} w''$ .
- $(j - n, g_1, g_2, g'_1, g'_2) \in \mathbf{Q}_\varepsilon^{\Pi, R[\rho \mapsto r]} w', w''$ .
- $(j - n, h_1^\circ, h_2^\circ, e'_1, e'_2, h'_1, h'_2) \in \mathbf{safe}_{\tau, \varepsilon}^{\Pi, \Lambda, \rho, A', R[\rho \mapsto r]} w^\circ, w''$ .

Now we must have  $r \in \text{dom}(w'')$  and letting  $A'' = (A \cap \text{dom}(w'')) \cup (\text{dom}(w'') \setminus \text{dom}(w'))$  gives  $r \in A''$  as well as  $A'' \setminus \{r\} = A'$ . With  $A''$  as the choice of new anonymous regions, the obligations in proving the termination branch of (3) are easily met; we just remark that we rely on the induction hypothesis to establish

$$(j - n, h_1^\circ, h_2^\circ, e'_1, e'_2, h'_1, h'_2) \in \mathbf{safe}_{\tau, \varepsilon - \rho}^{\Pi, \Lambda, A'', R} w^\circ, w''$$

◀

There are many details to the proof, but the idea is simple: we simultaneously view the computation in two ways, both with  $\rho$  as a private region and with  $\rho$  masked out. From the safety of former, we then get safety of the latter. In the masked out case,  $\rho$  is no longer a private region variable and the region name  $r$  associated with  $\rho$  joins the anonymous regions. In most cases, however, (region names associated with) private regions and anonymous regions are treated the same, so there is work to do only when anonymous regions are considered in isolation.

**Lemma 6.**  $\cdot | \Pi, \Lambda | \Gamma \models e_1 \preceq e_2 : \tau, \varepsilon$  implies  $\Pi | \Lambda | \Gamma \models \text{atomic } e_1 \preceq \text{atomic } e_2 : \tau, \varepsilon$  if also  $\varepsilon \subseteq \text{rds } \varepsilon \cap \text{wrs } \varepsilon$ .



**Proof.** Let  $a \in \mathbb{N}$  be the required number of anonymous regions from the assumption. To prove the desired, we unroll the definition of the logical relation: choose  $a$ , pick arbitrary  $k^\circ \in \mathbb{N}$ ,  $w^\circ \in \mathbf{W}$ ,  $R : \Pi \cup \Lambda \leftrightarrow |w^\circ|$ ,  $A \subseteq \text{dom}(w^\circ)$ ,  $\gamma_1, \gamma_2 \in \mathcal{V}^{|\Gamma|}$  and  $h_1^\circ, h_2^\circ, h_1^{\circ\prime}, h_2^{\circ\prime} \in \mathcal{H}$ . Assume that  $R(\text{FRV}(\varepsilon)) \subseteq \text{dom}(w^\circ)$ ,  $A \# R(\Pi \cup \Lambda)$ ,  $|A| \geq a$ ,  $\forall r \in A. w^\circ(r) = \emptyset$ ,  $(k^\circ, \gamma_1, \gamma_2) \in \llbracket \Gamma \rrbracket^{R, w^\circ}$ ,  $h_1^{\circ\prime} \subseteq h_1^\circ$ ,  $h_2^{\circ\prime} \subseteq h_2^\circ$ , and  $(k^\circ, h_1^{\circ\prime}, h_2^{\circ\prime}) \in \mathbf{P}_{\varepsilon}^{\Lambda, R, w^\circ}$ . We must show that  $(k^\circ, h_1^{\circ\prime}, h_2^{\circ\prime}, \text{atomic } e_1[\gamma_1/\Gamma], \text{atomic } e_2[\gamma_2/\Gamma], h_1^\circ, h_2^\circ) \in \mathbf{safe}_{\tau, \varepsilon}^{\Pi, \Lambda, A, R, w^\circ, w^\circ}$ .

We need to generalize a bit to handle the loopy behavior of atomic: we prove that for any  $k \leq k^\circ$  and any  $w \in \mathbf{W}$  with  $\mathbf{envtran}^{\Pi, \Lambda, A, R, w^\circ, w}$  we have

$$(k, h_1^{\circ\prime}, h_2^{\circ\prime}, \text{atomic } e_1[\gamma_1/\Gamma], \text{atomic } e_2[\gamma_2/\Gamma], h_1^\circ, h_2^\circ) \in \mathbf{safe}_{\tau, \varepsilon}^{\Pi, \Lambda, A, R, w^\circ, w}.$$

This we do by well-founded induction on  $k$ . Unroll the definition of safety. Pick arbitrary  $j \leq k$ ,  $w' \in \mathbf{W}$  and  $g_1, g_2, f_1, f_2 \in \mathcal{H}$ . Assume that the prerequisites hold, i.e., that we have  $\mathbf{envtran}^{\Pi, \Lambda, A, R, w, w'}$ ,  $(j, g_1, g_2) \in \mathbf{P}_{\varepsilon}^{\Pi, R, w'}$  and  $(g_1, h_1, f_1, g_2, h_2, f_2) \in \mathbf{splits}^{\Pi, \Lambda, A, R, w'}$ . Atomic commands always have the possibility of looping; hence we need not consider the termination branch. So we get to the progress branch: we assume that there are  $e_1', h_1^\dagger$ ,  $\mu$  and  $n \leq j - 1$  such that  $(\text{atomic } e_1[\gamma_1/\Gamma] | g_1 \cdot h_1 \cdot f_1) \rightarrow_{\mu}^{n+1} (e_1' | h_1^\dagger)$  and we must match this.

Observe first that  $\mathbf{envtran}^{\Pi, \Lambda, A, R, w^\circ, w'}$  holds. If, now, the configuration loops to itself, we take no steps on the right hand side and conclude with the induction hypothesis. So we are left to consider the case where  $(e_1[\gamma_1/\Gamma] | g_1 \cdot h_1 \cdot f_1) \xrightarrow{\mu}^n (e_1' | h_1^\dagger)$  and  $e_1' \in \mathcal{V}$ . Let  $m \in \mathbb{N}$  be the number of reduction steps, we name the configurations  $(e_1^0 | h_1^0), (e_1^1 | h_1^1), \dots, (e_1^m | h_1^m)$  in order and pick  $n_1, n_2, \dots, n_m \in \mathbb{N}$  non-zero and  $\mu_1, \mu_2, \dots, \mu_m$  such that

$$(e_1[\gamma_1/\Gamma] | g_1 \cdot h_1 \cdot f_1) = (e_1^0 | h_1^0) \rightarrow_{\mu_1}^{n_1} (e_1^1 | h_1^1) \rightarrow_{\mu_2}^{n_2} \dots \rightarrow_{\mu_m}^{n_m} (e_1^m | h_1^m) = (e_1' | h_1^\dagger)$$

with  $n = n_1 + n_2 + \dots + n_m$  and  $\mu = \mu_1 \cup \mu_2 \cup \dots \cup \mu_m$ . For convenience, we furthermore write  $N_j = \sum_{i=j+1}^m n_i$  for each  $0 \leq j \leq m$ .

It is time for the crux: for any  $0 \leq j \leq m$  there are  $w'', A' \subseteq \text{dom}(w'')$ ,  $h_1', e_2', h_2'$  such that

- $(e_2[\gamma_2/\Gamma] | g_2 \cdot h_2 \cdot f_2) \xrightarrow{*} (e_2' | h_2' \cdot f_2)$
- $\mathbf{selftran}^{\Pi, \Lambda, A, R, w', w''}$
- $h_1' \cdot f = h_1^j$
- $\bigcup_{i=1}^j \mu_i \setminus (\text{dom}_1^{\text{R}(\Pi \cup \Lambda)}(w'') \setminus \text{dom}_1^{\text{R}(\Pi \cup \Lambda)}(w')) \in \mathbf{effs}_{\varepsilon, h_1'}^{A', R, w''}$
- $(\llbracket \cdot \rrbracket, h_1', f_1, \llbracket \cdot \rrbracket, h_2', f_2) \in \mathbf{splits}^{\Pi, \Lambda, A', R, w''}$
- $(N_j + j - n, g_1 \cdot h_1, g_2 \cdot h_2, e_1^j, e_2', h_1', h_2') \in \mathbf{safe}_{\tau, \varepsilon}^{\Pi, \Lambda, A', R, w', w''}$ .

Notice the abuse of notation in item four: we do not really remove locations, we remove any actual effects tagged with the listed locations. The base case follows from the assumption of the lemma; the induction simply by unrolling the safety. In the end, we are able to use the properties in the case  $j = m$  to produce a right hand side reduction in the overall proof. There are many details to this; here we just remark that it would be more more pleasant to have

$$\bigcup_{i=1}^j \mu_i \in \mathbf{effs}_{\varepsilon, h_1'}^{A', R, w''}$$

as item four. Unfortunately, that would break in the inductive step, and we are stuck with the more complex version, which again forces the side condition on the rule. ◀

**Lemma 7.**  $\Pi, \Lambda \mid \cdot \mid \Gamma \models e_1 \preceq e_2 : \tau, \varepsilon$  and  $\Pi, \Lambda \mid \cdot \mid \Gamma \models e_1^\dagger \preceq e_2^\dagger : \tau^\dagger, \varepsilon^\dagger$  together implies  $\Pi \mid \Lambda \mid \Gamma \models \text{par } e_1$  and  $e_1^\dagger \preceq \text{par } e_2$  and  $e_2^\dagger : \tau \times \tau^\dagger, \varepsilon \cup \varepsilon^\dagger$ .

**Proof.** Let  $a_1, a_2 \in \mathbb{N}$  be the required numbers of anonymous regions from the respective assumptions. To prove the desired, we unroll the definition of the logical relation: choose  $a_1 + a_2$ , pick arbitrary  $k^\circ \in \mathbb{N}$ ,  $w^\circ \in \mathbf{W}$ ,  $R : \Pi \cup \Lambda \leftrightarrow |w^\circ|$ ,  $A^\circ \subseteq \text{dom}(w^\circ)$ ,  $\gamma_1, \gamma_2 \in \mathcal{V}^{|\Gamma|}$  and  $h_1^\circ, h_2^\circ, h_1^{\prime\circ}, h_2^{\prime\circ} \in \mathcal{H}$ . Assume that  $R(\text{FRV}(\varepsilon - \rho)) \subseteq \text{dom}(w^\circ)$ ,  $A^\circ \# R(\Pi \cup \Lambda)$ ,  $|A^\circ| \geq a_1 + a_2$ ,  $\forall r \in A^\circ. w^\circ(r) = \emptyset$ ,  $(k^\circ, \gamma_1, \gamma_2) \in \llbracket \Gamma \rrbracket^R w^\circ$ ,  $h_i^{\prime\circ} \subseteq h_i^\circ$ , and  $(k^\circ, h_1^{\prime\circ}, h_2^{\prime\circ}) \in \mathbf{P}_{\varepsilon_1 \cup \varepsilon_2}^{\Lambda, R} w^\circ$ . We must show that  $(k^\circ, h_1^{\prime\circ}, h_2^{\prime\circ}, \text{par } e_1$  and  $e_1^\dagger[\gamma_1/\Gamma], \text{par } e_2$  and  $e_2^\dagger[\gamma_2/\Gamma], h_1^\circ, h_2^\circ) \in \mathbf{safe}_{\tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2}^{\Pi, \Lambda, A^\circ, R} w^\circ, w^\circ$ .

We need, obviously, to make use of the assumptions of the lemma. To this end, we claim that to have

$$(k, h_1^{\prime\circ}, h_2^{\prime\circ}, \text{par } e_1^1 \text{ and } e_2^1, \text{par } e_2^2 \text{ and } e_1^2, h_1 \cdot h_1^1 \cdot h_1^2, h_2 \cdot h_2^1 \cdot h_2^2) \in \mathbf{safe}_{\tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2}^{\Pi, \Lambda, A_1 \cup A_2, R} w^\circ, w, \quad (5)$$

it suffices to know that

$$(k, h_1, h_2) \in \mathbf{P}_{\varepsilon_1 \cup \varepsilon_2}^{\Lambda, R} w \quad (6)$$

$$(h, h_1^{\prime\circ}, h_2^{\prime\circ}, h_1, h_2) \in \mathbf{Q}_{\varepsilon_1 \cup \varepsilon_2}^{\Lambda, R} w^\circ, w \quad (7)$$

$$(k, \emptyset, \emptyset, e_1^1, e_2^1, h_1^1, h_2^1) \in \mathbf{safe}_{\tau_1, \varepsilon_1}^{\Pi \cup \Lambda, \cdot, A_1, R} w^\circ, w \quad (8)$$

$$(k, \emptyset, \emptyset, e_1^2, e_2^2, h_1^2, h_2^2) \in \mathbf{safe}_{\tau_2, \varepsilon_2}^{\Pi \cup \Lambda, \cdot, A_2, R} w^\circ, w, \quad (9)$$

whenever  $A_1 \# A_2$  and  $\text{dom}_i^{A_j}(w) \subseteq \text{dom}(h_i^j)$  holds. Observe, that if we can prove this, then we are done by splitting  $A^\circ$  into two parts of appropriate sizes followed by an easy application of the assumptions of the lemma. The remaining part is to prove the claim, this we do by complete induction on  $k$ .

Assume the claim holds for all naturals strictly less than  $k \in \mathbb{N}$ ; we will now prove it also holds for  $k$ . To prove 5 we unroll the definition of safety. Pick arbitrary  $j \leq k$ ,  $w' \in \mathbf{W}$  and  $g_1, g_2, f_1, f_2 \in \mathcal{H}$ . Assume the prerequisites hold, i.e., that we have  $\mathbf{envtran}_{\varepsilon_1 \cup \varepsilon_2}^{\Pi, \Lambda, A_1 \cup A_2, R} w, w', (j, g_1, g_2) \in \mathbf{P}_{\varepsilon_1 \cup \varepsilon_2}^{\Pi, R} w'$  and  $(g_1, h_1 \cdot h_1^1 \cdot h_1^2, f_1, g_2, h_2 \cdot h_2^1 \cdot h_2^2, f_2) \in \mathbf{splits}_{\varepsilon_1 \cup \varepsilon_2}^{\Pi, \Lambda, A_1 \cup A_2, R} w'$ . Observe that we can instantiate assumptions 8 and 9 by taking  $g_i \cdot h_i$  as the public part of the heap, and adding the spare private part to the frame.

We follow the termination branch first, and assume that  $\text{irr}(\text{par } e_1^1 \text{ and } e_2^1 | g_1 \cdot h_1 \cdot h_1^1 \cdot h_1^2)$ . This means each of the subexpressions is also irreducible so from safety assumptions we learn, among other things, that  $e_1^1, e_2^1 \in \mathcal{V}$ . However, this means that  $(\text{par } e_1^1 \text{ and } e_2^1 | g_1 \cdot h_1 \cdot h_1^1 \cdot h_1^2) \mapsto ((e_1^1, e_2^1) | g_1 \cdot h_1 \cdot h_1^1 \cdot h_1^2)$ , which contradicts the irreducibility assumption.

For the progress branch, we are left with three possibilities of reduction: either both  $e_1^1$  and  $e_2^1$  are values, or the reduction happens inside one of them. We omit the case for reduction inside  $e_2^1$ , since it is completely symmetric to the one for  $e_1^1$ ; the other two cases follow.

Assume that  $(e_1^1 | g_1 \cdot h_1 \cdot h_1^1 \cdot h_1^2 \cdot f_1) \rightarrow_\mu^n (e_1^1 | h_1^1)$  for some  $n \leq j$ . From 8 we get  $e_2^1, w'', h_1', h_2', g_1', g_2', h_1^{\prime 1}$  and  $h_2^{\prime 1}$ , along with the following properties:

- $(e_2^1 | g_2 \cdot h_2 \cdot h_2^1 \cdot h_2^2 \cdot f_2) \mapsto^* (e_2^1 | g_2' \cdot h_2' \cdot h_2^{\prime 1} \cdot h_2^{\prime 2} \cdot f_2)$ ,
- $\mathbf{selftran}_{\varepsilon_1 \cup \varepsilon_2}^{\Pi \cup \Lambda, \cdot, A_1, R} w', w''$ ,
- $A_1' = (A_1 \cap \text{dom}(w'')) \cup (\text{dom}(w'') \setminus \text{dom}(w'))$ ,
- $h_1^{\prime 1} = g_1' \cdot h_1' \cdot h_1^{\prime 1} \cdot h_1^{\prime 2} \cdot f_1$ ,
- $\mu \in \mathbf{effs}_{\varepsilon_2, h_1^{\prime 1}}^{A_1', R} w''$

- $(g'_1 \cdot h'_1, h'_1, f_1 \cdot h_1^2, g'_2 \cdot h'_2, h_2^1, f_2 \cdot h_2^2) \in \mathbf{splits}^{\Pi \cup \Lambda, \cdot, A'_2, R} w''$ ,
- $(j - n, g_1 \cdot h_1, g_2 \cdot h_2, g'_1 \cdot h'_1, g'_2 \cdot h'_2) \in \mathbf{Q}_{\varepsilon_1}^{\Pi \cup \Lambda, R} w', w''$ ,
- $(j - n, \emptyset, \emptyset, e_1^1, e_2^1, h_1^1, h_1^2) \in \mathbf{safe}_{\tau_1, \varepsilon_1}^{\Pi \cup \Lambda, \cdot, A'_2, R} w^\circ, w''$ .

Obviously, we can now replay the reduction sequence for the whole right-hand-side expression. By postcondition separation and postcondition weakening lemmas, we can also split up the postcondition into separate parts for  $\Pi$  and  $\Lambda$ , the first of which is needed by progress branch. The other obligations of the progress branch are simple, leaving us with the final safety requirement:  $(j - n, h_1^\circ, h_2^\circ, \text{par } e_1^1 \text{ and } e_2^1, \text{par } e_2^1 \text{ and } e_2^2, h_1^1 \cdot h_1^1 \cdot h_1^2, h_2^1 \cdot h_2^1 \cdot h_2^2) \in \mathbf{safe}_{\tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2}^{\Pi, \Lambda, A'_1 \cup A_2, R} w^\circ, w''$ . This follows by induction hypothesis, since  $n > 0$ , if we can show the four assumptions. 6 and 7 hold due to downwards-closure and composition lemmas for pre- and post-condition, and the  $\Lambda$  part of the assumption above, and 8 we have verbatim as an assumption. This leaves us only 9 to show. To show it, notice that safety is both downwards closed in the step index, and closed under environment transitions and observe that we have  $\mathbf{envtran}^{\Pi \cup \Lambda, \cdot, A_2, R} w, w''$  by composing the two world transitions above. This, in conjunction with the original assumption, suffices to finish this part of the proof.

The final case we consider is when  $e_1^1, e_1^2 \in \mathcal{V}$  and  $(\text{par } e_1^1 \text{ and } e_1^2 | g_1 \cdot h_1 \cdot h_1^1 \cdot h_1^2) \rightarrow_{\emptyset}^1 (\langle e_1^1, e_1^2 \rangle | g_1 \cdot h_1 \cdot h_1^1 \cdot h_1^2)$ . In this case, we start with instantiating 8 and from the termination branch obtain:

- $(e_2^1 | g_2 \cdot h_2 \cdot h_2^1 \cdot h_2^2 \cdot f_2) \mapsto^* (e_2^1 | g'_2 \cdot h'_2 \cdot h_2^1 \cdot h_2^2 \cdot f_2)$ ,
- $\mathbf{selftran}^{\Pi \cup \Lambda, \cdot, A_1, R} w', w''$ ,
- $\emptyset = (A_1 \cap \text{dom}(w'')) \cup (\text{dom}(w'') \setminus \text{dom}(w'))$ ,
- $g_1 \cdot h_1 \cdot h_1^1 = g'_1 \cdot h'_1 \cdot h_1^1$ ,
- $(j, e_1^1, e_2^1) \in \llbracket \tau_1 \rrbracket^R (w'')$ ,
- $(j, g_1 \cdot h_1, g_2 \cdot h_2, g'_1 \cdot h'_1, g'_2 \cdot h'_2) \in \mathbf{Q}_{\varepsilon_1}^{\Pi \cup \Lambda, R} w', w''$ .

Now, with the new heaps, we can instantiate 9, and get:

- $(e_2^2 | g'_2 \cdot h'_2 \cdot h_2^1 \cdot h_2^2 \cdot f_2) \mapsto^* (e_2^2 | g''_2 \cdot h''_2 \cdot h_2^1 \cdot h_2^2 \cdot f_2)$ ,
- $\mathbf{selftran}^{\Pi \cup \Lambda, \cdot, A_2, R} w'', w'''$ ,
- $\emptyset = (A_2 \cap \text{dom}(w''')) \cup (\text{dom}(w''') \setminus \text{dom}(w''))$ ,
- $g'_1 \cdot h'_1 \cdot h_1^2 = g''_1 \cdot h''_1 \cdot h_1^2 s$ ,
- $(j, e_2^1, e_2^2) \in \llbracket \tau_2 \rrbracket^R (w''')$ ,
- $(j, g'_1 \cdot h'_1, g'_2 \cdot h'_2, g''_1 \cdot h''_1, g''_2 \cdot h''_2) \in \mathbf{Q}_{\varepsilon_2}^{\Pi \cup \Lambda, R} w'', w'''$ .

Now we can build the complete reduction for the right-hand side:  $(\text{par } e_2^1 \text{ and } e_2^2 | g_2 \cdot h_2 \cdot h_2^1 \cdot h_2^2 \cdot f_2) \mapsto^* (\langle e_2^1, e_2^2 \rangle | g''_2 \cdot h''_2 \cdot h_2^1 \cdot h_2^2 \cdot f_2)$ , and prove the rest of the required obligations: self-transition out of the composition of the two given, postcondition of public heap fragments by postcondition separation, composition and weakening, and the rest by simple manipulations. This leaves the final safety of two pairs to be proven. We unroll the definition again, taking  $i \leq j - 1$ , this time only considering the termination branch, since the left-hand side is a value, and do not do any reductions. This leaves us with proving two interesting obligations. The first one is  $(i, \langle e_1^1, e_1^2, \cdot \rangle, \langle e_2^1, e_2^2 \rangle) \in \llbracket \tau_1 \times \tau_2 \rrbracket^R (w^\dagger)$ . However, since  $w^\dagger \sqsupseteq w'''$ , we can easily show this by definition of the denotation of product types, downwards- and future-world-closure, and the value assumptions obtained above. The other interesting obligation is showing the postcondition for the local,  $\Lambda$  part of the heap. This is done by using postcondition separation, composition and weakening on the postcondition properties obtained from assumptions and by assumption 7 itself.  $\blacktriangleleft$

## A.7 Proof of the Parallelization Theorem

The proof of the Parallelization Theorem is quite tricky. We sketch the idea of the proof in the next section, then proceed with the technical formulation of potential safety and Catch-up Lemma, as well as formal proofs, in the subsequent one.

### A.7.1 Informal overview

The left-to-right direction of the theorem is not too hard, intuitively since sequential sequencing clearly can always be mimicked by parallel reduction. Thus, the hard part of the proof is to show the direction

$$\Pi \mid \Lambda \mid \Gamma \models \text{par } e_1 \text{ and } e_2 \lesssim_{\downarrow} \langle e_1, e_2 \rangle : \tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2. \quad (10)$$

The challenge is that if we take a step of reduction on the left, then we cannot always immediately mimick it on the right, because if it is a step in (a derivative of)  $e_2$  then we might not be ready to make it on the right, since we haven't yet started reducing in  $e_2$  on the right (we only start reducing  $e_2$  on the right when we are done with all the  $e_1$  reductions).

We address this issue by formulating a notion of *potential safety* to relate (a derivative of)  $e_2$  on the left with (a derivative of)  $e_2$  on the right. Safety, and also potential safety, really involves configurations (expressions and heaps) rather than just expression), but to keep the overview, we will continue just using expression notation in the high-level sketch here. Potential safety of  $e_2$  and  $\tilde{e}_2$  then allows  $\tilde{e}_2$  to *catch up* by doing some reductions to reach  $\tilde{e}'_2$ , and then  $e_2$  and  $\tilde{e}'_2$  should be safe (in the usual sense).

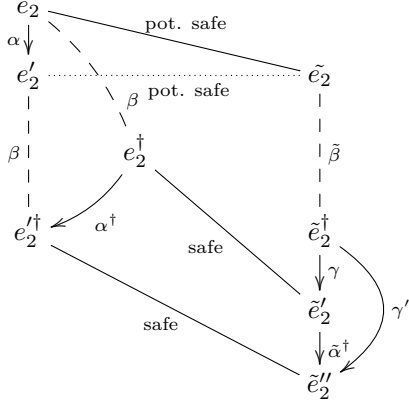
The idea of the proof of (10) is then to show (for suitable  $e_1, \tilde{e}_1, e_2, \tilde{e}_2$  satisfying conditions corresponding to those in the parallelization theorem) that

$$\begin{aligned} & \text{safe}(e_1, \tilde{e}_1) \wedge \text{potentially-safe}(e_2, \tilde{e}_2) \\ & \implies \text{safe}(\langle \text{par } e_1 \text{ and } e_2 \rangle, \langle \tilde{e}_1, \tilde{e}_2 \rangle). \end{aligned} \quad (11)$$

Now, for this to work, potential safety actually needs to allow for some transitions by  $e_1$  and  $\tilde{e}_1$  — and the context — but regulated by the overall assumptions regarding effects, etc., of the parallelization theorem. Thus potential safety of  $e_2$  and  $\tilde{e}_2$  actually says that for any good (in the sense that it is governed by overall assumptions regarding effects, etc.) pair of transitions,  $\beta$  and  $\tilde{\beta}$ , going from  $e_2$  to  $e'_2$  and  $\tilde{e}_2$  to  $\tilde{e}'_2$  respectively, there is some sequence of catching-up reductions, call them  $\gamma$ , that reaches a  $\tilde{e}'_2$  such that  $e'_2$  and  $\tilde{e}'_2$  are safe. Diagrammatically:

$$\begin{array}{ccc} e_2 & \xrightarrow{\text{pot. safe}} & \tilde{e}_2 \\ \beta \downarrow & & \downarrow \tilde{\beta} \\ e'_2 & & \tilde{e}'_2 \\ & \searrow \text{safe} & \downarrow \gamma \\ & & \tilde{e}'_2 \end{array}$$

To show (11), we proceed by induction on the index used in the safety predicate. There are three cases to consider, corresponding to possible reductions of  $\text{par } e_1 \text{ and } e_2$ . The case where  $e_1$  reduces is not too hard since we have defined potential safety to accommodate such reductions (the  $\beta$ 's). The case where  $e_1$  and  $e_2$  are both values is not too hard either, intuitively since we can just run  $\tilde{e}_2$  until we reach safety with  $e_2$  at which point the proof becomes simple. The case where  $e_2$  reduces, say to  $e'_2$  by a reduction  $\alpha$ , is more tricky. We proceed by showing potential safety of  $e'_2$  and  $\tilde{e}_2$  (argument outline follows below), and that



■ **Figure 10** Proving parallelization: recovering potential safety after  $\alpha$  transition.

$e_1$  and  $\tilde{e}_1$  stay safe (this follows from the assumption on effects, etc.), which allows us to conclude by induction.

The argument for the potential safety of  $e'_2$  and  $\tilde{e}_2$  is outlined in Figure 10. By assumption  $e_2$  and  $\tilde{e}_2$  are potentially safe, as depicted by the topmost arrow, and we need to show the potential safety depicted by the dashed arrow. Hence we consider any  $\beta$  and  $\tilde{\beta}$  as shown, and we need to show that there exists a catching-up reduction  $\gamma'$  from  $\tilde{e}_2^\dagger$  to  $\tilde{e}_2''$ . By properties of the instrumented operational semantics, we can commute  $\alpha$  and  $\beta$  to get  $\beta$  and  $\alpha^\dagger$ , as shown in the diagram (essentially, this is Lemma 2). Hence by the assumption of potential safety of  $e_2$  and  $\tilde{e}_2$ , we get safety of  $e_2^\dagger$  and  $\tilde{e}_2^\dagger$ , as shown. Using this safety and the  $\alpha^\dagger$  reduction, we get a reduction  $\tilde{\alpha}^\dagger$  and the final required safety of  $e_2^\dagger$  and  $\tilde{e}_2''$ . Thus the required catching-up reduction  $\gamma'$  is  $\gamma$  followed by  $\tilde{\alpha}^\dagger$ . Note that the latter makes intuitive sense; we did a reduction on the left in  $e_2$  and now we have correspondingly extended the catching-up reduction to be done on the right.

Formally, the Catch-up lemma in Appendix A.7 gives the precise formulation and proof of (11); item 3 in that lemma gives the precise technical formulation of potential safety.

### A.7.2 Technical details

We start by stating and proving the catch-up lemma mentioned in the explanation in the paper, that covers most of the proof of the theorem.

► **Lemma 17 (Catch-up).** *Assuming that*

1.  $(k, h_1, h_2) \in \mathbf{P}_{\varepsilon_1}^{\Lambda, R} w$ ,
2.  $(k, \emptyset, \emptyset, e_1, \tilde{e}_1, h_{11}, h_{12}) \in \mathbf{safe}_{\tau_1, \varepsilon_1}^{\Pi \cup \Lambda, \cdot, A_1, R} w^\circ, w$  and
3. for any  $g_1, g_2, h_1^\dagger, h_2^\dagger, f_1, f_2, w', j \leq k$  such that
  - $\mathbf{envtran}^{\Pi \cup \Lambda, \cdot, A_2, R} w, w'$
  - $(g_1 \cdot h_1^\dagger, h_{21}, f_1, g_2 \cdot h_2^\dagger, \emptyset, f_2) \in \mathbf{splits}^{\Pi \cup \Lambda, \cdot, A_2, R} w'$
  - $(j, h_1, h_2, h_1^\dagger, h_2^\dagger) \in \mathbf{Q}_{\varepsilon_1}^{\Lambda, R} w, w'$

there exist  $h'_1, h'_2, g'_1, g'_2, h'_{21}, h'_{22}, \tilde{e}'_1, w''$  such that

- $\mathbf{selftran}^{\Pi \cup \Lambda, \cdot, A_2, R} w', w''$
- $(\tilde{e}_2 | h_2^\dagger \cdot g_2 \cdot f_2) \mapsto^* (e'_2 | h'_2 \cdot g'_2 \cdot h'_{22} \cdot f_2)$
- $h_1^\dagger \cdot g_1 \cdot h_{21} = h'_1 \cdot g'_1 \cdot h'_{21}$
- $(k, g_1, g_2, g'_1, g'_2) \in \mathbf{Q}_{\varepsilon_2}^{\Pi, R} w', w''$

- $(k, h'_1, h'_2) \in \mathbf{P}_{\varepsilon_2}^{\Lambda, R} w''$
- $(j, h_1^\circ, h_2^\circ, h'_1, h'_2) \in \mathbf{Q}_{\varepsilon_1 \cup \varepsilon_2}^{\Lambda, R} w^\circ, w''$
- $(k, \emptyset, \emptyset, e_2, \tilde{e}'_2, h'_{21}, h_{22}) \in \mathbf{safe}_{\tau_2, \varepsilon_2}^{\Pi \cup \Lambda, \cdot, A'_2, R} w^\circ, w''$ ,

we have  $(k, h_1^\circ, h_2^\circ, \text{par } e_1 \text{ and } e_2, \langle \tilde{e}_1, \tilde{e}_2 \rangle, h_1 \cdot h_{11} \cdot h_{21}, h_2 \cdot h_{12}) \in \mathbf{safe}_{\tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2}^{\Pi, \Lambda, A_1 + A_2, R} w^\circ, w$ , as long as  $\text{rds } \varepsilon_1 \cup \text{wrs } \varepsilon_1 \cup \text{rds } \varepsilon_2 \cup \text{wrs } \varepsilon_2 \subseteq \Lambda$ ,  $\text{rds } \varepsilon_1 \# \text{wrs } \varepsilon_2$ ,  $\text{rds } \varepsilon_2 \# (\text{wrs } \varepsilon_1 \cup \text{als } \varepsilon_1)$ ,  $\text{wrs } \varepsilon_1 \# \text{wrs } \varepsilon_2$ ,  $\text{dom}(h_{i1}) \supseteq \text{dom}_1^{A_i}(w)$ ,  $\text{dom}(h_{12}) \supseteq \text{dom}_2^{A_1}(w)$  and  $\text{dom}_2^{A_2}(w) = \emptyset$ .

**Proof.** The proof proceeds by well-founded induction on  $k$ . Take any  $j \leq k, g_1, g_2, f_1, f_2, w_1$  such that

- $\mathbf{envtran}^{\Pi, \Lambda, A_1 + A_2, R} w, w_1$
- $(j, g_1, g_2) \in \mathbf{P}_{\varepsilon_1 \cup \varepsilon_2}^{\Pi, R} w_1$
- $(g_1, h_1 \cdot h_{11} \cdot h_{21}, f_1, g_2, h_2 \cdot h_{12}, f_2) \in \mathbf{splits}^{\Pi, \Lambda, A_1 + A_2, R} w_1$

There are two branches to consider. However,  $\text{par } e_1 \text{ and } e_2$  can always reduce, so the termination branch is trivial. In the progress case, we proceed by case analysis on the reduction to get three subcases.

1.  $(e_1 \mid g_1 \cdot h_1 \cdot h_{11} \cdot h_{21} \cdot f_1) \rightarrow_\mu^n (e'_1 \mid h_1^\dagger)$

In this case we can use the assumption (2), which we instantiate with  $j \leq k, g_1 \cdot h_1, g_2 \cdot h_2, f_1 \cdot h_{21}, f_2, w_1$ . We need to show the prerequisites:

- $\mathbf{envtran}^{\Pi \cup \Lambda, \cdot, A_1, R} w, w_1$ , which holds by env-transition weakening
- $(j + 1, g_1 \cdot h_1, g_2 \cdot h_2) \in \mathbf{P}_{\varepsilon_1}^{\Pi \cup \Lambda, R} w_1$ , which holds by assumptions, precondition composition, precondition weakening and future-world closure of precondition
- $(g_1 \cdot h_1, h_{11}, f_1 \cdot h_{21}, g_2 \cdot h_2, h_{12}, f_2) \in \mathbf{splits}^{\Pi \cup \Lambda, \cdot, A_1, R} w_1$ , which holds by massaging the assumption about splits.

From the progress branch we obtain the following properties:

- $(\tilde{e}_1 \mid g_2 \cdot h_2 \cdot h_{12} \cdot f_2) \mapsto^* (e'_1 \mid g'_2 \cdot h'_2 \cdot h'_{12} \cdot f_2)$
- $\mathbf{selftran}^{\Pi \cup \Lambda, \cdot, A_1, R} w_1, w_2$
- $A'_1 = (A_1 \cap \text{dom}(w_2)) \cup (\text{dom}(w_2) \setminus \text{dom}(w_1))$
- $h_1^\dagger = g'_1 \cdot h'_1 \cdot h'_{11} \cdot f_1 \cdot h_{21}$
- $\mu \in \mathbf{effs}_{\varepsilon_1, h'_{11}}^{A'_1, R} w_2$
- $(g'_1 \cdot h'_1, h'_{11}, f_1 \cdot h_{21}, g'_2 \cdot h'_2, h'_{12}, f_2) \in \mathbf{splits}^{\Pi \cup \Lambda, \cdot, A'_1, R} w_2$
- $(j - n, g_1 \cdot h_1, g_2 \cdot h_2, g'_1 \cdot h'_1, g'_2 \cdot h'_2) \in \mathbf{Q}_{\varepsilon_1}^{\Pi \cup \Lambda, R} w_1, w_2$
- $(j - n, \emptyset, \emptyset, e'_1, \tilde{e}'_1, h'_{11}, h'_{12}) \in \mathbf{safe}_{\tau_1, \varepsilon_1}^{\Pi \cup \Lambda, \cdot, A'_1, R} w^\circ, w_2$

From these we can conclude that:

- $(\langle \tilde{e}_1, \tilde{e}_2 \rangle \mid g_2 \cdot h_2 \cdot h_{12} \cdot f_2) \mapsto^* (\langle \tilde{e}'_1, \tilde{e}'_2 \rangle \mid g'_2 \cdot h'_2 \cdot h'_{12} \cdot f_2)$  (recall pairs of expressions is syntactic sugar)
- $\mathbf{selftran}^{\Pi, \Lambda, A_1 + A_2, R} w_1, w_2$ , by self-transition strengthening
- $A'_1 + A_2 = ((A_1 + A_2) \cap \text{dom}(w_2)) \cup (\text{dom}(w_2) \setminus \text{dom}(w_1))$ , by self-transition restrictions
- $(g'_1, h'_1 \cdot h'_{11} \cdot h_{21}, f_1, g'_2, h'_2 \cdot h'_{12}, f_2) \in \mathbf{splits}^{\Pi, \Lambda, A'_1 + A_2, R} w_2$ , by massaging the splits assumption
- $\mu \in \mathbf{effs}_{\varepsilon_1 \cup \varepsilon_2, h'_1 \cdot h'_{11} \cdot h_{21}}^{A'_1 + A_2, R} w_2$ , by effs strengthening
- $(j - n, g_1, g_2, g'_1, g'_2) \in \mathbf{Q}_{\varepsilon_1 \cup \varepsilon_2}^{\Pi, R} w_1, w_2$ , by postcondition separation and postcondition weakening
- $(j - n, h'_1, h'_2) \in \mathbf{P}_{\varepsilon_1}^{\Lambda, R} w_2$ , by postcondition separation and precondition composition
- $\mathbf{envtran}^{\Pi \cup \Lambda, \cdot, A_2, R} w, w_2$ , by envtransition composition and relation of env-transition and self-transition
- $(j - n, h_1, h_2, h'_1, h'_2) \in \mathbf{Q}_{\varepsilon_1}^{\Lambda, R} w, w_2$ , by postcondition separation

Since  $0 < n \leq j$ , we can now use the induction hypothesis to discharge the final safety obligation: we have already obtained preconditions (1) and (2), while (3) is closed under the world transition that we have made (see the final two properties above).

2.  $(e_2 \mid g_1 \cdot h_1 \cdot h_{11} \cdot h_{21} \cdot f_1) \rightarrow_\mu^n (e'_2 \mid h_1^\dagger)$

First, we need to establish that  $h_1^\dagger = g_1 \cdot h'_1 \cdot h'_{21} \cdot f_1 \cdot h_{11}$ ,  $\text{dom}(h'_1) = \text{dom}(h_1)$ ,  $\mu \in \mathbf{effs}_{\varepsilon_2, h'_{21}}^{A, R} w$  and  $(j - n, h'_1, h_2) \in \mathbf{P}_{\varepsilon_1}^{A, R} w_1$ . This can be obtained by using the assumption (3), as follows. Take  $g_1, g_2, h_1, h_2, f_1 \cdot h_{11}, f_2 \cdot h_{12}, w_1, k \leq k$  as the universally quantified variables in (3). We need to show:

- **envtran** $^{\Pi \cup \Lambda, \cdot, A_2, R} w_1, w_1$ , which holds by reflexivity
- $(g_1 \cdot h_1, h_{21}, f_1 \cdot h_{11}, g_2 \cdot h_2, \emptyset, f_2 \cdot h_{12}) \in \mathbf{splits}^{\Pi \cup \Lambda, \cdot, A_2, R} w_1$ , by massaging the splits we have assumed earlier
- $(k, h_1, h_2, h_1, h_2) \in \mathbf{Q}_{\varepsilon_1}^{A, R} w_1, w_1$ , again, by reflexivity.

From this we obtain a bunch of properties, by unfolding safety. However, we are only interested in a few, namely that  $h_1^\dagger = g_1^p \cdot h_1^l \cdot h_1^p \cdot f_1 \cdot h_{11}$ , and  $\mu \in \mathbf{effs}_{\varepsilon_2, h_1^p}^{A_2^\dagger, R} w^\dagger$  and **selftran** $^{\Pi \cup \Lambda, \cdot, A_2, R} w_1, w^\dagger$ . Let  $h_1^\ddagger = g_1^p \cdot h_1^l \cdot h_1^p$ .

Now, as witnesses for the progress case, we take  $e_2, w_1, A_1 + A_2, h'_1 \cdot h_{11} \cdot h'_{21}, h_2 \cdot h_{12}, g_1, g_2$ , where  $h'_1 = h_1^\ddagger \upharpoonright_{\text{dom}_1^R(\Lambda)(w_1)}$  and  $h'_{21} = h_1^\ddagger \upharpoonright_{\text{dom}_1(w_1) \setminus \text{dom}_1^R(\Pi \cup \Lambda)(w_1)}$ . It suffices to show:

- $(\tilde{e}_2 \mid g_2 \cdot h_2 \cdot h_{12} \cdot f_2) \mapsto^* (\tilde{e}_2 \mid g_2 \cdot h_2 \cdot h_{12} \cdot f_2)$  by reflexivity
- **selftran** $^{\Pi, \Lambda, A_1 + A_2, R} w_1, w_1$  by reflexivity
- $h_1^\dagger = g_1 \cdot h'_1 \cdot h_{11} \cdot h'_{21} \cdot f_1$  by the restriction on actual effects ( $\mu$ ) we know  $g_1$  didn't change, and we have an assumption to finish this off
- $(g_1, h'_1 \cdot h_{11} \cdot h'_{21}, f_1, g_2, h_2 \cdot h_{12}, f_2) \in \mathbf{splits}^{\Pi, \Lambda, A_1 + A_2, R} w_1$  by assumption
- $\mu \in \mathbf{effs}_{\varepsilon_2, h'_{21}}^{A_2, R} w_1$ , since any region killed between  $w_1$  and  $w^\dagger$  has to be in  $A_2$ ,  $\text{dom}(h_1^p) \subseteq \text{dom}(h'_{21})$ , and we have an assumption for the other effects inside the world
- $(j - n, g_1, g_2, g_1, g_2) \in \mathbf{Q}_{\varepsilon_1 \cup \varepsilon_2}^{\Pi, R} w_1, w_1$  by reflexivity,

so it remains to show that  $(j - n, h_1^\circ, h_2^\circ, \text{par } e_1 \text{ and } e'_2, (\tilde{e}_1, \tilde{e}_2), h'_1 \cdot h_{11} \cdot h'_{21}, h_2 \cdot h_{12}) \in \mathbf{safe}_{\tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2}^{\Pi, \Lambda, A_1 + A_2, R} w^\circ, w_1$  holds. We can use induction hypothesis to prove it, provided we show the three assumptions. The first one  $((j - n, h'_1, h_2) \in \mathbf{P}_{\varepsilon_1}^{A, R} w_1)$  is simple, since  $\mu$  is confined to  $\varepsilon_2$ , the second one  $((j - n, \emptyset, \emptyset, e_1, \tilde{e}_1, h_{11}, h_{12}) \in \mathbf{safe}_{\tau_1, \varepsilon_1}^{\Pi \cup \Lambda, \cdot, A_1, R} w^\circ, w_1)$  holds by env-transition closure and downwards closure of **safe**. This leaves us the last precondition to show. To this end, we take  $h_1^\dagger, h_2^\dagger, g_1^\dagger, g_2^\dagger, f_1^\dagger, f_2^\dagger, i \leq j - n$  and  $w_2$ , such that

- **envtran** $^{\Pi \cup \Lambda, \cdot, A_2, R} w_1, w_2$
- $(i, h_1^\dagger, h_2, h_1^\dagger, h_2^\dagger) \in \mathbf{Q}_{\varepsilon_1}^{A, R} w_1, w_2$
- $(g_1^\dagger \cdot h_1^\dagger, h'_{21}, f_1^\dagger, g_2^\dagger \cdot h_2^\dagger, \emptyset, f_2^\dagger) \in \mathbf{splits}^{\Pi \cup \Lambda, \cdot, A_2, R} w_2$

Now we can instantiate the original assumption (3) with  $h_1[h_1^\dagger/h'_1], h_2^\dagger, g_1^\dagger, g_2^\dagger, f_1^\dagger, f_2^\dagger, i \leq k$  and  $w_2$ , where the heap substitution notation means that any change (allocation or update) from  $h'_1$  to  $h_1^\dagger$  should be replayed on  $h_1$  (this is well-specified, since  $\text{dom}(h_1) = \text{dom}(h'_1)$ ). We need to prove the following:

- **envtran** $^{\Pi \cup \Lambda, \cdot, A_2, R} w, w_2$ , by envtran-comp and envtran-weaken
- $(i, h_1, h_2, h_1[h_1^\dagger/h'_1], h_2^\dagger) \in \mathbf{Q}_{\varepsilon_1}^{A, R} w, w_2$ , by definition of heap substitution, and assumption
- $(g_1^\dagger \cdot h_1[h_1^\dagger/h'_1], h_{21}, f_1^\dagger, g_2^\dagger \cdot h_2^\dagger, \emptyset, f_2^\dagger) \in \mathbf{splits}^{\Pi \cup \Lambda, \cdot, A_2, R} w_2$  by assumption, given that  $\text{dom}(h_1^\dagger) = \text{dom}(h_1[h_1^\dagger/h'_1])$

By proving these properties, we obtain the following:

- $(k, h_1^\dagger, h_2^\dagger) \in \mathbf{P}_{\varepsilon_2}^{A, R} w_3$
- **selftran** $^{\Pi \cup \Lambda, \cdot, A_2, R} w_2, w_3$

- $(\tilde{e}_2 | h_2^\dagger \cdot g_2^\dagger \cdot f_2^\dagger) \mapsto^* (\tilde{e}'_2 | h_2^\dagger \cdot g_2^\dagger \cdot h_{22} \cdot f_2^\dagger)$
- $h_1[h_1^\dagger/h_1'] \cdot g_1^\dagger \cdot h_{21} = h_1^\dagger \cdot g_1^\dagger \cdot h_{21}^\dagger$
- $(k, g_1^\dagger, g_2^\dagger, g_1^\dagger, g_2^\dagger) \in \mathbf{Q}_{\varepsilon_2}^{\Pi, R} w_2, w_3$
- $(i, h_1^\circ, h_2^\circ, h_1^\dagger, h_2^\dagger) \in \mathbf{Q}_{\varepsilon_1 \cup \varepsilon_2}^{\Lambda, R} w^\circ, w_3$
- $(g_1^\dagger \cdot h_1^\dagger, h_{21}^\dagger, f_1^\dagger, g_2^\dagger \cdot h_2^\dagger, h_{22}, f_2^\dagger) \in \mathbf{plits}^{\Pi \cup \Lambda, \cdot, A_2', R} w_3$
- $(k, \emptyset, \emptyset, e_2, e_2', h_{21}^\dagger, h_{22}) \in \mathbf{safe}_{\tau_2, \varepsilon_2}^{\Pi \cup \Lambda, \cdot, A_2', R} w^\circ, w_3$

Recall that  $\mu \in \mathbf{eff}_{\varepsilon_2, h_{21}'}^{A_2, R} w_1$ , and  $(e_2 | g_1 \cdot h_1 \cdot h_{11} \cdot h_{21} \cdot f_1) \rightarrow_\mu^n (e_2' | g_1 \cdot h_1' \cdot h_{21}' \cdot f_1 \cdot h_{11})$ . Hence, we can use Lemma 2 to get

$$(e_2 | g_1^\dagger \cdot h_1[h_1^\dagger/h_1'] \cdot h_{21} \cdot f_1^\dagger) \rightarrow_\mu^n (e_2' | g_1^\dagger \cdot h_1^\dagger \cdot h_{21}' \cdot f_1^\dagger),$$

which allows us to instantiate the safety predicate above. Note that this argument is precisely the commutation of  $\alpha$  and  $\beta$  in the explanation in Section 4.1 (cf. Figure 10).

We need to show:

- $\mathbf{envtran}^{\Pi \cup \Lambda, \cdot, A_2', R} w_3, w_3$ , by reflexivity
- $(j, g_1^\dagger, g_2^\dagger) \in \mathbf{P}_{\varepsilon_2}^{\Pi, R} w_3$ , trivial since  $\text{rds } \varepsilon_2 \# \Pi$
- $(j, h_1^\dagger, h_2^\dagger) \in \mathbf{P}_{\varepsilon_2}^{\Lambda, R} w_3$ , by assumption
- $(g_1^\dagger \cdot h_1^\dagger, h_{21}^\dagger, f_1 \cdot h_{11} \cdot f_1^\dagger, g_2^\dagger \cdot h_2^\dagger, h_{22}, f_2^\dagger) \in \mathbf{plits}^{\Pi \cup \Lambda, \cdot, A_2', R} w_3$ , by massaging the assumption.

These facts, along with the reduction step, give us the following:

- $\mathbf{selftran}^{\Pi \cup \Lambda, \cdot, A_2', R} w_3, w_4$
- $(\tilde{e}'_2 | h_2^\dagger \cdot g_2^\dagger \cdot h_{22} \cdot f_2^\dagger) \mapsto^* (\tilde{e}''_2 | g_2^f \cdot h_2^f \cdot h_{22}^f \cdot f_2^\dagger)$
- $g_1^\dagger \cdot h_1^\dagger \cdot h_{21}' = g_1^f \cdot h_1^f \cdot h_{21}^f$
- $(g_1^f \cdot h_1^f, h_{21}^f, f_1^\dagger, g_2^f \cdot h_2^f, h_{22}^f, f_2^\dagger) \in \mathbf{plits}^{\Pi \cup \Lambda, \cdot, A_2', R} w_4$
- $(j - n, g_1^\dagger, g_2^\dagger, h_1^\dagger, h_2^\dagger, g_1^f \cdot h_1^f, g_2^f \cdot h_2^f) \in \mathbf{Q}_{\varepsilon_2}^{\Pi \cup \Lambda, R} w_3, w_4$
- $(j - n, \emptyset, \emptyset, e_2', e_2'', h_{21}^f, h_{22}^f) \in \mathbf{safe}_{\tau_2, \varepsilon_2}^{\Pi \cup \Lambda, \cdot, A_2', R} w^\circ, w_4$

Now we can show the remaining goals:

- $\mathbf{selftran}^{\Pi \cup \Lambda, \cdot, A_2, R} w_2, w_4$ , by self-transition composition
  - $(\tilde{e}_2 | g_2^\dagger \cdot h_2^\dagger \cdot f_2^\dagger) \mapsto^* (\tilde{e}''_2 | g_2^f \cdot h_2^f \cdot h_{22}^f \cdot f_2^\dagger)$ , by transitivity
  - $h_1^\dagger \cdot g_1^\dagger \cdot h_{21}' = g_1^f \cdot h_1^f \cdot h_{21}^f$ , by assumption
  - $(j - n, g_1^\dagger, g_2^\dagger, g_1^f, g_2^f) \in \mathbf{Q}_{\varepsilon_2}^{\Pi, R} w_2, w_4$ , by postcondition composition
  - $(j - n, h_1^f, h_2^f) \in \mathbf{P}_{\varepsilon_2}^{\Lambda, R} w_4$ , by precondition composition
  - $(i, h_1^\circ, h_2^\circ, h_1^f, h_2^f) \in \mathbf{Q}_{\varepsilon_1 \cup \varepsilon_2}^{\Lambda, R} w^\circ, w_4$ , by postcondition composition
  - $(j - n, g_1^\dagger \cdot h_1^\dagger, g_2^\dagger \cdot h_2^\dagger, g_1^f \cdot h_1^f, g_2^f \cdot h_2^f) \in \mathbf{Q}_{\varepsilon_2}^{\Pi \cup \Lambda, R} w_3, w_4$ , by assumption
  - $(j - n, \emptyset, \emptyset, e_2', e_2'', h_{21}^f, h_{22}^f) \in \mathbf{safe}_{\tau_2, \varepsilon_2}^{\Pi \cup \Lambda, \cdot, A_2', R} w^\circ, w_4$ , by assumption,
- which ends this case.

### 3. $(\text{par } e_1 \text{ and } e_2 | g_1 \cdot h_1 \cdot h_{11} \cdot h_{21} \cdot f_1) \rightarrow_\emptyset^1 ((e_1, e_2) | g_1 \cdot h_1 \cdot h_{11} \cdot h_{21} \cdot f_1)$

This means in particular that  $\text{irr}(e_i | g_1 \cdot h_1 \cdot h_{11} \cdot h_{21} \cdot f_1)$ . We only need to consider the case where  $j \geq 1$ . We start by using the assumption (2). We need to provide:

- $\mathbf{envtran}^{\Pi \cup \Lambda, \cdot, A_1, R} w, w_1$ , by env-transition weakening
- $(j, g_1 \cdot h_1, g_2 \cdot h_2) \in \mathbf{P}_{\varepsilon_1}^{\Pi \cup \Lambda, R} w_1$ , by precondition separation, assumption and (1)
- $(g_1 \cdot h_1, h_{11}, f_1 \cdot h_{21}, g_2 \cdot h_2, h_{12}, f_2) \in \mathbf{plits}^{\Pi \cup \Lambda, \cdot, A_1, R} w_1$ , by massaging the assumption about **splits**.

From the termination case we can now get:

- $(\tilde{e}_1 | g_2 \cdot h_2 \cdot h_{12} \cdot f_2) \mapsto^* (\tilde{e}'_1 | g_2' \cdot h_2' \cdot h_{12}' \cdot f_2)$
- $\mathbf{selftran}^{\Pi \cup \Lambda, \cdot, A_2, R} w_1, w_2$
- $g_1 \cdot h_1 \cdot h_{11} = g_1' \cdot h_1' \cdot h_{11}'$
- $(g_1' \cdot h_1', h_{11}', f_1 \cdot h_{21}, g_2' \cdot h_2', h_{12}', f_2) \in \mathbf{plits}^{\Pi \cup \Lambda, \cdot, \emptyset, R} w_2$



- $(j, g_1, g_2, g'_1, g'_2) \in \mathbf{Q}_{\varepsilon_1}^{\Pi, R} w_1, w_2$
- $(j, h_1, h_2, h'_1, h'_2) \in \mathbf{Q}_{\varepsilon_1}^{\Lambda, R} w_1, w_2$
- $(j, e_1, \tilde{e}'_1) \in \llbracket \tau_1 \rrbracket^R w_2$

Note that  $\tilde{e}'_1$  is irreducible, as it's a value. At this point we can use assumption (3). We need to show:

- **envtran** $^{\Pi \cup \Lambda, \cdot, A_2, R} w, w_2$ , by relation between self-transition and env-transition and env-transition weakening
- $(j, h_1, h_2, h'_1, h'_2) \in \mathbf{Q}_{\varepsilon_1}^{\Lambda, R} w, w_2$ , by assumption
- $(g'_1 \cdot h'_1, h_{21}, f_1 \cdot h'_{11}, g'_2 \cdot h'_2, \emptyset, f_2 \cdot h'_{12}) \in \mathbf{splits}^{\Pi \cup \Lambda, \cdot, A_2, R} w_2$ , by **splits** above.

By these facts we are able to obtain the following:

- **selftran** $^{\Pi \cup \Lambda, \cdot, A_2, R} w_2, w_3$
- $(\tilde{e}_2 | g'_2 \cdot h'_2 \cdot f_2 \cdot h'_{12}) \mapsto^* (\tilde{e}_2 | g''_2 \cdot h''_2 \cdot h_{22} \cdot f_2 \cdot h'_{12})$
- $h'_1 \cdot g'_1 \cdot h_{21} = h''_1 \cdot g''_1 \cdot h'_{21}$
- $(k, g'_1, g'_2, g''_1, g''_2) \in \mathbf{Q}_{\varepsilon_2}^{\Pi, R} w_2, w_3$
- $(k, h''_1, h''_2) \in \mathbf{P}_{\varepsilon_2}^{\Lambda, R} w_3$
- $(j, h_1^\circ, h_2^\circ, h''_1, h''_2) \in \mathbf{Q}_{\varepsilon_1 \cup \varepsilon_2}^{\Lambda, R} w^\circ, w_3$
- $(k, \emptyset, \emptyset, e_2, \tilde{e}'_2, h''_{21}, h_{22}) \in \mathbf{safe}_{\tau_2, \varepsilon_2}^{\Pi \cup \Lambda, \cdot, A'_2, R} w^\circ, w_3$
- $(g''_1 \cdot h''_1, h''_{21}, f_1 \cdot h'_{11}, g''_2 \cdot h''_2, h_{22}, f_2 \cdot h'_{12}) \in \mathbf{splits}^{\Pi \cup \Lambda, \cdot, A'_2, R} w_3$ .

Finally, we can instantiate the safety predicate obtained above. We need to show:

- **envtran** $^{\Pi \cup \Lambda, \cdot, A'_2, R} w_3, w_3$ , by reflexivity
- $(j, h''_1, h''_2) \in \mathbf{P}_{\varepsilon_2}^{\Lambda, R} w_3$ , by assumption
- $(j, g''_1, g''_2) \in \mathbf{P}_{\varepsilon_2}^{\Pi, R} w_3$ , trivial, since  $\text{rds } \varepsilon_2 \# \Pi$
- $(g''_1 \cdot h''_1, h''_{21}, f_1 \cdot h'_{11}, g''_2 \cdot h''_2, h_{22}, f_2 \cdot h'_{12}) \in \mathbf{splits}^{\Pi \cup \Lambda, \cdot, A'_2, R} w_3$ , by assumption.

From the termination branch (by irreducibility of  $e_2$ ), we get:

- **selftran** $^{\Pi \cup \Lambda, \cdot, A'_2, R} w_3, w_4$
- $(\tilde{e}'_2 | g''_2 \cdot h''_2 \cdot h_{22} \cdot f_2 \cdot h'_{12}) \mapsto^* (\tilde{e}'_2 | g'''_2 \cdot h'''_2 \cdot h'''_{22} \cdot f_2 \cdot h'_{12})$
- $g''_1 \cdot h''_1 \cdot h''_{21} = g'''_1 \cdot h'''_1 \cdot h'''_{21}$
- $(j, g''_1, g''_2, g'''_1, g'''_2) \in \mathbf{Q}_{\varepsilon_2}^{\Pi, R} w_3, w_4$
- $(j, h''_1, h''_2, h'''_1, h'''_2) \in \mathbf{Q}_{\varepsilon_2}^{\Lambda, R} w_3, w_4$
- $(j, e_2, \tilde{e}'_2) \in \llbracket \tau_2 \rrbracket^R w_4$
- $(g'''_1 \cdot h'''_1, h'''_{21}, f_1 \cdot h'_{11}, g'''_2 \cdot h'''_2, h'''_{22}, f_2 \cdot h'_{12}) \in \mathbf{splits}^{\Pi \cup \Lambda, \cdot, \emptyset, R} w_4$

At this point we can finally provide witnesses and prove the remaining properties:

- **selftran** $^{\Pi, \Lambda, A_1 + A_2, R} w_1, w_4$ , by weakening and composition of previous self-transitions
- $(\langle \tilde{e}_1, \tilde{e}_2 \rangle | g_2 \cdot h_2 \cdot h_{12} \cdot f_2) \mapsto^* (\langle \tilde{e}'_1, \tilde{e}'_2 \rangle | g'''_2 \cdot h'''_2 \cdot h'''_{22} \cdot h'_{12} \cdot f_2)$ , by transitivity and irreducibility of  $\tilde{e}'_1$
- $\emptyset \in \mathbf{effs}_{\varepsilon_1 \cup \varepsilon_2, h'''_1 \cdot h'_{11} \cdot h'''_{21}}^{\emptyset, R} w_4$ , which is trivially true
- $g_1 \cdot h_1 \cdot h_{11} \cdot h_{21} = g'''_1 \cdot h'''_1 \cdot h_{11}' \cdot h'''_{21}$ , by congruence closure of equalities
- $(g'''_1, h'''_1 \cdot h'''_{21} \cdot h'_{11}, f_1, g'''_2, h'''_2 \cdot h'''_{22} \cdot h'_{12}, f_2) \in \mathbf{splits}^{\Pi, \Lambda, \emptyset, R} w_4$ , by massaging the assumption above
- $(j-1, g_1, g_2, g'''_1, g'''_2) \in \mathbf{Q}_{\varepsilon_1 \cup \varepsilon_2}^{\Pi, R} w_1, w_4$ , by weakening and composition of postcondition assumptions

This leaves us with the final obligation to prove:  $(j-1, h_1^\circ, h_2^\circ, \langle e_1, e_2 \rangle, \langle \tilde{e}'_1, \tilde{e}'_2 \rangle, h'''_1 \cdot h'_{11} \cdot h'''_{21}, h'''_2 \cdot h'_{12} \cdot h'''_{22}) \in \mathbf{safe}_{\tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2}^{\Pi, \Lambda, \emptyset, R} w^\circ, w_4$ . To do this, we assume:

- **envtran** $^{\Pi, \Lambda, \emptyset, R} w_4, w_5$
- $(i, g_1^\dagger, g_2^\dagger) \in \mathbf{P}_{\varepsilon_1 \cup \varepsilon_2}^{\Pi, R} w_5$
- $(g_1^\dagger, h'''_1 \cdot h'_{11} \cdot h'''_{21}, f_1^\dagger, g_2^\dagger, h'''_2 \cdot h'_{12} \cdot h'''_{22}, f_2^\dagger) \in \mathbf{splits}^{\Pi, \Lambda, \emptyset, R} w_5$

Obviously, both expressions are irreducible, so we take the same values as witnesses. Most obligations are proved by reflexivity or assumption, the ones left are:

- $(i, \langle e_1, e_2 \rangle, \langle \tilde{e}_1', \tilde{e}_2'' \rangle) \in \llbracket \tau_1 \times \tau_2 \rrbracket^R w_5$ , which holds by downwards- and future-world-closure of type denotation, and definition of product types, and
- $(i, h_1^\circ, h_2^\circ, h_1''', h_2''') \in \mathbf{Q}_{\varepsilon_1 \cup \varepsilon_2}^{\Lambda, R} w^\circ, w_5$ , which holds by weakening, composition and downwards closure of previous postcondition assumptions.

This ends the proof. ◀

► **Lemma 18** (Parallelization). *Assuming that*

1.  $\Pi, \Lambda \parallel \Gamma \vdash e_1 : \tau_1, \varepsilon_1$ ,
2.  $\Pi, \Lambda \parallel \Gamma \vdash e_2 : \tau_2, \varepsilon_2$ ,
3.  $\text{rds } \varepsilon_1 \cup \text{wrs } \varepsilon_1 \cup \text{rds } \varepsilon_2 \cup \text{wrs } \varepsilon_2 \subseteq \Lambda$ ,
4.  $\text{rds } \varepsilon_1 \cap \text{wrs } \varepsilon_2 = \text{rds } \varepsilon_2 \cap (\text{wrs } \varepsilon_1 \cup \text{als } \varepsilon_1) = \text{wrs } \varepsilon_1 \cap \text{wrs } \varepsilon_2 = \emptyset$ ,

*the following property holds:*

$$\Pi \mid \Lambda \mid \Gamma \models \langle e_1, e_2 \rangle \cong \text{par } e_1 \text{ and } e_2 : \tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2.$$

**Proof.** The proof consists of two parts, for two directions of contextual approximation. For both these direction we need assumptions about  $e_1$  and  $e_2$ , that follow by the fundamental theorem of logical relations:

$$\Pi, \Lambda \parallel \Gamma \models e_1 \preceq e_1 : \tau_1, \varepsilon_1 \qquad \Pi, \Lambda \parallel \Gamma \models e_2 \preceq e_2 : \tau_2, \varepsilon_2$$

1. To show:  $\Pi \mid \Lambda \mid \Gamma \models \langle e_1, e_2 \rangle \preceq \text{par } e_1 \text{ and } e_2 : \tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2$ .

We start this part, by taking  $a_1$  and  $a_2$ , the witnesses from the logical relations above, and setting  $a = a_1 + a_2$  as the witness for the logical relation. Now we can assume the initial conditions:

- $R(\text{FRV}(\varepsilon_1 \cup \varepsilon_2)) \subseteq \text{dom}(w)$
- $A_1 + A_2 \# R(\Pi \cup \Lambda)$
- $|A_i| \geq a_i$ , for  $i \in \{1, 2\}$
- $\forall r \in A_1 + A_2. w(r) = \emptyset$
- $(k, \gamma_1, \gamma_2) \in \llbracket \Gamma \rrbracket^R w$
- $h_i^\circ \subseteq h_i$ , for  $i \in \{1, 2\}$
- $(k, h_1^\circ, h_2^\circ) \in \mathbf{P}_{\varepsilon_1 \cup \varepsilon_2}^{\Lambda, R} w$ .

What remains to be shown is safety:  $(k, h_1^\circ, h_2^\circ, \langle e_1^1, e_2^1 \rangle, \text{par } e_1^2 \text{ and } e_2^2, h_1, h_2) \in \text{safe}_{\tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2}^{\Pi, \Lambda, A_1 + A_2, R} w, w$ , where  $e_i^j = e_i[\gamma_j/\Gamma]$ . We can also use the initial conditions above to obtain:

$$(k, \emptyset, \emptyset, e_i^1, e_i^2, \emptyset, \emptyset) \in \text{safe}_{\tau_i, \varepsilon_i}^{\Pi \cup \Lambda, \cdot, A_i, R} w, w.$$

We proceed by well-founded induction, keeping the additional assumption that  $(k, h_1 \upharpoonright_{\text{dom}_1^{\text{R}(\Lambda)}(w)}, h_2 \upharpoonright_{\text{dom}_2^{\text{R}(\Lambda)}(w)}) \in \mathbf{P}_{\varepsilon_1 \cup \varepsilon_2}^{\Lambda, R} w$ . After unrolling safety, there are three possibilities. Either  $e_i^1$  are both irreducible, in which case we can recover the final reductions for  $e_i^2$ , replay those, and reduce the parallel composition to the pair of values related to the left-hand-side pair, or at least one of  $e_i^1$  can reduce. In the latter cases, we can use the respective assumption to obtain the matching reduction, perform it, and use the induction hypothesis (with downwards- and envtran-closure ensuring that the other safety assumption matches and the precondition assumption fulfilled due to  $\Lambda$  being public). The details are very similar to the compatibility of parallel composition, and so are omitted.

2. To show:  $\Pi \mid \Lambda \mid \Gamma \models \text{par } e_1 \text{ and } e_2 \preceq \langle e_1, e_2 \rangle : \tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2$ .

We proceed in the similar manner, by taking the witnesses, providing  $a_1 + a_2$  as the witness for the whole computation, and obtaining the safety obligations: show that  $(k, h_1^\circ, h_2^\circ, \text{par } e_1^1 \text{ and } e_2^1, \langle e_1^2, e_2^2 \rangle, h_1, h_2) \in \mathbf{safe}_{\tau_1 \times \tau_2, \varepsilon_1 \cup \varepsilon_2}^{\Pi, \Lambda, A_1 + A_2, R} w, w$  provided that  $(k, \emptyset, \emptyset, e_i^1, e_i^2, \emptyset, \emptyset) \in \mathbf{safe}_{\tau_i, \varepsilon_i}^{\Pi \cup \Lambda, \cdot, A_i, R} w, w$ , where  $e_i^j = e_i[\gamma_j/\Gamma]$ . However, to show it we use the catch-up lemma proved above. To do this, we need to show its three preconditions (we let  $h_1^l = h_1 \upharpoonright_{\text{dom}_1^{R(\Lambda)}(w)}$ ,  $h_{11} = h_1 \setminus h_1^l$ ,  $h_{12} = \emptyset$ , and similar for  $h_2$ ):

- $(k, h_1^l, h_2^l) \in \mathbf{P}_{\varepsilon_1}^{\Lambda, R} w$ , which holds by pre-weakening,
- $(k, \emptyset, \emptyset, e_1^1, e_1^2, h_{11}, h_{12}) \in \mathbf{safe}_{\tau_1, \varepsilon_1}^{\Pi \cup \Lambda, \cdot, A_1, R} w, w$ , which is easy to show since  $h_{1i}$  is just some random heap outside the control of the world (regions in  $A_1$  are empty),
- and the final catching-up precondition. To show this property, take any  $g_1, g_2, h_1^\dagger, h_2^\dagger, f_1, f_2, w', j \leq k$  such that:
  - $\mathbf{envtran}^{\Pi \cup \Lambda, \cdot, A_2, R} w, w'$
  - $(g_1 \cdot h_1^\dagger, \emptyset, f_1, g_2 \cdot h_2^\dagger, \emptyset, f_2) \in \mathbf{splits}^{\Pi \cup \Lambda, \cdot, A_2, R} w'$
  - $(j, h_1, h_2, h_1^\dagger, h_2^\dagger) \in \mathbf{Q}_{\varepsilon_1}^{\Lambda, R} w, w'$

We take the witnesses to be the precise same things, so most of the obligations hold by reflexivity. The ones left are:

- $(j, h_1^\circ, h_2^\circ, h_1^\dagger, h_2^\dagger) \in \mathbf{Q}_{\varepsilon_1 \cup \varepsilon_2}^{\Lambda, R} w, w'$ , which by post-weaken, since  $h_i^\circ = h_i^l$ ,
- $(k, \emptyset, \emptyset, e_2^1, e_2^2, \emptyset, \emptyset) \in \mathbf{safe}_{\tau_2, \varepsilon_2}^{\Pi \cup \Lambda, \cdot, A_2, R} w, w'$ , which holds by envtran-closure of the safety assumption, and
- $(k, h_1^\dagger, h_2^\dagger) \in \mathbf{P}_{\varepsilon_2}^{\Lambda, R} w'$ , which holds because of disjointness condition on  $\varepsilon_1$  and  $\varepsilon_2$  — when cut down to the  $\varepsilon_2$ -read regions,  $h_i^\dagger = h_i^\circ$ , for which we have the right assumption.

This ends the proof. ◀



# A Separation Logic for Fictional Sequential Consistency: Technical Appendix

Filip Sieczkowski  
IT University of Copenhagen  
fisi@itu.dk

Kasper Svendsen  
Aarhus University  
ksvendsen@cs.au.dk

Lars Birkedal  
Aarhus University  
birkedal@cs.au.dk

November 28, 2013

## 1 The Language

We build our logic for a simple, class-based programming language. For simplicity of both semantics and the logic, the language uses let-bindings and expressions, but it stays relatively low-level by ensuring that all the values are machine-word size. The values include variables, natural numbers, booleans, unit, object references (pointers), the null pointer and the special variable **this**. The expressions include values, let bindings, conditionals, constructor and method calls, field reads and writes, atomic compare-and-swap expression and a fork call. The syntax of values and expressions is shown in Figure 1. Note that for simplicity classes can have only one constructor each, and that it has to end with **this**, which ensures that constructors return the newly allocated value. Actual object references  $o$  cannot occur in the text of the program, but need to be values for the operational semantics to be able to reduce.

In Figure 2 we present the semantic domains: the interesting part is the machine state, which in addition to the heap also contains a pool of store buffers, one per thread, and the program state. The latter is only consistent if the thread pool contains the same threads that the store buffer pool contains — or if the program has already gone wrong.

Following the Views framework [1], the operational semantics is split into two components: a thread-local small-step semantics labeled with actions that occur during the step, and action semantics that defines the effect of the action on the machine state. Below, the possible actions are given, and in the Figure 3 we present the thread-local semantics. The interpretation of actions is shown in Figure 6.

$$a \in Act ::= \varepsilon \mid read(t, o, f, v) \mid write(t, o, f, v) \mid cas(t, o, f, v_o, v_n, r) \mid new(t, C, o) \\ \mid type(t, o, T) \mid fork(t, o, T, t') \mid \dagger$$

$\text{ClassDef} ::= \text{class } C = \{\overline{T} \overline{f}; \text{CDef}; \overline{\text{MDef}}\}$	(Class definition)
$\text{MDef} ::= T \ m(\overline{T} \ x) = e$	(Method definition)
$\text{CDef} ::= C(\overline{T} \ x) = (e; \text{this})$	(Constructor definition)
$\text{Val } \ni v ::= x \mid \text{null} \mid \text{this} \mid o \mid n \mid b \mid ()$	(Values)
$\text{Expr } \ni e ::= v \mid \text{let } x = e_1 \text{ in } e_2 \mid \text{if } v \text{ then } e_1 \text{ else } e_2$	(Expressions)
$\quad \mid \text{new } C(\overline{v}) \mid v.f \mid v.f := v \mid v.m(\overline{v}) \mid \text{CAS}(v_1.f, v_2, v_3) \mid \text{fork}(v.m)$	
$\text{E} ::= [] \mid \text{let } x = E \text{ in } e$	(Evaluation Contexts)

Figure 1: The syntax of the programming language.

$t \in TId$	(Thread Identifiers)
$o \in OId$	(Object Identifiers)
$v \in Val \stackrel{\text{def}}{=} \{\text{null}\} + OId + \mathbb{N} + \mathbf{2} + \mathbf{1}$	(Semantic Values)
$OHeap \stackrel{\text{def}}{=} OId \times FName \rightarrow_{fin} Val$	(Object Heaps)
$THeap \stackrel{\text{def}}{=} OId \rightarrow_{fin} CName$	(Type Heaps)
$h \in Heap \stackrel{\text{def}}{=} OHeap \times THeap$	(Heaps)
$s \in SBuffer \stackrel{\text{def}}{=} \text{seq}(OId \times FName \times Val)$	(Store Buffers)
$T \in TPool \stackrel{\text{def}}{=} TId \rightarrow_{fin} Expr$	(Thread Pool)
$U \in SPool \stackrel{\text{def}}{=} TId \rightarrow_{fin} SBuffer$	(Store Buffer Pool)
$\mu \in MSt \stackrel{\text{def}}{=} Heap \times SPool + \{\zeta\}$	(Memory State)
$PSt \stackrel{\text{def}}{=} \{(\mu, T) \in MSt \times TPool \mid$	(Program State)
$\quad \forall h \in Heap, U \in SPool. \mu = (h, U) \Rightarrow \text{dom } U = \text{dom } T$	

Figure 2: Semantic domains

Note that the action semantics uses two helper functions: lookup and flush. These are used to determine the result of reading a field from a thread's perspective, and to update the state of the heap with a contents of the buffer. These functions are defined in Figure 5.

$$\begin{array}{c}
\frac{}{(t, E[\mathbf{let} \ x = v \ \mathbf{in} \ e]) \xrightarrow{\varepsilon} \{(t, E[e[v/x]])\}} \text{LET} \\
\frac{}{(t, E[\mathbf{if} \ \mathbf{true} \ \mathbf{then} \ e_1 \ \mathbf{else} \ e_2]) \xrightarrow{\varepsilon} \{(t, E[e_1])\}} \text{IF-TRUE} \\
\frac{}{(t, E[\mathbf{if} \ \mathbf{false} \ \mathbf{then} \ e_1 \ \mathbf{else} \ e_2]) \xrightarrow{\varepsilon} \{(t, E[e_2])\}} \text{IF-FALSE} \\
\frac{}{(t, E[\mathbf{CAS}(o.f, v_o, v_n)]) \xrightarrow{\text{cas}(t,o,f,v_o,v_n,r)} \{(t, E[r])\}} \text{CAS} \\
\frac{\text{ctorBody}(C) = (C(\overline{\mathbf{T}} \ x) = e)}{(t, E[\mathbf{new} \ C(\bar{v})]) \xrightarrow{\text{new}(t,C,o)} \{(t, E[e[o, \bar{v}/\mathbf{this}, \bar{x}])\}} \text{NEW} \\
\frac{}{(t, E[o.f]) \xrightarrow{\text{read}(t,o,f,v)} \{(t, E[v])\}} \text{READ} \\
\frac{}{(t, E[o.f := v]) \xrightarrow{\text{write}(t,o,f,v)} \{(t, E[()])\}} \text{WRITE} \\
\frac{\text{body}(C, m) = (\mathbf{T} \ m(\overline{\mathbf{T}} \ x) = e)}{(t, E[o.m(\bar{v})]) \xrightarrow{\text{type}(t,o,C)} \{(t, E[e[o, \bar{v}/\mathbf{this}, \bar{x}])\}} \text{CALL} \\
\frac{\text{body}(C, m) = (\mathbf{unit} \ m() = e) \quad t \neq t'}{(t, E[\mathbf{fork}(o.m)]) \xrightarrow{\text{fork}(t,o,C,t')} \{(t, E[()]), (t', e[o/\mathbf{this}])\}} \text{FORK} \\
\frac{}{(t, e) \xrightarrow{\text{flush}(t)} \{(t, e)\}} \text{FLUSH}
\end{array}$$

Figure 3: Thread-local semantics — the non-fault cases.

Finally, the complete semantics proceeds by reducing one of the threads using the thread-local semantics, then interpreting the resulting action with the action semantics, and reducing to a memory state in the resultant set, as in Figure 4. Note how in some cases, notably read, this might require “guessing” the return value, and checking that the guess was right using the action semantics.

$$\frac{t \in \text{dom } T \quad (t, T(t)) \xrightarrow{a} T' \quad \mu' \in \llbracket a \rrbracket(\mu)}{(\mu, T) \rightarrow (\mu', (T - t) + T')} \text{STEP}$$

Figure 4: Single step program evaluation

$$\begin{aligned}
& \text{lookup}(-, -, -, -) : \text{Old} \times \text{FName} \times \text{SBuffer} \times \text{Heap} \rightarrow \text{Val} \\
& \text{lookup}(o, f, \alpha, h) \stackrel{\text{def}}{=} \begin{cases} \perp & \text{when } (o, f) \notin \text{dom}(h) \\ v & \text{when } h(o, f) = v \wedge \forall v' \in \text{Val}. (o, f, v') \notin \alpha \\ & \text{or } \exists \beta, \gamma \in \text{SBuffer}. \alpha = \beta \cdot (o, f, v) \cdot \gamma \wedge \\ & \forall v' \in \text{Val}. (o, f, v') \notin \gamma \end{cases} \\
& \text{flush}(-, -) : \text{Heap} \times \text{SBuffer} \rightarrow \text{Heap} \\
& \text{flush}(h, \varepsilon) \stackrel{\text{def}}{=} h \\
& \text{flush}(h, (o, f, v) \cdot \alpha) \stackrel{\text{def}}{=} \begin{cases} \text{flush}(h[(o, f) \mapsto v], \alpha) & \text{when } (o, f) \in \text{dom}(h) \\ \text{flush}(h, \alpha) & \text{otherwise} \end{cases}
\end{aligned}$$

Figure 5: Generic mathematical functions used throughout the report.

Finally, we show one of the simple lemmas about operational semantics, that we need in the process of building the logic. The decomposition lemma states that the only terminal program state that is not a fault is the one, in which all the store buffers are flushed, and all the threads have computed their respective values.

**Lemma 1** (Decompose). *For any consistent, non-error program state  $((h, U), T)$  one of the following holds:*

- *For any thread  $t \in \text{dom}(T)$  there exists a value  $v \in \text{Val}$  such that  $T(t) = v$  and  $U(t) = \varepsilon$*
- *There exists a (possibly error) state  $(\mu', T')$  such that  $(\mu, T) \rightarrow (\mu', T')$ .*

## 2 Model

The construction in this section follows closely the one presented in iCAP. As in that case, the non-recursive parts of the model are defined in the category of Sets, **Sets**, while the recursive ones — in the topos of trees,  $\mathcal{S}$ .

### 2.1 Local States, Shared States and the Treatment of Store Buffers

Local states consist of the usual, physical heap and the state transition capability map. Shared states consist of an abstract state and a labelled transition system for each shared region. For the buffered updates we use the physical store buffers defined in Section 1.



$$\begin{aligned}
\llbracket \varepsilon \rrbracket(h, U) &= \{(h, U)\} \\
\llbracket \zeta \rrbracket(h, U) &= \{\zeta\} \\
\llbracket \text{read}(t, o, f, v) \rrbracket(h, U) &= \begin{cases} \{(h, U)\} & \text{if } (o, f) \in \text{dom } h \text{ and } \text{lookup}(o, f, U(t), h) = v \\ \emptyset & \text{if } t \notin \text{dom } U \text{ or } (o, f) \in \text{dom } h \text{ and } \text{lookup}(o, f, U(t), h) \neq v \\ \{\zeta\} & \text{if } (o, f) \notin \text{dom } h \end{cases} \\
\llbracket \text{write}(t, o, f, v) \rrbracket(h, U) &= \begin{cases} \{(h, U[t \mapsto U(t) \cdot (o, f, v)])\} & \text{if } (o, f) \in \text{dom } h \text{ and } t \in \text{dom } U \\ \{\zeta\} & \text{if } (o, f) \notin \text{dom } h \\ \emptyset & \text{if } t \notin \text{dom } U \end{cases} \\
\llbracket \text{new}(t, T, o) \rrbracket(h, U) &= \begin{cases} \{(h[(o, \bar{f}) \mapsto \mathbf{null}, o \mapsto T], U)\} & \text{if } o \notin \text{dom } h \text{ and } \text{fields}(T) = \bar{f} \\ \emptyset & \text{if } o \in \text{dom } h \\ \{\zeta\} & \text{if } \text{fields}(T) \text{ undefined} \end{cases} \\
\llbracket \text{cas}(t, o, f, v_o, v_n, r) \rrbracket(h, U) &= \begin{cases} \{(\text{flush}(h, U(t) \cdot (o, f, v_n)), U[t \mapsto \varepsilon])\} & \text{if } (o, f) \in \text{dom } h, r = \mathbf{true} \\ & \text{and } \text{lookup}(o, f, U(t), h) = v_o \\ \{(h, U)\} & \text{if } (o, f) \in \text{dom } h, r = \mathbf{false} \\ & \text{and } \text{lookup}(o, f, U(t), h) \neq v_o \\ \{\zeta\} & \text{if } (o, f) \notin \text{dom } h \\ \emptyset & \text{otherwise} \end{cases} \\
\llbracket \text{type}(t, o, C) \rrbracket(h, U) &= \begin{cases} \{(h, U)\} & \text{if } o \in \text{dom}_t h \text{ and } h_t(o) = C \\ \emptyset & \text{if } o \in \text{dom}_t h \text{ and } h_t(o) \neq C \\ \{\zeta\} & \text{if } o \notin \text{dom}_t h \end{cases} \\
\llbracket \text{fork}(t, o, T, t') \rrbracket(h, U) &= \begin{cases} \{(h, U[t' \mapsto \varepsilon])\} & \text{if } o \in \text{dom}_t h, h_t(o) = C \text{ and } t' \notin \text{dom } U \\ \emptyset & \text{if } o \in \text{dom}_t h \text{ and } t' \in \text{dom } U \text{ or } h_t(o) \neq C \\ \{\zeta\} & \text{if } o \notin \text{dom}_t h \end{cases} \\
\llbracket \text{flush}(t) \rrbracket(h, U) &= \begin{cases} \{(h[(o, f) \mapsto v], U[t \mapsto \alpha])\} & \text{if } U(t) = (o, f, v) \cdot \alpha \text{ and } (o, f) \in \text{dom } h \\ \emptyset & \text{if } t \notin \text{dom}(U), U(t) = \varepsilon, \text{ or } (o, f) \notin \text{dom } h \end{cases} \\
\llbracket a \rrbracket(\zeta) &= \{\zeta\}
\end{aligned}$$

Figure 6: Action semantics

$$\begin{aligned}
Perm &\stackrel{\text{def}}{=} \{q \in \mathbb{Q} \mid 0 \leq q \leq 1\} && \text{(Permissions)} \\
Perm^+ &\stackrel{\text{def}}{=} \{q \in Perm \mid 0 < q\} && \text{(Non-zero permissions)} \\
Cap &\stackrel{\text{def}}{=} \{f \in (RId \times AId) \rightarrow Perm \mid \exists R \subseteq_{fin} RId. \\
&\quad \forall r \in RId \setminus R. \forall \alpha \in AId. f(r, \alpha) = 0\} && \text{(Capabilities)} \\
l \in LState &\stackrel{\text{def}}{=} Heap \times Cap && \text{(Local States)} \\
LTS &\stackrel{\text{def}}{=} AId \rightarrow \mathcal{P}(SId \times SId) && \text{(Transition Systems)} \\
s \in SState &\stackrel{\text{def}}{=} RId \multimap_{fin} (SId \times LTS) && \text{(Shared States)} \\
AState &\stackrel{\text{def}}{=} LState \times SState \times SPool
\end{aligned}$$

where  $PVal$  denotes the least set such that

$$PVal \cong Val + Strings + (PVal \times PVal).$$

We write  $l.h$  and  $l.c$  to refer to the appropriate components of the local state  $l$ , while  $s(r).s$  and  $s(r).p$  to refer to the state identifier and transition system components of region  $r$  in a shared state  $s$ .

### Local State Composition

$$\begin{aligned}
\bullet_{LState} = \{(l_1, l_2, l_r) \mid &\text{dom}(l_1.h) \cap \text{dom}(l_2.h) = \emptyset \wedge (\forall t \in \Delta(RId \times AId). l_1.c(t) + l_2.c(t) \leq 1) \\
&\wedge l_r.h = l_1.h \cup l_2.h \wedge l_r.c = l_1.c + l_2.c\},
\end{aligned}$$

where  $+$  on finite maps denotes pointwise addition.

The relation is functional, so it induces a partial function

$$\bullet_{LState} : \Delta(LState) \times \Delta(LState) \multimap \Delta(LState) \in \mathcal{S}$$

**Lemma 2.**  $\forall l_1, l_2, l_3 \in \Delta(LState). \triangleright (l_1 = l_2 \bullet_{LState} l_3) \Rightarrow (l_1 = l_2 \bullet_{LState} l_3 \vee \triangleright \perp)$

*Proof.* Using the theory of upwards-closure from the iCAP technical report [2] □

### Action Allowed

$$\begin{aligned}
\text{act}_o &: LState \times RId \rightarrow \mathcal{P}(AId) \in \mathbf{Sets} \\
\text{act}_o(l, r) &\stackrel{\text{def}}{=} \{\alpha \in AId \mid l.c(r, \alpha) < 1\}
\end{aligned}$$

### Update Allowed

$$\begin{aligned}
\text{upd}_o &: LState \times RId \times LTS \rightarrow \mathcal{P}(SId \times SId) \in \mathbf{Sets} \\
\text{upd}_o(l, r, p) &\stackrel{\text{def}}{=} \{(s_1, s_2) \in SId \times SId \mid \exists \alpha \in \text{act}_o(l, r). (s_1, s_2) \in p(\alpha)\}
\end{aligned}$$

## 2.2 Interference Relation

In our setup, interference is more complicated than in plain iCAP, due to the existence of store buffered updates. While in iCAP the interference concerned only the shared regions and transitions they are allowed to make, we also have to consider the three kinds of actions the environment can take with respect to the buffers: the allocation of a new store buffer, a buffered update written into an existing one, or a flushing of an update from store buffer to the heap. The interference is thus defined as follows.

$$\begin{aligned}
(l, U_1) R_n (l, U_2) &\text{ iff } \text{dom}(U_1) \subseteq \text{dom}(U_2) \wedge \forall t \in \text{dom}(U_2) \setminus \text{dom}(U_1). U_2(t) = \varepsilon \wedge \\
&\quad \forall t \in \text{dom}(U_1). U_2(t) = U_1(t) \\
(l, U_1) R_w (l, U_2) &\text{ iff } \text{dom}(U_1) = \text{dom}(U_2) \wedge \\
&\quad \exists t \in \text{dom}(U_2), o \in OId, f \in FName, v \in Val. (o, f) \notin \text{dom}(l.h) \wedge \\
&\quad U_2(t) = U_1(t) \cdot (o, f, v) \wedge \forall t' \in \text{dom}(U_2). t \neq t' \Rightarrow U_1(t) = U_2(t) \\
(l_1, U_1) R_f (l_2, U_2) &\text{ iff } \text{dom}(U_1) = \text{dom}(U_2) \wedge \\
&\quad \exists t \in \text{dom}(U_1), o \in OId, f \in FName, v \in Val. U_1(t) = (o, f, v) \cdot U_2(t) \\
&\quad l_2 = \text{flush}(l_1, (o, f, v)) \wedge \forall t' \in \text{dom}(U_2). t \neq t' \Rightarrow U_1(t) = U_2(t) \\
R_{sb} &\stackrel{\text{def}}{=} (R_n \cup R_w \cup R_f)^* \\
(l_1, s_1) R_r^A (l_2, s_2) &\text{ iff } U_1 = U_2 \wedge l_1 \leq l_2 \wedge \text{dom}(s_1) \subseteq \text{dom}(s_2) \wedge \\
&\quad \forall r \in \text{dom}(s_1). (s_1(r).s, s_2(r).s) \in \left( \bigcup_{\alpha \in A, l_1.c(r, \alpha) < 1} S(r).t(\alpha) \right)^* \\
(l_1, s_1, U_1) R'_A (l_2, s_2, U_2) &\text{ iff } ((l_1, s_1) R_r^A (l_2, s_2) \wedge U_1 = U_2) \vee \\
&\quad ((l_1, U_1) R_{sb} (l_2, U_2) \wedge (s_1 \leq s_2)) \\
R_A &\stackrel{\text{def}}{=} (R'_A)^*
\end{aligned}$$

In the above  $R^*$  denotes the reflexive, transitive closure of  $R$ . We use  $R$  as a shorthand for  $R_{AId}$ .

## 2.3 Region Interpretations

We follow the iCAP treatment to give the concrete interpretation to the abstract shared states in the shared regions — and thus define the type of region interpretations as a guarded recursive type in the topos of trees,  $\mathcal{S}$ . We take  $RIntr$  to be an object in  $\mathcal{S}$  that satisfies the isomorphism

$$i : RIntr \cong \blacktriangleright((\Delta(SId) \times (\Delta(RId) \rightarrow_{fin} RIntr)) \rightarrow_{mon} \mathcal{P}^\uparrow(\Delta(AState))),$$

where  $\Delta(AState)$  is upwards-closed with respect to the interference relation  $R$ , and the ordering on action models,  $AMod \stackrel{\text{def}}{=} \Delta(RId) \rightarrow_{fin} RIntr$ , is given as

$$\zeta_1 \leq \zeta_2 \stackrel{\text{def}}{=} \text{dom}(\zeta_1) \subseteq \text{dom}(\zeta_2) \wedge \forall r \in \text{dom}(\zeta_1). \zeta_1(r) = \zeta_2(r).$$

The ordering on  $\Delta(SId)$  is discrete.

Now we can define the maps between  $\Delta(SId) \times AMod \rightarrow_{mon} \mathcal{P}^\uparrow(\Delta(AState))$  and  $RIntr$  as follows:

$$\begin{aligned}
lam &: (\Delta(SId) \times AMod \rightarrow_{mon} \mathcal{P}^\uparrow(\Delta(AState))) \rightarrow RIntr \\
lam &\stackrel{\text{def}}{=} i^{-1} \circ next \\
app &: RIntr \rightarrow (\Delta(SId) \times AMod \rightarrow_{mon} \mathcal{P}^\uparrow(\Delta(AState))) \\
app &\stackrel{\text{def}}{=} \lambda x : RIntr. \lambda(s, \zeta) : \Delta(SId) \times (AMod). \lambda a : \Delta(AState). \\
&\quad succ(J(J(i(x))(next(s), next(\zeta)))(next(a)))
\end{aligned}$$

**Lemma 3.**  $app \circ lam = \triangleright$

## 2.4 Instrumented States

Instrumented states consist of a concrete local state and store buffer pool, an abstract state transition system for each shared region, and concrete interpretation of each abstract shared state:

$$\mathcal{M} \stackrel{\text{def}}{=} \Delta(LState) \times \Delta(SState) \times \Delta(SPool) \times AMod.$$

We use  $m.l$ ,  $m.s$ ,  $m.U$  and  $m.a$  to refer to the appropriate components of the instrumented state  $m : \mathcal{M}$ .

We also define a subset of  $\mathcal{M}$  that omits the store buffers, which we will sometimes use:

$$\mathcal{H} \stackrel{\text{def}}{=} \Delta(LState) \times \Delta(SState) \times AMod.$$

We use analogous projections whenever needed.

**Propositions** We define two spaces of propositions,  $Prop_{TSO}$ , which denotes the general assertions, and  $Prop_{SC}$ , for assertions that do not mention the store buffers.

$$\begin{aligned}
Prop_{TSO} &\stackrel{\text{def}}{=} \mathcal{P}^\uparrow(\mathcal{M}) \\
Prop_{SC} &\stackrel{\text{def}}{=} \mathcal{P}^\uparrow(\mathcal{H}),
\end{aligned}$$

where the orderings are pointwise extension orderings.

We define an interpretation of  $Prop_{SC}$  in  $Prop_{TSO}$  as the propositions that have a “locally flushed” subset:

$$\begin{aligned}
\text{lfid} &\subseteq LState \times SPool \in \mathbf{Sets} \\
\text{lfid}(l, U) &\text{ iff } \forall t \in TId. \forall (o, f) \in \text{dom}(l). \forall v \in Val. (o, f, v) \notin U(t) \\
\lceil - \rceil &: Prop_{SC} \rightarrow Prop_{TSO} \in \mathcal{S} \\
\lceil p \rceil &\stackrel{\text{def}}{=} \{(l, s, U, \zeta) \in \mathcal{M} \mid \exists l' \leq l. (l', s, \zeta) \in p \wedge \text{lfid}(l', U)\}
\end{aligned}$$

Interpreted in this way  $Prop_{SC}$  defines a subset of well-behaved propositions that do not depend on the state of the buffers.

Erasure

$$\begin{aligned} \llbracket (s, \zeta) \rrbracket_r &\stackrel{\text{def}}{=} \{(l, U) \in LState \times SPool \mid (l, s, U) \in \text{app}(\zeta(r))(s(r).s, \zeta)\} \\ \llbracket (l, s, U, \zeta) \rrbracket_A &\stackrel{\text{def}}{=} \{(h, U) \in Heap \times SPool \mid \\ &\quad \exists l' \in LState, sr \in \text{dom}(s) \cap A \rightarrow LState. \\ &\quad h = l'.h \wedge l' = l \bullet \prod_{r \in \text{dom}(s) \cap A} sr(r) \wedge \\ &\quad \forall r \in \text{dom}(s) \cap A. (sr(r), U) \in \llbracket (s, \zeta) \rrbracket_r\} \end{aligned}$$

As usual, we use  $\llbracket m \rrbracket$  as a shorthand for  $\llbracket m \rrbracket_{RId}$ .

Composition

$$\bullet \mathcal{M} \stackrel{\text{def}}{=} \bullet LState \times \bullet = \times \bullet = \times \bullet =$$

## 2.5 Logical Connectives

Given below are the definitions for connectives in  $Prop_{TSO}$ . The definitions for  $Prop_{SC}$  are analogous.

$$\begin{aligned} emp &\stackrel{\text{def}}{=} \mathcal{M} \\ p * q &\stackrel{\text{def}}{=} \{m \in \mathcal{M} \mid \exists m_1, m_2 \in \mathcal{M}. m = m_1 \bullet m_2 \wedge m_1 \in p \wedge m_2 \in q\} \\ p \Rightarrow q &\stackrel{\text{def}}{=} \{m \in \mathcal{M} \mid \forall m' \geq m. m' \in p \Rightarrow m' \in q\} \\ \exists_\tau(p) &\stackrel{\text{def}}{=} \{m \in \mathcal{M} \mid \exists x : \tau. m \in p(x)\} \\ \forall_\tau(p) &\stackrel{\text{def}}{=} \{m \in \mathcal{M} \mid \forall x : \tau. m \in p(x)\} \end{aligned}$$

**Lemma 4.** *Let  $X$  be a total and inhabited object in  $\mathcal{S}$ . Then*

$$\forall p \in X \rightarrow Prop_{TSO}. \triangleright \exists_X(p) \iff \exists_X(\lambda x. \triangleright p(x))$$

*Proof.* Let  $m \in \mathcal{M}$  such that  $m \in \triangleright \exists_X(p)$ . Then  $\triangleright(m \in \exists_X(p))$  and thus

$$\triangleright(\exists x : X. m \in p(x)).$$

By totality and inhabitation,  $\exists x : X. \triangleright(m \in p(x))$  and thus  $m \in \exists_X(\lambda x. \triangleright p(x))$ .

Likewise, let  $m \in \mathcal{M}$  such that  $m \in \exists_X(\lambda x. \triangleright p(x))$  then  $\exists x : X. \triangleright(m \in p(x))$  and thus  $\triangleright(\exists x : X. m \in p(x))$ .  $\square$

**Properties of the Embedding** In order to use the  $\ulcorner \urcorner$  embedding to define the two separation logics, we need it to satisfy certain properties. These are formalized in the next few lemmas.

**Lemma 5.** *Embedding preserves the entailment relation, i.e., for any  $p, q \in Prop_{SC}$ ,*

$$(p \subseteq q) \iff (\ulcorner p \urcorner \subseteq \ulcorner q \urcorner).$$

*Proof* ( $\Rightarrow$ ). Assume  $p \subseteq q$ , and take any  $(l, s, U, \zeta) \in \ulcorner p \urcorner$ . By definition, this gives us  $l' \leq l$ , such that  $(l', s, \zeta) \in p$  and  $\text{lfid}(l', U)$ . By assumption, we have  $(l', s, \zeta) \in q$ , so we can take  $l'$  as a witness, and the properties hold trivially.  $\square$

*Proof* ( $\Leftarrow$ ). Assume  $\ulcorner p \urcorner \subseteq \ulcorner q \urcorner$ , and take  $(l, s, \zeta) \in p$ . This means  $(l, s, \emptyset, \zeta) \in \ulcorner p \urcorner$ , since  $\text{lfid}(l, \emptyset)$  holds trivially, and so we get  $(l, s, \emptyset, \zeta) \in \ulcorner q \urcorner$ . From this, we obtain  $l' \leq l$  such that  $(l', s, \zeta) \in q$ , but since  $q$  is upwards closed,  $(l, s, \zeta) \in q$  holds.  $\square$

**Lemma 6.** *For any  $p \in \text{Prop}_{TSO}$  and any  $m \in p$  there exists a local state  $l' \leq m.l$  such that  $\text{lfid}(l', m.U)$  and  $l'$  is the maximal state with this property (wrt the extension ordering).*

*Proof.* Take any  $p \in \text{Prop}_{TSO}$  and any  $(l, s, U, \zeta) \in p$ . Let  $s$  be the set of all the locations in the store-buffer pool  $U$ : formally,

$$(o, f) \in s \text{ iff } \exists t \in \text{dom}(U). \exists v \in \text{Val}. (o, f, v) \in U(t).$$

Take the state  $l' = l \upharpoonright_{\text{dom}(l) \setminus s}$  as the witness. Both  $l' \leq l$  and  $\text{lfid}(l', U)$  clearly hold, so it suffices to show that this is the biggest state with this property. To this end, take any  $l'' \leq l'$  such that  $\text{lfid}(l'', U)$ . We need to show that  $l'' \leq l'$ . To this end, take any  $(o, f) \in \text{dom}(l'')$ . Since both local states are smaller than  $l$  under extension ordering, it suffices to show that  $(o, f) \in \text{dom}(l')$  – the values  $(o, f)$  gets mapped to will have to agree. Clearly, we have  $(o, f) \in \text{dom}(l)$ , since  $l'' \leq l$ , so it suffices to show that  $(o, f) \notin s$ . To this end, assume  $(o, f) \in s$ . This gives us  $t \in \text{dom}(U)$  and  $v \in \text{Val}$  such that  $(o, f, v) \in U(t)$ . However, since  $\text{lfid}(l'', U)$ , we know that  $(o, f, v) \notin U(t)$ , which ends the proof.  $\square$

**Lemma 7.** *Embedding preserves limits, i.e., for any  $\tau$  and  $p : \tau \rightarrow \text{Prop}_{SC}$ ,*

$$\forall_\tau(\ulcorner p \urcorner) = \ulcorner \forall_\tau(p) \urcorner.$$

*Proof* ( $\subseteq$ ). Take any  $(l, s, U, \zeta) \in \forall_\tau(\ulcorner p \urcorner)$ . By definition of universal quantifier, this means that  $\forall x : \tau. (l, s, U, \zeta) \in \ulcorner p(x) \urcorner$ . By Lemma 6, we get  $l'$ , the greatest local state such that  $l' \leq m.l$  and  $\text{lfid}(l', U)$ . Now it suffices to show that  $(l', s, \zeta) \in \forall_\tau(p)$ . Take any  $x : \tau$ . From the assumption, we get  $l'' \leq l$  such that  $\text{lfid}(l'', U)$  and  $(l'', s, \zeta) \in p(x)$ . However by the universal property of  $l'$ , we know that  $l'' \leq l'$ , and so, by upwards closure  $(l', s, \zeta) \in p$ .  $\square$

*Proof* ( $\supseteq$ ). Take any  $(l, s, U, \zeta) \in \ulcorner \forall_\tau(p) \urcorner$ . By definition, this gives us  $l' \leq l$  such that  $\text{lfid}(l', U)$  and  $\forall x : \tau. (l', s, \zeta) \in p(x)$ . Taking any  $x : \tau$ , and picking  $l'$  as the witness trivially finishes the proof.  $\square$

**Lemma 8.** *Embedding preserves colimits, i.e., for any  $\tau$  and  $p : \tau \rightarrow \text{Prop}_{SC}$ ,*

$$\exists_\tau(\ulcorner p \urcorner) = \ulcorner \exists_\tau(p) \urcorner.$$

*Proof* ( $\subseteq$ ). Take any  $(l, s, U, \zeta) \in \exists_\tau(\ulcorner p \urcorner)$ . Unrolling the definitions, this gives us  $x : \tau$  and  $l' \leq l$  such that  $\text{lfid}(l', U)$  and  $(l', s, \zeta) \in p(x)$ . Taking  $l'$  and  $x$  as witnesses provides all the properties we need to finish the proof.  $\square$

*Proof* ( $\supseteq$ ). Take any  $(l, s, U, \zeta) \in \lceil \exists_\tau(p) \rceil$ . Unrolling the definitions gives us  $l' \leq l$  such that  $\text{ld}(l', U)$  and  $x : \tau$  such that  $(l', s, \zeta) \in p(x)$ . Taking  $x$  and  $l'$  as witnesses provides all the needed properties.  $\square$

**Lemma 9.** *Embedding preserves separating conjunctions, i.e., for any  $p, q \in Prop_{SC}$ ,*

$$\lceil p \rceil * \lceil q \rceil = \lceil p * q \rceil.$$

*Proof* ( $\subseteq$ ). Take any  $(l, s, U, \zeta) \in \lceil p \rceil * \lceil q \rceil$ . This means there exist  $l_1, l_2 \in \Delta(LState)$  such that  $l_1 \bullet l_2 = l$ ,  $(l_1, s, U, \zeta) \in \lceil p \rceil$  and  $(l_2, s, U, \zeta) \in \lceil q \rceil$ . By definition of the embedding, this gets us  $l'_1 \leq l_1$  and  $l'_2 \leq l_2$  such that  $(l'_1, s, \zeta) \in p$ ,  $(l'_2, s, \zeta) \in q$ ,  $\text{ld}(l'_1, U)$  and  $\text{ld}(l'_2, U)$ . Pick  $l'_1 \bullet l'_2 \leq l$  as a witness. Obviously,  $\text{ld}(l'_1 \bullet l'_2, U)$ , and we can split the state as  $l'_1$  and  $l'_2$  to show that  $(l'_1, s, \zeta) \in p$  and  $(l'_2, s, \zeta) \in q$ .  $\square$

*Proof* ( $\supseteq$ ). Take any  $(l, s, U, \zeta) \in \lceil p * q \rceil$ . This means that there exists  $l' \leq l$ ,  $l_1$  and  $l_2$  such that  $\text{ld}(l', U)$ ,  $l' = l_1 \bullet l_2$ ,  $(l_1, s, \zeta) \in p$  and  $(l_2, s, \zeta) \in q$ . Pick as witnesses  $l_1$  and  $l'' = l \upharpoonright_{\text{dom}(l) \setminus \text{dom}(l_1)}$ . These states compose to give  $l$ , so we need to show  $(l_1, s, U, \zeta) \in \lceil p \rceil$  and  $(l'', s, U, \zeta) \in \lceil q \rceil$ . For the first of these, pick  $l_1$  as the witness. Since  $l_1 \leq l'$ , we have  $\text{ld}(l_1, U)$ , and we have  $(l_1, s, \zeta) \in p$  by assumption. For the second goal, pick  $l_2$  as the witness. Since  $l_2$  is compatible with  $l_1$ , we get  $l_2 \leq l''$ , and since  $l_2 \leq l' \text{ — } \text{ld}(l_2, U)$ . By assumption we also have  $(l_2, s, \zeta) \in q$ .  $\square$

Validity

$$\text{valid}(p) \stackrel{\text{def}}{=} \forall m \in \mathcal{M}. m \in p$$

Region Assertion and Interpretation

$$\begin{aligned} \text{region} &: \mathcal{P}(\Delta(SId)) \times \Delta(LTS) \times \Delta(RId) \rightarrow Prop_{SC} \in \mathcal{S} \\ \text{region}(X, p, r) &\stackrel{\text{def}}{=} \{(l, s, \zeta) \in \mathcal{H} \mid s(r).s \in X \wedge s(r).p = p\} \\ \text{rintr} &: (\Delta(SId) \rightarrow Prop_{TSO}) \times \Delta(RId) \rightarrow Prop_{SC} \in \mathcal{S} \\ \text{rintr}(I, r) &\stackrel{\text{def}}{=} \{(l, s, \zeta) \in \mathcal{H} \mid r \in \text{dom}(\zeta) \wedge \forall x \in \Delta(SId), \zeta' \geq \zeta. \\ &\quad \text{app}(\zeta(r))(x, \zeta') = \triangleright(\lambda(l, s, U).I(x)(l, s, U, \zeta'))\} \end{aligned}$$

Note that although technically these assertions are in  $Prop_{SC}$ , we will often omit the  $\lceil - \rceil$  interpretation when using them, since they do not depend at all on the local state, so one can always pick the empty local state as the appropriate witness.

Action Permission

$$[\alpha]_\pi^r \stackrel{\text{def}}{=} \{(l, s, \zeta) \in \mathcal{H} \mid \pi \leq l.c(r, \alpha)\}$$

## 2.6 Stability

$$\begin{aligned} \text{stable}(p) &\stackrel{\text{def}}{=} R_A(p) \subseteq p \\ \text{bstable}(p) &\stackrel{\text{def}}{=} R_{sb}(p) \subseteq p \\ \text{rstable}(p) &\stackrel{\text{def}}{=} R_r^A(p) \subseteq p \end{aligned}$$

As usual, we use  $\text{stable}$  as a shorthand for  $\text{stable}_{RId}$

**Lemma 10.**  $\forall p \in Prop_{TSO}. \text{stable}(p) \iff \text{bstable}(p) \wedge \text{rstable}(p)$

**Lemma 11.**  $R(\text{emp}) \subseteq \text{emp} \wedge \forall p, q \in Prop_{TSO}. R(p * q) \subseteq R(p) * R(q)$

**Lemma 12.**  $\forall A \in \mathcal{P}(\Delta(RId)), p \in Prop_{TSO}. \text{stable}_A(p) \Rightarrow \text{stable}_A(\triangleright p)$

*Proof.* Assume  $m_1 \in \triangleright p$  and  $m_1 R_A m_2$ . By assumption,  $R_A(p) \subseteq p$ , and thus, by monotonicity of  $\triangleright$ ,

$$\triangleright(\forall m_1, m_2 \in \mathcal{M}. m_1 \in p \wedge m_1 R_A m_2 \Rightarrow m_2 \in p).$$

Thus, since  $\mathcal{M}$  is total and inhabited,

$$\forall m_1, m_2 \in \mathcal{M}. m_1 \in \triangleright p \wedge \triangleright(m_1 R_A m_2) \Rightarrow m_2 \in \triangleright p,$$

and so  $m_2 \in \triangleright p$ . □

## 2.7 View Shift

$$p \sqsubseteq_A q \stackrel{\text{def}}{=} \forall r \in Prop_{TSO}. \text{stable}(r) \Rightarrow [p * r]_A \subseteq [q * r]_A$$

Again, we write  $\sqsubseteq$  as a shorthand for  $\sqsubseteq_{RId}$ .

**Lemma 13.**  $\forall p, q, r \in Prop_{TSO}. \text{stable}(r) \wedge p \sqsubseteq_A q \Rightarrow p * r \sqsubseteq_A q * r$

**Lemma 14.**  $\forall p_1, p_2, q_1, q_2 \in Prop_{TSO}. p_1 \sqsubseteq q_1 \wedge p_2 \sqsubseteq q_2 \Rightarrow p_1 * p_2 \sqsubseteq q_1 * q_2$

## 2.8 Atomic Satisfaction

$a \text{ sat}_A \{p\} \{q\} \stackrel{\text{def}}{=} \forall r \in Prop_{TSO}, m \in \mathcal{M}, h \in \text{Heap}, U \in \text{SPool}.$

$$m \in p * \triangleright r \wedge (h, U) \in [m]_A \wedge \text{stable}(r) \Rightarrow$$

$$(\triangleright \not\vdash \llbracket a \rrbracket(h, U)) \wedge$$

$$\forall (h', U') \in \llbracket a \rrbracket(h, U). \exists m' \in \mathcal{M}. \triangleright (m' \in q * r \wedge (h', U') \in [m']_A)$$

We write  $a \text{ sat} \{p\} \{q\}$  for  $a \text{ sat}_{RId} \{p\} \{q\}$ , as usual.



**Lemma 15.**

$$\forall A \in \mathcal{P}(\Delta(RId)). \forall a \in \Delta(Act). \forall p, q \in Prop_{TSO}. \triangleright \perp \Rightarrow a \text{ sat}_A \{p\} \{q\}$$

*Proof.* Follows easily as  $\triangleright \perp$  implies  $\triangleright \not\vdash \notin \llbracket a \rrbracket(h, U)$  and  $\triangleright(m' \in q * r \wedge (h', U') \in \llbracket m' \rrbracket_A)$  for any  $m', h'$  and  $U'$ , and  $\mathcal{M}$  is inhabited.  $\square$

**Lemma 16** (Atomic-Shift).

$$\begin{aligned} &\forall A \in \mathcal{P}(\Delta(RId)), p_1, p_2, q_1, q_2 \in Prop_{TSO}. \\ &p_1 \sqsubseteq_A p_2 \wedge a \text{ sat}_A \{p_2\} \{q_2\} \wedge \triangleright(q_2 \sqsubseteq_A q_1) \Rightarrow a \text{ sat}_A \{p_1\} \{q_1\} \end{aligned}$$

*Proof.* Assume  $r \in Prop_{TSO}$ ,  $m \in \mathcal{M}$ ,  $h, h' \in Heap$ , and  $U, U' \in SPool$  such that

$$\text{stable}_A(r) \quad m \in p_1 * \triangleright r \quad (h, U) \in \llbracket m \rrbracket_A \quad (h', U') \in \llbracket a \rrbracket(h, U).$$

By Lemma 12 it follows that  $\text{stable}_A(\triangleright r)$  and thus  $(h, U) \in \llbracket p * \triangleright r \rrbracket_A \subseteq \llbracket q * \triangleright r \rrbracket_A$ . Hence, there exists  $m' \in \mathcal{M}$  such that  $\triangleright(m' \in q_2 * r \wedge (h', U') \in \llbracket m' \rrbracket_A)$ , and so

$$\triangleright(m' \in \llbracket q_2 * r \rrbracket_A \subseteq \llbracket q_1 * r \rrbracket_A),$$

which ends the proof.  $\square$

**Lemma 17** (Atomic-Frame).

$$\begin{aligned} &\forall A \in \mathcal{P}(\Delta(RId)), p, q, r \in Prop_{TSO}. \\ &\text{stable}_A(r) \wedge a \text{ sat}_A \{p\} \{q\} \Rightarrow a \text{ sat}_A \{p * \triangleright r\} \{q * r\} \end{aligned}$$

*Proof.* Assume  $r' \in Prop_{TSO}$ ,  $m \in \mathcal{M}$ ,  $h, h' \in Heap$  and  $U, U' \in SPool$  such that

$$\text{stable}_A(r') \quad m \in p * \triangleright r * \triangleright r' \quad (h, U) \in \llbracket m \rrbracket_A \quad (h', U') \in \llbracket a \rrbracket(h, U).$$

This means that  $m \in p * \triangleright(r * r')$ . Since  $\text{stable}_A(r * r')$ , from definition of atomic satisfiability it follows that there exists an  $m' \in \mathcal{M}$  such that

$$\triangleright(m' \in q * r * c \wedge (h', U') \in \llbracket m' \rrbracket_A),$$

which ends the proof.  $\square$

**Corollary 1.**

$$\begin{aligned} &\forall A \in \mathcal{P}(\Delta(RId)), p, q, r \in Prop_{TSO}. \\ &\text{stable}_A(r) \wedge a \text{ sat}_A \{p\} \{q\} \Rightarrow a \text{ sat}_A \{p * r\} \{q * r\} \end{aligned}$$

**Lemma 18** (Id-Refl).  $\forall A \in \mathcal{P}(\Delta(RId)), p \in Prop_{TSO}. \varepsilon \text{ sat}_A \{p\} \{p\}$

*Proof.* Take a  $r \in Prop_{TSO}$ ,  $m \in \mathcal{M}$ ,  $h \in Heap$  and  $U \in SPool$  such that

$$\text{stable}_A(r) \quad m \in p * \triangleright r \quad (h, U) \in \lfloor m \rfloor_A$$

Since  $\llbracket \varepsilon \rrbracket(h, U) = \{(h, U)\}$ , so clearly  $\not\vdash \notin \llbracket \varepsilon \rrbracket(h, U)$ . Now, it suffices to exhibit  $m'$  such that  $(h, U) \in \lfloor m' \rfloor_A$  and  $\triangleright(m' \in p * r)$ , which follow simply by monotnicity of  $\triangleright$  if we choose  $m' = m$ .  $\square$

**Lemma 19** (Flush-Refl). *For any  $A \in \mathcal{P}(\Delta(RId))$ ,  $p \in Prop_{TSO}$ ,  $t \in TId$ ,*

$$\text{stable}_A(p) \Rightarrow \text{flush}(t) \text{ sat}_A \{p\} \{p\}.$$

*Proof.* Take an  $r \in Prop_{TSO}$ ,  $m \in \mathcal{M}$ ,  $h \in Heap$  and  $U \in SPool$  such that

$$\text{stable}_A(r) \quad m \in p * \triangleright r \quad (h, U) \in \lfloor m \rfloor_A.$$

Since *flush* never faults, we only need to consider the postcondition branch. Take any  $(h', U') \in \llbracket \text{flush}(t) \rrbracket(h, U)$ . By definition this means that there exist  $(o, f) \in \text{dom}(h)$  and  $v$  such that  $h' = h[(o, f) \mapsto v]$ ,  $U(t) = (o, f, v) \cdot U'(t)$  and  $U(t') = U'(t')$  for any  $t' \neq t$ . As a witness we give  $m' = (\text{flush}(m.l, (o, f, v)), m.s, U', m.\zeta)$ . Since  $(m.l, U) R_f (m'.l, U')$  and  $\text{stable}_A(p * \triangleright r)$ , we easily have  $\triangleright(m' \in p * r)$  by monotonicity of  $\triangleright$ . This means we only need to show that  $(h', U') \in \lfloor m' \rfloor_A$ . From  $(h, U) \in \lfloor m \rfloor_A$  we get  $l_h$  and  $sr$  such that  $h = l_h.h$ ,  $l_h = m.l \bullet \prod_{r \in \text{dom}(m.s) \cap A} sr(r)$  and  $(sr(r), U) \in \lfloor (m.s, m.\zeta) \rfloor_r$  for any  $r \in \text{dom}(m.s) \cap A$ . Take  $l'_h = \text{flush}(l_h, (o, f, v))$  and  $sr'(r) = \text{flush}(sr(r), (o, f, v))$  as the witnesses. Since *flush* commutes over composition, we get

$$\begin{aligned} \text{flush}(l_h, (o, f, v)) &= \text{flush}(m.l \bullet \prod_{r \in \text{dom}(m.s) \cap A} sr(r), (o, f, v)) \\ &= m.l' \bullet \prod_{r \in \text{dom}(m.s) \cap A} sr'(r). \end{aligned}$$

Since, for every  $r \in \text{dom}(m.s) \cap A$  we also have  $(sr(r), U) R_f (sr'(r), U')$  and the erasures of shared regions are closed under interference (since they are defined by the application of an action model), we get  $(sr'(r), U') \in \lfloor (m.s, m.\zeta) \rfloor_r$  for each  $r \in \text{dom}(m.s) \cap A$ , which ends the proof.  $\square$

## 2.9 Specification Embedding

$$\begin{aligned} \text{spec} : \Omega &\rightarrow Prop_{TSO} \in \mathcal{S} \\ \text{spec}(s) &\stackrel{\text{def}}{=} \{m \in \mathcal{M} \mid s\} \end{aligned}$$

**Lemma 20** (Sat-Spec-Later).

$$\begin{aligned} \forall a \in Act, p, q \in Prop_{TSO}, s \in \Omega. \\ a \text{ sat } \{p * \text{spec}(\triangleright s)\} \{q\} &\Rightarrow a \text{ sat } \{p * \text{spec}(\triangleright s)\} \{q * \text{spec}(s)\} \end{aligned}$$

*Proof.* Take  $r \in Prop_{TSO}$ ,  $m \in \mathcal{M}$ ,  $h, h' \in Heap$  and  $U, U' \in SPool$  such that

$$\text{stable}(r) \quad m \in p * \text{spec}(\triangleright s) * \triangleright r \quad (h, U) \in [m]_A \quad (h', U') \in \llbracket a \rrbracket(h, U).$$

By our assumption, there exists an  $m' \in \mathcal{M}$ , such that  $\triangleright(m' \in q * r)$  and  $\triangleright(h' \in [m'])$ . Hence,  $\triangleright((m' \in q * r) \wedge s)$ , and so  $\triangleright(m' \in q * \text{spec}(s) * r)$ .  $\square$

**Lemma 21** (Later-Star).  $\triangleright(p * q) = (\triangleright p) * (\triangleright q)$

*Proof* ( $\subseteq$ ). Assume  $m \in \triangleright(p * q)$ , and so  $\triangleright(m \in p * q)$ , and, unfolding the definition of  $*$ ,

$$\triangleright(\exists m_1, m_2 \in \mathcal{M}. m = m_1 \bullet m_2 \wedge m_1 \in p \wedge m_2 \in q).$$

Now, unfolding the definition of  $\bullet$ , we get

$$\triangleright(\exists l_1, l_2 \in \Delta(LState). m.l = l_1 \bullet_{LState} l_2 \wedge (l_1, m.s, m.U, m.a) \in p \wedge (l_2, m.s, m.U, m.a) \in q).$$

This means we can take the witnesses  $l_1$  and  $l_2$ , such that

$$\triangleright(m.l = l_1 \bullet_{LState} l_2) \quad \triangleright((l_1, m.s, m.U, m.a) \in p) \quad \triangleright((l_2, m.s, m.U, m.a) \in q).$$

From the first of these properties, by Lemma 2 we get  $m.l = l_1 \bullet_{LState} l_2 \vee \triangleright \perp$ . We inspect the two cases.

Assume, first, that  $m.l = l_1 \bullet_{LState} l_2$ . Then,  $m = (l_1, m.s, m.U, m.a) \bullet (l_2, m.s, m.U, m.a)$ , and so  $m \in (\triangleright p) * (\triangleright q)$ .

In the other case, assume  $\triangleright \perp$ . Now, take the splitting of  $m = m \bullet (\varepsilon, m.s, m.U, m.a)$ . Since  $\triangleright \perp$ ,  $\triangleright((\varepsilon, m.s, m.U, m.a) \in q)$ , and so  $m \in (\triangleright p) * (\triangleright q)$ .  $\square$

*Proof* ( $\supseteq$ ). Follows by monotonicity of  $\triangleright$ .  $\square$

## 2.10 Thread Safety

We define safety by guarded recursion:

$$\begin{aligned} \text{safe} &: \mathcal{P}((TId \times Expr) \times Prop_{TSO} \times (Val \rightarrow Prop_{TSO})) \in \mathcal{S} \\ \text{safe} &\stackrel{\text{def}}{=} \text{fix } f. \widehat{\text{safe}}(f), \end{aligned}$$

where

$$\begin{aligned} \widehat{\text{safe}}(f)((t, e), p, q) &\stackrel{\text{def}}{=} (\exists v \in Val. e = v \wedge p \sqsubseteq q(v)) \vee \\ &\forall T \in \Delta(TPool), a \in \Delta(Act). (t, e) \xrightarrow{a} T \Rightarrow \\ &\exists p' : \text{dom}(T) \rightarrow Prop_{TSO}. t \in \text{dom}(T) \wedge \\ &(\forall t \in \text{dom}(T). \triangleright \text{stable}(p'(t))) \wedge \\ &a \text{ sat } \{p\} \{ \otimes_{t \in \text{dom}(T)} p'(t) \} \wedge \\ &\triangleright f((t, T(t)), p'(t), q) \wedge \\ &\triangleright \forall t' \in \text{dom}(T - t). f(t', p'(t'), \lambda v. \top) \end{aligned}$$

**Lemma 22** (Stable-Closed).

$$\begin{aligned} \forall X \in \mathcal{P}(\Delta(SId)), p \in \Delta(LTS), r \in \Delta(RId). \\ (\forall \alpha \in AId. \alpha \neq \alpha_i \Rightarrow p(\alpha)(X) \subseteq X) \Rightarrow \text{stable}(\text{region}(X, p, r) * \otimes_i [\alpha_i]_1^r) \end{aligned}$$

**Lemma 23** (ViewShift-Weaken).

$$\forall A, B \in \mathcal{P}(\Delta(RId)), p, q \in \text{Prop}_{TSO}. A \subseteq B \wedge p \sqsubseteq_A q \Rightarrow p \sqsubseteq_B q$$

**Lemma 24** (ViewShift-Trans).

$$\forall p, q, r \in \text{Prop}_{TSO}, A \in \mathcal{P}(\Delta(RId)). p \sqsubseteq_A q \wedge q \sqsubseteq_A r \Rightarrow p \sqsubseteq_A r$$

**Lemma 25** (RegInterp-WF).

$$\begin{aligned} \forall p \in \Delta(SId) \rightarrow \text{Prop}_{TSO}. \text{stable}(p) \Rightarrow \\ (\lambda(x, \zeta). \{(l, s, U) \mid (l, s, U, \zeta) \in p(x)\} \in \Delta(SId) \times AMod \rightarrow_{\text{mon}} \mathcal{P}^\uparrow(AState)) \end{aligned}$$

*Proof (Monotonicity).* Take  $x \in \Delta(SId)$ ,  $\zeta_1, \zeta_2 \in AMod$  such that  $\zeta_1 \leq \zeta_2$ . By upwards-closure of  $p$ , for any  $l \in LState$ ,  $s \in SState$  and  $U \in SPool$  such that  $(l, s, U, \zeta_1) \in p(x)$ ,  $(l, s, U, \zeta_2) \in p(x)$  also holds.  $\square$

*Proof (Upwards-closure).* Take  $x \in \Delta(SId)$ ,  $\zeta \in AMod$ ,  $l_1, l_2 \in LState$ ,  $s_1, s_2 \in SState$  and  $U_1, U_2 \in SPool$  such that  $(l_1, s_1, U_1, \zeta) \in p(x)$  and  $(l_1, s_1, U_1) R (l_2, s_2, U_2)$ . By stability of  $p$  it follows that  $(l_2, s_2, U_2, \zeta) \in p(x)$ .  $\square$

**Lemma 26** (Safe-Sat). *For any thread  $t \in \Delta(TId)$ , closed expression  $e \in \Delta(Expr)$ ,  $p \in \text{Prop}_{TSO}$  and  $q \in \Delta(Val) \rightarrow \text{Prop}_{TSO}$  we have*

$$\begin{aligned} (\forall a \in \Delta(Act). \forall v \in \Delta(Val). (t, e) \xrightarrow{a} \{(t, v)\} \Rightarrow a \text{ sat } \{p\} \{q(v)\}) \wedge \\ \text{stable}(p) \wedge (\forall v \in \Delta(Val). \text{stable}(q(v))) \wedge \text{atomic}(e) \\ \Rightarrow \text{safe}((t, e), p, q). \end{aligned}$$

*Proof.* The proof proceeds by Löb induction, which we will need to handle the flushing of the store buffer that can occur before the atomic command is actually executed.

Unfolding the definition of safety, we find that, since  $e$  is atomic, we should take the “progress” case. Take any  $T \in TPool$  and  $a \in Act$  such that  $(t, e) \xrightarrow{a} T$ . Since  $e$  is atomic, we only need to consider the cases when  $a$  is a *read*, *write*, *cas* or *flush* action. We proceed by case analysis.

First, let us consider the flushing. By definition of the semantics, we know that  $T = \{(t, e)\}$ . We take the map  $p'$  to be the map given by  $t \mapsto p$ . Thus, stability of  $p'$  follows by stability of  $p$ , atomic satisfaction of the *flush*( $t$ ) action – by Lemma 19, and since we did not fork any threads we only need to show that  $\triangleright \text{safe}((t, e), p, q)$  – which follows from our Löb inductive hypothesis.

In the other three cases, we have to consider proper reductions. In all of these cases  $T$  has the same shape: we have  $T = \{(t, v)\}$  for some value  $v$ . Thus, we can take the map  $p'$

to be given by  $t \mapsto q(v)$ . Then, the stability of  $p'$  follows by stability of  $q$ , and the atomic satisfaction by the lemma's assumption. Since we do not allocate any new threads, this only leaves us with showing that  $\triangleright \text{safe}((t, v), q(v), q)$ , which holds trivially by taking the "value" branch of the safety and showing that  $\triangleright(q(v) \sqsubseteq q(v))$  by monotonicity of  $\triangleright$  and reflexivity of the view-shift.  $\square$

**Lemma 27** (Region-Alloc).

$$\begin{aligned} & \forall F \in \mathcal{P}(\Delta(RId)), T \in \Delta(LTS), X \in \mathcal{P}(\Delta(SId)), x \in X. \\ & \forall I \in \Delta(RId) \rightarrow \Delta(SId) \rightarrow Prop_{TSO}, P \in Prop_{TSO}, A, B \in \mathcal{P}(\Delta(AId)) \\ & (\forall r \in \Delta(RId), s \in \Delta(SId). \text{stable}(I(r)(s))) \wedge F \text{ infinite} \wedge A, B \text{ finite} \wedge \\ & \text{valid}(\forall n \in F. P * \otimes_{\alpha \in A} [\alpha]_1^n \Rightarrow \triangleright I(n)(x)) \wedge A \cap B = \emptyset \\ & \Rightarrow P \sqsubseteq_F \exists n \in F. \text{region}(X, T, n) * \text{rintr}(I(n), n) * \otimes_{\alpha \in B} [\alpha]_1^n \end{aligned}$$

*Proof.* By the definition of the view shift, we take any  $r \in Prop_{TSO}$ ,  $m \in \mathcal{M}$ ,  $h \in \text{Heap}$ ,  $U \in \text{SPool}$  such that  $m \in P * r$  and  $(h, U) \in \lfloor m \rfloor_F$ . Hence, there have to exist  $l_1, l_2, l_c \in \text{LState}$ ,  $s \in \text{SState}$ ,  $\zeta \in \text{AMod}$  and  $sr : \text{dom}(s) \cap F \rightarrow \text{LState}$  such that

$$\begin{aligned} h &= l_c \bullet h & l_c &= l_1 \bullet_{\text{LState}} l_2 \bullet \prod_{r \in \text{dom}(s) \cap F} sr(r) & m &= (l_1 \bullet_{\text{LState}} l_2, s, U, \zeta) \\ \forall r \in \text{dom}(s) \cap F. & (sr(r), U) \in \lfloor (s, \zeta) \rfloor_r & (l_1, s, U, \zeta) &\in P & (l_2, s, U, \zeta) &\in r \end{aligned}$$

Pick an  $n \in U$  such that  $\forall \alpha \in AId. l_c.c(n, \alpha) = 0$  and  $n \notin \text{dom}(s) \cup \text{dom}(\zeta)$ . Now, let  $s' = s[n \mapsto (x, T)]$ ,  $\zeta' = \zeta[n \mapsto \text{lam}(\lambda(y, \zeta). \{(l, s, U) \in \text{AState} \mid (l, s, U, \zeta) \in I(n)(y)\})]$ . Note that  $\zeta'$  is well-defined, by virtue of Lemma 25. We have  $s \leq s'$  and  $\zeta \leq \zeta'$ , and so  $(l_1, s', U, \zeta') \in P$  and  $(l_2, s', U, \zeta') \in r$ . Furthermore,  $((\varepsilon, [(n, A) \mapsto 1]), s', U, \zeta') \in \otimes_{\alpha \in A} [\alpha]_1^n$ , and so, by the validity assumption,

$$(l_1 \bullet_{\text{LState}} (\varepsilon, [(n, A) \mapsto 1]), s', U, \zeta') \in \triangleright I(n)(x).$$

Now, clearly  $(\varepsilon, s', U, \zeta') \in \text{region}(X, T, n)$ . Furthermore, for any  $\zeta'' \geq \zeta'$  and  $y \in \text{SId}$ ,

$$\begin{aligned} \text{app}(\zeta'(n))(y, \zeta'') &= \text{app}(\text{lam}(\lambda(y, \zeta). \{(l, s, U) \mid (l, s, U, \zeta) \in I(n)(y)\}))(y, \zeta'') \\ &= (\triangleright(\lambda(y, \zeta). \{(l, s, U) \mid (l, s, U, \zeta) \in I(n)(y)\}))(y, \zeta'') \\ &= \{(l, s, U) \mid (l, s, U, \zeta'') \in \triangleright I(n)(y)\}. \end{aligned}$$

Thus,  $(\varepsilon, s', U, \zeta'') \in \text{rintr}(I, n)$  also holds. Hence, we have

$$(l_2 \bullet_{\text{LState}} (\varepsilon, [(n, B) \mapsto 1]), s', U, \zeta') \in \text{region}(X, T, n) * \text{rintr}(I, n) * (\otimes_{\alpha \in B} [\alpha]_1^n) * r.$$

Lastly, we also get  $(h, U) \in \lfloor (l_2 \bullet_{\text{LState}} (\varepsilon, [(n, B) \mapsto 1]), s', U, \zeta') \rfloor_F$ , which ends the proof.  $\square$

**Lemma 28.**

$$\begin{aligned}
& \forall A \in \mathcal{P}(\Delta(RId)), X, Y \in \mathcal{P}(\Delta(SId)), T \in \Delta(LTS), n \in \Delta(RId). \\
& \forall I \in \Delta(SId) \rightarrow Prop_{TSO}, p, q \in Prop_{TSO}, \alpha \in \Delta(AId), \pi \in \Delta(Perm^+), f : X \rightarrow Y. \\
& \text{stable}(p) \wedge n \in A \wedge (\forall x \in X. (x, f(x)) \in T(\alpha) \vee f(x) = x) \wedge \\
& (\forall x \in X. p * (\triangleright I(x)) * [\alpha]_{\pi}^n \sqsubseteq_{A \setminus \{n\}} q * \triangleright I(f(x))) \\
& \Rightarrow (\text{region}(X, T, n) * \text{rintr}(I, n) * p * [\alpha]_{\pi}^n \sqsubseteq_A \text{region}(Y, T, n) * q)
\end{aligned}$$

*Proof.* Take  $r \in Prop_{TSO}$ ,  $m \in \mathcal{M}$ ,  $h \in Heap$  and  $U \in SPool$  such that

$$\text{stable}(r) \quad m \in \text{region}(X, T, n) * \text{rintr}(I, n) * p * [\alpha]_{\pi}^n * r \quad (h, U) \in [m]_A.$$

This means, there exist  $l_1, l_2, l_3, l_4, l_5 \in LState$ ,  $s \in SState$ ,  $U' \in SPool$  and  $\zeta \in AMod$  such that

$$\begin{aligned}
& (l_1, s, U', \zeta) \in \text{region}(X, T, n) \quad (l_2, s, U', \zeta) \in \text{rintr}(I, n) \\
& (l_3, s, U', \zeta) \in p \quad (l_4, s, U', \zeta) \in [\alpha]_{\pi}^n \quad (l_5, s, U', \zeta) \in r \\
& m = (l_1 \bullet_{LState} l_2 \bullet_{LState} l_3 \bullet_{LState} l_4 \bullet_{LState} l_5, s, U', \zeta).
\end{aligned}$$

Let  $l = m.l$ . By unfolding  $(h, U) \in [m]_A$  we get that  $U = U'$ , and  $l_c \in LState$  and  $sr : \text{dom}(s) \cap A \rightarrow LState$  such that

$$h = l_c.h \quad l_c = l \bullet_{LState} \prod_{r \in \text{dom}(s) \cap A} sr(r) \quad \forall r \in \text{dom}(s) \cap A. (sr(r), U) \in [(s, \zeta)]_r.$$

Since  $(l_1, s, U, \zeta) \in \text{region}(X, T, n)$ , we know  $n \in \text{dom}(s)$ ,  $s(n).s \in X$  and  $s(n).p = T$  hold. Since  $n \in \text{dom}(s) \cap A$ , we get  $sr(n) \in [(s, \zeta)]_n$  and  $\pi_2(sr(n)) = U$ , and so  $(sr(n), s, U) \in \text{app}(\zeta(n))(s(n).s, \zeta)$ . Unfolding  $(l_2, s, U, \zeta) \in \text{rintr}(I, n)$ , we get that

$$\forall x \in \Delta(SId), \zeta' \geq \zeta. \text{app}(\zeta(n))(x, \zeta') = \triangleright(\lambda(l, s, U). I(x)(l, s, U, \zeta')).$$

From this, instantiating with  $s(n).s$  and  $\zeta$ , we get that  $\triangleright(I(s(n).s)(sr(n), s, U, \zeta))$ . By assumption, we know that  $(s(n).s, f(s(n).s)) \in T(\alpha)$  (in the case where  $f(s(n).s) = s(n).s$  it suffices to use the assumed view-shift). Furthermore, since  $(l_4, s, \zeta) \in [\alpha]_{\pi}^n$  it follows that  $\pi \leq l_4.c(n, \alpha)$  and thus  $l_3.c(n, \alpha) < 1$ ,  $sr(n).c(n, \alpha) < 1$  and  $l_5.c(n, \alpha) < 1$ . Thus,

$$(l_3, s, U) R_A (l_3, s', U) \quad (l_5, s, U) R_A (l_5, s', U) \quad (sr(n), s, U) R_A (sr(n), s', U),$$

where  $s' = s[n \mapsto (f(s(n).s), s(n).p)]$ . Hence, by stability of  $p$ ,  $\triangleright I(s(n).s)$  and  $r$ , it follows that

$$(l_3, s', U, \zeta) \in p \quad (sr(n), s', U, \zeta) \in \triangleright I(s(n).s) \quad (l_4, s', U, \zeta) \in [\alpha]_{\pi}^n \quad (l_5, s', U, \zeta) \in r.$$

Furthermore, for every  $r \in \text{dom}(s) \cap (A \setminus \{n\})$ ,  $sr(r).c(n, \alpha) < 1$ , so

$$(sr(r), s, U) R_A (sr(r), s', U),$$

and so, by stability of region interpretations,

$$\forall r \in \text{dom}(s') \cap (A \setminus \{n\}). (sr(r), U) \in \llbracket (s', \zeta) \rrbracket_r.$$

Thus, we get that

$$(l \bullet_{LState} \pi_1(sr(n)), s', U, \zeta) \in \text{region}(\{f(s(n).s)\}, T, n) * (\triangleright I(s(n).s)) * [\alpha]_{\pi}^n * p * r * \{m \in \mathcal{M} \mid \zeta \leq m.a\}$$

and

$$(h, U) \in \llbracket (l \bullet_{LState} sr(n), s', U, \zeta) \rrbracket_{A \setminus \{n\}}.$$

Now we can use the assumed view-shift, from which it follows that there exists an  $m' \in \mathcal{M}$  such that

$$m' \in q * (\triangleright I(f(s(n).s))) * r * \text{region}(\{f(s(n).s)\}, T, n) * \{m \in \mathcal{M} \mid \zeta \leq m.a\} \\ (h, U) \in \llbracket m' \rrbracket_{A \setminus \{n\}}.$$

Thus, there exist  $l'_1, l'_2, l'_3, l'_4, l'_5 \in LState$ ,  $s'' \in SState$  and  $\zeta' \in AMod$  such that

$$m' = (l'_1 \bullet_{LState} l'_2 \bullet_{LState} l'_3 \bullet_{LState} l'_4 \bullet_{LState} l'_5, s'', U, \zeta') \\ (l'_1, s'', U, \zeta') \in q \quad (l'_2, s'', U, \zeta') \in \triangleright I(f(s(n).s)) \quad (l'_3, s'', U, \zeta') \in r \\ (l'_4, s'', U, \zeta') \in \text{region}(\{f(s(n).s)\}, T, n) \quad (l'_5, s'', U, \zeta') \in \{m \in \mathcal{M} \mid \zeta \leq m.a\}.$$

Hence,  $s''(n).s = f(s(n).s) = s'(n).s$ . From  $(h, U) \in \llbracket m' \rrbracket_{A \setminus \{n\}}$ , it follows that there exist  $l'_c \in LState$  and  $sr' : \text{dom}(s'') \cap (A \setminus \{n\}) \rightarrow LState$  such that

$$l'_c = l'_1 \bullet_{LState} l'_2 \bullet_{LState} l'_3 \bullet_{LState} l'_4 \bullet_{LState} l'_5 \bullet_{LState} \otimes_{r \in \text{dom}(s'') \cap (A \setminus \{n\})} \pi_1(sr'(r)) \\ h = l'_c.h \quad \forall r \in \text{dom}(s'') \cap (A \setminus \{n\}). (sr'(r), U) \in \llbracket (s'', \zeta') \rrbracket_r.$$

From  $(l'_5, s'', U, \zeta') \in \{m \in \mathcal{M} \mid \zeta \leq m.a\}$  it follows that  $\zeta \leq \zeta'$ , and thus

$$\forall x \in \Delta(SId). \text{app}(\zeta'(n))(x, \zeta') = \triangleright(\lambda(l, s, U). I(x)(l, s, U, \zeta')).$$

Hence,  $(l'_2, s'', U) \in \text{app}(\zeta'(n))(f(s(n).s), \zeta')$ , and so  $(l'_2, U) \in \llbracket (s'', \zeta') \rrbracket_n$ . Thus, we have  $(h, U) \in \llbracket \text{region}(Y, T, n) * q \rrbracket_A$ , which ends the proof.  $\square$

**Lemma 29.**

$$\forall A \in \mathcal{P}(\Delta(RId)). \forall X \in \mathcal{P}(\Delta(SId)). \forall T \in \Delta(LTS). \forall n \in \Delta(RId).$$

$$\forall I \in \Delta(SId) \rightarrow \text{Prop}_{TSO}. \forall p, p' \in \text{Prop}_{TSO}. \forall q \in \Delta(SId) \rightarrow \text{Prop}_{TSO}. \forall B \in \mathcal{P}(\Delta(AId)).$$

$$\forall f : \Delta(SId) \rightarrow \mathcal{P}(\Delta(SId)). \forall g : \Delta(AId) \rightarrow \Delta(\text{Perm}^+).$$

$$(\forall y \in \Delta(SId). \text{stable}(q(y))) \wedge n \in A \wedge (\forall x \in X. \forall y \in f(x). (x, y) \in T(B)^*) \wedge$$

$$(\forall x \in X. a \text{ sat}_{A \setminus \{n\}} \{p * (\triangleright I(x))\}) \left\{ \exists y \in f(x). q(y) * \otimes_{\alpha \in B} [\alpha]_{g(\alpha)}^n * \triangleright I(y) \right\}$$

$$\Rightarrow a \text{ sat}_A \{ \text{region}(X, T, I, n) * p \} \left\{ \exists y \in \Delta(SId). \text{region}(\{y\}, T, n) * \otimes_{\alpha \in B} [\alpha]_{g(\alpha)}^n * q(y) \right\}$$

*Proof.* Assume  $r \in Prop_{TSO}$ ,  $m \in \mathcal{M}$ ,  $h, h' \in \Delta(Heap)$  and  $U, U' \in SPool$  such that

$$m \in region(X, T, I, n) * p * \triangleright r \quad (h, U) \in [m]_A \quad (h', U') \in \llbracket a \rrbracket_A(h, U) \quad stable(r).$$

Thus, there exist  $l_1, l_2, l_3, l_4 \in LState$ ,  $s \in SState$  and  $\zeta \in AMod$  such that

$$(l_1, s, U, \zeta) \in region(X, T, n) \quad (l_2, s, U, \zeta) \in rintr(I, n) \quad (l_3, s, U, \zeta) \in p \\ (l_4, s, U, \zeta) \in \triangleright r \quad m = (l_1 \bullet l_2 \bullet l_3 \bullet l_4 \bullet l_5, s, U, \zeta).$$

Let  $l = l_1 \bullet l_2 \bullet l_3 \bullet l_4$ . Unfolding  $(h, U) \in [m]_A$  we get an  $l_c \in LState$  and  $sr : (\text{dom}(s) \cap A) \rightarrow LState$  such that

$$h = l_c.h \quad l_c = l \bullet \otimes_{r \in (\text{dom}(s) \cap A)} sr(r) \quad \forall r \in (\text{dom}(s) \cap A). (sr(r), U) \in [(s, \zeta)]_r.$$

Since  $(l_1, s, U, \zeta) \in region(X, T, n)$ , it follows that  $n \in \text{dom}(s)$ ,  $s(n).s \in X$  and  $s(n).p = T$ . Since we also know that  $n \in A$ , it follows that  $(sr(n), U) \in [(s, \zeta)]_n$ , and so  $(sr(n), s, U) \in app(\zeta(n))(s(n).s, \zeta)$ . Unfolding  $(l_2, s, U, \zeta) \in rintr(I, n)$  it follows that

$$\forall x \in \Delta(SId). \forall \zeta' \geq \zeta. app(\zeta(n))(x, \zeta') = \triangleright(\lambda(l, s, U). I(x)(l, s, U, \zeta')).$$

Thus, in particular,  $app(\zeta(n))(s(n).s, \zeta) = \triangleright(\lambda(l, s, U). I(s(n).s)(l, s, U, \zeta))$ , and so

$$\triangleright I(s(n).s)(sr(n), s, U, \zeta).$$

The high-level plan is now to use the assumption with a slightly modified state to account for opening of the region, and afterwards use the state obtained from the assumption to build a final instrumented state. To this end, we instantiate the assumption with a frame  $r' = r * region(\{s(n).s\}, T, n) * \{m \in \mathcal{M} \mid m.\zeta \geq \zeta\}$  and an instrumented state  $\bar{m} = (l \bullet sr(n), s, U, \zeta)$ . The erasure  $(h, U) \in [\bar{m}]_{A \setminus \{n\}}$  holds easily, since  $n \notin A \setminus \{n\}$  but  $sr(n)$  is in local state. We also know that  $(l_3, s, U, \zeta) \in p$ ,  $(sr(n), s, U, \zeta) \in \triangleright(I(s(n).s))$  and  $(l_4, s, U, \zeta) \text{ in } \triangleright r$ , so it suffices to show that

$$(l_1 \bullet l_2, s, U, \zeta) \in \triangleright(region(\{s(n).s\}, T, n) * \{m \in \mathcal{M} \mid m.\zeta \geq \zeta\}),$$

which holds by monotonicity of later and the definition. The frame is also stable, since we cannot change the state of region  $n$ .

This means we get a state  $\bar{m}'$  such that  $(h', U') \in \triangleright[\bar{m}']_{A \setminus \{n\}}$  and

$$\bar{m}' \in \triangleright(\exists y \in f(x) * q(y) * \otimes_{\alpha \in B} [\alpha]_{g(\alpha)}^n * \triangleright I(y) * r').$$

Thus, we can take a  $y \in \Delta(SId)$  such that  $\triangleright(y \in f(x))$ , and local states  $l'_1, l'_2, l'_3$  and  $l'_4$ , such that

$$\bar{m}' = (l'_1 \bullet l'_2 \bullet l'_3 \bullet l'_4, s', U', \zeta') \\ (l'_1, s', U', \zeta') \in \triangleright q(y) \quad (l'_2, s', U', \zeta') \in \triangleright \otimes_{\alpha \in B} [\alpha]_{g(\alpha)}^n \quad \triangleright ((l'_3, s', U', \zeta') \in \triangleright I(y)) \\ (l'_4, s', U', \zeta') \in \triangleright r \quad \triangleright (s'(n).s = s(n).s) \quad \triangleright (s'(n).p = T) \quad \triangleright (\zeta' \geq \zeta).$$



Now, we can take  $m' = (l'_1 \bullet l'_2 \bullet l'_4, s'[n \mapsto (y, T)], U', \zeta')$  as the witness. We need to show that  $\triangleright((h', U') \in \lfloor m' \rfloor_A)$  and  $m' \in \triangleright(\exists y \in \Delta(\text{SID}). \text{region}(\{y\}, T, n) * q(y) * \otimes_{\alpha \in B} [\alpha]_{g(\alpha)}^n * r)$ . For the latter, we take  $y$  as the witness and split the local state four ways into  $\varepsilon$ ,  $l'_1$ ,  $l'_2$  and  $l'_4$ . The region assertion holds by definition, and the assertion about permissions by the earlier property, since it doesn't depend on the shared state. That leaves us with  $r$  and  $q(y)$ . Both these assertions are stable, so it suffices to show that  $x$  and  $y$  are in the interference relation from their perspective. To show this, notice that the permissions in  $B$  suffice to transition from  $x$  to  $y$  – and these permissions are outside both  $l'_1$  and  $l'_4$ . Thus, we can get  $(l'_1, s'[n \mapsto (y, T)], U', \zeta') \in \triangleright q(y)$  and  $(l'_4, s'[n \mapsto (y, T)], U', \zeta') \in \triangleright r$ .

This leaves us with proving that  $(h', U') \in \lfloor m' \rfloor_A$ . To this end, take  $\bar{l}$  and  $\bar{s}r$  :  $\text{dom}(s') \cap (A \setminus \{n\}) \rightarrow \text{LState}$  – the witnesses of  $(h', U') \in \lfloor \bar{m}' \rfloor_{A \setminus \{n\}}$ . We know that

$$\begin{aligned} h' &= \bar{l}.h \quad \bar{l} = l'_1 \bullet l'_2 \bullet l'_3 \bullet l'_4 \bullet \prod_{r \in \text{dom}(s) \cap (A \setminus \{n\})} \bar{s}r(r) \\ \forall r \in \text{dom}(s) \cap (A \setminus \{n\}). \quad &\triangleright((\bar{s}r(r), U') \in \lfloor (s', \zeta') \rfloor_r) \end{aligned}$$

Take as witnesses  $l' = \bar{l}$  and  $sr' = \bar{s}r[n \mapsto l'_3]$ . Clearly  $sr'$  has the appropriate domain, and the first two conditions hold. For the final one, take any  $r \in \text{dom}(s) \cap A$ . We have two cases to consider: either  $r = n$ , or  $r \in \text{dom}(s) \cap (A \setminus \{n\})$ . In the latter, we can use the same argument as for  $r$  and  $q(y)$  to show that the update of shared state is an allowed interference for the region  $r$ . In the former, we need to show that  $\triangleright(l'_3, U') \in \lfloor (s'[n \mapsto (y, T)], \zeta') \rfloor_r$ . Unrolling the definition, it suffices to show that  $\triangleright(l'_3, s'[n \mapsto (y, T)], U') \in \text{app}(\zeta'(r))(y, \zeta')$ . However, since  $\triangleright(\zeta' \geq \zeta)$ , by an earlier equation that followed from region interpretation, it suffices to show that  $\triangleright(l'_3, s'[n \mapsto (y, T)], U', \zeta') \in \triangleright I(y)$ , which follows by the same stability argument as the other assertions.  $\square$

**Lemma 30.** *For any  $p \in \text{Prop}_{SC}$ ,  $(l, s, U, \zeta) \in \ulcorner p \urcorner$ ,  $(l', U') \in R_{sb}(l, U)$ ,  $(l', s, U', \zeta) \in \ulcorner p \urcorner$ .*

*Proof.* To show the lemma holds, it suffices to show  $p$  is closed under each of  $R_n$ ,  $R_w$  and  $R_f$ . These proofs are presented in the following paragraphs.

$(R_n)$  Take any  $(l, s, U, \zeta) \in \ulcorner p \urcorner$  and  $(l', U') \in R_n(l, U)$ . By definition of  $R_n$  we have that  $l = l'$ ,  $\forall t \in \text{dom}(U') \setminus \text{dom}(U)$ .  $U'(t) = \varepsilon$  and  $\forall t \in \text{dom}(U)$ .  $U'(t) = U(t)$ . From definition of  $\ulcorner - \urcorner$  we get a state  $l_1 \leq l$  such that  $(l_1, s, \zeta) \in p$  and

$$\forall t \in \text{dom}(U). \forall (o, f) \in \text{dom}(l_1). \forall v \in \text{Val}. (o, f, v) \notin U(t).$$

We take the same  $l_1$  as a witness, and so we only need to show that  $(o, f, v) \notin U'(t)$  for any  $t \in \text{dom}(U')$ ,  $(o, f) \in \text{dom}(l_1)$  and  $v \in \text{Val}$ . If  $t \in \text{dom}(U)$ , the property holds, since in that case  $U(t) = U'(t)$  and we know it held for  $U$ . If  $t \notin \text{dom}(U)$ , on the other hand, we know  $U'(t) = \varepsilon$ , so in particular  $(o, f, v) \notin U'(t)$ .

$(R_w)$  Take any  $(l, s, U, \zeta) \in \ulcorner p \urcorner$  and  $(l', U') \in R_w(l, U)$ . By definition of  $R_w$  we know that  $l = l'$ ,  $\text{dom}(U) = \text{dom}(U')$  and we get  $t \in \text{dom}(U)$ ,  $o \in \text{OId}$ ,  $f \in \text{FName}$ , and

$v \in Val$  such that  $(o, f) \notin \text{dom}(l)$ ,  $U'(t) = U(t) \cdot (o, f, v)$  and  $U'(t') = U(t')$  for any  $t' \neq t$ . From the definition of  $\ulcorner - \urcorner$  we get a local state  $l_1 \leq l$  such that  $(l_1, s, \zeta) \in p$  and

$$\forall t \in \text{dom}(U). \forall (o, f) \in \text{dom}(l_1). \forall v \in Val. (o, f, v) \notin U(t).$$

Again, we take  $l_1$  as the witness, which means we only need to prove  $(o', f', v') \notin U(t')$  for any  $(o', f') \in \text{dom}(l_1)$ ,  $t' \in \text{dom}(U)$ . Clearly, it suffices to check that this holds for the newly added update,  $(o, f, v)$  in thread  $t$ , since the rest of the contents of  $U'$  is inherited from  $U$ . However, we know that  $(o, f) \notin \text{dom}(l)$ , so clearly  $(o, f) \notin \text{dom}(l_1)$ , so the property holds.

$(R_f)$  Take any  $(l, s, U, \zeta) \in \ulcorner p \urcorner$  and  $(l', U') \in R_w(l, U)$ . By definition of  $R_w$  we know that  $\text{dom}(U) = \text{dom}(U')$  and we get  $t \in \text{dom}(U)$ ,  $o \in OId$ ,  $f \in FName$  and  $v \in Val$  such that  $l' = \text{flush}(l, (o, f, v))$ ,  $U(t) = (o, f, v) \cdot U'(t)$  and  $U(t') = U'(t')$  for any  $t' \neq t$ . As before, from the definition of  $\ulcorner - \urcorner$  we get an  $l_1 \leq l$  such that  $(l_1, s, \zeta) \in p$  and there are no updates to the domain of  $l_1$ . From the latter property, we know that  $(o, f) \notin \text{dom}(l_1)$ , and so  $l_1 \leq l'$ . This allows us to take  $l_1$  as the witness, and in this case both properties hold trivially, since the updates in  $U'$  are a subset of those in  $U$ .  $\square$

## 2.11 Thread-pool Evaluation

$$eval : \Delta(MSt) \times \Delta(TPool) \times \mathcal{P}(\Delta(MSt) \times \Delta(TPool)) \rightarrow \Omega$$

$$eval(\mu, T, q) \stackrel{\text{def}}{=} (\text{irr}(\mu, T) \wedge (\mu, T) \in q) \vee (\forall T', \mu'. (\mu, T) \rightarrow (\mu', T') \Rightarrow \triangleright eval(\mu', T', q))$$

**Lemma 31** (Safe-Eval). *For any thread pool  $(\mu, T) \in \Delta(PSt)$  and families of assertions  $p : \text{dom}(T) \rightarrow Prop_{TSO}$ ,  $q : \text{dom}(T) \rightarrow Val \rightarrow Prop_{TSO}$ , if*

$$\begin{aligned} \forall t \in \text{dom}(T). \text{stable}(p(t)) \quad \forall t \in \text{dom}(T). \forall v \in Val. \text{stable}(q(t)(v)) \\ \forall t \in \text{dom}(T). \text{safe}((t, T(t)), p(t), q(t)) \quad \mu \in \llbracket \otimes_{t \in \text{dom}(T)} p(t) \rrbracket \end{aligned}$$

then

$$eval(\mu, T, \lambda(\mu', T'). \mu' \in \llbracket \otimes_{t \in \text{dom}(T)} q(t)(T'(t)) \rrbracket).$$

*Proof.* By Löb induction. Since we know  $\mu$  is an erasure of a state, we know there exist  $h$  and  $U$  such that  $\mu = (h, U)$  and  $\text{dom}(U) = \text{dom}(T)$ . By Lemma 1, there are two cases: either we are in a terminal state, or there exist a possible reduction. We proceed by inspecting these two.

Let us first consider the terminal case. By Lemma 1, we know that for any thread  $t \in \text{dom}(T)$ ,  $T(t)$  is a value, and  $U(t) = \varepsilon$ . Thus,  $\text{irr}(\mu, T)$  holds, and so we pick the left branch. We are left with showing that  $(h, U) \in \llbracket \otimes_{t \in \text{dom}(T)} q(t)(T(t)) \rrbracket$ . However, since we know that no thread can proceed, we also learn, from the safety assumption, that for any thread we have  $p(t) \sqsubseteq q(t)(T(t))$ . This shows us how to proceed: we should successively apply the view-shifts for each consecutive thread, while treating the separating conjunctions of the threads already shifted and these yet to be shifted as a

frame. Formally we need to weaken our assumption to say that for any thread pools  $T_1$  and  $T_2$  such that  $T_1 \uplus T_2 = T$  if  $(h, U) \in \llbracket \otimes_{t \in \text{dom}(T_1)} q(t)(T_1(t)) * \otimes_{t \in \text{dom}(T_2)} p(t) \rrbracket$ , then  $(h, U) \in \llbracket \otimes_{t \in \text{dom}(T)} q(t)(T(t)) \rrbracket$ , and proceed by induction on the size of  $T_2$ . In the base case  $T_2$  is empty, so  $T_1 = T$ , and the lemma holds trivially. Otherwise, we know that  $T_2 = (t, v) \uplus T'_2$  for some  $T'_2$ ,  $t$  and  $v$ . Thus, since all of  $p$  and  $q$  are stable, we can apply the view-shift for thread  $t$ : we know that  $(h, U) \in \llbracket p(t) * \otimes_{t \in \text{dom}(T_1)} q(t)(T_1(t)) * \otimes_{t \in \text{dom}(T'_2)} p(t) \rrbracket$ , so we get that  $(h, U) \in \llbracket q(t)(v) * \otimes_{t \in \text{dom}(T_1)} q(t)(T_1(t)) * \otimes_{t \in \text{dom}(T'_2)} p(t) \rrbracket$ . Now, to use the induction hypothesis for the smaller  $T'_2$  we only need to take  $T'_1 = T_1 \uplus (t, v)$ . Thus, this part of the proof is finished.

Let us now consider the case where a reduction could occur. The decomposition lemma gives us a state  $(\mu', T')$  such that  $((h, U), T) \rightarrow (\mu', T')$ . By definition of single step evaluation, this means that there is a thread  $t \in \text{dom}(T)$ , an action  $a \in \text{Act}$ , and a thread pool  $T' \in \text{TPool}$ , such that  $(t, T(t)) \xrightarrow{a} T''$ ,  $\mu' \in \llbracket a \rrbracket(\mu)$  and  $T' = (T - t) \uplus T''$ . We will consider the case when  $a = \text{flush}(t)$  separately, since this can occur even if  $T(t)$  is a value; all the other actions will follow from the safety of thread  $t$ .

Let us assume  $a = \text{flush}(t)$ . In this case we can take the instrumented state  $m$  that mediates the erasure and use the Lemma 19 to get a state  $m' \in \triangleright \otimes_{t \in \text{dom}(T)} p(t)$  such that  $\mu' \in \triangleright \llbracket m' \rrbracket$ . The remaining obligation is easily discharged by induction hypothesis and monotonicity of  $\triangleright$ .

Now, let us consider the case when  $a$  is not a flush action. In this case, we know that  $T(t)$  can not be a value, so from the assumption of safety of  $t$  we get  $p' : \text{dom}(T'') \rightarrow \text{Prop}_{TSO}$  and the following facts:

$$\begin{aligned} t \in \text{dom}(T'') \quad \forall t' \in \text{dom}(T''). \quad & \triangleright \text{stable}(p'(t')) \quad a \text{ sat } \{p(t)\} \{ \otimes_{t' \in \text{dom}(T'')} p'(t') \} \\ & \triangleright \text{safe}((t, T''(t)), p'(t), q(t)) \quad \forall t' \in \text{dom}(T'' - t). \quad \text{safe}((t', T''(t')), p'(t'), \lambda \_ . \top). \end{aligned}$$

As in the previous case, we can now take the mediating instrumented state  $m$  and use the atomic satisfaction, taking  $\otimes_{t' \in \text{dom}(T-t)} p(t')$  as a frame, to obtain a state  $m'$  such that

$$\mu' \in \triangleright \llbracket m' \rrbracket \quad m' \in \triangleright \otimes_{t' \in \text{dom}(T-t)} p(t') * \otimes_{t' \in \text{dom}(T'')} p'(t')$$

This, along with the previously obtained safety and stability properties allows us to use the induction hypothesis to conclude that

$$\triangleright \text{eval}(\mu', T', \lambda(\mu'', T'')). \quad \mu'' \in \llbracket \otimes_{t \in \text{dom}(T)} q(t)(T''(t)) * \otimes_{t \in \text{dom}(T''-t)} \top \rrbracket.$$

However, since we know that  $p * \top \subseteq \top$  for any assertion  $p$ , this gives us the required obligation that

$$\triangleright \text{eval}(\mu', T', \lambda(\mu'', T'')). \quad \mu'' \in \llbracket \otimes_{t \in \text{dom}(T)} q(t)(T''(t)) \rrbracket,$$

which ends the proof. □

Additional logical connectives

$$\begin{aligned}
\lceil p \text{ in } t \rceil &\stackrel{\text{def}}{=} \{(l, s, U, \zeta) \in \mathcal{M} \mid (\text{flush}(l, U(t)), s, U[t \mapsto \varepsilon], \zeta) \in \lceil p \rceil\} \\
p \mathcal{U}_t q &\stackrel{\text{def}}{=} \{(l, s, U, \zeta) \in \mathcal{M} \mid \exists \alpha, \beta \in \Delta(\text{SBuffer}). \exists o \in \Delta(\text{OID}). \exists f \in \Delta(\text{FName}). \exists v \in \Delta(\text{Val}). \\
&\quad U(t) = \alpha \cdot (o, f, v) \cdot \beta \wedge (l, s, U[t \mapsto \alpha], \zeta) \in p \wedge \\
&\quad (\text{flush}(l, \alpha \cdot (o, f, v)), s, U[t \mapsto \beta], \zeta) \in q\} \\
\text{cl}_A(p) &\stackrel{\text{def}}{=} \{m \in \mathcal{M} \mid \exists m' \in p. m' (\leq \cup R_A)^* m\}
\end{aligned}$$

Semantic read and write judgments

$$\begin{aligned}
p \vdash_w o.f \mapsto v, q &\stackrel{\text{def}}{=} \forall t \in \text{TId}. \forall (l, s, U, \zeta) \in p(t). t \in \text{dom}(U) \Rightarrow \triangleright \text{lookup}(o.f, U(t), l) = v \wedge \\
&\quad \triangleright (\text{flush}(l, U(t)), s, U[t \mapsto \varepsilon], \zeta) \in \lceil q(t) * o.f \mapsto v \rceil \\
p \vdash_r^m o.f \mapsto v &\stackrel{\text{def}}{=} \forall t \in \text{TId}. \forall m \in p(t). t \in \text{dom}(U) \Rightarrow \triangleright \text{lookup}(o.f, m.U(t), m.l) = v \\
p \vdash_r^o o.f \mapsto v &\stackrel{\text{def}}{=} \forall t, t' \in \text{TId}. \forall m \in p(t'). t \neq t' \wedge t \in \text{dom}(U) \Rightarrow \\
&\quad \triangleright \text{lookup}(o.f, m.U(t), m.l) = v
\end{aligned}$$

### 3 Logic

The specification logic is given by the specification entailment judgment,

$$\Gamma \mid \Phi \vdash S,$$

where  $S$  is a specification  $\Gamma$  — a logical variable context, and  $\Phi$  — a specification context. The specification logic extends a standard higher-order intuitionistic logic. Similarly, both of the assertion logics, SC and TSO, are given by the assertion logic entailment judgments, of the form

$$\Gamma; \Delta \mid \Phi \mid P \vdash_L Q,$$

where  $L$  ranges over  $\{SC, TSO\}$ ,  $P$  and  $Q$  are the assertions of the appropriate logic,  $\Gamma$  is the logical variable context,  $\Delta$  — the program variable context, and  $\Phi$  — the specification context. The SC logic is a standard higher order assertion logic, while the TSO logic extends the higher-order intuitionistic separation logic with several additional proof rules for the TSO connectives.

#### 3.1 TSO Assertion Logic

Listed below are some of the proof rules of the TSO assertion logic. For the sake of clarity, we do not present the introduction and elimination rules for the standard connectives of higher-order intuitionistic separation logic here.

$$\begin{array}{c}
\frac{\Gamma; \Delta \mid \Phi \vdash \text{stable}(P) \wedge \text{stable}(Q)}{\Gamma; \Delta \mid \Phi \mid (P \mathcal{U}_t Q) \vdash_{TSO} (P \mathcal{U}_t Q) \vee Q} \text{STAB-}\mathcal{U} \\
\frac{}{\Gamma; \Delta \mid \Phi \mid P \vdash_{TSO} (P)} \text{STAB-I} \quad \frac{\Gamma; \Delta \mid \Phi \vdash \text{stable}(P)}{\Gamma; \Delta \mid \Phi \mid (P) \vdash_{TSO} P} \text{STAB-S} \\
\frac{}{\Gamma; \Delta \mid \Phi \mid \ulcorner P \urcorner \vdash_{TSO} \ulcorner P \text{ in } t \urcorner} \text{W-HINTERP} \quad \frac{}{\Gamma; \Delta \mid P \mathcal{U}_t Q \vdash_{TSO} \ulcorner Q \text{ in } t \urcorner} \text{W-}\mathcal{U} \\
\frac{\Gamma; \Delta \mid \Phi \mid P_1 \vdash_{TSO} P_2 \quad \Gamma; \Delta \mid \Phi \mid Q_1 \vdash_{TSO} Q_2}{\Gamma; \Delta \mid \Phi \mid P_1 \mathcal{U}_t Q_1 \vdash_{TSO} P_2 \mathcal{U}_t Q_2} \text{CONS-}\mathcal{U} \\
\frac{\Gamma; \Delta \mid \Phi \vdash \text{b-stable}(P)}{\Gamma; \Delta \mid \Phi \mid P * (Q \mathcal{U}_t R) \vdash_{TSO} Q \mathcal{U}_t (P * R)} \text{*}\mathcal{U}
\end{array}$$

### 3.2 Syntactic Sugar

We provide some syntactic sugar for more convenient reasoning within the TSO logic, where the pre- and postconditions are parametrized with the current thread identifier.

$$\begin{aligned}
\bar{P} &\equiv \lambda t \in \text{TId}. \ulcorner P \text{ in } t \urcorner \\
P \mathcal{U} Q &\equiv \lambda t \in \text{TId}. P(t) \mathcal{U}_t Q(t) \\
\bar{P} \mathcal{U} Q &\equiv \lambda t \in \text{TId}. \exists t' \in \text{TId}. t \neq t' * P(t') \mathcal{U}_{t'} Q(t') \\
\text{stable}(P) &\equiv \text{r-stable}(P) \wedge \text{b-stable}(P)
\end{aligned}$$

### 3.3 Read and Write Rules

These judgments determine what values can be read from or written to the local state. The write judgment also provides the value of the field (for use of the compare-and-swap rule), as well as information about the state that will become flushed at the time the write action reaches the heap. The read judgment comes in two forms, dependent on whether the assertion currently inspected seen from the perspective of the thread which performs the read or some other thread's.

The forms of these judgments are  $\Gamma; \Delta \mid P \vdash_r^s x.f \mapsto v$  and  $\Gamma; \Delta \mid P \vdash_w x.f \mapsto v, Q$  respectively, where  $s$ , ranging over the set  $\{m, o\}$ , is the flag. Note that  $Q$  above is an  $\text{Prop}_{\text{SC}}$ .

$$\begin{array}{c}
\frac{}{\Gamma; \Delta \mid \triangleright x.f \mapsto v \vdash_r^m x.f \mapsto v} \text{R-SELF} \quad \frac{}{\Gamma; \Delta \mid \triangleright^\Gamma x.f \mapsto v^\Gamma \vdash_r^o x.f \mapsto v} \text{R-OTHER} \\
\frac{\Gamma; \Delta \mid Q \vdash_r^m x.f \mapsto v}{\Gamma; \Delta \mid P \mathcal{U} Q \vdash_r^m x.f \mapsto v} \text{R-}\mathcal{U}\text{-M} \quad \frac{\Gamma; \Delta \mid P \vdash_r^o x.f \mapsto v}{\Gamma; \Delta \mid P \mathcal{U} Q \vdash_r^o x.f \mapsto v} \text{R-}\mathcal{U}\text{-O} \\
\frac{\Gamma; \Delta \mid P \vdash_r^o x.f \mapsto v}{\Gamma; \Delta \mid P \bar{\mathcal{U}} Q \vdash_r^m x.f \mapsto v} \text{R-}\bar{\mathcal{U}} \quad \frac{\Gamma; \Delta \mid \cdot \mid P \vdash Q \quad \Gamma; \Delta \mid Q \vdash_r^s x.f \mapsto v}{\Gamma; \Delta \mid P \vdash_r^s x.f \mapsto v} \text{R-CONS} \\
\frac{}{\Gamma; \Delta \mid \triangleright x.f \mapsto v \vdash_w x.f \mapsto v, \top} \text{W-AX} \quad \frac{\Gamma; \Delta \mid P \vdash_w x.f \mapsto v, R}{\Gamma; \Delta \mid P * \triangleright \bar{Q} \vdash_w x.f \mapsto v, Q * R} \text{W-*} \\
\frac{\Gamma; \Delta \mid Q \vdash_w x.f \mapsto v, R}{\Gamma; \Delta \mid P \mathcal{U} Q \vdash_w x.f \mapsto v, R} \text{W-}\mathcal{U} \quad \frac{\Gamma; \Delta \mid \cdot \mid P \vdash Q \quad \Gamma; \Delta \mid Q \vdash_w x.f \mapsto v, R}{\Gamma; \Delta \mid P \vdash_w x.f \mapsto v, R} \text{W-CONS}
\end{array}$$

### 3.4 Specification Logic

In the following we present the rules for reasoning within the specification logic. The introduction and elimination rules for standard higher-order intuitionistic logic are omitted, since they are completely standard.

#### 3.4.1 Reasoning about Hoare triples

##### Atomic Command Rules

$$\begin{array}{c}
\frac{\Gamma; \Delta \mid P \vdash_r^m x.f \mapsto v}{\Gamma; \Delta \mid \Phi \vdash \langle P \rangle x.f \langle r. P * r = v \rangle^C} \text{A-READ} \\
\frac{\Gamma; \Delta \mid P \vdash_w x.f \mapsto v', Q}{\Gamma; \Delta \mid \Phi \vdash \langle P \rangle x.f := v \langle \_ . P \mathcal{U} (\Gamma Q * x.f \mapsto v^\Gamma) \rangle^C} \text{A-WRITE} \\
\frac{\Gamma; \Delta \mid P \vdash_r^m x.f \mapsto v \quad v \neq v_2}{\Gamma; \Delta \mid \Phi \vdash \langle P \rangle \mathbf{CAS}(x.f, v_1, v_2) \langle r. r = \mathbf{false} * P \rangle^C} \text{A-CAS-FALSE} \\
\frac{\Gamma; \Delta \mid P \vdash_w x.f \mapsto v, Q}{\Gamma; \Delta \mid \Phi \vdash \langle P \rangle \mathbf{CAS}(x.f, v', v) \langle r. r = \mathbf{true} * \Gamma Q * x.f \mapsto v'^\Gamma \rangle^C} \text{A-CAS-TRUE} \\
\Gamma, \Delta \mid \Phi \vdash \forall y, r. \text{stable}(Q(y, r)) \\
\Gamma, \Delta \mid \Phi \vdash n \in C \quad \Gamma, \Delta \mid \Phi \vdash \forall x \in X. \forall y \in f(x). (x, y) \in (T(A))^* \\
\frac{\Gamma \mid \Phi \vdash \forall x \in X. (\Delta). \langle P * \triangleright I(x) \rangle e \langle r. \exists y \in f(x). Q(y, r) * \otimes_{\alpha \in A} [\alpha]_{g(\alpha)}^n * \triangleright I(y) \rangle^{C \setminus \{n\}}}{\Gamma \mid \Phi \vdash (\Delta). \langle P * \text{region}(X, T, I, n) \rangle} \text{A-REGION} \\
e \\
\langle r. \exists y. Q(y, r) * \otimes_{\alpha \in A} [\alpha]_{g(\alpha)}^n * \text{region}(\{y\}, T, I, n) \rangle^C \\
\frac{\Gamma \mid \Phi \vdash (\Delta). \langle P \rangle e \langle Q \rangle^C \quad \text{atomic}(e)}{\Gamma \mid \Phi \vdash (\Delta). \{P\} e \{Q\}} \text{A-START}
\end{array}$$

### Structural Rules

$$\begin{array}{c}
\frac{\Gamma \mid \Phi \vdash P \sqsubseteq^A Q \quad \Gamma \mid \Phi \vdash \text{stable}(R)}{\Gamma \mid \Phi \vdash P * R \sqsubseteq^A Q * R} \text{V-FRAME} \\
\frac{\Gamma \mid \Phi \vdash (\Delta). \langle P \rangle \text{ e } \langle Q \rangle \quad \Gamma, \Delta \mid \Phi \vdash \text{stable}(R)}{\Gamma \mid \Phi \vdash (\Delta). \{P * \triangleright R\} \text{ e } \{Q * R\}} \text{A-FRAME} \\
\frac{\Gamma \mid \Phi \vdash (\Delta). \{P\} \text{ e } \{Q\} \quad \Gamma, \Delta \mid \Phi \vdash \text{stable}(R)}{\Gamma \mid \Phi \vdash (\Delta). \{P * R\} \text{ e } \{Q * R\}} \text{FRAME} \\
\frac{\Gamma \mid \Phi \vdash P_1 \sqsubseteq^A P_2 \quad \Gamma \mid \Phi \vdash (\Delta). \langle P_2 \rangle \text{ e } \langle Q_2 \rangle^A \quad \Gamma \mid \Phi \vdash Q_2 \sqsubseteq^A Q_1}{\Gamma \mid \Phi \vdash (\Delta). \langle P_1 \rangle \text{ e } \langle Q_1 \rangle^A} \text{A-CONS} \\
\frac{\Gamma \mid \Phi \vdash P_1 \sqsubseteq P_2 \quad \Gamma \mid \Phi \vdash (\Delta). \{P_2\} \text{ e } \{Q_2\} \quad \Gamma \mid \Phi \vdash Q_2 \sqsubseteq Q_1}{\Gamma \mid \Phi \vdash (\Delta). \{P_1\} \text{ e } \{Q_1\}} \text{CONS}
\end{array}$$

### Other Rules for Commands

$$\begin{array}{c}
\frac{\Gamma, \bar{y}, \mathbf{this} \mid \Phi \vdash \text{stable}(P) \wedge \text{stable}(Q) \quad \Gamma \mid \Phi \vdash \triangleright(C::m : (\bar{y}). \{P\} \{Q\}) \quad \Gamma \vdash C : \text{Class}}{\Gamma \mid \Phi(\Delta). \{P[\bar{v}/\bar{y}, x/\mathbf{this}] * x : C\} x.m(\bar{v}) \{Q[\bar{v}/\bar{y}, x/\mathbf{this}]\}} \text{CALL} \\
\frac{\Gamma, \mathbf{this} \mid \Phi \vdash \text{stable}(P) \wedge \text{stable}(Q) \quad \Gamma \mid \Phi \vdash \triangleright(C::m : (-). \{\ulcorner P \urcorner\} \{Q\}) \quad \Gamma \vdash C : \text{Class}}{\Gamma \mid \Phi \vdash (\Delta). \{\ulcorner P[x/\mathbf{this}] \urcorner * y : C\} \mathbf{fork}(y.m) \{\ulcorner \top \urcorner\}} \text{FORK} \\
\frac{\Gamma, \bar{x} \mid \Phi \vdash \text{stable}(P) \wedge \text{stable}(Q) \quad \Gamma \mid \Phi \vdash \triangleright(C : (\bar{x}). \{P\} \{Q\}) \quad \Gamma \vdash C : \text{Class}}{\Gamma \mid \Phi(\Delta). \{P[\bar{v}/\bar{x}]\} \mathbf{new} C(\bar{v}) \{Q[\bar{v}/\bar{x}]\}} \text{NEW} \\
\frac{\Gamma \mid \Phi \vdash (\Delta). \{P\} \text{ e}_1 \{r. Q(r)\} \quad \Gamma \mid \Phi \vdash (\Delta, x). \{Q(x)\} \text{ e}_2 \{r. R(r)\}}{\Gamma \mid \Phi \vdash (\Delta). \{P\} \mathbf{let} x = \text{e}_1 \mathbf{in} \text{e}_2 \{r. R(r)\}} \text{BIND} \\
\frac{\Gamma; - \vdash v : \text{Val}}{\Gamma \mid \Phi \vdash \{\ulcorner \top \urcorner\} v \{r. r = v\}} \text{VAL}
\end{array}$$

### View-Shifts

$$\begin{array}{c}
\frac{\Gamma \mid \Phi \vdash P \sqsubseteq^A Q \quad \Gamma \mid \Phi \vdash Q \sqsubseteq^A R}{\Gamma \mid \Phi \vdash P \sqsubseteq^A R} \text{VTRANS} \quad \frac{\Gamma \mid \Phi \mid P \vdash Q}{\Gamma \mid \Phi \vdash P \sqsubseteq^A Q} \text{VIMPL} \\
\frac{\Gamma \mid \Phi \vdash x \in X \quad \Gamma \mid \Phi \vdash \forall n \in C. \forall s. \text{up}(I(n)(s)) \quad \Gamma \mid \Phi \vdash A \text{ and } B \text{ are finite} \quad \Gamma \mid \Phi \vdash C \text{ is infinite} \quad \Gamma \mid \Phi \vdash \forall n \in C. P * \otimes_{\alpha \in A} [\alpha]_1^n \Rightarrow \triangleright I(n)(x)}{\Gamma \mid \Phi \vdash \forall n \in C. \forall s. \text{stable}(I(n)(s)) \quad \Gamma \mid \Phi \vdash A \cap B = \emptyset} \text{VALLOC} \\
\frac{\Gamma \mid \Phi \vdash P \sqsubseteq^C \exists n \in C. \mathbf{region}(X, T, I(n), n) * \otimes_{\alpha \in B} [\alpha]_1^n}{}
\end{array}$$

$$\frac{\begin{array}{c} \Gamma \mid \Phi \vdash \text{stable}(P) \quad \Gamma \mid \Phi \vdash \text{stable}(Q) \\ \Gamma \mid \Phi \vdash n \in A \quad \Gamma \mid \Phi \vdash \forall x \in X. f(x) \in Y \\ \Gamma \mid \Phi \vdash \forall x \in X. (x, f(x)) \in T(\alpha) \vee f(x) = x \\ \Gamma \mid \Phi \vdash \forall x \in X. P * \triangleright I(x) * [\alpha]_{\pi}^n \sqsubseteq^{A \setminus \{n\}} Q * \triangleright I(f(x)) \end{array}}{\Gamma \mid \Phi \vdash \text{region}(X, T, I, n) * P * [\alpha]_{\pi}^n \sqsubseteq^A \text{region}(Y, T, I, n) * Q} \text{VOPEN}$$

Derived Rules – Hoare triples for the SC Logic

$$\frac{\Gamma; - \vdash v : \text{Val} \quad x \in \Delta}{\Gamma \mid \Phi \vdash (\Delta).[x.f \mapsto v] \text{ x.f } [r. r = v * x.f \mapsto v]} \text{S-READ}$$

$$\frac{\Gamma; - \vdash v : \text{Val} \quad x \in \Delta}{\Gamma \mid \Phi \vdash (\Delta).[x.f \mapsto \_] \text{ x.f } := v [r. x.f \mapsto v]} \text{S-WRITE}$$

$$\frac{\Gamma, \bar{y}, \mathbf{this}, r \mid \Phi \vdash r\text{-stable}(P) \wedge r\text{-stable}(Q) \quad \Gamma \mid \Phi \vdash \triangleright (C::m : (\bar{y}). [P] [r. Q]) \quad \Gamma \vdash C : \text{Class}}{\Gamma \mid \Phi(\Delta).[P[\bar{v}/\bar{y}, x/\mathbf{this}] * x : C] \text{ x.m}(\bar{v}) [r. Q[\bar{v}/\bar{y}, x/\mathbf{this}]}] \text{S-CALL}$$

$$\frac{\Gamma, \bar{x}, r \mid \Phi \vdash \text{stable}(P) \wedge \text{stable}(Q) \quad \Gamma \mid \Phi \vdash \triangleright (C : (\bar{x}). [P] [r. Q]) \quad \Gamma \vdash C : \text{Class}}{\Gamma \mid \Phi(\Delta).[P[\bar{v}/\bar{x}]] \mathbf{new} C(\bar{v}) [r. Q[\bar{v}/\bar{x}]]} \text{S-NEW}$$

$$\frac{\Gamma \mid \Phi \vdash (\Delta).[P] e [r. Q] \quad \Gamma, \Delta \mid \Phi \vdash r\text{-stable}(R)}{\Gamma \mid \Phi \vdash (\Delta).[P * R] e [r. Q * R]} \text{S-FRAME}$$

$$\frac{\Gamma; \Delta \mid \Phi \mid P_1 \vdash_{SC} P_2 \quad \Gamma \mid \Phi \vdash (\Delta).[P_2] e [r. Q_2] \quad \Gamma, r; \Delta \mid \Phi \mid Q_2 \vdash_{SC} Q_1}{\Gamma \mid \Phi \vdash (\Delta).[P_1] e [r. Q_1]} \text{S-CONS}$$

$$\frac{\Gamma; \Delta \mid \Phi \mid P \vdash v = \mathbf{true} \quad \Gamma \mid \Phi \vdash (\Delta).[P] e_1 [r. Q]}{\Gamma \mid \Phi \vdash (\Delta).[P] \mathbf{if} v \mathbf{then} e_1 \mathbf{else} e_2 [r. Q]} \text{S-IF-T}$$

$$\frac{\Gamma; \Delta \mid \Phi \mid P \vdash v = \mathbf{false} \quad \Gamma \mid \Phi \vdash (\Delta).[P] e_2 [r. Q]}{\Gamma \mid \Phi \vdash (\Delta).[P] \mathbf{if} v \mathbf{then} e_1 \mathbf{else} e_2 [r. Q]} \text{S-IF-F}$$

$$\frac{\Gamma \mid \Phi \vdash (\Delta).[P] e_1 [r. Q(r)] \quad \Gamma \mid \Phi \vdash (\Delta, x).[Q(x)] e_2 [r. R(r)]}{\Gamma \mid \Phi \vdash (\Delta).[P] \mathbf{let} x = e_1 \mathbf{in} e_2 [r. R(r)]} \text{S-BIND}$$

$$\frac{\Gamma; - \vdash v : \text{Val}}{\Gamma \mid \Phi \vdash [\top] v [r. r = v]} \text{S-VAL} \quad \frac{\Gamma \mid \Phi \vdash (\Delta).\{\bar{P}\} e \{\bar{Q}\}}{\Gamma \mid \Phi \vdash (\Delta).[P] e [r. Q(r)]} \text{S-SHIFT}$$

### 3.4.2 Later operator

In the following proof rules, the rules for distribution over quantifiers and binary connectives work for both SC- and TSO-level connectives (except for  $\mathcal{U}$  operator, which is only defined on the TSO level), hence we omit the subscripts on the entailments.



$$\begin{array}{c}
\frac{\Gamma \mid \Phi, \triangleright S \vdash S}{\Gamma \mid \Phi \vdash S} \text{SLOB} \quad \frac{}{\Gamma \mid \Phi \vdash S \Rightarrow \triangleright S} \text{SMONO} \\
\frac{op \in \{\wedge, \vee, *, \mathcal{U}_t\}}{\Gamma; \Delta \mid \Phi \mid \triangleright(P \text{ op } Q) \dashv\vdash (\triangleright P) \text{ op } (\triangleright Q)} \text{LBIN} \quad \frac{E \in \{\ulcorner - \urcorner, \ulcorner - \text{ in } t \urcorner\}}{\Gamma; \Delta \mid \Phi \mid \triangleright(E(P)) \dashv\vdash E(\triangleright P)} \text{LEMB} \\
\frac{}{\Gamma; \Delta \mid \Phi \mid \triangleright(P \Rightarrow Q) \vdash (\triangleright P) \Rightarrow (\triangleright Q)} \text{LIMPL} \quad \frac{}{\Gamma; \Delta \mid \Phi \mid \triangleright(P * Q) \vdash (\triangleright P) * (\triangleright Q)} \text{LWAND} \\
\frac{Q \in \{\forall, \exists\}}{\Gamma; \Delta \mid \Phi \mid \triangleright(Qx : \tau. P(x)) \dashv\vdash Qx : \tau. \triangleright P(x)} \text{LQUANT}
\end{array}$$

### 3.4.3 Reasoning about stability

Stability consists of two components: stability under region interference (**r-stable**) and stability under store-buffer interference (**b-stable**). The first of these is defined for both level of assertions, the second only for  $\text{Prop}_{\text{TSO}}$ . Note how both embeddings preserve stability under regions and *grant* stability under store-buffer interference, and that the explicit stabilization does indeed provide stability.

$$\begin{array}{c}
\frac{\Gamma \mid \Phi \vdash \forall \alpha \notin A. \forall x \in X. T(\alpha)(x) \subseteq X}{\Gamma \mid \Phi \vdash \text{r-stable}(\text{region}(X, T, n) * \otimes_{\alpha \in A} [\alpha]_1^n)} \\
\frac{\Gamma \vdash l : \text{Ald} \rightarrow \mathcal{P}(\text{Sld} \times \text{Sld}) \quad \Gamma \vdash n : \text{Rld}}{\Gamma \mid \Phi \vdash \text{r-stable}(\text{rintr}(l, n))} \\
\frac{c \in \{\top, \perp, \text{emp}\}}{\Gamma \mid \Phi \vdash \text{stable}(c)} \quad \frac{}{\Gamma \mid \Phi \vdash \text{stable}(x.f \mapsto v)} \quad \frac{}{\Gamma \mid \Phi \vdash \text{stable}(x : C)} \\
\frac{}{\Gamma \mid \Phi \vdash \text{stable}(\llbracket P \rrbracket)} \quad \frac{}{\Gamma \mid \Phi \vdash \text{b-stable}(\ulcorner P \urcorner)} \quad \frac{}{\Gamma \mid \Phi \vdash \text{b-stable}(\ulcorner P \text{ in } t \urcorner)} \\
\frac{\Gamma \mid \Phi \vdash \text{r-stable}(P)}{\Gamma \mid \Phi \vdash \text{r-stable}(\ulcorner P \urcorner)} \quad \frac{\Gamma \mid \Phi \vdash \text{r-stable}(P)}{\Gamma \mid \Phi \vdash \text{r-stable}(\ulcorner P \text{ in } t \urcorner)}
\end{array}$$

$$\begin{array}{c}
\frac{\Gamma \mid \Phi \vdash \text{b-stable}(P) \quad \Gamma \mid \Phi \vdash \text{b-stable}(Q) \quad op \in \{\wedge, \vee, *\}}{\Gamma \mid \Phi \vdash \text{b-stable}(P \text{ op } Q)} \\
\frac{\Gamma \mid \Phi \vdash \forall x : \tau. \text{b-stable}(P(x)) \quad \Gamma \mid \Phi \vdash \forall x : \tau. \text{b-stable}(P(x))}{\Gamma \mid \Phi \vdash \text{b-stable}(\forall x : \tau. P(x))} \quad \frac{\Gamma \mid \Phi \vdash \forall x : \tau. \text{b-stable}(P(x))}{\Gamma \mid \Phi \vdash \text{b-stable}(\exists x : \tau. P(x))} \\
\frac{\Gamma \mid \Phi \vdash \text{r-stable}(P) \quad \Gamma \mid \Phi \vdash \text{r-stable}(Q) \quad op \in \{\wedge, \vee, *\}}{\Gamma \mid \Phi \vdash \text{r-stable}(P \text{ op } Q)} \\
\frac{\Gamma \mid \Phi \vdash \forall x : \tau. \text{r-stable}(P(x)) \quad \Gamma \mid \Phi \vdash \forall x : \tau. \text{r-stable}(P(x))}{\Gamma \mid \Phi \vdash \text{r-stable}(\forall x : \tau. P(x))} \quad \frac{\Gamma \mid \Phi \vdash \forall x : \tau. \text{r-stable}(P(x))}{\Gamma \mid \Phi \vdash \text{r-stable}(\exists x : \tau. P(x))}
\end{array}$$

### 3.4.4 Method verification

To verify a method or a constructor, we verify its body:

$$\frac{\Gamma; \Delta, \mathbf{this} \vdash P : \text{Tld} \rightarrow \text{Prop}_{\text{Tso}} \quad \Gamma; \Delta, \mathbf{this} \vdash Q : \text{Tld} \rightarrow \text{Val} \rightarrow \text{Prop}_{\text{Tso}} \quad \text{body}(C, m) = (\text{T } m(\Delta) = e) \quad \Gamma \mid \Phi \vdash (\Delta, \mathbf{this}).\{P * \mathbf{this} : C\} e \{Q\}}{\Gamma \mid \Phi \vdash C::m : (\Delta). \{P\} \{Q\}} \\
\text{ctorBody}(C) = (C(\Delta) = e) \\
\frac{\Gamma; \Delta \vdash P : \text{Tld} \rightarrow \text{Prop}_{\text{Tso}} \quad \Gamma; \Delta \vdash Q : \text{Tld} \rightarrow \text{Val} \rightarrow \text{Prop}_{\text{Tso}} \quad \Gamma \mid \Phi \vdash (\Delta, \mathbf{this}).\{P * \otimes_{f \in \text{fields}(C)} \ulcorner \mathbf{this}.f \mapsto \text{null} \urcorner * \mathbf{this} : C\} e \{Q\}}{\Gamma \mid \Phi \vdash C : (\Delta). \{P\} \{Q\}}$$

Note that constructors are syntactically required to return the value of the newly allocated object.

## 4 Interpretation

### 4.1 Types

$$\begin{aligned}
\llbracket \vdash 1 : \text{Type} \rrbracket &= 1 \\
\llbracket \vdash \tau \rightarrow \sigma : \text{Type} \rrbracket &= \llbracket \vdash \tau : \text{Type} \rrbracket \rightarrow \llbracket \vdash \sigma : \text{Type} \rrbracket \\
\llbracket \vdash \tau \times \sigma : \text{Type} \rrbracket &= \llbracket \vdash \tau : \text{Type} \rrbracket \times \llbracket \vdash \sigma : \text{Type} \rrbracket \\
\llbracket \vdash \tau + \sigma : \text{Type} \rrbracket &= \llbracket \vdash \tau : \text{Type} \rrbracket + \llbracket \vdash \sigma : \text{Type} \rrbracket \\
\llbracket \vdash \mathcal{P}(\tau) : \text{Type} \rrbracket &= \mathcal{P}(\llbracket \vdash \tau : \text{Type} \rrbracket) \\
\llbracket \vdash \text{Prop}_{\text{Tso}} : \text{Type} \rrbracket &= \text{Prop}_{\text{Tso}} \\
\llbracket \vdash \text{Prop}_{\text{sc}} : \text{Type} \rrbracket &= \text{Prop}_{\text{sc}} \\
\llbracket \vdash \text{Spec} : \text{Type} \rrbracket &= \Omega \\
\llbracket \vdash \Delta(X) : \text{Type} \rrbracket &= \Delta(X)
\end{aligned}$$

## 4.2 Contexts

$$\begin{aligned} \llbracket \Gamma, x : \tau \rrbracket &= \llbracket \Gamma \rrbracket \times \llbracket \vdash \tau : \text{Type} \rrbracket & \llbracket \varepsilon \rrbracket &= 1 \\ \llbracket \Delta, x : \text{Val} \rrbracket &= \llbracket \Delta \rrbracket \times \llbracket \vdash \text{Val} : \text{Type} \rrbracket & \llbracket \varepsilon \rrbracket &= 1 \end{aligned}$$

## 4.3 Lambda calculus

$$\begin{aligned} \llbracket \Gamma; \Delta \vdash x : \tau \rrbracket(\gamma, \delta) &= \pi_x(\vartheta) \\ \llbracket \Gamma; \Delta \vdash x : \text{Val} \rrbracket(\gamma, \delta) &= \pi_x(\delta) \\ \llbracket \Gamma; \Delta \vdash \lambda x : \tau. M : \tau \rightarrow \sigma \rrbracket(\gamma, \delta) &= \lambda v \in \llbracket \vdash \tau : \text{Type} \rrbracket. \llbracket \Gamma, x : \tau; \Delta \vdash M : \sigma \rrbracket((\vartheta, v), \theta) \\ \llbracket \Gamma; \Delta \vdash M N : \sigma \rrbracket(\gamma, \delta) &= (\llbracket \Gamma; \Delta \vdash M : \tau \rightarrow \sigma \rrbracket(\gamma, \delta))(\llbracket \Gamma; \Delta \vdash N : \tau \rrbracket(\gamma, \delta)) \end{aligned}$$

## 4.4 TSO Assertion Logic

$$\begin{aligned} \llbracket \Gamma; \Delta \vdash \perp : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) &= \emptyset \\ \llbracket \Gamma; \Delta \vdash \top : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) &= \mathcal{M} \\ \llbracket \Gamma; \Delta \vdash P \wedge Q : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) &= \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) \cap \llbracket \Gamma; \Delta \vdash Q : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) \\ \llbracket \Gamma; \Delta \vdash P \vee Q : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) &= \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) \cup \llbracket \Gamma; \Delta \vdash Q : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) \\ \llbracket \Gamma; \Delta \vdash P \Rightarrow Q : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) &= \{m \in \mathcal{M} \mid \forall n \geq m. \\ &\quad n \in \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) \Rightarrow \\ &\quad n \in \llbracket \Gamma; \Delta \vdash Q : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta)\} \\ \llbracket \Gamma; \Delta \vdash \forall x : \tau. P : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) &= \bigcap_{v \in \llbracket \vdash \tau : \text{Type} \rrbracket} \llbracket \Gamma, x : \tau; \Delta \vdash P : \text{Prop}_{\text{TSO}} \rrbracket((\gamma, v), \delta) \\ \llbracket \Gamma; \Delta \vdash \exists x : \tau. P : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) &= \bigcup_{v \in \llbracket \vdash \tau : \text{Type} \rrbracket} \llbracket \Gamma, x : \tau; \Delta \vdash P : \text{Prop}_{\text{TSO}} \rrbracket((\gamma, v), \delta) \\ \llbracket \Gamma; \Delta \vdash \ulcorner P \urcorner : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) &\stackrel{\text{def}}{=} \text{let } p = \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) \text{ in } \ulcorner p \urcorner \\ \llbracket \Gamma; \Delta \vdash \ulcorner P \text{ in } t \urcorner : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) &\stackrel{\text{def}}{=} \text{let } p = \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) \text{ in } \ulcorner p \text{ in } t \urcorner \\ \llbracket \Gamma; \Delta \vdash P \mathcal{U}_t Q : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) &\stackrel{\text{def}}{=} \text{let } p = \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) \text{ in} \\ &\quad \text{let } q = \llbracket \Gamma; \Delta \vdash Q : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) \text{ in} \\ &\quad p \mathcal{U}_t q \\ \llbracket \Gamma; \Delta \vdash \langle P \rangle : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) &\stackrel{\text{def}}{=} \text{let } p = \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) \text{ in } \text{cl}(p) \\ \llbracket \Gamma; \Delta \vdash \text{emp} : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) &= \mathcal{M} \\ \llbracket \Gamma; \Delta \vdash P * Q : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) &= \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) * \llbracket \Gamma; \Delta \vdash Q : \text{Prop}_{\text{TSO}} \rrbracket(\gamma, \delta) \end{aligned}$$

## 4.5 SC Assertion Logic

$$\begin{aligned}
\llbracket \Gamma; \Delta \vdash \perp : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) &= \emptyset \\
\llbracket \Gamma; \Delta \vdash \top : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) &= \mathcal{H} \\
\llbracket \Gamma; \Delta \vdash P \wedge Q : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) &= \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) \cap \llbracket \Gamma; \Delta \vdash Q : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) \\
\llbracket \Gamma; \Delta \vdash P \vee Q : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) &= \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) \cup \llbracket \Gamma; \Delta \vdash Q : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) \\
\llbracket \Gamma; \Delta \vdash P \Rightarrow Q : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) &= \{m \in \mathcal{H} \mid \forall n \geq m. \\
&\quad n \in \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) \Rightarrow \\
&\quad n \in \llbracket \Gamma; \Delta \vdash Q : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta)\} \\
\llbracket \Gamma; \Delta \vdash \forall x : \tau. P : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) &= \bigcap_{v \in \llbracket \vdash \tau : \text{Type} \rrbracket} \llbracket \Gamma, x : \tau; \Delta \vdash P : \text{Prop}_{\text{SC}} \rrbracket((\gamma, v), \delta) \\
\llbracket \Gamma; \Delta \vdash \exists x : \tau. P : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) &= \bigcup_{v \in \llbracket \vdash \tau : \text{Type} \rrbracket} \llbracket \Gamma, x : \tau; \Delta \vdash P : \text{Prop}_{\text{SC}} \rrbracket((\gamma, v), \delta) \\
\llbracket \Gamma; \Delta \vdash \text{emp} : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) &= \mathcal{H} \\
\llbracket \Gamma; \Delta \vdash P * Q : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) &= \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) * \llbracket \Gamma; \Delta \vdash Q : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) \\
\llbracket \Gamma; \Delta \vdash x.f \mapsto v : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) &= \text{let } o = \llbracket \Gamma; \Delta \vdash x : \text{Val} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \text{let } f = \llbracket \Gamma; \Delta \vdash f : \text{Field} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \text{let } v = \llbracket \Gamma; \Delta \vdash v : \text{Val} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \{m \in \mathcal{H} \mid m.l(o, f) = v\} \\
\llbracket \Gamma; \Delta \vdash \text{region}(M, N, R) : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) &= \\
&\quad \{m \in \mathcal{M} \mid \exists T : \Delta(\text{AId} \rightarrow \mathcal{P}(\text{SId} \times \text{SId})) \\
&\quad \quad \text{let } M = \llbracket \Gamma; \Delta \vdash M : \mathcal{P}(\text{SId}) \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \quad \text{let } R = \llbracket \Gamma; \Delta \vdash R : \text{RId} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \quad \phi(T) = \llbracket \Gamma; \Delta \vdash T : \text{AId} \rightarrow \mathcal{P}(\text{SId} \times \text{SId}) \rrbracket(\gamma, \delta) \wedge \\
&\quad \quad m \in \text{region}(M, T, R)\} \\
\llbracket \Gamma; \Delta \vdash \text{rintr}(I, R) : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) &= \text{rintr}(\llbracket \Gamma; \Delta \vdash I : \text{SId} \rightarrow \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \delta), \llbracket \Gamma; \Delta \vdash R : \text{RId} \rrbracket(\gamma, \delta)) \\
\llbracket \Gamma; \Delta \vdash [A]_P^R : \text{Prop}_{\text{SC}} \rrbracket(\gamma, \delta) &= \text{action}(\llbracket \Gamma; \Delta \vdash A : \text{AId} \rrbracket(\gamma, \delta), \\
&\quad \llbracket \Gamma; \Delta \vdash R : \text{RId} \rrbracket(\gamma, \delta), \\
&\quad \llbracket \Gamma; \Delta \vdash P : \text{Perm} \rrbracket(\gamma, \delta))
\end{aligned}$$

where  $\phi$  embeds predicates from our set-theoretic metalogic in the topos of trees. For more details, consult the iCAP technical report [2].

## 4.6 Specification Logic

$$\begin{aligned}
\llbracket \Gamma \vdash \perp : \text{Spec} \rrbracket(\gamma) &= \perp \\
\llbracket \Gamma \vdash \top : \text{Spec} \rrbracket(\gamma) &= \top \\
\llbracket \Gamma \vdash S \wedge T : \text{Spec} \rrbracket(\gamma) &= \llbracket \Gamma \vdash S : \text{Spec} \rrbracket(\gamma) \wedge \llbracket \Gamma \vdash T : \text{Spec} \rrbracket(\gamma) \\
\llbracket \Gamma \vdash S \vee T : \text{Spec} \rrbracket(\gamma) &= \llbracket \Gamma \vdash S : \text{Spec} \rrbracket(\gamma) \vee \llbracket \Gamma \vdash T : \text{Spec} \rrbracket(\gamma) \\
\llbracket \Gamma \vdash S \Rightarrow T : \text{Spec} \rrbracket(\gamma) &= \llbracket \Gamma \vdash S : \text{Spec} \rrbracket(\gamma) \Rightarrow \llbracket \Gamma \vdash T : \text{Spec} \rrbracket(\gamma)
\end{aligned}$$

$$\begin{aligned}
\llbracket \Gamma \vdash \forall x : \tau. S : \text{Spec} \rrbracket(\gamma) &= \forall v \in \llbracket \vdash \tau : \text{Type} \rrbracket. \llbracket \Gamma, x : \tau \vdash S : \text{Spec} \rrbracket((\gamma, v)) \\
\llbracket \Gamma \vdash \exists x : \tau. S : \text{Spec} \rrbracket(\gamma) &= \exists v \in \llbracket \vdash \tau : \text{Type} \rrbracket. \llbracket \Gamma, x : \tau \vdash S : \text{Spec} \rrbracket((\gamma, v)) \\
\llbracket \Gamma \vdash M =_{\tau} N : \text{Spec} \rrbracket(\gamma) &= \llbracket \Gamma \vdash M : \tau \rrbracket(\gamma) = \llbracket \Gamma \vdash N : \tau \rrbracket(\gamma)
\end{aligned}$$

$$\begin{aligned}
\llbracket \Gamma \vdash (\Delta). \{P\} e \{Q\} \rrbracket(\gamma) &\stackrel{\text{def}}{=} \forall t \in \text{TId}. \forall \delta \in \llbracket \Delta \rrbracket. \\
&\quad \text{let } p = \llbracket \Gamma; \Delta \vdash P : \text{TId} \rightarrow \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \text{let } q = \llbracket \Gamma; \Delta \vdash Q : \text{TId} \rightarrow \text{Val} \rightarrow \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \text{safe}((t, \delta(e)), p(t), q(t))
\end{aligned}$$

$$\begin{aligned}
\llbracket \Gamma \vdash (\Delta). \langle P \rangle e \langle Q \rangle^A \rrbracket(\gamma) &\stackrel{\text{def}}{=} \forall t \in \text{TId}. \forall \delta \in \llbracket \Delta \rrbracket. \forall a \in \text{Act}. \forall v \in \text{Val}. \\
&\quad \text{let } p = \llbracket \Gamma; \Delta \vdash P : \text{TId} \rightarrow \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \text{let } q = \llbracket \Gamma; \Delta \vdash Q : \text{TId} \rightarrow \text{Val} \rightarrow \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad (t, \delta(e)) \xrightarrow{a} \{(t, v)\} \Rightarrow a \text{ sat}_A \{p(t)\} \{q(t)(v)\}
\end{aligned}$$

$$\begin{aligned}
\llbracket \Gamma \vdash r\text{-stable}(P) : \text{Spec} \rrbracket(\gamma) &= \text{let } p = \llbracket \Gamma; \varepsilon \vdash P : \text{Prop}_{\text{sc}} \rrbracket(\gamma, \varepsilon) \text{ in } r\text{stable}(p) \\
\llbracket \Gamma \vdash b\text{-stable}(P) : \text{Spec} \rrbracket(\gamma) &= \text{let } p = \llbracket \Gamma; \varepsilon \vdash P : \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \varepsilon) \text{ in } b\text{stable}(p) \\
\llbracket \Gamma \vdash P \sqsubseteq^R Q : \text{Spec} \rrbracket(\gamma) &= \text{let } p = \llbracket \Gamma; \varepsilon \vdash P : \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \varepsilon) \text{ in} \\
&\quad \text{let } q = \llbracket \Gamma; \varepsilon \vdash Q : \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \varepsilon) \text{ in} \\
&\quad \text{let } R = \llbracket \Gamma \vdash R : \mathcal{P}(\text{RId}) \rrbracket(\gamma) \text{ in} \\
&\quad p \sqsubseteq_R q
\end{aligned}$$

## 4.7 Judgments

$$\begin{aligned}
\llbracket \Gamma; \Delta \mid P \vdash_w x.f \mapsto v, Q \rrbracket(\gamma, \delta) &= \text{let } p = \llbracket \Gamma; \Delta \vdash P : TId \rightarrow \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \text{let } q = \llbracket \Gamma; \Delta \vdash Q : \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \text{let } o = \llbracket \Gamma; \Delta \vdash x : \text{Val} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \text{let } v = \llbracket \Gamma; \Delta \vdash v : \text{Val} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad p \vdash_w o.f \mapsto v, q \\
\llbracket \Gamma; \Delta \mid P \vdash_r^m x.f \mapsto v \rrbracket(\gamma, \delta) &= \text{let } p = \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \text{let } o = \llbracket \Gamma; \Delta \vdash x : \text{Val} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \text{let } v = \llbracket \Gamma; \Delta \vdash v : \text{Val} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad p \vdash_r^m o.f \mapsto v \\
\llbracket \Gamma; \Delta \mid p \vdash_r^o x.f \mapsto v \rrbracket(\gamma, \delta) &= \text{let } p = \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \text{let } o = \llbracket \Gamma; \Delta \vdash x : \text{Val} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad \text{let } v = \llbracket \Gamma; \Delta \vdash v : \text{Val} \rrbracket(\gamma, \delta) \text{ in} \\
&\quad p \vdash_r^o o.f \mapsto v
\end{aligned}$$

## 4.8 Embeddings and later operators

$$\begin{aligned}
\llbracket \Gamma \vdash \text{valid}(P) : \text{Spec} \rrbracket(\gamma) &= \text{valid}(\llbracket \Gamma; - \vdash P : \text{Prop}_{\text{Tso}} \rrbracket(\gamma)) \\
\llbracket \Gamma \vdash \triangleright S : \text{Spec} \rrbracket(\gamma) &= \triangleright(\llbracket \Gamma \vdash S : \text{Spec} \rrbracket(\gamma))
\end{aligned}$$

$$\begin{aligned}
\llbracket \Gamma; \Delta \vdash \text{spec}(S) : \text{Prop}_{\text{sc}} \rrbracket(\gamma, \delta) &= \text{spec}(\llbracket \Gamma \vdash S : \text{Spec} \rrbracket(\gamma)) \\
\llbracket \Gamma; \Delta \vdash \triangleright P : \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \delta) &= \triangleright(\llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \delta)) \\
\llbracket \Gamma; \Delta \vdash \triangleright P : \text{Prop}_{\text{sc}} \rrbracket(\gamma, \delta) &= \triangleright(\llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{sc}} \rrbracket(\gamma, \delta))
\end{aligned}$$

## 4.9 Entailment

$$\begin{aligned}
\llbracket \Gamma; \Delta \mid \Phi \mid P \vdash_{\text{Tso}} Q \rrbracket &= \\
\forall \gamma \in \llbracket \Gamma \rrbracket. \forall \delta \in \llbracket \Delta \rrbracket. \llbracket \Phi \rrbracket(\gamma) \Rightarrow \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \delta) \subseteq \llbracket \Gamma; \Delta \vdash Q : \text{Prop}_{\text{Tso}} \rrbracket(\gamma, \delta)
\end{aligned}$$

$$\begin{aligned}
\llbracket \Gamma; \Delta \mid \Phi \mid P \vdash_{\text{sc}} Q \rrbracket &= \\
\forall \gamma \in \llbracket \Gamma \rrbracket. \forall \delta \in \llbracket \Delta \rrbracket. \llbracket \Phi \rrbracket(\gamma) \Rightarrow \llbracket \Gamma; \Delta \vdash P : \text{Prop}_{\text{sc}} \rrbracket(\gamma, \delta) \subseteq \llbracket \Gamma; \Delta \vdash Q : \text{Prop}_{\text{sc}} \rrbracket(\gamma, \delta)
\end{aligned}$$

$$\llbracket \Gamma \mid S_1, \dots, S_n \vdash T \rrbracket = \forall \gamma \in \llbracket \Gamma \rrbracket. \left( \bigwedge_{i \in \{1, \dots, n\}} \llbracket \Gamma \vdash S_i : \text{Spec} \rrbracket(\gamma) \right) \Rightarrow \llbracket \Gamma \vdash T : \text{Spec} \rrbracket(\gamma)$$

## 5 Soundness

**Lemma 32.**

$$\forall p \in Prop_{SC}. \triangleright \lceil p \rceil = \lceil \triangleright p \rceil$$

*Proof* ( $\subseteq$ ). Assume  $\triangleright(l, s, U, \zeta) \in \lceil p \rceil$ . Then by definition of the embedding,

$$\triangleright(\exists l' \in \Delta(LState). l' \leq l \wedge (l', s, U, \zeta) \in p \wedge \text{lfid}(l', U))$$

Since later commutes over existentials over constant sets and conjunctions and  $l' \leq l$  and  $\text{lfid}$  simply reduces to (universally quantified) equalities on constant sets, which are upwards-closed, the above is equivalent to:

$$\exists l' \in \Delta(LState). (l' \leq l \vee \triangleright \perp) \wedge (l', s, U, \zeta) \in \triangleright p \wedge (\text{lfid}(l', U) \vee \triangleright \perp)$$

In case  $\triangleright \perp$  then  $(l, s, U, \zeta) \in \lceil \triangleright p \rceil$  holds trivially, by taking  $l'$  to be the empty local state, which is trivially smaller than  $l$ . Otherwise, the conclusion follows directly from the assumptions.  $\square$

*Proof* ( $\supseteq$ ). Follows easily by monotonicity of  $\triangleright$ .  $\square$

**Lemma 33.**

$$\forall p, q \in Prop_{TSO}. \triangleright (p \mathcal{U}_t q) = (\triangleright p) \mathcal{U}_t (\triangleright q)$$

*Proof* ( $\subseteq$ ). Assume  $(l, s, U, \zeta) \in \triangleright (p \mathcal{U}_t q)$ . Hence, by the definition of  $\mathcal{U}_t$

$$\begin{aligned} &\triangleright(\exists \alpha, \beta \in \Delta(SBuffer). \exists o \in \Delta(OId). \exists f \in \Delta(FName). \exists v \in \Delta(Val). \\ &U(t) = \alpha \cdot (o, f, v) \cdot \beta \wedge (l, s, U[t \mapsto \alpha], \zeta) \in p \wedge \\ &(\text{flush}(l, \alpha \cdot (o, f, v)), s, U[t \mapsto \beta], \zeta) \in q) \end{aligned}$$

Since constant sets are trivially total and  $\triangleright$  commutes over  $\wedge$  this reduces to

$$\begin{aligned} &\exists \alpha, \beta \in \Delta(SBuffer). \exists o \in \Delta(OId). \exists f \in \Delta(FName). \exists v \in \Delta(Val). \\ &\triangleright (U(t) = \alpha \cdot (o, f, v) \cdot \beta \wedge (l, s, U[t \mapsto \alpha], \zeta) \in \triangleright p \wedge \\ &(\text{flush}(l, \alpha \cdot (o, f, v)), s, U[t \mapsto \beta], \zeta) \in \triangleright q) \end{aligned}$$

Since equality on constant sets is upwards-closed, this is equivalent to

$$\begin{aligned} &\exists \alpha, \beta \in \Delta(SBuffer). \exists o \in \Delta(OId). \exists f \in \Delta(FName). \exists v \in \Delta(Val). \\ &U(t) = \alpha \cdot (o, f, v) \cdot \beta \wedge (l, s, U[t \mapsto \alpha], \zeta) \in \triangleright p \wedge \\ &(\text{flush}(l, \alpha \cdot (o, f, v)), s, U[t \mapsto \beta], \zeta) \in \triangleright q \end{aligned}$$

Hence  $(l, s, U, \zeta) \in (\triangleright p) \mathcal{U}_t (\triangleright q)$ .  $\square$

*Proof* ( $\supseteq$ ). Assume  $(l, s, U, \zeta) \in (\triangleright p) \mathcal{U}_t (\triangleright q)$ . Then

$$\begin{aligned} & \exists \alpha, \beta \in \Delta(\text{SBuffer}). \exists o \in \Delta(\text{OId}). \exists f \in \Delta(\text{FName}). \exists v \in \Delta(\text{Val}). \\ & U(t) = \alpha \cdot (o, f, v) \cdot \beta \wedge (l, s, U[t \mapsto \alpha], \zeta) \in \triangleright p \wedge \\ & (\text{flush}(l, \alpha \cdot (o, f, v)), s, U[t \mapsto \beta], \zeta) \in \triangleright q \end{aligned}$$

by monotonicity of  $\triangleright$  and by commuting  $\triangleright$  over  $\wedge$  and  $\exists$ , it follows that

$$\begin{aligned} & \triangleright(\exists \alpha, \beta \in \Delta(\text{SBuffer}). \exists o \in \Delta(\text{OId}). \exists f \in \Delta(\text{FName}). \exists v \in \Delta(\text{Val}). \\ & U(t) = \alpha \cdot (o, f, v) \cdot \beta \wedge (l, s, U[t \mapsto \alpha], \zeta) \in p \wedge \\ & (\text{flush}(l, \alpha \cdot (o, f, v)), s, U[t \mapsto \beta], \zeta) \in q) \end{aligned}$$

Thus  $(l, s, U, \zeta) \in \triangleright(p \mathcal{U}_t q)$ . □

**Lemma 34.** For any  $A \in \mathcal{P}(\Delta(\text{RId}))$ ,  $p \in \Delta(\text{TId}) \rightarrow \text{Prop}_{\text{TSO}}$ ,  $f \in \Delta(\text{FName})$ ,  $v_1, v_2 \in \Delta(\text{Val})$ ,  $t \in \Delta(\text{TId})$  and  $o \in \Delta(\text{OId})$ ,

$$p \vdash_r^m o.f \mapsto v_1 \Rightarrow \text{read}(t, o, f, v_2) \text{ sat}_A \{p(t)\} \{p(t) * v_1 = v_2\}$$

*Proof.* Let  $r \in \text{Prop}_{\text{TSO}}$ ,  $m \in \mathcal{M}$ ,  $h \in \Delta(\text{Heap})$ , and  $U \in \Delta(\text{SPool})$  such that

$$m \in p(t) * \triangleright r \qquad (h, U) \in [m]_A \qquad \text{stable}(r)_A.$$

Hence, there exists  $m_1, m_2 \in \mathcal{M}$  such that  $m = m_1 \bullet m_2$ ,  $m_1 \in p(t)$  and  $m_2 \in \triangleright r$ . If  $t \notin \text{dom}(U)$ ,  $\llbracket \text{read}(t, o, f, v_2) \rrbracket(h, U) = \emptyset$  and the conclusion holds vacuously. Assume, then, that  $t \in \text{dom}(U)$ . Now, from the definition of  $\vdash_r^m$  it follows that

$$\triangleright \text{lookup}(o.f, m_1.U(t), m_1.l) = v_1.$$

Since equalities on terms of constant sets are upwards-closed it follows that

$$\text{lookup}(o.f, m_1.U(t), m_1.l) = v_1 \vee \triangleright \perp$$

The  $\triangleright \perp$  case follows trivially as the definition of atomic satisfaction only requires the post-condition to hold later. In the  $\text{lookup}(o.f, m_1.U(t), m_1.l) = v_1$  case,  $o.f \in \text{dom } h$ , and so  $\not\vdash \notin \llbracket \text{read}(t, o, f, v_2) \rrbracket(h, U)$ . Now we can take any  $(h', U') \in \llbracket \text{read}(t, o, f, v_2) \rrbracket(h, U)$ . By definition, it follows that  $h = h'$ ,  $U = U'$ , and  $v_1 = v_2$ . We pick  $m'$  to be  $m$  and thus have to show that

$$\triangleright(m \in p(t) * v_1 = v_1 * r) \qquad \triangleright((h, U) \in [m]_A)$$

Both follow easily from monotonicity of  $\triangleright$  and the fact that  $\varepsilon \in v_1 = v_1$ . □

**Lemma 35.** For any  $A \in \mathcal{P}(\Delta(\text{RId}))$ ,  $p \in \Delta(\text{TId}) \rightarrow \text{Prop}_{\text{TSO}}$ ,  $q \in \Delta(\text{TId}) \rightarrow \text{Prop}_{\text{SC}}$ ,  $f \in \Delta(\text{FName})$ ,  $v_1, v_2 \in \Delta(\text{Val})$ ,  $t \in \Delta(\text{TId})$  and  $o \in \Delta(\text{OId})$ ,

$$p \vdash_w o.f \mapsto v_2, q \Rightarrow \text{write}(t, o, f, v_1) \text{ sat}_A \{p(t)\} \{p(t) \mathcal{U}_t (\ulcorner q(t) * o.f \mapsto v_1 \urcorner)\}$$



*Proof.* Let  $r \in Prop_{TSO}$ ,  $m \in \mathcal{M}$ ,  $h \in \Delta(Heap)$  and  $U \in \Delta(SPool)$  such that

$$m \in p(t) * \triangleright r \qquad (h, U) \in [m]_A \qquad \text{stable}(r)$$

Hence, there exists  $m_1, m_2 \in \mathcal{M}$  such that  $m = m_1 \bullet m_2$ ,  $m_1 \in p(t)$  and  $m_2 \in \triangleright r$ . Let

$$m = (l, s, U, \zeta) \qquad m_1 = (l_1, s, U, \zeta) \qquad m_2 = (l_2, s, U, \zeta)$$

Analogously to the read case, we only need to consider the case when  $t \in \text{dom}(U)$ . From the definition of  $\vdash_w$  it follows that

$$\triangleright \text{lookup}(o.f, U(t), l_1) = v_2 \quad \triangleright (\text{flush}(l_1, U(t)), s_1, U[t \mapsto \varepsilon], \zeta_1) \in \ulcorner q(t) * o.f \mapsto v_2 \urcorner.$$

Since equalities on constant sets are upwards-closed, it follows that

$$\text{lookup}(o.f, U(t), l_1) = v_2 \vee \triangleright \perp$$

As the  $\triangleright \perp$  case follows easily from Lemma 15, assume  $\text{lookup}(o.f, U(t), l_1) = v_2$ . Then  $(o, f) \in \text{dom}(h)$  and thus  $\not\vdash \llbracket \text{write}(t, o, f, v_1) \rrbracket (h, U)$ . Now, we can take any  $(h', U') \in \llbracket \text{write}(t, o, f, v_1) \rrbracket (h, U)$ . By definition of action semantics, this means that  $h' = h$  and  $U' = U[t \mapsto U(t) \cdot (o, f, v_1)]$ .

Take  $m' = (l, s, U', \zeta)$ . Since  $(o, f) \in \text{dom}(l_1)$ ,  $\triangleright ((l_2, s, U', \zeta) \in r)$  and  $(h, U') \in [m']_A$  follow by stability of  $r$  and interpretations of the shared regions. Hence, it suffices to show that  $\triangleright (l_1, s, U', \zeta) \in p(t) \mathcal{U}_t (\ulcorner q(t) * o.f \mapsto v_1 \urcorner)$ . To show this, we pick the final update in  $U'(t)$ , and thus need to show

$$\triangleright (l_1, s, U, \zeta) \in p(t) \quad \triangleright (\text{flush}(l_1, U(t) \cdot (o, f, v_1)), s, U[t \mapsto \varepsilon], \zeta) \in \ulcorner q(t) * o.f \mapsto v_1 \urcorner$$

The first property follows from our assumptions and monotonicity of later; the second is more involved. Since the interpretation of  $Prop_{SC}$  distributes over separating conjunction, we know there exist  $l_3, l_4 \in \Delta(LState)$ , such that  $\triangleright l_3 \bullet l_4 = \text{flush}(l_1, U(t))$ ,  $\triangleright (l_3, s_1, U[t \mapsto \varepsilon], \zeta) \in \ulcorner q(t) \urcorner$  and  $\triangleright (l_4, s_1, U[t \mapsto \varepsilon], \zeta) \in \ulcorner o.f \mapsto v_1 \urcorner$ . We can now show, using the properties of flush, that

$$\text{flush}(l_1, U(t) \cdot (o, f, v_1)) = \text{flush}(l_3, (o, f, v_1)) \bullet \text{flush}(l_4, (o, f, v_1)).$$

Since  $\triangleright (o, f) \in \text{dom}(l_4)$  has to hold, we can also conclude easily that

$$\triangleright (\text{flush}(l_4, (o, f, v_1)), s, U[t \mapsto \varepsilon], \zeta) \in \ulcorner o.f \mapsto v_1 \urcorner.$$

However, since  $\triangleright (o, f) \notin \text{dom}(l_3)$ ,  $\triangleright (l_3, U[t \mapsto \varepsilon]) R_{sb} (\text{flush}(l_3, (o, f, v_1)), U[t \mapsto \varepsilon])$ , and by Lemma 30 and monotonicity of later,  $\ulcorner q(t) \urcorner$  is stable under store-buffer interference later, which concludes the proof.  $\square$

**Lemma 36.** For any  $A \in \mathcal{P}(\Delta(RId))$ ,  $p \in \Delta(TId) \rightarrow Prop_{TSO}$ ,  $t \in \Delta(TId)$ ,  $o \in \Delta(OId)$ ,  $f \in \Delta(FName)$ ,  $v, v_o, v_n \in \Delta(Val)$ ,  $x \in \Delta(\mathbf{2})$ ,

$$p \vdash_r^m o.f \mapsto v \wedge v \neq v_o \Rightarrow \text{cas}(t, o, f, v_o, v_n, x) \text{ sat}_A \{p(t)\} \{p(t) * x = \mathbf{false}\}$$

*Proof.* Let  $r \in Prop_{TSO}$ ,  $m \in \mathcal{M}$ ,  $h \in \Delta(Heap)$ , and  $U \in \Delta(SPool)$  such that

$$m \in p(t) * \triangleright r \qquad (h, U) \in [m]_A \qquad \text{stable}_A(r).$$

Hence, there exists  $m_1, m_2 \in \mathcal{M}$  such that  $m = m_1 \bullet m_2$ ,  $m_1 \in p(t)$  and  $m_2 \in \triangleright r$ . As with read and write actions, if  $t \notin \text{dom}(U)$ , the statement is vacuously true, since in that case  $\llbracket cas(t, o, f, v_o, v_n, x) \rrbracket = \emptyset$ . Thus, we know that  $t \in \text{dom}(U)$ , and so, by definition of  $\vdash_r^m$ , we get  $\triangleright \text{lookup}(o.f, U(t), m_1.l) = v$ . As equalities on constant sets are upwards-closed, it follows that

$$\text{lookup}(o.f, U(t), m_1.l) = v \vee \triangleright \perp$$

As the conclusion follows trivially in the  $\triangleright \perp$  case (by Lemma 15), assume

$$\text{lookup}(o.f, U(t), m_1.l) = v.$$

This means that  $(o, f) \in \text{dom}(m_1.l) \subseteq \text{dom}(h)$ , and so  $\not\vdash \llbracket cas(t, o, f, v_o, v_n, x) \rrbracket$ .

Take any  $(h', U') \in \llbracket cas(t, o, f, v_o, v_n, x) \rrbracket(h, U)$ . Since  $v \neq v_o$ , there is only one non-vacuous case, in which we learn that  $h' = h$ ,  $U' = U$  and  $x = \mathbf{false}$ . Thus, we pick  $m' = m$ , so  $(h, U) \in [m]_A$  holds by assumption. To show that

$$\triangleright(m \in p(t) * \mathbf{false} = \mathbf{false} * r),$$

we pick the same splitting of  $m$ , which gives us  $m_2 \in \triangleright r$ . Now we can split  $m_1$  as  $m_1 \bullet (\varepsilon, m.s, m.U, m.\zeta)$ , and both branches follow easily by monotonicity of  $\triangleright$ .  $\square$

**Lemma 37.** For any  $A \in \mathcal{P}(\Delta(RId))$ ,  $p \in \Delta(TId) \rightarrow Prop_{TSO}$ ,  $q \in \Delta(TId) \rightarrow Prop_{SC}$ ,  $t \in \Delta(TId)$ ,  $o \in \Delta(OId)$ ,  $f \in \Delta(FName)$ ,  $v_o, v_n \in \Delta(Val)$ ,  $x \in \Delta(\mathbf{2})$ ,

$$p \vdash_w o.f \mapsto v_o, q \Rightarrow cas(t, o, f, v_o, v_n, x) \text{ sat}_A \{p(t)\} \{\ulcorner q(t) * o.f \mapsto v_n^\neg * x = \mathbf{true} \urcorner\}$$

*Proof.* Let  $r \in Prop_{TSO}$ ,  $m \in \mathcal{M}$ ,  $h \in \Delta(Heap)$  and  $U \in \Delta(SPool)$  such that

$$m \in p(t) * \triangleright r \qquad (h, U) \in [m]_A \qquad \text{stable}_A(r)$$

Hence, there exists  $m_1, m_2 \in \mathcal{M}$  such that  $m = m_1 \bullet m_2$ ,  $m_1 \in p(t)$  and  $m_2 \in \triangleright r$ . Let

$$m = (l, s, U, \zeta) \qquad m_1 = (l_1, s, U, \zeta) \qquad m_2 = (l_2, s, U, \zeta)$$

As in the other cases, if  $t \notin \text{dom}(U)$ , the statement holds vacuously. In the case when  $t \in \text{dom}(U)$ , by definition of  $\vdash_w$  we get  $\triangleright \text{lookup}(o.f, U(t), l_1) = v_o$  and

$$\triangleright(\text{flush}(l_1, U(t)), s, U[t \mapsto \varepsilon], \zeta) \in \ulcorner q(t) * o.f \mapsto v_o^\neg \urcorner.$$

Since equalities on constant sets are upwards-closed it follows that

$$\text{lookup}(o.f, U(t), l_1) = v_o \vee \triangleright \perp$$

As the conclusion follows trivially by Lemma 15 in the  $\triangleright \perp$  case, assume  $\text{lookup}(o.f, U(t), l_1) = v_o$ . Since  $(o, f) \in \text{dom}(l_1) \subseteq \text{dom}(h)$ , we can conclude that  $\zeta \notin \llbracket \text{cas}(t, o, f, v_o, v_n, x) \rrbracket$ , which leaves us with the second conjunct.

Take any  $(h', U') \in \llbracket \text{cas}(t, o, f, v_o, v_n, x) \rrbracket(h, U)$ . The only non-vacuous case is the successful compare-and-swap action, in which case  $x = \mathbf{true}$ ,  $U' = U[t \mapsto \varepsilon]$ , and  $h' = \text{flush}(h, U(t) \cdot (o, f, v_n))$ . Pick  $m' = (\text{flush}(l, U(t) \cdot (o, f, v_n)), s, U[t \mapsto \varepsilon], \zeta)$ . Since flushing distributes over  $\bullet_{LState}$ , we can easily show that

$$(\text{flush}(l_2, U(t) \cdot (o, f, v_n)), s, U[t \mapsto \varepsilon], \zeta) \in \triangleright(r)$$

by using stability of the latter assertion. By using the same argument as in the case of write, and Lemma 30, we can also show

$$\triangleright(\text{flush}(l_1, U(t) \cdot (o, f, v_n)), s, U[t \mapsto \varepsilon], \zeta) \in \ulcorner q(t) * o.f \mapsto v_n \urcorner * \mathbf{true} = \mathbf{true}.$$

What remains is showing that  $(h', U') \in \llbracket m' \rrbracket_A$ . However, we can pick the state  $l'_n = \text{flush}(l', U(t) \cdot (o, f, v))$  and the map  $sr_n(r) = (\text{flush}(\pi_1(sr(r)), U(t) \cdot (o, f, v)), U')$ , where  $l'$  and  $sr$  are the witnesses of  $(h, U) \in \llbracket m \rrbracket_A$ . This leaves us with showing  $sr_n(r) \in \llbracket (s, \zeta) \rrbracket_r$  for  $r \in \text{dom } s \cap A$ . However, since  $(o, f) \notin \text{dom}(\pi_1(sr(r)))$ , we know that  $sr(r) R_{sb} sr_n(r)$ . Also, since  $\llbracket (s, \zeta) \rrbracket_r$  is defined as application of the action model, it's upwards-closed under interference, and so  $sr_n(r) \in \llbracket (s, \zeta) \rrbracket_r$ .  $\square$

## 5.1 Read and Write Judgments

**Lemma 38** (W-Ax-Sound). *For any  $o \in \Delta(OId)$ ,  $f \in \Delta(FName)$ ,  $v \in \Delta(Val)$*

$$(\lambda t \in TId. \triangleright \ulcorner o.f \mapsto v \mathbf{in} t \urcorner) \vdash_w o.f \mapsto v, \top.$$

*Proof.* Take any  $t \in \Delta(TId)$  and  $\triangleright(l, s, U, \zeta) \in \ulcorner o.f \mapsto v \mathbf{in} t \urcorner$  such that  $t \in \text{dom}(U)$ . By definition, this means that  $\triangleright(\text{flush}(l, U(t)), s, U[t \mapsto \varepsilon], \zeta) \in \ulcorner o.f \mapsto v \urcorner$ . This implies both that

$$\triangleright(\text{flush}(l, U(t)), s, U[t \mapsto \varepsilon], \zeta) \in \ulcorner \top * o.f \mapsto v \urcorner$$

and, by definition of points-to predicate, that  $\triangleright(\text{flush}(l, U(t))).h(x, f) = v$ , which means that

$$\triangleright \text{lookup}(x.f, U(t), l) = v,$$

which in turn ends the proof.  $\square$

**Lemma 39** (W-\*Sound). *For any  $p \in \Delta(TId) \rightarrow Prop_{TSO}$ ,  $q \in Prop_{SC}$ ,  $r \in Prop_{SC}$ ,  $o \in \Delta(OId)$ ,  $f \in \Delta(FName)$ ,  $v \in \Delta(Val)$*

$$p \vdash_w o.f \mapsto v, r \Rightarrow (\lambda t \in TId. p(t) * \ulcorner q \mathbf{in} t \urcorner) \vdash_w o.f \mapsto v, q * r.$$

*Proof.* Take any  $t \in \Delta(TId)$  and  $m \in p(t) * \ulcorner q \mathbf{in} t \urcorner$  such that  $t \in \text{dom}(m.U)$ . By definition, there exist states  $m_1$  and  $m_2$  such that  $m = m_1 \bullet m_2$ ,  $m_1 \in p(t)$  and  $m_2 \in \ulcorner q \mathbf{in} t \urcorner$ . Let  $(l_1, s, U, \zeta) = m_1$  and  $(l_2, s, U, \zeta) = m_2$ . From the assumption we get that

$\triangleright\text{lookup}(o.f, U(t), l_1) = v$  and  $\triangleright(\text{flush}(l_1, U(t)), s, U[t \mapsto \varepsilon], \zeta) \in \ulcorner r * o.f \mapsto v \urcorner$ . Since  $m.l = l_1 \bullet_{LState} l_2$ , this means that  $\triangleright\text{lookup}(o.f, m.U(t), m.l) = v$ . For the remaining property, note that

$$\text{flush}(l_1 \bullet_{LState} l_2, \alpha) = \text{flush}(l_1, \alpha) \bullet_{LState} \text{flush}(l_2, \alpha),$$

so it suffices to show that  $\triangleright(\text{flush}(l_2, U(t)), s, U[t \mapsto \varepsilon], \zeta) \in \ulcorner q \urcorner$ , which follows directly from  $m_2 \in \triangleright^\ulcorner q \urcorner \text{ in } t^\urcorner$ .  $\square$

**Lemma 40** (W- $\mathcal{U}$ -Sound). *For any  $p \in \Delta(TId) \rightarrow Prop_{TSO}$ ,  $q \in \Delta(TId) \rightarrow Prop_{TSO}$ ,  $r \in Prop_{SC}$ ,  $o \in \Delta(OId)$ ,  $f \in \Delta(FName)$ ,  $v \in \Delta(Val)$ ,*

$$q \vdash_w o.f \mapsto v, r \Rightarrow (\lambda t \in TId. p(t) \mathcal{U}_t q(t)) \vdash_w o.f \mapsto v, r.$$

*Proof.* Take any  $t \in \Delta(TId)$  and  $(l, s, U, \zeta) \in p(t) \mathcal{U}_t q(t)$  such that  $t \in \text{dom}(U)$ . By definition of  $\mathcal{U}_t$ , this means that there exist  $\alpha$  and  $\beta$  such that  $U(t) = \alpha \cdot \beta$  and  $(\text{flush}(l, \alpha), s, U[t \mapsto \beta], \zeta) \in q(t)$  (we fold the update from the definition into  $\alpha$ , since it's not of interest in this proof). By instantiating the assumption with the latter, we get that  $\triangleright\text{lookup}(o.f, \beta, \text{flush}(l, \alpha)) = v$  and  $\triangleright(\text{flush}(\text{flush}(l, \alpha), \beta), s, U[t \mapsto \varepsilon], \zeta) \in \ulcorner r * o.f \mapsto v \urcorner$ . By the properties of flush and lookup, this means that  $\triangleright\text{lookup}(o.f, \alpha \cdot \beta, l) = v$  and

$$\triangleright(\text{flush}(l, \alpha \cdot \beta), s, U[t \mapsto \varepsilon], \zeta) \in \ulcorner r * o.f \mapsto v \urcorner,$$

which ends the proof.  $\square$

**Lemma 41** (W-Cons-Sound). *For any  $p \in \Delta(TId) \rightarrow Prop_{TSO}$ ,  $q \in \Delta(TId) \rightarrow Prop_{TSO}$ ,  $r \in Prop_{SC}$ ,  $o \in \Delta(OId)$ ,  $f \in \Delta(FName)$ ,  $v \in \Delta(Val)$ ,*

$$q \vdash_w o.f \mapsto v, r \wedge (\forall t \in TId. p(t) \subseteq q(t)) \Rightarrow p \vdash_w o.f \mapsto v, r.$$

*Proof.* Take any  $t \in \Delta(TId)$  and  $m \in p(t)$ , such that  $t \in \text{dom}(m.U)$ . By the entailment assumption, this means  $m \in q(t)$ , so  $\triangleright\text{lookup}(o.f, m.U(t), m.l) = v$  and

$$\triangleright(\text{flush}(m.l, m.U(t)), m.s, m.U[t \mapsto \varepsilon], \zeta) \in \ulcorner r * o.f \mapsto v \urcorner,$$

which ends the proof.  $\square$

**Lemma 42** (R-Self-Sound). *For any  $o \in \Delta(OId)$ ,  $f \in \Delta(FName)$ ,  $v \in \Delta(Val)$ ,*

$$(\lambda t. \triangleright^\ulcorner o.f \mapsto v \text{ in } t^\urcorner) \vdash_r^m o.f \mapsto v.$$

*Proof.* Take any  $t \in TId$  and  $(l, s, U, \zeta) \in p(t)$  such that  $t \in \text{dom}(U)$ . We have

$$\triangleright(\text{flush}(l, U(t)), s, U[t \mapsto \varepsilon], \zeta) \in \ulcorner o.f \mapsto v \urcorner.$$

By definition, this means we have  $\triangleright(l' \leq \text{flush}(l, U(t)))$  such that  $\triangleright(l', s, \zeta) \in o.f \mapsto v$ , and so  $\triangleright o.f \in \text{dom}(l')$ . From this, using properties of flush and lookup, we can conclude that  $\triangleright\text{lookup}(o.f, U(t), l) = v$ , which ends the proof.  $\square$

**Lemma 43** (R- $\mathcal{U}$ -M-Sound). *For any  $p, q \in \Delta(TId) \rightarrow Prop_{TSO}$ ,  $o \in \Delta(OId)$ ,  $f \in \Delta(FName)$ ,  $v \in \Delta(Val)$ ,*

$$q \vdash_{\mathbb{R}}^m o.f \mapsto v \Rightarrow (\lambda t \in TId. p(t) \mathcal{U}_t q(t)) \vdash_{\mathbb{R}}^m o.f \mapsto v.$$

*Proof.* Take any  $t \in \Delta(TId)$  and  $(l, s, U, \zeta) \in p(t) \mathcal{U}_t q(t)$  such that  $t \in \text{dom}(U)$ . By definition of  $\mathcal{U}_t$  we have  $\alpha$  and  $\beta$  such that  $U(t) = \alpha \cdot \beta$  and

$$(\text{flush}(l, \alpha), s, U[t \mapsto \beta], \zeta) \in q(t).$$

This allows us to use the assumption, from which we can conclude that

$$\triangleright \text{lookup}(o.f, \beta, \text{flush}(l, \alpha)) = v,$$

and so, using the properties of lookup and flush,  $\triangleright \text{lookup}(o.f, \alpha \cdot \beta, l) = v$ , which ends the proof.  $\square$

**Lemma 44** (R- $\bar{\mathcal{U}}$ -Sound). *For any  $p, q \in \Delta(TId) \rightarrow Prop_{TSO}$ ,  $o \in \Delta(OId)$ ,  $f \in \Delta(FName)$ ,  $v \in \Delta(Val)$ ,*

$$p \vdash_{\mathbb{R}}^o o.f \mapsto v \Rightarrow (\lambda t \in TId. \exists t' \in \Delta(TId). t \neq t' * p(t') \mathcal{U}_{t'} q(t')) \vdash_{\mathbb{R}}^m o.f \mapsto v.$$

*Proof.* Take any  $t \in \Delta(TId)$  and  $m \in \exists t' \in \Delta(TId). t \neq t' * p(t') \mathcal{U}_{t'} q(t')$  such that  $t \in \text{dom}(m.U)$ . By definition, this gives us  $t' \in \Delta(TId)$  and states  $(l_1, s, U, \zeta) \in t \neq t'$ ,  $(l_2, s, U, \zeta) \in p(t') \mathcal{U}_{t'} q(t')$ , such that  $m = (l_1 \bullet_{LState} l_2, s, U, \zeta)$ . By definition of  $\mathcal{U}_t$  we get  $\alpha$  and  $\beta$  such that  $U(t') = \alpha \cdot \beta$  and  $(l_2, s, U[t' \mapsto \alpha]) \in p(t')$  (this time we fold the update into  $\beta$ ). Now we can use the assumption, obtaining

$$\triangleright \text{lookup}(o.f, U[t' \mapsto \alpha](t), l_2) = v,$$

from which it easily follows, since  $t \neq t'$ , that  $\triangleright \text{lookup}(o.f, U(t), l) = v$ , which ends the proof.  $\square$

**Lemma 45** (R-Cons-M-Sound). *For any  $p, q \in \Delta(TId) \rightarrow Prop_{TSO}$ ,  $o \in \Delta(OId)$ ,  $f \in \Delta(FName)$ ,  $v \in \Delta(Val)$ ,*

$$q \vdash_{\mathbb{R}}^m o.f \mapsto v \wedge (\forall t. p(t) \subseteq q(t)) \Rightarrow p \vdash_{\mathbb{R}}^m o.f \mapsto v.$$

*Proof.* Take any  $t \in \Delta(TId)$  and  $m \in p(t)$  such that  $t \in \text{dom}(m.U)$ . By the entailment assumption  $m \in q(t)$ , so from the other assumption we get  $\triangleright \text{lookup}(o.f, m.U(t), m.l) = v$ , which ends the proof.  $\square$

**Lemma 46** (R-Other-Sound). *For any  $o \in \Delta(OId)$ ,  $f \in \Delta(FName)$ ,  $v \in \Delta(Val)$ ,*

$$\triangleright^{\ulcorner} o.f \mapsto v^{\urcorner} \vdash_{\mathbb{R}}^o o.f \mapsto v.$$

*Proof.* Take any  $t, t' \in \Delta(TId)$  and  $m \in \triangleright^{\ulcorner} o.f \mapsto v^{\urcorner}$  such that  $t \neq t'$  and  $t \in \text{dom}(m.U)$ . By definition of  $\ulcorner - \urcorner$ , we know that  $\triangleright m.l(o, f) = v$  and

$$\triangleright (\forall t \in \text{dom}(m.U). \forall v \in \text{Val}. (o, f, v) \notin m.U(t)).$$

This means that  $\triangleright(o, f, v) \notin m.U(t)$  for any  $v$ , and so

$$\triangleright(\text{lookup}(o.f, m.U(t), m.l) = m.l(o, f) = v).$$

□

**Lemma 47** (R- $\mathcal{U}$ -O-Sound). *For any  $p, q \in \Delta(TId) \rightarrow \text{Prop}_{TSO}$ ,  $o \in \Delta(OId)$ ,  $f \in \Delta(FName)$ ,  $v \in \Delta(Val)$ ,*

$$p \vdash_r^o o.f \mapsto v \Rightarrow (\lambda t \in TId. p(t) \mathcal{U}_t q(t)) \vdash_r^o o.f \mapsto v.$$

*Proof.* Take any  $t, t' \in \Delta(TId)$ ,  $(l, s, U, \zeta) \in p(t') \mathcal{U}_t q(t')$  such that  $t \neq t'$  and  $t \in \text{dom}(U)$ . By definition of  $\mathcal{U}_t$ , we get  $\alpha$  and  $\beta$  such that  $U(t') = \alpha \cdot \beta$  and  $(l, s, U[t' \mapsto \alpha], \zeta) \in p(t')$ . Thus, we can use the assumption to conclude that  $\triangleright \text{lookup}(o.f, U[t' \mapsto \alpha](t), l) = v$  and, since  $t \neq t'$ ,  $\triangleright \text{lookup}(o.f, U(t), l) = v$ . □

**Lemma 48** (R-Cons-O-Sound). *For any  $p, q \in \Delta(TId) \rightarrow \text{Prop}_{TSO}$ ,  $o \in \Delta(OId)$ ,  $f \in \Delta(FName)$ ,  $v \in \Delta(Val)$ ,*

$$q \vdash_r^o o.f \mapsto v \wedge (\forall t. p(t) \subseteq q(t)) \Rightarrow p \vdash_r^o o.f \mapsto v.$$

*Proof.* Take any  $t, t' \in \Delta(TId)$  and  $m \in p(t')$  such that  $t \neq t'$  and  $t \in \text{dom}(U)$ . By the entailment assumption, this means  $m \in q(t')$ , so we can use the read-judgment assumption to conclude that  $\triangleright \text{lookup}(o.f, m.U(t), m.l) = v$ . □

**Theorem 1** (Soundness). *If  $\Gamma \mid \Phi \vdash S$  then  $\llbracket \Gamma \mid \Phi \vdash S \rrbracket$ .*

## 6 Verification of a spin-lock in the TSO logic

In this section we verify the spin-lock implementation using the TSO logic. We assume the reader has read the accompanying paper.

### 6.1 Specification

$$\exists isLock, locked : \text{Prop}_{SC} \times \text{Val} \rightarrow \text{Prop}_{SC}.$$

$$\forall R : \text{Prop}_{SC}. \text{stable}(R) \Rightarrow$$

$$\begin{aligned} & \{\bar{R}\} \text{Lock}() \{r. isLock(R, r)\} \\ & \wedge \{isLock(R, \mathbf{this})\} \text{Lock.acquire}() \{locked(R, \mathbf{this}) * \ulcorner R \urcorner\} \\ & \wedge \{locked(R, \mathbf{this}) * \bar{R}\} \text{Lock.release}() \{\top\} \\ & \wedge \text{valid}(\forall x : \text{Val}. isLock(R, x) \Leftrightarrow isLock(R, x) * isLock(R, x)) \\ & \wedge \forall x : \text{Val}. \text{stable}(isLock(R, x)) \wedge \text{stable}(locked(R, x)) \end{aligned}$$

Formally, the *isLock* and *locked* assertions are also embedded using  $\overline{\phantom{x}}$ , however as mentioned earlier, region assertions are independent of the local state and we thus elide the embeddings to make the proofs more readable.

## 6.2 Predicate Definitions

$$isLock(R, x) = \exists n : \text{Rld. } [ACQ]_{-}^n * \text{region}(\{U, L\}, T_{lock}, I(x, R, n), n)$$

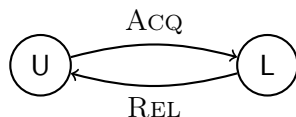
$$locked(R, x) = \exists n : \text{Rld. } [ACQ]_{-}^n * [REL]_{1}^n * \text{region}(\{L\}, T_{lock}, I(x, R, n), n)$$

where

$$I(x, R, n)(U) = \exists t : \text{Tld. } (\ulcorner x.locked \mapsto \text{true} \urcorner \mathcal{U}_t \ulcorner x.locked \mapsto \text{false} * R * [REL]_{1}^n \urcorner)$$

$$I(x, R, n)(L) = \ulcorner x.locked \mapsto \text{true} \urcorner$$

and  $T_{lock}$  refers to the following transition system



## 6.3 Proof outline

In this section we sketch proofs for each of the spin-lock methods in iCAP-TSO. Interestingly, the *structure* of the proof outlines differ in crucial ways from the corresponding proofs in iCAP (which assumes a strong memory model): In the case of acquire, even if the shared lock region is in the unlocked abstract state, it might still appear to be locked from the point of view of a client attempting to acquire the lock (if the lock was last released by a different thread, and release has not made its way to main memory yet), and the client will thus be unable to acquire the lock.

### Constructor

```

class Lock {
  bool locked;

  Lock() =
  {  $\overline{R} * \ulcorner \text{this.locked} \mapsto \text{null} \urcorner$  }
  <  $\overline{R} * \ulcorner \text{this.locked} \mapsto \text{null} \urcorner$  >
  CAS(this.locked, false, null);
  < ( $\ulcorner R * \text{this.locked} \mapsto \text{null} \urcorner$ ) >
  < region( $\{U, L\}$ ,  $T_{lock}$ ,  $I(\text{this}, R, n)$ ,  $n$ ) *  $[ACQ]_{-}^n$  >
  <  $isLock(R, \text{this})$  >
  {  $isLock(R, \text{this})$  }
  this
  {  $r. isLock(R, r)$  }

```

**Acquire.** The following is a proof sketch of the acquire method. As expected, most of the interesting reasoning concerns the atomic compare-and-swap expression that attempts to acquire the lock.

```

acquire() =
  {isLock(R, this)}
  {region({U, L}, Tlock, I(this, R, n), n) * [ACQ]-n}
  let y = CAS(this.locked, true, false) in
  {(y = true * region({L}, Tlock, I(this, R, n), n) * [ACQ]-n * [REL]1n *  $\ulcorner$ R $\urcorner$ )  $\vee$ 
  (y = false * region({U, L}, Tlock, I(this, R, n), n) * [ACQ]-n)}
  if y then
  {region({L}, Tlock, I(this, R, n), n) * [ACQ]-n * [REL]1n *  $\ulcorner$ R $\urcorner$ }
  {locked(R, this) *  $\ulcorner$ R $\urcorner$ }
  ()
  else
  {region({U, L}, Tlock, I(this, R, n), n) * [ACQ]-n}
  {isLock(R, this)}
  acquire()
  {locked(R, this) *  $\ulcorner$ R $\urcorner$ }

```

To verify the atomic compare-and-swap we use the **ATOMIC** rule to open the shared lock region. Since the pre-condition asserts that the lock is either locked or unlocked ( $region(\{U, L\}, \dots)$ ), we get two proof obligations, corresponding to each case:

```

<[ACQ]-n *  $\triangleright$ I(this, R, n)(U)>
  CAS(this.locked, true, false)
<r.  $\exists y \in \{U, L\}. [ACQ]-n *  $\triangleright$ I(this, R, n)(y) * Q(y, r, n)$ >

```

where

$$Q(y, r, n) = (y = L * [REL]<sub>1</sub><sup>n</sup> *  $\ulcorner$ R $\urcorner$  * r = true)  $\vee$  (y = U * r = false)$$

and

```

<[ACQ]-n *  $\triangleright$ I(this, R, n)(L)>
  CAS(this.locked, true, false)
<r. [ACQ]-n *  $\triangleright$ I(this, R, n)(L) * r = false>

```

We start with the (easy) second proof obligation:

```

<[ACQ]-n *  $\triangleright$  $\ulcorner$ this.locked  $\mapsto$  true $\urcorner$ >
   $\langle \triangleright \ulcorner$ this.locked  $\mapsto$  true $\urcorner$   $\rangle$ 
  CAS(this.locked, true, false)
  <r.  $\triangleright \ulcorner$ this.locked  $\mapsto$  true $\urcorner$  * r = false>
  <r. [ACQ]-n *  $\triangleright \ulcorner$ this.locked  $\mapsto$  true $\urcorner$  * r = false>

```



By rule A-CAS-FALSE it thus suffices to prove that

$$\triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \vdash_r^m \mathbf{this.locked} \mapsto \text{true},$$

which follows easily by rules R-SELF and R-CONS (to weaken  $\ulcorner - \urcorner$  to  $\bar{-}$ ).

To prove the first proof obligation, we first use STAB- $\mathcal{U}$  to do case analysis on whether or not the last release has made its way to main memory yet. Then we commute  $\triangleright$  all the way into the pre-condition, before finally doing case analysis on whether (if the last release is still pending) the last release is in our store buffer or not.

$$\begin{aligned} & \langle [\text{ACQ}]_n^n * \triangleright (\exists t : \text{Tld. } (\ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U}_t \ulcorner \mathbf{this.locked} \mapsto \text{false} * \text{R} * [\text{REL}]_1^n \urcorner)) \rangle \\ & \langle \triangleright (\exists t : \text{Tld. } (\ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U}_t \ulcorner \mathbf{this.locked} \mapsto \text{false} * \text{R} * [\text{REL}]_1^n \urcorner)) \rangle \\ & \langle \exists t : \text{Tld. } \triangleright (\ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U}_t \ulcorner \mathbf{this.locked} \mapsto \text{false} * \text{R} * [\text{REL}]_1^n \urcorner)) \rangle \\ & \langle \triangleright (\ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U}_t \ulcorner \mathbf{this.locked} \mapsto \text{false} * \text{R} * [\text{REL}]_1^n \urcorner)) \rangle \\ & \langle \triangleright ((\ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U}_t (\ulcorner \mathbf{this.locked} \mapsto \text{false} * \text{R} * [\text{REL}]_1^n \urcorner)) \vee \\ & \quad (\ulcorner \mathbf{this.locked} \mapsto \text{false} * \text{R} * [\text{REL}]_1^n \urcorner)) \rangle \\ & \langle \triangleright ((\ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U}_t \ulcorner \mathbf{this.locked} \mapsto \text{false} * \text{R} * [\text{REL}]_1^n \urcorner) \vee \\ & \quad \ulcorner \mathbf{this.locked} \mapsto \text{false} * \text{R} * [\text{REL}]_1^n \urcorner)) \rangle \\ & \langle \langle \langle \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U}_t \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner \text{R} \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner \rangle \vee \\ & \quad \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner \text{R} \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner \rangle \rangle \vee \\ & \langle \langle \langle \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U} \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner \text{R} \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner \rangle \vee \\ & \quad (\triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \bar{\mathcal{U}} \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner \text{R} \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner) \rangle \vee \\ & \quad \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner \text{R} \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner) \rangle \end{aligned}$$

**CAS(this.locked, true, false)**

$$\begin{aligned} & \langle r. \exists y \in \{\text{U}, \text{L}\}. \triangleright I(\mathbf{this}, \text{R}, n)(y) * Q(y, r, n) \rangle \\ & \langle r. [\text{ACQ}]_n^n * \exists y \in \{\text{U}, \text{L}\}. \triangleright I(\mathbf{this}, \text{R}, n)(y) * Q(y, r, n) \rangle \\ & \langle r. \exists y \in \{\bar{\text{U}}, \bar{\text{L}}\}. [\text{ACQ}]_n^n * \triangleright I(\mathbf{this}, \text{R}, n)(y) * Q(y, r, n) \rangle \end{aligned}$$

This leaves us with the following three proof obligations:

$$\begin{aligned} & \langle \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U} \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner \text{R} \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner \rangle \\ & \quad \mathbf{CAS}(\mathbf{this.locked}, \text{true}, \text{false}) \\ & \langle r. \exists y \in \{\text{U}, \text{L}\}. \triangleright I(\mathbf{this}, \text{R}, n)(y) * Q(y, r, n) \rangle \end{aligned}$$

and

$$\begin{aligned} & \langle \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \bar{\mathcal{U}} \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner \text{R} \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner \rangle \\ & \quad \mathbf{CAS}(\mathbf{this.locked}, \text{true}, \text{false}) \\ & \langle r. \exists y \in \{\text{U}, \text{L}\}. \triangleright I(\mathbf{this}, \text{R}, n)(y) * Q(y, r, n) \rangle \end{aligned}$$

and

$$\begin{aligned} & \langle \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner \text{R} \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner \rangle \\ & \quad \mathbf{CAS}(\mathbf{this.locked}, \text{true}, \text{false}) \\ & \langle r. \exists y \in \{\text{U}, \text{L}\}. \triangleright I(\mathbf{this}, \text{R}, n)(y) * Q(y, r, n) \rangle \end{aligned}$$

corresponding to the three cases:

- that the last release is pending in our store buffer, or
- the last release is pending in some other thread's store buffer, or
- the last release has already made its way to main memory.

In the first and the last case, the lock is objectively *and* subjectively unlocked, and we can thus acquire the lock and transition to the locked state. However, in the second case, the lock is objectively unlocked, but subjectively locked, and we thus remain in the unlocked state. We thus strengthen the post-conditions of these three proof obligations as follows:

$$\begin{aligned} &\langle \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U} \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner R \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner \rangle \\ &\quad \text{CAS}(\mathbf{this.locked}, \text{true}, \text{false}) \\ &\langle \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner * [\text{REL}]_1^n * \ulcorner R \urcorner * r = \text{true} \rangle \\ &\langle r. \triangleright I(\mathbf{this}, R, n)(L) * Q(L, r, n) \rangle \end{aligned}$$

and

$$\begin{aligned} &\langle \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \bar{\mathcal{U}} \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner R \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner \rangle \\ &\quad \text{CAS}(\mathbf{this.locked}, \text{true}, \text{false}) \\ &\langle r. \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \bar{\mathcal{U}} \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner R \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner * r = \text{false} \rangle \\ &\langle r. \triangleright (\exists t : \text{TId}. (\ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U}_t \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * R * [\text{REL}]_1^n)) * r = \text{false} \rangle \\ &\langle r. \triangleright I(\mathbf{this}, R, n)(U) * Q(U, r, n) \rangle \end{aligned}$$

and

$$\begin{aligned} &\langle \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner R \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner \rangle \\ &\quad \text{CAS}(\mathbf{this.locked}, \text{true}, \text{false}) \\ &\langle r. \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner * [\text{REL}]_1^n * \ulcorner R \urcorner * r = \text{true} \rangle \\ &\langle r. \triangleright I(\mathbf{this}, R, n)(L) * Q(L, r, n) \rangle \end{aligned}$$

By rules A-CAS-FALSE and A-CAS-TRUE, these three proof obligations reduce to the following three read/write proof obligations:

$$\begin{aligned} &\triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U} \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner R \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner \\ &\quad \vdash_w \mathbf{this.locked} \mapsto \text{false}, [\text{REL}]_1^n * R \end{aligned}$$

and

$$\triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \bar{\mathcal{U}} \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner R \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner \vdash_r^m \mathbf{this.locked} \mapsto \text{true}$$

and

$$\triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * \triangleright \ulcorner R \urcorner * \triangleright \ulcorner [\text{REL}]_1^n \urcorner \vdash_w \mathbf{this.locked} \mapsto \text{false}, [\text{REL}]_1^n * R$$

In the first and last case we can use rule W-CONS to weaken  $\ulcorner - \urcorner$  to  $\bar{\quad}$  and W-\* to frame off R and [REL] and remove the later in front of these. It thus suffices to prove:

$$\begin{aligned} &\triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U} \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner \vdash_w \mathbf{this.locked} \mapsto \text{false}, \top \\ &\quad \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \bar{\mathcal{U}} \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner \vdash_r^m \mathbf{this.locked} \mapsto \text{true} \\ &\quad \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner \vdash_w \mathbf{this.locked} \mapsto \text{false}, \top \end{aligned}$$

The last proof obligation follows easily from W-AX and W-CONS. By rules W- $\mathcal{U}$  and R- $\overline{\mathcal{U}}$ , the first two obligations reduce to:

$$\begin{aligned} \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner \vdash_w \mathbf{this.locked} \mapsto \text{false}, \top \\ \triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \vdash_r^o \mathbf{this.locked} \mapsto \text{true} \end{aligned}$$

These follow easily using rules W-AX, W-CONS and R-OTHER, R-CONS, respectively.

**Release.** Below we sketch the proof outline of release:

```

release() =
{locked(R, this) *  $\overline{R}$ }
{region({L},  $\top_{lock}$ ,  $I(\mathbf{this}, R, n), n$ ) * [ACQ] $^n$  * [REL] $^n_1$  *  $\overline{R}$ }
  { $\triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner$  * [ACQ] $^n$  * [REL] $^n_1$  *  $\overline{R}$ }
    { $\triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner$  * [REL] $^n_1$  *  $\overline{R}$ }
      this.locked := false
        {( $\triangleright \ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner$  * [REL] $^n_1$  *  $\overline{R}$ )  $\mathcal{U}$  ([REL] $^n_1$  *  $\ulcorner R \urcorner$  *  $\ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner$ )}
        { $\triangleright (\exists t : \text{TId. } (\ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U}_t \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * R * [REL] $^n_1$ ))$ }
        { $\triangleright (\exists t : \text{TId. } (\ulcorner \mathbf{this.locked} \mapsto \text{true} \urcorner \mathcal{U}_t \ulcorner \mathbf{this.locked} \mapsto \text{false} \urcorner * R * [REL] $^n_1$ )) * [ACQ] $^n$ }$ }
        { $\triangleright I(\mathbf{this}, R, n)(U) * [ACQ] $^n$ }$ }
      {region({U, L},  $\top_{lock}$ ,  $I(\mathbf{this}, R, n), n$ ) * [ACQ] $^n$ }
    {isLock(R, this)}

```

The innermost proof obligation follows by A-WRITE.

## References

- [1] T. Dinsdale-Young, L. Birkedal, P. Gardner, M. Parkinson, and H. Yang. Views: Compositional Reasoning for Concurrent Programs. In *Proceedings of POPL*, 2013.
- [2] K. Svendsen and L. Birkedal. Impredicative Concurrent Abstract Predicates. Technical report, 2013. Available at [www.kasv.dk/articles/icap-tr.pdf](http://www.kasv.dk/articles/icap-tr.pdf).