

# Strong Logic for Weak Memory: Reasoning About Release-Acquire Consistency in Iris

Jan-Oliver Kaiser, Hoang-Hai Dang, Derek Dreyer, Ori Lahav, and Viktor Vafeiadis

MPI-SWS, Saarbrücken and Kaiserslautern, Germany\*  
{janno, haidang, dreyer, orilahav, viktor}@mpi-sws.org

---

## Abstract

The field of *concurrent separation logics* (CSLs) has recently undergone two exciting developments: (1) the *Iris framework* for encoding and unifying advanced higher-order CSLs and formalizing them in Coq, and (2) the adaptation of CSLs to account for *weak memory models*, notably C11’s release-acquire (RA) consistency. Unfortunately, these developments are seemingly incompatible, since Iris only applies to languages with an operational interleaving semantics, while C11 is defined by a declarative (axiomatic) semantics. In this paper, we show that, on the contrary, it is not only feasible but useful to marry these developments together. Our first step is to provide a novel operational characterization of RA+NA, the fragment of C11 containing RA accesses and “non-atomic” (normal data) accesses. Instantiating Iris with this semantics, we then derive higher-order variants of two prominent RA+NA logics, GPS and RSL. Finally, we deploy these derived logics in order to perform the first mechanical verifications (in Coq) of several interesting case studies of RA+NA programming. In a nutshell, we provide the first foundationally verified framework for proving programs correct under C11’s weak-memory semantics.

**1998 ACM Subject Classification** F.3.1 Specifying and Verifying and Reasoning about Programs; F.3.2 Semantics of Programming Languages

**Keywords and phrases** Weak memory models, release-acquire, concurrency, separation logic

**Digital Object Identifier** [10.4230/LIPIcs.ECOOP.2017.17](https://doi.org/10.4230/LIPIcs.ECOOP.2017.17)

## 1 Introduction

*Separation logic* [24] is a refinement of Hoare logic with an intrinsic notion of *ownership*: whereas an assertion in Hoare logic denotes a fact about the global machine state, an assertion in separation logic denotes ownership of (and knowledge about) a *piece* of that state, and the *separating conjunction*  $P * Q$  denotes that the assertions  $P$  and  $Q$  own disjoint pieces of state. This ownership reading of assertions is useful for giving “local” (or “small-footprint”) specifications for primitive commands, which are much easier to compose soundly into specifications for larger programs. Moreover, as O’Hearn was the first to observe [23], separation logic is also eminently suitable for *concurrent* programs. In particular, ownership provides a direct and convenient way of explaining how synchronization mechanisms serve to transfer control of shared state between threads. Although O’Hearn’s original concurrent version of separation logic was geared toward reasoning about coarse-grained synchronization via semaphores, the subsequent decade of research into *concurrent separation logics* (CSLs) has shown that ownership and separation are just as useful for reasoning about more fine-grained

---

\* Saarland Informatics Campus.



and low-level synchronization mechanisms, such as those employed in the implementations of non-blocking data structures [34, 10, 6, 31, 29, 22, 5].

In this paper, we consider two of the most recent, boundary-pushing developments in concurrent separation logics: (1) the *Iris framework* for encoding and unifying advanced higher-order CSLs and formalizing them in Coq [13, 12, 15, 16], and (2) the adaptation of CSLs to account for *weak memory models*, notably C11’s *release-acquire* (RA) consistency [33, 32, 7, 19]. Although these developments have thus far (for reasons explained below) appeared to be incompatible, we show that in fact they are not! Quite the contrary: we demonstrate that it is not only feasible but useful to marry them together, and in so doing, provide the first foundationally verified framework for proving programs correct under C11’s weak memory semantics.

## 1.1 Iris: A Unifying Framework for Concurrent Separation Logics

After O’Hearn’s original CSL, there came a steady stream of “new and improved” CSLs appearing on at least a yearly basis. Unfortunately, as these new CSLs grew ever more expressive, they also grew increasingly complex, baking in increasingly sophisticated proof rules as primitive, with the relationships and compatibility between different proof rules (*e.g.*, whether they could be soundly combined in one logic) remaining unclear.

The central source of complexity in most existing CSLs lies in their mechanisms for controlling *interference* between threads accessing shared state, which have evolved from Jones’s *rely-guarantee* [11] to the much more sophisticated and elaborate *protocol* mechanisms appearing in logics like CaReSL [31], iCAP [29], and TaDA [5]. In an attempt to consolidate the field, Jung *et al.* developed **Iris** [13, 12, 15], a logic with the express goal of showing that even the fanciest of these interference-control mechanisms could be encoded via a combination of two orthogonal “off-the-shelf” ingredients: (1) *partial commutative monoids* (PCMs) for formalizing protocols on shared state, and (2) *invariants* for enforcing them. Invariants are an old and ubiquitous concept in program verification, and PCMs have been used in a number of prior logics to represent different kinds of *ghost* (or *auxiliary*) state *i.e.*, logical state that is manipulated as part of the proof of a program but is not manipulated directly by the program itself. Jung *et al.*’s observation was that in fact these two simple mechanisms are all you need: using just PCMs and invariants, one can *derive* a variety of powerful forms of protocol-based reasoning from prior CSLs *within* Iris, and by virtue of working in a unified framework, these derived mechanisms are automatically compatible (different mechanisms can be used soundly to verify different modules in a program). Iris also goes beyond most prior CSLs by supporting higher-order quantification and *impredicative invariants*—invariants that can talk recursively about the existence of (other) invariants—which are crucial for reasoning about languages with higher-order state (*e.g.*, Rust).

In the past, the complexity of CSLs was further exacerbated by the fact that (until very recently [26]) they only supported manual and error-prone “pencil-and-paper” proofs. The initial version of Iris [13] was no exception: the *soundness* of the core logic was verified in Coq, but the Coq development provided no support for *using* the logic (either to encode other logics or to verify programs interactively). However, in the past year, Krebbers *et al.* [16] have developed IPM, an interactive proof mode geared toward using Iris as a proof development environment for verifying concurrent programs within Coq. With IPM, Iris has begun the transition to a more practically useful proof tool, and is already being deployed effectively for larger verification efforts, *e.g.*, in the RustBelt project [9].

## 1.2 Separation Logics for Release-Acquire Consistency

Iris is a “generic” logical framework in that it is parameterized over the programming language in question—it merely requires, like the vast majority of prior work on concurrent program verification, that the language have an *operational, interleaving semantics*, typically known as a *sequentially consistent* (SC) semantics [20]. Under SC, threads take turns accessing the shared memory, and updates to memory are immediately visible to all other threads.

SC semantics has the benefit that it is easy to define and manipulate formally, but it is also woefully unrealistic: no serious language guarantees a fully SC semantics, because of the significant performance costs associated with maintaining the fiction of a single, globally consistent view of memory on modern multi-core architectures. One of the reasons for this discrepancy between the theory and the reality of concurrent programming is that, until relatively recently, formal accounts of more realistic—so-called *weak* (or *relaxed*)—memory models for concurrent programming languages were not available. However, in the past decade, great progress has been made on formalizing weak memory models, with a notable high point being the formalization of the C/C++11 memory model (hereafter, C11) [3].

In response to this development, a number of verification researchers have followed suit by building new verification tools—program logics, model checkers, testing frameworks, etc.—that account for these more realistic memory models. In particular, Vafeiadis and collaborators have thus far developed several different separation logics for C11, including RSL [33] and GPS [32]. The main focus of these logics is on RA+NA, an important fragment of C11 consisting of *release-acquire* (RA) accesses and *non-atomic* (NA) accesses. RA accesses are useful because they support a common idiom of message-passing synchronization at low cost compared to SC. NA accesses are intended for “normal” data accesses and are even more efficiently implementable than RA accesses, with the proviso that they are not permitted to race (*i.e.*, races on non-atomics cause the entire program to have undefined behavior).

A major challenge that Vafeiadis *et al.* had to overcome was the fact that C11 is defined using a radically different semantics than SC. Specifically, it is defined by a *declarative* (or *axiomatic*) semantics, in which the allowed behaviors of a program are defined by enumerating candidate executions (represented as “event graphs”) and then restricting attention to the executions that obey various coherence axioms. In building separation logics for C11, Vafeiadis *et al.* were thus not able to use the standard model of Hoare-style program specifications from prior separation logics because notions like “the machine states before and after executing a command  $C$ ” do not have a clear meaning in C11’s declarative semantics.

To account for this radically different type of semantics, they were instead forced to essentially throw away the “separation-logic textbook” and come up with an entirely new, non-standard model of separation logic in terms of predicates on event graphs. While groundbreaking, this approach has had several downsides. Firstly, certain essential mechanisms of SC-based separation logic (such as ghost state), which are easy to justify in standard models, became very difficult to justify in the new event-graph-based models of RA+NA logics. Secondly, the complexity of these new models has made them challenging to adapt and extend, and their non-standard nature has posed a major accessibility hurdle for researchers accustomed to traditional models of separation logic. Last but not least, although the soundness of these logics has been verified formally in Coq, there has thus far been no tool support for *using* the logics to prove programs correct under RA+NA semantics.

### 1.3 Our Contributions

Given our above description, it may seem that the Iris framework’s reliance on interleaving semantics renders it fundamentally inapplicable to reasoning about C11’s weak-memory semantics. In this paper, we show that this is not the case at all—not only is it possible to derive RA+NA logics like GPS and RSL within Iris, but there are several tangible benefits to doing so. Deriving such logics within Iris:

- Lets us take advantage of the rich features of the Iris host logic (*e.g.*, separation, invariants) when proving soundness of the derived logics, thereby significantly lifting the abstraction level at which those soundness proofs are carried out (compared to prior work).
- Allows us to support some very useful features in our derived logics by directly importing them from Iris. Such features include PCM-based ghost state, higher-order impredicative quantification, and Iris’s interactive proof mode in Coq. By virtue of being encoded in Iris, our derived logics inherit these features for free.
- Makes it easy to experiment with the derived logics and quickly develop new and useful extensions (*e.g.*, single-writer protocols, see below).
- Makes it possible to soundly compose proofs from different derived logics, since they are all carried out in the uniform framework of Iris.

Our first step (§2) is to avoid the essential complicating factor—C11’s declarative account of RA+NA—and instead work with an operational account. Building closely on Lahav *et al.*’s recently proposed “strong release-acquire” (SRA) semantics [18], we define a novel, operational, interleaving semantics for RA+NA. Our operational account of the RA fragment of the language is very similar to Lahav *et al.*’s operational account of SRA in that it models writing and reading of memory via the sending and receiving of timestamped messages; the main difference is that the RA rule for assigning timestamps is slightly more liberal. Our account of NA is new, though; it uses timestamps to model races on non-atomics as stuck (unsafe) machine states. We have proven that, under the reasonable restriction that programs do not mix RMWs (atomic updates) and non-atomic reads at the same location, our semantics is equivalent to the standard declarative semantics of the RA+NA fragment of C11.

Next, since our new semantics for RA+NA is an interleaving semantics, we can instantiate Iris with it. In §3, we review the basic reasoning mechanisms of Iris, and show how to use them to derive small-footprint proof rules for reasoning about RA+NA programs. We apply these rules to verify a simple message-passing example of RA+NA programming in Iris. However, as will become clear, reasoning directly with the Iris primitive mechanisms is rather too low-level and a more abstract logic is needed.

In §4, we present iGPS, a higher-order variant of Turon *et al.*’s GPS logic [32], which supports much higher-level reasoning about RA+NA programs. Unlike the original GPS, iGPS is derived within Iris on top of the small-footprint proof rules from §3. It also extends GPS with *single-writer protocols*, an extremely useful feature that simplifies proofs of RA+NA programs in the common case where there are no write-write races on atomic accesses.

In §5, we briefly describe some other contributions, including iRSL, a higher-order variant of RSL [33] derived within Iris, and several case studies that we have verified using iGPS and iRSL in Coq. These examples showcase one of the major advantages of working in the Iris framework: our ability to verify weak-memory programs, foundationally and mechanically, with the same degree of ease that was previously only possible for SC programs.

Finally, in §6, we conclude with related work.

## 2 Release-Acquire and Non-Atomics

In this section, we introduce our operational semantics for RA+NA, which we then use as the machine for our working language  $\lambda_{RN}$ . Subsequent sections will show how to build a logic for  $\lambda_{RN}$  using Iris.

C11 provides several memory access modes, each ensuring a different degree of consistency. In this paper we focus on RA+NA, the fragment of C11 consisting only of release-acquire (RA) and non-atomic (NA) accesses. Non-atomic accesses (which we denote with “[na]”) are the default type of memory accesses, intended to be used for normal data rather than for synchronization. Thus, C11 forbids any data races on non-atomic accesses, and programs that may have such races are considered buggy (they have undefined semantics). In contrast, RA accesses (which we denote with “[at]” for *atomic*) are permitted to race, but provide just enough consistency guarantees to enable the well-known *message passing* (MP) idiom:

$$\begin{array}{l} x_{[na]} := 0; y_{[na]} := 0; \\ x_{[na]} := 37; \quad \parallel \quad \mathbf{repeat} \ y_{[at]}; \\ y_{[at]} := 1 \quad \parallel \quad x_{[na]} \end{array}$$

Initially, both variables  $x$  and  $y$  are set to 0. The first thread will initialize  $x$  to 37 (non-atomically) and then set the variable  $y$  to 1 (via a release write) as a way of sending a message to the second thread that  $x$  has been properly initialized and is ready for consumption. The second thread will repeatedly read  $y$  (via an acquire read) until it observes  $y \neq 0$ , at which point—thanks to release-acquire semantics—it will know that it can safely access  $x$ . Summing up, the use of RA here ensures that the non-atomic write to  $x$  in the first thread “happens before” the non-atomic read of  $x$  in the second thread—*i.e.*, that they do not race—and furthermore that the read of  $x$  will return 37.

The formal semantics of RA+NA is “declarative”, formulated as a set of constraints on execution graphs. We will instead now present an alternative *operational* semantics of RA+NA. Our operational semantics is not completely coherent with C11’s for programs that mix atomic and non-atomic accesses to the same location (although the semantics of such programs is already known to be problematic [2]—see §6 for further discussion of this point). However, for the large class of programs that do not mix atomic updates (like CAS) and non-atomic reads at the same location, our semantics is provably equivalent to C11’s declarative semantics. This class of programs includes all C11 programs considered (and verified) in this paper. (For formal details of the correspondence between our semantics and C11’s, see [Appendix A](#).) We will first start with the pure RA fragment, and then add a “race detector” for non-atomic accesses.

### 2.1 Release-Acquire

Our operational semantics for RA starts from the observation that in RA—in contrast to a standard heap language—different threads have a different view of what the state is. Accordingly, we need to keep track of past write events as they might still be relevant for some subset of threads. Moreover, we need to keep writes to the same location in a total order enforced in C11 under the name *modification order* (*mo* for short). Finally, we also need to keep track of each thread’s “progress” in terms of which writes are visible to it, as this determines what a thread may read and where its writes may end up.

For the *mo* order, the RA machine manages for each location a totally ordered set of timestamps  $t \in \text{Time} \triangleq \mathbb{N}$ . Each write of some value  $v$  to a location  $\ell$  gets assigned a

$$\begin{array}{c}
\text{THREAD-READ} \\
\frac{(\ell, v, t, V) \in M \quad T(\pi)(\ell) \leq t}{(M, T) \xrightarrow{\langle \text{Read}, \ell, v \rangle}^\pi (M, T[\pi \mapsto T(\pi) \sqcup V])} \\
\\
\text{THREAD-WRITE} \\
\frac{\neg \exists v', V. (\ell, v', t, V) \in M \quad T(\pi)(\ell) < t \quad V' = T(\pi)[\ell \mapsto t]}{(M, T) \xrightarrow{\langle \text{Write}, \ell, v \rangle}^\pi (M \cup \{(\ell, v, t, V')\}, T[\pi \mapsto V'])} \\
\\
\text{THREAD-UPDATE} \\
\frac{(\ell, v_o, t, V) \in M \quad T(\pi)(\ell) \leq t \quad \neg \exists v, V. (\ell, v, t+1, V) \in M \quad V' = T(\pi)[\ell \mapsto t+1] \sqcup V}{(M, T) \xrightarrow{\langle \text{Update}, \ell, v_o, v_n \rangle}^\pi (M \cup \{(\ell, v_n, t+1, V')\}, T[\pi \mapsto V'])} \\
\\
\text{THREAD-FORK} \qquad \qquad \qquad \text{THREAD-UNINITIALIZED} \\
\frac{\rho \notin \text{dom}(T)}{(M, T) \xrightarrow{\langle \text{Fork}, \rho \rangle}^\pi (M, T[\rho \mapsto T(\pi)])} \qquad \frac{\varepsilon \in \{ \langle \text{Read}, \ell, v \rangle, \langle \text{Update}, \ell, v_o, v_n \rangle \} \quad T(\pi)(\ell) = \perp}{(M, T) \xrightarrow{\varepsilon}^\pi \perp_{\text{uninit}}}
\end{array}$$

■ **Figure 1** Per-thread reductions for RA without NA.

timestamp (that is unique for  $\ell$ ), resulting in a write event  $\omega \in \text{Event} \triangleq \text{Loc} \times \text{Val} \times \text{Time}$ , where  $v \in \text{Val} \triangleq \mathbb{Z}$ .<sup>1</sup> Using timestamps, the thread’s “progress” is represented by a *view*,  $V \in \text{View} \triangleq \text{Loc} \stackrel{\text{fin}}{\times} \text{Time}$ , which records the timestamp of the most recent write event observed by the thread for every location. To enable communication between threads, every write event is augmented with the writing thread’s view, yielding a message  $m \in \text{Msg} \triangleq \text{Event} \times \text{View}$ . The machine state  $\sigma$  comprises a message pool (called *memory*) and a view for every thread.

► **Definition 1** (Simplified Physical State). Let  $\sigma \in \Sigma \triangleq (\mathcal{P}(\text{Msg}) \times (\text{ThreadId} \stackrel{\text{fin}}{\times} \text{View})) \uplus \{\perp_{\text{uninit}}\}$  represent physical machine states, where  $\perp_{\text{uninit}}$  represents an error state. We write  $M$  and  $T$  to denote the two components of a non-error state.

The  $\lambda_{\text{RN}}$  language’s reductions are factored into *expression reductions*, concerned with the evaluation of the language’s expressions, and *machine reductions*, concerned with how the execution of an expression affects the machine state. We will define the expression reductions later when we formally define  $\lambda_{\text{RN}}$ . Here we focus on the machine reductions.

We define event labels  $\varepsilon \in \mathcal{E} \triangleq \{ \langle \text{Read}, \ell, v \rangle, \langle \text{Write}, \ell, v \rangle, \langle \text{Update}, \ell, v_o, v_n \rangle, \langle \text{Fork}, \rho \rangle \}$ , representing reads, writes, atomic updates (RMW’s), and forks (with  $\rho$  being the newly created thread id), respectively. The reductions are defined by a set of local, per-thread reductions  $\xrightarrow{\varepsilon}^\pi \subseteq \Sigma \times \Sigma$  given in **Figure 1**, where  $\pi$  represents the current thread’s id.

A write (**THREAD-WRITE**) picks an *unused* timestamp  $t$  for location  $\ell$  that is greater than the thread’s view of  $\ell$ , updates the thread’s view to the new view  $V'$  that includes  $t$ , and adds a corresponding message to the memory. A read (**THREAD-READ**) incorporates the view  $V$  of the message that it reads into the thread’s own view.<sup>2</sup> Note that the message being read

<sup>1</sup> The full semantics also supports allocation, which induces an allocation value  $A$ . We do not mention it here for the sake of simplicity.

<sup>2</sup> Here, we use the join operator  $\sqcup$  on views:  $(V_1 \sqcup V_2)(\ell) = \max\{V_1(\ell), V_2(\ell)\}$  if  $\ell \in \text{dom}(V_1) \cap \text{dom}(V_2)$ ;  $(V_1 \sqcup V_2)(\ell) = V_i(\ell)$  if  $\ell \in \text{dom}(V_i) \setminus \text{dom}(V_j)$ ; and  $(V_1 \sqcup V_2)(\ell)$  is undefined if  $\ell \notin \text{dom}(V_1) \cup \text{dom}(V_2)$ .



is required to have a timestamp that is not smaller than the thread’s view of the relevant location. Updates (**THREAD-UPDATE**) combine reading and writing in one step. In addition, updates must “pick”  $t + 1$  as a timestamp for the new message, where  $t$  is the timestamp of the read message. This implies, in particular, that two different updates cannot read the same message, and corresponds to C11’s atomicity condition, which requires every update to read from its *mo*-immediate predecessor. **THREAD-FORK** adds a new thread whose view is copied from its parent. Finally, **THREAD-UNINITIALIZED** detects reads from uninitialized locations, and moves to the error state  $\perp_{\text{uninit}}$ .

## Functional correctness of MP

With the operational semantics of RA, we can now sketch why MP (assuming for now that all its accesses are RA) is functionally correct, *i.e.*, why the read of  $x$  by the second thread will return 37 when the program terminates. The write of 37 to  $x$  is recorded at a view  $V_{37}$ , which is then included in the view  $V_1$  of the write of 1 to  $y$  by the first thread. When the second thread reads 1 from  $y$ , its local view is updated to incorporate  $V_1$  (and thus also  $V_{37}$ ). A read from  $x$  is now guaranteed to read from the message setting  $x$  to 37 or from a more recent one, but no more recent one exists. Consequently, the return value will be 37.

## 2.2 Non-Atomics

Formally, C11 defines a data race as two memory accesses to the same location—of which at least one is a write and at least one is non-atomic—that are not ordered by “happens-before.” A program that exhibits data races in *some* of its execution graphs is called racy, and its behavior is considered undefined. We now show how to account for non-atomics and data races in the context of our operational semantics.

Let us first extend the set of physical states by another error state  $\perp_{\text{race}}$ , whose intent is captured by the following correspondence: a program is racy if and only if at least one of its machine executions can reach  $\perp_{\text{race}}$  (stated and proved formally in [Appendix A](#)).

To detect data races during the execution of a program, we add an additional component to the physical state: the *non-atomic view*  $N$ , which tracks the timestamp of the most recent non-atomic write to every location. Then, we place the following restrictions on all atomic and non-atomic operations (if violated, the program will enter the  $\perp_{\text{race}}$  state):

- To perform any access (atomic or non-atomic) to a location  $\ell$ , a thread  $\pi$  must have observed the most recent non-atomic write to  $\ell$ , *i.e.*,  $N(\ell) \leq T(\pi)(\ell)$ .
- A thread  $\pi$  can only perform a non-atomic read from a location  $\ell$  if it has observed the most recent (atomic or non-atomic) write to  $\ell$ , *i.e.*,  $\nexists t, (\ell, \_, t, \_) \in M. T(\pi)(\ell) < t$ .

In addition to these restrictions, we require non-atomic writes to pick timestamps greater than all existing timestamps of messages of the same location. Intuitively, these restrictions enforce that each non-atomic write to  $\ell$  starts a new “era” in  $\ell$ ’s timestamps, after which any attempt to access writes from a previous era (or to write with a timestamp from a previous era) constitutes a race. Note that there is an asymmetry between non-atomic reads and writes: non-atomic writes to  $\ell$  are allowed even when the thread has not observed the most recent write to  $\ell$ , as it is only required to observe the most recent *non-atomic* write to  $\ell$ . One might fear that this fails to detect the case when a non-atomic write is racing with a concurrent atomic write (and the atomic write happens first); but in this case the race will be detected in a different execution where the non-atomic write happens first (and the atomic write enters the  $\perp_{\text{race}}$  state), so the program will nevertheless be declared racy.

$$\begin{array}{c}
\text{READ} \\
\frac{(\ell, v, t, V) \in M \quad T(\pi)(\ell) \leq t \quad \alpha = \text{na} \Rightarrow \forall v', t', V'. (\ell, v', t', V') \in M \Rightarrow t' \leq T(\pi)(\ell)}{(M, T, N) \xrightarrow{\langle \text{Read}_{\alpha}, \ell, v \rangle}^{\pi} (M, T[\pi \mapsto T(\pi) \sqcup V], N)} \\
\\
\text{WRITE-AT} \\
\frac{\neg \exists v', V. (\ell, v', t, V) \in M \quad N(\ell) \leq T(\pi)(\ell) < t \quad V' = T(\pi)[\ell \mapsto t]}{(M, T, N) \xrightarrow{\langle \text{Write}_{\text{at}}, \ell, v \rangle}^{\pi} (M \cup \{(\ell, v, t, V')\}, T[\pi \mapsto V'], N)} \\
\\
\text{WRITE-NA} \\
\frac{\neg \exists v', V. (\ell, v', t, V) \in M \quad N(\ell) \leq T(\pi)(\ell) \quad V' = T(\pi)[\ell \mapsto t] \quad \forall v', t', V. (\ell, v', t', V) \in M \Rightarrow t' < t}{(M, T, N) \xrightarrow{\langle \text{Write}_{\text{na}}, \ell, v \rangle}^{\pi} (M \cup \{(\ell, v, t, V')\}, T[\pi \mapsto V'], N[\ell \mapsto t])} \\
\\
\text{UPDATE} \\
\frac{(\ell, v_o, t, V) \in M \quad N(\ell) \leq T(\pi)(\ell) \leq t \quad \neg \exists v, V. (\ell, v, t+1, V) \in M \quad V' = T(\pi)[\ell \mapsto t+1] \sqcup V}{(M, T, N) \xrightarrow{\langle \text{Update}, \ell, v_o, v_n \rangle}^{\pi} (M \cup \{(\ell, v_n, t+1, V')\}, T[\pi \mapsto V'], N)} \\
\\
\text{FORK} \\
\frac{\rho \notin \text{dom}(T)}{(M, T, N) \xrightarrow{\langle \text{Fork}, \rho \rangle}^{\pi} (M, T[\rho \mapsto T(\pi)], N)} \\
\\
\text{RACE-I} \\
\frac{\varepsilon \in \{\langle \text{Read}_{\alpha}, \ell, v \rangle, \langle \text{Write}_{\alpha}, \ell, v \rangle, \langle \text{Update}, \ell, v_o, v_n \rangle\} \quad T(\pi)(\ell) < N(\ell)}{(M, T, N) \xrightarrow{\varepsilon}^{\pi} \perp_{\text{race}}} \\
\\
\text{RACE-II} \\
\frac{\exists v', t', V'. (\ell, v', t', V') \in M \wedge T(\pi)(\ell) < t'}{(M, T, N) \xrightarrow{\langle \text{Read}_{\text{na}}, \ell, v \rangle}^{\pi} \perp_{\text{race}}} \\
\\
\text{UNINITIALIZED} \\
\frac{\varepsilon \in \{\langle \text{Read}_{\alpha}, \ell, v \rangle, \langle \text{Update}, \ell, v_o, v_n \rangle\} \quad T(\pi)(\ell) = \perp}{(M, T, N) \xrightarrow{\varepsilon}^{\pi} \perp_{\text{uninit}}}
\end{array}$$

■ **Figure 2** Per-thread reductions for the RA+NA machine.

Revisiting MP, we note that it is safe to have non-atomic accesses to  $x$ : the write is performed while the left thread is necessarily aware of the most recent non-atomic write to  $x$  (the initialization); and the read is performed while the right thread is necessarily aware of the most recent write to  $x$ , whose timestamp was incorporated into the right thread's view when it read  $y = 1$ .

Figure 2 presents the full operational semantics. It is based on the following definition of a physical state.

► **Definition 2** (Physical State).

Let  $\sigma \in \Sigma \triangleq \left( \mathcal{P}(\text{Msg}) \times (\text{ThreadId} \stackrel{\text{fin}}{\times} \text{View}) \times \text{View} \right) \uplus \{ \perp_{\text{race}}, \perp_{\text{uninit}} \}$  represent physical machine states. We write  $M$ ,  $T$ , and  $N$  to denote the components of a non-error state. The initial physical state, denoted  $\sigma_{\text{init}}$ , is given by  $(\emptyset, [0 \mapsto \emptyset], \emptyset)$ .

## 2.3 The $\lambda_{\text{RN}}$ language

$\lambda_{\text{RN}}$  is a standard lambda calculus with recursive functions, forks, and references with atomic and non-atomic accesses. The **repeat** construct that we have used in MP can be defined in terms of recursive functions. The interesting part of the language and its expression reductions is given in Figure 3. The expression reduction relation  $(e \xrightarrow{\varepsilon} e', \bar{e}_f)$  has four components: the original expression  $e$ , an (optional) machine memory event  $\varepsilon$ , the resulting



$$\begin{array}{l}
v \in \text{Val} ::= () \mid z \in \mathbb{Z} \mid \ell \in \text{Loc} \mid \mathbf{fix} (f, x). e \\
\alpha \in \text{Access} ::= \text{at} \mid \text{na} \\
e \in \text{Expr} ::= v \mid e_1 e_2 \mid \mathbf{if} z \mathbf{then} e_1 \mathbf{else} e_2 \mid \mathbf{fork} e \mid \ell_{[\alpha]} \mid \ell_{[\alpha]} := v \mid \mathbf{cas}(\ell, v, v) \mid \dots
\end{array}$$

$\ell_{[\alpha]}$	$\xrightarrow{\langle \text{Read}_{\alpha}, \ell, v \rangle}$	$v, \text{nil}$
$\ell_{[\alpha]} := v$	$\xrightarrow{\langle \text{Write}_{\alpha}, \ell, v \rangle}$	$() , \text{nil}$
$\mathbf{cas}(\ell, v_o, v_n)$	$\xrightarrow{\langle \text{Update}, \ell, v_o, v_n \rangle}$	$1, \text{nil}$
$\mathbf{cas}(\ell, v_o, v_n)$	$\xrightarrow{\langle \text{Read}_{\text{at}}, \ell, v \rangle}$	$0, \text{nil} \quad \text{if } v \neq v_o$
$\mathbf{fork} e$	$\xrightarrow{\langle \text{Fork}, \rho \rangle}$	$() , [e]$
$(\mathbf{fix} (f, x). e) v$	$\rightarrow$	$e[(\mathbf{fix} (f, x). e)/f][v/x], \text{nil}$
$\mathbf{if} z \mathbf{then} e_1 \mathbf{else} e_2$	$\rightarrow$	$e_1, \text{nil} \quad \text{if } z \neq 0$
$\mathbf{if} z \mathbf{then} e_1 \mathbf{else} e_2$	$\rightarrow$	$e_2, \text{nil} \quad \text{if } z = 0$
		...

■ **Figure 3** Main  $\lambda_{\text{RN}}$  expressions and expression reductions.

$\frac{\text{COMBINED-PURE} \quad e \rightarrow e', \text{nil}}{\sigma; e \rightarrow^{\pi} \sigma; e', \text{nil}}$	$\frac{\text{COMBINED-MEM} \quad \begin{array}{l} \varepsilon \neq \langle \text{Fork}, - \rangle \\ e \xrightarrow{\varepsilon} e', \text{nil} \quad \sigma \xrightarrow{\varepsilon, \pi} \sigma' \end{array}}{\sigma; e \xrightarrow{\varepsilon, \pi} \sigma'; e', \text{nil}}$	$\frac{\text{COMBINED-FORK} \quad \begin{array}{l} e \xrightarrow{\langle \text{Fork}, \rho \rangle} e', [e_f] \quad \sigma \xrightarrow{\langle \text{Fork}, \rho \rangle} \pi \sigma' \end{array}}{\sigma; e \xrightarrow{\langle \text{Fork}, \rho \rangle} \pi \sigma'; e', [e_f]}$
$\frac{\text{THREADPOOL-RED-PURE} \quad \begin{array}{l} \sigma; \mathcal{TS}(\pi) \rightarrow^{\pi} \sigma'; e', \text{nil} \end{array}}{\sigma; \mathcal{TS} \rightarrow^{\pi} \sigma'; \mathcal{TS}[\pi \mapsto e']}$	$\frac{\text{THREADPOOL-RED-MEM} \quad \begin{array}{l} \sigma; \mathcal{TS}(\pi) \xrightarrow{\varepsilon, \pi} \sigma'; e', \text{nil} \end{array}}{\sigma; \mathcal{TS} \xrightarrow{\varepsilon, \pi} \sigma'; \mathcal{TS}[\pi \mapsto e']}$	
	$\frac{\text{THREADPOOL-RED-FORK} \quad \begin{array}{l} \sigma; \mathcal{TS}(\pi) \xrightarrow{\langle \text{Fork}, \rho \rangle} \pi \sigma'; e', [e_f] \end{array}}{\sigma; \mathcal{TS} \xrightarrow{\langle \text{Fork}, \rho \rangle} \pi \sigma'; \mathcal{TS}[\pi \mapsto e] \uplus [\rho \mapsto e_f]}$	

■ **Figure 4** Threadpool reductions.

expression  $e'$ , and a list of newly created threads  $\bar{e}_f$ . Only the rule for **fork**  $e$  creates a new thread (*i.e.*, a singleton list  $[e]$ ), while all other reductions produce an empty list (*i.e.*,  $\text{nil}$ ).

The *per-thread language reductions*  $(\sigma; e \xrightarrow{\varepsilon, \pi} \sigma'; e', \bar{e}_f)$  are then the combination of the expression reductions and the machine reductions, given by the COMBINED-\* rules in Figure 4. Non-stateful reductions (COMBINED-PURE) simply defer to the expression reductions, while stateful reductions (COMBINED-MEM and COMBINED-FORK) use the event label  $\varepsilon$  and the thread id  $\pi$  to tie the expression and machine reductions together correctly. These per-thread reductions then are lifted in a straightforward manner to the full (threadpool) reductions.

### 3 Iris

Iris is a generic framework for constructing concurrent separation logics. One can instantiate the framework with any language that has an operational interleaving semantics, and then easily derive time-tested reasoning principles for one's target logic, including various “protocol” mechanisms for controlling interference. Figure 5 provides an excerpt of Iris syntax.

$$P ::= \text{False} \mid \text{True} \mid P \Rightarrow Q \mid P \wedge Q \mid P \vee Q \mid P * Q \mid P \multimap Q \mid \exists x. P \mid \forall x. P \mid \mu x. P \\ \mid \text{Phys}(\sigma) \mid \overset{\gamma}{\square} a \mid \triangleright P \mid \square P \mid \overline{P}^{\mathcal{N}} \mid P \overset{\mathcal{N}_1}{\Rightarrow} \overset{\mathcal{N}_2}{Q} \mid \{P\} e \{x. Q\}_{\mathcal{N}} \mid \dots$$

■ **Figure 5** An excerpt of Iris syntax.

Iris supports the common connectives ( $\text{False}$ ,  $\text{True}$ ,  $\Rightarrow$ ,  $\wedge$ ,  $\vee$ ,  $*$ ,  $\multimap$ ,  $\exists$ ,  $\forall$ ,  $\mu$ ) and proof rules standard in higher-order separation logics. Iris’s extended set of constructs includes physical state ownership  $\text{Phys}(\sigma)$ , ghost state ownership  $\overset{\gamma}{\square} a$ , the later  $\triangleright$  and always  $\square$  modalities, invariants  $\overline{P}^{\mathcal{N}}$ , view shifts  $P \overset{\mathcal{N}_1}{\Rightarrow} \overset{\mathcal{N}_2}{Q}$ , and Hoare triples  $\{P\} e \{x. Q\}_{\mathcal{N}}$ . We will first explain these constructs via a running example, in which we use Iris to verify the MP example in a simple, sequentially consistent language called  $\lambda_{\text{SC}}$ . This will not only illustrate how one can derive within Iris a target logic for a language defined by an operational semantics, but will also serve as a warm-up for our subsequent explanation of how we can instantiate Iris to reason about weak memory.

### Road map

The process of instantiating Iris to derive new logics follows a simple pattern, which is worth articulating up front:

1. When we first instantiate Iris, the only primitive assertion we get about the state of the program is the *physical state ownership* assertion  $\text{Phys}(\sigma)$ , which asserts that  $\sigma$  is the current global state of the machine. Together with this assertion we also get for free a bunch of *large-footprint* specifications for the primitive commands of the language, based directly on their operational semantics. For example, the primitive specification we get for updating a location in  $\lambda_{\text{SC}}$  will be  $\{\text{Phys}(\sigma)\} \ell := v \{\text{Phys}(\sigma[\ell \mapsto v])\}$ .
2. Of course, one of the main points of separation logic is to be able to reason modularly using *local* assertions about the machine state, such as the *points-to* assertion,  $\ell \hookrightarrow v$ , and correspondingly give *small-footprint* specifications of the primitive commands, such as  $\{\ell \hookrightarrow w\} \ell := v \{\ell \hookrightarrow v\}$ . In Iris, such local assertions are not baked into the logic, but rather are encodable using *ghost state ownership* assertions, and the user of the logic has a great deal of flexibility concerning how these assertions are defined. In the case of the points-to assertion,  $\ell \hookrightarrow v$ , we will define this assertion so as to represent the *knowledge* that  $\ell$  currently points to  $v$  and the *rights* to read and write  $\ell$ .
3. On its own, a local, user-defined ghost state assertion like the points-to assertion is merely a *representation* of knowledge and rights. In order to give meaning to such a ghost state assertion—*i.e.*, to make sure it is in sync with the primitive physical state assertion—we establish an *invariant* tying the assertions together. In the case of the points-to assertion, this invariant will enforce that when a thread owns the ghost state assertion  $\ell \hookrightarrow v$ , its “knowledge” that  $\ell$  currently points to  $v$  in the physical machine state is actually correct.

**In short**, *ghost state assertions represent* local knowledge and rights concerning the machine state, and *invariants enforce* that ghost state assertions mean what they say they mean.

### 3.1 Iris by Example

Our example programming language  $\lambda_{\text{SC}}$  is a standard lambda calculus with references. It is basically the same as  $\lambda_{\text{RN}}$ , except that all accesses are sequentially consistent, and races are permitted (they do not induce stuckness). The language’s physical state is a heap  $\sigma$ , which is a finite map from allocated locations to values. The main heap-related reductions of  $\lambda_{\text{SC}}$  are given in [Figure 6](#). When we instantiate Iris with  $\lambda_{\text{SC}}$ ’s operational semantics,

$$\begin{aligned}
(!\ell, \sigma) &\rightarrow (v, \sigma) && \text{if } \sigma(\ell) = v \\
(\ell := v, \sigma) &\rightarrow ((), \sigma[\ell \mapsto v]) && \text{if } \ell \in \text{dom}(\sigma) \\
&\dots
\end{aligned}$$

■ **Figure 6** Main heap-related reductions of the  $\lambda_{\text{SC}}$  language.

(as explained in the above road map) what we get automatically from Iris are the following *large-footprint* Hoare triples concerning the physical state ownership assertion  $\text{Phys}(\sigma)$ :

$$\begin{array}{ll}
\text{PHYS-HEAP-READ} & \text{PHYS-HEAP-WRITE} \\
\{\text{Phys}(\sigma) * \sigma(\ell) = v\} !\ell \{z. z = v * \text{Phys}(\sigma)\} & \{\text{Phys}(\sigma)\} \ell := v \{\text{Phys}(\sigma[\ell \mapsto v])\}
\end{array}$$

Note that  $z$  in the first triple binds the return value of the expression  $!\ell$ . In the second triple, the expression returns the unit value, so we elide the binder.

### 3.1.1 Encoding Separation Logic for $\lambda_{\text{SC}}$

We would now like to encode these *small-footprint* Hoare triples for  $\lambda_{\text{SC}}$ :

$$\begin{array}{ll}
\text{HEAP-READ} & \text{HEAP-WRITE} \\
\{\ell \hookrightarrow v\} !\ell \{z. z = v * \ell \hookrightarrow v\} & \{\ell \hookrightarrow w\} \ell := v \{\ell \hookrightarrow v\}
\end{array}$$

The first step is to define the points-to assertion,  $\ell \hookrightarrow v$ , using Iris's *ghost state*.

#### Ghost state and partial commutative monoids

Ghost state is non-physical state that is only used as part of a program verification but is not itself part of the machine state. In Iris, ghost state is formalized using partial commutative monoids (PCMs).<sup>3</sup> The assertion  $\overline{[a : \overline{M}]^\gamma}$  asserts the ownership of the ghost resource  $a$  for an instance  $\gamma$  of the PCM  $M$ . Separating conjunction for ghost state assertions simply lifts the PCM composition operation to the assertion level:  $\overline{[a : \overline{M}]^\gamma} * \overline{[b : \overline{M}]^\gamma} \iff \overline{[a \cdot b : \overline{M}]^\gamma}$ . If two PCM fragments are not compatible (i.e. their composition is not defined), then it is not possible to own both of them at the same time, i.e., if  $a \cdot b = \perp$  then  $\overline{[a]^\gamma} * \overline{[b]^\gamma} \Rightarrow \text{False}$ .<sup>4</sup> In order to maintain consistency of the logic, therefore, changes to ghost state are restricted to *frame-preserving* updates, in which a PCM fragment  $a$  can only be updated to  $b$  if  $b$  preserves compatibility with any other fragments in the environment (the *frame*):

$$\frac{\text{GHOST-UPDATE} \quad \forall a_f. a \cdot a_f \neq \perp \Rightarrow b \cdot a_f \neq \perp}{\overline{[a]^\gamma} \Rightarrow \overline{[b]^\gamma}}$$

Ghost updates belong to the set of *logical computations*, or in Iris terminology, *view shifts*. A view shift  $P \Rightarrow Q$  represents the capability of transforming a resource satisfying  $P$  into a resource satisfying  $Q$  without changing the physical state.

<sup>3</sup> Actually, ghost state in Iris is based on the more general mechanism of “cameras” (aka step-indexed resource algebras), which can support a more general form of higher-order ghost state [12].

<sup>4</sup> In the rest of the paper we also suppress the PCM  $M$  in  $\overline{[a : \overline{M}]^\gamma}$  when it can be inferred in context.

### A PCM for heaps

As a step towards defining  $\ell \hookrightarrow v$ , let us now construct a PCM called **HEAP** that has the same basic structure as the physical heap, but allows splitting and recombination. (We will ultimately need a slightly more sophisticated PCM to define  $\ell \hookrightarrow v$ , but **HEAP** is an important part of the construction.) **HEAP** is a finite partial map from locations to values, with the empty heap as its unit element, and the composition on heaps is defined as disjoint union (*i.e.*, union if the heaps have disjoint domain, and undefined otherwise). The composition implies that the singleton heap  $[\ell := v]$  does not combine with itself, so it can only be uniquely owned, and it also represents the permission required to update  $\ell$ :

$$\begin{array}{c} \text{GHOST-HEAP-EXCLUSIVE} \\ \{[\ell := v]\}^\gamma * \{[\ell := w]\}^\gamma \vdash \text{False} \end{array} \qquad \begin{array}{c} \text{GHOST-HEAP-UPDATE} \\ \{[\ell := w]\}^\gamma \Rightarrow \{[\ell := v]\}^\gamma \end{array}$$

The singleton heap  $[\ell := v]$  therefore has the desired properties for defining the local assertion  $\ell \hookrightarrow v$ , but unfortunately it is still not quite enough: we also need some way to tie this ghost state assertion to the underlying physical state of the program. Toward this end, we employ Iris's *invariants*.

### Invariants

Invariants in Iris can be thought of as assertions that hold of *some* shared resource at all times, although the choice of which shared resource satisfies them is allowed to vary over time. The Iris invariant assertion  $\boxed{P}^{\mathcal{N}}$  stipulates that  $P$  is an invariant. The resource that satisfies it is shared with all threads, and thus any thread can access it freely in a single physical step<sup>5</sup>: it can *open* the invariant and gain local ownership of the resource for the duration of the operation, so long as it can *close* the invariant by relinquishing ownership of some (potentially different) resource satisfying  $P$  at the end of the operation. For bookkeeping purposes—specifically, to ensure that we do not unsoundly open the same invariant more than once in a nested fashion—invariants in Iris are named, and the  $\mathcal{N}$  in the above invariant assertion is a *namespace* (set of names) from which the name of the invariant must come.

Invariants belong to the set of *persistent* assertions, denoted with the always modality  $\Box$ . The assertion  $\Box P$  establishes the *knowledge* that  $P$  holds without any ownership, and therefore holds forever after. Putting resources into an invariant is thus a common way to share or transfer ownership through the use of freely distributable knowledge.

Meanwhile, the actions of opening and closing invariants belong to the set of logical computations, or view shifts. To account for invariants, view shifts are extended with namespaces as well:  $P \overset{\mathcal{N}_1}{\Rightarrow} \overset{\mathcal{N}_2}{Q}$  asserts that, assuming the invariants named in  $\mathcal{N}_1$  hold before the view shift, then the invariants named in  $\mathcal{N}_2$  hold after the view shift. Opening and closing of invariants are then formalized as follows:

$$\begin{array}{c} \text{INV-OPEN} \\ \boxed{P}^{\mathcal{N}} \vdash \text{True} \overset{\mathcal{N}}{\Rightarrow} \emptyset P \end{array} \qquad \begin{array}{c} \text{INV-CLOSE} \\ \boxed{P}^{\mathcal{N}} \vdash P \emptyset \overset{\mathcal{N}}{\Rightarrow} \text{True} \end{array}$$

**INV-OPEN** allows a thread to open the invariant  $\boxed{P}^{\mathcal{N}}$  and gain ownership of  $P$  but prevents it from doing so more than once. Only after applying **INV-CLOSE** and re-establishing the invariant will the thread be able to open it again.

<sup>5</sup> In Iris terminology, a resource in an invariant can be accessed within an *atomic* operation, which is an operation that takes only a single physical step of execution. We do not use the term here to avoid confusion with C11 atomics.

**Note:** The **INV-OPEN** rule as stated here is only sound if  $P$  talks about ownership (of physical or ghost state) and not about the existence of other invariants. In general, however, Iris makes no such restriction; rather, it supports *impredicative invariants*, meaning that  $P$  can be an arbitrary assertion. In order to avoid paradoxes caused by impredicative circularities (like the one described in [15]), the fully general version of this rule in Iris requires that  $P$  be guarded by the step-indexed *later* modality ( $\triangleright$ ). Fortunately, in most cases the  $\triangleright$ 's can be stripped away automatically (the Iris proof mode in Coq provides support for doing this) and do not play an interesting role in proofs. To focus the presentation of this paper, we will therefore suppress further discussion of the  $\triangleright$  modality.

Hoare triples in Iris are also annotated with invariant namespaces, since Hoare triples combine both physical and logical computations. A Hoare triple  $\{P\} e \{x. Q\}_{\mathcal{N}}$  with a namespace  $\mathcal{N}$  implies that if the invariants in  $\mathcal{N}$  hold before the expression's execution, then they will be preserved between every step and also after its execution. Consequently, when verifying any single physical step of computation, we are free to open the invariants in  $\mathcal{N}$  so long as we immediately close them. This reasoning is encapsulated in the following “atomic rule of consequence”:

$$\frac{\text{ACONSQ} \quad P \stackrel{\mathcal{N} \uplus \mathcal{N}'}{\Rightarrow} P' \quad \{P'\} e \{v. Q'\}_{\mathcal{N}} \quad \forall v. Q' \stackrel{\mathcal{N}}{\Rightarrow} \mathcal{N} \uplus \mathcal{N}' Q \quad e \text{ takes 1 physical step}}{\{P\} e \{v. Q\}_{\mathcal{N} \uplus \mathcal{N}'}}$$

Since bookkeeping of namespaces is largely a tedious detail (and one which Coq will force us to get right), we will for the remainder of the paper suppress namespaces from definitions and proofs. We will always use disjoint namespaces to ensure correctness in opening invariants.

### Linking physical and ghost state using invariants and the “authoritative” PCM

Now, returning to our example, the key idea is to use an invariant to tie the physical state assertion together with local ghost state assertions, thereby giving them meaning. To achieve this, we will employ an extremely useful construction called the *authoritative* PCM [13].

Given a base PCM  $\mathcal{M}$ , the authoritative PCM  $\text{AUTH}(\mathcal{M})$  has two types of elements: authoritative  $\bullet a$  and non-authoritative  $\circ a$  (for  $a \in \mathcal{M}$ ). For any instance  $\gamma$ ,  $\llbracket \bullet a \rrbracket^{\gamma}$  is *exclusive* (i.e.,  $\llbracket \bullet a \rrbracket^{\gamma} * \llbracket \bullet a \rrbracket^{\gamma} \vdash \text{False}$ ), and is the main point of reference for all the non-authoritative fragments, in the sense that any ownable fragment  $\llbracket \circ b \rrbracket^{\gamma}$  must have  $b$  *included* in  $a$ , that is  $\exists c. a = b \cdot c$ . The PCM's update therefore requires more: if one wants to update  $b$  to  $b'$ , it has not only to ensure  $b'$  is compatible with  $c$ , but also has to update  $a$  to  $a' = b' \cdot c$ . These properties are summarized in the following two rules:


$$\begin{array}{c} \text{AUTH-AGREE} \\ \llbracket \bullet a \rrbracket^{\gamma} * \llbracket \circ b \rrbracket^{\gamma} \vdash \exists c. a = b \cdot c \end{array} \quad \begin{array}{c} \text{AUTH-UPDATE} \\ \frac{b' \cdot c \neq \perp}{\llbracket \bullet b \cdot c \rrbracket^{\gamma} * \llbracket \circ b \rrbracket^{\gamma} \Rightarrow \llbracket \bullet b' \cdot c \rrbracket^{\gamma} * \llbracket \circ b \rrbracket^{\gamma}} \end{array}$$

With these two rules in hand, we can derive the following rules for operations on  $\text{AUTH}(\text{HEAP})$ :

$$\begin{array}{c} \text{AGHOST-HEAP-EXCLUSIVE} \\ \llbracket \circ [\ell := v] \rrbracket^{\gamma} * \llbracket \circ [\ell := w] \rrbracket^{\gamma} \vdash \text{False} \end{array} \quad \begin{array}{c} \text{AGHOST-HEAP-AGREE} \\ \llbracket \bullet \sigma \rrbracket^{\gamma} * \llbracket \circ [\ell := v] \rrbracket^{\gamma} \vdash \sigma(\ell) = v \end{array}$$

$$\begin{array}{c} \text{AGHOST-HEAP-UPDATE} \\ \llbracket \bullet \sigma \rrbracket^{\gamma} * \llbracket \circ [\ell := w] \rrbracket^{\gamma} \Rightarrow \llbracket \bullet \sigma[\ell \mapsto v] \rrbracket^{\gamma} * \llbracket \circ [\ell := v] \rrbracket^{\gamma} \end{array}$$

We are now ready to establish the invariant  $\llbracket \exists \sigma. \text{Phys}(\sigma) * \llbracket \bullet \sigma \rrbracket^{\gamma} \rrbracket$ , which binds together the physical state ownership and the authoritative ghost heap ownership. With the invariant



■ **Figure 7** Verification of MP in  $\lambda_{\text{SC}}$ .

in place, **AGHOST-HEAP-AGREE** implies that if a thread owns the singleton ghost heap  $\boxed{\circ}[\ell := v]$  locally, then, in combination with the invariant, it is guaranteed that  $\ell$  currently has value  $v$  in the physical heap. **AGHOST-HEAP-EXCLUSIVE** and **AGHOST-HEAP-UPDATE** ensure that only the one thread who owns  $\boxed{\circ}[\ell := v]$  can make updates to the contents of  $\ell$ .

The points-to assertion is then defined as  $\ell \hookrightarrow v \triangleq \boxed{\circ}[\ell := v]$ , and we can easily prove the small-footprint triples from the beginning of this section by combining the above rules for authoritative ghost heaps with those for opening and closing invariants.

### 3.1.2 Verifying MP in $\lambda_{\text{SC}}$

Using the small-footprint triples, we are ready to verify MP in  $\lambda_{\text{SC}}$ . We discuss the proof in a bit of detail here, since later on we will show how to adapt this proof to verify MP under weak-memory semantics.

The proof of MP is given in **Figure 7**. As a proof convention, we only mention persistent assertions (like invariants) once and use them freely later, since they are always true after being established. The proof works essentially as follows.

First of all, both threads want to operate on  $y$  simultaneously, so we need to put ownership of  $y$  into an invariant  $\boxed{\text{Inv}_y}$ . This invariant says that  $y$  is in one of two states—0 or 1. We can establish the invariant right after the initialization of  $y$  (the write of 0 to  $y$ ), because  $y$  is in state 0 at that moment. The first thread is responsible for setting  $y$  to state 1. When the second thread observes that  $y$  is in state 1, it will expect to be able to gain ownership of  $x \hookrightarrow 37$ . To achieve this, in state 1,  $\text{Inv}_y$  asserts the existence of another invariant  $\boxed{\text{Inv}_x}$  concerning  $x$ , and it is this latter invariant that we use to transfer ownership of location  $x$  from the first thread to the second thread.

To understand  $\text{Inv}_x$ , it helps to have seen the film *Raiders of the Lost Ark*, or at least the first few minutes of it, in which Indiana Jones (played by Harrison Ford) attempts to steal a precious golden idol from an ancient Peruvian temple—without setting off booby traps—by swapping it for a similarly weighted bag of sand. Unfortunately for him, the temple detects his ruse and tries to kill him. But we can play a similar trick, and Iris will be perfectly happy! In our case, the “golden idol” is  $x \hookrightarrow 37$ , which is transferred into the invariant  $\boxed{\text{Inv}_x}$  when it is established by the first thread. The “bag of sand” is a “token”  $\boxed{\diamond}$  (a uniquely ownable piece of ghost state) that is given to the second thread at the beginning of its execution.  $\text{Inv}_x$  simply asserts that it owns either the golden idol or the bag of sand. Thus, when the second thread learns of the existence of  $\text{Inv}_x$ , it can safely use the invariant opening and closing rules to swap the bag of sand in its possession for the golden idol owned by  $\text{Inv}_x$ , and thereafter claim local ownership of  $x \hookrightarrow 37$ .





$$\begin{array}{c}
 \text{BASE-NA-READ} \\
 \boxed{\text{PSInv}} \vdash \{\text{Seen}(\pi, V) * \text{Hist}(\ell, h) * \text{init}(h, V) * \text{na}(h, V)\} \\
 \ell_{[\text{na}]}, \pi \\
 \{v. \text{Seen}(\pi, V) * \text{Hist}(\ell, h) * \text{na}(h, V) * \max(h).v = v\} \\
 \\
 \text{BASE-NA-WRITE} \\
 \boxed{\text{PSInv}} \vdash \{\text{Seen}(\pi, V) * \text{Hist}(\ell, h) * \text{alloc}(h, V)\} \\
 \ell_{[\text{na}]} := v, \pi \\
 \{\exists V' \sqsupseteq V, t, h' = \{(v, t, V')\}. \text{Seen}(\pi, V') * \text{Hist}(\ell, h') * \text{na}(h', V') * \text{init}(h', V')\} \\
 \\
 \text{BASE-AT-READ} \\
 \boxed{\text{PSInv}} \vdash \{\text{Seen}(\pi, V) * \text{Hist}(\ell, h) * \text{init}(h, V)\} \\
 \ell_{[\text{at}]}, \pi \\
 \{v. \exists t_1, V_1, V' \sqsupseteq V \sqcup V_1. \text{Seen}(\pi, V') * \text{Hist}(\ell, h) * (v, t_1, V_1) \in h * t_1 \geq V(\ell)\} \\
 \\
 \text{BASE-AT-WRITE} \\
 \boxed{\text{PSInv}} \vdash \{\text{Seen}(\pi, V) * \text{Hist}(\ell, h) * \text{alloc}(h, V)\} \\
 \ell_{[\text{at}]} := v, \pi \\
 \{\exists V' \sqsupseteq V, t, h' = h \uplus \{(v, t, V')\}. \text{Seen}(\pi, V') * \text{Hist}(\ell, h') * \text{init}(h', V')\}
 \end{array}$$

■ **Figure 9** Selected Hoare triples of  $\lambda_{\text{RN}}$  base logic.

It is important to note that the history ownership assertion  $\text{Hist}(\ell, h)$  does not record the full history of  $\ell$ , but only write events of the *current era* (see §2.2), *i.e.*, only events as recent as or more recent than the last non-atomic write to  $\ell$ . This is reflected in the last condition of  $\text{HInv}$ :  $m.t \geq \sigma.N(\ell)$ . Defining  $\text{Hist}(\ell, h)$  this way helps to simplify the job of the user of the logic in establishing the absence of races: by construction, it is impossible to even attempt to read (racily) from write events before the current era. In order to preserve this property of  $\text{Hist}(\ell, h)$ , a non-atomic write  $(v_{\text{na}}, t_{\text{na}}, V_{\text{na}})$  must completely remove all write events currently in  $h$  (which would be racy to access now), and replace it with a new history  $h'$  that contains only the newly created non-atomic write event:  $h' = \{(v_{\text{na}}, t_{\text{na}}, V_{\text{na}})\}$ .

With the physical state shared and its ghost counterpart splittable, we are ready to derive the small-footprint Hoare triples, which constitute a base logic that is powerful enough to verify programs in  $\lambda_{\text{RN}}$ . A selected set of these triples is given in Figure 9. The general pattern of these rules is that a thread  $\pi$  needs to own the history  $h$  of a location  $\ell$  and its own thread view  $V$  in order to operate on  $\ell$ . Additionally,  $\pi$  needs to show certain relations between  $h$  and  $V$  in order to guarantee the safety of its operations. These relations are represented by the  $\text{alloc}$ ,  $\text{init}$ , and  $\text{na}$  predicates. The  $\text{alloc}(h, V)$  predicate (resp.  $\text{init}(h, V)$ ) asserts that  $V$  has observed an event in  $h$  which ensures that  $\ell$  is allocated (resp. initialized). The  $\text{na}(h, V)$  predicate asserts that  $V$  has observed the *mo*-latest write event of  $h$ , which is needed to do a non-atomic read. Note that all of these predicates require that  $V$  has seen *some* event from  $h$ —*i.e.*, an event from the current era of  $\ell$ —which, as discussed above, is a prerequisite for non-raciness.

These base logic rules provide a very concise explanation of  $\lambda_{\text{RN}}$ 's operational semantics. **BASE-AT-READ**, for example, requires the current view's knowledge of  $\ell$  being initialized and ensures that the new updated view  $V'$  is at least the join of the local view  $V$  and the view  $V_1$  of the write event that the thread reads from. This event must be from  $h$  and not *mo*-earlier

*Invariants:*

$$\begin{aligned} \text{Inv}_y(V_0) &\triangleq \exists h. \text{Hist}(y, h) * (0, \_, V_0) \in h * \forall V_1, v_1 \neq 0. (v_1, \_, V_1) \in h \Rightarrow \exists V_{37} \sqsubseteq V_1. \boxed{\text{Inv}_x(V_{37})} \\ \text{Inv}_x(V_{37}) &\triangleq \boxed{\diamond} \vee \text{Hist}(x, [(37, \_, V_{37})]) \end{aligned}$$

*Thread 1 proof outline:*

$$\begin{aligned} &\{ \text{Seen}(\pi, V_0) * \text{Hist}(x, [(0, \_, V_x)]) * V_x \sqsubseteq V_0 * \boxed{\text{Inv}_y(V_0)} \} \\ &x_{[\text{na}]} := 37 \\ &\{ \exists V_{37} \sqsupseteq V_0. \text{Seen}(\pi, V_{37}) * \text{Hist}(x, [(37, \_, V_{37})]) \} \\ &\{ \text{Seen}(\pi, V_{37}) * \boxed{\text{Inv}_x(V_{37})} \} \\ &\text{open } \text{Inv}_y \left\{ \begin{array}{l} \{ \text{Seen}(\pi, V_{37}) * \exists h. \text{Hist}(y, h) * \dots \} \\ y_{[\text{at}]} := 1 \\ \{ \exists V_1 \sqsupseteq V_{37}. \text{Seen}(\pi, V_1) * \text{Hist}(y, h \uplus [(1, \_, V_1)]) * \boxed{\text{Inv}_x(V_{37})} \} \end{array} \right\} \\ &\{ \text{Seen}(\pi, V_1) * \boxed{\text{Inv}_y(V_0)} \} \end{aligned}$$

*Thread 2 proof outline:*

$$\begin{aligned} &\{ \text{Seen}(\pi, V_0) * \boxed{\text{Inv}_y(V_0)} * \boxed{\diamond} \} \\ &\text{repeat } y_{[\text{at}]}; \\ &\{ \exists V_1, V_{37}, V_2. V_2 \sqsupseteq V_1 \sqsupseteq V_{37} * \text{Seen}(\pi, V_2) * \boxed{\text{Inv}_x(V_{37})} * \boxed{\diamond} \} \\ &\{ \text{Seen}(\pi, V_2) * V_{37} \sqsubseteq V_2 * \text{Hist}(x, [(37, \_, V_{37})]) \} \\ &x_{[\text{na}]} \\ &\{ z. \text{Seen}(\pi, V_2) * z = 37 * \text{Hist}(x, [(37, \_, V_{37})]) \} \end{aligned}$$

■ **Figure 10** Verification of MP in  $\lambda_{\text{RN}}$  base logic.

than the write event observed by the thread previously. **BASE-NA-READ** is similar, except that it requires that the current view must have observed the *mo*-latest write event to  $\ell$ , and therefore reads from that write event, which we denote by  $\max(h)$ . **BASE-NA-WRITE** preserves the **HInv** invariant by proactively dropping from the history all the old write events, which are *mo*-before this non-atomic write. Notice that, unlike **BASE-NA-READ**, **BASE-NA-WRITE** does not require  $\text{na}(h, V)$ , as explained in §2.2.

### 3.2.2 MP in $\lambda_{\text{RN}}$

We show that the base logic is enough to verify MP in  $\lambda_{\text{RN}}$ . The invariants and the proof, given in Figure 10, follow closely those used for MP in  $\lambda_{\text{SC}}$ . The singleton heap ownership in  $\lambda_{\text{SC}}$  is replaced with the history ownership, and extra conditions on view extension are added to reflect the view updates inherent in  $\lambda_{\text{RN}}$ . More specifically, in  $\text{Inv}_y$ , we enforce that any write of a non-zero value<sup>6</sup> to  $y$  be at a view  $V_1$  that extends  $V_{37}$ , which is the view at which the write of 37 to  $x$  is made by the first thread. Consequently, when the second thread observes  $V_1$  (by reading  $y$  to be non-zero), it must have also observed  $V_{37}$ , and thus can read  $x = 37$ , using the Indiana Jones invariant  $\text{Inv}_x$ . The extra conditions on  $V_0$  ensure that  $y$  is initialized with 0 at  $V_0$ , so that the second thread (having observed  $V_0$ ) can safely read  $y$ .

We have shown that the base logic is powerful enough to verify MP in  $\lambda_{\text{RN}}$ , and in principle it is powerful enough to verify many other realistic weak memory programs that

<sup>6</sup> In the MP example this value is always 1.

are expressible in  $\lambda_{RN}$  as well. However, it is also clear that the base logic is not abstract enough: one has to burden oneself with keeping track of the low-level details of changes to locations' histories and threads' views. What we really want is a way of abstracting away from those low-level details and finding simple high-level reasoning principles for  $\lambda_{RN}$ , the type of reasoning principles supported by RA+NA logics like GPS and RSL. We will now see how such high-level principles can in fact be derived on top of our low-level base logic.

## 4 iGPS

Vafeiadis *et al.* have introduced several logics for C11, and two in particular that were focused on RA+NA: GPS [32] and RSL [33]. The key difference between these logics and the RA+NA base logic presented in §3.2.1 is that, in GPS and RSL, the user does not reason explicitly about views—instead, the assertions of these logics are *implicitly* predicated over the view of the thread asserting them. This helps to hide much tedious reasoning about views, and leads naturally to a model of assertions as predicates over views (§4.3).

We have encoded iGPS and iRSL, variants of both GPS and RSL, in Iris, and we will focus here on iGPS, since it is the more sophisticated of the two logics. (We briefly describe iRSL in §5.) GPS introduced several useful abstractions that were not present in RSL: PCM-based ghost state, single-location protocols, and escrows. In encoding GPS in Iris, as noted in the introduction, we get PCM-based ghost state completely for free, just by working in Iris. Below, we will describe the other features, and how we account for them in Iris.

Note that our goal with iGPS is not to slavishly imitate all details of GPS, but to provide proof rules very similar to GPS's that are both strong enough to support all the examples from the GPS papers [32, 30] and significantly easier to prove sound within Iris. To this end, we slightly restrict select rules, but only in a way that does not impact their utility on known case studies. We discuss the differences from the original rules in detail in §6. Furthermore, in §4.2, we show how iGPS supports an additional feature—*single-writer protocols*—that was not supported by the original GPS and that significantly simplifies proofs.

### 4.1 Key Features of GPS

In this section, we explain the key features of GPS and how they are formalized in iGPS (besides PCM-based ghost state, which is directly imported into iGPS from Iris). A selected set of iGPS proof rules is given in Figure 11.

#### Non-atomics

Since non-atomic locations may not be raced on, GPS (and iGPS) reason about them much in the same way that locations are reasoned about in standard separation logic: using the points-to assertion,  $\ell \hookrightarrow v$ . Note that the proof rules for reading (iGPS-NA-READ) and writing (iGPS-NA-WRITE) and the exclusivity property (iGPS-NA-EXCLUSIVE) are equivalent to those from the logic for  $\lambda_{SC}$  (§3.1.1). Additionally, GPS (and iGPS) support fractional ownership of non-atomics to allow such locations to be read (but not written) by multiple threads at once. We omit the rules of fractional ownership for brevity.

#### Protocols

To reason about RA atomics, we need a mechanism for controlling interference on such accesses. Toward this end, CSLs for SC have supported a variety of *protocol* mechanisms, which control how shared state may evolve over time, and several of the more recent logics [31, 29, 22]

$$\begin{array}{c}
\text{iGPS-NA-READ} \quad \quad \quad \text{iGPS-NA-WRITE} \quad \quad \quad \text{iGPS-NA-EXCLUSIVE} \\
\{\ell \hookrightarrow v\} \ell_{[\text{na}]} \{w. w = v * \ell \hookrightarrow v\} \quad \quad \quad \{\ell \hookrightarrow \_ \} \ell_{[\text{na}]} := v \{\ell \hookrightarrow v\} \quad \quad \quad \ell \hookrightarrow v * \ell \hookrightarrow w \Rightarrow \perp \\
\\
\text{iGPS-READ} \quad \quad \quad \text{iGPS-WRITE} \\
\frac{\forall s' \sqsupseteq s, v. P * \tau_{\text{read}}(s', v) \Rightarrow Q}{\{\ell : s \mid \tau\} * P \ell_{[\text{at}]} \{v. \exists s' \sqsupseteq s. \{\ell : s' \mid \tau\} * Q\}} \quad \quad \quad \frac{(\forall s''. s' \sqsupseteq s'') \quad P \Rightarrow \tau_{\text{full}}(s', v) * Q}{\{\ell : s \mid \tau\} * P \ell_{[\text{at}]} := v \{\ell : s' \mid \tau\} * Q} \\
\\
\text{iGPS-CAS} \\
\frac{\forall s' \sqsupseteq s. P * \tau_{\text{full}}(s', v_o) \Rightarrow \exists s'' \sqsupseteq s'. \tau_{\text{full}}(s'', v_n) * Q \quad \forall s' \sqsupseteq s. P * \tau_{\text{read}}(s', v_o) \Rightarrow R}{\{\ell : s \mid \tau\} * P \text{ cas}(\ell, v_o, v_n) \{v. \exists s'' \sqsupseteq s. \{\ell : s'' \mid \tau\} * ((v = 1 \wedge Q) \vee (v = 0 \wedge R))\}} \\
\\
\text{iGPS-PERSISTENT} \quad \quad \quad \text{iGPS-AGREE} \\
\{\ell : s \mid \tau\} \Rightarrow \square \{\ell : s \mid \tau\} \quad \quad \quad \{\ell : s_1 \mid \tau\} * \{\ell : s_2 \mid \tau\} \Rightarrow s_1 \sqsubseteq s_2 \vee s_2 \sqsubseteq s_1 \\
\\
\text{iGPS-ESCROW-INTRO} \quad \quad \quad \text{iGPS-ESCROW-ELIM} \quad \quad \quad \text{iGPS-ESCROW-PERSISTENT} \\
Q \Rightarrow [P \rightsquigarrow Q] \quad \quad \quad P \wedge [P \rightsquigarrow Q] \Rightarrow Q \quad \quad \quad [P \rightsquigarrow Q] \Rightarrow \square [P \rightsquigarrow Q]
\end{array}$$

■ **Figure 11** iGPS proof rules for non-atomics, protocols, and escrows.

employ *state transition systems* (STs) to formalize such protocols. Crucially, protocols enforce *irreversibility*: the state of an STS protocol can only make forward progress over the course of a proof. For example, in §3.1.2, a protocol could enforce that the variable  $y$  could only progress from 0 to 1 but not back again. (We did not need to enforce that property to verify the MP example, but it is useful to be able to in general.) In Iris, protocols are encoded using a combination of invariants and ghost state.

Under weak memory, invariants and protocols are unsound in general because they require a single coherent history of updates to all locations. GPS showed how to partially restore protocol reasoning for weak memory with *single-location protocols*: protocols which restrict the evolution of a single shared location. Intuitively, single-location protocols are sound due to the per-location coherence property of C11 (often called “SC per location”): the writes to any single location are totally ordered (by *mo*). In particular, they maintain the invariant that the order of protocol states associated with writes is consistent with their timestamp (*mo*) order. If write event  $x$  to location  $\ell$  with associated protocol state  $s_x$  is *mo*-before write event  $y$  (to the same location) with protocol state  $s_y$ , then  $s_x$  will be before  $s_y$  in protocol order. Thus, once a thread has observed that the protocol on  $\ell$  has reached state  $s_x$ , it can from that point on only observe the protocol on  $\ell$  to be in states that are accessible from  $s_x$ . This fulfills the expectation that protocol transitions are irreversible.

GPS protocols come equipped with an interpretation predicate which specifies the resources held by the protocol depending on the protocol’s state and the location’s value. The primitive operations on an atomic location serve to transfer resources in and out of its protocol:

- Writes may transfer resources into the protocol, but may not transfer anything out.
- Reads may not transfer any resources into the protocol, but they may transfer “knowledge” (*i.e.*, duplicable resources) out of it. They are restricted to transferring out duplicable resources because there may be many reads of the same write event.
- Updates (RMWs), by virtue of the physical synchronization they provide, may transfer resources both in and out of the protocol.

In iGPS, we represent protocols in a slightly different way, using *two* interpretations: a “full” interpretation, and a duplicable “read” interpretation that is implied by the full

interpretation. The intuition is that the read interpretation is used for reads (since they may only transfer duplicable resources out of the protocol) and the full interpretation is used for the other operations.

We will now present the formal definition of protocols and the corresponding proof rules.

► **Definition 3 (Protocols).** A protocol  $\tau$  comprises a non-empty state set  $\mathbb{S}$ , a reflexive, transitive transition relation  $\sqsubseteq \subseteq \mathbb{S} \times \mathbb{S}$ , and two interpretation predicates  $\tau_m(\cdot, \cdot) \in \mathbb{S} \times \text{Val} \rightarrow \text{Prop}$  with  $m \in \{\text{read}, \text{full}\}$  representing read and full interpretations, respectively. The interpretation predicate has to fulfill the following two laws:

$$\tau_{\text{full}}(s, v) \Rightarrow \tau_{\text{full}}(s, v) * \tau_{\text{read}}(s, v) \qquad \tau_{\text{read}}(s, v) \Rightarrow \tau_{\text{read}}(s, v) * \tau_{\text{read}}(s, v)$$

We write  $\boxed{\ell : s \mid \tau}$  (as in GPS) to denote the persistent assertion that  $\ell$  is governed by protocol  $\tau$  and that the protocol has been observed in state  $s$ .

The first rule, **iGPS-AGREE**, represents the guarantee that every protocol is always in a state consistent with all observations, *i.e.*, that all observed states can be linearly ordered w.r.t. to the transition relation  $\sqsubseteq$ .

**iGPS-READ** enables reading from a location governed by protocol  $\tau$ —allowing the user to observe a future state  $s'$  of whatever state  $s$  it has previously observed, and providing the associated read interpretation  $\tau_{\text{read}}$ .

Writes to the location are subject to **iGPS-WRITE**, which allows the user to move the protocol to a “final state”  $s'$ —*i.e.*, a state accessible from every state in the protocol—so long as they can provide  $\tau_{\text{full}}$  for  $s'$ . This rule may seem very weak, since it forces the protocol into a final state, but this weakness derives from the need to handle the general case where there can be write-write races. Write-write races allow only very limited reasoning: both writes have to prove that their protocol state is later in protocol order to the other one. The write rule presented here solves this problem in a very simple way (see §6 for a comparison with the original GPS rule). In §4.2, we present a much stronger write rule that is optimized for the common case where there are no write-write races on the location.

Finally, **iGPS-CAS** governs updates. Its two premises represent the success and failure case, respectively. If the operation succeeds, the value read,  $v_o$ , is guaranteed to belong to a future state  $s'$ . The user picks the new state  $s''$  depending on  $s'$  and establishes  $\tau_{\text{full}}(s'', v_n)$ , making use of  $\tau_{\text{full}}(s', v_o)$ . In case of a failure, the rule degenerates to **iGPS-READ**.

## Escrows

A limitation of GPS protocols is that they offer no way to transfer ownership of (non-duplicable) resources from one thread to another unless the receiving thread performs physical synchronization via an update operation. For example, in our MP example, there is no update operation, and yet we want to transfer ownership of the non-atomic location  $x$  from the first thread to the second. For such an example, an additional mechanism for ownership transfer is required.

This motivates *escrows*, a mechanism for *logical* synchronization which, unlike protocols, is not tied to physical locations.

► **Definition 4 (Escrows).** An escrow  $[P \rightsquigarrow Q]$  consists of a *guard* resource  $P$  and a *payload* resource  $Q$  to be transferred. The guard resource  $P$  must be exclusive, *i.e.*  $P * P \Rightarrow \perp$ . The escrow assertion itself is persistent knowledge (freely duplicable).

The idea of escrows is really just a slight generalization of the “Indiana Jones invariant”  $\text{Inv}_x$  that we used in the proof of the MP example from §3.1.2. Following the explanation

$$\begin{array}{l}
\mathbf{XE}(x) \triangleq \llbracket \diamond \rrbracket \rightsquigarrow x \hookrightarrow 37 \\
\mathbf{YP}(x)(0, v) \triangleq v = 0 \\
\mathbf{YP}(x)(1, v) \triangleq v = 1 * \mathbf{XE}(x)
\end{array}
\quad
\left\{ \begin{array}{l}
\{x \hookrightarrow 0 * \boxed{y : 0 \mid \mathbf{YP}(x)}\} \\
x_{[\text{na}]} := 37 \\
\{x \hookrightarrow 37\} \\
\{\mathbf{XE}(x) * \boxed{y : 0 \mid \mathbf{YP}(x)}\} \\
y_{[\text{at}]} := 1 \\
\boxed{y : 1 \mid \mathbf{YP}(x)}
\end{array} \right\}
\quad
\left\| \left\{ \begin{array}{l}
\boxed{y : 0 \mid \mathbf{YP}(x)} * \llbracket \diamond \rrbracket \\
\mathbf{repeat} \ y_{[\text{at}]}; \\
\boxed{y : 1 \mid \mathbf{YP}(x)} * \mathbf{XE}(x) * \llbracket \diamond \rrbracket \\
\{x \hookrightarrow 37\} \\
x_{[\text{na}]} \\
\{z. z = 37 * x \hookrightarrow 37\}
\end{array} \right\}
\right.$$

■ **Figure 12** Verification of MP with iGPS.

there, the payload resource  $Q$  is the “golden idol”, the guard resource  $P$  is the “bag of sand”, and the escrow allows one to swap  $P$  for  $Q$ . The restriction on exclusivity of  $P$  ensures that this swap can only be performed once.

The proof rules for escrows follow the above intuition. **iGPS-ESCROW-INTRO** places the payload resource  $Q$  in escrow. Any thread that learns of the existence of this escrow can then use **iGPS-ESCROW-ELIM** to trade ownership of the guard resource  $P$  for  $Q$ .<sup>7</sup>

## Message passing in iGPS

The verification of MP using iGPS is given in **Figure 12**. Although the verification is sound for  $\lambda_{\text{RN}}$ , it is much simpler than the proof we carried out in the base logic of Iris, and is in fact very close in structure to the SC verification of MP in  $\lambda_{\text{SC}}$ . In particular, the Indiana Jones invariant  $\text{Inv}_x$  has now become an escrow  $\mathbf{XE}$ , and the invariant  $\text{Inv}_y$  has now become a (2-state) iGPS protocol  $\mathbf{YP}$ , but otherwise the steps are almost the same. The abstraction of iGPS has relieved us from the burden of reasoning with history and view updates explicitly.

## 4.2 Single-Writer Protocols

As we observed above, the iGPS protocol write rule suffers from a restriction: the user has to transition to a final state. This restriction is not present in CSLs for SC, which let the user pick the future state depending on the current one, much as **iGPS-CAS** does. Fortunately, in the common case when there are no write-write races to the location, this restriction can be lifted by *single-writer protocols*, a novel invention of iGPS.

A single-writer protocol splits the protocol assertion into two parts: an exclusive writer assertion  $\boxed{\ell : s \mid \tau}_W$  and a persistent reader assertion  $\boxed{\ell : s \mid \tau}_R$ . Owning the writer assertion provides both the permission to change the state as well as the guarantee that no one else can change it. Owning the reader assertion only allows reads. Thus, the reader assertion represents a lower bound on the protocol state whereas the state contained in the writer assertion is exactly the most recent state of the protocol. The full proof rules for single-writer protocols are given in **Figure 13**, and of these, the write rule **iGPS-SW-WRITE** is the most important. The writer knows exactly what the current state is, and is free to pick the next state accordingly.

<sup>7</sup> The rule given for elimination is only sound if  $Q$  is “timeless”, meaning that it only describes ownership of (ghost) state, not knowledge about protocols or escrows, as is the case in our message passing example. A more general rule, which returns  $Q$  under the later modality, is given in **Appendix C**.

$$\begin{array}{c}
\text{iGPS-SW-EXCLUSIVE-WRITER} \\
\boxed{\ell : s_1 \mid \tau}_W * \boxed{\ell : s_2 \mid \tau}_W \Rightarrow \perp \\
\\
\text{iGPS-SW-AGREE} \\
\boxed{\ell : s_1 \mid \tau}_R * \boxed{\ell : s_2 \mid \tau}_R \Rightarrow s_1 \sqsubseteq s_2 \vee s_2 \sqsubseteq s_1 \\
\\
\text{iGPS-SW-MAX} \\
\boxed{\ell : s_1 \mid \tau}_W * \boxed{\ell : s_2 \mid \tau}_R \Rightarrow s_1 \sqsupseteq s_2 \\
\\
\text{iGPS-SW-READ-EXCLUSIVE} \\
\left\{ \boxed{\ell : s \mid \tau}_W \right\} \ell_{[\text{at}]} \left\{ v. \boxed{\ell : s \mid \tau}_W * \tau_{\text{read}}(s, v) \right\} \\
\\
\text{iGPS-SW-READ} \\
\frac{\forall s' \sqsupseteq s, v. P * \tau_{\text{read}}(s', v) \Rightarrow Q}{\left\{ \boxed{\ell : s \mid \tau}_R * P \right\} \ell_{[\text{at}]} \left\{ v. \exists s' \sqsupseteq s. \boxed{\ell : s' \mid \tau}_R * Q \right\}} \\
\\
\text{iGPS-SW-WRITE} \\
\frac{P * \boxed{\ell : s'' \mid \tau}_W * \tau_{\text{full}}(s, \_) \Rightarrow \tau_{\text{full}}(s'', v) * Q \quad s'' \sqsupseteq s}{\left\{ \boxed{\ell : s \mid \tau}_W * P \right\} \ell_{[\text{at}]} := v \left\{ \boxed{\ell : s'' \mid \tau}_R * Q \right\}}
\end{array}$$

■ **Figure 13** A selection of single-writer proof rules.

### Applications of single-writer protocols

Single-writer protocols provide more explicitly intuitive and concise proofs over normal protocols when there are no *write-write* races, *i.e.*, only one thread is writing to the location. This may mean that there is exactly one writer in the whole program, or (perhaps more typically) that the programmer is using sufficient synchronization to ensure that there is exactly one active writer at a time. Such is the case for several headlining examples verified in GPS, including *circular buffer*, *bounded ticket lock*, and *read-copy update* [32, 30].

In the original GPS proofs for these examples, the lack of single-writer protocols meant that the proofs had to employ a significant amount of tedious ghost state (mostly in the form of so-called “protocol tokens”) to formalize the fact that the writing thread knew exactly which state the protocol had to be in at any given time. Using single-writer protocols, this reasoning is immediate from the **iGPS-SW-WRITE** rule. By removing the need for such boilerplate ghost state, single-writer protocols simplify and clarify the proofs of these examples.

An intriguing feature of the **iGPS-SW-WRITE** rule is that it is possible for the writer to relinquish ownership of the exclusive writer permission *while doing the write itself*. (This is why the writer permission appears in the precondition of the premise.) This extra flexibility is particularly useful when reasoning about RA implementations of locks (such as the bounded ticket lock). When the lock holder releases the lock (typically with a release write), this feature allows them to also give up their permission to do further release writes to the lock, so that it can be transferred to the next thread that acquires the lock.

## 4.3 The Model of iGPS

We now briefly describe the model of iGPS assertions. **Figure 14** contains an excerpt of the encoding of the standard assertions and connectives of CSL as well as non-atomics and escrows. The somewhat more involved model of protocols is detailed in **Appendix F**.

We model iGPS assertions as monotone predicates over views. The view parameter represents the current view of the thread making the assertion. The monotonicity requirement is motivated by the observation that the view of a thread only grows over the execution of a program. To ensure properties like the frame rule, it is therefore crucial that simply adding information to a view does not invalidate previously held (*e.g.*, framed) assertions. As a consequence of this requirement, we explicitly monotonize the encoding when necessary.



$$\begin{array}{ll}
\llbracket a_i \rrbracket^\gamma \triangleq \lambda \_ . \llbracket a_i \rrbracket^\gamma & \llbracket \{P\} e \{v. Q\} \rrbracket \triangleq \lambda V. \forall V' \sqsupseteq V, \pi. \\
\llbracket \Box P \rrbracket \triangleq \lambda V. \Box \llbracket P \rrbracket(V) & \{\text{PSInv}\} * \text{Seen}(\pi, V') * \llbracket P \rrbracket(V') \\
\llbracket P * Q \rrbracket \triangleq \lambda V. \llbracket P \rrbracket(V) * \llbracket Q \rrbracket(V) & (e, \pi) \\
\llbracket P \Rightarrow Q \rrbracket \triangleq \lambda V. \forall V' \sqsupseteq V. \llbracket P \rrbracket(V') \Rightarrow \llbracket Q \rrbracket(V') & \{v. \exists V'' \sqsupseteq V'. \text{Seen}(\pi, V'') * \llbracket Q \rrbracket(V'')\} \\
\llbracket \ell \hookrightarrow v \rrbracket \triangleq \lambda V. \exists V_{\text{na}} \sqsubseteq V. \text{Hist}(\ell, \{(v, \_ , V_{\text{na}})\}) & \\
\llbracket P \rightsquigarrow Q \rrbracket \triangleq \lambda V. \exists V_0 \sqsubseteq V. \llbracket P \rrbracket(V_0) \vee \llbracket Q \rrbracket(V_0) & 
\end{array}$$

■ **Figure 14** Definition of iGPS assertions.

We benefit greatly from the support offered by the surrounding logic. As a result, the model is extremely simple, with the lion’s share of connectives being translated in a purely structural way and the remaining ones making direct use of ambient Iris connectives. The most interesting encodings are those of Hoare triples, non-atomics, escrows, and protocols. We now discuss these in more detail.

The model of Hoare triples embodies the intuition behind our encoding of iGPS assertions as view predicates: the view at which we operate is that of the local thread. In the encoding, we achieve this by quantifying over a view  $V'$  and tying  $V'$  to the physical view of the thread  $\pi$  via the  $\text{Seen}(\pi, V')$  assertion and to the original pre-condition  $P$ , which is required to hold at  $V'$ . As the thread’s physical view may evolve during the execution of the expression  $e$ , the triple returns an extended view  $V'' \sqsupseteq V'$  and the corresponding  $\text{Seen}(\pi, V'')$  assertion, together with the post-condition  $Q$ , which is guaranteed to be valid at  $V''$ .

The encoding of non-atomics is particularly simple due to the properties of the  $\text{Hist}$  assertion. As all writes in the history have to be *mo*-after the most recent non-atomic write, the history becomes a singleton when the location is used non-atomically. To tie the local view  $V$  to the view of the non-atomic write  $V_{\text{na}}$ , we simply demand that  $V$  extend  $V_{\text{na}}$ .

We encode escrows with a single, simple invariant, which holds either the guard resource  $P$  or the payload resource  $Q$ . The view  $V_0$  at which the invariant owns one of these resources is the view used to initialize the escrow. Knowing an escrow at a local view  $V$  thus reduces to knowing that  $V_0 \sqsubseteq V$ .

## Protocols

The model of iGPS protocols consists of two parts: a protocol invariant, and local protocol assertions given out to clients. The protocol invariant owns the location’s history as well as the *logical history*  $\Delta$ , which tracks the transition history of protocol states and is always kept in agreement with the location’s history. Additionally, the invariant owns  $\tau_{\text{read}}$  for all writes in the history and  $\tau_{\text{full}}$  for select ones, depending on the protocol’s type. For example, in normal protocols,  $\tau_{\text{full}}$  is only stored for *CAS-able* write events, justifying that only an update can access the full interpretation of the write event that it reads from (see **iGPS-CAS**). Meanwhile, local protocol assertions hold knowledge about the logical history, which gives effectively lower bounds on the current state of the protocol, and by the protocol invariant, indirectly implies knowledge about the location’s history. In the case of single-writer protocols, the writer assertion also owns the exclusive right to change the state. This construction also relies on the authoritative PCM (see §3.1.1). More details are given in **Appendix F**.

## Soundness of iGPS

The soundness of iGPS is expressed by the following theorem.

► **Theorem 1 (Adequacy).** *For any expression  $e$ , physical state  $\sigma'$ , and meta-level predicate on values  $\Phi$ , we have*

$$\begin{aligned} (\vdash \{\top\} e \{v. \Phi(v)\}) &\Rightarrow \forall \mathcal{TS}. \sigma_{\text{init}}; [0 \mapsto e] \rightarrow^* \sigma'; \mathcal{TS} \Rightarrow \\ &(\mathcal{TS}(0) \in \text{Val} \Rightarrow \Phi(\mathcal{TS}(0))) \\ &\wedge \forall \rho \in \text{dom}(\mathcal{TS}). \mathcal{TS}(\rho) \in \text{Val} \vee \exists \sigma'', e'', e''_f. \mathcal{TS}(\rho) \rightarrow^\rho \sigma'', e'', e''_f \end{aligned}$$

The theorem connects iGPS program specifications ( $\vdash \{\top\} e \{v. \Phi(v)\}$ ) and the program’s possible executions, and provides two guarantees:

1. If the original thread with id 0 terminates with a value  $v$ , we know that  $\Phi(v)$  holds.
2. For any pair of a state  $\sigma'$  and a threadpool  $\mathcal{TS}$  reachable from the initial state  $\sigma_{\text{init}}$  (see [Definition 2](#)) and the initial threadpool  $[0 \mapsto e]$ , we know that any thread  $\rho$  in  $\mathcal{TS}$  either has terminated with a value or can still be reduced in the state  $\sigma'$ .

The proof follows from the adequacy theorem of Iris.

## 5 Other Contributions

In our Coq development accompanying this paper, we make several additional contributions that we briefly summarize here.

### An RSL encoding

Using the same base logic from [§3.2.1](#), we have mechanized iRSL, a higher-order variant of RSL [\[33\]](#). RSL focuses on the message-passing style transferring of resources through release-write/acquire-read pairs. The two main assertions of RSL are  $\text{Rel}(\ell, Q)$  and  $\text{Acq}(\ell, Q)$ , representing the permission to write to and read from  $\ell$ , respectively. The resource  $Q(v)$  is released by writing  $v$  to  $\ell$ , and then acquired by reading  $v$  from  $\ell$ . Consequently, to support this MP-like mechanism, the encoding’s model shares a great deal with the model of iGPS, namely the full vs. read interpretation construction for protocols.

One particular feature of RSL, however, demands special attention. Although simpler than GPS, RSL has the extra ability to split the receiver predicate  $Q$  into *smaller* predicates, so that different acquire reads of the same value  $v$  can acquire different parts of the transferred resource:  $\text{Acq}(\ell, \lambda v. Q_1(v) * Q_2(v)) \Rightarrow \text{Acq}(\ell, Q_1) * \text{Acq}(\ell, Q_2)$ . It is not obvious how to prove this sound when the splitting is completely arbitrary. Fortunately, a similar pattern, called the *barrier* pattern, has been addressed by Jung *et al.* [\[12\]](#), who propose the mechanism of “higher-order ghost state” to support such splitting. Our iRSL model basically extends Jung *et al.*’s barrier proof with a more complex (Iris) protocol to carefully manage resource splitting.

The encoding in Iris also provides us with useful extensions to the logic. Without extra work, iRSL naturally supports PCM-based ghost state and higher-order assertions, which were not available in the original RSL. The encoding shows that our approach has the right, reusable foundations to construct different logics for RA+NA.

In our RSL encoding, assertions are encoded as view predicates and proof rules are proven sound with respect to the base logic—in the same way as our GPS encoding. This allows us to soundly combine RSL and GPS reasoning principles in the same proof *at no additional cost*. It is even possible to design iGPS protocols whose state interpretation mentions iRSL assertions and vice versa. Of course, at a single point in time a location can only be governed by either iGPS or iRSL, as they represent incompatible modes of ownership transfer.

## Allocation and deallocation

We have also incorporated support for memory allocation and deallocation into our RA+NA operational semantics. Since C11 is not clear about the semantics of allocation and deallocation, we take the liberty of defining them as reasonably as possible: in short, allocation and deallocation behave as non-atomic writes with special values  $A$  and  $D$ , respectively.

## Fractional protocols

So far, all protocols presented are permanent: once the protocols are established, they govern their locations forever. This poses two interesting questions: 1) Can we change the protocol which governs a location? and 2) How can we deallocate a location governed by a protocol? To support these features, we derive, with little modification to the iGPS model, *fractional protocols*, whose protocol assertions also assert the permission to even *use* the protocol. Initially, a protocol  $\tau$  for a location  $\ell$  will be established with the full fraction, and then it will be distributed to those who want to use  $\tau$ . Later, when the full fraction is recollected, one can *disable* the protocol (since no one else can use it), regain the *raw* ownership of  $\ell$ , and then deallocate  $\ell$ —or more interestingly, establish a new protocol for  $\ell$ ! These *recollectable* protocols open up possibilities for verifying programs that do custom memory reclamation, *e.g.*, RCU (see below). In the current Coq development, we have created fractional versions of both normal and single-writer protocols.

## Mechanization and Case Studies

Our Coq mechanization employs a shallow embedding of iGPS (and iRSL), making critical use of the Iris Proof Mode [16]. In its current form, the proof mode is specific to the algebra of Iris and offers no additional support for embedded logics like our encoding of iGPS assertions. There are two consequences to this: 1) Unlike in the paper presentation of iGPS, where thread views are completely hidden in the (Iris) model of the logic, in our Coq proofs thread views are visible. However, they are also unobtrusive: all assertions in a given proof context always hold at the current thread’s local view, and the view only changes when the thread takes a step. Thus, while the views are visible, they are always manipulated and kept in sync in a very straightforward way, which we mostly automate with a set of simple tactics. 2) iGPS assertions cannot always be manipulated directly by the proof mode. We sometimes have to unfold our embedding of iGPS assertions—but not in the statements of our lemmas and theorems—to make explicit the underlying Iris assertions so that the proof mode can operate on them. As our embedding is very simple, this has little additional overhead. In particular, all lemmas and theorems stated at the iGPS level remain easily applicable even to the unfolded definitions at the Iris level.

We have verified all of the standard examples that have been proven in previous work. The simplest of these is the *spin-lock* example, proven in iRSL. More interestingly, using iGPS, we have also mechanized the *message passing*, *circular buffer*, *bounded ticket lock*, and *Michael-Scott lock-free queue* examples, which were only verified by hand in the original GPS paper. We have also verified a variant of the *read-copy update* (RCU) technique employed in the Linux kernel, following the proof of Tassarotti *et al.* [30]. The RCU proof is the most substantial example in iGPS, which simplifies the original proof in GPS by using fractional single-writer protocols that allow garbage collection. To our knowledge, our development provides the very first mechanized proofs of the circular buffer, bounded ticket lock, Michael-Scott queue, and RCU examples in a weak-memory setting.

## 6 Related Work

This paper demonstrates one of the first major applications of the Iris framework. Other recent applications include Krebbers *et al.* [16], who developed the interactive Coq proof mode for Iris that we rely on heavily in this paper, and Krogh-Jespersen *et al.* [17], who use Iris to encode a logical-relations model of a relational model of a type-and-effect system for a higher-order, concurrent programming language. Neither of those papers considers weak memory.

There are a number of program logics for weak memory models [27, 4, 1, 19], some of which have mechanized soundness proofs [33, 7, 25], but none of which provide real support for mechanized proofs of weak-memory programs in the way that we do.

FSL++ [8], an extension of FSL [7] (with ghost state) and RSL (with relaxed accesses), was used to mechanize a proof of an implementation of atomic reference counters based on the one in Rust’s `Arc` library. The proof is done by applying the basic laws of separation logic, resulting in painful manual work. Our approach alleviates a great deal of such tedium using the Iris proof mode. As of now, however, iGPS cannot reason about relaxed accesses.

More recently, RSL, FSL, and FSL++ have been encoded in Viper [21] to provide an automated verification approach to weak memory programs [28]. The encodings, however, axiomatize all proof rules without providing soundness guarantees, and are specific to the style of RSL and its FSL descendants. Our base logic, in contrast, is not tied to any specific surface logic and can be used to develop and prove sound different surface logics. It remains to be seen if the more expressive GPS protocols can be encoded in Viper.

iGPS is based on the GPS logic [32] and supports all reasoning mechanisms of GPS. However, the exact rules of iGPS differ in various small ways from the original ones in GPS. These differences stem from pragmatic choices made to simplify the soundness proof of iGPS. A particularly noteworthy difference from GPS appears in the premises of the **iGPS-READ** and **iGPS-WRITE** rules: the split of the protocol interpretations  $\tau$  into  $\tau_{\text{read}}$  and  $\tau_{\text{full}}$ . We show the original GPS rules below for comparison.

$$\begin{array}{c}
 \text{GPS-READ} \\
 \frac{\forall s' \sqsupseteq s. P * \tau(s', v) \Rightarrow \Box Q}{\left\{ \boxed{\ell : s} \boxed{\tau} * P \right\} \ell_{[\text{at}]} \left\{ v, \boxed{\ell : s'} \boxed{\tau} * Q \right\}} \\
 \\
 \text{GPS-WRITE} \\
 \frac{\forall s' \sqsupseteq s. P * \tau(s', v) \Rightarrow s' \sqsubseteq s'' \quad P \Rightarrow \tau(s'', w) * Q}{\left\{ \boxed{\ell : s} \boxed{\tau} * P \right\} \ell_{[\text{at}]} := w \left\{ \boxed{\ell : s''} \boxed{\tau} * Q \right\}}
 \end{array}$$

The interpretation  $\tau(s, v)$  in these two GPS rules is equivalent to  $\tau_{\text{full}}(s, v)$  in iGPS. In GPS, the user can gain access to the full interpretation of the “current” protocol state ( $s'$ ), but only to obtain some *knowledge*, not to consume (*i.e.*, transfer out of the protocol) any non-duplicable resources owned by that interpretation. This is enforced in **GPS-READ** by guarding  $Q$  with an always modality  $\Box$ , and in **GPS-WRITE** by requiring the user to establish the interpretation of the new write using only the local resource  $P$ . In contrast, iGPS does not provide the user with the full interpretation, but only the read interpretation  $\tau_{\text{read}}$ . This weakens, for example, the **iGPS-WRITE** rule in comparison with **GPS-WRITE**, because the user cannot use  $\tau_{\text{full}}$  to show  $s' \sqsubseteq s''$ .

The reason for this discrepancy between GPS and iGPS is that the soundness proof of GPS reasons about an entire program execution graph at once. With its bird’s-eye view of the entire execution, GPS can, for the duration of a step, assemble resources that have been transferred elsewhere in the graph to re-construct  $\tau_{\text{full}}$  for the user. The soundness of iGPS, on the other hand, is established in a simpler and more local manner, without involving reasoning about the full execution of the program. We avoid the global soundness argument of GPS and instead provide a pragmatic solution which—judging from our success in porting

GPS examples—is effectively as strong as GPS and, at the same time, makes for a very simple soundness proof: we maintain the duplicable  $\tau_{\text{read}}$  for all (past) events and can thus easily provide it to the client of **iGPS-READ**.

Essentially, the reason our rules are as effective (if not more so) than those of GPS is that GPS provides one-size-fits-all rules, which are applicable to both programs with write-write races and those without, whereas we provide special support for the common case where there are no write-write races. For programs without those races, the rules provided by GPS are quite cumbersome to use and often require additional ghost state for bookkeeping. iGPS instead supports an optimized write rule for the common case in which there are no such races, via single-writer protocols. For the remaining cases, the rather simple-minded rule **iGPS-WRITE** appears to suffice in all the examples we have considered thus far.

Our operational semantics for RA draws heavily on Lahav *et al.*'s semantics for SRA [18]—a stronger variant of RA, which is equivalent to it in the absence of write-write races. SRA was developed to provide an intuitive operational characterization which is as efficiently implementable as RA. However, as we observe here, moving to an operational characterization does not in fact require any strengthening of the RA semantics (even the slight strengthening of SRA). The main difference between the operational semantics of SRA and the one we give for RA is that writes in SRA always take a globally maximal timestamp, whereas in RA they need not do so. The canonical example demonstrating this difference is the 2+2W example (see Lahav *et al.* [18] for more details).

Going beyond Lahav *et al.*, we offer the first operational account of the interaction of RA and non-atomic accesses. Our semantics corresponds to C11's for programs that do not mix atomic RMW operations and non-atomic reads at the same location. We feel this is a reasonable restriction, given that C11's treatment of programs mixing atomic and non-atomic accesses is already known to be problematic [2]. Our semantics does not correspond to C11's for arbitrary programs, as evidenced by the following example:

$$\mathbf{cas}(x, 0, 1) \parallel \begin{array}{l} x_{[\text{at}]} := 2; \\ a := x_{[\text{na}]} \end{array}$$

C11 considers this program racy because, if the CAS succeeds, the first thread's update of  $x$  to 1 and the second thread's non-atomic read of  $x$  are not in a happens-before relation. In contrast, our semantics does not consider this a race because the non-atomic read is always guaranteed to read from the previous write with value 2. We find the C11 semantics for this program to be rather unintuitive, but leave a more thorough investigation of the issue to future work.

Our RA semantics may also be considered a close derivative of Kang *et al.*'s “promising” semantics [14], which is geared toward solving a broader problem with the full C11 model (the so-called “out-of-thin-air” problem). We look forward to using Iris to construct program logics for this promising semantics.

## Acknowledgements

We would like to thank Mohit Vyas for spotting a mistake in our original proof of correspondence to C11, and Mark Batty for helpful conversations.

This research was supported in part by a European Research Council (ERC) Consolidator Grant for the project “RustBelt”, funded under the European Union's Horizon 2020 Framework Programme (grant agreement no. 683289).

## References

- 1 Tatsuya Abe and Toshiyuki Maeda. Observation-based concurrent program logic for relaxed memory consistency models. In *APLAS*, pages 63–84, 2016.
- 2 Mark Batty, Kayvan Memarian, Kyndylan Nienhuis, Jean Pichon-Pharabod, and Peter Sewell. The problem of programming language concurrency semantics. In *ESOP*, 2015.
- 3 Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber. Mathematizing C++ concurrency. In *POPL*, pages 55–66, 2011.
- 4 Richard Bornat, Jade Alglave, and Matthew J. Parkinson. New lace and arsenic: adventures in weak memory with a program logic. *CoRR*, abs/1512.01416, 2015.
- 5 Pedro da Rocha Pinto, Thomas Dinsdale-Young, and Philippa Gardner. TaDA: A logic for time and data abstraction. In *ECOOP*, 2014.
- 6 T. Dinsdale-Young, M. Dodds, P. Gardner, M. Parkinson, and V. Vafeiadis. Concurrent abstract predicates. In *ECOOP 2010*, volume 6183 of *LNCS*, pages 504–528. Springer, 2010.
- 7 Marko Doko and Viktor Vafeiadis. A program logic for C11 memory fences. In *VMCAI*, volume 9583 of *Lecture Notes in Computer Science*, pages 413–430. Springer, 2016.
- 8 Marko Doko and Viktor Vafeiadis. Tackling real-life relaxed concurrency with FSL++. In *ESOP*, 2017.
- 9 Derek Dreyer. The RustBelt project. <http://plv.mpi-sws.org/rustbelt/>.
- 10 Xinyu Feng. Local rely-guarantee reasoning. In *POPL*, pages 315–327, 2009.
- 11 C. B. Jones. Tentative steps toward a development method for interfering programs. *TOPLAS*, 5(4):596–619, 1983.
- 12 Ralf Jung, Robbert Krebbers, Lars Birkedal, and Derek Dreyer. Higher-order ghost state. In *ICFP*, pages 256–269, 2016.
- 13 Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. Iris: Monoids and invariants as an orthogonal basis for concurrent reasoning. In *POPL*, pages 637–650, 2015.
- 14 Jeehoon Kang, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, and Derek Dreyer. A promising semantics for relaxed-memory concurrency. In *POPL*, pages 175–189, 2017.
- 15 Robbert Krebbers, Ralf Jung, Aleš Bizjak, Jacques-Henri Jourdan, Derek Dreyer, and Lars Birkedal. The essence of higher-order concurrent separation logic. In *ESOP*, 2017.
- 16 Robbert Krebbers, Amin Timany, and Lars Birkedal. Interactive proofs in higher-order concurrent separation logic. In *POPL*, pages 205–217, 2017.
- 17 Morten Krogh-Jespersen, Kasper Svendsen, and Lars Birkedal. A relational model of types-and-effects in higher-order concurrent separation logic. In *POPL*, pages 218–231, 2017.
- 18 Ori Lahav, Nick Giannarakis, and Viktor Vafeiadis. Taming release-acquire consistency. In *POPL*, POPL 2016, pages 649–662. ACM, 2016.
- 19 Ori Lahav and Viktor Vafeiadis. Owicki-Gries reasoning for weak memory models. In *Automata, Languages, and Programming, ICALP 2015*, volume 9135 of *LNCS*, pages 311–323. Springer, 2015.
- 20 Leslie Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Computers*, 28(9):690–691, 1979.
- 21 Peter Müller, Malte Schwerhoff, and Alexander J. Summers. Viper: A verification infrastructure for permission-based reasoning. In *VMCAI*, pages 41–62, 2016.
- 22 Aleksandar Nanevski, Ruy Ley-Wild, Ilya Sergey, and Germán Andrés Delbianco. Communicating state transition systems for fine-grained concurrent resources. In *ESOP*, 2014.
- 23 Peter W. O’Hearn. Resources, concurrency, and local reasoning. *Theor. Comput. Sci.*, 375(1-3):271–307, 2007.
- 24 John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *LICS*, 2002.

- 25 Tom Ridge. A rely-guarantee proof system for x86-TSO. In *VSTTE 2010*, volume 6217 of *LNCS*, pages 55–70. Springer, 2010.
- 26 Ilya Sergey, Aleksandar Nanevski, and Anindya Banerjee. Mechanized verification of fine-grained concurrent programs. In *PLDI*, pages 77–87, 2015.
- 27 Filip Sieczkowski, Kasper Svendsen, Lars Birkedal, and Jean Pichon-Pharabod. A separation logic for fictional sequential consistency. In *ESOP 2015*, volume 9032 of *LNCS*, pages 736–761. Springer, 2015.
- 28 Alexander Summers. Personal communication, 2017.
- 29 Kasper Svendsen and Lars Birkedal. Impredicative concurrent abstract predicates. In *ESOP*, 2014.
- 30 Joseph Tassarotti, Derek Dreyer, and Viktor Vafeiadis. Verifying read-copy-update in a logic for weak memory. In *36th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2015*, pages 110–120. ACM, 2015.
- 31 Aaron Turon, Derek Dreyer, and Lars Birkedal. Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency. In *ICFP*. ACM, 2013.
- 32 Aaron Turon, Viktor Vafeiadis, and Derek Dreyer. GPS: Navigating weak memory with ghosts, protocols, and separation. In *OOPSLA*, OOPSLA 2014, pages 691–707. ACM, 2014.
- 33 Viktor Vafeiadis and Chinmay Narayan. Relaxed separation logic: A program logic for C11 concurrency. In *OOPSLA 2013*, pages 867–884. ACM, 2013.
- 34 Viktor Vafeiadis and Matthew Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR 2007*, volume 4703 of *LNCS*, pages 256–271. Springer, 2007.



## A Operational Semantics for RA+NA

We assume that programs do not include atomic update (RMW) to a location that is also read non-atomically.

### A.1 Declarative Semantics

We define threadpool reduction that generates *traces*:

$$\begin{array}{c}
 \text{TRACE-RED-SILENT} \\
 \frac{\mathcal{TS}(\pi) \rightarrow e, \text{nil}}{\mathcal{TS} \xrightarrow{\epsilon}^{\pi} \mathcal{TS}[\pi \mapsto e]}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{TRACE-RED-MEM} \\
 \frac{\mathcal{TS}(\pi) \xrightarrow{\epsilon} e, \text{nil}}{\mathcal{TS} \xrightarrow{\epsilon}^{\pi} \mathcal{TS}[\pi \mapsto e]}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{TRACE-RED-FORK} \\
 \frac{\mathcal{TS}(\pi) \xrightarrow{\langle \text{Fork}, \rho \rangle} e, e_f}{\mathcal{TS} \xrightarrow{\langle \text{Fork}, \rho \rangle}^{\pi} \mathcal{TS}[\pi \mapsto e] \uplus [\rho \mapsto e_f]}
 \end{array}$$

We write  $\mathcal{TS} \Rightarrow^{\pi} \mathcal{TS}'$  if  $\mathcal{TS} \xrightarrow{x}^{\pi} \mathcal{TS}'$  for some transition label  $x$ ;  $\mathcal{TS} \xrightarrow{x} \mathcal{TS}'$  if  $\mathcal{TS} \xrightarrow{x}^{\pi} \mathcal{TS}'$  for some thread identifier  $\pi$ ; and  $\mathcal{TS} \Rightarrow \mathcal{TS}'$  if  $\mathcal{TS} \xrightarrow{x}^{\pi} \mathcal{TS}'$  for some transition label  $x$  and thread identifier  $\pi$ . A threadpool is called *final* if  $\mathcal{TS}(\pi) \in \text{Val}$  for every  $\pi \in \text{dom}(\mathcal{TS})$ .

► **Definition 2.** A *trace* is a sequence of pairs  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$ . We say that  $tr = \langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$  is a trace of an expression  $e$  if

$$[0 \mapsto e] \xrightarrow{\epsilon}^* \xrightarrow{\varepsilon_1}^{\pi_1} \xrightarrow{\epsilon}^* \dots \xrightarrow{\epsilon}^* \xrightarrow{\varepsilon_n}^{\pi_n} \xrightarrow{\epsilon}^* \mathcal{TS}$$

for some thread  $\pi$  and threadpool  $\mathcal{TS}$ . When  $\mathcal{TS}$  is final, we call  $tr$  a *full trace*.

- **Lemma 3 (Receptiveness).** Let  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$  be a trace of an expression  $e$ .
- If  $\varepsilon_n = \langle \text{Read}_{\alpha}, \ell, v \rangle$ , then for every value  $v'$ ,  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \langle \text{Read}_{\alpha}, \ell, v' \rangle, \pi_n \rangle$  is a trace of  $e$ .
  - If  $\varepsilon_n = \langle \text{Update}, \ell, v_o, v_n \rangle$ , then for every value  $v'_o$ , either  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \langle \text{Read}_{\text{at}}, \ell, v'_o \rangle, \pi_n \rangle$  is a trace of  $e$ , or there exists a value  $v'_n$  such that  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \langle \text{Update}, \ell, v'_o, v'_n \rangle, \pi_n \rangle$  is a trace of  $e$ .

► **Definition 4.** A trace  $tr = \langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$  induces a partial order  $\text{sb}(tr)$  (*sequenced-before*) on  $\{1, \dots, n\}$  that is given by:

$$\begin{array}{c}
 \frac{k_1 < k_2 \quad \pi_{k_1} = \pi_{k_2}}{\langle k_1, k_2 \rangle \in \text{sb}(tr)}
 \qquad
 \frac{k_1 < k_2 \quad \varepsilon_{k_1} = \langle \text{Fork}, \pi_{k_2} \rangle}{\langle k_1, k_2 \rangle \in \text{sb}(tr)} \\
 \\
 \frac{\langle k_1, k_2 \rangle \in \text{sb}(tr) \quad \langle k_2, k_3 \rangle \in \text{sb}(tr)}{\langle k_1, k_3 \rangle \in \text{sb}(tr)}
 \end{array}$$

- **Lemma 5.** Let  $tr$  be a trace of an expression  $e$ . Then:
- Any prefix of  $tr$  is a trace of  $e$ .
  - Any permutation  $tr'$  of  $tr$  that respects  $\text{sb}(tr)$  (*i.e.*,  $\text{sb}(tr') = \text{sb}(tr)$ ) is a trace of  $e$ .

► **Definition 6.** An *execution graph*  $G$  consists of:

1. a finite set of nodes  $\mathbf{E} \subseteq \mathbb{N}$  (called *events*).
2. a function  $\text{tid}$  assigning a thread identifier to every event in  $\mathbf{E}$ . We denote by  $\mathbf{E}^{\pi}$  the set of events in  $a \in \mathbf{E}$  with  $\text{tid}(a) = \pi$ .

3. a function  $\text{lab}$  assigning a *label* to every event in  $\mathbf{E}$ . Labels take one of the following forms:  $\langle \text{Read}_\alpha, \ell, v \rangle$ ,  $\langle \text{Write}_\alpha, \ell, v \rangle$ ,  $\langle \text{Update}, \ell, v_o, v_n \rangle$ , or  $\langle \text{Fork}, \pi \rangle$ , where  $\alpha \in \{\text{at}, \text{na}\}$  is the access mode,  $\ell$  is the location accessed, and  $v, v_o, v_n$  are read/written values, and  $\pi$  is a thread identifier. The function  $\text{lab}$  naturally induces functions  $\text{typ}$ ,  $\text{mod}$ ,  $\text{loc}$ ,  $\text{val}_r$ , and  $\text{val}_w$  that return (when applicable) the type ( $\text{FORK}, \text{R}, \text{W}$  or  $\text{U}$ ), access mode (in particular,  $\text{mod}(a) = \text{at}$  if  $\text{typ}(a) = \text{U}$ ), location, and read/written values of an event. For  $\mathbf{T} \in \{\text{S}, \text{R}, \text{W}, \text{U}\}$ ,  $\mathbf{T}$  denotes the set  $\{a \in \mathbf{E} \mid \text{typ}(a) = \mathbf{T}\}$ . We also concatenate the event sets notations, use subscripts to denote the accessed location, and superscripts for access modes (e.g.,  $\text{WU}_\ell^{\text{at}}$  denotes all events  $a \in \text{W} \cup \text{U}$  with  $\text{loc}(a) = \ell$  and  $\text{mod}(a) = \text{at}$ ).
4. a strict partial order  $\text{sb} \subseteq \mathbf{E} \times \mathbf{E}$ , called *sequenced-before*, (a.k.a. *program order*).
5. a binary relation  $\text{rf} \subseteq \mathbf{E} \times \mathbf{E}$  called *reads-from*, satisfying: (i)  $a \in \text{WU}$  and  $b \in \text{RU}$  for every  $\langle a, b \rangle \in \text{rf}$ ; (ii)  $\text{loc}(a) = \text{loc}(b)$  and  $\text{val}_w(a) = \text{val}_r(b)$  for every  $\langle a, b \rangle \in \text{rf}$ ; and (iii)  $a_1 = a_2$  whenever  $\langle a_1, b \rangle, \langle a_2, b \rangle \in \text{rf}$ .
6. a strict partial order  $\text{mo}$  on  $\text{WU}$ , called *modification order*, which is a disjoint union of relations  $\{\text{mo}_\ell\}_{\ell \in \text{Loc}}$ , such that each  $\text{mo}_\ell$  is a strict total order on  $\text{WU}_\ell$ .

In what follows, to resolve ambiguities, we may include a prefix “ $G$ .” to refer to the components of an execution graph  $G$ .

► **Definition 7.** An execution graph  $G$  follows a trace  $tr = \langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$  if the following hold:

- $\mathbf{E}$  consists of  $n$  events  $a_1, \dots, a_n$  with  $\text{lab}(a_k) = \varepsilon_k$  and  $\text{tid}(a_k) = \pi_k$  for every  $1 \leq k \leq n$ .
- $\text{sb} = \{\langle a_{k_1}, a_{k_2} \rangle \mid \langle k_1, k_2 \rangle \in \text{sb}(tr)\}$ .

We call  $G$  an *execution graph of an expression  $e$*  if  $G$  follows some trace of  $e$ .

Note that  $\text{rf}$  and  $\text{mo}$  are not restricted.

► **Definition 8.** Given an execution graph  $G$ , the relations  $\text{sw}$  (*synchronization*) and  $\text{hb}$  (*happens-before*) are given by:

$$\text{sw} = \{\langle a, b \rangle \in \text{rf} \mid \text{mod}(a) = \text{mod}(b) = \text{at}\} \quad \text{hb} = (\text{sb} \cup \text{sw})^+$$

► **Definition 9** (Consistency). An execution graph  $G$  is *consistent* if the following hold:

- For every  $b \in \text{RU}$ , if there exists some  $a \in \text{WU}_{\text{loc}(b)}$  such that  $\langle a, b \rangle \in \text{hb}^?$ ;  $\text{sb}$ , then  $\langle c, b \rangle \in \text{rf}$  for some  $c \in \mathbf{E}$ .
- $\text{sb} \cup \text{rf}$  is acyclic.
- $\text{rf}; \text{hb}$  is irreflexive.
- $\text{mo}; \text{rf}; \text{hb}$  is irreflexive
- $\text{mo}; \text{hb}$  is irreflexive.
- $\text{mo}; \text{hb}; \text{rf}^{-1}$  is irreflexive.
- $\text{mo}; \text{rf}; \text{hb}; \text{rf}^{-1}$  is irreflexive.
- $(\text{rf}; [\text{U}]) \cap (\text{mo}; \text{mo}) = \emptyset$ .

► **Definition 10.** A trace  $tr$  is *allowed* for an expression  $e$  if  $tr$  is a trace of  $e$  and there exists a consistent execution graph  $G$  that follows  $tr$ .

► **Definition 11** (Conflict). Two events  $a$  and  $b$  are called *conflicting* in an execution graph  $G$  if  $a, b \in \mathbf{E}$ ,  $\{\text{typ}(a), \text{typ}(b)\} \cap \{\text{W}, \text{U}\} \neq \emptyset$ ,  $a \neq b$ , and  $\text{loc}(a) = \text{loc}(b)$ .

► **Definition 12** (Races). A pair  $\langle a, b \rangle$  is called a *race* in  $G$  if  $a$  and  $b$  are conflicting events in  $G$ , and  $\langle a, b \rangle \notin \text{hb} \cup \text{hb}^{-1}$ . An execution graph  $G$  is called *racy* if there is some race  $\langle a, b \rangle$  in  $G$  with  $\text{na} \in \{\text{mod}(a), \text{mod}(b)\}$ .

► **Definition 13** (Initialization). A location  $\ell$  is called *uninitialized* in an execution graph  $G$  if there exists some  $b \in \text{RU}_\ell$  such that  $\langle a, b \rangle \notin \text{hb}^?; \text{sb}$  for every  $a \in \text{WU}_\ell$ . An execution graph  $G$  is called *uninitialized* if some location is uninitialized in  $G$ .

► **Definition 14** (Buggy expression). An expression  $e$  is called *buggy* if some consistent execution graph of  $e$  is either racy or uninitialized.

The following definition and lemma are useful in the sequel.

► **Definition 15.** Given an execution graph  $G$ , we denote by  $\text{mo}^{\text{na}}$  the relation  $\text{mo} \cap ((W^{\text{na}} \times E) \cup (E \times W^{\text{na}}))$ .

► **Lemma 16.** Let  $G$  be a non-racy consistent execution graph. Then,  $\text{hb} = (\text{sb} \cup \text{rf} \cup \text{mo}^{\text{na}})^+$ .

**Proof.** Since  $G$  is not racy, we have  $\text{rf} \cup \text{mo}^{\text{na}} \subseteq \text{hb} \cup \text{hb}^{-1}$ . Since  $G$  is consistent,  $\text{rf}; \text{hb}$  and  $\text{mo}^{\text{na}}; \text{hb}$  are irreflexive, and so  $\text{rf} \cup \text{mo}^{\text{na}} \subseteq \text{hb}$ . It easily follows that  $\text{hb} = (\text{sb} \cup \text{rf} \cup \text{mo}^{\text{na}})^+$ . ◀

### A.1.1 Alternative declarative semantics

Next, we present an alternative declarative semantics in which  $\text{hb}$  is substituted by  $(\text{sb} \cup \text{rf})^+$ , and races require stronger conditions as defined below. This semantics simplifies our proof below of the equivalence to the operational semantics.

► **Definition 17** (Strong consistency). *Strong consistency* is defined exactly as consistency with  $(\text{sb} \cup \text{rf})^+$  instead of  $\text{hb}$ .

Observing that  $\text{hb} \subseteq (\text{sb} \cup \text{rf})^+$  in every execution graph, it is easy to see that strong consistency implies consistency. In addition, strong consistency can be simplified as follows:

► **Lemma 18.** An execution graph  $G$  is strongly consistent iff the following hold:

- For every  $b \in \text{RU}$ , if there exists some  $a \in \text{WU}_{\text{loc}(b)}$  such that  $\langle a, b \rangle \in (\text{sb} \cup \text{rf})^*$ ;  $\text{sb}$ , then  $\langle c, b \rangle \in \text{rf}$  for some  $c \in \text{E}$ .
- $\text{sb} \cup \text{rf}$  is acyclic.
- $\text{mo}; (\text{sb} \cup \text{rf})^+$  is irreflexive.
- $\text{mo}; (\text{sb} \cup \text{rf})^+; \text{rf}^{-1}$  is irreflexive.
- $(\text{rf}; [\text{U}]) \cap (\text{mo}; \text{mo}) = \emptyset$ .

► **Definition 19** (Strongly racy). An execution graph  $G$  is called *strongly racy* if there exist  $a, b \in \text{E}$  such that the following hold:

- $a \neq b$ .
- $\text{loc}(a) = \text{loc}(b)$ .
- $\langle a, b \rangle \notin \text{mo}^?$ ;  $(\text{sb} \cup \text{rf})^+$  and  $\langle b, a \rangle \notin (\text{sb} \cup \text{rf})^+$ .
- Either  $(a \in W^{\text{at}}$  and  $b \in R^{\text{na}})$  or  $a \in W^{\text{na}}$ .

► **Definition 20** (Strong initialization). A location  $\ell$  is called *strongly uninitialized* in an execution graph  $G$  if there exists some  $b \in \text{RU}_\ell$  such that  $\langle a, b \rangle \notin (\text{sb} \cup \text{rf})^*; \text{sb}$  for every  $a \in \text{WU}_\ell$ . An execution graph  $G$  is called *strongly uninitialized* if some location is strongly uninitialized in  $G$ .

► **Notation 1.** Let  $G$  be an execution graph. For a set  $E \subseteq \text{E}$ , we denote by  $G \cap E$  the execution graph  $G'$  given by:  $G'.\text{E} = E$ ,  $G'.\text{tid} = G.\text{tid}|_E$ ,  $G'.\text{lab} = G.\text{lab}|_E$ ,  $G'.\text{sb} = G.\text{sb} \cap (E \times E)$ ,  $G'.\text{rf} = G.\text{rf} \cap (E \times E)$ , and  $G'.\text{mo} = G.\text{mo} \cap (E \times E)$ . Given  $a \in \text{E}$ , we write  $G \setminus \{a\}$  for  $G \cap (E \setminus \{a\})$ . Given  $R \subseteq \text{E} \times \text{E}$ , by *R-prefix* of  $G$  we mean any execution graph  $G'$  such that  $G' = G \cap E$  for some set  $E$  satisfying  $a \in E$  whenever  $b \in E$  and  $\langle a, b \rangle \in R$ .

► **Lemma 21** (Adding writes). Let  $G$  be an execution graph and  $a \in \text{W}$  be an  $(\text{sb} \cup \text{rf})^+$ -maximal event in  $\text{E}$ . Suppose that  $G \setminus \{a\}$  is strongly consistent. Let  $G'$  be the execution graph that is obtained from  $G$  by making  $a$  an  $\text{mo}$ -maximal event. Then,  $G'$  is strongly consistent.

► **Lemma 22** (Adding reads). Let  $G$  be an execution graph and  $a \in \mathbf{R}$  be an  $(\text{sb} \cup \text{rf})^+$ -maximal event in  $\mathbf{E}$ . Suppose that  $G \setminus \{a\}$  is strongly consistent, and that  $\langle b, a \rangle \in (\text{sb} \cup \text{rf})^*$ ;  $\text{sb}$  for some  $b \in \mathbf{WU}_{\text{loc}(a)}$ . Then, there exists a strongly consistent execution graph  $G'$  that satisfies the following:

- $G'.\mathbf{E} = G.\mathbf{E}$ ,  $G'.\text{tid} = G.\text{tid}$ ,  $G'.\text{sb} = G.\text{sb}$ , and  $G'.\text{mo} = G.\text{mo}$ .
- $G'.\text{lab}(b) = G.\text{lab}(b)$  for every  $b \in G.\mathbf{E} \setminus \{a\}$  and  $G'.\text{lab}(a) = \langle \text{Read}_{G.\text{mod}(a)}, G.\text{loc}(a), v \rangle$  for some value  $v$ .
- $G'.\text{rf} \cap (G.\mathbf{E} \times (G.\mathbf{E} \setminus \{a\})) = G.\text{rf} \cap (G.\mathbf{E} \times (G.\mathbf{E} \setminus \{a\}))$ .
- $G'.\text{rf} \cap (G.\mathbf{E} \times \{a\}) \subseteq (G'.\text{sb} \cup G'.\text{rf})^*$ ;  $G'.\text{sb}$ .

**Proof.** Let  $B = \{b \in \mathbf{WU}_{\text{loc}(a)} \mid \langle b, a \rangle \in (\text{sb} \cup \text{rf})^*; \text{sb}\}$ , and let  $b = \max_{\text{mo}} B$ . Let  $G'$  be the execution graph obtained from  $G$  by replacing the value read at  $a$  by  $\text{val}_{\mathbf{w}}(b)$ , and adding the  $\text{rf}$ -edge  $\langle b, a \rangle$  (instead of the  $\text{rf}$ -edge entering  $a$  in  $G$ , if such exists). It is straightforward to see that  $G'$  is a strongly consistent. ◀

► **Lemma 23** (Adding updates). Let  $G$  be an execution graph and  $a \in \mathbf{U}$  be an  $(\text{sb} \cup \text{rf})^*$ -maximal event in  $\mathbf{E}$ . Suppose that  $G \setminus \{a\}$  is strongly consistent, and that  $\langle b, a \rangle \in (\text{sb} \cup \text{rf})^*$ ;  $\text{sb}$  for some  $b \in \mathbf{WU}_{\text{loc}(a)}$ . Then, there exists a strongly consistent execution graph  $G'$  that satisfies the following:

- $G'.\mathbf{E} = G.\mathbf{E}$ ,  $G'.\text{tid} = G.\text{tid}$ ,  $G'.\text{sb} = G.\text{sb}$ .
- $G'.\text{lab}(b) = G.\text{lab}(b)$  for every  $b \in G.\mathbf{E} \setminus \{a\}$  and  $G'.\text{lab}(a) = \langle \text{Update}, G.\text{loc}(a), v, v' \rangle$  for some values  $v, v'$ .
- $G'.\text{rf} \cap (G.\mathbf{E} \times (G.\mathbf{E} \setminus \{a\})) = G.\text{rf} \cap (G.\mathbf{E} \times (G.\mathbf{E} \setminus \{a\}))$ .
- $G'.\text{rf} \cap (G.\mathbf{E} \times \{a\}) \subseteq ((G'.\text{sb} \cup G'.\text{rf})^+; G'.\text{sb}) \cup (G.\mathbf{WU}^{\text{at}} \times \{a\})$ .
- $G'.\text{mo} \cap ((G.\mathbf{E} \setminus \{a\}) \times (G.\mathbf{E} \setminus \{a\})) = G.\text{mo} \cap ((G.\mathbf{E} \setminus \{a\}) \times (G.\mathbf{E} \setminus \{a\}))$ .

**Proof.** Let  $B = (\mathbf{WU}_{\ell}^{\text{at}} \setminus \{a\}) \cup \{b \in \mathbf{W}_{\ell}^{\text{na}} \mid \langle b, a \rangle \in (\text{sb} \cup \text{rf})^*; \text{sb}\}$ , where  $\ell = \text{loc}(a)$ . Let  $b = \max_{\text{mo}} B$ , and let  $G'$  be the execution graph obtained from  $G$  by (i) replacing the value read at  $a$  by  $\text{val}_{\mathbf{w}}(b)$ ; (ii) adding the  $\text{rf}$ -edge  $\langle b, a \rangle$  (instead of the  $\text{rf}$ -edge entering  $a$  in  $G$ , if such exists); and (iii) making  $a$  to be the  $\text{mo}$ -successor of  $b$ . It is straightforward to see that  $G'$  is a strongly consistent. ◀

► **Lemma 24** (Modifying updates). Let  $G$  be a strongly consistent execution graph, and let  $a \in \mathbf{U}$  be an  $(\text{sb} \cup \text{rf})^+$ -maximal event in  $\mathbf{E}$ . Then, any execution graph  $G'$  that is obtained from  $G$  by changing the value written at  $a$ , or by making  $a$  a read event, is strongly consistent.

► **Lemma 25.** Let  $G$  be a racy consistent execution graph of an expression  $e$ . Then, some strongly consistent execution graph of  $e$  is either strongly racy or strongly uninitialized.

**Proof.** Let  $a_1, \dots, a_n$  be an enumeration of  $\mathbf{E}$  that respects  $\text{sb} \cup \text{rf}$ , and let  $k$  be the minimal index such that  $G \cap \{a_1, \dots, a_k\}$  is racy.

First, we claim that for every event  $a$ , if  $\langle b, a_k \rangle \in (\text{sb} \cup \text{rf} \cup \text{mo}^{\text{na}})^*$ ;  $\text{sb}$ , then  $\langle b, a_k \rangle \in \text{hb}$ . Indeed, suppose that  $\langle b, a_k \rangle \in (\text{sb} \cup \text{rf} \cup \text{mo}^{\text{na}})^*$ ;  $\text{sb}$ . Let  $G' = G \cap \{a_1, \dots, a_{k-1}\}$ . The minimality of  $k$  ensures that  $G'$  is not racy, and by Lemma 16, it follows that  $\langle b, a_k \rangle \in G.\text{hb}^?; G.\text{sb} \subseteq G.\text{hb}$ .

Next, let  $j$  be the minimal index such that  $\langle a_j, a_k \rangle$  is a race in  $G$  and  $\text{na} \in \{\text{mod}(a_j), \text{mod}(a_k)\}$ . Let  $E = \{a \in \mathbf{E} \mid \langle a, a_k \rangle \in (\text{sb} \cup \text{rf})^* \vee \langle a, a_j \rangle \in (\text{sb} \cup \text{rf})^*\}$ , and let  $G_0 = G \cap E$ . Note that  $a_j$  has no outgoing  $\text{sb} \cup \text{rf}$ -edges in  $G_0$  except (possibly) an  $\text{rf}$ -edge to  $a_k$  (otherwise,  $\langle a_j, a_k \rangle \in (G.\text{sb} \cup G.\text{rf})^*; G.\text{sb}$ , and so  $\langle a_j, a_k \rangle \in G.\text{hb}$ , which contradicts the fact that  $\langle a_j, a_k \rangle$  is a race in  $G$ ). It is also easy to verify that  $G_0$  is a strongly consistent execution graph of  $e$ . If  $G_0$  is strongly uninitialized, then we are done. Otherwise, if  $a_k \in G_0.\text{RU}$  then,

we have  $\langle b, a_k \rangle \in (G_0.\text{sb} \cup G_0.\text{rf})^*; G_0.\text{sb}$  for some  $b \in G_0.\text{WU}_{\text{loc}(a_k)}$ . Consider the possible cases:

- $a_k \in \text{R}$ : Then, we must have  $a_j \in \text{W}$ . Let  $G'$  be a strongly consistent execution graph that satisfies the properties mentioned in Lemma 22 applied on  $G_0$  and  $a_k$ . The receptiveness property ensures that  $G'$  is an execution graph of  $e$ . Let  $G''$  be a strongly consistent execution graph that satisfies the properties mentioned in Lemma 21 applied on  $G'$  and  $a_j$ . Since  $G''.\text{E} = G'.\text{E}$ ,  $G''.\text{lab} = G'.\text{lab}$  and  $G''.\text{sb} = G'.\text{sb}$ , we have that  $G''$  is an execution graph of  $e$  as well. It follows that  $\langle a_j, a_k \rangle \notin G''.\text{mo}^?; (G''.\text{sb} \cup G''.\text{rf})^+$ , and  $\langle a_k, a_j \rangle \notin (G''.\text{sb} \cup G''.\text{rf})^+$ . In addition, since  $\langle a_j, a_k \rangle$  is a race in  $G$ , we have  $G''.\text{loc}(a_j) = G''.\text{loc}(a_k)$ . Hence,  $G''$  is a strongly racy strongly consistent execution graph of  $e$ .
- $a_k \in \text{W}$ : Let  $G'$  be a strongly consistent execution graph that satisfies the properties mentioned in Lemma 21 applied on  $G_0$  and  $a_k$ . As in the previous case,  $G'$  is an execution graph of  $e$ . It follows that  $\langle a_k, a_j \rangle \notin G'.\text{mo}^?; (G'.\text{sb} \cup G'.\text{rf})^+$ , and  $\langle a_j, a_k \rangle \notin (G'.\text{sb} \cup G'.\text{rf})^+$ . In addition, since  $\langle a_j, a_k \rangle$  is a race in  $G$ , we have  $G'.\text{loc}(a_j) = G'.\text{loc}(a_k)$ . Hence,  $G'$  is a strongly racy strongly consistent execution graph of  $e$ .
- $a_k \in \text{U}$ : Then, we must have  $a_j \in \text{W}^{\text{na}}$ . Let  $G'$  be a strongly consistent execution graph that satisfies the properties mentioned in Lemma 23 applied on  $G_0$  and  $a_k$ . The receptiveness property ensures that some execution graph  $G^*$  obtained from  $G'$  by possibly making  $a_k$  a read event is an execution graph of  $e$ . By Lemma 24,  $G^*$  is also strongly consistent. Note also that we do not have  $\langle a_j, a_k \rangle \in G^*.\text{rf}$ . Let  $G''$  be a strongly consistent execution graph that satisfies the properties mentioned in Lemma 21 applied on  $G^*$  and  $a_j$ . Since  $G''.\text{E} = G^*.\text{E}$ ,  $G''.\text{lab} = G^*.\text{lab}$  and  $G''.\text{sb} = G^*.\text{sb}$ , we have that  $G''$  is an execution graph of  $e$  as well. It follows that  $\langle a_j, a_k \rangle \notin G''.\text{mo}^?; (G''.\text{sb} \cup G''.\text{rf})^+$ , and  $\langle a_k, a_j \rangle \notin (G''.\text{sb} \cup G''.\text{rf})^+$ . In addition, since  $\langle a_j, a_k \rangle$  is a race in  $G$ , we have  $G''.\text{loc}(a_j) = G''.\text{loc}(a_k)$ . Hence,  $G''$  is a strongly racy strongly consistent execution graph of  $e$ . ◀

► **Lemma 26.** Let  $G$  be an uninitialized consistent execution graph of an expression  $e$ . Then, some strongly consistent execution graph of  $e$  is either strongly racy or strongly uninitialized.

**Proof.** Let  $\ell$  be an uninitialized location in  $G$ , and let  $b \in \text{RU}_\ell$  such that  $\langle a, b \rangle \notin \text{hb}^?; \text{sb}$  for every  $a \in \text{WU}_\ell$ . If  $G$  is racy, then we are done by Lemma 25. Otherwise, by Lemma 16, we have  $\text{hb} = (\text{sb} \cup \text{rf} \cup \text{mo}^{\text{na}})^+$ . Hence,  $\langle a, b \rangle \notin (\text{sb} \cup \text{rf})^*; \text{sb}$  for every  $a \in \text{WU}_\ell$ , and so  $G$  is strongly uninitialized. ◀

► **Corollary 27.** Let  $e$  be an expression. Then:

- If  $e$  is not buggy, then a trace  $tr$  is allowed for  $e$  iff  $tr$  is a trace of  $e$  and there exists a strongly consistent execution graph  $G$  that follows  $tr$ .
- $e$  is buggy iff some strongly consistent execution graph of  $e$  is either strongly racy or strongly uninitialized.

**Proof.** Easily follows from our definitions, Lemma 25, Lemma 26, and Lemma 16:

- Suppose that  $e$  is not buggy. If  $tr$  is a trace of  $e$ , and  $G$  is a strongly consistent execution graph that follows  $tr$ , then  $G$  is also consistent, and, by definition,  $tr$  be a trace that is allowed for  $e$ . For the converse, let  $tr$  be a trace that is allowed for  $e$ . Then, by definition,  $tr$  is a trace of  $e$  and there exists a consistent execution graph  $G$  that follows  $tr$ . Since  $e$  is not buggy,  $G$  is not racy. By Lemma 16, we have  $(G.\text{sb} \cup G.\text{rf})^+ \subseteq G.\text{hb}$ , and hence  $G$  is strongly consistent.

- If some strongly consistent execution graph  $G$  of  $e$  is either strongly racy or strongly uninitialized, then  $G$  is also a consistent execution graph of  $e$  which is either racy or uninitialized, and hence  $e$  is buggy. For the converse, suppose that  $e$  is buggy. By definition, there exists a consistent execution graph  $G$  of  $e$  that is either racy or uninitialized. Then, by Lemma 25 and Lemma 26, some strongly consistent execution graph of  $e$  is either strongly racy or strongly uninitialized. ◀

► **Lemma 28.** Let  $G$  be a strongly consistent execution graph that is not strongly racy. Then,  $\text{mo}^{\text{na}} \subseteq (\text{sb} \cup \text{rf})^+$ .

**Proof.** Let  $\langle a, b \rangle \in \text{mo}^{\text{na}}$ . First, if  $b \in \mathbb{W}^{\text{na}}$ , then since  $G$  is not strongly racy, and  $\text{mo}; (\text{sb} \cup \text{rf})^+$  is irreflexive, we must have  $\langle a, b \rangle \in (\text{sb} \cup \text{rf})^+$ . Otherwise, we have  $a \in \mathbb{W}^{\text{na}}$ . Let  $b_1, \dots, b_n$  be an enumeration of  $\mathbb{E}$  that respects  $\text{sb} \cup \text{rf}$ . We prove by induction on  $k$  that  $\langle a, b_k \rangle \in \text{mo}^{\text{na}}$  implies  $\langle a, b_k \rangle \in (\text{sb} \cup \text{rf})^+$ . Suppose the claim hold for all  $k < m$ , we prove it for  $m$ . Assume that  $\langle a, b_m \rangle \in \text{mo}^{\text{na}}$ . Since  $G$  is not strongly racy and  $\text{mo}; (\text{sb} \cup \text{rf})^+$  is irreflexive, we have  $\langle a, b_m \rangle \in \text{mo}^?; (\text{sb} \cup \text{rf})^+$ . Let  $b_i \in \mathbb{E}$  such that  $\langle a, b_i \rangle \in \text{mo}^?$  and  $\langle b_i, b_m \rangle \in (\text{sb} \cup \text{rf})^+$ . If  $a = b_i$ , then we are done. Otherwise, by the induction hypothesis,  $\langle a, b_i \rangle \in (\text{sb} \cup \text{rf})^+$ , and again it follows that  $\langle a, b_m \rangle \in (\text{sb} \cup \text{rf})^+$ . ◀

## A.2 First operational semantics

We present an operational semantics, equivalent to the declarative RA+NA semantics, in the form of an abstract machine whose states are execution graphs, incrementally growing during the run. The initial execution graph, denoted  $G_0$ , is the empty execution graph.

► **Notation 2.** For two execution graphs  $G$  and  $G'$ , thread identifiers  $\pi, \rho$ , and a label  $\varepsilon$ , we write  $G' \in \text{Add}(G, \pi, \rho, \varepsilon)$  if there exists an event  $a$  such that the following hold:

- $G'.\mathbb{E} = G.\mathbb{E} \uplus \{a\}$ ,  $G'.\text{tid} = G.\text{tid} \cup \{a \mapsto \rho\}$ , and  $G'.\text{lab} = G.\text{lab} \cup \{a \mapsto \varepsilon\}$ .
- $G'.\text{sb} = (G.\text{sb} \cup (G.\mathbb{E}^\pi \times \{a\}))^+$ .
- $G'.\text{rf} \supseteq G.\text{rf}$ .
- $G'.\text{mo} \supseteq G.\text{mo}$  and if  $\varepsilon = \langle \text{Write}_{\text{na}}, -, - \rangle$ , then  $a$  is  $G'.$ mo-maximal.

The machine steps are based on execution graph steps:

$$\begin{array}{c}
 \varepsilon \in \{ \langle \text{Read}_\alpha, \ell, v \rangle, \langle \text{Write}_\alpha, \ell, v \rangle, \langle \text{Update}, \ell, v_o, v_n \rangle \} \\
 \forall a \in \mathbb{W}_\ell^{\text{na}}. \exists b \in \mathbb{E}^\pi. \langle a, b \rangle \in \text{mo}^?; (\text{sb} \cup \text{rf})^* \\
 \varepsilon = \langle \text{Read}_{\text{na}}, \ell, v \rangle \Rightarrow \forall a \in \mathbb{W}_\ell. \exists b \in \mathbb{E}^\pi. \langle a, b \rangle \in \text{mo}^?; (\text{sb} \cup \text{rf})^* \\
 \varepsilon \in \{ \langle \text{Read}_\alpha, \ell, v \rangle, \langle \text{Update}, \ell, v_o, v_n \rangle \} \Rightarrow \exists a \in \mathbb{WU}_\ell, b \in \mathbb{E}^\pi. \langle a, b \rangle \in (\text{sb} \cup \text{rf})^* \\
 G' \in \text{Add}(G, \pi, \pi, \varepsilon) \quad G' \text{ is strongly consistent} \\
 \hline
 G \xrightarrow{\varepsilon, \pi} G'
 \end{array}$$
  

$$\begin{array}{c}
 \text{RACE-I} \\
 \varepsilon \in \{ \langle \text{Read}_\alpha, \ell, v \rangle, \langle \text{Write}_\alpha, \ell, v \rangle, \langle \text{Update}, \ell, v_o, v_n \rangle \} \\
 \exists a \in \mathbb{W}_\ell^{\text{na}}. \forall b \in \mathbb{E}^\pi. \langle a, b \rangle \notin \text{mo}^?; (\text{sb} \cup \text{rf})^* \\
 \hline
 G \xrightarrow{\varepsilon, \pi} \perp_{\text{race}}
 \end{array}$$
  

$$\begin{array}{c}
 \text{RACE-II} \\
 \exists a \in \mathbb{W}_\ell. \forall b \in \mathbb{E}^\pi. \langle a, b \rangle \notin \text{mo}^?; (\text{sb} \cup \text{rf})^* \\
 \hline
 G \xrightarrow{\langle \text{Read}_{\text{na}}, \ell, v \rangle, \pi} \perp_{\text{race}}
 \end{array}$$
  

$$\begin{array}{c}
 \text{UNINITIALIZED} \\
 \varepsilon \in \{ \langle \text{Read}_\alpha, \ell, v \rangle, \langle \text{Update}, \ell, v_o, v_n \rangle \} \\
 \neg \exists a \in \mathbb{WU}_\ell, b \in \mathbb{E}^\pi. \langle a, b \rangle \in (\text{sb} \cup \text{rf})^* \\
 \hline
 G \xrightarrow{\varepsilon, \pi} \perp_{\text{uninit}}
 \end{array}$$

The machine steps are combined with expression reductions and lifted to thread-pool reductions exactly as in [Figure 4](#) (using execution graphs  $G$  instead of physical states  $\sigma$ ), defining reductions of the form  $G; \mathcal{TS} \xrightarrow{x, \pi} G'; \mathcal{TS}'$ . As before, we write  $G; \mathcal{TS} \rightarrow^\pi G'; \mathcal{TS}'$  if  $G; \mathcal{TS} \xrightarrow{x, \pi} G'; \mathcal{TS}'$  for some transition label  $x$ ;  $G; \mathcal{TS} \xrightarrow{x} G'; \mathcal{TS}'$  if  $G; \mathcal{TS} \xrightarrow{x, \pi} G'; \mathcal{TS}'$  for some thread identifier  $\pi$ ; and  $G; \mathcal{TS} \rightarrow G'; \mathcal{TS}'$  if  $G; \mathcal{TS} \xrightarrow{x, \pi} G'; \mathcal{TS}'$  for some transition label  $x$  and thread identifier  $\pi$ .

► **Lemma 29.** Suppose that  $G_0 \xrightarrow{\varepsilon_1, \pi_1} \dots \xrightarrow{\varepsilon_n, \pi_n} G$  for some trace  $tr = \langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$ . Then  $G$  is a strongly consistent execution graph that follows  $tr$ .

**Proof.** Easily follows from our definitions (note that strong consistency is trivially preserved in fork steps). ◀

► **Lemma 30.** Let  $G$  be a strongly consistent execution graph of an expression  $e$ . Suppose that  $G$  is not strongly racy or strongly uninitialized. Then, there exists a trace  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$  of  $e$ , such that  $G_0 \xrightarrow{\varepsilon_1, \pi_1} \dots \xrightarrow{\varepsilon_n, \pi_n} G$ .

**Proof.** By [Lemma 28](#),  $G.\text{sb} \cup G.\text{rf} \cup G.\text{mo}^{\text{na}}$  is acyclic. Let  $a_1, \dots, a_n$  be an enumeration of  $G'.\text{E}$  that respects  $(G'.\text{sb} \cup G'.\text{rf} \cup G'.\text{mo}^{\text{na}})^+$ . For every  $1 \leq i \leq n$ , let  $\pi_i = \text{tid}(a_i)$  and  $\varepsilon_i = \text{lab}(a_i)$ . Then,  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$  is a trace of  $e$ , and it is easy to see that by following this enumeration, i.e., adding  $a_i$  in step  $i$  for every  $1 \leq i \leq n$ , we obtain that  $G_0 \xrightarrow{\varepsilon_1, \pi_1} \dots \xrightarrow{\varepsilon_n, \pi_n} G'$ . In particular, note that the fact that  $G$  is strongly consistent entails that any  $(G.\text{sb} \cup G.\text{rf})^+$ -prefix of  $G$  is also strongly consistent. ◀

► **Theorem 31.** Let  $e$  be a non-buggy expression. Then, a trace  $tr$  is allowed for  $e$  iff  $tr$  has a permutation  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$  that respects  $\text{sb}(tr)$ , and

$$G_0; [0 \mapsto e] \xrightarrow{\varepsilon}^* \xrightarrow{\varepsilon_1, \pi_1} \xrightarrow{\varepsilon}^* \dots \xrightarrow{\varepsilon}^* \xrightarrow{\varepsilon_n, \pi_n} \xrightarrow{\varepsilon}^* G; \mathcal{TS}$$

for some execution graph  $G$  and threadpool  $\mathcal{TS}$ .

**Proof.** Follows from our definitions [Cor. 27](#), [Lemma 29](#) and [Lemma 30](#):

( $\Rightarrow$ ) Suppose that a trace  $tr$  is allowed for  $e$ . By [Cor. 27](#),  $tr$  is a trace of  $e$  and there exists a strongly consistent execution graph  $G$  that follows  $tr$ . Furthermore, [Cor. 27](#) entails that  $G$  is not strongly racy or strongly uninitialized. By [Lemma 30](#), there exists a trace  $tr' = \langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$  of  $e$ , such that  $G_0 \xrightarrow{\varepsilon_1, \pi_1} \dots \xrightarrow{\varepsilon_n, \pi_n} G$ . By definition,  $tr'$  must be a permutation of  $tr$  that respects  $\text{sb}(tr)$ .

( $\Leftarrow$ ) Suppose that  $tr$  has a permutation  $tr' = \langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$  that respects  $\text{sb}(tr)$ , and  $G_0; [0 \mapsto e] \xrightarrow{\varepsilon}^* \xrightarrow{\varepsilon_1, \pi_1} \xrightarrow{\varepsilon}^* \dots \xrightarrow{\varepsilon}^* \xrightarrow{\varepsilon_n, \pi_n} \xrightarrow{\varepsilon}^* G; \mathcal{TS}$  for some execution graph  $G$  and threadpool  $\mathcal{TS}$ . By definition,  $tr'$  is a trace of  $e$ . By [Lemma 29](#),  $G$  is a strongly consistent execution graph that follows  $tr'$ . Since  $tr'$  is a permutation of  $tr$  that respects  $\text{sb}(tr)$ ,  $G$  follows  $tr$ , and by [Lemma 5](#),  $tr$  is a trace of  $e$ . Hence,  $tr$  is allowed for  $e$ . ◀

► **Theorem 32.** An expression  $e$  is buggy iff  $G_0; [0 \mapsto e] \rightarrow^* \perp_{\text{race}}$  or  $G_0; [0 \mapsto e] \rightarrow^* \perp_{\text{uninit}}$ .

**Proof.**

( $\Rightarrow$ ) Suppose that  $e$  is buggy. By [Cor. 27](#), some strongly consistent execution graph  $G$  of  $e$  is strongly racy or strongly uninitialized. Let  $a_1, \dots, a_n$  be an enumeration of  $G.\text{E}$  that respects  $\text{sb} \cup \text{rf}$ . Let  $k$  be the minimal index such that  $G \cap \{a_1, \dots, a_k\}$  is strongly racy or strongly uninitialized. Consider the two cases:

- $G \cap \{a_1, \dots, a_k\}$  is strongly racy. Let  $j < k$  be the minimal index such that  $\text{loc}(a_k) = \text{loc}(a_j)$  and one of the following hold:



1.  $\langle a_k, a_j \rangle \notin \text{mo}^?; (\text{sb} \cup \text{rf})^+$  and  $\langle a_j, a_k \rangle \notin (\text{sb} \cup \text{rf})^+$ , and either  $(a_k \in \text{W}^{\text{at}}$  and  $a_j \in \text{R}^{\text{na}}$ ) or  $a_k \in \text{W}^{\text{na}}$ .
2.  $\langle a_j, a_k \rangle \notin \text{mo}^?; (\text{sb} \cup \text{rf})^+$  and  $\langle a_k, a_j \rangle \notin (\text{sb} \cup \text{rf})^+$ , and either  $(a_j \in \text{W}^{\text{at}}$  and  $a_k \in \text{R}^{\text{na}}$ ) or  $a_j \in \text{W}^{\text{na}}$ .

Distinguish the following cases:

- $a_k \in \text{W}^{\text{na}}$  and  $\langle a_k, a_j \rangle \notin \text{mo}^?; (\text{sb} \cup \text{rf})^+$ :  
Let  $B = \{a \in \text{E} \mid \langle a, a_j \rangle \in (G.\text{sb} \cup G.\text{rf})^+ \vee \langle a, a_k \rangle \in (G.\text{sb} \cup G.\text{rf})^*\}$ , and let  $G' = G \cap B$ . Then,  $G'$  is strongly consistent, not strongly racy, and not strongly uninitialized, and by [Lemma 30](#),  $G_0 \xrightarrow{\varepsilon_1, \pi_1} \dots \xrightarrow{\varepsilon_n, \pi_n} G'$  for some trace  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$  of  $e$ . Since  $\langle a_k, a_j \rangle \notin G.\text{mo}^?; (G.\text{sb} \cup G.\text{rf})^+$ , we have that  $\langle a_k, b \rangle \notin G'.\text{mo}^?; (G'.\text{sb} \cup G'.\text{rf})^+$  for every  $b \in G'.\text{E}^{\text{tid}(a_j)}$ . Since  $a_k \in G'.\text{W}_{\text{loc}(a_k)}^{\text{na}}$ , we have  $G' \xrightarrow{\text{lab}(a_j), \text{tid}(a_j)} \perp_{\text{race}}$ . Furthermore, it is easy to see that  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle, \langle \text{lab}(a_j), \text{tid}(a_j) \rangle$  is also a trace of  $e$ , and so  $G_0; [0 \mapsto e] \rightarrow^* \perp_{\text{race}}$ .
- $a_j \in \text{W}^{\text{na}}$  and  $\langle a_j, a_k \rangle \notin \text{mo}^?; (\text{sb} \cup \text{rf})^+$ :  
This case is handled symmetrically.
- $a_k \in \text{W}^{\text{at}}$ ,  $a_j \in \text{R}^{\text{na}}$ , and  $\langle a_k, a_j \rangle \notin \text{mo}^?; (\text{sb} \cup \text{rf})^+$ :  
Let  $B = \{a \in \text{E} \mid \langle a, a_j \rangle \in (G.\text{sb} \cup G.\text{rf})^+ \vee \langle a, a_k \rangle \in (G.\text{sb} \cup G.\text{rf})^*\}$ , and let  $G' = G \cap B$ . Then,  $G'$  is strongly consistent, not strongly racy, and not strongly uninitialized, and by [Lemma 30](#),  $G_0 \xrightarrow{\varepsilon_1, \pi_1} \dots \xrightarrow{\varepsilon_n, \pi_n} G'$  for some trace  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$  of  $e$ . Since  $\langle a_k, a_j \rangle \notin G.\text{mo}^?; (G.\text{sb} \cup G.\text{rf})^+$ , we have that  $\langle a_k, b \rangle \notin G'.\text{mo}^?; (G'.\text{sb} \cup G'.\text{rf})^+$  for every  $b \in G'.\text{E}^{\text{tid}(a_j)}$ . Since  $a_k \in G'.\text{W}_{\text{loc}(a_k)}$ , we have  $G' \xrightarrow{\text{lab}(a_j), \text{tid}(a_j)} \perp_{\text{race}}$ . Furthermore, it is easy to see that  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle, \langle \text{lab}(a_j), \text{tid}(a_j) \rangle$  is also a trace of  $e$ , and so  $G_0; [0 \mapsto e] \rightarrow^* \perp_{\text{race}}$ .
- $a_j \in \text{W}^{\text{at}}$ ,  $a_k \in \text{R}^{\text{na}}$ , and  $\langle a_j, a_k \rangle \notin \text{mo}^?; (\text{sb} \cup \text{rf})^+$ :  
This case is handled symmetrically.
- $G \cap \{a_1, \dots, a_k\}$  is strongly uninitialized: Then,  $a_k \in \text{RU}$  and  $\langle a, a_k \rangle \notin (\text{sb} \cup \text{rf})^*; \text{sb}$  for every  $a \in \text{WU}_{\text{loc}(a_k)}$ . Let  $B = \{a \in \text{E} \mid \langle a, a_k \rangle \in (G.\text{sb} \cup G.\text{rf})^+\}$ , and let  $G' = G \cap B$ . Then,  $G'$  is a  $(\text{sb} \cup \text{rf})^+$ -prefix of  $G$ , it is strongly consistent, not strongly racy, and not strongly uninitialized, and by [Lemma 30](#),  $G_0 \xrightarrow{\varepsilon_1, \pi_1} \dots \xrightarrow{\varepsilon_n, \pi_n} G'$  for some trace  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$  of  $e$ . Since  $\langle a, a_k \rangle \notin (G.\text{sb} \cup G.\text{rf})^*; G.\text{sb}$  for every  $a \in G.\text{WU}_{\text{loc}(a_k)}$ , we have  $G' \xrightarrow{\text{lab}(a_k), \text{tid}(a_k)} \perp_{\text{uninit}}$ . Furthermore, it is easy to see that  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle, \langle \text{lab}(a_k), \text{tid}(a_k) \rangle$  is also a trace of  $e$ , and so  $G_0; [0 \mapsto e] \rightarrow^* \perp_{\text{uninit}}$ .

( $\Leftarrow$ ) Suppose that  $G_0; [0 \mapsto e] \rightarrow^* \perp_{\text{race}}$  or  $G_0; [0 \mapsto e] \rightarrow^* \perp_{\text{uninit}}$ . Then, there exists a trace  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$  of  $e$ , such that  $G_0 \xrightarrow{\varepsilon_1, \pi_1} \dots \xrightarrow{\varepsilon_n, \pi_n} \perp_{\text{race}}$  or  $G_0 \xrightarrow{\varepsilon_1, \pi_1} \dots \xrightarrow{\varepsilon_n, \pi_n} \perp_{\text{uninit}}$ . Let  $G_{n-1}$  such that  $G_0 \xrightarrow{\varepsilon_1, \pi_1} \dots \xrightarrow{\varepsilon_{n-1}, \pi_{n-1}} G_{n-1} \xrightarrow{\varepsilon_n, \pi_n} \perp_{\text{race}}$ , or  $G_0 \xrightarrow{\varepsilon_1, \pi_1} \dots \xrightarrow{\varepsilon_{n-1}, \pi_{n-1}} G_{n-1} \xrightarrow{\varepsilon_n, \pi_n} \perp_{\text{uninit}}$ . By [Lemma 29](#),  $G_{n-1}$  is a strongly consistent execution of  $e$ . Since  $G_{n-1} \xrightarrow{\varepsilon_n, \pi_n} \perp_{\text{race}}$  or  $G_{n-1} \xrightarrow{\varepsilon_n, \pi_n} \perp_{\text{uninit}}$ , at least one of the following hold:

- $\varepsilon_n = \langle \text{Read}_{\text{na}}, \ell, v \rangle$  and there exists  $a \in G_{n-1}.\text{WU}_{\ell}$ , such that  $\langle a, b \rangle \notin (G_{n-1}.\text{sb} \cup G_{n-1}.\text{rf})^+$  for every  $b \in G_{n-1}.\text{E}^{\pi_n}$ . Let  $G' \in \text{Add}(G_{n-1}, \pi_n, \pi_n, \varepsilon_n)$ , and let  $b$  denote the additional event in  $G'$ . Clearly,  $G'$  is an execution of  $e$ . Using [Lemma 22](#), as well as the receptiveness property of traces, we can obtain an execution  $G''$  of  $e$  that is consistent, and satisfies  $\langle a, b \rangle \notin G''.\text{hb}$ . Since  $\text{mod}(b) = \text{na}$ , it follows that  $G''$  is racy.
- $\varepsilon_n \in \{\langle \text{Read}_{\alpha}, \ell, v \rangle, \langle \text{Write}_{\alpha}, \ell, v \rangle, \langle \text{Update}, \ell, v_o, v_n \rangle\}$ , and there exists  $a \in G_{n-1}.\text{W}_{\ell}^{\text{na}}$ , such that  $\langle a, b \rangle \notin (G_{n-1}.\text{sb} \cup G_{n-1}.\text{rf})^+$  for every  $b \in G_{n-1}.\text{E}^{\pi_n}$ . Let  $G' \in \text{Add}(G_{n-1}, \pi_n, \pi_n, \varepsilon_n)$ , and let  $b$  denote the additional event in  $G'$ . Clearly,  $G'$  is an execution of  $e$ . Using [Lemma 21](#), [Lemma 22](#) and [Lemma 23](#), as well as the receptiveness property of traces, and

**Lemma 24** in case  $\varepsilon_n$  is an update, we can obtain an execution  $G''$  of  $e$  that is consistent, and satisfies  $\langle a, b \rangle \notin G''.\text{hb}$ . Since  $\text{mod}(a) = \text{na}$ , it follows that  $G''$  is racy.

- $\varepsilon_n \in \{\langle \text{Read}_\alpha, \ell, v \rangle, \langle \text{Update}, \ell, v_o, v_n \rangle\}$ , and there do not exist  $a \in G_{n-1}.\text{WU}_\ell$  and  $b \in G_{n-1}.\text{E}^\pi$  such that  $\langle a, b \rangle \in (G_{n-1}.\text{sb} \cup G_{n-1}.\text{rf})^*$ . Let  $G' \in \text{Add}(G_{n-1}, \pi_n, \pi_n, \varepsilon_n)$  with  $G'.\text{rf} = G.\text{rf}$ . Clearly,  $G'$  is an execution of  $e$ , and it is easy to see that it is consistent and uninitialized. ◀

### A.3 Timestamp machine

Next, we show that the timestamp machine presented in §2 is equivalent to the operational machine above.

► **Definition 33.** A *timestamp assignment* for an execution graph  $G$  is a function  $ts : \text{WU} \rightarrow \text{Time}$ , that satisfies  $ts(a) < ts(b)$  whenever  $\langle a, b \rangle \in \text{mo}$ . Given a timestamp assignment  $ts$  for  $G$  and an event  $a \in \text{W} \cup \text{U}$ , the view of an event  $a$  in  $G$  according to  $ts$ , denoted  $\text{view}(a, G, ts)$ , and the message induced by  $a$  in  $G$  according to  $ts$ , denoted  $\text{msg}(a, G, ts)$ , are given by:

$$\begin{aligned} \text{view}(a, G, ts) &= \lambda \ell. \max\{ts(b) \mid b \in \text{WU}_\ell, \langle b, a \rangle \in (\text{sb} \cup \text{rf})^*\} \\ \text{msg}(a, G, ts) &= (\text{loc}(a), \text{val}_w(a), ts(a), \text{view}(a, G, ts)) \end{aligned}$$

If  $\{ts(b) \mid b \in \text{WU}_\ell, \langle b, a \rangle \in (\text{sb} \cup \text{rf})^*\}$  is empty, we take  $\text{view}(a, G, ts)(\ell) = \perp$  (undefined).

► **Definition 34.** Let  $G$  be an execution graph and  $ts$  be a timestamp assignment  $ts$  for  $G$ . The physical state  $(M_G^{ts}, T_G^{ts}, N_G^{ts})$  is given by:

- $M_G^{ts} = \{\text{msg}(a, G, ts) \mid a \in \text{WU}\}$
- $T_G^{ts} = \lambda \pi. \lambda \ell. \max\{ts(a) \mid a \in \text{WU}_\ell, \exists b \in \text{E}^\pi. \langle a, b \rangle \in (\text{sb} \cup \text{rf})^*\}$  (If  $\{ts(a) \mid a \in \text{WU}_\ell, \exists b \in \text{E}^\pi. \langle a, b \rangle \in (\text{sb} \cup \text{rf})^*\}$  is empty, we take  $T_G^{ts}(\pi)(\ell) = \perp$ )
- $N_G^{ts} = \lambda \ell. \max\{ts(a) \mid a \in \text{W}_\ell^{\text{na}}\}$  (If  $\{ts(a) \mid a \in \text{W}_\ell^{\text{na}}\}$  is empty, we take  $N_G^{ts}(\ell) = \perp$ )

We say that  $G$  *relates* to a physical state  $(M, T, N)$ , denoted  $G \sim (M, T, N)$  if  $(M, T, N) = (M_G^{ts}, T_G^{ts}, N_G^{ts})$  for some timestamp assignment  $ts$  for  $G$ .

The following lemmas (tediously but) straightforwardly follow from the definitions.

► **Lemma 35.** Suppose that  $(M, T, N) \xrightarrow{\varepsilon}^\pi (M', T', N')$ , and let  $G$  be an execution graph such that  $G \sim (M, T, N)$ . Then,  $G \xrightarrow{\varepsilon}^\pi G'$  for some execution graph  $G'$  such that  $G' \sim (M', T', N')$ .

► **Lemma 36.** Suppose that  $(M, T, N) \xrightarrow{\varepsilon}^\pi \perp_{\text{race}}$ , and let  $G$  be an execution graph such that  $G \sim (M, T, N)$ . Then,  $G \xrightarrow{\varepsilon}^\pi \perp_{\text{race}}$ .

► **Lemma 37.** Suppose that  $(M, T, N) \xrightarrow{\varepsilon}^\pi \perp_{\text{uninit}}$ , and let  $G$  be an execution graph such that  $G \sim (M, T, N)$ . Then,  $G \xrightarrow{\varepsilon}^\pi \perp_{\text{uninit}}$ .

► **Lemma 38.** Suppose that  $G \xrightarrow{\varepsilon}^\pi G'$ , and let  $ts'$  be a timestamp assignment for  $G'$ . Let  $ts = ts'|_{G.\text{WU}}$ . Then,  $ts$  is a timestamp assignment for  $G$ , and  $(M_G^{ts}, T_G^{ts}, N_G^{ts}) \xrightarrow{\varepsilon}^\pi (M_{G'}^{ts'}, T_{G'}^{ts'}, N_{G'}^{ts'})$ .

► **Lemma 39.** Suppose that  $G_0 \xrightarrow{\varepsilon_1}^{\pi_1} \dots \xrightarrow{\varepsilon_n}^{\pi_n} G_n$ . Then,  $(M_0, T_0, N_0) \xrightarrow{\varepsilon_1}^{\pi_1} \dots \xrightarrow{\varepsilon_n}^{\pi_n} (M_n, T_n, N_n)$  for some physical state  $(M_n, T_n, N_n)$ .

**Proof.** Let  $ts_n$  be some timestamp assignment for  $G_n$ , and for every  $0 \leq k < n$ , let  $ts_k = ts_n|_{G_k.\text{WU}}$ . For every  $0 \leq k \leq n$ , let  $(M_k, T_k, N_k) = (M_{G_k}^{ts_k}, T_{G_k}^{ts_k}, N_{G_k}^{ts_k})$ . (Clearly, this definition of  $(M_k, T_k, N_k)$  agrees with the definition of the initial physical state in **Definition 2**). By **Lemma 38**, we have  $(M_{k-1}, T_{k-1}, N_{k-1}) \xrightarrow{\varepsilon_k}^{\pi_k} (M_k, T_k, N_k)$  for every  $1 \leq i \leq n$ . ◀

► **Lemma 40.** Suppose that  $G \xrightarrow{\varepsilon}^{\pi} \perp_{\text{race}}$ , and let  $(M, T, N)$  be a physical state such that  $G \sim (M, T, N)$ . Then,  $(M, T, N) \xrightarrow{\varepsilon}^{\pi} \perp_{\text{race}}$ .

**Proof.** Let  $ts$  be a timestamp assignment such that  $(M, T, N) = (M_G^{ts}, T_G^{ts}, N_G^{ts})$ . Consider the two possible cases:

- $\varepsilon \in \{\langle \text{Read}_{\alpha}, \ell, v \rangle, \langle \text{Write}_{\alpha}, \ell, v \rangle, \langle \text{Update}, \ell, v_o, v_n \rangle\}$ , and  $\exists a \in W_{\ell}^{\text{na}}$ .  $\forall b \in E^{\pi}$ .  $\langle a, b \rangle \notin \text{mo}^?; (\text{sb} \cup \text{rf})^*$ : We show that  $T(\pi)(\ell) < N(\ell)$  (and it would follow that  $(M, T, N) \xrightarrow{\varepsilon}^{\pi} \perp_{\text{race}}$ ). Let  $a \in W_{\ell}^{\text{na}}$  be the **mo**-maximal event such that  $\forall b \in E^{\pi}$ .  $\langle a, b \rangle \notin \text{mo}^?; (\text{sb} \cup \text{rf})^*$ . Then,  $N(\ell) \geq ts(a)$ . Hence, it suffices to show that  $T(\pi)(\ell) < ts(a)$ . Let  $c \in WU_{\ell}$  and  $b \in E^{\pi}$  such that  $\langle c, b \rangle \in (\text{sb} \cup \text{rf})^*$ . We show that  $ts(c) < ts(a)$ . Indeed, otherwise, we have  $ts(a) \leq ts(c)$ , and so  $\langle a, c \rangle \in \text{mo}^?$ . Hence,  $\langle a, b \rangle \in \text{mo}^?; (\text{sb} \cup \text{rf})^*$ , which contradicts our assumption.
- $\varepsilon = \langle \text{Read}_{\text{na}}, \ell, v \rangle$  and  $\exists a \in W_{\ell}$ .  $\forall b \in E^{\pi}$ .  $\langle a, b \rangle \notin \text{mo}^?; (\text{sb} \cup \text{rf})^*$ : We show that  $\exists v', t', V'. (\ell, v', t', V') \in M \wedge T(\pi)(\ell) < t'$  (and it would follow that  $(M, T, N) \xrightarrow{\varepsilon}^{\pi} \perp_{\text{race}}$ ). Let  $a \in W_{\ell}$  such that  $\forall b \in E^{\pi}$ .  $\langle a, b \rangle \notin \text{mo}^?; (\text{sb} \cup \text{rf})^*$ .  $\text{msg}(a, G, ts) = (\ell, v', t', V')$ . Then,  $(\ell, v', t', V') \in M$ , and it suffices to show that  $T(\pi)(\ell) < t'$ . Let  $c \in WU_{\ell}$  and  $b \in E^{\pi}$  such that  $\langle c, b \rangle \in (\text{sb} \cup \text{rf})^*$ . We show that  $ts(c) < t' = ts(a)$ . Indeed, otherwise, we have  $t(c) \geq t(a)$  and so  $\langle a, c \rangle \in \text{mo}^?$ . Hence,  $\langle a, b \rangle \in \text{mo}^?; (\text{sb} \cup \text{rf})^*$ , which contradicts our assumption. ◀

► **Lemma 41.** Suppose that  $G \xrightarrow{\varepsilon}^{\pi} \perp_{\text{unit}}$ , and let  $(M, T, N)$  be a physical state such that  $G \sim (M, T, N)$ . Then,  $(M, T, N) \xrightarrow{\varepsilon}^{\pi} \perp_{\text{unit}}$ .

**Proof.** Let  $ts$  be a timestamp assignment such that  $(M, T, N) = (M_G^{ts}, T_G^{ts}, N_G^{ts})$ . Since  $G \xrightarrow{\varepsilon}^{\pi} \perp_{\text{unit}}$ , we have  $\varepsilon \in \{\langle \text{Read}_{\alpha}, \ell, v \rangle, \langle \text{Update}, \ell, v_o, v_n \rangle\}$ , and  $\neg \exists a \in WU_{\ell}, b \in E^{\pi}$ .  $\langle a, b \rangle \in (\text{sb} \cup \text{rf})^*$ . By definition, we have  $T_G^{ts}(\pi)(\ell) = \perp$ , and so  $(M, T, N) \xrightarrow{\varepsilon}^{\pi} \perp_{\text{unit}}$ . ◀

Finally, by combining the results above, we obtain the following:

► **Corollary 42.** Let  $e$  be an expression.

- If  $e$  is not buggy, then a trace  $tr$  is allowed for  $e$  iff  $tr$  has a permutation  $\langle \varepsilon_1, \pi_1 \rangle, \dots, \langle \varepsilon_n, \pi_n \rangle$  that respects  $\text{sb}(tr)$ , and

$$(M_0, T_0, N_0); [0 \mapsto e] \xrightarrow{\varepsilon}^* \xrightarrow{\varepsilon_1 \mapsto \pi_1} \xrightarrow{\varepsilon}^* \dots \xrightarrow{\varepsilon}^* \xrightarrow{\varepsilon_n \mapsto \pi_n} \xrightarrow{\varepsilon}^* (M, T, N); \mathcal{TS}$$

for some thread  $\pi$ , physical state  $(M, T, N)$ , and threadpool  $\mathcal{TS}$ .

- $e$  is buggy iff  $(M_0, T_0, N_0); [0 \mapsto e] \rightarrow^* \perp_{\text{race}}$  or  $(M_0, T_0, N_0); [0 \mapsto e] \rightarrow^* \perp_{\text{unit}}$  for some thread  $\pi$ .

**B** Single-writer protocols' proof rules

$$\begin{array}{c}
\boxed{\ell : s \mid \tau}_R \Rightarrow \square \boxed{\ell : s \mid \tau}_R \qquad \boxed{\ell : s_1 \mid \tau}_W * \boxed{\ell : s_2 \mid \tau}_W \Rightarrow \perp \\
\boxed{\ell : s_1 \mid \tau}_R * \boxed{\ell : s_2 \mid \tau}_R \Rightarrow s_1 \sqsubseteq s_2 \vee s_2 \sqsubseteq s_1 \qquad \boxed{\ell : s_1 \mid \tau}_W * \boxed{\ell : s_2 \mid \tau}_R \Rightarrow s_2 \sqsubseteq s_1 \\
\boxed{\ell : s \mid \tau}_W \Rightarrow \boxed{\ell : s \mid \tau}_W * \boxed{\ell : s \mid \tau}_R \\
\frac{\boxed{\ell : s \mid \tau}_W \Rightarrow \triangleright \tau_{\text{full}}(s, v) * \triangleright \tau_{\text{read}}(s, v) * Q}{\{\ell \hookrightarrow \_ \} \ell_{[\text{na}]} := v \{Q\}} \\
\frac{\left\{ \boxed{\ell : s \mid \tau}_W \right\} \ell_{[\text{at}]} \left\{ v. \boxed{\ell : s \mid \tau}_W * \tau_{\text{read}}(s, v) \right\}}{\forall s' \sqsupseteq s, v. P * \tau_{\text{read}}(s', v) \Rightarrow Q} \\
\frac{\left\{ \boxed{\ell : s \mid \tau}_R * P \right\} \ell_{[\text{at}]} \left\{ v. \exists s' \sqsupseteq s. \boxed{\ell : s' \mid \tau}_R * Q \right\}}{P * \boxed{\ell : s'' \mid \tau}_W * \tau_{\text{full}}(s, \_) \Rightarrow \tau_{\text{full}}(s'', v) * Q \quad s'' \sqsupseteq s} \\
\frac{\left\{ \boxed{\ell : s \mid \tau}_W * P \right\} \ell_{[\text{at}]} := v \left\{ \boxed{\ell : s'' \mid \tau}_R * Q \right\}}{}
\end{array}$$

**C** Escrows and exchanges' proof rules

In addition to escrows, we have mechanized a weaker variant of *exchanges* [30]. Our variant is weaker in the sense that the resources put in the exchanges have to be justified at the same view. Unlike the general exchanges, our variant can be trivially proven sound. In order to enforce that the resources are put back at the same view as when they are taken out, we need to expose the *identity* of views in the logic. However, since we only need to know equalities of views, we just need to bind views to names and expose the names. We call those names *abstract views*. The assertion  $\underline{P}_\beta$  asserts that  $P$  holds at some abstract view  $\beta$ . To use  $P$  a thread needs the extra knowledge  $\text{Seen}(\beta)$  which states that it has observed the view with name  $\beta$ .

The assertion  $[P \rightsquigarrow Q]_\beta^\mathcal{N}$  asserts the knowledge of an exchange at some abstract view  $\beta$  under  $\mathcal{N}$ . Escrows, unsurprisingly, are just a special instance of exchanges.

$$\begin{array}{c}
\text{EXCHANGE-INTRO-LEFT} \qquad \text{EXCHANGE-INTRO-RIGHT} \\
\triangleright P \Rightarrow_{\mathcal{N}} \exists \beta. [P \rightsquigarrow Q]_\beta^\mathcal{N} \qquad \triangleright Q \Rightarrow_{\mathcal{N}} \exists \beta. [P \rightsquigarrow Q]_\beta^\mathcal{N} \\
\text{EXCHANGE-ELIM-LEFT} \qquad \text{EXCHANGE-ELIM-RIGHT} \\
\frac{P * P \Rightarrow \perp}{[P \rightsquigarrow Q]_\beta^\mathcal{N} * \triangleright \underline{P}_\beta \Rightarrow_{\mathcal{N}} \triangleright \underline{Q}_\beta} \qquad \frac{Q * Q \Rightarrow \perp}{[P \rightsquigarrow Q]_\beta^\mathcal{N} * \triangleright \underline{Q}_\beta \Rightarrow_{\mathcal{N}} \triangleright \underline{P}_\beta} \\
\text{ESCROW-INIT} \qquad \text{ESCROW-APPLY} \\
\triangleright Q \Rightarrow_{\mathcal{N}} [P \rightsquigarrow Q]^\mathcal{N} \qquad \frac{P * P \Rightarrow \perp}{[P \rightsquigarrow Q]^\mathcal{N} * \triangleright P \Rightarrow_{\mathcal{N}} \triangleright Q}
\end{array}$$

**D** RSL proof rules

RSL focuses on the message-passing style transferring of resources through release-write/acquire-read pairs. The two main assertions of RSL are  $\text{Rel}(\ell, Q)$ , representing the permission to write to  $\ell$ , and  $\text{Acq}(\ell, Q)$ , representing the permission to read from  $\ell$ . The resource  $Q(v)$  is transferred through the write/read pair of value  $v$ . The assertion  $\text{RMWAcq}(\ell, P)$  is another permission that is used to CAS on the location.

We present in [Figure 15](#) the proof rules of iRSL, which were already checked by Coq. The encoding extends the original assertions with a predicate  $D$ , which is duplicable, to represent the restricted interpretation of the underlying invariant governing  $\ell$  (in a similar fashion to what we did in iGPS). For more details please refer to the Coq development, which also contains a fractional variant.

$$\begin{array}{c}
 \text{RSL-WREL} \\
 \{\triangleright Q(v) * \triangleright D(v) * \text{Rel}(\ell, Q, D)\} \ell_{[\text{at}]} := v \{\text{Init}(\ell) * \text{Rel}(\ell, Q, D)\} \\
 \\
 \text{RSL-RACQ} \\
 \{\text{Init}(\ell) * \text{Acq}(\ell, R, D)\} \ell_{[\text{at}]} \{v. \triangleright R(v) * \triangleright D(v) * \text{Acq}(\ell, R[v := \top], D)\} \\
 \\
 \text{RSL-CAS} \\
 \frac{P * \triangleright Q(v_o) * \triangleright D(v_o) \Rightarrow_{-\mathcal{N}. \ell} \triangleright Q(v_n) * \triangleright D(v_n) * R(1, v_o) \quad \forall v \neq v_o. P * D(v) \Rightarrow R(0, v)}{\{\text{Init}(\ell) * \text{RMWAcq}(\ell, Q, D) * P\} \text{cas}(\ell, v_o, v_n) \{b. \exists v. R(b, v) * \text{Init}(\ell) * \text{RMWAcq}(\ell, Q, D)\}} \\
 \\
 \text{RSL-REL-SPLITJOIN} \\
 \text{Rel}(\ell, Q_1, D) * \text{Rel}(\ell, Q_2, D) \dashv\vdash \text{Rel}(\ell, \lambda v. Q_1(v) \vee Q_2(v), D) \\
 \\
 \text{RSL-ACQ-SPLITJOIN} \\
 \text{Acq}(\ell, R_1, D) * \text{Acq}(\ell, R_2, D) \Leftrightarrow \text{Acq}(\ell, \lambda v. R_1(v) * R_2(v), D) \\
 \\
 \text{RSL-RMW-ACQ-SPLIT} \\
 \text{RMWAcq}(\ell, P, D) \Rightarrow \text{RMWAcq}(\ell, P, D) * \text{Acq}(\ell, \text{True}, D) \\
 \\
 \text{RSL-INIT-PERSISTENT} \qquad \text{RSL-RMW-PERSISTENT} \\
 \text{Init}(\ell) \vdash \Box \text{Init}(\ell) \qquad \text{RMWAcq}(\ell, P, D) \vdash \Box \text{RMWAcq}(\ell, P, D) \\
 \\
 \text{RSL-REL-PERSISTENT} \\
 \text{Rel}(\ell, Q, D) \vdash \Box \text{Rel}(\ell, Q, D) \\
 \\
 \text{RSL-ALLOC-R} \\
 \{\top\} \mathbf{alloc}(1) \{\ell. \text{Rel}(\ell, P) * \text{Acq}(\ell, P)\} \\
 \\
 \text{RSL-ALLOC-M} \\
 \{\top\} \mathbf{alloc}(1) \{\ell. \text{Rel}(\ell, P) * \text{RMWAcq}(\ell, P)\} \\
 \\
 \text{RSL-INIT-R} \\
 \{\text{Own}(\ell) * \triangleright Q(v) * \triangleright D(v)\} \ell_{[\text{na}]} := v \{\text{Init}(\ell) * \text{Rel}(\ell, Q, D) * \text{Acq}(\ell, Q, D)\} \\
 \\
 \text{RSL-INIT-M} \\
 \{\text{Own}(\ell) * \triangleright Q(v) * \triangleright D(v)\} \ell_{[\text{na}]} := v \{\text{Init}(\ell) * \text{Rel}(\ell, Q, D) * \text{RMWAcq}(\ell, Q, D)\}
 \end{array}$$

■ **Figure 15** Proof rules of iRSL.

Let  $ns \triangleq 0, tc \triangleq 1$ .

<pre> <b>newLock()</b> <math>\triangleq</math>   <b>let</b> <math>x = \text{alloc}(2)</math> <b>in</b>     <math>x + ns_{[na]} := 0;</math>     <math>x + tc_{[na]} := 0;</math>     <math>x</math> </pre>	<pre> <b>lock</b>(<math>x</math>) <math>\triangleq</math>   <b>let</b> <math>y = \text{FAI}(x + tc)</math> <b>in</b>     <b>repeat</b>       <b>let</b> <math>z = x + ns_{[at]}</math> <b>in</b>         <math>y == z</math> </pre>	<pre> <b>unlock</b>(<math>x</math>) <math>\triangleq</math>   <b>let</b> <math>z = x + ns_{[at]}</math> <b>in</b>     <math>x + ns_{[at]} := (z + 1) \bmod \mathbf{C}</math> </pre>
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

■ **Figure 16** Bounded ticket lock in  $\lambda_{RN}$ .

## E

 Bounded ticket locks

Ticket locks are a fair locking mechanism found in the Linux kernel. The first formal proof of ticket locks in RA+NA was given in the appendix of the GPS paper. We follow that variant and its proof's intuition closely, although our proof, thanks to single-writer protocols, is significantly simpler.

The code of the *bounded* ticket lock variant in  $\lambda_{RN}$  is given in [Figure 16](#). The lock has two atomic fields, a *ticket* counter ( $x + tc$ ) and a *now-serving* counter ( $x + ns$ ), which are initialized with 0. A thread can acquire the lock by first asking for a ticket  $t$  from the ticket counter, which is a fetch-and-increment (**FAI**) operation that hands out the current value of the ticket counter as the new ticket, and increases the counter by 1. After obtaining the ticket, the thread waits for the now-serving counter to reach its ticket (it keeps checking  $x + ns$ ), by which time it is guaranteed that previous tickets have been *served*, so it is safe to acquire the lock. To release the lock, the thread only needs to write  $t + 1$  to the now-serving counter, signifying that the next ticket can now be served.

In the bounded variant of ticket locks, the counters are stored in a fixed-size machine word with the maximum unsigned value  $\mathbf{C}$ . Once the counters reach  $\mathbf{C} - 1$ , the next increment will reset them back to 0. This additional flavor of reality induces the main complication of the verification. In unbounded ticket locks, it is easy to maintain the invariant that no two threads have the same ticket. Here, however, tickets are being *reused*, and the ticket counter must ensure that a ticket  $t$  has already been served before re-issuing it. In order to circumvent this, the number of concurrent threads trying to acquire the lock is restricted at most  $\mathbf{C}$ . Additionally, before getting a new ticket, each thread must prove to the ticket counter that its most recent ticket has been served. Consequently, the ticket counter must track which tickets have been issued to which threads and by its best the value of the now-serving counter, in order to keep the number of tickets waiting under  $\mathbf{C}$ .

The now-serving counter is where the actual transfer of resources takes place: the lock resource  $P$  is transferred from the release write to  $x + ns$  (by the **unlock** call of the previous ticket holder) to the acquire read by the **lock** call of the current ticket holder. Since a read cannot get ownership, escrows are used instead to transfer  $P$ . The releasing thread creates an escrow to store  $P$  and deliver it to the holder of ticket  $t + 1$ . The escrow is acquired by the  $t + 1$  ticket holder when the loop of its call to **lock** stops. By that time the acquiring thread should also have the correct exclusive token to use the escrow and obtain  $P$ . The releasing thread saves the proof that it has released the lock, in order to present that to the ticket counter when it requests another ticket later.

Our proof significantly simplifies the original proof by:

- **Using single-writer protocols for the now-serving counter.** The original proof uses a normal protocol for the counter, in combination with two other ghost components



$$\begin{aligned}
\text{Perm}^\gamma(t) &\triangleq \{\{\{t\}, \circ \emptyset\} : \text{DISJ}(\mathbb{N}) \times \text{AUTH}(\mathbb{N} \stackrel{\text{fin}}{\dashv} \text{EX}(\mathbb{N} \cup \{-1\}))\}^\gamma \\
\text{AllTks}^\gamma(M) &\triangleq \{\{\emptyset, \bullet M\}\}^\gamma \\
\text{MyTkt}^\gamma [i \mapsto t] &\triangleq \{\{\emptyset, \circ [i := t]\}\}^\gamma \\
\text{Escs}(x, \gamma, n) &\triangleq \forall t \leq n. \left[ \text{Perm}^\gamma(t) \rightsquigarrow P * \boxed{x + \text{ns} : t \mid \text{NSP}(\gamma)}_W \right] \\
\text{NSP}(\gamma)(n, z) &\triangleq z = n \bmod \mathbf{C} * \square \text{Escs}(x, \gamma, n) \\
\text{Waiting}(M, n) &\triangleq \{i \in \text{dom}(M) \mid M(i) \geq n\} \\
\text{TCP}(\gamma, x)(\text{Inv}, y) &\triangleq \exists t, n, M. y = t \bmod \mathbf{C} \wedge n \leq t \wedge t \leq n + |\text{Waiting}(M, n)| \\
&\quad \wedge t = \max_i M(i) \wedge \text{dom}(M) = [0, \dots, \mathbf{C} - 1] \\
&\quad \wedge (\forall i, j \in \text{dom}(M). M(i) = M(j) \neq -1 \Rightarrow i = j) \\
&\quad * \text{AllTks}^\gamma(M) * \bigstar_{j \geq t} \text{Perm}^\gamma(j) * \boxed{x + \text{ns} : n \mid \text{NSP}(\gamma)}_R \\
\text{MayAcquire}(x) &\triangleq \exists \gamma, i, t. \text{MyTkt}^\gamma [i \mapsto t] * \boxed{x + \text{tc} : \text{Inv} \mid \text{TCP}(\gamma, x)} \\
&\quad * t \neq -1 \Rightarrow \boxed{x + \text{ns} : t + 1 \mid \text{NSP}(\gamma)}_R \\
\text{MayRelease}(x) &\triangleq \exists \gamma, i, t. \text{MyTkt}^\gamma [i \mapsto t] * \boxed{x + \text{tc} : \text{Inv} \mid \text{TCP}(\gamma, x)} * \boxed{x + \text{ns} : t \mid \text{NSP}(\gamma)}_W
\end{aligned}$$

■ **Figure 17** Proof setup for bounded ticket locks.

to ensure that the counter's value cannot surpass the ticket  $t$  held by the acquiring thread, and that once the lock is acquired, the value stays exactly  $t$  until the thread releases the lock by increasing it to  $t + 1$ . In our proof, we add the single-writer permission of  $x + \text{ns}$  to the lock's resources, so a thread also acquires the full-write ownership of  $x + \text{ns}$  when it acquires the lock, and releases the ownership by the same time with releasing the lock. This removes two components of the ghost monoid needed by the original proof to maintain knowledge about some upper-bound of the protocol's state.

- **Tracking only the most recent ticket of a thread.** The original proof tracks the complete history of issued tickets, which we found unnecessary. Instead, we only need to track the most recent ticket issued to a thread, which simplifies the tracking component of our ghost monoid. However, we do not claim that this results from the use of single-writer protocols: it seems that the simplification can also be applied to the original proof.

## E.1 Proof setup for bounded ticket locks

We explain the technical construction of the proof. The complete definitions of monoids, protocols and escrows needed are given in [Figure 17](#).

We use protocols  $\text{TCP}(\gamma, x)$  and  $\text{NSP}(\gamma)$  to govern the ticket counter and the now-serving counter respectively. The  $\text{TCP}(\gamma, x)$  protocol only has a single state  $\text{Inv}$ , while the states of  $\text{NSP}(\gamma)$  are natural numbers, which represent *abstract* tickets. Though working differently, both counters record the abstract ticket  $t$  they are currently working with, who requires the concrete counter value  $t_c$  to follow the equation  $t_c = t \bmod \mathbf{C}$ . The abstract ticket of each counter can only increase over time (the state relation is  $\leq$  for  $\text{NSP}(\gamma)$ ). By exploiting abstract tickets, we do not need to deal with reusing concrete tickets.

The interpretations of  $\text{NSP}(\gamma)$  and  $\text{TCP}(\gamma, x)$  depend respectively on the first and second components of a product monoid for our ghost states. The first component  $\text{DISJ}(\mathbb{N})$

is a *disjoint-union* monoid of natural numbers, whose composition is only defined for disjoint sets of natural numbers:  $X \cdot Y = X \cup Y$  if  $X \cap Y = \emptyset$ . It is used for abstract ticket ownerships. The second component  $\text{AUTH}(\mathbb{N} \stackrel{\text{fin}}{\mapsto} \text{EX}(\mathbb{N} \cup \{-1\}))$  is an authoritative monoid of a finite partial map from  $\mathbb{N}$  to exclusive elements from  $\mathbb{N} \cup \{-1\}$ . It is used for tracking issued abstract tickets for all participating threads. The composition works in the same way as that of the heap monoid in  $\lambda_{\text{SC}}$  and the history and thread-view monoids in  $\lambda_{\text{RN}}$ : it has the same rules for exclusiveness, agreement, and updates as those monoids. The composition of the product monoid is then lifted piecewise.

### The now-serving counter protocol

The assertion  $\text{Perm}^\gamma(t)$  embeds the ownership of the abstract ticket  $t$  in its first component (while the second component is an empty map which contains zero knowledge or ownership).  $\text{Perm}^\gamma(t)$  satisfies exclusiveness, realizing the requirement that a ticket can only be uniquely owned:

$$\text{Perm}^\gamma(t) * \text{Perm}^\gamma(t) \Rightarrow \text{False}$$

Exclusiveness of  $\text{Perm}^\gamma(t)$  allows us to define the escrow used to transfer the lock's resource, which contains  $P$  and the full single-writer permission of the now-serving counter. This is reflected in the definition of  $\mathbf{Escs}(x, \gamma, n)$ , which asserts the existence of the escrows for all abstract tickets up to  $n$ . Additionally,  $\mathbf{Escs}(x, \gamma, n)$  subsumes the proofs that all abstract tickets up to (but excluding)  $n$  have been served, due to the fact that the lock's resource can only be stored in exactly one of the escrows at a time, and that the verification maintains that the resource only moves from the escrow of some ticket  $t$  to that of  $t + 1$ . These proofs are needed when a thread updates the now-serving counter  $x + \text{ns}$ , to release the lock. This is enforced in  $\mathbf{NSP}(\gamma)$ . When increasing the counter with value  $z = n \bmod \mathbf{C}$  to  $(n + 1) \bmod \mathbf{C}$ , the lock-releasing thread moves the protocol state from  $n$  to  $n + 1$ , and must re-establish  $\mathbf{NSP}(\gamma)$  by allocating the escrow for the ticket  $n + 1$  of the next lock-acquiring thread.

### The ticket counter protocol

Tracking tickets is the main role of the ticket counter. The tracking  $M$  is a finite map from participant ids to their corresponding most recent abstract tickets. The domain of  $M$ , i.e. the set of ids, is restricted to the set  $[0, \dots, \mathbf{C} - 1]$ , since there are at most  $\mathbf{C}$  threads participating in the lock. If a thread has never obtained any ticket, then it is mapped to  $-1$ . The authoritative part of the map  $\bullet M$  is part of  $\text{AllTks}^\gamma(M)$ , which in turn is stored in the protocol  $\mathbf{TCP}(\gamma, x)$  that governs the ticket counter  $x + \text{tc}$ . The non-authoritative  $\circ M$  is split into  $\text{MyTkt}^\gamma [i \mapsto M(i)]$ 's for  $i \in [0, \dots, \mathbf{C} - 1]$ . Each  $\text{MyTkt}^\gamma [i \mapsto M(i)]$  represents the knowledge of the most recent ticket  $M(i)$ , as well as the exclusive permission required to acquire a new ticket for the thread  $i$ . These permissions are allocated when the lock is created, and then distributed externally to the threads that wish to use the lock.

$\mathbf{TCP}(\gamma, x)$  always keep its authoritative tracker  $\text{AllTks}^\gamma(M)$ . The maximum abstract ticket of  $M$  is  $t - 1$ , which is also the most recently issued ticket by the ticket counter. Meanwhile, the next ticket to be issued is  $t$ , and the current value  $y$  of the ticket counter must be the corresponding concrete ticket of  $t$ .  $M$  of course is injective, since no two threads should get the same ticket.  $\mathbf{TCP}(\gamma, x)$  also holds permissions  $\ast_{j \geq t} \text{Perm}^\gamma(j)$  of to-be-issued abstract tickets. When requesting a ticket, a thread obtains the abstract ticket  $t$  (and the concrete ticket  $y$ ), as well as  $\text{Perm}^\gamma(t)$  to use the escrow later. The thread also increases the

$$\begin{aligned}
& \left\{ \triangleright P \right\} \mathbf{newLock}() \left\{ x. \bigstar_{i < \mathbf{C}} \text{MayAcquire}(x) \right\} \\
& \left\{ \text{MayAcquire}(x) \right\} \mathbf{lock}(x) \left\{ \triangleright P * \text{MayRelease}(x) \right\} \\
& \left\{ \triangleright P * \text{MayRelease}(x) \right\} \mathbf{unlock}(x) \left\{ \text{MayAcquire}(x) \right\}
\end{aligned}$$

■ **Figure 18** Specification of bounded ticket locks.

ticket counter value (concretely and abstractly), and, by doing so, it must show that it has observed the now-serving counter to be at least some state that is not too far away from the ticket counter, which is realized by the conditions between  $t$  and  $n$  in the definition of **TCP**. These conditions say that the ticket counter is always ahead of the now-serving counter, but their distance cannot surpass the number of waiting tickets known to the ticket counter, i.e. those that have been issued but are still not reported to the ticket counter as served.

### Specification for bounded ticket lock

The specification that we need to verify is very simple, as given in [Figure 18](#).

$\text{MayAcquire}(x)$  is the permission for a thread to request a ticket and later acquire the lock. The thread must have the right  $\text{MyTkt}^\gamma [i \mapsto t]$  to change its most recent ticket  $t$ , and it must show that  $t$  has been served, by knowing  $\boxed{x + \text{ns} : t + 1 \mid \mathbf{NSP}(\gamma)}_R$ .

By acquiring the lock, a thread earns  $\triangleright P$  and  $\text{MayRelease}(x)$ , which is the permission to release the lock.  $\text{MayRelease}(x)$  includes the writer permission of  $\boxed{x + \text{ns} : t \mid \mathbf{NSP}(\gamma)}_W$  for the ticket  $t$  that was used to acquire the lock. The thread has to give up this writer permission to update the now-serving counter to  $t + 1$ , at the same time gaining only the reader permission for  $t + 1$ , which is enough to re-establish  $\text{MayAcquire}(x)$  that can be used for its next turn. Note that  $P$  is guarded by a later, since it was put in an escrow. Fortunately there should be always enough steps in a client program to remove the later.

The proof outlines are given [Figure 19](#), [Figure 20](#), and [Figure 21](#). For **newLock**, we only need to set up the default initial states of the protocols, for exactly  $\mathbf{C}$  participants. The proof of **unlock** simply exploit the power of the single writer protocol to release the writer permission while doing the write itself (see [§4.2](#)). We explain the proof of **lock** in more detail.

## E.2 Proof steps for lock

The proof outline for the **lock** procedure is given in [Figure 21](#). We explain the main reasoning below.

- With the permission  $\text{MyTkt}^\gamma [i \mapsto t_i]$ , a thread  $i$  acquires a new abstract ticket  $t$  (concretely  $y$ ) along with the permission  $\text{Perm}^\gamma(t)$ , and updates the tracker for  $i$  to  $t$ . It also increases the ticket counter concretely to  $(y + 1) \bmod \mathbf{C}$  along with an abstract ticket  $t + 1$ , by which the ticket counter demands  $\boxed{x + \text{ns} : n_0 \mid \mathbf{NSP}(\gamma)}_R$  for some good enough  $n_0$ . This demand is fulfilled either by what the thread has locally, or what is already in the interpretation previously. In either way the thread updates its observation of  $\mathbf{NSP}(\gamma)$  to  $n_0$ .

$$\begin{array}{l}
\{\triangleright P\} \\
\mathbf{let } x = \mathbf{alloc}(2) \mathbf{ in} \\
\{\triangleright P * x + \mathbf{ns} \hookrightarrow A * x + \mathbf{tc} \hookrightarrow A\} \\
\{\triangleright P * x + \mathbf{ns} \hookrightarrow A * x + \mathbf{tc} \hookrightarrow A * \exists \gamma, M. \bigstar_{t \geq 0} \text{Perm}^\gamma(t) * \text{AllTcts}(M_{-1}) * \bigstar_{i < \mathbf{C}} \text{MyTkt}^\gamma [i \mapsto -1]\} \\
x + \mathbf{ns}_{[\mathbf{na}]} := 0; \\
\left\{ \boxed{x + \mathbf{ns} : 0 \mid \text{NSP}(\gamma)}_W * x + \mathbf{tc} \hookrightarrow A * \bigstar_{t \geq 0} \text{Perm}^\gamma(t) * \text{AllTcts}(M_{-1}) * \bigstar_{i < \mathbf{C}} \text{MyTkt}^\gamma [i \mapsto -1] \right\} \\
x + \mathbf{tc}_{[\mathbf{na}]} := 0; \\
\left\{ \boxed{x + \mathbf{ns} : 0 \mid \text{NSP}(\gamma)}_R * \boxed{x + \mathbf{tc} : \text{Inv} \mid \text{TCP}(\gamma, x)}_R * \bigstar_{i < \mathbf{C}} \text{MyTkt}^\gamma [i \mapsto -1] \right\} \\
x \\
\left\{ \bigstar_{i < \mathbf{C}} \text{MayAcquire}(x) \right\}
\end{array}$$

■ **Figure 19** Proof outline for **newLock**

$$\begin{array}{l}
\{\triangleright P * \text{MayRelease}(x)\} \\
\left\{ \exists \gamma, i, t. \triangleright P * \text{MyTkt}^\gamma [i \mapsto t] * \boxed{x + \mathbf{tc} : \text{Inv} \mid \text{TCP}(\gamma, x)} * \boxed{x + \mathbf{ns} : t \mid \text{NSP}(\gamma)}_W \right\} \\
\text{Ctx: } \boxed{x + \mathbf{tc} : \text{Inv} \mid \text{TCP}(\gamma, x)} \\
\mathbf{let } z = x + \mathbf{ns}_{[\mathbf{at}]} \mathbf{ in} \\
\left\{ z = t \bmod \mathbf{C} * \triangleright P * \text{MyTkt}^\gamma [i \mapsto t] * \boxed{x + \mathbf{ns} : t \mid \text{NSP}(\gamma)}_W \right\} \\
x + \mathbf{ns}_{[\mathbf{at}]} := (z + 1) \bmod \mathbf{C} \\
\left\{ \text{MyTkt}^\gamma [i \mapsto t] * \boxed{x + \mathbf{ns} : t + 1 \mid \text{NSP}(\gamma)}_R \right\} \\
\{\text{MayAcquire}(x)\}
\end{array}$$

■ **Figure 20** Proof outline for **unlock**


$$\begin{array}{l}
 \{\text{MayAcquire}(x)\} \\
 \left\{ \exists \gamma, i, t_i. \text{MyTkt}^\gamma [i \mapsto t_i] * \boxed{x + tc : \text{Inv} \mid \text{TCP}(\gamma, x)} * t_i \neq -1 \Rightarrow \boxed{x + ns : t_i + 1 \mid \text{NSP}(\gamma)}_R \right\} \\
 \text{let } y = \text{FAI}(x + tc) \text{ in} \\
 \left\{ \exists t, n_0. y = t \bmod \mathbf{C} \wedge t < n_0 + \mathbf{C} * \text{MyTkt}^\gamma [i \mapsto t] * \text{Perm}^\gamma(t) * \triangleright \boxed{x + ns : n_0 \mid \text{NSP}(\gamma)}_R \right\} \\
 \text{Ctx: } y = t \bmod \mathbf{C} \wedge t < n_0 + \mathbf{C} \\
 \text{Ctx: } \boxed{x + tc : \text{Inv} \mid \text{TCP}(\gamma, x)} \\
 \text{repeat} \\
 \left\{ \exists n \geq n_0. \text{MyTkt}^\gamma [i \mapsto t] * \text{Perm}^\gamma(t) * \boxed{x + ns : n \mid \text{NSP}(\gamma)}_R \right\} \\
 \text{let } z = x + ns_{[\text{at}]} \text{ in} \\
 \left\{ \begin{array}{l} \exists n' \geq n. z = n' \bmod \mathbf{C} * \square \text{Escs}(x, \gamma, n') \\ * \text{MyTkt}^\gamma [i \mapsto t] * \text{Perm}^\gamma(t) * \boxed{x + ns : n' \mid \text{NSP}(\gamma)}_R \end{array} \right\} \\
 y == z \\
 \left\{ \begin{array}{l} \exists n \geq n_0. y = z * z = n \bmod \mathbf{C} * \square \text{Escs}(x, \gamma, n) \\ * \text{MyTkt}^\gamma [i \mapsto t] * \text{Perm}^\gamma(t) * \boxed{x + ns : n \mid \text{NSP}(\gamma)}_R \end{array} \right\} \\
 \text{Have } n \leq t < n_0 + \mathbf{C} \leq n + \mathbf{C} \text{ and } t \bmod \mathbf{C} = n \bmod \mathbf{C} \\
 \text{Ctx: } t = n \\
 \left\{ \text{MyTkt}^\gamma [i \mapsto t] * \text{Perm}^\gamma(t) * \left[ \text{Perm}^\gamma(t) \rightsquigarrow P * \boxed{x + ns : t \mid \text{NSP}(\gamma)}_W \right] \right\} \\
 \left\{ \text{MyTkt}^\gamma [i \mapsto t] * \triangleright P * \triangleright \boxed{x + ns : t \mid \text{NSP}(\gamma)}_W * \boxed{x + ns : t \mid \text{NSP}(\gamma)}_R \right\} \\
 \left\{ \triangleright P * \text{MyTkt}^\gamma [i \mapsto t] * \boxed{x + ns : t \mid \text{NSP}(\gamma)}_W \right\} \\
 \left\{ \triangleright P * \text{MayRelease}(x) \right\}
 \end{array}$$

■ **Figure 21** Proof outline for **lock**

- By entering the loop, the thread keeps updating its observation of  $\text{NSP}(\gamma)$  to some state  $n$  later than  $n_0$ . When the loop exits, the thread knows that the two concrete tickets  $y$  and  $z$  are the same, and this fact should imply that their corresponding abstract tickets  $t$  and  $n$  are also the same. To show this we need to use a subtle reasoning with the single writer protocol.

Since we already have  $t \bmod \mathbf{C} = n \bmod \mathbf{C}$  and  $t < n + \mathbf{C}$  by acquiring the ticket and loop-checking the counter, we only need to show  $n \leq t$  to show  $n = t$ . This matches the intuition that the now-serving counter cannot surpass the unserved ticket, since we have locally the permission  $\text{Perm}^\gamma(t)$ . We assume the contrary  $t < n$  and derive a contradiction. Note that by reading the now-serving counter we also obtain  $\text{Escs}(x, \gamma, n)$ . Then we can use  $\text{Perm}^\gamma(t)$  with the escrow for  $t$  to obtain  $\triangleright \boxed{x + ns : t \mid \text{NSP}(\gamma)}_W$ . Combining with our locally-known  $\boxed{x + ns : n \mid \text{NSP}(\gamma)}_R$ , it must be that  $n \leq t$ : *contradiction!*

- Knowing that  $n = t$ , we use the escrow with  $\text{Perm}^\gamma(t)$  to acquire the lock, and gain  $\triangleright P$ . We also gain  $\triangleright \boxed{x + ns : t \mid \text{NSP}(\gamma)}_W$ , which we combine with  $\boxed{x + ns : t \mid \text{NSP}(\gamma)}_R$  to remove the later, and complete the  $\text{MayRelease}(x)$  permission.



(a) Visualization of the iGPS protocol invariant.

$$\begin{aligned}
 \text{PrInv}_\tau(\ell) &\triangleq \exists \Delta, \delta_x \in \Delta. \\
 &\text{Hist}(\ell, \{(v, t, V) \mid (s, v, t, V) \in \Delta\}) \\
 &* \{ \bullet \Delta \cdot \circ \Delta \} : \text{LH} * \{ \bullet \delta_x : \text{XW} \} \\
 &* \left( *_{(s, v, \_, V) \in \Delta} \tau_{\text{read}}(s, v)(V) \right) \\
 &* \left( *_{(s, v, t, V) \in \Delta} (\nexists \delta \in \Delta. \delta.t = t + 1) \Rightarrow \tau_{\text{full}}(s, v)(V) \right) \\
 &* \text{Consistent}(\Delta) * \text{Unique}(\Delta) * \text{FinalInv}(\Delta)
 \end{aligned}$$

(b) Definition of the iGPS protocol invariant.

■ **Figure 22** The iGPS protocol invariant.

## F iGPS Protocol Model

The model of iGPS protocols consists of two parts: a global protocol invariant  $\text{PrInv}_\tau(\ell)$  shared by all protocol variants, and local protocol assertions given out to clients, which depend on the particular protocol variant. We will now explain both parts in detail. The definitions are given in [Figure 22b](#).<sup>8</sup>

The protocol invariant owns, among other resources, the location’s history  $\text{Hist}(\ell, h)$  – though  $h$  will be expressed differently, as we explain below. In addition to the physical history, the invariant also owns the authoritative fragment of the *logical history*  $\Delta$  – an augmented history that write events paired with protocol states.

The protocol invariant further owns the protocol interpretation for every state in  $\Delta$ . As all states are potentially targets of reads, we store  $\tau_{\text{read}}$  for every state in the history. To enable ownership transfer via CAS, we additionally need to store  $\tau_{\text{full}}$  for those states whose associated write event can still be target of an update operation.

This condition is sufficient for non-exclusive protocols. However, single-writer protocols offer an additional mechanism of transferring out resources from the protocol interpretation: exclusive writes. To support both protocol variants, the invariant tracks the most recent exclusive write  $(s_x, v_x, V_x)$  and stores  $\tau_{\text{full}}$  for  $\delta_x$  as well as any later write that can be target of CAS. In the case of non-exclusive protocols,  $\delta_x$  is set to the initial protocol write and all CAS-able writes maintain their full interpretation  $\tau_{\text{full}}$ .

We track  $\delta_x$  with an additional authoritative piece of ghost state. The protocol invariant holds the authoritative fragment  $\bullet \Delta_x$  while the non-authoritative fragment  $\circ \Delta_x$  is given to the writer – or put into the shared invariant  $\text{PrShared}_\tau$  in the case of non-exclusive protocols.

Apart from those resources, the protocol invariant imposes three conditions on the physical history  $h$  and the logical history  $\Delta$ . The most important one is *consistency*: the modification order on  $h$  must correspond to the protocol order on  $\Delta$ . The **Consistent** assertion captures exactly this property.

<sup>8</sup> The definitions differ slightly from the Coq mechanization to simplify the presentation.

$$\begin{aligned}
 \text{Reader}(\delta) &\triangleq \boxed{\circ(\circ\{\delta\}) : \text{LH}} \\
 \text{Writer}(\delta) &\triangleq \exists \Delta. \boxed{\circ(\bullet \Delta \cdot \circ \Delta) : \text{LH}} * \delta \in \Delta * \forall \delta' \in \Delta. \delta'.t \leq \delta.t \\
 \text{PrShared}_\tau(t_x) &\triangleq \exists \delta, \delta_x. \text{Writer}(\delta) * \boxed{\circ \delta_x : \text{XW}} * \delta_x.t = t_x. \\
 \boxed{\ell : s \mid \tau} &\triangleq \lambda V. \boxed{\text{PrInv}_\tau(\ell)} * \exists V_R \sqsubseteq V, t_x \leq V(\ell). \text{Reader}((s, \_, V_R)) * \boxed{\text{PrShared}_\tau(t_x)} \\
 \boxed{\ell : s \mid \tau}_R &\triangleq \lambda V. \boxed{\text{PrInv}_\tau(\ell)} * \exists V_R \sqsubseteq V. \text{Reader}((s, \_, \_, V_R)) \\
 \boxed{\ell : s \mid \tau}_W &\triangleq \lambda V. \boxed{\text{PrInv}_\tau(\ell)} * \exists V_W \sqsubseteq V, v_w, t_w. \text{Writer}((s, v_w, t_w, V_W)) * \boxed{(s, v_w, t_w, V_W) : \text{XW}}
 \end{aligned}$$

■ **Figure 23** Definition of the iGPS protocol assertions for non-exclusive protocols and single-writer protocols.

Additionally, we need to make sure that the logical history does not assign multiple protocol states to a write event. To this end, **Unique** asserts that  $\Delta$  contains any write event at most once. With this restriction in place, we can now express  $h$  in terms of  $\Delta$  by forgetting the protocol states in  $\Delta$ , i.e.  $h = \{(v, t, V) \mid (s, v, t, V) \in \Delta\}$ .

Finally, we need an additional piece of information on the logical history:  $\Delta$  contains only final states, i.e. states reachable from any other state, in the *disconnected segment*. We call this property **FinalInv**. The disconnected segment – illustrated in [Figure 22a](#) – spans all writes which are not connected to the most recent exclusive write. The write events contained in that segment share an important property: they could be *mo*-after newly inserted writes (which are always *mo*-after the most recent exclusive write). To re-establish **Consistent** for new writes, we need to be able to relate their protocol state to any existing *mo*-later write. **FinalInv** guarantees that the state associated with the newly inserted write comes before any state in the disconnected segment.

Before we turn to the local protocol assertions, we define the PCM's used in **PrInv** $_\tau$ . The PCM for the most recent exclusive write  $\delta_x$  is a simple AUTH construction similar to the ones shown in [§3](#), providing both an authoritative fragment for **PrInv** $_\tau$  as well as an exclusive non-authoritative fragment to be given to the writer or to **PrShared** $_\tau$ . The PCM for logical histories is more involved, as it needs to provide three kinds of fragments: (1) an authoritative fragment held by **PrInv** $_\tau$ , (2) an exclusive fragment given to the writer or to **PrShared** $_\tau$ , representing the right to change the logical history – and thus also the physical one – and (3) a persistent fragment used in the (reader) protocol assertion, representing a lower bound on protocol state. To fulfill these requirements, we construct the rather intimidating double-AUTH PCM LH below.

$$\begin{aligned}
 \text{LE} &\triangleq \text{State} \times \text{Val} \times \text{Time} \times \text{View} && \text{(logical event)} \\
 \text{LES} &\triangleq \mathcal{P}_\cup(\text{LE}) && \text{(logical event set)} \\
 \text{XW} &\triangleq \text{AUTH}(\text{EX}(\text{LE})) && \text{(most recent excl. event)} \\
 \text{LH} &\triangleq \text{AUTH}(\text{AUTH}(\text{LES})) && \text{(auth. logical event set)}
 \end{aligned}$$

With these definitions we can now define protocol assertions, which are presented in [Figure 23](#). As before, we elide ghost and invariant names, which add some bookkeeping overhead to the definitions.