Rely-Guarantee Reasoning for Causally Consistent Shared Memory*

Ori Lahav[1], Brijesh Dongol[2], and Heike Wehrheim[3]

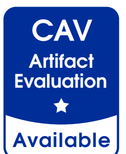

[1] Tel Aviv University, Tel Aviv, Israel
[2] University of Surrey, Guildford, UK
[3] University of Oldenburg, Oldenburg, Germany

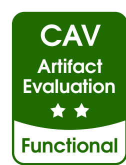

Abstract. Rely-guarantee (RG) is a highly influential compositional proof technique for concurrent programs, which was originally developed assuming a sequentially consistent shared memory. In this paper, we first generalize RG to make it parametric with respect to the underlying memory model by introducing an RG framework that is applicable to any model axiomatically characterized by Hoare triples. Second, we instantiate this framework for reasoning about concurrent programs under *causally consistent memory*, which is formulated using a recently proposed *potential-based* operational semantics, thereby providing the first reasoning technique for such semantics. The proposed program logic, which we call Piccolo, employs a novel assertion language allowing one to specify ordered sequences of states that each thread may reach. We employ Piccolo for multiple litmus tests, as well as for an adaptation of Peterson's algorithm for mutual exclusion to causally consistent memory.

1 Introduction

Rely-guarantee (RG) is a fundamental compositional proof technique for concurrent programs [21,48]. Each program component P is specified using *rely* and *guarantee* conditions, which means that P can tolerate any environment interference that follows its rely condition, and generate only interference included in its guarantee condition. Two components can be composed in parallel provided that the rely of each component agrees with the guarantee of the other.

The original RG framework and its soundness proof have assumed a sequentially consistent (SC) memory [33], which is unrealistic in modern processor architectures and programming languages. Nevertheless, the main principles behind RG are not at all specific for SC. Accordingly, our first main contribution, is to formally decouple the underlying memory model from the RG proof principles, by proposing a generic RG framework parametric in the input memory model.

* Lahav is supported by the Israel Science Foundation (grants 1566/18 and 814/22) and by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 851811). Dongol is supported by EPSRC grants EP/X015149/1, EP/V038915/1, EP/R025134/2, VeTSS, and ARC Discovery Grant DP190102142. Wehrheim is supported by the German Research Council DFG (project no. 467386514).

To do so, we assume that the underlying memory model is axiomatized by Hoare triples specifying pre- and postconditions on memory states for each primitive operation (e.g., loads and stores). This enables the formal development of RG-based logics for different shared memory models as instances of one framework, where all build on a uniform soundness infrastructure of the RG rules (e.g., for sequential and parallel composition), but employ different specialized assertions to describe the possible memory states, where specific soundness arguments are only needed for primitive memory operations.

The second contribution of this paper is an instance of the general RG framework for *causally consistent shared memory*. The latter stands for a family of wide-spread and well-studied memory models weaker than SC, which are sufficiently strong for implementing a variety of synchronization idioms [6,12,26]. Intuitively, unlike SC, causal consistency allows different threads to observe writes to memory in different orders, as long as they agree on the order of writes that are causally related. This concept can be formalized in multiple ways, and here we target a strong form of causal consistency, called *strong release-acquire* (SRA) [28,31] (and equivalent to "causal convergence" from [12]), which is a slight strengthening of the well-known release-acquire (RA) model (used by C/C++11). (The variants of causal consistency only differ for programs with write/write races [10,28], which are rather rare in practice.)

Our starting point for axiomatizing SRA as Hoare triples is the *potential-based* operational semantics of SRA, which was recently introduced with the goal of establishing the decidability of control state reachability under this model [27,28] (in contrast to undecidability under RA [1]). Unlike more standard presentations of weak memory models whose states record information about the *past* (e.g., in the form of store buffers containing executed writes before they are globally visible [36], partially ordered execution graphs [8,20,31], or collections of timestamped messages and thread views [11,16,17,23,25,47]), the states of the potential-based model track possible *futures* ascribing what sequences of observations each thread can perform. We find this approach to be a particularly appealing candidate for Hoare-style reasoning which would naturally generalize SC-based reasoning. Intuitively, while an assertion in SC specifies possible observations at a given program point, an assertion in a potential-based model should specify possible *sequences* of observations.

To pursue this direction, we introduce a novel assertion language, resembling temporal logics, which allows one to express properties of sequences of states. For instance, our assertions can express that a certain thread may currently read $\mathtt{x} = 0$, but it will have to read $\mathtt{x} = 1$ once it reads $\mathtt{y} = 1$. Then, we provide Hoare triples for SRA in this assertion language, and incorporate them in the general RG framework. The resulting program logic, which we call $\mathsf{Piccolo}$, provides a novel approach to reason on concurrent programs under causal consistency, which allows for simple and direct proofs, and, we believe, may constitute a basis for automation in the future.

$$\{\mathtt{y} \neq 1\}$$

Thread $\mathtt{T}_1$	**Thread** $\mathtt{T}_2$
$\{\mathit{True}\}$	$\{\mathtt{y} = 1 \Rightarrow \mathtt{x} = 1\}$
$1 : \mathtt{STORE}(\mathtt{x}, 1);$	$3 : \mathtt{a} := \mathtt{LOAD}(\mathtt{y});$
$\{\mathtt{x} = 1\}$	$\{\mathtt{a} = 1 \Rightarrow \mathtt{x} = 1\}$
$2 : \mathtt{STORE}(\mathtt{y}, 1)$	$4 : \mathtt{b} := \mathtt{LOAD}(\mathtt{x})$
$\{\mathit{True}\}$	$\{\mathtt{a} = 1 \Rightarrow \mathtt{b} = 1\}$

$$\{\mathtt{a} = 1 \Rightarrow \mathtt{b} = 1\}$$

Fig. 1. Message passing in SC

$$\{\mathtt{T}_0 \ltimes [\mathtt{y} \neq 1]\}$$

Thread $\mathtt{T}_1$	**Thread** $\mathtt{T}_2$
$\{\mathit{True}\}$	$\{\mathtt{T}_2 \ltimes [\mathtt{y} \neq 1] ; [\mathtt{x} = 1]\}$
$1 : \mathtt{STORE}(\mathtt{x}, 1);$	$3 : \mathtt{a} := \mathtt{LOAD}(\mathtt{y});$
$\{\mathtt{T}_1 \ltimes [\mathtt{x} = 1]\}$	$\{\mathtt{a} = 1 \Rightarrow \mathtt{T}_2 \ltimes [\mathtt{x} = 1]\}$
$2 : \mathtt{STORE}(\mathtt{y}, 1)$	$4 : \mathtt{b} := \mathtt{LOAD}(\mathtt{x})$
$\{\mathit{True}\}$	$\{\mathtt{a} = 1 \Rightarrow \mathtt{b} = 1\}$

$$\{\mathtt{a} = 1 \Rightarrow \mathtt{b} = 1\}$$

Fig. 2. Message passing in SRA

2 Motivating Example

To make our discussion concrete, consider the message passing program (MP) in Figures 1 and 2, comprising shared variables $\mathtt{x}$ and $\mathtt{y}$ and local registers $\mathtt{a}$ and $\mathtt{b}$. The proof outline in Fig. 1 assumes SC, whereas Fig. 2 assumes SRA. In both cases, at the end of the execution, we show that if $\mathtt{a}$ is 1, then $\mathtt{b}$ must also be 1. We use these examples to explain the two main concepts introduced in this paper: (i) a generic RG framework and (ii) its instantiation with a potential-focused assertion system that enables reasoning under SRA.

Rely-Guarantee. The proof outline in Fig. 1 can be read as an RG derivation:

1. Thread $\mathtt{T}_1$ locally establishes its postcondition when starting from any state that satisfies its precondition. This is trivial since its postcondition is *True*.
2. Thread $\mathtt{T}_1$ relies on the fact that its used assertions are *stable* w.r.t. interference from its environment. We formally capture this condition by a rely set $\mathcal{R}_1 \triangleq \{\mathit{True}, \mathtt{x} = 1\}$.
3. Thread $\mathtt{T}_1$ guarantees to its concurrent environment that its only interferences are $\mathtt{STORE}(\mathtt{x}, 1)$ and $\mathtt{STORE}(\mathtt{y}, 1)$, and furthermore that $\mathtt{STORE}(\mathtt{y}, 1)$ is only performed when $\mathtt{x} = 1$ holds. We formally capture this condition by a guarantee set $\mathcal{G}_1 \triangleq \{\{\mathit{True}\}\ \mathtt{T}_1 \mapsto \mathtt{STORE}(\mathtt{x}, 1), \{\mathtt{x} = 1\}\ \mathtt{T}_1 \mapsto \mathtt{STORE}(\mathtt{y}, 1)\}$, where each element is a command guarded by a precondition.
4. Thread $\mathtt{T}_2$ locally establishes its postcondition when starting from any state that satisfies its precondition. This is straightforward using standard Hoare rules for assignment and sequential composition.
5. Thread $\mathtt{T}_2$'s rely set is again obtained by collecting all the assertions used in its proof: $\mathcal{R}_2 \triangleq \{\mathtt{y} = 1 \Rightarrow \mathtt{x} = 1, \mathtt{a} = 1 \Rightarrow \mathtt{x} = 1, \mathtt{a} = 1 \Rightarrow \mathtt{b} = 1\}$. Indeed, the local reasoning for $\mathtt{T}_2$ needs all these assertions to be stable under the environment interference.
6. Thread $\mathtt{T}_2$'s guarantee set is given by:
$$\mathcal{G}_2 \triangleq \{\{\mathtt{y} = 1 \Rightarrow \mathtt{x} = 1\}\ \mathtt{T}_2 \mapsto \mathtt{a} := \mathtt{LOAD}(\mathtt{y}), \{\mathtt{a} = 1 \Rightarrow \mathtt{x} = 1\}\ \mathtt{T}_2 \mapsto \mathtt{b} := \mathtt{LOAD}(\mathtt{x})\}$$
7. To perform the parallel composition, $\langle \mathcal{R}_1, \mathcal{G}_1 \rangle$ and $\langle \mathcal{R}_2, \mathcal{G}_2 \rangle$ should be *non-interfering*. This involves showing that each $R \in \mathcal{R}_i$ is *stable* under each $G \in \mathcal{G}_j$ for $i \neq j$. That is, if $G = \{P\}\ \tau \mapsto c$, we require the Hoare triple

$\{P \cap R\}\ \tau \mapsto c\ \{R\}$ to hold. In this case, these proof obligations are straightforward to discharge using Hoare's assignment axiom (and is trivial for $i = 1$ and $j = 2$ since load instructions leave the memory intact).

Remark 1. Classical treatments of RG involve two related ideas [21]: (1) specifying a component by rely and guarantee conditions (together with standard pre- and postconditions); and (2) taking the relies and guarantees to be binary relations over states. Our approach adopts (1) but not (2). Thus, it can be seen as an RG presentation of the Owicki-Gries method [37], as was previously done in [32]. We have not observed an advantage for using binary relations in our examples, but the framework can be straightforwardly modified to do so.

Now, observe that substantial aspects of the above reasoning are *not* directly tied with SC. This includes the Hoare rules for compound commands (such as sequential composition above), the idea of specifying a thread using collections of stable rely assertions and guaranteed guarded primitive commands, and the non-interference condition for parallel composition. To carry out this generalization, we assume that we are provided an assertion language whose assertions are interpreted as *sets of memory states* (which can be much more involved than simple mappings of variables to values), and a set of valid Hoare triples for the primitive instructions. The latter is used for checking validity of primitive triples, (e.g., $\{P\}\ \mathtt{T}_1 \mapsto \mathtt{STORE}(\mathtt{x}, 1)\ \{Q\}$), as well as non-interference conditions (e.g., $\{P \cap R\}\ \mathtt{T}_1 \mapsto \mathtt{STORE}(\mathtt{x}, 1)\ \{R\}$). In §4, we present this generalization, and establish the soundness of RG principles independently of the memory model.

Potential-based reasoning. The second contribution of our work is an application of the above to develop a logic for a potential-based operational semantics that captures SRA. In this semantics every memory state records sequences of store mappings (from shared variables to values) that each thread may observe. For example, assuming all variables are initialized to 0, if $\mathtt{T}_1$ executed its code until completion before $\mathtt{T}_2$ even started (so under SC the memory state is the store $\{\mathtt{x} \mapsto 1, \mathtt{y} \mapsto 1\}$), we may reach the SRA state in which $\mathtt{T}_1$'s potential consists of one store $\{\mathtt{x} \mapsto 1, \mathtt{y} \mapsto 1\}$, and $\mathtt{T}_2$'s potential is the sequence of stores:

$$\langle \{\mathtt{x} \mapsto 0, \mathtt{y} \mapsto 0\}, \{\mathtt{x} \mapsto 1, \mathtt{y} \mapsto 0\}, \{\mathtt{x} \mapsto 1, \mathtt{y} \mapsto 1\} \rangle,$$

which captures the stores that $\mathtt{T}_2$ may observe in the order it may observe them. Naturally, potentials are *lossy* allowing threads to non-deterministically lose a subsequence of the current store sequence, so they can progress in their sequences. Thus, $\mathtt{T}_2$ can read 1 from y only after it loses the first two stores in its potential, and from this point on it can only read 1 from x. Now, one can see that *all* potentials of $\mathtt{T}_2$ at its initial program point are, in fact, subsequences of the above sequence (regardless of where $\mathtt{T}_1$ is), and conclude that $\mathtt{a} = 1 \Rightarrow \mathtt{b} = 1$ holds when $\mathtt{T}_2$ terminates.

To capture the above informal reasoning in a Hoare logic, we designed a new form of assertions capturing possible locally observable sequences of stores, rather than one global store, which can be seen as a restricted fragment of linear

temporal logic. The proof outline using these assertions is given in Fig. 2. In particular, $[\mathtt{x}=1]$ is satisfied by all store sequences in which every store maps x to 1, whereas $[\mathtt{y}\neq 1]\,;[\mathtt{x}=1]$ is satisfied by all store sequences that can be split into a (possibly empty) prefix whose value for y is not 1 followed by a (possibly empty) suffix whose value for x is 1. Assertions of the form $\tau \ltimes I$ state that the potential of thread τ includes only store sequences that satisfy I.

The first assertion of $\mathtt{T}_2$ is implied by the initial condition, $\mathtt{T}_0 \ltimes [\mathtt{y}\neq 1]$, since the potential of the parent thread $\mathtt{T}_0$ is inherited by the forked child threads and $\mathtt{T}_2 \ltimes [\mathtt{y}\neq 1]$ implies $\mathtt{T}_2 \ltimes [\mathtt{y}\neq 1]\,; I$ for any I. Moreover, $\mathtt{T}_2 \ltimes [\mathtt{y}\neq 1]\,;[\mathtt{x}=1]$ is preserved by (i) line 1 because writing 1 to x leaves $[\mathtt{y}\neq 1]$ unchanged and re-establishes $[\mathtt{x}=1]$; and (ii) line 2 because the semantics for SRA ensures that after reading 1 from y by $\mathtt{T}_2$, the thread $\mathtt{T}_2$ is confined by $\mathtt{T}_1$'s potential just before it wrote 1 to y, which has to satisfy the precondition $\mathtt{T}_1 \ltimes [\mathtt{x}=1]$. (SRA allows to update the other threads' potential only when the suffix of the potential after the update is observable by the writer thread.)

In §6 we formalize these arguments as Hoare rules for the primitive instructions, whose soundness is checked using the potential-based operational semantics and the interpretation of the assertion language. Finally, Piccolo is obtained by incorporating these Hoare rules in the general RG framework.

Remark 2. Our presentation of the potential-based semantics for SRA (fully presented in §5) deviates from the original one in [28], where it was called loSRA. The most crucial difference is that while loSRA's potentials consist of lists of per-location read options, our potentials consist of lists of *stores* assigning a value to every variable. (This is similar in spirit to the adaptation of load buffers for TSO [4,5] to snapshot buffers in [2]). Additionally, unlike loSRA, we disallow empty potential lists, require that the potentials of the different threads agree on the very last value to each location, and handle read-modify-write (RMW) instructions differently. We employed these modifications to loSRA as we observed that direct reasoning on loSRA states is rather unnatural and counterintuitive, as loSRA allows traces that *block* a thread from reading any value from certain locations (which cannot happen in the version we formulate). For example, a direct interpretation of our assertions over loSRA states would allow states in which $\tau \ltimes [x=v]$ and $\tau \ltimes [x\neq v]$ both hold (when τ does not have any option to read from x), while these assertions are naturally contradictory when interpreted on top of our modified SRA semantics. To establish confidence in the new potential-based semantics we have proved *in Coq* its equivalence to the standard execution-graph based semantics of SRA (over 5K lines of Coq proofs) [29].

3 Preliminaries: Syntax and Semantics

In this section we describe the underlying program language, leaving the shared-memory semantics parametric.

Syntax. The syntax of programs, given in Fig. 3, is mostly standard, comprising primitive (atomic) commands c and compound commands C. The non-standard

$$\begin{array}{rlrl} \textit{values} & v \in \mathsf{Val} = \{0, 1, ...\} & \textit{shared variables} & x, y \in \mathsf{Loc} = \{\mathtt{x}, \mathtt{y}, ...\} \\ \textit{local registers} & r \in \mathsf{Reg} = \{\mathtt{a}, \mathtt{b}, ...\} & \textit{thread identifiers} & \tau, \pi \in \mathsf{Tid} = \{\mathtt{T}_0, \mathtt{T}_1, ...\} \end{array}$$

$$\begin{array}{l} e ::= r \mid v \mid e + e \mid e = e \mid \neg e \mid e \wedge e \mid e \vee e \mid \ldots \\ c ::= r := e \mid \mathtt{STORE}(x, e) \mid r := \mathtt{LOAD}(x) \mid \mathtt{SWAP}(x, e) \qquad \tilde{c} ::= \langle c, \vec{r} := \vec{e} \rangle \\ C ::= c \mid \tilde{c} \mid \mathtt{skip} \mid C \,;\, C \mid \mathtt{if}\ e\ \mathtt{then}\ C\ \mathtt{else}\ C \mid \mathtt{while}\ e\ \mathtt{do}\ C \mid C\ {}^{\tau}||^{\tau}\ C \end{array}$$

Fig. 3. Program syntax

$$\frac{\gamma' = \gamma[r \mapsto \gamma(e)]}{r := e \gg \gamma \xrightarrow{\varepsilon} \gamma'} \qquad \frac{l = \mathtt{W}(x, \gamma(e))}{\mathtt{STORE}(x, e) \gg \gamma \xrightarrow{l} \gamma} \qquad \frac{l = \mathtt{R}(x, v) \qquad \gamma' = \gamma[r \mapsto v]}{r := \mathtt{LOAD}(x) \gg \gamma \xrightarrow{l} \gamma'}$$

$$\frac{l = \mathtt{RMW}(x, v, \gamma(e))}{\mathtt{SWAP}(x, e) \gg \gamma \xrightarrow{l} \gamma} \qquad \frac{c \gg \gamma \xrightarrow{l_\varepsilon} \gamma_0 \qquad r_1 := e_1 \gg \gamma_0 \xrightarrow{\varepsilon} \gamma_1 \quad \ldots \quad r_n := e_n \gg \gamma_{n-1} \xrightarrow{\varepsilon} \gamma_n}{\langle c, \langle r_1, \ldots, r_n \rangle := \langle e_1, \ldots, e_n \rangle \rangle \gg \gamma \xrightarrow{l_\varepsilon} \gamma_n}$$

Fig. 4. Small-step semantics of (instrumented) primitive commands ($\tilde{c} \gg \gamma \xrightarrow{l_\varepsilon} \gamma'$)

components are instrumented commands $\tilde{c}$, which are meant to atomically execute a primitive command c and a (multiple) assignment $\vec{r} := \vec{e}$. Such instructions are needed to support auxiliary (a.k.a. ghost) variables in RG proofs. In addition, `SWAP` (a.k.a. atomic exchange) is an example of an RMW instruction. For brevity, other standard RMW instructions, such as `FADD` and `CAS`, are omitted.

Unlike many weak memory models that only support top-level parallelism, we include dynamic thread creation via commands of the form $C_1\ {}^{\tau_1}||^{\tau_2}\ C_2$ that forks two threads named τ_1 and τ_2 that execute the commands C_1 and C_2, respectively. Each C_i may itself comprise further parallel compositions. Since thread identifiers are explicit, we require commands to be *well formed*. Let $\mathsf{Tid}(C)$ be the set of all thread identifiers that appear in C. A command C is *well formed*, denoted $\mathsf{wf}(C)$, if parallel compositions inside employ disjoint sets of thread identifiers. This notion is formally defined by induction on the structure of commands, with the only interesting case being $\mathsf{wf}(C_1\ {}^{\tau_1}||^{\tau_2}\ C_2)$ if $\mathsf{wf}(C_1) \wedge \mathsf{wf}(C_2) \wedge \tau_1 \neq \tau_2 \wedge \mathsf{Tid}(C_1) \cap \mathsf{Tid}(C_2) = \emptyset$.

Program semantics. We provide small-step operational semantics to commands independently of the memory system. To connect this semantics to a given memory system, its steps are instrumented with labels, as defined next.

Definition 1. A *label* l takes one of the following forms: a read $\mathtt{R}(x, v_\mathtt{R})$, a write $\mathtt{W}(x, v_\mathtt{W})$, a read-modify-write $\mathtt{RMW}(x, v_\mathtt{R}, v_\mathtt{W})$, a fork $\mathtt{FORK}(\tau_1, \tau_2)$, or a join $\mathtt{JOIN}(\tau_1, \tau_2)$, where $x \in \mathsf{Loc}$, $v_\mathtt{R}, v_\mathtt{W} \in \mathsf{Val}$, and $\tau_1, \tau_2 \in \mathsf{Tid}$. We denote by Lab the set of all labels.

$$\frac{\tilde{c} \gg \gamma \xrightarrow{l_\varepsilon} \gamma'}{\langle \tilde{c}, \gamma \rangle \xrightarrow{l_\varepsilon} \langle \mathtt{skip}, \gamma' \rangle} \qquad \frac{\langle C_1, \gamma \rangle \xrightarrow{l_\varepsilon} \langle C_1', \gamma' \rangle}{\langle C_1 \mathbin{;} C_2, \gamma \rangle \xrightarrow{l_\varepsilon} \langle C_1' \mathbin{;} C_2, \gamma' \rangle} \qquad \frac{}{\langle \mathtt{skip} \mathbin{;} C_2, \gamma \rangle \xrightarrow{\varepsilon} \langle C_2, \gamma \rangle}$$

$$\frac{\begin{array}{c}\gamma(e) = \mathit{true} \Rightarrow i = 1 \\ \gamma(e) \neq \mathit{true} \Rightarrow i = 2\end{array}}{\langle \mathtt{if}\ e\ \mathtt{then}\ C_1\ \mathtt{else}\ C_2, \gamma \rangle \xrightarrow{\varepsilon} \langle C_i, \gamma \rangle} \qquad \frac{C' = \mathtt{if}\ e\ \mathtt{then}\ (C \mathbin{;} \mathtt{while}\ e\ \mathtt{do}\ C)\ \mathtt{else\ skip}}{\langle \mathtt{while}\ e\ \mathtt{do}\ C, \gamma \rangle \xrightarrow{\varepsilon} \langle C', \gamma \rangle}$$

Fig. 5. Small-step semantics of commands ($\langle C, \gamma \rangle \xrightarrow{l_\varepsilon} \langle C', \gamma' \rangle$)

$$\frac{\langle C, \gamma \rangle \xrightarrow{l_\varepsilon} \langle C', \gamma' \rangle}{\begin{array}{l}\langle \mathcal{C}_0 \uplus \{\tau \mapsto C\}, \gamma \rangle \\ \xrightarrow{\tau, l_\varepsilon} \langle \mathcal{C}_0 \uplus \{\tau \mapsto C'\}, \gamma \rangle\end{array}} \qquad \frac{\begin{array}{c}\mathcal{C}(\tau) = C_1 \ {}^{\tau_1}\|^{\tau_2}\ C_2 \\ \tau_1 \notin \mathit{dom}(\mathcal{C}) \qquad \tau_2 \notin \mathit{dom}(\mathcal{C}) \\ l = \mathtt{FORK}(\tau_1, \tau_2) \\ \mathcal{C}' = \{\tau_1 \mapsto C_1, \tau_2 \mapsto C_2\}\end{array}}{\langle \mathcal{C}, \gamma \rangle \xrightarrow{\tau, l} \langle \mathcal{C} \uplus \mathcal{C}', \gamma \rangle} \qquad \frac{\begin{array}{c}\mathcal{C} = \left\{\begin{array}{l}\tau \mapsto C_1 \ {}^{\tau_1}\|^{\tau_2}\ C_2, \\ \tau_1 \mapsto \mathtt{skip}, \tau_2 \mapsto \mathtt{skip}\end{array}\right\} \\ l = \mathtt{JOIN}(\tau_1, \tau_2) \\ \mathcal{C}' = \{\tau \mapsto \mathtt{skip}\}\end{array}}{\langle \mathcal{C}_0 \uplus \mathcal{C}, \gamma \rangle \xrightarrow{\tau, l} \langle \mathcal{C}_0 \uplus \mathcal{C}', \gamma \rangle}$$

Fig. 6. Small-step semantics of command pools ($\langle \mathcal{C}, \gamma \rangle \xrightarrow{\tau, l_\varepsilon} \langle \mathcal{C}', \gamma' \rangle$)

Definition 2. A *register store* is a mapping $\gamma : \mathsf{Reg} \to \mathsf{Val}$. Register stores are extended to expressions as expected. We denote by Γ the set of all register stores.

The semantics of (instrumented) primitive commands is given in Fig. 4. Using this definition, the semantics of commands is given in Fig. 5. Its steps are of the form $\langle C, \gamma \rangle \xrightarrow{l_\varepsilon} \langle C', \gamma' \rangle$ where C and C' are commands, γ and γ' are register stores, and $l_\varepsilon \in \mathsf{Lab} \cup \{\varepsilon\}$ (ε denotes a thread internal step). We lift this semantics to *command pools* as follows.

Definition 3. A *command pool* is a non-empty partial function $\mathcal{C}$ from thread identifiers to commands, such that the following hold:
1. $\mathsf{Tid}(\mathcal{C}(\tau_1)) \cap \mathsf{Tid}(\mathcal{C}(\tau_2)) = \emptyset$ for every $\tau_1 \neq \tau_2$ in $\mathit{dom}(\mathcal{C})$.
2. $\tau \notin \mathsf{Tid}(\mathcal{C}(\tau))$ for every $\tau \in \mathit{dom}(\mathcal{C})$.

We write command pools as sets of the form $\{\tau_1 \mapsto C_1, \ldots, \tau_n \mapsto C_n\}$.

Steps for command pools are given in Fig. 6. They take the form $\langle \mathcal{C}, \gamma \rangle \xrightarrow{\tau, l_\varepsilon} \langle \mathcal{C}', \gamma' \rangle$, where $\mathcal{C}$ and $\mathcal{C}'$ are command pools, γ and γ' are register stores, and $\langle \tau : l_\varepsilon \rangle$ (with $\tau \in \mathsf{Tid}$ and $l_\varepsilon \in \mathsf{Lab} \cup \{\varepsilon\}$) is a *command transition label.*

Memory semantics. To give semantics to programs under a memory model, we synchronize the transitions of a command C with a memory system. We leave the memory system parametric, and assume that it is represented by a labeled transition system (LTS) $\mathcal{M}$ with set of states denoted by $\mathcal{M}.\mathtt{Q}$, and steps denoted by $\rightarrow_{\mathcal{M}}$. The transition labels of general memory system $\mathcal{M}$ consist of non-silent program transition labels (elements of $\mathsf{Tid} \times \mathsf{Lab}$) and a (disjoint) set $\mathcal{M}.\mathbf{\Theta}$ of internal memory actions, which is again left parametric (used, e.g., for memory-internal propagation of values).

Example 1. The simple memory system that guarantees sequential consistency is denoted here by SC. This memory system tracks the most recent value written to each variable and has no internal transitions ($\mathsf{SC}.\mathbf{\Theta} = \emptyset$). Formally, it is defined by $\mathsf{SC}.\mathtt{Q} \triangleq \mathsf{Loc} \to \mathsf{Val}$ and $\to_{\mathsf{SC}}$ is given by:

$$\frac{l = \mathtt{R}(x, v_{\mathtt{R}}) \quad m(x) = v_{\mathtt{R}}}{m \xrightarrow{\tau,l}_{\mathsf{SC}} m} \qquad \frac{l = \mathtt{W}(x, v_{\mathtt{W}}) \quad m' = m[x \mapsto v_{\mathtt{W}}]}{m \xrightarrow{\tau,l}_{\mathsf{SC}} m'} \qquad \frac{l = \mathtt{RMW}(x, v_{\mathtt{R}}, v_{\mathtt{W}}) \quad m(x) = v_{\mathtt{R}} \quad m' = m[x \mapsto v_{\mathtt{W}}]}{m \xrightarrow{\tau,l}_{\mathsf{SC}} m'} \qquad \frac{l \in \{\mathtt{FORK}(_,_), \mathtt{JOIN}(_,_)\}}{m \xrightarrow{\tau,l}_{\mathsf{SC}} m}$$

The composition of a program with a general memory system is defined next.

Definition 4. The *concurrent system* induced by a memory system $\mathcal{M}$, denoted by $\overline{\mathcal{M}}$, is the LTS whose transition labels are the elements of $(\mathsf{Tid} \times (\mathsf{Lab} \cup \{\varepsilon\})) \uplus \mathcal{M}.\mathbf{\Theta}$; states are triples of the form $\langle \mathcal{C}, \gamma, m \rangle$ where $\mathcal{C}$ is a command pool, γ is a register store, and $m \in \mathcal{M}.\mathtt{Q}$; and the transitions are "synchronized transitions" of the program and the memory system, using labels to decide what to synchronize on, formally given by:

$$\frac{\langle \mathcal{C}, \gamma \rangle \xrightarrow{\tau,l} \langle \mathcal{C}', \gamma' \rangle \quad l \in \mathsf{Lab} \quad m \xrightarrow{\tau,l}_{\mathcal{M}} m'}{\langle \mathcal{C}, \gamma, m \rangle \xrightarrow{\tau,l}_{\overline{\mathcal{M}}} \langle \mathcal{C}', \gamma', m' \rangle} \qquad \frac{\langle \mathcal{C}, \gamma \rangle \xrightarrow{\tau,\varepsilon} \langle \mathcal{C}', \gamma' \rangle}{\langle \mathcal{C}, \gamma, m \rangle \xrightarrow{\tau,\varepsilon}_{\overline{\mathcal{M}}} \langle \mathcal{C}', \gamma', m \rangle} \qquad \frac{\theta \in \mathcal{M}.\mathbf{\Theta} \quad m \xrightarrow{\theta}_{\mathcal{M}} m'}{\langle \mathcal{C}, \gamma, m \rangle \xrightarrow{\theta}_{\overline{\mathcal{M}}} \langle \mathcal{C}, \gamma, m' \rangle}$$

4 Generic Rely-Guarantee Reasoning

In this section we present our generic RG framework. Rather than committing to a specific assertion language, our reasoning principles apply on the *semantic level*, using sets of states instead of syntactic assertions. The structure of proofs still follows program structure, thereby retaining RG's compositionality. By doing so, we decouple the semantic insights of RG reasoning from a concrete syntax. Next, we present proof rules serving as blueprints for memory model specific proof systems. An instantiation of this blueprint requires lifting the semantic principles to syntactic ones. More specifically, it requires

1. a language with (a) concrete assertions for specifying sets of states and (b) operators that match operations on sets of states (like $\wedge$ matches $\cap$); and
2. sound Hoare triples for primitive commands.

Thus, each instance of the framework (for a specific memory system) is left with the task of identifying useful abstractions on states, as well as a suitable formalism, for making the generic semantic framework into a proof system.

RG judgments. We let $\mathcal{M}$ be an arbitrary memory system and $\Sigma_{\mathcal{M}} \triangleq \Gamma \times \mathcal{M}.\mathtt{Q}$. Properties of programs $\mathcal{C}$ are stated via *RG judgments*:

$$\mathcal{C}\ \underline{sat}_{\mathcal{M}}\ (P, \mathcal{R}, \mathcal{G}, Q)$$

where $P, Q \subseteq \Sigma_{\mathcal{M}}$, $\mathcal{R} \subseteq \mathcal{P}(\Sigma_{\mathcal{M}})$, and $\mathcal{G}$ is a set of *guarded commands*, each of which takes the form $\{G\}\ \tau \mapsto \alpha$, where $G \subseteq \Sigma_{\mathcal{M}}$ and α is either an (instrumented) primitive command $\tilde{c}$ or a fork/join label (of the form $\texttt{FORK}(\tau_1, \tau_2)$ or $\texttt{JOIN}(\tau_1, \tau_2)$). The latter is needed for considering the effect of forks and joins on the memory state.

Interpretation of RG judgments. RG judgments $\mathcal{C}\ \underline{sat}_{\mathcal{M}}\ (P, \mathcal{R}, \mathcal{G}, Q)$ state that a terminating run of $\mathcal{C}$ starting from a state in P, under any concurrent context whose transitions preserve each of the sets of states in $\mathcal{R}$, will end in a state in Q and perform only transitions contained in $\mathcal{G}$. To formally define this statement, following the standard model for RG, these judgments are interpreted on *computations* of programs. Computations arise from runs of the concurrent system (see Def. 4) by abstracting away from concrete transition labels and including arbitrary "environment transitions" representing steps of the concurrent context. We have:

- *Component* transitions of the form $\langle \mathcal{C}, \gamma, m \rangle \xrightarrow{\texttt{cmp}} \langle \mathcal{C}', \gamma', m' \rangle$.
- *Memory* transitions, which correspond to internal memory steps (labeled with $\theta \in \mathcal{M}.\mathbf{\Theta}$), of the form $\langle \mathcal{C}, \gamma, m \rangle \xrightarrow{\texttt{mem}} \langle \mathcal{C}, \gamma, m' \rangle$.
- *Environment* transitions of the form $\langle \mathcal{C}, \gamma, m \rangle \xrightarrow{\texttt{env}} \langle \mathcal{C}, \gamma', m' \rangle$.

Note that memory transitions do not occur in the classical RG presentation (since SC does not have internal memory actions).

A *computation* is a (potentially infinite) sequence

$$\xi = \langle \mathcal{C}_0, \gamma_0, m_0 \rangle \xrightarrow{a_1} \langle \mathcal{C}_1, \gamma_1, m_1 \rangle \xrightarrow{a_2} ...$$

with $a_i \in \{\texttt{cmp}, \texttt{env}, \texttt{mem}\}$. We let $\langle \mathcal{C}_{\mathsf{last}(\xi)}, \gamma_{\mathsf{last}(\xi)}, m_{\mathsf{last}(\xi)} \rangle$ denotes its last element, when ξ is finite. We say that ξ is a computation *of a command pool* $\mathcal{C}$ when $\mathcal{C}_0 = \mathcal{C}$ and for every $i \geq 0$:

- If $a_i = \texttt{cmp}$, then $\langle \mathcal{C}_i, \gamma_i, m_i \rangle \xrightarrow{\tau, l_\varepsilon}_{\mathcal{M}} \langle \mathcal{C}_{i+1}, \gamma_{i+1}, m_{i+1} \rangle$ for some $\tau \in \mathsf{Tid}$ and $l_\varepsilon \in \mathsf{Lab} \cup \{\varepsilon\}$.
- If $a_i = \texttt{mem}$, then $\langle \mathcal{C}_i, \gamma_i, m_i \rangle \xrightarrow{\theta}_{\mathcal{M}} \langle \mathcal{C}_{i+1}, \gamma_{i+1}, m_{i+1} \rangle$ for some $\theta \in \mathcal{M}.\mathbf{\Theta}$.

We denote by $\mathit{Comp}(\mathcal{C})$ the set of all computations of a command pool $\mathcal{C}$.

To define validity of RG judgments, we use the following definition.

Definition 5. Let $\xi = \langle \mathcal{C}_0, \gamma_0, m_0 \rangle \xrightarrow{a_1} \langle \mathcal{C}_1, \gamma_1, m_1 \rangle \xrightarrow{a_2} ...$ be a computation, and $\mathcal{C}\ \underline{sat}_{\mathcal{M}}\ (P, \mathcal{R}, \mathcal{G}, Q)$ an RG-judgment.

- ξ admits P if $\langle \gamma_0, m_0 \rangle \in P$.
- ξ admits $\mathcal{R}$ if $\langle \gamma_i, m_i \rangle \in R \Rightarrow \langle \gamma_{i+1}, m_{i+1} \rangle \in R$ for every $R \in \mathcal{R}$ and $i \geq 0$ with $a_{i+1} = \texttt{env}$.
- ξ admits $\mathcal{G}$ if for every $i \geq 0$ with $a_{i+1} = \texttt{cmp}$ and $\langle \gamma_i, m_i \rangle \neq \langle \gamma_{i+1}, m_{i+1} \rangle$ there exists $\{P\}\ \tau \mapsto \alpha \in \mathcal{G}$ such that $\langle \gamma_i, m_i \rangle \in P$ and
  - if $\alpha = \tilde{c}$ is an instrumented primitive command, then for some $l_\varepsilon \in \mathsf{Lab} \cup \{\varepsilon\}$, we have $\langle \{\tau \mapsto \tilde{c}\}, \gamma_i, m_i \rangle \xrightarrow{\tau, l_\varepsilon}_{\mathcal{M}} \langle \{\tau \mapsto \texttt{skip}\}, \gamma_{i+1}, m_{i+1} \rangle$

$$\textsc{skip}\ \frac{}{\{\tau \mapsto \texttt{skip}\}\ \underline{sat}_{\mathcal{M}}\ (P, \{P\}, \emptyset, P)} \qquad \textsc{com}\ \frac{\mathcal{M} \vDash \{P\}\ \tau \mapsto \tilde{c}\ \{Q\}}{\{\tau \mapsto \tilde{c}\}\ \underline{sat}_{\mathcal{M}}\ (P, \{P, Q\}, \{\{P\}\ \tau \mapsto \tilde{c}\}, Q)}$$

$$\textsc{seq}\ \frac{\{\tau \mapsto C_1\}\ \underline{sat}_{\mathcal{M}}\ (P, \mathcal{R}_1, \mathcal{G}_1, R) \qquad \{\tau \mapsto C_2\}\ \underline{sat}_{\mathcal{M}}\ (R, \mathcal{R}_2, \mathcal{G}_2, Q)}{\{\tau \mapsto C_1\ ; C_2\}\ \underline{sat}_{\mathcal{M}}\ (P, \mathcal{R}_1 \cup \mathcal{R}_2, \mathcal{G}_1 \cup \mathcal{G}_2, Q)}$$

$$\textsc{if}\ \frac{\{\tau \mapsto C_1\}\ \underline{sat}_{\mathcal{M}}\ (P \cap [\![e]\!], \mathcal{R}_1, \mathcal{G}_1, Q) \qquad \{\tau \mapsto C_2\}\ \underline{sat}_{\mathcal{M}}\ (P \setminus [\![e]\!], \mathcal{R}_2, \mathcal{G}_2, Q)}{\{\tau \mapsto \texttt{if}\ e\ \texttt{then}\ C_1\ \texttt{else}\ C_2\}\ \underline{sat}_{\mathcal{M}}\ (P, \mathcal{R}_1 \cup \mathcal{R}_2 \cup \{P\}, \mathcal{G}_1 \cup \mathcal{G}_2, Q)}$$

$$\textsc{while}\ \frac{P \setminus [\![e]\!] \subseteq Q \qquad \{\tau \mapsto C\}\ \underline{sat}_{\mathcal{M}}\ (P \cap [\![e]\!], \mathcal{R}, \mathcal{G}, P)}{\{\tau \mapsto \texttt{while}\ e\ \texttt{do}\ C\}\ \underline{sat}_{\mathcal{M}}\ (P, \mathcal{R} \cup \{P, Q\}, \mathcal{G}, Q)}$$

$$\textsc{par}\ \frac{\begin{array}{c}\{\tau_1 \mapsto C_1\}\ \underline{sat}_{\mathcal{M}}\ (P_1, \mathcal{R}_1, \mathcal{G}_1, Q_1) \qquad \{\tau_2 \mapsto C_2\}\ \underline{sat}_{\mathcal{M}}\ (P_2, \mathcal{R}_2, \mathcal{G}_2, Q_2) \\ P \subseteq P_1 \cap P_2 \qquad Q_1 \cap Q_2 \subseteq Q \qquad \langle \mathcal{R}_1, \mathcal{G}_1 \rangle \text{ and } \langle \mathcal{R}_2, \mathcal{G}_2 \rangle \text{ are non-interfering}\end{array}}{\{\tau_1 \mapsto C_1\} \uplus \{\tau_2 \mapsto C_2\}\ \underline{sat}_{\mathcal{M}}\ (P, \mathcal{R}_1 \cup \mathcal{R}_2 \cup \{P, Q\}, \mathcal{G}_1 \cup \mathcal{G}_2, Q)}$$

$$\textsc{fork-join}\ \frac{\begin{array}{c}\mathcal{M} \vDash \{P\}\ \tau \mapsto \texttt{FORK}(\tau_1, \tau_2)\ \{P'\} \qquad \mathcal{M} \vDash \{Q'\}\ \tau \mapsto \texttt{JOIN}(\tau_1, \tau_2)\ \{Q\} \\ \{\tau_1 \mapsto C_1\} \uplus \{\tau_2 \mapsto C_2\}\ \underline{sat}_{\mathcal{M}}\ (P', \mathcal{R}, \mathcal{G}, Q') \\ \mathcal{G}' = \mathcal{G} \cup \{\{P\}\ \tau \mapsto \texttt{FORK}(\tau_1, \tau_2), \{Q'\}\ \tau \mapsto \texttt{JOIN}(\tau_1, \tau_2)\}\end{array}}{\{\tau \mapsto C_1\ {}^{\tau_1}\|^{\tau_2}\ C_2\}\ \underline{sat}_{\mathcal{M}}\ (P, \mathcal{R} \cup \{P, Q\}, \mathcal{G}', Q)}$$

Fig. 7. Generic sequential RG proof rules (letting $[\![e]\!] = \{\langle \gamma, m \rangle \mid \gamma(e) = \mathit{true}\}$)

  - if $\alpha \in \{\texttt{FORK}(\tau_1, \tau_2), \texttt{JOIN}(\tau_1, \tau_2)\}$, then $m_i \xrightarrow{\tau, \alpha}_{\mathcal{M}} m_{i+1}$ and $\gamma_i = \gamma_{i+1}$.
- ξ admits Q if $\langle \gamma_{\mathsf{last}(\xi)}, m_{\mathsf{last}(\xi)} \rangle \in Q$ whenever ξ is finite and $\mathcal{C}_{\mathsf{last}(\xi)}(\tau) = \texttt{skip}$ for every $\tau \in \mathit{dom}(\mathcal{C}_{\mathsf{last}(\xi)})$.

We denote by $\mathsf{Assume}(P, \mathcal{R})$ the set of all computations that admit P and $\mathcal{R}$, and by $\mathsf{Commit}(\mathcal{G}, Q)$ the set of all computations that admit $\mathcal{G}$ and Q.

Then, *validity* of a judgment if defined as

$$\models \mathcal{C}\ \underline{sat}_{\mathcal{M}}\ (P, \mathcal{R}, \mathcal{G}, Q) \overset{\triangle}{\Leftrightarrow} \mathit{Comp}(\mathcal{C}) \cap \mathsf{Assume}(P, \mathcal{R}) \subseteq \mathsf{Commit}(\mathcal{G}, Q)$$

Memory triples. Our proof rules build on *memory triples*, which specify pre- and postconditions for primitive commands for a memory system $\mathcal{M}$.

Definition 6. A *memory triple for a memory system* $\mathcal{M}$ is a tuple of the form $\{P\}\ \tau \mapsto \alpha\ \{Q\}$, where $P, Q \subseteq \Sigma_{\mathcal{M}}$, $\tau \in \mathsf{Tid}$, and α is either an instrumented primitive command, a fork label, or a join label. A memory triple for $\mathcal{M}$ is *valid*, denoted by $\mathcal{M} \vDash \{P\}\ \tau \mapsto \alpha\ \{Q\}$, if the following hold for every $\langle \gamma, m \rangle \in P$, $\gamma' \in \Gamma$ and $m' \in \mathcal{M}.\mathtt{Q}$:

- if α is an instrumented primitive command and $\langle \{\tau \mapsto \alpha\}, \gamma, m \rangle \xrightarrow{\tau, l_\varepsilon}_{\mathcal{M}} \langle \{\tau \mapsto \texttt{skip}\}, \gamma', m' \rangle$ for some $l_\varepsilon \in \mathsf{Lab} \cup \{\varepsilon\}$, then $\langle \gamma', m' \rangle \in Q$.

– If $\alpha \in \{\texttt{FORK}(\tau_1, \tau_2), \texttt{JOIN}(\tau_1, \tau_2)\}$ and $m \xrightarrow{\tau,\alpha}_{\mathcal{M}} m'$, then $\langle \gamma, m' \rangle \in Q$.

Example 2. For the memory system SC introduced in Ex. 1, we have, e.g., memory triples of the form $\mathsf{SC} \vDash \{e(r := x)\}\ \tau \mapsto r := \texttt{LOAD}(x)\ \{e\}$ (where $e(r := x)$ is the expression e with all occurrences of r replaced by x).

RG proof rules. We aim at proof rules deriving valid RG judgments. Figure 7 lists (semantic) proof rules based on externally provided memory triples. These rules basically follows RG reasoning for sequential consistency. For example, rule SEQ states that RG judgments of commands C_1 and C_2 can be combined when the postcondition of C_1 and the precondition of C_2 agree, thereby uniting their relies and guarantees. Rule COM builds on memory triples. The rule PAR for parallel composition combines judgments for two components when their relies and guarantees are *non-interfering*. Intuitively speaking, this means that each of the assertions that each thread relied on for establishing its proof is preserved when applying any of the assignments collected in the guarantee set of the other thread. An example of non-interfering rely-guarantee pairs is given in step 7 in §2. Formally, non-interference is defined as follows:

Definition 7. Rely-guarantee pairs $\langle \mathcal{R}_1, \mathcal{G}_1 \rangle$ and $\langle \mathcal{R}_2, \mathcal{G}_2 \rangle$ are *non-interfering* if $\mathcal{M} \vDash \{R \cap P\}\ \tau \mapsto \alpha\ \{R\}$ holds for every $R \in \mathcal{R}_1$ and $\{P\}\ \tau \mapsto \alpha \in \mathcal{G}_2$, and similarly for every $R \in \mathcal{R}_2$ and $\{P\}\ \tau \mapsto \alpha \in \mathcal{G}_1$.

In turn, FORK-JOIN combines the proof of a parallel composition with proofs of fork and join steps (which may also affect the memory state). Note that the guarantees also involve guarded commands with FORK and JOIN labels.

Additional rules for consequence and introduction of auxiliary variables are elided here (they are similar to their SC counterparts), and provided in the extended version of this paper [30].

Soundness. To establish soundness of the above system we need an additional requirement regarding the internal memory transitions (for SC this closure vacuously holds as there are no such transitions). We require all relies in $\mathcal{R}$ to be *stable under internal memory transitions*, i.e. for $R \in \mathcal{R}$ we require

$$\forall \gamma, m, m', \theta \in \mathcal{M}.\mathbf{\Theta}.\ m \xrightarrow{\theta}_{\mathcal{M}} m' \Rightarrow (\langle \gamma, m \rangle \in R \Rightarrow \langle \gamma, m' \rangle \in R) \qquad (\mathsf{mem})$$

This condition is needed since the memory system can non-deterministically take its internal steps, and the component's proof has to be stable under such steps.

With this requirement, we are able to establish soundness. The proof, which generally follows [48] is is given in the extended version of this paper [30]. We write $\vdash \mathcal{C}\ \underline{sat}_{\mathcal{M}}\ (P, \mathcal{R}, \mathcal{G}, Q)$ for provability of a judgment using the semantic rules presented above.

Theorem 1 (Soundness). $\vdash C\ \underline{sat}_{\mathcal{M}}\ (P, \mathcal{R}, \mathcal{G}, Q) \implies \vDash C\ \underline{sat}_{\mathcal{M}}\ (P, \mathcal{R}, \mathcal{G}, Q)$.

5 Potential-based Memory System for SRA

In this section we present the potential-based semantics for Strong Release-Acquire (SRA), for which we develop a novel RG logic. Our semantics is based on the one in [27, 28], with certain adaptations to make it better suited for Hoare-style reasoning (see Remark 2).

In weak memory models, threads typically have different views of the shared memory. In SRA, we refer to a memory snapshot that a thread may observe as a *potential store*:

Definition 8. A *potential store* is a function $\delta : \mathsf{Loc} \to \mathsf{Val} \times \{\mathtt{R}, \mathtt{RMW}\} \times \mathsf{Tid}$. We write $\mathtt{val}(\delta(x))$, $\mathtt{rmw}(\delta(x))$, and $\mathtt{tid}(\delta(x))$ to retrieve the different components of $\delta(x)$. We denote by Δ the set of all potential stores.

Having $\delta(x) = \langle v, \mathtt{R}, \tau \rangle$ allows to read the value v from x (and further ascribes that this read reads from a write performed by thread τ, which is technically needed to properly characterize the SRA model). In turn, having $\delta(x) = \langle v, \mathtt{RMW}, \tau \rangle$ further allows to perform an RMW instruction that atomically reads and modifies x.

Potential stores are collected in *potential store lists* describing the values which can (potentially) be read and in what order.

Notation 9 Lists over an alphabet A are written as $L = a_1 \cdot \ldots \cdot a_n$ where $a_1, \ldots, a_n \in A$. We also use $\cdot$ to concatenate lists, and write $L[i]$ for the i'th element of L and $|L|$ for the length of L.

A *(potential) store list* is a finite sequence of potential stores ascribing a possible sequence of stores that a thread can observe, in the order it will observe them. The RMW-flags in these lists have to satisfy certain conditions: once the flag for a location is set, it remains set in the rest of the list; and the flag must be set at the end of the list. Formally, store lists are defined as follows.

Definition 10. A *store list* $L \in \mathcal{L}$ is a non-empty finite sequence of potential stores with *monotone RMW-flags* ending with an RMW, that is: for all $x \in \mathsf{Loc}$,

1. if $\mathtt{rmw}(L[i](x)) = \mathtt{RMW}$, then $\mathtt{rmw}(L[j](x)) = \mathtt{RMW}$ for every $i < j \leq |L|$, and
2. $\mathtt{rmw}(L[|L|](x)) = \mathtt{RMW}$.

Now, SRA states ($\mathsf{SRA}.\mathtt{Q}$) consist of *potential mappings* that assign potentials to threads as defined next.

Definition 11. A *potential* D is a non-empty set of potential store lists. A *potential mapping* is a function $\mathcal{D} : \mathsf{Tid} \rightharpoonup \mathcal{P}(\mathcal{L}) \setminus \{\emptyset\}$ that maps thread identifiers to potentials such that all lists agree on the very final potential store (that is: $L_1[|L_1|] = L_2[|L_2|]$ whenever $L_1 \in \mathcal{D}(\tau_1)$ and $L_2 \in \mathcal{D}(\tau_2)$).

These potential mappings are "lossy" meaning that potential stores can be arbitrarily dropped. In particular, dropping the first store in a list enables reading from the second. This is formally done by transitioning from a state $\mathcal{D}$ to a "smaller" state $\mathcal{D}'$ as defined next.

WRITE
$$\frac{\begin{array}{l}\forall L' \in \mathcal{D}'(\tau).\, \exists L \in \mathcal{D}(\tau).\, L' = L[x \mapsto \langle v_{\mathtt{W}}, \mathtt{RMW}, \tau\rangle] \\ \quad \forall \pi \in dom(\mathcal{D}) \setminus \{\tau\}, L' \in \mathcal{D}'(\pi).\, \exists L_0, L_1. \\ \qquad L_0 \cdot L_1 \in \mathcal{D}(\pi) \wedge L_1 \in \mathcal{D}(\tau) \wedge \\ \qquad L' = L_0[x\mapsto \mathtt{R}] \cdot L_1[x \mapsto \langle v_{\mathtt{W}}, \mathtt{RMW}, \tau\rangle]\end{array}}{\mathcal{D} \xrightarrow{\tau, \mathtt{W}(x, v_{\mathtt{W}})}_{\mathsf{SRA}} \mathcal{D}'}$$

LOSE
$$\frac{\mathcal{D}' \sqsubseteq \mathcal{D}}{\mathcal{D} \xrightarrow{\varepsilon}_{\mathsf{SRA}} \mathcal{D}'}$$

DUP
$$\frac{\mathcal{D} \preceq \mathcal{D}'}{\mathcal{D} \xrightarrow{\varepsilon}_{\mathsf{SRA}} \mathcal{D}'}$$

READ
$$\frac{\begin{array}{l}\exists \pi.\, \forall L \in \mathcal{D}(\tau).\, \mathtt{val}(L[1](x)) = v_{\mathtt{R}} \wedge \\ \qquad\qquad \mathtt{tid}(L[1](x)) = \pi\end{array}}{\mathcal{D} \xrightarrow{\tau, \mathtt{R}(x, v_{\mathtt{R}})}_{\mathsf{SRA}} \mathcal{D}}$$

RMW
$$\frac{\begin{array}{c}\forall L \in \mathcal{D}(\tau).\, \mathtt{rmw}(L[1](x)) = \mathtt{RMW} \\ \mathcal{D} \xrightarrow{\tau, \mathtt{R}(x, v_{\mathtt{R}})}_{\mathsf{SRA}} \mathcal{D} \qquad \mathcal{D} \xrightarrow{\tau, \mathtt{W}(x, v_{\mathtt{W}})}_{\mathsf{SRA}} \mathcal{D}'\end{array}}{\mathcal{D} \xrightarrow{\tau, \mathtt{RMW}(x, v_{\mathtt{R}}, v_{\mathtt{W}})}_{\mathsf{SRA}} \mathcal{D}'}$$

FORK
$$\frac{\begin{array}{c}\mathcal{D}_{\mathrm{new}} = \{\tau_1 \mapsto \mathcal{D}(\tau), \tau_2 \mapsto \mathcal{D}(\tau)\} \\ \mathcal{D}' = \mathcal{D}|_{dom(\mathcal{D}) \setminus \{\tau\}} \uplus \mathcal{D}_{\mathrm{new}}\end{array}}{\mathcal{D} \xrightarrow{\tau, \mathtt{FORK}(\tau_1, \tau_2)}_{\mathsf{SRA}} \mathcal{D}'}$$

JOIN
$$\frac{\begin{array}{c}\mathcal{D}_{\mathrm{new}} = \{\tau \mapsto \mathcal{D}(\tau_1) \cap \mathcal{D}(\tau_2)\} \\ \mathcal{D}' = \mathcal{D}|_{dom(\mathcal{D}) \setminus \{\tau_1, \tau_2\}} \uplus \mathcal{D}_{\mathrm{new}}\end{array}}{\mathcal{D} \xrightarrow{\tau, \mathtt{JOIN}(\tau_1, \tau_2)}_{\mathsf{SRA}} \mathcal{D}'}$$

Fig. 8. Steps of SRA (defining $\delta[x \mapsto \langle v, u, \tau\rangle](y) = \langle v, u, \tau\rangle$ if $y = x$ and $\delta(y)$ else, and $\delta[x \mapsto \mathtt{R}]$ to set all RMW-flags for x to $\mathtt{R}$; both pointwise lifted to lists)

Definition 12. The (overloaded) partial order $\sqsubseteq$ is defined as follows:

1. on potential store lists: $L' \sqsubseteq L$ if L' is a nonempty subsequence of L;
2. on potentials: $D' \sqsubseteq D$ if $\forall L' \in D'.\, \exists L \in D.\, L' \sqsubseteq L$;
3. on potential mappings: $\mathcal{D}' \sqsubseteq \mathcal{D}$ if $\mathcal{D}'(\tau) \sqsubseteq \mathcal{D}(\tau)$ for every $\tau \in dom(\mathcal{D})$.

We also define $L \preceq L'$ if L' is obtained from L by duplication of some stores (e.g., $\delta_1 \cdot \delta_2 \cdot \delta_3 \preceq \delta_1 \cdot \delta_2 \cdot \delta_2 \cdot \delta_3$). This is lifted to potential mappings as expected.

Figure 8 defines the transitions of SRA. The LOSE and DUP steps account for losing and duplication in potentials. Note that these are both internal memory transitions (required to preserve relies as of (mem)). The FORK and JOIN steps distribute potentials on forked threads and join them at the end. The READ step obtains its value from the first store in the lists of the potential of the reader, provided that all these lists agree on that value and the writer thread identifier. RMW steps atomically perform a read and a write step where the read is restricted to an RMW-marked entry.

Most of the complexity is left for the WRITE step. It updates to the new written value for the writer thread τ. For every other thread, it updates a *suffix* (L_1) of the store list with the new value. For guaranteeing causal consistency this updated suffix cannot be arbitrary: it has to be in the potential of the writer thread ($L_1 \in \mathcal{D}(\tau)$). This is the key to achieving the "shared-memory causality principle" of [28], which ensures causal consistency.

Example 3. Consider again the MP program from Fig. 2. After the initial fork step, threads $\mathtt{T}_1$ and $\mathtt{T}_2$ may have the following store list in their potentials:

$$\mathtt{L} = \begin{bmatrix} \mathtt{x} \mapsto \langle 0, \mathtt{RMW}, \mathtt{T}_0 \rangle \\ \mathtt{y} \mapsto \langle 0, \mathtt{RMW}, \mathtt{T}_0 \rangle \end{bmatrix} \cdot \begin{bmatrix} \mathtt{x} \mapsto \langle 0, \mathtt{RMW}, \mathtt{T}_0 \rangle \\ \mathtt{y} \mapsto \langle 0, \mathtt{RMW}, \mathtt{T}_0 \rangle \end{bmatrix} \cdot \begin{bmatrix} \mathtt{x} \mapsto \langle 0, \mathtt{RMW}, \mathtt{T}_0 \rangle \\ \mathtt{y} \mapsto \langle 0, \mathtt{RMW}, \mathtt{T}_0 \rangle \end{bmatrix}.$$

Then, $\mathtt{STORE}(\mathtt{x}, 1)$ by $\mathtt{T}_1$ can generate the following store list for $\mathtt{T}_2$:

$$\mathtt{L}_2 = \begin{bmatrix} \mathtt{x} \mapsto \langle 0, \mathtt{R}, \mathtt{T}_0 \rangle \\ \mathtt{y} \mapsto \langle 0, \mathtt{RMW}, \mathtt{T}_0 \rangle \end{bmatrix} \cdot \begin{bmatrix} \mathtt{x} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \\ \mathtt{y} \mapsto \langle 0, \mathtt{RMW}, \mathtt{T}_0 \rangle \end{bmatrix} \cdot \begin{bmatrix} \mathtt{x} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \\ \mathtt{y} \mapsto \langle 0, \mathtt{RMW}, \mathtt{T}_0 \rangle \end{bmatrix}.$$

Thus $\mathtt{T}_2$ keeps the possibility of reading the "old" value of $\mathtt{x}$. For $\mathtt{T}_1$ this is different: the model allows the writing thread to only see its new value of $\mathtt{x}$ and all entries for $\mathtt{x}$ in the store list are updated. Thus, for $\mathtt{T}_1$ we obtain store list

$$\mathtt{L}_1 = \begin{bmatrix} \mathtt{x} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \\ \mathtt{y} \mapsto \langle 0, \mathtt{RMW}, \mathtt{T}_0 \rangle \end{bmatrix} \cdot \begin{bmatrix} \mathtt{x} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \\ \mathtt{y} \mapsto \langle 0, \mathtt{RMW}, \mathtt{T}_0 \rangle \end{bmatrix} \cdot \begin{bmatrix} \mathtt{x} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \\ \mathtt{y} \mapsto \langle 0, \mathtt{RMW}, \mathtt{T}_0 \rangle \end{bmatrix}.$$

Next, when $\mathtt{T}_1$ executes $\mathtt{STORE}(\mathtt{y}, 1)$, again, the value for $\mathtt{y}$ has to be updated to 1 in $\mathtt{T}_1$ yielding

$$\mathtt{L}_1' = \begin{bmatrix} \mathtt{x} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_0 \rangle \\ \mathtt{y} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \end{bmatrix} \cdot \begin{bmatrix} \mathtt{x} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \\ \mathtt{y} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \end{bmatrix} \cdot \begin{bmatrix} \mathtt{x} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \\ \mathtt{y} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \end{bmatrix}.$$

For $\mathtt{T}_2$ the write step may change $\mathtt{L}_2$ to

$$\mathtt{L}_2' = \begin{bmatrix} \mathtt{x} \mapsto \langle 0, \mathtt{R}, \mathtt{T}_0 \rangle \\ \mathtt{y} \mapsto \langle 0, \mathtt{R}, \mathtt{T}_0 \rangle \end{bmatrix} \cdot \begin{bmatrix} \mathtt{x} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \\ \mathtt{y} \mapsto \langle 0, \mathtt{R}, \mathtt{T}_0 \rangle \end{bmatrix} \cdot \begin{bmatrix} \mathtt{x} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \\ \mathtt{y} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \end{bmatrix}.$$

Thus, thread $\mathtt{T}_2$ can still see the old values, or lose the prefix of its list and see the new values. Importantly, it cannot read 1 from $\mathtt{y}$ *and then* 0 from $\mathtt{x}$. Note that $\mathtt{STORE}(\mathtt{y}, 1)$ by $\mathtt{T}_1$ *cannot* modify $\mathtt{L}_2$ to the list

$$\mathtt{L}_2'' = \begin{bmatrix} \mathtt{x} \mapsto \langle 0, \mathtt{R}, \mathtt{T}_0 \rangle \\ \mathtt{y} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \end{bmatrix} \cdot \begin{bmatrix} \mathtt{x} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \\ \mathtt{y} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \end{bmatrix} \cdot \begin{bmatrix} \mathtt{x} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \\ \mathtt{y} \mapsto \langle 1, \mathtt{RMW}, \mathtt{T}_1 \rangle \end{bmatrix},$$

as it requires $\mathtt{T}_1$ to have $\mathtt{L}_2$ in its *own potential*. This models the intended semantics of message passing under causal consistency.

The next theorem establishes the equivalence of SRA as defined above and opSRA from [28], which is an (operational version of) the standard strong release-acquire declarative semantics [26,31]. (As a corollary, we obtain the equivalence between the potential-based system from [28] and the variant we define in this paper.)

Our notion of equivalence employed in the theorem is *trace equivalence*. We let a trace of a memory system be a sequence of transition labels, ignoring ε transitions, and consider traces of SRA starting from an initial state $\lambda\tau \in \{\mathtt{T}_1, \ldots, \mathtt{T}_N\}. \{\langle \lambda x. \langle 0, \mathtt{RMW}, \mathtt{T}_0 \rangle \rangle\}$ and traces of opSRA starting from the initial execution graph that consists of a write event to every location writing 0 by a distinguished initialization thread $\mathtt{T}_0$.

extended expressions	E	$::= e \mid x \mid \mathtt{R}(x) \mid E+E \mid \neg E \mid E \wedge E \mid \ldots$
interval assertions	I	$::= [E] \mid I\,;I \mid I \wedge I \mid I \vee I$
assertions	φ, ψ	$::= \tau \ltimes I \mid e \mid \varphi \wedge \varphi \mid \varphi \vee \varphi$

Fig. 9. Assertions of Piccolo

Theorem 2. *A trace is generated by* SRA *iff it is generated by opSRA.*

The proof is of this theorem is by simulation arguments (forward simulation in one direction and backward for the converse). It is mechanized in Coq [29]. The mechanized proof does not consider fork and join steps, but they can be straightforwardly added.

6 Program Logic

For the instantiation of our RG framework to SRA, we next (1) introduce the assertions of the logic Piccolo and (2) specify memory triples for Piccolo. Our logic is inspired by *interval logics* like Moszkowski's ITL [35] or duration calculus [13].

Syntax and semantics. Figure 9 gives the grammar of Piccolo. We base it on *extended expressions* which—besides registers—can also involve locations as well as expressions of the form $\mathtt{R}(x)$ (to indicate RMW-flag $\mathtt{R}$). Extended expressions E can hold on entire *intervals* of a store list (denoted $[E]$). Store lists can be split into intervals satisfying different interval expressions ($I_1\,;\ldots;I_n$) using the ";" operator (called "chop"). In turn, $\tau \ltimes I$ means that all store lists in τ's potential satisfy I. For an assertion φ, we let $fv(\varphi) \subseteq \mathsf{Reg} \cup \mathsf{Loc} \cup \mathsf{Tid}$ be the set of registers, locations and thread identifiers occurring in φ, and write $\mathtt{R}(x) \in \varphi$ to indicate that the term $\mathtt{R}(x)$ occurs in φ.

As an example consider again MP (Fig. 2). We would like to express that $\mathtt{T}_2$ upon seeing $\mathtt{y}$ to be 1 cannot see the old value 0 of $\mathtt{x}$ anymore. In Piccolo this is expressed as $\mathtt{T}_2 \ltimes [\mathtt{y} \neq 1]\,;[\mathtt{x} = 1]$: the store lists of $\mathtt{T}_2$ can be split into two intervals (one possibly empty), the first satisfying $\mathtt{y} \neq 1$ and the second $\mathtt{x} = 1$.

Formally, an assertion φ describes register stores coupled with SRA states:

Definition 13. Let γ be a register store, δ a potential store, L a store list, and $\mathcal{D}$ a potential mapping. We let $[\![e]\!]_{\langle\gamma,\delta\rangle} = \gamma(e)$, $[\![x]\!]_{\langle\gamma,\delta\rangle} = \delta(x)$, and $[\![\mathtt{R}(x)]\!]_{\langle\gamma,\delta\rangle} =$ $\mathsf{if}\ \mathtt{rmw}(\delta(x)) = \mathtt{R}\ \mathsf{then}\ \mathit{true}\ \mathsf{else}\ \mathit{false}$. The extension of this notation to any extended expression E is standard. The validity of assertions in $\langle\gamma,\mathcal{D}\rangle$, denoted by $\langle\gamma,\mathcal{D}\rangle \models \varphi$, is defined as follows:

1. $\langle\gamma,L\rangle \models [E]$ if $[\![E]\!]_{\langle\gamma,\delta\rangle} = \mathit{true}$ for every $\delta \in L$.
2. $\langle\gamma,L\rangle \models I_1\,;I_2$ if $\langle\gamma,L_1\rangle \models I_1$ and $\langle\gamma,L_2\rangle \models I_2$ for some (possibly empty) L_1 and L_2 such that $L = L_1 \cdot L_2$.
3. $\langle\gamma,L\rangle \models I_1 \wedge I_2$ if $\langle\gamma,L\rangle \models I_1$ and $\langle\gamma,L\rangle \models I_2$ (similarly for $\vee$).
4. $\langle\gamma,\mathcal{D}\rangle \models \tau \ltimes I$ if $\langle\gamma,L\rangle \models I$ for every $L \in \mathcal{D}(\tau)$.
5. $\langle\gamma,\mathcal{D}\rangle \models e$ if $\gamma(e) = \mathit{true}$.

Assumption	Pre	Command	Post	Reference
	$\{\varphi(r := e)\}$	$\tau \mapsto r := e$	$\{\varphi\}$	Subst-asgn
$x \notin fv(\varphi)$	$\{\varphi\}$	$\tau \mapsto \texttt{WRITE}(x, e)$	$\{\varphi\}$	Stable-wr
$r \notin fv(\varphi)$	$\{\varphi\}$	$\tau \mapsto r := \texttt{LOAD}(x)$	$\{\varphi\}$	Stable-ld
$\tau \notin fv(\varphi)$	$\{\varphi\}$	$\tau \mapsto \texttt{FORK}(\tau_1, \tau_2)$	$\{\varphi\}$	Stable-fork
$\tau \notin fv(\varphi)$	$\{\varphi\}$	$\tau \mapsto \texttt{JOIN}(\tau_1, \tau_2)$	$\{\varphi\}$	Stable-join
	$\{e \wedge \tau \ltimes I\}$	$\tau \mapsto \texttt{FORK}(\tau_1, \tau_2)$	$\{e \wedge \tau_1 \ltimes I \wedge \tau_2 \ltimes I\}$	Fork
	$\{e \wedge \tau_1 \ltimes I \wedge \tau_2 \ltimes I\}$	$\tau \mapsto \texttt{JOIN}(\tau_1, \tau_2)$	$\{e \wedge \tau \ltimes I\}$	Join
	$\{\mathit{True}\}$	$\tau \mapsto \texttt{WRITE}(x, e)$	$\{\tau \ltimes [x = e]\}$	Wr-own
$\texttt{R}(x) \notin I$	$\{\pi \ltimes I\}$	$\tau \mapsto \texttt{WRITE}(x, e)$	$\{\pi \ltimes (I \wedge [\texttt{R}(x)]) \,;\, [x = e]\}$	Wr-other-1
$x \notin fv(I_\tau)$, $\texttt{R}(x) \notin I$	$\{\tau \ltimes I_\tau \wedge \pi \ltimes I \,;\, I_\tau\}$	$\tau \mapsto \texttt{WRITE}(x, e)$	$\{\pi \ltimes I \,;\, I_\tau\}$	Wr-other-2
$x \notin fv(I_\tau)$	$\{\tau \ltimes I_\tau\}$	$\tau \mapsto \texttt{WRITE}(x, e)$	$\{\pi \ltimes [\texttt{R}(x)] \,;\, I_\tau\}$	Wr-other-3
$x \notin fv(I)$	$\{\tau \ltimes [\texttt{R}(x)] \,;\, I\}$	$\tau \mapsto \texttt{SWAP}(x, e)$	$\{\tau \ltimes I\}$	Swap-skip

Fig. 10. Memory triples for Piccolo using $\texttt{WRITE} \in \{\texttt{SWAP}, \texttt{STORE}\}$ and assuming $\tau \neq \pi$

6. $\langle \gamma, \mathcal{D} \rangle \models \varphi_1 \wedge \varphi_2$ if $\langle \gamma, \mathcal{D} \rangle \models \varphi_1$ and $\langle \gamma, \mathcal{D} \rangle \models \varphi_2$ (similarly for $\vee$).

Note that with $\wedge$ and $\vee$ as well as negation on expressions,[4] the logic provides the operators on sets of states necessary for an instantiation of our RG framework. Further, the requirements from SRA states guarantee certain properties:

- For $\varphi_1 = \tau \ltimes [E^\tau_1] \,;\, ... \,;\, [E^\tau_n]$ and $\varphi_2 = \pi \ltimes [E^\pi_1] \,;\, ... \,;\, [E^\pi_m]$: if $E^\tau_i \wedge E^\pi_j \Rightarrow \mathit{False}$ for all $1 \leq i \leq n$ and $1 \leq j \leq m$, then $\varphi_1 \wedge \varphi_2 \Rightarrow \mathit{False}$ (follows from the fact that all lists in potentials are non-empty and agree on the last store).
- If $\langle \gamma, \mathcal{D} \rangle \models \tau \ltimes [\texttt{R}(x)] \,;\, [E]$, then every list $L \in \mathcal{D}(\tau)$ contains a non-empty suffix satisfying E (since all lists have to end with RMW-flags set on).

All assertions are preserved by steps LOSE and DUP. This stability is required by our RG framework (condition (mem))[5]. Stability is achieved here because negations occur on the level of (simple) expressions only (e.g., we cannot have $\neg(\tau \ltimes [x = v])$, meaning that τ must have a store in its potential whose value for x is not v, which would not be stable under LOSE).

Proposition 1. *If $\langle \gamma, \mathcal{D} \rangle \models \varphi$ and $\mathcal{D} \xrightarrow{\varepsilon}_{\mathsf{SRA}} \mathcal{D}'$, then $\langle \gamma, \mathcal{D}' \rangle \models \varphi$.*

Memory triples. Assertions in Piccolo describe sets of states, thus can be used to formulate memory triples. Figure 10 gives the base triples for the different primitive instructions.

We see the standard SC rule of assignment (Subst-asgn) for registers followed by a number of stability rules detailing when assertions are not affected by instructions. Axioms Fork and Join describe the transfer of properties from forking thread to forked threads and back.

[4] Negation just occurs on the level of simple expressions e which is sufficient for calculating $P \setminus [\![e]\!]$ required in rules IF and WHILE.

[5] Such stability requirements are also common to other reasoning techniques for weak memory models, e.g., [19].

The next four axioms in the table concern write instructions (either SWAP or STORE). They reflect the semantics of writing in SRA: (1) In the writer thread τ all stores in all lists get updated (axiom Wr-own). Other threads π will have (2) their lists being split into "old" values for x with R flag and the new value for x (Wr-other-1), (3) properties (expressed as I_τ) of suffixes of lists being preserved when the writing thread satisfies the same properties (Wr-other-2) and (4) their lists consisting of R-accesses to x followed by properties of the writer (Wr-other-3). The last axiom concerns SWAP only: as it can only read from store entries marked as RMW it discards intervals satisfying $[\mathtt{R}(x)]$.

Example 4. We employ the axioms for showing one proof step for MP, namely one pair in the non-interference check of the rely $\mathcal{R}_2$ of $\mathtt{T}_2$ with respect to the guarantees $\mathcal{G}_1$ of $\mathtt{T}_1$:

$$\{\mathtt{T}_2 \ltimes [\mathtt{y} \neq 1] \mathbin{;} [\mathtt{x} = 1] \wedge \mathtt{T}_1 \ltimes [\mathtt{x} = 1]\}\ \mathtt{T}_1 \mapsto \mathtt{STORE}(\mathtt{x}, 1)\ \{\mathtt{T}_2 \ltimes [\mathtt{y} \neq 1] \mathbin{;} [\mathtt{x} = 1]\}$$

By taking I_τ to be $[\mathtt{x} = 1]$, this is an instance of Wr-other-2.

In addition to the axioms above, we use a *shift* rule for load instructions:

$$\textsc{Ld-shift}\ \frac{\{\tau \ltimes I\}\ \tau \mapsto r := \mathtt{LOAD}(x)\ \{\psi\}}{\{\tau \ltimes [(e \wedge E)(r := x)] \mathbin{;} I\}\ \tau \mapsto r := \mathtt{LOAD}(x)\ \{(e \wedge \tau \ltimes [E]; I) \vee \psi\}}$$

A load instruction reads from the first store in the lists, however, if the list satisfying $[(e \wedge E)(r := x)]$ in $[(e \wedge E)(r := x)] \mathbin{;} I$ is empty, it reads from a list satisfying I. The shift rule for LOAD puts this shifting to next stores into a proof rule. Like the standard Hoare rule Subst-asgn, Ld-shift employs backward substitution.

Example 5. We exemplify rule Ld-shift on another proof step of example MP, one for local correctness of $\mathtt{T}_2$:

$$\{\mathtt{T}_2 \ltimes [\mathtt{y} \neq 1] \mathbin{;} [\mathtt{x} = 1]\}\ \mathtt{T}_2 \mapsto \mathtt{a} := \mathtt{LOAD}(\mathtt{y})\ \{\mathtt{a} = 1 \Rightarrow \mathtt{T}_2 \ltimes [\mathtt{x} = 1]\}$$

From axiom Stable-ld we get $\{\mathtt{T}_2 \ltimes [\mathtt{x} = 1]\}\ \mathtt{T}_2 \mapsto \mathtt{a} := \mathtt{LOAD}(\mathtt{y})\ \{\mathtt{T}_2 \ltimes [\mathtt{x} = 1]\}$. We obtain $\{\mathtt{T}_2 \ltimes [\mathtt{y} \neq 1] \mathbin{;} [\mathtt{x} = 1]\}\ \mathtt{T}_2 \mapsto \mathtt{a} := \mathtt{LOAD}(\mathtt{y})\ \{\mathtt{a} \neq 1 \vee \mathtt{T}_2 \ltimes [\mathtt{x} = 1]\}$ using the former as premise for Ld-shift.

In addition, we include the standard conjunction, disjunction and consequence rules of Hoare logic. For instrumented primitive commands we employ the following rule:

$$\textsc{Instr}\ \frac{\{\psi_0\}\ \tau \mapsto c\ \{\psi_1\} \qquad \{\psi_1\}\ \tau \mapsto r_1 := e_1\ \{\psi_2\} \ldots \{\psi_{n-1}\}\ \tau \mapsto r_n := e_n\ \{\psi_n\}}{\{\psi_0\}\ \tau \mapsto \langle c, \langle r_1, \ldots, r_n\rangle := \langle e_1, \ldots, e_n\rangle\rangle\ \{\psi_n\}}$$

Finally, it can be shown that all triples derivable from axioms and rules are valid memory triples.

Lemma 1. *If a* Piccolo *memory triple is derivable,* $\vdash_{\mathsf{Piccolo}} \{\varphi\}\ \tau \mapsto \alpha\ \{\psi\}$*, then* $\mathsf{SRA} \vDash \{\{\langle\gamma, \mathcal{D}\rangle \mid \langle\gamma, \mathcal{D}\rangle \models \varphi\}\}\ \tau \mapsto \alpha\ \{\{\langle\gamma, \mathcal{D}\rangle \mid \langle\gamma, \mathcal{D}\rangle \models \psi\}\}$.

$\{\mathtt{T}_0 \ltimes [\mathtt{x} \neq 2]\}$

Thread $\mathtt{T}_1$	**Thread** $\mathtt{T}_2$
$\{\mathtt{T}_1 \ltimes I^{\mathtt{x}}_{0}\}$	$\{\mathtt{T}_2 \ltimes I^{\mathtt{x}}_{012}\}$
$1 : \mathtt{STORE}(\mathtt{x}, 1);$	$3 : \mathtt{a} := \mathtt{LOAD}(\mathtt{x});$
$\{\mathtt{T}_1 \ltimes I^{\mathtt{x}}_{1}\}$	$\{\mathtt{a} = 2 \Rightarrow \mathtt{T}_2 \ltimes I^{\mathtt{x}}_{2}\}$
$2 : \mathtt{STORE}(\mathtt{x}, 2)$	$4 : \mathtt{b} := \mathtt{LOAD}(\mathtt{x})$
$\{\mathit{True}\}$	$\{\mathtt{a} = 2 \Rightarrow \mathtt{b} = 2\}$

$\{\mathtt{a} = 2 \Rightarrow \mathtt{b} \neq 1\}$

Fig. 11. RRC for two threads (a.k.a. CoRR0)

$\{\mathtt{T}_0 \ltimes I^{\mathtt{x}}_{0}\}$

Thread $\mathtt{T}_1$	**Thread** $\mathtt{T}_2$	**Thread** $\mathtt{T}_3$	**Thread** $\mathtt{T}_4$
$\{\mathtt{T}_1 \ltimes I^{\mathtt{x}}_{02}\}$	$\{\mathtt{T}_1 \ltimes I^{\mathtt{x}}_{01}\}$	$\{\mathtt{T}_3 \ltimes (I^{\mathtt{x}}_{012} \vee I^{\mathtt{x}}_{021})\}$	$\{\mathtt{T}_4 \ltimes (I^{\mathtt{x}}_{012} \vee I^{\mathtt{x}}_{021})\}$
$1 : \mathtt{STORE}(\mathtt{x}, 1)$	$2 : \mathtt{STORE}(\mathtt{x}, 2)$	$3 : \mathtt{a} := \mathtt{LOAD}(\mathtt{x});$	$5 : \mathtt{c} := \mathtt{LOAD}(\mathtt{x});$
$\{\mathit{True}\}$	$\{\mathit{True}\}$	$\{\mathtt{a} = 2 \Rightarrow \mathtt{T}_3 \ltimes I^{\mathtt{x}}_{21}\}$	$\{\mathtt{c} = 1 \Rightarrow \mathtt{T}_4 \ltimes I^{\mathtt{x}}_{12}\}$
		$4 : \mathtt{b} := \mathtt{LOAD}(\mathtt{x})$	$6 : \mathtt{d} := \mathtt{LOAD}(\mathtt{x})$
		$\{\langle \mathtt{a}, \mathtt{b}\rangle = \langle 2, 1\rangle \Rightarrow \mathtt{T}_3 \ltimes I^{\mathtt{x}}_{1}\}$	$\{\langle \mathtt{c}, \mathtt{d}\rangle = \langle 1, 2\rangle \Rightarrow \mathtt{T}_4 \ltimes I^{\mathtt{x}}_{2}\}$

$\{\langle \mathtt{a}, \mathtt{b}\rangle = \langle 2, 1\rangle \Rightarrow \langle \mathtt{c}, \mathtt{d}\rangle \neq \langle 1, 2\rangle\}$

Fig. 12. RRC for four threads (a.k.a. CoRR2)

7 Examples

We discuss examples verified in Piccolo. Additional examples can be found in the extended version of this paper [30].

Coherence. We provide two coherence examples in Figures 11 and 12, using the notation $I^x_{v_1 v_2 \ldots v_n} = [x = v_1] ; [x = v_2] ; \ldots ; [x = v_n]$. Fig. 11 enforces an ordering on writes to the shared location $\mathtt{x}$ on thread $\mathtt{T}_1$. The postcondition guarantees that after reading the second write, thread $\mathtt{T}_2$ cannot read from the first. Fig. 12 is similar, but the writes to $\mathtt{x}$ occur on two different threads. The postcondition of the program guarantees that the two different threads agree on the order of the writes. In particular if one reading thread (here $\mathtt{T}_3$) sees the value 2 then 1, it is impossible for the other reading thread (here $\mathtt{T}_4$) to see 1 then 2.

Potential assertions provide a compact and intuitive mechanism for reasoning, e.g., in Fig. 11, the precondition of line 3 precisely expresses the order of values available to thread $\mathtt{T}_2$. This presents an improvement over view-based assertions [16], which required a separate set of assertions to encode write order.

Peterson's algorithm. Figure 13 shows Peterson's algorithm for implementing mutual exclusion for two threads [38] together with Piccolo assertions. We depict only the code of thread $\mathtt{T}_1$. Thread $\mathtt{T}_2$ is symmetric. A third thread $\mathtt{T}_3$ is assumed stopping the other two threads at an arbitrary point in time. We use $\mathtt{do}\ C\ \mathtt{until}\ e$ as a shorthand for $C\ ;\ \mathtt{while}\ e\ \mathtt{do}\ C$. For correctness under

Thread $\mathtt{T}_1$

$\{\neg a_1 \wedge \neg a_2 \wedge mx_1 = 0\}$
while $\neg$stop do $\{\neg a_1 \wedge (\neg a_2 \vee T_1 \ltimes [R(turn)] ; [flag_2])\}$
1 : $\mathtt{STORE}(flag_1, \mathit{true})$; $\{\neg a_1 \wedge T_1 \ltimes [flag_1] \wedge (\neg a_2 \vee T_1 \ltimes [R(turn)] ; [flag_2])\}$
2 : $\langle \mathtt{SWAP}(turn, 2); a_1 := \mathit{true}\rangle$;
3 : do $\{a_1 \wedge (\neg a_2 \vee T_1 \ltimes [flag_2 \wedge turn \neq 1] \vee P)\}$
4 : $fl_1 := \mathtt{LOAD}(flag_2)$; $\{a_1 \wedge (\neg a_2 \vee (fl_1 \wedge T_1 \ltimes [flag_2 \wedge turn \neq 1]) \vee P)\}$
5 : $tu_1 := \mathtt{LOAD}(turn)$; $\{a_1 \wedge (\neg a_2 \vee (fl_1 \wedge tu_1 \neq 1 \wedge T_1 \ltimes [flag_2 \wedge turn \neq 1]) \vee P)\}$
6 : until $\neg fl_1 \vee (tu_1 = 1)$; $\{a_1 \wedge (\neg a_2 \vee P)\}$
7 : $\mathtt{STORE}(cs, \bot)$; $\{a_1 \wedge (\neg a_2 \vee P)\}$
8 : $\mathtt{STORE}(cs, 0)$; $\{T_1 \ltimes [cs = 0] \wedge a_1 \wedge (\neg a_2 \vee P)\}$
9 : $mx_1 := \mathtt{LOAD}(cs)$; $\{mx_1 = 0 \wedge a_1 \wedge (\neg a_2 \vee P)\}$
10 : $\langle \mathtt{STORE}(flag_1, 0); a_1 := \mathit{false}\rangle$
$\{mx_1 = 0\}$

Fig. 13. Peterson's algorithm, where $P = T_1 \ltimes [R(turn)] ; [flag_2 \wedge turn = 1]$. Thread $\mathtt{T}_2$ is symmetric and we assume a stopper thread $\mathtt{T}_3$ that sets stop to *true*.

SRA, all accesses to the shared variable turn are via a SWAP, which ensures that turn behaves like an SC variable.

Correctness is encoded via registers mx_1 and mx_2 into which the contents of shared variable cs is loaded. Mutual exclusion should guarantee both registers to be 0. Thus neither threads should ever be able to read cs to be $\bot$ (as stored in line 7). The proof (like the associated SC proof in [9]) introduces auxiliary variables a_1 and a_2. Variable a_i is initially *false*, set to *true* when a thread $\mathtt{T}_i$ has performed its swap, and back to *false* when $\mathtt{T}_i$ completes.

Once again potentials provide convenient mechanisms for reasoning about the interactions between the two threads. For example, the assertion $T_1 \ltimes [R(turn)] ; [flag_2]$ in the precondition of line 2 encapsulates the idea that an RMW on turn (via $\mathtt{SWAP}(turn, 2)$) must read from a state in which $flag_2$ holds, allowing us to establish $T_1 \ltimes [flag_2]$ as a postcondition (using the axiom Swap-skip). We obtain disjunct $T_1 \ltimes [flag_2 \wedge turn \neq 1]$ after additionally applying Wr-own.

8 Discussion, Related and Future Work

Previous RG-like logics provided ad-hoc solutions for other concrete memory models such as x86-TSO and C/C++11 [11,16,17,32,39,40,47]. These approaches established soundness of the proposed logic with an ad-hoc proof that couples together memory and thread transitions. We believe that these logics can be formulated in our proposed general RG framework (which will require extensions to other memory operations such as fences).

Moreover, Owicki-Gries logics for different fragments of the C11 memory model [16, 17, 47] used specialized assertions over the underlying view-based semantics. These include *conditional-view assertion* (enabling reasoning about MP), and *value-order* (enabling reasoning about coherence). Both types of assertions are special cases of the potential-based assertions of Piccolo.

Ridge [40] presents an RG reasoning technique tailored to x86-TSO, treating the write buffers in TSO architectures as threads whose steps have to preserve relies. This is similar to our notion of stability of relies under internal memory transitions. Ridge moreover allows to have memory-model specific assertions (e.g., on the contents of write buffers).

The OGRA logic [32] for Release-Acquire (which is slightly weaker form of causal consistency compared to SRA studied in this paper) takes a different approach, which cannot be directly handled in our framework. It employs simple SC-like assertions at the price of having a non-standard non-interference condition which require a stronger form of stability.

Coughlin et al. [14, 15] provide an RG reasoning technique for weak memory models with a semantics defined in terms of *reordering relations* (on instructions). They study both multicopy and non-multicopy atomic architectures, but in all models, the rely-guarantee assertions are interpreted over SC.

Schellhorn et al. [41] develop a framework that extends ITL with a compositional interleaving operator, enabling proof decomposition using RG rules. Each interval represents a sequence of states, strictly alternating between program and environment actions (which may be a skip action). This work is radically different from ours since (1) their states are interpreted using a standard SC semantics, and (2) their intervals represent an *entire execution* of a command as well the interference from the environment while executing that command.

Under SC, rely-guarantee was combined with separation logic [44, 46], which allows the powerful synergy of reasoning using stable invariants (as in rely-guarantee) and ownership transfer (as in concurrent separation logic). It is interesting to study a combination of our RG framework with concurrent separation logics for weak memory models, such as [43, 45].

Other works have studied the decidability of verification for causal consistency models. In work preceding the potential-based SRA model [28], Abdulla et al. [1] show that verification under RA is undecidable. In other work, Abdulla et al. [3] show that the reachability problem under TSO remains decidable for systems with dynamic thread creation. Investigating this question under SRA is an interesting topic for future work.

Finally, the spirit of our generic approach is similar to Iris [22], Views [18], Ogre and Pythia [7], the work of Ponce de León et al. [34], and recent axiomatic characterizations of weak memory reasoning [19], which all aim to provide a *generic* framework that can be instantiated to underlying semantics.

In the future we are interested in automating the reasoning in Piccolo, starting from automatically checking for validity of program derivations (using, e.g., SMT solvers for specialised theories of sequences or strings [24, 42]), and, including, more ambitiously, synthesizing appropriate Piccolo invariants.

References

1. Abdulla, P.A., Arora, J., Atig, M.F., Krishna, S.N.: Verification of programs under the release-acquire semantics. In: PLDI. pp. 1117–1132. ACM (2019), https://doi.org/10.1145/3314221.3314649

2. Abdulla, P.A., Atig, M.F., Bouajjani, A., Kumar, K.N., Saivasan, P.: Deciding reachability under persistent x86-tso. Proc. ACM Program. Lang. **5**(POPL), 1–32 (2021), https://doi.org/10.1145/3434337
3. Abdulla, P.A., Atig, M.F., Bouajjani, A., Kumar, K.N., Saivasan, P.: Verifying reachability for TSO programs with dynamic thread creation. In: NETYS. LNCS, vol. 13464, pp. 283–300. Springer (2022), https://doi.org/10.1007/978-3-031-17436-0_19
4. Abdulla, P.A., Atig, M.F., Bouajjani, A., Ngo, T.P.: The benefits of duality in verifying concurrent programs under TSO. In: CONCUR. LIPIcs, vol. 59, pp. 5:1–5:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2016), https://doi.org/10.4230/LIPIcs.CONCUR.2016.5
5. Abdulla, P.A., Atig, M.F., Bouajjani, A., Ngo, T.P.: A load-buffer semantics for total store ordering. Log. Methods Comput. Sci. **14**(1) (2018), https://doi.org/10.23638/LMCS-14(1:9)2018
6. Ahamad, M., Neiger, G., Burns, J.E., Kohli, P., Hutto, P.W.: Causal memory: Definitions, implementation, and programming. Distributed Comput. **9**(1), 37–49 (1995), https://doi.org/10.1007/BF01784241
7. Alglave, J., Cousot, P.: Ogre and Pythia: an invariance proof method for weak consistency models. In: Castagna, G., Gordon, A.D. (eds.) POPL. pp. 3–18. ACM (2017), https://doi.org/10.1145/3009837.3009883
8. Alglave, J., Maranget, L., Tautschnig, M.: Herding cats: Modelling, simulation, testing, and data mining for weak memory. ACM Trans. Program. Lang. Syst. **36**(2), 7:1–7:74 (2014), https://doi.org/10.1145/2627752
9. Apt, K.R., de Boer, F.S., Olderog, E.: Verification of Sequential and Concurrent Programs. Texts in Computer Science, Springer (2009), https://doi.org/10.1007/978-1-84882-745-5
10. Beillahi, S.M., Bouajjani, A., Enea, C.: Robustness against transactional causal consistency. Log. Methods Comput. Sci. **17**(1) (2021), https://lmcs.episciences.org/7149
11. Bila, E.V., Dongol, B., Lahav, O., Raad, A., Wickerson, J.: View-Based Owicki-Gries Reasoning for Persistent x86-TSO. In: ESOP. LNCS, vol. 13240, pp. 234–261. Springer (2022), https://doi.org/10.1007/978-3-030-99336-8_9
12. Bouajjani, A., Enea, C., Guerraoui, R., Hamza, J.: On verifying causal consistency. In: POPL. pp. 626–638. ACM (2017), https://doi.org/10.1145/3009837.3009888
13. Chaochen, Z., Hoare, C.A.R., Ravn, A.P.: A calculus of durations. Inf. Process. Lett. **40**(5), 269–276 (1991), https://doi.org/10.1016/0020-0190(91)90122-X
14. Coughlin, N., Winter, K., Smith, G.: Rely/guarantee reasoning for multicopy atomic weak memory models. In: FM. LNCS, vol. 13047, pp. 292–310. Springer (2021), https://doi.org/10.1007/978-3-030-90870-6_16
15. Coughlin, N., Winter, K., Smith, G.: Compositional reasoning for non-multicopy atomic architectures. Form. Asp. Comput. (dec 2022), https://doi.org/10.1145/3574137
16. Dalvandi, S., Doherty, S., Dongol, B., Wehrheim, H.: Owicki-Gries Reasoning for C11 RAR. In: ECOOP. LIPIcs, vol. 166, pp. 11:1–11:26. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020), https://doi.org/10.4230/LIPIcs.ECOOP.2020.11
17. Dalvandi, S., Dongol, B., Doherty, S., Wehrheim, H.: Integrating Owicki-Gries for C11-Style Memory Models into Isabelle/HOL. J. Autom. Reason. **66**(1), 141–171 (2022), https://doi.org/10.1007/s10817-021-09610-2

18. Dinsdale-Young, T., Birkedal, L., Gardner, P., Parkinson, M.J., Yang, H.: Views: compositional reasoning for concurrent programs. In: POPL. pp. 287–300. ACM (2013), https://doi.org/10.1145/2429069.2429104
19. Doherty, S., Dalvandi, S., Dongol, B., Wehrheim, H.: Unifying Operational Weak Memory Verification: An Axiomatic Approach. ACM Trans. Comput. Log. **23**(4), 27:1–27:39 (2022), https://doi.org/10.1145/3545117
20. Doherty, S., Dongol, B., Wehrheim, H., Derrick, J.: Verifying C11 programs operationally. In: PPoPP. pp. 355–365. ACM (2019), https://doi.org/10.1145/3293883.3295702
21. Jones, C.B.: Tentative steps toward a development method for interfering programs. ACM Trans. Program. Lang. Syst. **5**(4), 596–619 (1983), https://doi.org/10.1145/69575.69577
22. Jung, R., Krebbers, R., Jourdan, J., Bizjak, A., Birkedal, L., Dreyer, D.: Iris from the ground up: A modular foundation for higher-order concurrent separation logic. J. Funct. Program. **28**, e20 (2018), https://doi.org/10.1017/S0956796818000151
23. Kaiser, J., Dang, H., Dreyer, D., Lahav, O., Vafeiadis, V.: Strong logic for weak memory: Reasoning about release-acquire consistency in iris. In: ECOOP. LIPIcs, vol. 74, pp. 17:1–17:29. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2017), https://doi.org/10.4230/LIPIcs.ECOOP.2017.17
24. Kan, S., Lin, A.W., Rümmer, P., Schrader, M.: CertiStr: a certified string solver. In: CPP. pp. 210–224. ACM (2022), https://doi.org/10.1145/3497775.3503691
25. Kang, J., Hur, C., Lahav, O., Vafeiadis, V., Dreyer, D.: A promising semantics for relaxed-memory concurrency. In: POPL. pp. 175–189. ACM (2017), https://doi.org/10.1145/3009837.3009850
26. Lahav, O.: Verification under causally consistent shared memory. ACM SIGLOG News **6**(2), 43–56 (2019), https://doi.org/10.1145/3326938.3326942
27. Lahav, O., Boker, U.: Decidable verification under a causally consistent shared memory. In: PLDI. pp. 211–226. ACM (2020), https://doi.org/10.1145/3385412.3385966
28. Lahav, O., Boker, U.: What's Decidable About Causally Consistent Shared Memory? ACM Trans. Program. Lang. Syst. **44**(2), 8:1–8:55 (2022), https://doi.org/10.1145/3505273
29. Lahav, O., Dongol, B., Wehrheim, H.: Artifact: Rely-guarantee reasoning for causally consistent shared memory. Zenodo (2023), https://doi.org/10.5281/zenodo.7929646
30. Lahav, O., Dongol, B., Wehrheim, H.: Rely-guarantee reasoning for causally consistent shared memory (extended version) (2023), https://doi.org/10.48550/arXiv.2305.08486
31. Lahav, O., Giannarakis, N., Vafeiadis, V.: Taming release-acquire consistency. In: POPL. pp. 649–662. ACM (2016), https://doi.org/10.1145/2837614.2837643
32. Lahav, O., Vafeiadis, V.: Owicki-Gries Reasoning for Weak Memory Models. In: ICALP. LNCS, vol. 9135, pp. 311–323. Springer (2015), https://doi.org/10.1007/978-3-662-47666-6_25
33. Lamport, L.: How to make a multiprocessor computer that correctly executes multiprocess programs. IEEE Trans. Computers **28**(9), 690–691 (1979), https://doi.org/10.1109/TC.1979.1675439
34. de León, H.P., Furbach, F., Heljanko, K., Meyer, R.: BMC with memory models as modules. In: FMCAD. pp. 1–9. IEEE (2018), https://doi.org/10.23919/FMCAD.2018.8603021

35. Moszkowski, B.C.: A complete axiom system for propositional interval temporal logic with infinite time. Log. Methods Comput. Sci. **8**(3) (2012), https://doi.org/10.2168/LMCS-8(3:10)2012
36. Owens, S., Sarkar, S., Sewell, P.: A better x86 memory model: x86-tso. In: TPHOLs. LNCS, vol. 5674, pp. 391–407. Springer (2009), https://doi.org/10.1007/978-3-642-03359-9_27
37. Owicki, S.S., Gries, D.: An Axiomatic Proof Technique for Parallel Programs I. Acta Informatica **6**, 319–340 (1976), https://doi.org/10.1007/BF00268134
38. Peterson, G.L.: Myths about the mutual exclusion problem. Inf. Process. Lett. **12**(3), 115–116 (1981)
39. Raad, A., Lahav, O., Vafeiadis, V.: Persistent Owicki-Gries reasoning: a program logic for reasoning about persistent programs on Intel-x86. Proc. ACM Program. Lang. **4**(OOPSLA), 151:1–151:28 (2020), https://doi.org/10.1145/3428219
40. Ridge, T.: A Rely-Guarantee Proof System for x86-TSO. In: VSTTE. LNCS, vol. 6217, pp. 55–70. Springer (2010), https://doi.org/10.1007/978-3-642-15057-9_4
41. Schellhorn, G., Tofan, B., Ernst, G., Pfähler, J., Reif, W.: RGITL: A temporal logic framework for compositional reasoning about interleaved programs. Ann. Math. Artif. Intell. **71**(1-3), 131–174 (2014), https://doi.org/10.1007/s10472-013-9389-z
42. Sheng, Y., Nötzli, A., Reynolds, A., Zohar, Y., Dill, D.L., Grieskamp, W., Park, J., Qadeer, S., Barrett, C.W., Tinelli, C.: Reasoning about vectors using an SMT theory of sequences. In: IJCAR. LNCS, vol. 13385, pp. 125–143. Springer (2022), https://doi.org/10.1007/978-3-031-10769-6_9
43. Svendsen, K., Pichon-Pharabod, J., Doko, M., Lahav, O., Vafeiadis, V.: A separation logic for a promising semantics. In: Ahmed, A. (ed.) ESOP. LNCS, vol. 10801, pp. 357–384. Springer (2018), https://doi.org/10.1007/978-3-319-89884-1_13
44. Vafeiadis, V.: Modular fine-grained concurrency verification. Ph.D. thesis, University of Cambridge, UK (2008), https://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.612221
45. Vafeiadis, V., Narayan, C.: Relaxed separation logic: a program logic for C11 concurrency. In: OOPSLA. pp. 867–884. ACM (2013), https://doi.org/10.1145/2509136.2509532
46. Vafeiadis, V., Parkinson, M.J.: A marriage of rely/guarantee and separation logic. In: CONCUR. LNCS, vol. 4703, pp. 256–271. Springer (2007), https://doi.org/10.1007/978-3-540-74407-8_18
47. Wright, D., Batty, M., Dongol, B.: Owicki-Gries Reasoning for C11 Programs with Relaxed Dependencies. In: FM. LNCS, vol. 13047, pp. 237–254. Springer (2021), https://doi.org/10.1007/978-3-030-90870-6_13
48. Xu, Q., de Roever, W.P., He, J.: The Rely-Guarantee Method for Verifying Shared Variable Concurrent Programs. Formal Aspects Comput. **9**(2), 149–174 (1997), https://doi.org/10.1007/BF01211617