



Robustness against Release/Acquire Semantics

Ori Lahav
Tel Aviv University
Israel
orilahav@tau.ac.il

Roy Margalit
Tel Aviv University
Israel
roy.margalit@cs.tau.ac.il

Abstract

We present an algorithm for automatically checking robustness of concurrent programs against C/C++11 release/acquire semantics, namely verifying that all program behaviors under release/acquire are allowed by sequential consistency. Our approach reduces robustness verification to a reachability problem under (instrumented) sequential consistency. We have implemented our algorithm in a prototype tool called *Rocker* and applied it to several challenging concurrent algorithms. To the best of our knowledge, this is the first precise method for verifying robustness against a high-level programming language weak memory semantics.

CCS Concepts • **Theory of computation** → *Verification by model checking; Concurrent algorithms; Program semantics; Program verification; Program analysis*; • **Software and its engineering** → *Software verification*.

Keywords weak memory models, C/C++11, release/acquire, robustness

ACM Reference Format:

Ori Lahav and Roy Margalit. 2019. Robustness against Release/Acquire Semantics. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '19)*, June 22–26, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3314221.3314604>

1 Introduction

Release/acquire (RA), the fragment of the C/C++11 memory model [14] consisting of release stores, acquire loads and acquire-release read-modify-writes (RMWs), is a particularly useful and well-behaved weak memory model [36]. It is weaker than sequential consistency (SC) [40] and allows higher performance implementations. For example, x86-TSO [50] provides RA “for free” (its memory model is stronger than RA), and POWER [45] implements RA using

‘lightweight sync’ instructions rather than more expensive ‘full sync’ instructions which are needed for SC.

At the same time, since RA is designed to support the common “message passing” synchronization idiom, the guarantees provided by RA suffice to implement various fundamental concurrent algorithms and synchronization mechanisms. In fact, many useful programs are actually *robust against RA*—the behaviors they exhibit under RA semantics are also allowed under SC—or can be made robust by placing few *SC-fences* or by strengthening certain reads and writes to be RMW operations. Such modifications are sometimes necessary, with the best known example being Dekker’s mutual exclusion algorithm, whose RA (non-SC) behavior is harmful for its correctness.

A natural question is thus whether one can automatically verify robustness against RA. Our main contribution is a decision procedure for this problem. Besides our theoretical interest, we believe that this result can facilitate the development of concurrent algorithms for RA. In particular, if we are able to verify robustness against RA, various programs designed for SC may be directly ported and verified with more ordinary techniques assuming SC. Further, robustness of non-robust programs may be enforced (by placing *SC-fences* or RMW operations), and verifying the robustness of the strengthened program.

To precisely state our result, it is crucial to carefully define what constitutes a behavior of a concurrent program under SC and under RA, which in turn determines what robustness means. Here, it is natural to use operational presentations of SC and RA as memory subsystems, formulated as labeled transition systems (for RA one could use the timestamp machine introduced in [33]). Then, program behaviors correspond to program states that are reachable when linked with each of the memory subsystems. More precisely, thinking about a concurrent program as a labeled transition system (whose states comprise of the values of the thread-local program counters and variables), one may identify SC (RA) program behaviors with the set of states of the program that are reachable in its runs when synchronized with runs of the SC (RA) memory subsystems. This definition of program behavior leads to what is known as *state robustness*, and corresponds to typical safety properties verification using local assertions and global invariants that relate values of local variables and program counters.

Nevertheless, following [24, Thm. 2.12], it is easy to show that verifying state robustness against RA is as hard as the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PLDI '19, June 22–26, 2019, Phoenix, AZ, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6712-7/19/06.

<https://doi.org/10.1145/3314221.3314604>

general state reachability problem under RA. The latter problem was recently shown to be undecidable [2]. Thus we resort to a more informative definition of a behavior, leading to a stronger notion of robustness. By doing so, we follow works on robustness against hardware models, TSO in particular (e.g., [17, 19]), where state robustness—like state reachability—is non-primitive recursive [11, 12]. For this matter, we use formulations of SC and RA as labeled transition systems whose states are (C/C++11-like) *execution graphs*. Execution graphs keep track of the full partially ordered history of the run (and thus in this presentation both SC and RA are infinite state systems), including the reads-from mapping (mapping each read to the write it read from) and the modification order (a total order on writes to the same location). The difference between SC and RA is then reduced to the transitions they allow. For instance, when adding reads to the execution graph, SC requires that it reads from the write that is maximal in the modification order, while RA places much weaker restrictions. Now, we can identify program behaviors with pairs of states of both the program and the memory subsystem that are reachable in their synchronized runs. We refer to the robustness notion induced by this definition as *execution-graph robustness*.

Our main contribution is a decision procedure that checks whether a given concurrent program is execution-graph robust against RA. To achieve this, we show how this verification problem can be reduced to a state reachability problem under a (finite state) instrumented SC memory. Roughly speaking, this memory keeps track of the relevant parts of the generated execution graph and uses this information for monitoring that RA execution graphs cannot diverge from SC ones. We prove that our approach is sound and precise. In particular, it follows that this verification problem for programs with bounded data domain is PSPACE-complete.

Our approach can be straightforwardly extended to handle C/C++11's *non-atomic accesses*. A data-race on a non-atomic access is considered an undefined behavior, and, thus, robustness of a program should also imply that it has no data-races on non-atomic accesses. Since robust programs have only SC executions, checking for data-races can be done using standard techniques. For completeness, we incorporated these checks in our method simultaneously to the verification of robustness against RA.

We have implemented our method in a prototype tool, called *Rocker*, using Spin [31] as a back-end model checker under SC. We used *Rocker* to verify the robustness of several concurrent algorithms, including Peterson's mutual execution adaptations for RA [57], sequence locks [16] and user-mode read-copy-update (RCU) implementations [26]. In particular, we observe that execution-graph robustness is a useful property, allowing one, in many cases, to think in terms of SC while running on a weaker model.

The rest of this paper is structured as follows. In §2 we formally present the programming language and the notion

of state robustness. In §3 we present the RA concurrency semantics. In §4 we define execution-graph robustness against RA. In §5 we present our decision procedure. In §6 we extend the decision procedure to support non-atomic accesses. In §7 we discuss the implementation and our experiments with it. In §8 we discuss related work. Finally, in §9 we conclude and outline directions for future work. Additional material and proofs for the claims of this paper are available in [1]. The prototype implementation and the examples it was tested on are available in the artifact accompanying this paper.

2 Preliminaries: State Robustness

Given a (binary) relation R , $dom(R)$ and $codom(R)$ denote its domain and codomain, and $R^?$, R^+ , and R^* denote its reflexive, transitive, and reflexive-transitive closures. The inverse of a relation R is denoted by R^{-1} , and the (left) composition of two relations R_1, R_2 is denoted by $R_1 ; R_2$. We denote by $[A]$ the identity relation on a set A . In particular, $[A] ; R ; [B] = R \cap (A \times B)$. For a strict total order R , we write $R|_{imm}$ to denote the set of *immediate R-edges*, i.e., $R|_{imm} = R \setminus (R ; R)$.

2.1 Programming Language

Let Val, Loc, Reg be finite sets of values, (shared) locations, and register names. We assume that Val contains a distinguished value 0, used as the initial value for all locations. Figure 1 presents our toy programming language. Its expressions are constructed from registers (local variables) and values. Instructions include assignments and conditional branching, as well as memory operations. Intuitively speaking, an assignment $r := e$ assigns the value of e to register r (involving no memory access); $\text{if } e \text{ goto } n$ jumps to line n of the program iff the value of e is not 0; a write $x := e$ stores the value of e in x ; a read $r := x$ loads the value of x to register r ; $r := \text{FADD}(x, e)$ atomically increments x by the value of e and loads the old value of x to r ; and $r := \text{CAS}(x, e_R \rightarrow e_W)$ atomically loads the value of x to r , compares it to the value of e_R , and if the two values are equal, replaces the value of x by the value of e_W .

The less standard instructions `wait` and `BCAS` are blocking: `wait(x = e)` blocks the current thread until it manages to load the value of e from x ; and `BCAS(x, eR → eW)` blocks the current thread until it performs a successful CAS of x from the value of e_R (to the value of e_W). These instructions can be easily implemented using loops (e.g., $L : r := x; \text{if } r \neq e \text{ goto } L$ with fresh r for `wait(x = e)`). Nevertheless, as we demonstrate in the end of this section, including them as primitives leads to a more expressive notion of robustness.

In turn, a sequential program $S \in \text{SProg}$ is a finite map from \mathbb{N} to instructions (we assume that $0 \in dom(S)$), and a concurrent program P is a top-level parallel composition of sequential programs, defined as a mapping from a finite set $\text{Tid} \subseteq \mathbb{N}$ of thread identifiers to SProg. In our examples, we often write sequential programs as sequences of instructions

$v \in \text{Val}$	Values	$\text{Exp} \ni e ::= r \mid v \mid e + e \mid e = e \mid e \neq e \mid \dots$	Sequential programs:
$x \in \text{Loc}$	Locations	$\text{Inst} \ni \text{inst} ::= r := e \mid \text{if } e \text{ goto } n \mid x := e \mid r := x \mid$	$S \in \text{SProg} \triangleq \mathbb{N} \xrightarrow{\text{fin}} \text{Inst}$
$r \in \text{Reg}$	Registers	$r := \text{FADD}(x, e) \mid r := \text{CAS}(x, e \rightarrow e)$	Concurrent programs:
$\tau \in \text{Tid} \subseteq \mathbb{N}$	Thread identifiers	$\text{wait}(x = e) \mid \text{BCAS}(x, e \rightarrow e)$	$P : \text{Tid} \rightarrow \text{SProg}$

Figure 1. Domains and programming language syntax.

$\frac{S(pc) = r := e}{\Phi' = \Phi[r \mapsto \Phi(e)]}$	$\frac{S(pc) = \text{if } e \text{ goto } n}{\Phi(e) \neq 0}$	$\frac{S(pc) = \text{if } e \text{ goto } n}{\Phi(e) = 0}$	$\frac{S(pc) = x := e}{l = W(x, \Phi(e))}$	$\frac{S(pc) = r := x}{l = R(x, v) \quad \Phi' = \Phi[r \mapsto v]}$
$\langle pc, \Phi \rangle \xrightarrow{\epsilon} \langle pc + 1, \Phi' \rangle$	$\langle pc, \Phi \rangle \xrightarrow{\epsilon} \langle n, \Phi \rangle$	$\langle pc, \Phi \rangle \xrightarrow{\epsilon} \langle pc + 1, \Phi \rangle$	$\langle pc, \Phi \rangle \xrightarrow{l} \langle pc + 1, \Phi \rangle$	$\langle pc, \Phi \rangle \xrightarrow{l} \langle pc + 1, \Phi' \rangle$
$\frac{S(pc) = r := \text{FADD}(x, e)}{l = \text{RMW}(x, v, v + \Phi(e))}$	$\frac{S(pc) = r := \text{CAS}(x, e_R \rightarrow e_W)}{l = \text{RMW}(x, \Phi(e_R), \Phi(e_W))}$	$\frac{S(pc) = r := \text{CAS}(x, e_R \rightarrow e_W)}{l = R(x, v) \quad v \neq \Phi(e_R)}$	$\frac{S(pc) = \text{wait}(x = e)}{l = R(x, \Phi(e))}$	$\frac{S(pc) = \text{BCAS}(x, e_R \rightarrow e_W)}{l = \text{RMW}(x, \Phi(e_R), \Phi(e_W))}$
$\langle pc, \Phi \rangle \xrightarrow{l} \langle pc + 1, \Phi' \rangle$	$\langle pc, \Phi \rangle \xrightarrow{l} \langle pc + 1, \Phi' \rangle$	$\langle pc, \Phi \rangle \xrightarrow{l} \langle pc + 1, \Phi' \rangle$	$\langle pc, \Phi \rangle \xrightarrow{l} \langle pc + 1, \Phi \rangle$	$\langle pc, \Phi \rangle \xrightarrow{l} \langle pc + 1, \Phi \rangle$

Figure 2. Transitions of LTS induced by a sequential program $S \in \text{SProg}$.

delimited by line breaks, use ‘||’ for parallel composition, and refer to the program threads as τ_1, τ_2, \dots following their left-to-right order in the program listing.

2.2 From Programs to Transition Systems

A *labeled transition system* (LTS) A over an alphabet Σ is a tuple $\langle Q, q_0, \rightarrow \rangle$, where Q is a set of *states*, $q_0 \in Q$ is the *initial state*, and $\rightarrow \subseteq Q \times \Sigma \times Q$ is a set of *transitions*. We write $\xrightarrow{\sigma}$ for the relation $\{\langle q, q' \rangle \mid \langle q, \sigma, q' \rangle \in \rightarrow\}$, and \rightarrow for $\bigcup_{\sigma \in \Sigma} \xrightarrow{\sigma}$. We denote by $A.Q$, $A.q_0$ and \rightarrow_A the three components of an LTS A . A state $q \in A.Q$ is called *reachable* in A if $A.q_0 \xrightarrow_A^* q$. A symbol $\sigma \in \Sigma$ is *enabled* in q (alternatively, q *enables* σ) if $q \xrightarrow_A^\sigma q'$ for some q' . A sequence $\sigma_1, \dots, \sigma_n$ is a *trace* of A if $A.q_0 \xrightarrow_A^{\sigma_1} \dots \xrightarrow_A^{\sigma_n} q$ for some q .

Definition 2.1. A *label* $l \in \text{Lab}$ is either $R(x, v_R)$ (read label), $W(x, v_W)$ (write label), or $\text{RMW}(x, v_R, v_W)$ (RMW label), where $x \in \text{Loc}$ and $v_R, v_W \in \text{Val}$. The functions typ , loc , val_R , and val_W return (when applicable) the type (R/W/RMW), location, read value, and written value of a given label.

A sequential program $S \in \text{SProg}$ induces an LTS over $\text{Lab} \cup \{\epsilon\}$, whose states are pairs $\langle pc, \Phi \rangle$ where $pc \in \mathbb{N}$ (called *program counter*) and $\Phi : \text{Reg} \rightarrow \text{Val}$ (called *store*, and extended to expressions in the obvious way). Its initial state is $\langle 0, \lambda r. 0 \rangle$, and its transitions are given in Fig. 2, following the informal description above of the language constructs. In the sequel we identify sequential programs with their induced LTSs (when writing, e.g., $S.Q$ and \rightarrow_S).

Example 2.2. We present the LTS induced by a simple sequential program S . Let $\text{Val} = \{0, \dots, 4\}$, $\text{Loc} = \{x\}$ and $\text{Reg} = \{r\}$. We use $+$ to denote the possibly overflowing sum (e.g., $2 + 4 = 1$), and evaluate expressions of the form $r < e$

to be 1 if $\Phi(r) < \Phi(e)$ and 0 otherwise.

$0 : r := r + 1$
 $1 : \text{if } r < 2 \text{ goto } 0$
 $2 : x := r$
 $S.Q = \{0, 1, 2, 3\} \times \{[r \mapsto v] \mid v \in \text{Val}\}$
 $S.q_0 = \langle 0, [r \mapsto 0] \rangle$

\rightarrow_S is given by:

$\{\langle 0, [r \mapsto v] \rangle \xrightarrow{\epsilon_S} \langle 1, [r \mapsto v + 1] \rangle \mid v \in \text{Val}\} \cup$
 $\{\langle 1, [r \mapsto v] \rangle \xrightarrow{\epsilon_S} \langle 0, [r \mapsto v] \rangle \mid v < 2\} \cup$
 $\{\langle 1, [r \mapsto v] \rangle \xrightarrow{\epsilon_S} \langle 2, [r \mapsto v] \rangle \mid v \geq 2\} \cup$
 $\{\langle 2, [r \mapsto v] \rangle \xrightarrow{W(x, v)} \langle 3, [r \mapsto v] \rangle \mid v \in \text{Val}\}$

A concurrent program P induces an LTS over the alphabet $\text{Tid} \times (\text{Lab} \cup \{\epsilon\})$. Its states are tuples in $\prod_{\tau \in \text{Tid}} P(\tau).Q$; its initial state is $\lambda \tau. P(\tau).q_0$; and its transitions are interleaved transitions of P 's components, given by:

$$\frac{q_\tau \xrightarrow{l_\epsilon}_{P(\tau)} q'_\tau \quad \forall \pi \neq \tau. q_\pi = q'_\pi}{\lambda \pi. q_\pi \xrightarrow{\langle \tau, l_\epsilon \rangle} \lambda \pi. q'_\pi}$$

In the sequel we identify concurrent programs with their induced LTSs. We often use vector notation (e.g., \bar{q}) to denote states of concurrent programs.

2.3 Concurrent Systems and State Robustness

To give semantics to concurrent programs, we synchronize them with *memory subsystems*, as defined next.

Definition 2.3. A *memory subsystem* is a (possibly infinite) LTS over the alphabet $\mathbb{N} \times \text{Lab}$.

The labels here are pairs in $\mathbb{N} \times \text{Lab}$ representing the thread identifier and the label of the performed operation.¹

¹This formulation suffices for the purposes of this paper. In a broader context, memory subsystems may also employ internal memory actions, such as propagation from local stores to the main memory in TSO. Extending the definitions to a more general notion of robustness is straightforward.

The most well-known memory subsystem is the one of sequential consistency, denoted here by SC. This memory subsystem simply tracks the most recent value written to each location. Formally, it is defined by $SC.Q \triangleq Loc \rightarrow Val$, $SC.q_0 \triangleq \lambda x. 0$, and \rightarrow_{SC} is given by:

$$\frac{M' = M[x \mapsto v_W] \quad l = W(x, v_W)}{M \xrightarrow{\langle \tau, l \rangle}_{SC} M'} \quad \frac{M(x) = v_R \quad l = R(x, v_R)}{M \xrightarrow{\langle \tau, l \rangle}_{SC} M} \quad \frac{M(x) = v_R \quad M' = M[x \mapsto v_W] \quad l = RMW(x, v_R, v_W)}{M \xrightarrow{\langle \tau, l \rangle}_{SC} M'}$$

Note that SC is oblivious to the thread that takes the action (we have $M \xrightarrow{\langle \tau, l \rangle}_{SC} M'$ iff $M \xrightarrow{\langle \pi, l \rangle}_{SC} M'$).

By synchronizing a concurrent program and a memory subsystem, we obtain a *concurrent system* as defined next.

Definition 2.4. A *concurrent system* is a pair, denoted P_M , where P is a concurrent program and M is a memory subsystem. A concurrent system P_M induces an LTS over $Tid \times Lab$ whose states are pairs in $P.Q \times M.Q$; its initial state is $\langle P.q_0, M.q_0 \rangle$; and its transitions are given by:

$$\frac{\bar{q} \xrightarrow{\langle \tau, \epsilon \rangle^*}_P \xrightarrow{\langle \tau, l \rangle}_P \xrightarrow{\langle \tau, \epsilon \rangle^*}_P \bar{q}' \quad q_M \xrightarrow{\langle \tau, l \rangle}_M q'_M}{\langle \bar{q}, q_M \rangle \xrightarrow{\langle \tau, l \rangle}_{P_M} \langle \bar{q}', q'_M \rangle}$$

In the sequel we identify concurrent systems with their induced state machines.

We can now define state robustness against a given memory subsystem. This definition essentially identifies the behaviors of a program P under a memory subsystem M with the first projection of the states that are reachable in P_M .

Definition 2.5. A state \bar{q} of a concurrent program P is *reachable under a memory subsystem M* if $\langle \bar{q}, q_M \rangle$ is reachable in the concurrent system P_M for some $q_M \in M.Q$.

Definition 2.6. A concurrent program P is *state robust against a memory subsystem M* if every reachable state of P under M is also reachable under SC.

We can now demonstrate the reason for including the blocking instructions `wait` and BCAS as primitives. Consider the following implementations of a “global barrier”:

$$\begin{array}{l} 0 : x := 1 \\ 1 : r_1 := y \\ 2 : \text{if } r_1 \neq 1 \\ \quad \text{goto } 1 \end{array} \parallel \begin{array}{l} 0 : y := 1 \\ 1 : r_2 := x \quad x := 1 \\ 2 : \text{if } r_2 \neq 1 \quad \text{wait}(y = 1) \\ \quad \text{goto } 1 \end{array} \parallel \begin{array}{l} y := 1 \\ \text{wait}(x = 1) \end{array} \quad (\text{BAR})$$

While the two programs are functionally equivalent, only the right program may be state robust against memory subsystems M that allow reading of “stale values” (such as RA and TSO). Indeed, the state in which both threads are in their last program line ($pc_1 = pc_2 = 2$) after reading 0 ($\Phi_1(r_1) = \Phi_2(r_2) = 0$) is reachable for the program on the left under such memory subsystem, but clearly not under SC. In many cases, such robustness violations are not harmful for

the safety of the program, as they only imply that under weak memory the program may remain longer waiting in the busy loop.² A corresponding state is not reachable for the program on the right, and thus, using the blocking wait instruction, one may mask such benign robustness violations.

Similar benign robustness violations when using CAS, e.g., in spin loops, can be avoided using the BCAS primitive. Handling blocking instructions is essential to establish robustness of some interesting examples (e.g., RCU), without having more fences than actually necessary for program correctness.

3 Release/Acquire Semantics

In this section, we introduce the RA memory subsystem. RA’s original presentation, as a fragment of C/C++11 [14], is declarative (a.k.a. axiomatic), i.e., it is formulated as a collection of formal consistency constraints that are used to filter candidate execution graphs. In our proofs we use such a presentation (see [1, §A]), but for the current purpose we need to define RA as an LTS. The declarative RA semantics can be easily “operationalized”, as was done, e.g., in [54], so that consistent execution graphs are incrementally constructed. We will need this presentation as well (see §4.2), but since execution-graph semantics is often considered unintuitive, we present here an equivalent operational model, due to [33], which is perhaps more natural as an operational semantics for readers unfamiliar with the declarative style.

The memory in the RA operational model is a set of timestamped messages, which record all previously executed writes. *Timestamps* are taken to be natural numbers, $\text{Time} \triangleq \mathbb{N}$. A timestamp and a location uniquely identify a message (that is, there cannot coexist in memory two messages of the same location and timestamp). Each thread maintains its *view* of the memory, where $T \in \text{View}$ is a function $Loc \rightarrow \text{Time}$. The thread’s view places lower bounds on the set of messages that the thread may read, as well as the timestamps it may pick when adding new messages to memory. Messages carry views as well, which record the thread’s view at the time the message was added to memory. When a message is read, its view is incorporated into the thread view, which, roughly speaking, ensures that the thread becomes aware of whatever the message it reads was aware of.

Formally, a *message* $m \in \text{Msg}$ is a tuple of the form $\langle x=v@t, T \rangle$ where $x \in Loc$, $v \in Val$, $t \in \text{Time}$, and $T \in \text{View}$. The states of the RA memory subsystem are given by $RA.Q \triangleq \mathcal{P}(\text{Msg}) \times (\mathbb{N} \rightarrow \text{View})$ (it consists of memory and thread views), with the initial state being $RA.q_0 \triangleq \langle \{ \langle x=0@0, T_0 \rangle \mid x \in Loc \}, \lambda n. T_0 \rangle$, where $T_0 \triangleq \lambda x. 0$ denotes the initial view.

²Without liveness guarantees, this program may not terminate under weak memory semantics. In this paper, as most existing work on weak memory specification and verification, we focus on finite traces and safety properties.

$$\begin{array}{c}
\neg\exists v', T'. \langle x=v'\text{@}t, T' \rangle \in M \\
\mathcal{T}(\tau)(x) < t \\
T = \mathcal{T}(\tau)[x \mapsto t] \\
M' = M \cup \{\langle x=v\text{@}t, T \rangle\} \\
\mathcal{T}' = \mathcal{T}[\tau \mapsto T] \\
l = W(x, v) \\
\hline
\langle M, \mathcal{T} \rangle \xrightarrow{\langle \tau, l \rangle}_{\text{RA}} \langle M', \mathcal{T}' \rangle
\end{array}
\quad
\begin{array}{c}
\langle x=v\text{@}t, T \rangle \in M \\
\mathcal{T}(\tau)(x) \leq t \\
\mathcal{T}' = \mathcal{T}[\tau \mapsto \mathcal{T}(\tau) \sqcup T] \\
l = R(x, v) \\
\hline
\langle M, \mathcal{T} \rangle \xrightarrow{\langle \tau, l \rangle}_{\text{RA}} \langle M, \mathcal{T}' \rangle
\end{array}$$

$$\begin{array}{c}
\langle x=v_R\text{@}t, T_R \rangle \in M \quad \mathcal{T}(\tau)(x) \leq t \\
\neg\exists v, T. \langle x=v\text{@}t + 1, T \rangle \in M \\
T_W = \mathcal{T}(\tau)[x \mapsto t + 1] \sqcup T_R \\
M' = M \cup \{\langle x=v_W\text{@}t + 1, T_W \rangle\} \quad \mathcal{T}' = \mathcal{T}[\tau \mapsto T_W] \\
l = \text{RMW}(x, v_R, v_W) \\
\hline
\langle M, \mathcal{T} \rangle \xrightarrow{\langle \tau, l \rangle}_{\text{RA}} \langle M', \mathcal{T}' \rangle
\end{array}$$

Figure 3. Transitions of the RA memory subsystem.

The transitions of RA are given in Fig. 3, where \sqcup denotes pointwise maximum ($T_1 \sqcup T_2 = \lambda x. \max\{T_1(x), T_2(x)\}$). To perform a write to x , thread τ (1) picks a timestamp that is available for x in the current memory and is greater than the timestamp in τ 's view for x ; (2) updates its view to include the new timestamp; (3) adds a message to the memory carrying τ 's (updated) view. In turn, to read from x , τ may pick any message of x in the memory whose timestamp is not lower than the timestamp in τ 's view for x . The view of the read message is incorporated in τ 's view. Finally, RMWs are obtained as an atomic combination of a read and a write, but crucially require that the timestamp of the added message is the successor of the timestamp of the read message. This guarantees that distinct RMWs never read from the same message (see Ex. 3.5 below).

Next, we provide simple examples of runs of concurrent programs under the RA memory subsystem, and analyze their robustness. When writing views, we often write only their non-zero elements.

Example 3.1 (Store buffer). The following program is the simplest example of a weak behavior allowed by RA:

$$\begin{array}{l}
x := 1 \quad \Big\| \quad y := 1 \\
a := y // 0 \quad \Big\| \quad b := x // 0
\end{array} \quad (\text{SB})$$

Here and henceforth, we use comment annotations to denote a particular program state. In this example, the annotations denote the state in which both program counters point to the end of the program, and the values of a and b are both 0. To reach this state under RA (cf. Def. 2.5), τ_1 may run first: add $\langle x=1\text{@}1, [x \mapsto 1] \rangle$ to the memory (this does not affect the view of τ_2), and read the initial message $\langle y=0\text{@}0, T_0 \rangle$. Then, τ_2 adds $\langle y=1\text{@}1, [y \mapsto 1] \rangle$ to the memory, and it is free to read the initial message $\langle x=0\text{@}0, T_0 \rangle$. Under SC, this state

is clearly unreachable, and thus, this program is not state robust against RA (cf. Def. 2.6).

Example 3.2 (Message passing). RA is designed to support “flag-based” synchronization. That is, the following annotated behavior is *disallowed* under RA:

$$\begin{array}{l}
x := 1 \quad \Big\| \quad a := y // 1 \\
y := 1 \quad \Big\| \quad b := x // 0
\end{array} \quad (\text{MP})$$

Indeed, τ_2 can read 1 for y , only after τ_1 executed the two writes adding messages $m_x = \langle x=1\text{@}t_x, [x \mapsto t_x] \rangle$ and $m_y = \langle y=1\text{@}t_y, [x \mapsto t_x, y \mapsto t_y] \rangle$ to the memory with $t_x, t_y > 0$. When reading m_y , τ_2 increases its view of x to be t_x , and then, since $t_x > 0$, it is unable to read the initial message of x , and must read m_x . Hence, it can be easily seen that this program is state robust against RA. This example also shows that a stronger definition of robustness, which requires that P_{SC} and P_{RA} have the same traces, is too strong to be of any use. Indeed, the transition $\langle \tau_2, R(y, 0) \rangle$ is allowed under RA also after τ_1 performed its two writes, and thus, such stronger condition would deem this program as non-robust.

Example 3.3 (Independent reads of independent writes). Unlike TSO, RA is *non-multi-copy-atomic*. That is, different threads may observe different stores in different orders. Thus, RA allows the following behavior:

$$x := 1 \quad \Big\| \quad \begin{array}{l} a := x // 1 \\ b := y // 0 \end{array} \quad \Big\| \quad \begin{array}{l} c := y // 1 \\ d := x // 0 \end{array} \quad \Big\| \quad y := 1 \quad (\text{IRIW})$$

Indeed, nothing in RA forbids a run in which the two writers finished their execution, and then τ_2 picks the message written by τ_1 for x and the initialization message for y , while τ_3 picks the message written by τ_4 for y and the initialization message for x . The corresponding program state is unreachable under SC, and, thus, this program is not state robust against RA. (It is, nevertheless, robust against TSO.)

Example 3.4. Unlike the SRA model [36], under RA, write steps do not have to choose a *globally* maximal timestamp. Thus, the following outcome is allowed [56], and the program is not state robust against RA (it is robust against TSO):

$$\begin{array}{l}
x := 1 \\
y := 2 \\
a := y // 1
\end{array} \quad \Big\| \quad \begin{array}{l}
y := 1 \\
x := 2 \\
a := x // 1
\end{array} \quad (2+2W)$$

Indeed, to execute both writes, τ_1 may add the messages $m_1^x = \langle x=1\text{@}2, [x \mapsto 2] \rangle$ and $m_2^y = \langle y=2\text{@}1, [x \mapsto 2, y \mapsto 1] \rangle$, and τ_2 may add the messages $m_1^y = \langle y=1\text{@}2, [y \mapsto 2] \rangle$ and $m_2^x = \langle x=2\text{@}1, [x \mapsto 1, y \mapsto 2] \rangle$. Now, τ_1 's view for y is 1 and it may read m_1^y , and τ_2 's view for x is 1 and it may read m_1^x .

Example 3.5. A crucial property of RMWs is that two (successful) RMWs never read from the same message. Indeed, this allows the standard implementation of lock acquisition using RMWs. This property is guaranteed in RA by forcing RMWs to use $t + 1$ as the timestamp for the added message, where t is the timestamp of the message that was read. To

see how this works consider the following (robust) program (the annotated behavior is disallowed under RA):

$$a := \text{CAS}(x, 0 \rightarrow 1) // 0 \parallel b := \text{CAS}(x, 0 \rightarrow 1) // 0 \quad (2\text{RMW})$$

W.l.o.g., if τ_1 runs first, it reads from the initialization message $\langle x=0@0, T_0 \rangle$ (it is the only message of x in the memory), and it is forced to add a message *with timestamp* 1, namely $\langle x=1@1, [x \mapsto 1] \rangle$. When τ_2 runs, it may *not* read from the initialization message, as it will again require adding a message of x with timestamp 1, but such a message already exists in memory. Thus, it may only read from the message that was added by τ_1 , and the CAS will fail.

Example 3.6. RMW operations to a distinguished otherwise-unused location can force synchronization, practically serving as *SC-fences* [36, 37] (in fact, this is how we encode SC-fences in our programming language). To see this, consider the following modification of the **SB** program:

$$\begin{array}{l} x := 1 \\ r := \text{FADD}(f, 0) \\ a := y // 0 \end{array} \parallel \begin{array}{l} y := 1 \\ r := \text{FADD}(f, 0) \\ b := x // 0 \end{array} \quad (\text{SB+RMWs})$$

Here, the annotated program behavior is disallowed under RA, and, consequently, this program is state robust against RA. Indeed, suppose, w.l.o.g., that τ_1 executes the FADD first and adds the message $m = \langle f=0@1, [x \mapsto t_x, f \mapsto 1] \rangle$ (where $t_x > 0$). When τ_2 executes its FADD, it has to read m , and update its view of x to t_x . Then, when it reads x it may not pick the initial message. It is crucial to use the same location in both FADDs: unlike TSO, under RA a single barrier (equivalently, a single FADD instruction to an otherwise-unused location) has no effect.

Finally, note that SC is clearly stronger than RA:

Lemma 3.7. *If a state \bar{q} of a concurrent program P is reachable under SC, then it is also reachable under RA.*

Proof. RA can simulate SC: in read (and RMW) steps, read the message with the maximal timestamp; and in write steps, pick t to be greater than the maximal timestamp of the messages of the written location. \square

4 Execution-Graph Robustness

While state robustness is a natural criterion, it is also very fragile and hard to test. For instance, if we replace the two written values in the **SB** program (Ex. 3.1) by 0's (writing once again the initial value), then the program becomes state robust, simply because reachable program states cannot distinguish runs under RA from runs under SC. Similarly, if we remove the two final reads in the **2+2W** program (Ex. 3.4), we obtain a “vacuously” state robust program. In this section, we present a stronger notion of robustness, which we call *execution-graph robustness*. (In particular, these two examples are not execution-graph robust.) In §5, we show how execution-graph robustness can be decided. This leads to a sound verification algorithm for state robustness.

Execution-graph robustness is based on different presentations of the SC and RA memory subsystems, which we denote by SCG and RAG, whose states are execution graphs capturing (partially ordered) histories of executed actions. The fact that the states of SCG and RAG are the same mathematical objects allows us to easily compare program behaviors under the two memory subsystems. In the rest of this section, we present SCG and RAG, and define execution-graph robustness. First, we define execution graphs, starting with their nodes, called *events*.

Definition 4.1. An *event* $e \in \text{Event}$ is a tuple $\langle \tau, s, l \rangle \in (\mathbb{N} \uplus \{\perp\}) \times \mathbb{N} \times \text{Lab}$, where τ is a thread identifier (or \perp for initialization events), s is a serial number inside each thread (0 for initialization events), and l is a label (as defined in Def. 2.1). The functions tid , sn , and lab return the thread identifier, serial number, and label of an event. The functions typ , loc , val_R , and val_W are lifted to events in the obvious way. We use R, W, RMW for the following sets of events:

$$\begin{aligned} R &\triangleq \{e \mid \text{typ}(e) \in \{R, \text{RMW}\}\} & W &\triangleq \{e \mid \text{typ}(e) \in \{W, \text{RMW}\}\} \\ \text{RMW} &\triangleq \{e \mid \text{typ}(e) = \text{RMW}\} \end{aligned}$$

We employ subscripts and superscripts to restrict sets of events to certain location and thread identifier (e.g., $W_x = \{w \in W \mid \text{loc}(w) = x\}$ and $E^\tau = \{e \in E \mid \text{tid}(e) = \tau\}$).

Definition 4.2. The set Init of *initialization events* is given by $\text{Init} \triangleq \{\langle \perp, 0, W(x, 0) \rangle \mid x \in \text{Loc}\}$. We say that a set $E \subseteq \text{Event}$ is *initialized* if $\text{Init} \subseteq E$, and $\text{tid}(e) \neq \perp$ and $\text{sn}(e) \neq 0$ for every $e \in E \setminus \text{Init}$.

Our representation of events induces a *sequenced-before* partial order on events, where $e_1 < e_2$ holds iff ($e_1 \in \text{Init}$ and $e_2 \notin \text{Init}$) or ($e_1 \notin \text{Init}$, $e_2 \notin \text{Init}$, $\text{tid}(e_1) = \text{tid}(e_2)$, and $\text{sn}(e_1) < \text{sn}(e_2)$). That is, initialization events precede all non-initialization events, while events of the same thread are ordered according to their serial numbers.

In turn, an execution graph consists of a set of events, a *reads-from* mapping that determines the write event from which each read reads its value, and a *modification order* which totally orders the writes to each location. In terms of the model in §3, the modification order represents the timestamp order on messages to each location.

Definition 4.3. An *execution graph* $G \in \text{EGraph}$ is a tuple $\langle E, rf, mo \rangle$ where:

1. E is an initialized finite set of events.
2. rf , called *reads-from*, is a relation on E satisfying:
 - If $\langle w, r \rangle \in rf$ then $w \in W, r \in R, \text{loc}(w) = \text{loc}(r), \text{val}_W(w) = \text{val}_R(r)$, and $w \neq r$.
 - $w_1 = w_2$ whenever $\langle w_1, r \rangle, \langle w_2, r \rangle \in rf$ (each read reads from at most one write).
 - $E \cap R \subseteq \text{codom}(rf)$ (each read reads from some write).
3. mo , called *modification order*, is a disjoint union of relations $\{mo_x\}_{x \in \text{Loc}}$, such that each mo_x is a strict total order on $E \cap W_x$.

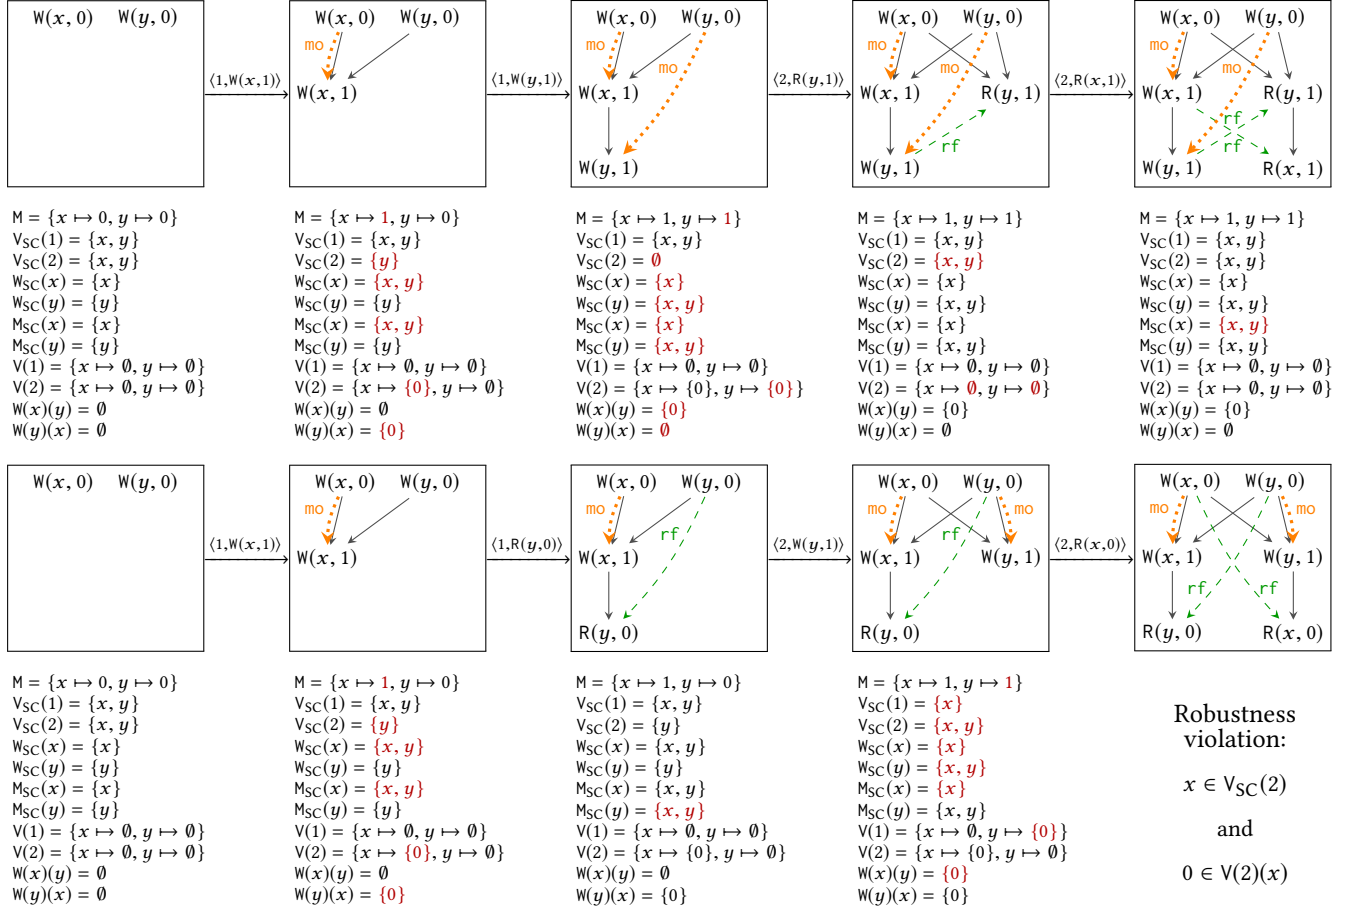


Figure 4. Illustrations of runs: (i) of SCG for the MP program (Ex. 3.2); and (ii) of RAG for the SB program (Ex. 3.1). Each illustration is followed by the corresponding run of SCM for monitoring robustness as described in §5 (deltas from the previous state are colored).

We denote the components of G by $G.E$, $G.rf$ and $G.mo$. We use $G.po$ to denote the restriction of the order on events to $G.E$ ($G.po \triangleq [G.E]; <; [G.E]$). In addition, for a set $E' \subseteq \text{Event}$, we write $G.E'$ for $G.E \cap E'$ (e.g., $G.W = G.E \cap W$).

Next, we define a general execution-graph-based memory subsystem, called FG (standing for “Free Graphs”). Later, the memory subsystems SCG and RAG are defined as restrictions of FG. To define FG, we use the following notation that extends a given graph G with a new event e , placed last in its thread, and either reading from a designated event w or placed as the immediate successor of w in the modification order. When e is an RMW event, it is *both* reading from w , and placed as the immediate successor of w in the modification order. This is in accordance with the usual atomicity restriction in declarative semantics, according to which RMWs read from their immediate *mo*-predecessors.

Notation 4.4. For $G \in \text{EGraph}$, $\tau \in \mathbb{N}$, $l \in \text{Lab}$ and $w \in W$, $\text{add}(G, \tau, l, w)$ denotes the triple $\langle E', rf', mo' \rangle$ defined as

follows, where $e = \langle \tau, \max\{\text{sn}(e) \mid e \in G.E^\tau\} + 1, l \rangle$:

$$E' = G.E \cup \{e\} \quad rf' = \begin{cases} G.rf \cup \{\langle w, e \rangle\} & e \in R \\ G.rf & \text{otherwise} \end{cases}$$

$$mo' = \begin{cases} G.mo \cup \text{dom}(G.mo^?; [\{w\}]) \times \{e\} & e \in W \\ \cup \{e\} \times \text{codom}([\{w\}]; G.mo) & \\ G.mo & \text{otherwise} \end{cases}$$

Definition 4.5. The initial execution graph G_0 is given by $G_0 \triangleq \langle \text{Init}, \emptyset, \emptyset \rangle$. The memory subsystem FG is defined by $\text{FG.Q} \triangleq \text{EGraph}$, $\text{FG.q}_0 \triangleq G_0$, and \rightarrow_{FG} is defined as follows:

$$\frac{w \in G.W_{\text{loc}(l)} \quad \text{typ}(l) \in \{R, \text{RMW}\} \implies \text{val}_w(w) = \text{val}_R(l)}{G \xrightarrow{\langle \tau, l \rangle}_{\text{FG}} \text{add}(G, \tau, l, w)}$$

The conditions in FG’s step ensure that $\text{add}(G, \tau, l, w)$ is indeed an execution graph: *mo* should only relate events in W of the same location; and *rf* goes from W to R only between

events of the same location and matching values. Below, we refer to the write w in such steps as the *predecessor write*.

Next, we present the memory subsystems SCG and RAG. Both are *based on execution graphs*: $\text{SCG.Q} = \text{RAG.Q} \triangleq \text{EGraph}$; $\text{SCG.q}_0 = \text{RAG.q}_0 \triangleq G_0$; and $\xrightarrow{\langle \tau, l \rangle}_{\text{SCG}} \subseteq \xrightarrow{\langle \tau, l \rangle}_{\text{FG}}$ and $\xrightarrow{\langle \tau, l \rangle}_{\text{RAG}} \subseteq \xrightarrow{\langle \tau, l \rangle}_{\text{FG}}$ for every $\tau \in \mathbb{N}$ and $l \in \text{Lab}$.

4.1 The Memory Subsystem SCG

The steps of SCG are uniformly given by:

$$\frac{\text{typ}(l) \in \{\text{R}, \text{RMW}\} \implies \text{val}_w(G.w_{\text{loc}(l)}^{\max}) = \text{val}_R(l)}{G \xrightarrow{\langle \tau, l \rangle}_{\text{SCG}} \text{add}(G, \tau, l, G.w_{\text{loc}(l)}^{\max})}$$

where $G.w_x^{\max}$ denotes the $G.\text{mo}$ maximal write to x in G ($G.w_x^{\max} \triangleq \max_{G.\text{mo}} G.W_x$).

SCG steps require the predecessor write to be $G.w_{\text{loc}(l)}^{\max}$: added write events are placed last in $G.\text{mo}$, and read events read from the latest added write. Figure 4 illustrates an example of a run of the MP program (Ex. 3.2) under SCG.

Lemma 4.6. *SCG and SC have the same traces.*

Proof (outline). Define the *memory* of a given $G \in \text{EGraph}$ by $M(G) \triangleq \lambda x. \text{val}_w(G.w_x^{\max})$. It is easy to show that $\text{SC.q}_0 = M(G_0)$ and $\{\langle M(G), G \mid G \in \text{EGraph} \rangle\}$ is a bisimulation relation between SC and SCG. \square

4.2 The Memory Subsystem RAG

To define the transitions of RAG, we use the following standard derived “happens-before” relation:

$$G.\text{hb} \triangleq (G.\text{po} \cup G.\text{rf})^+$$

Roughly speaking, $G.\text{hb}$ abstracts RA’s execution order: any run of the timestamp machine in §3 follows some linearization of hb , and, conversely, all linearizations of hb induce runs of the timestamp machine. Using hb , the steps of RAG are uniformly given by:

$$\frac{\begin{array}{l} w \in G.W_{\text{loc}(l)} \\ \text{typ}(l) \in \{\text{R}, \text{RMW}\} \implies \text{val}_w(w) = \text{val}_R(l) \\ w \notin \text{dom}(G.\text{mo}; G.\text{hb}^?; [G.E^\tau]) \end{array}}{\text{typ}(l) \in \{\text{W}, \text{RMW}\} \implies w \notin \text{dom}(G.\text{mo}|_{\text{imm}}; [\text{RMW}])} \\ G \xrightarrow{\langle \tau, l \rangle}_{\text{RAG}} \text{add}(G, \tau, l, w)$$

The first two conditions in the step are the general conditions of FG (see Def. 4.5). The third and fourth conditions restrict the choice of the predecessor write w . Unlike in SCG, w is not necessarily $G.w_{\text{loc}(l)}^{\max}$. Instead, it is subject to two conditions. First, the thread that takes the action must not have observed an mo -later write, where observed writes are writes that have a (possibly empty) hb -path to (some event of) the thread ($w \notin \text{dom}(G.\text{mo}; G.\text{hb}^?; [G.E^\tau])$). Referring to the timestamp machine, this is in accordance with the choice of the message to read in read steps and the new added messages in write

steps (their timestamp cannot be smaller than the last timestamp observed by the thread for the location). Second, when writing (by a write or an RMW), the predecessor write w cannot be the immediate mo -predecessor of some (other) RMW event ($w \notin \text{dom}(G.\text{mo}|_{\text{imm}}; [\text{RMW}])$). In the timestamp machine, this corresponds to the fact that timestamp of the message added by an RMW must be the immediate successor of the timestamp of the message read by the RMW. Note that in graphs generated by RAG, RMWs always read from their immediate mo -predecessor ($G.\text{rf}; [\text{RMW}] = G.\text{mo}|_{\text{imm}}; [\text{RMW}]$), which is the usual atomicity condition in declarative weak memory semantics.

It is easy to see that SCG is more restrictive than RAG (and thus, the run of SCG for the MP program in Fig. 4 is also allowed under RAG):

Lemma 4.7. *If $G \xrightarrow{\langle \tau, l \rangle}_{\text{SCG}} G'$, then $G \xrightarrow{\langle \tau, l \rangle}_{\text{RAG}} G'$.*

Proof. Pick $w = G.w_{\text{loc}(l)}^{\max}$ as the predecessor write. By definition we have $w \in G.W_{\text{loc}(l)}$, $w \notin \text{dom}(G.\text{mo}; G.\text{hb}^?; [G.E^\tau])$, and $w \notin \text{dom}(G.\text{mo}|_{\text{imm}}; [\text{RMW}])$. \square

Figure 4 illustrates an example of a run of the SB program (Ex. 3.1) under RAG. The last step there is disallowed by SCG—the predecessor write is not the mo -maximal one.

Next, we state the equivalence between RAG and RA. A proof outline is provided in [1, §B].

Lemma 4.8. *RAG and RA have the same traces.*

4.3 Execution-Graph Robustness

Next, we define execution-graph robustness and show that it implies state robustness.

Definition 4.9. A concurrent program P is *execution-graph robust* against RA if every reachable state $\langle \bar{q}, G \rangle$ in the concurrent system P_{RAG} is also reachable in P_{SCG} .

Proposition 4.10. *If P is execution-graph robust against RA then it is state robust against RA.*

Proof. Let \bar{q} be a state of P that is reachable under RA. Let $\langle M, \mathcal{T} \rangle \in \text{RA.Q}$ such that $\langle \bar{q}, \langle M, \mathcal{T} \rangle \rangle$ is reachable in P_{RA} . By Lemma 4.8, $\langle \bar{q}, G \rangle$ is reachable in P_{RAG} for some G . Since P is execution-graph robust against RA, it follows that $\langle \bar{q}, G \rangle$ is reachable in P_{SCG} . By Lemma 4.6, $\langle \bar{q}, M \rangle$ is reachable in P_{SC} for some $M \in \text{SC.Q}$, and so \bar{q} is reachable under SC. \square

Execution-graph robustness, as we demonstrate below, is not overly strong for establishing state robustness in a variety of concurrent algorithms. In particular, the state robust litmus tests mentioned in §3 (MP, 2RMW, SB+RMWs) are also execution-graph robust.

5 Verifying Execution-Graph Robustness

In this section, we present our approach to the verification of execution-graph robustness against RA. First, Thm. 5.1 below

reduces this problem to reachability of certain configurations in P_{SCG} . To state this theorem, we require two more standard derived relations in execution graphs:

$$G.fr \triangleq (G.rf^{-1} ; G.mo) \setminus [G.E] \text{ (from-read/reads-before)}$$

$$G.hb_{SC} \triangleq (G.hb \cup G.mo \cup G.fr)^+ \text{ (SC-happens-before)}$$

The *from-read* relation, fr , relates every read event r to all writes that are mo -later than the write that r reads from (identity is subtracted to avoid self loops in RMW events). The *SC-happens-before* relation, $G.hb_{SC}$, following [51], abstracts SC's execution order: to yield certain execution G , the SCG memory subsystem must follow $G.hb_{SC}$. Thus, runs of SCG can yield an execution graph G iff $G.hb_{SC}$ is irreflexive.

Theorem 5.1. *Let P be a concurrent program. Call a tuple $\langle \bar{q}, G, \tau, l, w \rangle \in P.Q \times EGraph \times Tid \times Lab \times Event$ a non-robustness witness for P if the following hold:*

- $\langle \bar{q}, G \rangle$ is reachable in the concurrent system P_{SCG} .
- \bar{q} enables $\langle \tau, l \rangle$ (in the LTS induced by P).
- $w \neq G.w_{loc(l)}^{max}$.
- $G \xrightarrow{\langle \tau, l \rangle}_{RAG} add(G, \tau, l, w)$.
- $G.w_{loc(l)}^{max} \in dom(G.hb_{SC} ; [G.E^\tau])$.

Then, P is execution-graph robust against RA iff there does not exist a non-robustness witness for P .

Theorem 5.1 reduces execution-graph robustness of a program P to the existence of a reachable state in the concurrent system P_{SCG} that satisfies certain properties. More precisely, P is not robust iff there exist a reachable state $\langle \bar{q}, G \rangle$ of P_{SCG} and a transition $\langle \tau, l \rangle$ that is enabled in \bar{q} , such that: (a) there is an hb_{SC} -path in G from $w_{loc(l)}^{max}$ to (some event of) thread τ ; and (b) G enables the transition $\langle \tau, l \rangle$ in RAG with a predecessor write $w \neq G.w_{loc(l)}^{max}$.

The proof is given in [1, §A]. Roughly speaking, we utilize purely declarative presentations of SCG and RAG, and show that the existence of a non-robustness witness allows RA executions to diverge w.r.t. SC ones, and that given a “minimal” such divergence, one can construct a non-robustness witness. The latter has generally a similar structure to proofs establishing the DRF (data-race-freedom) guarantee [6, 29].

We note that DRF for RA can be easily obtained as a corollary of **Thm. 5.1**. Indeed, if a program P is race-free (under SC), then all reachable states $\langle \bar{q}, G \rangle$ in P_{SCG} satisfy $G.mo \cup G.fr \subseteq G.hb$. It follows that $G.hb_{SC} \subseteq G.hb$, and thus, $G.w_{loc(l)}^{max} \in dom(G.hb_{SC} ; [G.E^\tau])$ implies that only $G.w_{loc(l)}^{max}$ may serve as the predecessor write in RAG transitions from G . Therefore, P cannot have a non-robustness witness, and **Thm. 5.1** ensures that it is execution-graph robust.

Similarly, it follows that a program with no concurrent writes under SC cannot have weak behaviors allowed by RA (as was established in [7] for a certain variant of causal consistency). Indeed, if P has no concurrent writes (under SC), then all reachable states $\langle \bar{q}, G \rangle$ in P_{SCG} satisfy $[W]$;

$G.hb_{SC} \subseteq G.hb$ (use hb to reach the last write in the hb_{SC} -path and from that point on no mo and fr edges are used). Again, **Thm. 5.1** ensures that P is execution-graph robust.

It remains to show that the condition in **Thm. 5.1** can be automatically checked. Since SCG is not finite (execution graphs of programs with loops may grow unboundedly), we cannot naively explore traces of P_{SCG} . The key idea is to define a *finite* memory subsystem, which we call SCM (for SC with Monitors), that simulates SCG (so that they have the same traces) and precisely track the properties of SCG's execution graphs that are needed for monitoring the above condition.

Next, we gradually present SCM's states, which are composed of eight components in total, and the transitions between them. **Figure 4** provides detailed examples of runs of SCM for the **MP** and **SB** programs, together with the corresponding runs of SCG. Below, we use I as a metavariable for states of SCM and write $I(G)$ for the SCM state that corresponds to an execution graph G .

Memory (I.M). The basic building block for SCM is the (finite) memory subsystem SC whose states are simple location-value mappings (see §2.3). Thus, a state I of SCM has a memory component, denoted $I.M$, which is a function from Loc to Val storing the value written by $G.w_x^{max}$ for every location x . Formally, we have

$$I(G).M = \lambda x. val_w(G.w_x^{max}).$$

The transitions of SCM are subject to the same constraints as SC with respect to this component. The other components of the states of SCM are used to track more properties of G , and do not restrict SCM's traces. Thus, the fact that SCM has the same traces as SCG directly follows from **Lemma 4.6**.

hb_{SC} -tracking ($I.V_{SC}, I.M_{SC}, I.W_{SC}$). For checking condition (a) above, we need to know for every thread τ and location x whether τ is “ hb_{SC} -aware” of w_x^{max} . To include and maintain this information in a state I of SCM, we use three components, denoted by $I.V_{SC}$, $I.M_{SC}$ and $I.W_{SC}$.

The first, $I.V_{SC}$, is a function in $Tid \rightarrow \mathcal{P}(Loc)$ tracking precisely this property. Formally, we have:

$$I(G).V_{SC} = \lambda \tau. \{x \mid G.w_x^{max} \in dom(G.hb_{SC} ; [Init \cup G.E^\tau])\}.$$

Having $x \in I(G).V_{SC}(\tau)$ means that τ is hb_{SC} -aware of w_x^{max} , i.e., $G.w_x^{max}$ is an initialization write (of which all threads are aware) or $\langle G.w_x^{max}, e \rangle \in G.hb_{SC} ?$ for some $e \in G.E^\tau$.

In turn, to maintain $I.V_{SC}$, we include two additional components, $I.M_{SC}$ and $I.W_{SC}$, both of which are functions in $Loc \rightarrow \mathcal{P}(Loc)$. Consider first an SCG-step that adds a write (or RMW) event w to location x in thread τ . Following SCG, w is placed it last in mo , which means that every event accessing x becomes hb_{SC} before w (writes to x have mo to w and reads from x have fr to w). In turn, the thread τ in which w is added will have (additional) hb_{SC} -paths from every w_y^{max} that previously had an hb_{SC} -path to some event accessing x .

	$\langle \tau, W(x, v) \rangle$ or $\langle \tau, \text{RMW}(x, v_R, v_W) \rangle$	$\langle \tau, R(x, v) \rangle$
$V'_{\text{SC}} = \lambda\tau.$	$\begin{cases} V_{\text{SC}}(\tau) \cup M_{\text{SC}}(x) & \pi = \tau \\ V_{\text{SC}}(\pi) \setminus \{x\} & \pi \neq \tau \end{cases}$	$\begin{cases} V_{\text{SC}}(\tau) \cup W_{\text{SC}}(x) & \pi = \tau \\ V_{\text{SC}}(\pi) & \pi \neq \tau \end{cases}$
$M'_{\text{SC}} = \lambda y.$	$\begin{cases} M_{\text{SC}}(x) \cup V_{\text{SC}}(\tau) & y = x \\ M_{\text{SC}}(y) \setminus \{x\} & y \neq x \end{cases}$	$\begin{cases} M_{\text{SC}}(x) \cup V_{\text{SC}}(\tau) & y = x \\ M_{\text{SC}}(y) & y \neq x \end{cases}$
$W'_{\text{SC}} = \lambda y.$	$\begin{cases} M_{\text{SC}}(x) \cup V_{\text{SC}}(\tau) & y = x \\ W_{\text{SC}}(y) \setminus \{x\} & y \neq x \end{cases}$	$W_{\text{SC}}(y)$

Figure 5. Maintaining V_{SC} , M_{SC} and W_{SC} in SCM transitions.

To properly reflect this in $I.V_{\text{SC}}(\tau)$, we maintain $I.M_{\text{SC}}$ that tracks for every $x \in \text{Loc}$ the set of locations y such that w_y^{max} has an hb_{SC} -path to some event accessing x . In steps that write to x in thread τ , we incorporate $I.M_{\text{SC}}(x)$ into $I.V_{\text{SC}}(\tau)$.

Second, similarly, when an SCG-step adds a read event r of location x in thread τ , it reads from w_x^{max} , and so we have $\langle w_x^{\text{max}}, r \rangle \in \text{hb}_{\text{SC}}$. In turn, thread τ will have (additional) hb_{SC} -paths from every w_y^{max} that previously had an hb_{SC} -path to w_x^{max} . Accordingly, the $I.W_{\text{SC}}$ component tracks for every $x \in \text{Loc}$ the set of locations y such that $G.w_y^{\text{max}}$ has an hb_{SC} -path to w_x^{max} . In steps that read from x in thread τ , we incorporate $I.W_{\text{SC}}(x)$ into $I.V_{\text{SC}}(\tau)$. Note that, while $y \in I.M_{\text{SC}}(x)$ iff w_y^{max} has an hb_{SC} -path to *some event* accessing x , we have $y \in I.W_{\text{SC}}(x)$ iff w_y^{max} has an hb_{SC} -path to w_x^{max} (equivalently, to some *write* event accessing x). This implies, in particular, that we always have $I.W_{\text{SC}}(x) \subseteq I.M_{\text{SC}}(x)$.

Formally, the meaning of these two “helper” components is given by:

$$I(G).M_{\text{SC}} = \lambda x. \{y \mid G.w_y^{\text{max}} \in \text{dom}(G.\text{hb}_{\text{SC}}^?; [G.E_x])\}$$

$$I(G).W_{\text{SC}} = \lambda x. \{y \mid \langle G.w_y^{\text{max}}, G.w_x^{\text{max}} \rangle \in G.\text{hb}_{\text{SC}}^?\}$$

Initially, we take $\text{SCM.q}_0.V_{\text{SC}} = \lambda\tau.$ Loc and $\text{SCM.q}_0.M_{\text{SC}} = \text{SCM.q}_0.W_{\text{SC}} = \lambda x. \{x\}$.

Figure 5 presents the maintenance of $I.V_{\text{SC}}$, $I.M_{\text{SC}}$ and $I.W_{\text{SC}}$ (primed components denote the corresponding components after the transition mentioned in the column headers). In particular, note that when a write (or RMW) to x is performed it becomes the new w_x^{max} and it has no hb_{SC} -paths to other events in the execution graph. Thus, we remove x from $I.V_{\text{SC}}(\pi)$ for every thread π except for the one that performed the write, as well as from $I.M_{\text{SC}}(y)$ and $I.W_{\text{SC}}(y)$ for every $y \neq x$. In addition, when accessing location x in thread τ , $I.M_{\text{SC}}(x)$ inherits $I.V_{\text{SC}}(\tau)$ (every event that had hb_{SC} -path to thread τ now has hb_{SC} -path to an event accessing x), and when writing to location x in thread τ , $I.W_{\text{SC}}(x)$ inherits $I.V_{\text{SC}}(\tau)$ (every event that had hb_{SC} -path to thread τ now has hb_{SC} -path to w_x^{max}).

RAG-tracking ($I.V$, $I.W$, $I.V_{\text{RMW}}$, $I.W_{\text{RMW}}$). It remains to extend the instrumentation, so that we can check for every thread τ and label l , whether the transition $\langle \tau, l \rangle$ is enabled in RAG with a predecessor write that is not $w_{\text{loc}(l)}^{\text{max}}$ (condition (b)

above). For this matter, we include four additional components in the state I of SCM. Two of them, $I.V$ and $I.V_{\text{RMW}}$, are functions in $(\text{Tid} \times \text{Loc}) \rightarrow \mathcal{P}(\text{Val})$ and are the ones used to check the above condition. The other two, $I.W$ and $I.W_{\text{RMW}}$, are functions in $(\text{Loc} \times \text{Loc}) \rightarrow \mathcal{P}(\text{Val})$ and, as before, are used to properly maintain $I.V$ and $I.V_{\text{RMW}}$.

To understand these components, recall the transition of RAG in §4.2:

Read: Consider first a read transition $\langle \tau, l \rangle$ with $l = R(x, v_R)$. By definition, an execution graph G enables $\langle \tau, l \rangle$ with a predecessor write $w \in G.W_x$ if $\text{val}_W(w) = v_R$ and $w \notin \text{dom}(G.\text{mo}; G.\text{hb}^?; [G.E^\tau])$. To be able to check this condition, we use $I.V$ to track for every $\tau \in \text{Tid}$ and $x \in \text{Loc}$ the set of values that are written by some $w \in G.W_x$ that is not $G.w_x^{\text{max}}$ and satisfies $w \notin \text{dom}(G.\text{mo}; G.\text{hb}^?; [G.E^\tau])$. Then, to check condition (b) above, we check whether $v_R \in I.V(\tau)(x)$. In other words, $I.V(\tau)(x)$ tracks the set of values that can be read by thread τ from x under RAG, excluding the case of reading from w_x^{max} (which is also allowed by SCG).

As before, to maintain $I.V$, we use another component in I . When thread τ reads (or performs an RMW to) x in a transition of SCG, it induces an $\text{mo}; \text{hb}$ -path to thread τ from any write that had $\text{mo}; \text{hb}$ -path to w_x^{max} . Thus, after such transition, $I.V(\tau)(y)$ should be restricted to values written by some $w \in G.W_y$ such that $\langle w, G.w_x^{\text{max}} \rangle \notin G.\text{mo}; G.\text{hb}^?$. Accordingly, $I.W$ tracks for every pair $x, y \in \text{Loc}$ the set of values that are written by some write $w \in G.W_y$ that is not $G.w_y^{\text{max}}$ and satisfies $\langle w, G.w_x^{\text{max}} \rangle \notin G.\text{mo}; G.\text{hb}^?$.

Write and RMW: A write (or RMW) transition is similar, but it is subject to an additional constraint in RAG: the predecessor write w should not be an mo -immediate predecessor of an RMW event in G (equivalently, w should not be read by an RMW event). For this condition, we use $I.V_{\text{RMW}}$, that, as $I.V$, tracks for every $\tau \in \text{Tid}$ and $x \in \text{Loc}$ the set of values that are written by some $w \in G.W_x$ that is not $G.w_x^{\text{max}}$ and satisfies $w \notin \text{dom}(G.\text{mo}; G.\text{hb}^?; [G.E^\tau])$, but further requires that $w \notin \text{dom}(G.\text{mo}|_{\text{imm}}; [\text{RMW}])$. To maintain $I.V_{\text{RMW}}$, we use $I.W_{\text{RMW}}$, which is similar to $I.W$ with the same additional condition on w (i.e., $w \notin \text{dom}(G.\text{mo}|_{\text{imm}}; [\text{RMW}])$).

Formally, the meaning of these components is given by:

$$I(G).V = \lambda\tau, x. \text{val}_W[W \setminus \text{dom}(R; [G.E^\tau])]$$

$$I(G).W = \lambda y, x. \text{val}_W[W \setminus \text{dom}(R; [\{G.w_y^{\text{max}}\}])]$$

$$I(G).V_{\text{RMW}} = \lambda\tau, x. \text{val}_W[W \setminus \text{dom}(R; [G.E^\tau] \cup R_{\text{RMW}})]$$

$$I(G).W_{\text{RMW}} = \lambda y, x. \text{val}_W[W \setminus \text{dom}(R; [\{G.w_y^{\text{max}}\} \cup R_{\text{RMW}}])]$$

where $W = G.W_x \setminus \{G.w_x^{\text{max}}\}$, $R = G.\text{mo}; G.\text{hb}^?$ and $R_{\text{RMW}} = G.\text{mo}|_{\text{imm}}; [\text{RMW}]$ (the function val_W is extended to sets of events in the obvious way). Initially, since each location has only one write in the initial graph, these four components all return the empty set of values. Figure 6 presents our maintenance of these components.

	$\langle \tau, W(x, v) \rangle$ where $v_R = M(x)$	$\langle \tau, R(x, v) \rangle$	$\langle \tau, RMW(x, v_R, v_W) \rangle$
$V' = \lambda \pi, y.$	$\begin{cases} \emptyset & \pi = \tau, y = x \\ V(\pi)(x) \cup \{v_R\} & \pi \neq \tau, y = x \\ V(\pi)(y) & y \neq x \end{cases}$	$\begin{cases} V(\tau)(y) \cap W(x)(y) & \pi = \tau \\ V(\pi)(y) & \pi \neq \tau \end{cases}$	$\begin{cases} V(\tau)(y) \cap W(x)(y) & \pi = \tau \\ V(\pi)(x) \cup \{v_R\} & \pi \neq \tau, y = x \\ V(\pi)(y) & \pi \neq \tau, y \neq x \end{cases}$
$W' = \lambda z, y.$	$\begin{cases} V(\tau)(y) & z = x, y \neq x \\ W(z)(x) \cup \{v_R\} & z \neq x, y = x \\ W(z)(y) & \text{otherwise} \end{cases}$	$W(z)(y)$	$\begin{cases} W(x)(y) \cap V(\tau)(y) & z = x, y \neq x \\ W(z)(x) \cup \{v_R\} & z \neq x, y = x \\ W(z)(y) & \text{otherwise} \end{cases}$
$V'_{RMW} = \lambda \pi, y.$	$\begin{cases} \emptyset & \pi = \tau, y = x \\ V_{RMW}(\pi)(x) \cup \{v_R\} & \pi \neq \tau, y = x \\ V_{RMW}(\pi)(y) & y \neq x \end{cases}$		$\begin{cases} V_{RMW}(\tau)(y) \cap W_{RMW}(x)(y) & \pi = \tau \\ V_{RMW}(\pi)(y) & \pi \neq \tau \end{cases}$
$W'_{RMW} = \lambda z, y.$	$\begin{cases} V_{RMW}(\tau)(y) & z = x, y \neq x \\ W_{RMW}(z)(x) \cup \{v_R\} & z \neq x, y = x \\ W_{RMW}(z)(y) & \text{otherwise} \end{cases}$	$W_{RMW}(z)(y)$	$\begin{cases} W_{RMW}(x)(y) \cap V_{RMW}(\tau)(y) & z = x, y \neq x \\ W_{RMW}(z)(y) & \text{otherwise} \end{cases}$

Figure 6. Maintaining V , W , V_{RMW} , and W_{RMW} in SCM transitions.

Putting all pieces together, the states of SCM are tuples $I = \langle M, V_{SC}, M_{SC}, W_{SC}, V, W, V_{RMW}, M_{RMW} \rangle$. Its transitions are obtained by instrumenting the transitions of SC (which govern the M component) with the transformations in Figures 5 and 6. The next lemma (which we proved in Coq) ensures that they track the intended properties.

Lemma 5.2. *The following hold:*

- $SCM.q_0 = I(G_0)$.
- If $G \xrightarrow{\langle \tau, l \rangle}_{SCG} G'$, then $I(G) \xrightarrow{\langle \tau, l \rangle}_{SCM} I(G')$.
- If $I(G) \xrightarrow{\langle \tau, l \rangle}_{SCM} I'$, then $G \xrightarrow{\langle \tau, l \rangle}_{SCG} G'$ and $I(G') = I'$ for some $G' \in EGraph$.

Our main result easily follows from Thm. 5.1 and Lemma 5.2:

Theorem 5.3. *P is execution-graph robust against RA iff for every reachable state $\langle \bar{q}, I \rangle$ in P_{SCM} , the following hold for every $\langle \tau, l \rangle$ that is enabled in \bar{q} and satisfies $\text{loc}(l) \in I.V_{SC}(\tau)$, where $x = \text{loc}(l)$ and $v_R = \text{val}_R(l)$:*

- if $\text{typ}(l) = W$ then $I.V_{RMW}(\tau)(x) = \emptyset$.
- if $\text{typ}(l) = R$ then $v_R \notin I.V(\tau)(x)$.
- if $\text{typ}(l) = RMW$ then $v_R \notin I.V_{RMW}(\tau)(x)$.

PSPACE-completeness (assuming bounded data domain as we defined in §2) easily follows:

Corollary 5.4. *Verifying execution-graph robustness against RA for a given input program is PSPACE-complete.*

Proof (outline). For the upper bound, we can (gradually) guess a run of P_{SCM} and check the conditions of Thm. 5.3 at each step. The memory required for storing a state is polynomial in the size of P . The lower bound is established as the one in [19] for TSO, by a reduction from reachability under SC (which is PSPACE-complete [35]): A program can be made robust by adding fences (as in Ex. 3.6) between every two instructions, and an artificial robustness violation (e.g., in the form of SB) can be added when the target state is reached. \square

Note that for verifying robustness we generate one reachability query, and since we only monitor traces, we do not add additional non-determinism w.r.t. reachability under SC. However, the instrumentation in SCM creates dependencies between instructions (e.g., both a write to x and a write to $y \neq x$ require to update the bit representing $y \in M_{SC}(x)$), which may hinder partial order reduction.

5.1 Abstract Value Management

The V and V_{RMW} (and, consequently, W and W_{RMW}) components in SCM states are often “too elaborate” for what is actually needed to verify robustness. For example, for a program P without CAS, wait and BCAS instructions, whether P_{RAG} enables a transition or not does not depend on the value being read. In such case, we only need to check whether $I.V(\tau)(x)$ is empty (for reads) and whether $I.V_{RMW}(\tau)(x)$ is empty (for writes and RMWs). More generally, we only need to track values that may affect P_{RAG} transitions (e.g., block a thread from executing or make an RMW succeed). Next, we use this observation to reduce the metadata size in SCM. To do so, we first define *critical values*.

Definition 5.5. A value $v \in \text{Val}$ is called a *critical value* of $x \in \text{Loc}$ in a sequential program S if at least one of the following hold for some $q \in S.Q$: (1) q enables $R(x, v)$ but there exists v' such that q does not enable $R(x, v')$ and $RMW(x, v', v_W)$ for every v_W ; (2) q enables $RMW(x, v, v_W)$ for some $v_W \in \text{Val}$ but there exists v' such that q does not enable $RMW(x, v', v'_W)$ for every v'_W . We call v a critical value of x in a (concurrent) program P if it is a critical value of x in $P(\tau)$ for some $\tau \in \text{Tid}$, and denote by $\text{Val}(P, x)$ the set of critical values of x in P .

For instance, if $\text{wait}(x = 1)$ is included in a program P then 1 is a critical value of x in P . Similarly, $r := \text{CAS}(x, 0 \rightarrow 1)$ (e.g., for implementing spin locks) makes 0 a critical value of x . A program without CAS, wait and BCAS instructions has

no critical values. On the other hand, in a program including an instruction like $r := \text{CAS}(x, r' \rightarrow e)$ (where the expected value is not a constant), we have $\text{Val}(P, x) = \text{Val}$ (in which case, our proposed optimization does not change anything).

Now, the $V, V_{\text{RMW}}, W, W_{\text{RMW}}$ components can be restricted to record information only about the critical values (so, we have $V, V_{\text{RMW}} : \text{Tid} \rightarrow \prod_{x \in \text{Loc}} \mathcal{P}(\text{Val}(P, x))$ and $W, W_{\text{RMW}} : \text{Loc} \rightarrow \prod_{x \in \text{Loc}} \mathcal{P}(\text{Val}(P, x))$), and additional components $\text{CV}, \text{CV}_{\text{RMW}} : \text{Tid} \rightarrow \mathcal{P}(\text{Loc})$ and $\text{CW}, \text{CW}_{\text{RMW}} : \text{Loc} \rightarrow \mathcal{P}(\text{Loc})$ (disjunctively) summarize all non-critical values. The latter are formally interpreted as follows (using the interpretations above):

$$\begin{aligned} I(G).\text{CV} &= \lambda\tau. \{y \mid I(G).V(\tau)(y) \setminus \text{Val}(P, y) \neq \emptyset\} \\ I(G).\text{CV}_{\text{RMW}} &= \lambda\tau. \{y \mid I(G).\text{CV}_{\text{RMW}}(\tau)(y) \setminus \text{Val}(P, y) \neq \emptyset\} \\ I(G).\text{CW} &= \lambda x. \{y \mid I(G).W(x)(y) \setminus \text{Val}(P, y) \neq \emptyset\} \\ I(G).\text{CW}_{\text{RMW}} &= \lambda x. \{y \mid I(G).W_{\text{RMW}}(x)(y) \setminus \text{Val}(P, y) \neq \emptyset\} \end{aligned}$$

That is, $\text{CV}(\tau)$ (respectively, $\text{CV}_{\text{RMW}}(\tau)$) contains all locations y for which there exist at least one non-critical value that is written by a non-**mo**-maximal write to y that can serve as the predecessor write in an RAG read (respectively, write or RMW) step. The maintenance of these components (given in [1, §C]) is straightforwardly derived from the maintenance of $V, V_{\text{RMW}}, W, W_{\text{RMW}}$.

In turn, three conditions are added to [Thm. 5.3](#):

- if $\text{typ}(l) = W$ then $x \notin I.\text{CV}_{\text{RMW}}(\tau)$.
- if $\text{typ}(l) = R$ and $v_R \notin \text{Val}(P, x)$ then $x \notin I.\text{CV}(\tau)$.
- if $\text{typ}(l) = \text{RMW}$ and $v_R \notin \text{Val}(P, x)$ then $x \notin I.\text{CV}_{\text{RMW}}(\tau)$.

This construction results in smaller instrumentation (and fewer operations to maintain the instrumentation), where the size (number of bits) of the monitoring metadata is

$$3|\text{Tid}||\text{Loc}| + 4|\text{Loc}|^2 + 2(|\text{Tid}| + |\text{Loc}|) \sum_{x \in \text{Loc}} |\text{Val}(P, x)|.$$

In particular, for programs without **CAS**, **wait** and **BCAS** instructions the metadata size is $3|\text{Tid}||\text{Loc}| + 4|\text{Loc}|^2$, while in the worst case (when all values are critical) we will have $|\text{Loc}|(|\text{Tid}| + 2|\text{Loc}| + 2|\text{Val}|(|\text{Tid}| + |\text{Loc}|))$. In some of the examples we checked, this optimization dramatically reduce the verification time (e.g., the ‘ticketlock4’ example in §7 is x9 faster). In addition, it may be beneficial for programs with infinite data domains but finite sets of critical values, where the (generally undecidable) reachability problem in P_{SCM} can be solved using abstraction techniques. (This is left for future work.)

6 Extension with Non-atomic Accesses

In this section, we describe an extension of our approach to handle C/C++’s non-atomic accesses, typically used for “data variables” (unlike “synchronization variables”). A data-race on a non-atomic access is considered an undefined behavior, and thus non-atomic accesses allow very efficient implementation. In turn, robustness of a program should imply that it has no data-races on non-atomic accesses.

For this extension, we assume that $\text{Loc} = \text{Loc}_{\text{ra}} \uplus \text{Loc}_{\text{na}}$ is composed from a set of *release/acquire* locations and a disjoint set of *non-atomic* locations (we do not consider release/acquire and non-atomic accesses to the same location). The programming language [Fig. 1](#) is extended with instructions $x_{\text{na}} := e$ and $r := x_{\text{na}}$ for $x_{\text{na}} \in \text{Loc}_{\text{na}}$, $e \in \text{Exp}$, and $r \in \text{Reg}$. The rest of the instructions only apply to locations in Loc_{ra} (in particular, there are no RMW instructions for non-atomic locations).

The SC and SCG systems ignore the type of the location, while RAG is extended to detect races on non-atomic locations. We refer to the extended memory subsystem as RAG+NA. The state of RAG+NA are execution graphs (as in RAG) as well as a special state, denoted by \perp , that the system enters once a race is detected. To define RAG+NA’s transitions, **hb** is modified so that only **rf**-edges on release/acquire accesses synchronize:

$$G.\text{hb} \triangleq (G.\text{po} \cup \bigcup_{x \in \text{Loc}_{\text{ra}}} [W_x]; G.\text{rf}; [R_x])^+$$

Now, the transitions of RAG+NA extend the transitions of RAG (which govern the release/acquire locations) with the following steps for non-atomic accesses:

$$\frac{x_{\text{na}} = \text{loc}(l) \quad x_{\text{na}} \in \text{Loc}_{\text{na}} \quad \text{typ}(l) = R \implies \text{val}_W(G.w_{x_{\text{na}}}^{\max}) = \text{val}_R(l) \quad G.w_{x_{\text{na}}}^{\max} \in \text{dom}(G.\text{hb}^?; [G.E^\tau])}{G \xrightarrow{\langle \tau, l \rangle}_{\text{RAG+NA}} \text{add}(G, \tau, l, G.w_x^{\max})}$$

$$\frac{\text{loc}(l) \in \text{Loc}_{\text{na}} \quad G.w_{\text{loc}(l)}^{\max} \notin \text{dom}(G.\text{hb}^?; [G.E^\tau])}{G \xrightarrow{\langle \tau, l \rangle}_{\text{RAG+NA}} \perp}$$

Thus, for a thread to successfully perform a non-atomic access to location x_{na} , it must have observed (in **hb**) the **mo**-maximal (equivalently, **hb**-maximal) write to x_{na} . Otherwise, the system moves to the \perp state.

Execution-graph robustness against RAG+NA is defined just as against RA (cf. [Def. 4.9](#)), and it implies state robustness against RAG+NA. Since P_{SCG} never reaches states of the form $\langle \bar{q}, \perp \rangle$, execution-graph robustness against RAG+NA implies that such states are not reachable in $P_{\text{RAG+NA}}$. Next, [Theorem 5.1](#) is extended as follows:

Definition 6.1. A state \bar{q} of a concurrent program is *racy* if \bar{q} enables both $\langle \tau, l_1 \rangle$ and $\langle \pi, l_2 \rangle$ for some $\tau \neq \pi$ and $l_1, l_2 \in \text{Lab}$ with $\text{loc}(l_1) = \text{loc}(l_2) \in \text{Loc}_{\text{na}}$ and $W \in \{\text{typ}(l_1), \text{typ}(l_2)\}$.

Theorem 6.2. A concurrent program P is *execution-graph robust against RAG+NA* iff there does not exist a non-robustness witness $\langle \bar{q}, G, \tau, l, w \rangle$ for P with $\text{loc}(l) \in \text{Loc}_{\text{ra}}$ (as defined in [Thm. 5.1](#)), and there does not exist a reachable state $\langle \bar{q}, G \rangle$ in P_{SCG} such that \bar{q} is *racy*.

The SCM system can be easily adapted for monitoring the conditions of [Thm. 6.2](#). The memory component in SCM’s

Program	Res	#T	LoC	Time	SC	Trencher TSO	
						Res	Time
barrier (BAR)	✓	2	11	1.6 (100%)	1.1	✗*	-
dekker-sc	✗	2	43	4.2 (100%)	1.3	✗	5.9
dekker-tso	✓	2	49	5.2 (100%)	1.3	✓	5.9
peterson-sc	✗	2	28	2.5 (100%)	1.2	✗	5.6
peterson-tso	✓	2	30	3.3 (100%)	1.3	✓	5.6
peterson-ra	✓	2	44	5.8 (100%)	1.2	✓	5.8
peterson-ra-dmitriy	✓	2	36	4.3 (100%)	1.2	✓	5.5
peterson-ra-bratosz	✗	2	28	3.4 (100%)	1.1	✗	5.6
lambert2-sc	✗	2	65	9.1 (100%)	1.3	✗	8.0
lambert2-tso	✗	2	69	13.7 (100%)	1.3	✓	8.2
lambert2-ra	✓	2	79	18.9 (99%)	1.4	✓	7.8
lambert2-3-ra	✓	3	123	215.6 (21%)	6.1	✗*	-
spinlock	✓	2	34	1.6 (100%)	1.2	✓	5.4
spinlock4	✓	4	66	6.4 (80%)	1.6	✓	6.8
ticketlock	✓	2	25	2.6 (100%)	1.1	✓	5.8
ticketlock4	✓	4	49	22.6 (25%)	7.5	✓	23.4
seqlock	✓	4	49	20.7 (16%)	3.4	✓	8.9
nbw-w-lr-rl	✓	4	50	5.7 (100%)	1.2	✓	8.6
rcu	✓	4	74	67.6 (10%)	2.2	✗*	-
rcu-offline	✓	3	215	137.9 (50%)	18.3	✗*	-
cilk-the-wsq-sc	✗	2	57	5.0 (100%)	1.2	✗	9.6
cilk-the-wsq-tso	✓	2	59	6.1 (100%)	1.3	✓	11.7
chase-lev-sc	✗	3	55	3.8 (100%)	29.5	✗	15.3
chase-lev-tso	✗	3	57	4.9 (100%)	31.3	✓	128.1
chase-lev-ra	✓	3	61	67.1 (8%)	38.1	✓	108.3

Figure 7. Experiments with *Rocker*

states is extended in the obvious way to track the latest value of non-atomic locations as well. Since non-atomic instructions do not affect inter-thread synchronization, the monitoring instrumentation in §5 requires no change (it only applies to the locations in Loc_{ra}). Since SCM and SCG have the same traces, the additional condition about races can be checked on SCM runs.

7 Implementation and Evaluation

We implemented our algorithm in a prototype tool called *Rocker* (for RObustness CheCker), which uses Spin [31] as a back-end model checker. The implementation and the examples it was tested on are available in the artifact accompanying this paper. *Rocker* takes as input a program in our toy programming language, and converts it to Promela code (Spin’s input language) with appropriate instrumentation and assertions that check for execution-graph robustness against RA. Thus, our implementation is actually using the SC memory subsystem, and implements the monitoring of SCM by instrumenting the input program. When a robustness violation is detected, one can use Spin’s output to see the trace leading to this violation. In addition, since in any case we explore traces of the input program under SC, *Rocker* allows one to include standard assertions, which will be verified as well by the model checker.

We performed a series of experiments on litmus tests, examples from [5, 17], and additional concurrent algorithms. Figure 7 summarizes the running times on some of the examples when executed on an Intel® Core™ i5-6300U CPU @ 2.40GHz GNU/Linux machine. Columns ‘Res’, ‘#T’, and ‘LoC’ respectively present the robustness of the input program, the number of threads, and total number of lines of code. Column ‘Time’ shows the verification time (in seconds), and the percentage of that time that was dedicated to compiling Spin’s verifier (using gcc with -O2). The latter often completely dominates the total time. Generating the input for Spin is negligibly fast (< 0.1s), as well as Spin’s verifier generation in C (< 0.2s). Column ‘SC’ provides, for the sake of comparison, the verification duration using Spin with no instrumentation whatsoever. In this mode, only the assertions in the input are verified assuming SC semantics.

For some of the examples Fig. 7 provides several versions of the same algorithm: The ‘-sc’ suffix denotes an original algorithm as designed for SC; the ‘-tso’ suffix denotes its strengthening with fences to ensure robustness against TSO; and, when needed, the ‘-ra’ suffix is a further strengthening that ensures robustness against RA. For instance, it is well known that Peterson mutual exclusion algorithm (‘peterson-sc’) is not robust against relaxed memory. For TSO, placing one fence in each thread suffices to ensure robustness. For RA, more fences are needed (‘peterson-ra’). Alternatively, as noted in [57], one may replace certain write operations by RMWs (‘peterson-dmitriy’). The choice of these writes is critical—*Rocker* correctly identified that a different version is incorrect (‘peterson-bratosz’). Other algorithms, which were designed with relaxed memory considerations in mind, e.g., Seqlocks [16] and a user-level RCU [26], do not require fences at all. Note that we have also verified robustness of a more involved RCU implementation (‘rcu-offline’), where the writer is not a unique thread, and threads may declare that they are going offline, stop the communication with the writer and return online later on.

Finally, column ‘Trencher’ provides the (total) running time of Trencher, a tool for verifying robustness against TSO [17], which also uses Spin for model checking. (A newer version of Trencher that implements its own model checker crashed on some of these examples.) Their notion of robustness is similar to execution-graph robustness, but it should be noted that *Rocker* and Trencher solve different problems: TSO and RA are fundamentally different models, where RA is weaker and non-multi-copy atomic. Thus, this comparison is of limited significance (see also §8). The input language is different as well. In particular, Trencher does not handle blocking instructions. For this reason, Trencher reports some examples as non-robust (marked with *), while no additional fences are needed for them to function correctly under TSO. We note that Trencher can be used in parallel to *Rocker* for verifying robustness against RA: a violation detected by Trencher implies non-robustness against RA.

8 Related Work

Robustness against weak memory semantics was studied before for *hardware* models, especially in the context of automatically enforcing robustness by inserting memory fences and other synchronization primitives (see, e.g., [9, 21, 24, 25], as well as [8] for a practical approximate generic approach).

In particular, robustness against TSO (and its PSO variant) [10, 32, 50] received considerable attention, e.g., [4, 5, 17–19, 22, 23, 30, 41–43, 49]. Generally speaking, the closest to our approach is Burckhardt and Musuvathi [22], implemented in a tool called Sober, which reduces robustness against TSO to reachability under SC in an instrumented program that verifies that TSO executions cannot diverge w.r.t. SC ones.³ In addition, verifying (trace based) robustness against TSO was shown by Bouajjani et al. [19] to be PSPACE complete—the same complexity, as we show, as verifying execution-graph robustness against RA.

Except for the fact that RA is strictly weaker than TSO (see the 2+2W and IRIW programs above), there are crucial differences between TSO and RA that do not allow one to apply the approaches developed for TSO when targeting RA. First, TSO's operational model provides a simple description of its runs, identifying a TSO run with an SC run where global effects of write instructions may be delayed. This presentation of TSO plays a key role in the characterization, verification and enforcement of robustness against TSO (see, e.g., [5, 17–19]). RA does not admit a similar presentation, and in fact, since RA is non-multi-copy-atomic (see Ex. 3.3), unlike TSO, RA cannot be explained by program transformations (instruction reorderings and eliminations) on top of SC [38]. Second, RMW operations in RA provide much weaker guarantees than in TSO, where even a failed CAS (when a CAS instruction is included as a primitive, as in [43]) serves as a memory fence. As described in §5, handling RMWs in RA (where, in particular, a failed CAS is nothing more than a plain read) requires certain technical novelties.

Less work was devoted to robustness against a *programming language* concurrency semantics. The well-known DRF guarantee [6, 29] is a simple robustness criterion, e.g., for a strengthened version of C11 [13, 39], but it is too weak, as (low-level) synchronizations naturally involve data-races, and often do not imply non-robustness. Meshman et al. [46] proposed an (approximate and incomplete) method that uses CDSchecker [48] for restricting non-SC behaviors of C11 programs. For a particular class of “server client programs”, it was shown in [36] that certain simple fence insertion strategy ensures robustness. However, in this paper we are interested in precise robustness verification for arbitrary programs.

Verification under RA has also received significant attention. This includes works on program logics, e.g., [27, 33, 37,

53–55], which require manual proofs, and (bounded) model checkers, e.g., [3, 34, 48], which provide limited guarantees for programs with loops. These methods can be used to verify programs that are not necessarily robust against RA. The verification problem of programs with loops under RA (i.e., given a program P and a state $\bar{q} \in P.Q$, is \bar{q} reachable under the concurrent system P_{RA} ?) was recently shown to be undecidable [2]. (For TSO, this problem is decidable but non-primitive recursive [11, 12].) As shown in [24, Thm. 2.12], this immediately entails the undecidability of state robustness.

Finally, robustness was also studied, e.g., in [15, 20, 28, 47], in the context of distributed systems, where SC is replaced by *serializability*. Unlike the current work, these works are focused on practical over-approximations, and do not provide provably precise general verification methods.

9 Conclusion

We have presented a method to verify execution-graph robustness against release/acquire concurrency semantics, in particular, establishing the decidability of this problem. Our method works by exploring only runs of the program under SC while monitoring certain properties for the detection of robustness violations. We believe that our result can play an important role in verification and development of concurrent algorithms for weak memory semantics, alongside with other existing methods.

In the future, we plan to study the applicability of our approach for different and extended models, such as RC11 [39], WeakRC11 [34], SRA [36], as well as transactional consistency models, such as PSI [52]. In addition, we are interested in deriving efficient and precise methods for automatic robustness enforcement (such as fence insertion) as were developed before for hardware models; as well as in handling parametrized programs with arbitrary number of threads.

Acknowledgments

We thank the PLDI'19 reviewers for their helpful feedback. This research was supported by the Israel Science Foundation (grant number 5166651), and by Len Blavatnik and the Blavatnik Family foundation. The first author was also supported by the Alon Young Faculty Fellowship.

References

- [1] Ori Lahav and Roy Margalit. 2019. Supplementary material for this paper. <https://www.cs.tau.ac.il/~orilahav/papers/pldi19full.pdf>
- [2] Parosh Aziz Abdulla, Jatin Arora, Mohamed Faouzi Atig, and Shankaranarayanan Krishna. 2019. Verification of programs under the release-acquire semantics. In *PLDI (to appear)*.
- [3] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Bengt Jonsson, and Tuan Phong Ngo. 2018. Optimal stateless model checking under the release-acquire semantics. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 135 (Oct. 2018), 29 pages. <https://doi.org/10.1145/3276505>
- [4] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Magnus Lång, and Tuan Phong Ngo. 2015. Precise and sound automatic fence insertion

³However, as Burnim et al. [23] observed (and as was verified in [44]), the declarative TSO model in [22] is broken (it mishandles internal reads-from edges), rendering Sober unsound.

- procedure under PSO. In *NETYS*. Springer International Publishing, Cham, 32–47.
- [5] Parosh Aziz Abdulla, Mohamed Faouzi Atig, and Tuan-Phong Ngo. 2015. The best of both worlds: Trading efficiency and optimality in fence insertion for TSO. In *ESOP*. Springer-Verlag New York, Inc., New York, 308–332. https://doi.org/10.1007/978-3-662-46669-8_13
- [6] Sarita V. Adve and Mark D. Hill. 1990. Weak ordering—a new definition. In *ISCA*. ACM, New York, 2–14. <https://doi.org/10.1145/325164.325100>
- [7] Mustaque Ahamad, Gil Neiger, James E. Burns, Prince Kohli, and Phillip W. Hutto. 1995. Causal memory: definitions, implementation, and programming. *Distributed Computing* 9, 1 (1995), 37–49.
- [8] Jade Alglave, Daniel Kroening, Vincent Nimal, and Daniel Poetzl. 2017. Don't sit on the fence: a static analysis approach to automatic fence insertion. *ACM Trans. Program. Lang. Syst.* 39, 2, Article 6 (May 2017), 38 pages. <https://doi.org/10.1145/2994593>
- [9] Jade Alglave and Luc Maranget. 2011. Stability in weak memory models. In *CAV*. Springer-Verlag, Berlin, Heidelberg, 50–66. <http://dl.acm.org/citation.cfm?id=2032305.2032311>
- [10] Jade Alglave, Luc Maranget, and Michael Tautschnig. 2014. Herding cats: modelling, simulation, testing, and data mining for weak memory. *ACM Trans. Program. Lang. Syst.* 36, 2, Article 7 (July 2014), 74 pages. <https://doi.org/10.1145/2627752>
- [11] Mohamed Faouzi Atig, Ahmed Bouajjani, Sebastian Burckhardt, and Madanlal Musuvathi. 2010. On the verification problem for weak memory models. In *POPL*. ACM, New York, 7–18. <https://doi.org/10.1145/1706299.1706303>
- [12] Mohamed Faouzi Atig, Ahmed Bouajjani, Sebastian Burckhardt, and Madanlal Musuvathi. 2012. What's decidable about weak memory models?. In *ESOP*. Springer-Verlag, Berlin, Heidelberg, 26–46. https://doi.org/10.1007/978-3-642-28869-2_2
- [13] Mark Batty, Kayvan Memarian, Kyndylan Nienhuis, Jean Pichon-Pharabod, and Peter Sewell. 2015. The problem of programming language concurrency semantics. In *ESOP*. Springer, Berlin, Heidelberg, 283–307. https://doi.org/10.1007/978-3-662-46669-8_12
- [14] Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber. 2011. Mathematizing C++ concurrency. In *POPL*. ACM, New York, 55–66. <https://doi.org/10.1145/1925844.1926394>
- [15] Giovanni Bernardi and Alexey Gotsman. 2016. Robustness against consistency models with atomic visibility. In *CONCUR*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 7:1–7:15. <https://doi.org/10.4230/LIPLcs.CONCUR.2016.7>
- [16] Hans-J. Boehm. 2012. Can Seqlocks get along with programming language memory models?. In *MSPC*. ACM, New York, 12–20. <https://doi.org/10.1145/2247684.2247688>
- [17] Ahmed Bouajjani, Egor Derevenetc, and Roland Meyer. 2013. Checking and enforcing robustness against TSO. In *ESOP*. Springer-Verlag, Berlin, Heidelberg, 533–553. https://doi.org/10.1007/978-3-642-37036-6_29
- [18] Ahmed Bouajjani, Constantin Enea, Suha Orhun Mutluergil, and Serdar Tasiran. 2018. Reasoning about TSO programs using reduction and abstraction. In *CAV*. Springer, Cham, 336–353.
- [19] Ahmed Bouajjani, Roland Meyer, and Eike Möhlmann. 2011. Deciding robustness against total store ordering. In *ICALP*. Springer, Berlin, Heidelberg, 428–440.
- [20] Lucas Brutschy, Dimitar Dimitrov, Peter Müller, and Martin Vechev. 2018. Static serializability analysis for causal consistency. In *PLDI*. ACM, New York, 90–104. <https://doi.org/10.1145/3192366.3192415>
- [21] Sebastian Burckhardt, Rajeev Alur, and Milo M. K. Martin. 2007. CheckFence: Checking consistency of concurrent data types on relaxed memory models. In *PLDI*. ACM, New York, 12–21. <https://doi.org/10.1145/1250734.1250737>
- [22] Sebastian Burckhardt and Madanlal Musuvathi. 2008. Effective program verification for relaxed memory models. In *CAV*. Springer-Verlag, Berlin, Heidelberg, 107–120. https://doi.org/10.1007/978-3-540-70545-1_12
- [23] Jacob Burnim, Koushik Sen, and Christos Stergiou. 2011. Sound and complete monitoring of sequential consistency for relaxed memory models. In *TACAS*. Springer, Berlin, Heidelberg, 11–25.
- [24] Egor Derevenetc. 2015. *Robustness against relaxed memory models*. Ph.D. Dissertation. University of Kaiserslautern. <http://kluedo.ub.uni-kl.de/frontdoor/index/index/docId/4074>
- [25] Egor Derevenetc and Roland Meyer. 2014. Robustness against Power is PSpace-complete. In *ICALP*. Springer, Berlin, Heidelberg, 158–170.
- [26] Mathieu Desnoyers, Paul E. McKenney, Alan S. Stern, Michel R. Dagenais, and Jonathan Walpole. 2012. User-level implementations of read-copy update. *IEEE Trans. Parallel Distrib. Syst.* 23, 2 (Feb. 2012), 375–382. <https://doi.org/10.1109/TPDS.2011.159>
- [27] Simon Doherty, Brijesh Dongol, Heike Wehrheim, and John Derrick. 2019. Verifying C11 programs operationally. In *PPoPP*. ACM, New York, 355–365. <https://doi.org/10.1145/3293883.3295702>
- [28] Alan Fekete, Dimitrios Liarokapis, Elizabeth O'Neil, Patrick O'Neil, and Dennis Shasha. 2005. Making snapshot isolation serializable. *ACM Trans. Database Syst.* 30, 2 (June 2005), 492–528. <https://doi.org/10.1145/1071610.1071615>
- [29] Kourosh Gharachorloo, Sarita V. Adve, Anoop Gupta, John L. Hennessy, and Mark D. Hill. 1992. Programming for different memory consistency models. *J. Parallel and Distrib. Comput.* 15, 4 (1992), 399 – 407. [https://doi.org/10.1016/0743-7315\(92\)90052-O](https://doi.org/10.1016/0743-7315(92)90052-O)
- [30] Alexey Gotsman, Madanlal Musuvathi, and Hongseok Yang. 2012. Show no weakness: sequentially consistent specifications of TSO libraries. In *DISC*. Springer-Verlag, Berlin, Heidelberg, 31–45. https://doi.org/10.1007/978-3-642-33651-5_3
- [31] Gerard J. Holzmann. 1997. The model checker SPIN. *IEEE Transactions on software engineering* 23, 5 (1997), 279–295.
- [32] SPARC International Inc. 1994. *The SPARC architecture manual (version 9)*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA.
- [33] Jan-Oliver Kaiser, Hoang-Hai Dang, Derek Dreyer, Ori Lahav, and Viktor Vafeiadis. 2017. Strong logic for weak memory: Reasoning about release-acquire consistency in Iris. In *ECOOP*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 17:1–17:29. <https://doi.org/10.4230/LIPLcs.ECOOP.2017.17>
- [34] Michalis Kokologiannakis, Ori Lahav, Konstantinos Sagonas, and Viktor Vafeiadis. 2017. Effective stateless model checking for C/C++ concurrency. *Proc. ACM Program. Lang.* 2, POPL, Article 17 (Dec. 2017), 32 pages. <https://doi.org/10.1145/3158105>
- [35] Dexter Kozen. 1977. Lower bounds for natural proof systems. In *SFCS*. IEEE Computer Society, Washington, 254–266. <https://doi.org/10.1109/SFCS.1977.16>
- [36] Ori Lahav, Nick Giannarakis, and Viktor Vafeiadis. 2016. Taming release-acquire consistency. In *POPL*. ACM, New York, 649–662. <https://doi.org/10.1145/2837614.2837643>
- [37] Ori Lahav and Viktor Vafeiadis. 2015. Owicki-Gries reasoning for weak memory models. In *ICALP*. Springer-Verlag, Berlin, Heidelberg, 311–323. https://doi.org/10.1007/978-3-662-47666-6_25
- [38] Ori Lahav and Viktor Vafeiadis. 2016. Explaining relaxed memory models with program transformations. In *FM*. Springer, Cham, 479–495. https://doi.org/10.1007/978-3-319-48989-6_29
- [39] Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, and Derek Dreyer. 2017. Repairing sequential consistency in C/C++11. In *PLDI*. ACM, New York, 618–632. <https://doi.org/10.1145/3062341.3062352>
- [40] Leslie Lamport. 1979. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Computers* 28, 9 (1979), 690–691.
- [41] Alexander Linden and Pierre Wolper. 2011. A verification-based approach to memory fence insertion in relaxed memory systems. In *SPIN*. Springer-Verlag, Berlin, Heidelberg, 144–160. <http://dl.acm.org/citation.cfm?id=2032692.2032707>
- [42] Alexander Linden and Pierre Wolper. 2013. A verification-based approach to memory fence insertion in PSO memory systems. In *TACAS*. Springer-Verlag, Berlin, Heidelberg, 339–353. https://doi.org/10.1007/978-3-540-70545-1_12

- 978-3-642-36742-7_24
- [43] Feng Liu, Nayden Nedev, Nedyalko Prasadnikov, Martin Vechev, and Eran Yahav. 2012. Dynamic synthesis for relaxed memory models. In *PLDI*. ACM, New York, 429–440. <https://doi.org/10.1145/2254064.2254115>
- [44] Sela Mador-Haim, Rajeev Alur, and Milo M K. Martin. 2010. Generating litmus tests for contrasting memory consistency models. In *CAV*. Springer-Verlag, Berlin, Heidelberg, 273–287. https://doi.org/10.1007/978-3-642-14295-6_26
- [45] Luc Maranget, Susmit Sarkar, and Peter Sewell. 2012. A tutorial introduction to the ARM and POWER relaxed memory models. <http://www.cl.cam.ac.uk/~pes20/ppc-supplemental/test7.pdf>.
- [46] Yuri Meshman, Noam Rinetzky, and Eran Yahav. 2015. Pattern-based synthesis of synchronization for the C++ memory model. In *FMCAD*. FMCAD Inc, Austin, TX, 120–127. <http://dl.acm.org/citation.cfm?id=2893529.2893552>
- [47] Kartik Nagar and Suresh Jagannathan. 2018. Automated detection of serializability violations under weak consistency. In *CONCUR 2018*, Vol. 118. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 41:1–41:18. <https://doi.org/10.4230/LIPIcs.CONCUR.2018.41>
- [48] Brian Norris and Brian Demsky. 2013. CDSchecker: checking concurrent data structures written with C/C++ atomics. In *OOPSLA*. ACM, New York, 131–150. <https://doi.org/10.1145/2509136.2509514>
- [49] Scott Owens. 2010. Reasoning about the implementation of concurrency abstractions on x86-TSO. In *ECOOP*. Springer-Verlag, Berlin, Heidelberg, 478–503.
- [50] Scott Owens, Susmit Sarkar, and Peter Sewell. 2009. A better x86 memory model: x86-TSO. In *TPHOLS*. Springer, Heidelberg, 391–407. https://doi.org/10.1007/978-3-642-03359-9_27
- [51] Dennis Shasha and Marc Snir. 1988. Efficient and correct execution of parallel programs that share memory. *ACM Trans. Program. Lang. Syst.* 10, 2 (April 1988), 282–312. <https://doi.org/10.1145/42190.42277>
- [52] Yair Sovran, Russell Power, Marcos K. Aguilera, and Jinyang Li. 2011. Transactional storage for geo-replicated systems. In *SOSP*. ACM, New York, 385–400. <https://doi.org/10.1145/2043556.2043592>
- [53] Kasper Svendsen, Jean Pichon-Pharabod, Marko Doko, Ori Lahav, and Viktor Vafeiadis. 2018. A separation logic for a promising semantics. In *ESOP*. Springer International Publishing, Cham, 357–384.
- [54] Aaron Turon, Viktor Vafeiadis, and Derek Dreyer. 2014. GPS: Navigating weak memory with ghosts, protocols, and separation. In *OOPSLA*. ACM, New York, 691–707. <https://doi.org/10.1145/2660193.2660243>
- [55] Viktor Vafeiadis and Chinmay Narayan. 2013. Relaxed separation logic: A program logic for C11 concurrency. In *OOPSLA*. ACM, New York, 867–884. <https://doi.org/10.1145/2509136.2509532>
- [56] John Wickerson, Mark Batty, Tyler Sorensen, and George A. Constantinides. 2017. Automatically comparing memory consistency models. In *POPL*. ACM, New York, 190–204. <https://doi.org/10.1145/3009837.3009838>
- [57] Anthony Williams. 2008. Peterson’s lock with C++0x atomics. Retrieved October 26, 2018 from https://www.justsoftwaresolutions.co.uk/threading/petersons_lock_with_C++0x_atomics.html