

Verifying Observational Robustness against a C11-Style Memory Model

ROY MARGALIT, Tel Aviv University, Israel

ORI LAHAV, Tel Aviv University, Israel

We study the problem of verifying the robustness of concurrent programs against a C11-style memory model that includes relaxed accesses and release/acquire accesses and fences, and show that this verification problem can be reduced to a standard reachability problem under sequential consistency. We further observe that existing robustness notions do not allow the verification of programs that use speculative reads as in the sequence lock mechanism, and introduce a novel “observational robustness” property that fills this gap. In turn, we show how to soundly check for observational robustness. We have implemented our method and applied it to several challenging concurrent algorithms, demonstrating the applicability of our approach. To the best of our knowledge, this is the first method for verifying robustness against a programming language concurrency model that includes relaxed accesses and release/acquire fences.

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

Additional Key Words and Phrases: weak memory models, C/C++11, robustness, shared-memory concurrency

ACM Reference Format:

Roy Margalit and Ori Lahav. 2021. Verifying Observational Robustness against a C11-Style Memory Model. *Proc. ACM Program. Lang.* 5, POPL, Article 4 (January 2021), 33 pages. <https://doi.org/10.1145/3434285>

1 INTRODUCTION

Programming and reasoning about concurrent programs under memory models weaker than sequential consistency (SC) is nutritiously hard and error-prone. When programmers manage to avoid data races, the DRF (data-race-freedom) guarantee allows them to safely imagine a heaven of SC multiprocessors and SC-preserving compilers [Adve and Hill 1990]. For a variety of programs, however, avoiding races altogether is too restrictive and costly. Interestingly, often this does not mean that the program may behave under a weak memory model differently than it would behave under SC. Accordingly, there is a need to develop *robustness* criteria (i.e., conditions that ensure that a given program has only SC behaviors), which are more precise than DRF, and accompany them with robustness verification methods and tools. Safety verification under a weak memory model can be then reduced to robustness verification plus safety verification assuming SC [Burckhardt and Musuvathi 2008].

Precise robustness criteria and robustness verification tools were previously developed for hardware models (x86-TSO, see, e.g., [Bouajjani et al. 2013, 2011; Owens 2010]), and POWER [Derevenec and Meyer 2014]), and recently for the release/acquire fragment (RA) of C11 [Lahav and Margalit 2019]. In this work we are interested in a larger fragment of C11, which also includes relaxed

Authors' addresses: Roy Margalit, Tel Aviv University, Israel, roy.margalit@cs.tau.ac.il; Ori Lahav, Tel Aviv University, Israel, orilahav@tau.ac.il.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2021 Copyright held by the owner/author(s).

2475-1421/2021/1-ART4

<https://doi.org/10.1145/3434285>

accesses and release/acquire fences. These provide more fine-grained control on the required synchronization, which allows for better performance. The price to pay is that we get much farther than SC: relaxed accesses do not maintain basic causality as they can be reordered¹ (see, e.g., Ex. 4.2 below); release and acquire fences separate the synchronization points from the actual writes and reads [Batty et al. 2011]; and *release sequences*, a special mechanism for using relaxed read-modify-write instructions (RMWs) for synchronization, are hard to reason about [Doko and Vafeiadis 2017]. This makes programming more difficult and error-prone (see, e.g., the case of Peterson’s lock described in §7), which, in fact, may (and, we believe, often does) cause programmers to refrain from using relaxed accesses altogether.

In this work, we extend the results of Lahav and Margalit [2019] to a larger fragment of C11, which also includes relaxed reads/writes/RMWs and release/acquire fences. We show that robustness against this model is reduced to a reachability problem under an instrumented SC semantics.² Roughly speaking, while the generated execution histories (a.k.a. *execution graphs*) are unbounded (since programs may have loops), we show that a finite collection of (rather intricate) properties of the generated execution history suffices for (sound and precise) monitoring of states where threads may observe (either by reading or by overwriting) values under the weak semantics that cannot be observed under SC. Based on this reduction, we obtain a decision procedure for robustness, which we have also implemented and experimented with, using SPIN [Holzmann 1997] for the reachability analysis under SC. In particular, it follows that for programs with bounded data domain robustness verification against the model studied here is PSPACE-complete (just like robustness verification against x86-TSO or RA).

The factors that make programming with relaxed accesses to be difficult also make robustness verification for relaxed accesses to be technically challenging. Indeed, the instrumented semantics that we develop requires much finer distinctions maintaining multiple views for each thread, which in turn allow us to compare the observable values under SC to those observable in the relaxed semantics. Nevertheless, while the instrumentation is heavier than the one in [Lahav and Margalit 2019], our experiments show that this has only a minor effect on the practical verification times.

While handling relaxed accesses, we have identified a drawback of all (to the best of our knowledge) existing robustness criteria, which becomes more important for robustness against a semantics that includes relaxed accesses. It is related to the question of what exactly constitutes a program behavior (and, in turn, what is a robustness violation?). The simplest answer identifies program behaviors with sets of reachable program states, where program states consist of the values of the different program counters and local variables. However, Derevenetc [2015] observed that, while desirable, solving the robustness verification problem for a state-based robustness notion is as hard as full-fledged verification under the weak semantics. The latter can be extremely difficult from a computational complexity point of view (non-primitive recursive for TSO [Atig et al. 2010, 2012] and undecidable for RA [Abdulla et al. 2019]). Thus, researchers have resorted to stronger robustness definitions, in which program behaviors are taken to be sets of “annotated” histories of the executions of the programs (called *traces* or *execution graphs*), and robustness requires that all such histories generated by the program under the weak semantics can be also generated under SC.

However, we observe here that both state-based and history-based robustness notions forbid the benign programming idiom of *speculative* reads as used in sequence lock algorithms. These algorithms used in the Linux kernel were identified as a main use of relaxed accesses [Sinclair

¹Like all verification works, we are unable to work with the original C11 model, which allows unrestricted cycles in the union of “program order” and “reads-from”, and thus exhibits “out-of-thin-air” values and even fails to provide the most basic DRF guarantee. We follow [Boehm and Demsky 2014; Lahav et al. 2017] and disallow such cycles, which invalidates store-after-load reordering. Section 3 discusses further technical differences between the model used here and C11.

²Equivalently, as we do in our implementation, one can think about reachability under SC of an instrumented program.

et al. 2017], and considered as probably “the most challenging of common programming idioms to implement with modern programming language memory models” [Boehm 2012]. Roughly speaking, they are locking mechanisms in which relaxed accesses are used to efficiently speculatively read some shared data, requiring a validation mechanism (implemented with stronger accesses or fences) before the read values are being used. If validation fails, the speculations are disposed and retried. Observing stale speculative values under relaxed semantics (which cannot be observed under SC) may be perfectly fine, but it does imply a reachable intermediate program state that is not reachable under SC, resulting in a robustness violation. Nevertheless, if the validation mechanism can be proved to throw away stale values before they are actually being used, this mechanism should be considered as a benign temporary robustness violation that can never harm the program’s safety.

To support verification of such mechanisms, we propose a novel notion of robustness, which we call *observational robustness*. It is a history-based notion (to make it computationally attainable), but it allows the program to have some non-SC histories provided that they can be “turned” into SC histories without affecting any value that the program used. To do so, we include a dependency relation (akin to the one used in modern hardware models, see, e.g., [Flur et al. 2016; Podkopaev et al. 2019]) in the program’s execution graphs, and allow observationally robust programs to read stale values (violating SC) as long as no further operation depends on these values.

As we did for usual robustness, we show that observational robustness can be verified by inspecting *only* SC executions of the program, which, again, allows the reduction to reachability in an instrumented SC semantics. These results require certain technical novelties as we have to carefully generate an SC execution from any execution representing an observational robustness violation (that could have diverged from SC early on), as well as to devise a specialized taint tracking mechanism for identifying dependencies on stale values. We note that our observational robustness verification is sound but, unlike our result for usual (history-based) robustness, imprecise. Still, we show that it verifies the challenging examples of [Boehm 2012] deeming them as robust and safe. We leave the problem of developing a more precise analysis (as well as identifying the robustness notion that our analysis is precise for) to future work.

To conclude, the main contributions of this paper are summarized as follows:

- (1) We develop a sound and precise robustness verification method (by reduction to reachability in an instrumented SC semantics) for a C11-style model that includes relaxed accesses and release/acquire fences.
- (2) We introduce a notion of *observational robustness* that allows one to reason about sequence-locks-like programs using speculative reads, and provide a sound verification method for this robustness notion.
- (3) We provide an implementation of our approach and present its application on multiple challenging examples.

Outline. The rest of this paper is organized as follows. In §2 we present the programming language that we assume in this paper. In §3 we present the C11-style memory model that we study in this paper and discuss its relation to (R)C11. In §4 we present our robustness notions and the key theorems that enable their verification. In §5 we present the reduction of robustness verification to a reachability problem under SC. In §7 we discuss the implementation and our experiments with it. In §8 we present related work and conclude.

Additional Material. Proofs of the theorems in the paper are given in the its accompanying technical appendix available at [Margalit and Lahav 2020]. The prototype implementation and the examples it was tested on are available in the artifact accompanying this paper (available at <https://dl.acm.org/doi/10.1145/3410267/full/>).

$v, u \in \text{Val}$	Values	$o_X \in \text{Mod}_X$ where $X \in \{R, W, \text{RMW}, F\}$	Access modes
$x, y, z \in \text{Loc} \subseteq \{x, y, \dots\}$	Locations	$sPr \in \text{SProg} \triangleq \{0, 1, \dots, N\} \rightarrow \text{Inst}$	Sequential programs
$r \in \text{Reg} \subseteq \{a, b, \dots\}$	Registers	$Pr : \text{Tid} \rightarrow \text{SProg}$	(Concurrent) programs
$\tau, \pi \in \text{Tid} \subseteq \{T_1, T_2, \dots\}$	Thread identifiers		

$$e ::= r \mid v \mid e + e \mid e = e \mid e \neq e \mid \dots$$

$$\text{Inst} \ni \text{inst} ::= r := e \mid \text{if } e \text{ goto } pc_1, \dots, pc_n \mid \text{assert}(e)$$

$$\mid x.\text{store}(e, o_W) \mid r := x.\text{load}(o_R)$$

$$\mid r := x.\text{FADD}(e, o_{\text{RMW}}) \mid r := x.\text{CAS}(e, e, o_{\text{RMW}}, o_R) \mid \text{fence}(o_F)$$

$$\mid \text{wait}(x = e, o_R) \mid x.\text{BCAS}(e, e, o_{\text{RMW}})$$

Fig. 1. Domains and programming language syntax.

2 CONCURRENT PROGRAMMING LANGUAGE

In this section we formulate the syntax of our C-style toy programming language (§2.1), and present the interpretation of concurrent programs as labeled transition systems synchronized with general memory systems (§2.2). We start with several preliminary conventions and notations.

Relations. Given a (binary) relation R , $\text{dom}(R)$ and $\text{codom}(R)$ denote its domain and codomain, and R^2 , 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, $[A] ; R ; [B] = R \cap (A \times B)$.

Labeled Transition Systems. A labeled transition system (LTS) A is a tuple $\langle Q, \Sigma, q_0, \rightarrow \rangle$, where Q is a set of states, Σ is an alphabet, $q_0 \in Q$ 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 $A.Q$, $A.\Sigma$, $A.q_0$, and \rightarrow_A the 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 $\sigma \in A.\Sigma$ is *enabled* in $q \in A.Q$ (alternatively, q *enables* σ) if $q \xrightarrow_A \sigma q'$ for some $q' \in A.Q$. For a finite sequence $tr \in A.\Sigma^*$ (i.e., a function from $\{1, \dots, N\}$ to $A.\Sigma$ for some $N \geq 0$), we write \xrightarrow_A^{tr} for the composition $\xrightarrow_A^{tr(1)} ; \dots ; \xrightarrow_A^{tr(|tr|)}$. A sequence $tr \in A.\Sigma^*$ such that $A.q_0 \xrightarrow_A^{tr} q$ for some $q \in A.Q$ is called a *trace* of A .

2.1 Language Syntax

We assume *finite* sets of values, (shared) memory locations, register names, and thread identifiers: $\text{Val} \subseteq \mathbb{N}$, $\text{Loc} \subseteq \{x, y, \dots\}$, $\text{Reg} \subseteq \{a, b, \dots\}$, and $\text{Tid} \subseteq \{T_1, T_2, \dots\}$. We assume that Val contains a distinguished value 0, used as the initial value for all locations. Memory instructions in our language employ *access modes*, taken from the set $\text{Mod} \triangleq \{\text{rlx}, \text{rel}, \text{acq}, \text{acqrel}\}$, which determine their “consistency level”.³ Its following subsets provide the possible read, write, RMW, and fence modes respectively:

$$\begin{aligned} \text{Mod}_R &\triangleq \{\text{rlx}, \text{acq}\} & \text{Mod}_{\text{RMW}} &\triangleq \{\text{rlx}, \text{acq}, \text{rel}, \text{acqrel}\} \\ \text{Mod}_W &\triangleq \{\text{rlx}, \text{rel}\} & \text{Mod}_F &\triangleq \{\text{acq}, \text{rel}, \text{acqrel}\} \end{aligned}$$

A partial order \sqsubseteq on Mod , which intuitively orders access modes according to their strength, is defined by (the reflexive transitive closure of):

$$\text{rlx} \sqsubseteq \text{acq} \sqsubseteq \text{acqrel} \quad \text{and} \quad \text{rlx} \sqsubseteq \text{rel} \sqsubseteq \text{acqrel}$$

³In §6 we extend our results to support C11-style *non-atomic accesses* with *na* access mode.

Figure 1 presents our toy programming language. Its expressions are constructed from registers (local variables) and values. Instructions include assignments, memory accesses and release/acquire fences. SC-fences are not taken as primitives in this language but rather included as a syntactic sugar; see §3.1 and Remark 1. We include non-deterministic conditional jumps instructions (if e goto pc_1, \dots, pc_n jumps to some program counter pc_k if e is not 0) from which usual conditionals and loops can be implemented. Assertions can be used for safety verification (assert(e) fails if e is 0). Memory accesses consist of loads, stores, and RMWs, which are either (never failing) fetch-and-adds (FADD) or (possibly failing) compare-and-swaps (CAS). Both RMW operations return the value that was read before its update. The language also includes two kinds of blocking instructions: wait($x = e, o_R$) blocks the execution of the thread until it can read the value of e from x (with mode $o_R \in \text{Mod}_R$); and $x.\text{BCAS}(e, e, o_{\text{RMW}})$ is an (always successful) blocking CAS. The blocking instructions are implementable using busy loops, but, as observed in [Lahav and Margalit 2019], they make the robustness analysis more expressive as they allow one to mask benign robustness violations: it often happens that stale values can be read in the busy loop, while the use of the blocking instructions instead of the busy loops ensures robustness.

Sequential and concurrent programs are defined as follows.

Definition 2.1. A sequential program sPr is a function from a set of the form $\{0, 1, \dots, N\}$ (the possible values of the program counter) to instructions. We denote by $S\text{Prog}$ the set of all sequential programs. A concurrent program Pr is a mapping from Tid to $S\text{Prog}$. For simplicity, we assume that the sets of registers mentioned in the components of concurrent programs are pairwise disjoint.

Concurrent programs (which we also shortly call *programs*) are top-level parallel compositions of sequential programs. In our examples, we often write sequential programs as sequences of instructions delimited by “;” or line breaks, and use ‘||’ for parallel composition.

2.2 From Programs to Labeled Transition Systems

Next, we show how programs are read as labeled transition systems. At this stage, the memory system, responsible of determining which values can be read at each point, is left parametric. To define the alphabet of these LTSs we use *event labels* and *sequential transition labels*:

Definition 2.2. An event label l takes one of the following forms: $R(o_R, x, v_R)$ (read label), $W(o_W, x, v_W)$ (write label), $\text{RMW}(o_{\text{RMW}}, x, v_R, v_W)$ (RMW label), $R^*(o_R, x, v_R)$ (wait/failed CAS label), $F(o_F)$ (fence label), or CTRL (control label), where $x \in \text{Loc}$, $v_R, v_W \in \text{Val}$, and $o_X \in \text{Mod}_X$ (for $X \in \{R, W, \text{RMW}, F\}$). The functions typ , loc , val_R , val_W , and mod return (when applicable) the type ($R/W/\text{RMW}/R^*/F/\text{CTRL}$), location (x), read value (v_R), written value (v_W), and access mode (o_X) of a given event label. We denote by Lab the set of all event labels.

Definition 2.3. A sequential transition label L is a triple of the form $\langle l_\varepsilon, R_{\text{in}}, R_{\text{out}} \rangle$, where $l_\varepsilon \in \text{Lab} \cup \{\varepsilon\}$ and $R_{\text{in}}, R_{\text{out}} \subseteq \text{Reg}$. The functions lab , R_{in} , and R_{out} return the components of a given sequential transition label. The functions typ , loc , val_R , val_W , and mod are lifted to sequential transition labels in the obvious way (undefined when $\text{lab}(L) = \varepsilon$). We denote by TLab the set of all sequential transition labels.

Event labels record the memory operation performed at each execution step. Distinguishing normal reads and R^* -reads which arise from wait and (failing) CAS instructions is instrumental in our observational robustness notion (we will not allow “speculative” R^* -reads). In turn, sequential transition labels also include silent program steps (with $\text{lab}(L) = \varepsilon$), and further record the input registers R_{in} (registers that were read in each step) and the output registers R_{out} (registers that were written to in each step). Together with the fact that we have explicit CTRL-labels, this allows us

$$\begin{array}{c}
\frac{sPr(pc) = r := e \quad \phi' = \phi[r \mapsto \phi(e)] \quad R_{in} = \text{Reg}(e) \quad R_{out} = \{r\}}{\langle pc, \phi \rangle \xrightarrow{\varepsilon, R_{in}, R_{out}}_{sPr} \langle pc + 1, \phi' \rangle} \\
\frac{sPr(pc) = \text{if } e \text{ goto } pc_1, \dots, pc_n \quad \phi(e) \neq 0 \implies pc' \in \{pc_1, \dots, pc_n\} \quad \phi(e) = 0 \implies pc' = pc + 1 \quad l = \text{CTRL} \quad R_{in} = \text{Reg}(e)}{\langle pc, \phi \rangle \xrightarrow{L_{R_{in}}, 0}_{sPr} \langle pc', \phi \rangle} \\
\frac{sPr(pc) = \text{assert}(e) \quad \phi(e) \neq 0 \implies pc' = pc' + 1 \quad \phi(e) = 0 \implies pc' = \perp \quad l = \text{CTRL} \quad R_{in} = \text{Reg}(e)}{\langle pc, \phi \rangle \xrightarrow{L_{R_{in}}, 0}_{sPr} \langle pc', \phi \rangle} \\
\frac{sPr(pc) = x.\text{store}(e, o_W) \quad l = W(o_W, x, \phi(e)) \quad R_{in} = \text{Reg}(e)}{\langle pc, \phi \rangle \xrightarrow{L_{R_{in}}, 0}_{sPr} \langle pc + 1, \phi \rangle} \\
\frac{sPr(pc) = r := x.\text{load}(o_R) \quad l = R(o_R, x, v) \quad \phi' = \phi[r \mapsto v] \quad R_{out} = \{r\}}{\langle pc, \phi \rangle \xrightarrow{L, 0, R_{out}}_{sPr} \langle pc + 1, \phi' \rangle} \\
\frac{sPr(pc) = r := x.\text{FADD}(e, o_{RMW}) \quad l = \text{RMW}(o_{RMW}, x, v, v + \phi(e)) \quad \phi' = \phi[r \mapsto v] \quad R_{in} = \text{Reg}(e) \quad R_{out} = \{r\}}{\langle pc, \phi \rangle \xrightarrow{L_{R_{in}}, R_{out}}_{sPr} \langle pc + 1, \phi' \rangle} \\
\frac{sPr(pc) = r := x.\text{CAS}(e_R, e_W, o_{RMW}, o_R) \quad l = \text{RMW}(o_{RMW}, x, \phi(e_R), \phi(e_W)) \quad \phi' = \phi[r \mapsto \phi(e_R)] \quad R_{in} = \text{Reg}(e_R) \cup \text{Reg}(e_W) \quad R_{out} = \{r\}}{\langle pc, \phi \rangle \xrightarrow{L_{R_{in}}, R_{out}}_{sPr} \langle pc + 1, \phi' \rangle} \\
\frac{sPr(pc) = r := x.\text{CAS}(e_R, e_W, o_{RMW}, o_R) \quad l = R^*(o_R, x, v) \quad v \neq \phi(e_R) \quad \phi' = \phi[r \mapsto v] \quad R_{in} = \text{Reg}(e_R) \quad R_{out} = \{r\}}{\langle pc, \phi \rangle \xrightarrow{L_{R_{in}}, R_{out}}_{sPr} \langle pc + 1, \phi' \rangle} \\
\frac{sPr(pc) = \text{fence}(o_F) \quad l = F(o_F)}{\langle pc, \phi \rangle \xrightarrow{L, 0, 0}_{sPr} \langle pc + 1, \phi \rangle} \\
\frac{sPr(pc) = \text{wait}(x = e, o_R) \quad l = R^*(o_R, x, \phi(e)) \quad R_{in} = \text{Reg}(e)}{\langle pc, \phi \rangle \xrightarrow{L_{R_{in}}, 0}_{sPr} \langle pc + 1, \phi \rangle} \\
\frac{sPr(pc) = x.\text{BCAS}(e_R, e_W, o_{RMW}) \quad l = \text{RMW}(o_{RMW}, x, \phi(e_R), \phi(e_W)) \quad R_{in} = \text{Reg}(e_R) \cup \text{Reg}(e_W)}{\langle pc, \phi \rangle \xrightarrow{L_{R_{in}}, 0}_{sPr} \langle pc + 1, \phi \rangle}
\end{array}$$

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

to observe dependencies between actions in the LTS induced by a program, with no need to refer back to the code itself. Note that for the language constructs introduced above, R_{out} will always be either empty or a singleton.

The next definition formally associates sequential programs with LTSs. In the sequel we identify sequential programs with their induced LTSs (when writing, e.g., $sPr.Q$ and \rightarrow_{sPr}).

Definition 2.4. The LTS (over the alphabet TLab) induced by sPr is defined as follows:

- Its states are pairs $q = \langle pc, \phi \rangle$, where $pc \in \{0, \dots, N\} \cup \{\perp\}$ is the *program counter* (where $N \geq 0$ and \perp denotes an “error state”) and $\phi : \text{Reg} \rightarrow \text{Val}$ is the *local store* assigning values to registers. We assume that ϕ is extended to apply on expressions e in a standard way that reflects the (unspecified) expression semantics.
- The initial state is $\langle 0, \lambda r. 0 \rangle$.
- The transitions are given in Fig. 2 (providing formal meaning to the descriptions above).

Example 2.5. We present the LTS induced by a simple sequential program. Let $\text{Val} = \{0, \dots, 4\}$, $\text{Loc} = \{x\}$ and $\text{Reg} = \{a\}$. We evaluate expressions of the form $r = e$ to be 1 if $\phi(r) = \phi(e)$ and 0 otherwise.

$$\begin{array}{l}
sPr \triangleq \\
0 : x.\text{store}(1, \text{rel}) \\
1 : a := x.\text{load}(\text{acq}) \\
2 : \text{assert}(a = 1) \\
sPr.Q = \{0, 1, 2, 3, \perp\} \times \{[a \mapsto v] \mid v \in \text{Val}\} \quad sPr.Q_0 = \langle 0, [a \mapsto 0] \rangle \\
\rightarrow_{sPr} = \{ \langle 0, [a \mapsto v] \rangle \xrightarrow{W(\text{rel}, x, 1), 0, 0}_{sPr} \langle 1, [a \mapsto v] \rangle \mid v \in \text{Val} \} \cup \\
\{ \langle 1, [a \mapsto v] \rangle \xrightarrow{R(\text{acq}, x, u), 0, \{a\}}_{sPr} \langle 2, [a \mapsto u] \rangle \mid v, u \in \text{Val} \} \cup \\
\{ \langle 2, [a \mapsto 1] \rangle \xrightarrow{\text{CTRL}, \{a\}, 0}_{sPr} \langle 3, [a \mapsto 1] \rangle \} \cup \\
\{ \langle 2, [a \mapsto v] \rangle \xrightarrow{\text{CTRL}, \{a\}, 0}_{sPr} \langle \perp, [a \mapsto v] \rangle \mid v \in \text{Val} \setminus \{1\} \}
\end{array}$$

The last definition is straightforwardly lifted to the concurrent setting. (Again, we will identify concurrent programs with their induced LTSs.)

Definition 2.6. The LTS induced by a (concurrent) program Pr is the LTS over the alphabet $\text{Tid} \times \text{TLab}$ whose states are tuples $\bar{q} \in \prod_{\tau \in \text{Tid}} Pr(\tau).Q$; its initial state is $\lambda\tau. Pr(\tau).q_0$; and its transitions are interleaved transitions of Pr 's components, formally given by:

$$\frac{\bar{q}(\tau) \xrightarrow{L}_{Pr(\tau)} q'}{\bar{q} \xrightarrow{\tau, L}_{Pr} \bar{q}[\tau \mapsto q']}$$

On the level of programs, the values read by loads and RMWs are completely arbitrary. They are constrained by synchronizing a program with a *memory system*, and obtaining a *concurrent system*, as defined next.

Definition 2.7. A *memory system* \mathcal{M} is an LTS over the alphabet $\text{Tid} \times \text{TLab}$. The concurrent system induced by a program Pr and a memory system \mathcal{M} , denoted by $Pr \parallel \mathcal{M}$, is the LTS over the alphabet $\text{Tid} \times \text{TLab}$ whose states are pairs in $Pr.Q \times \mathcal{M}.Q$; its initial state is $\langle Pr.q_0, \mathcal{M}.q_0 \rangle$; and its transitions are synchronized transitions of Pr and \mathcal{M} , formally given by:

$$\frac{\bar{q} \xrightarrow{\tau, L}_{Pr} \bar{q}' \quad q_M \xrightarrow{\tau, L}_M q'_M}{\langle \bar{q}, q_M \rangle \xrightarrow{\tau, L}_{Pr \parallel \mathcal{M}} \langle \bar{q}', q'_M \rangle}$$

The most well-known memory system is the one of sequential consistency, denoted here by SC. This memory system simply tracks the most recent value written to each location. Formally, it is defined as follows.

Definition 2.8. The memory system SC is given by: $\text{SC}.Q \triangleq \text{Loc} \rightarrow \text{Val}$, $\text{SC}.q_0 \triangleq \lambda x \in \text{Loc}. 0$ and \rightarrow_{SC} is given by:

$$\frac{\text{typ}(L) \in \{R, R^*\} \quad M(\text{loc}(L)) = \text{val}_R(L)}{M \xrightarrow{\tau, L}_{\text{SC}} M} \quad \frac{\text{typ}(L) = W \quad M' = M[\text{loc}(L) \mapsto \text{val}_W(L)]}{M \xrightarrow{\tau, L}_{\text{SC}} M'} \quad \frac{\text{typ}(L) = \text{RMW} \quad M(\text{loc}(L)) = \text{val}_R(L) \quad M' = M[\text{loc}(L) \mapsto \text{val}_W(L)]}{M \xrightarrow{\tau, L}_{\text{SC}} M'} \quad \frac{\text{typ}(L) \notin \{R, W, \text{RMW}, R^*\}}{M \xrightarrow{\tau, L}_{\text{SC}} M}$$

Finally, we can define what it means for a program to fail under a memory system \mathcal{M} .

Definition 2.9. A state $\bar{q} \in Pr.Q$ is *failing* if $\bar{q}(\tau) = \langle \perp, \phi \rangle$ for some $\tau \in \text{Tid}$ and local store ϕ .

Definition 2.10. A program Pr *may fail under a memory system* \mathcal{M} if $\langle \bar{q}, q_M \rangle$ is reachable in $Pr \parallel \mathcal{M}$ for some failing state $\bar{q} \in Pr.Q$ and $q_M \in \mathcal{M}.Q$.

Our assumption that the program is finite-data (in particular, we assume that Val is finite) ensures that the question whether a given program may fail under SC is decidable (it is PSPACE-complete [Kozen 1977]). The goal of robustness analysis against a memory system \mathcal{M} (weaker than SC) is to reduce the question of whether a given program Pr may fail under \mathcal{M} to the question of whether Pr may fail under SC.

3 DECLARATIVE CONCURRENCY SEMANTICS USING EXECUTION GRAPHS

In this section we formulate the C11-style memory model, which we call RC20, that we consider in this paper. Like C11 [Batty et al. 2011] and its rectified version RC11 [Lahav et al. 2017], RC20 is defined declaratively—it identifies program behaviors with partially ordered program traces, called *execution graphs*, that meet certain consistency constraints. In §3.1 we formally define execution graphs and present RC20's consistency notion. In §3.2, we formulate the correspondence between programs and sets of execution graphs.

3.1 Execution Graphs and RC20-Consistency

We define execution graphs, starting from their ingredients: an underlying set of *events*, a *labeling function* assigning an event label to each event, a *dependency relation* between events in the same thread, a *reads-from* mapping determining the write event from which each read event reads its value, and a *modification order* totally ordering the writes to each location.

Definition 3.1. An event e is a pair $\langle \tau, sn \rangle$, where $\tau \in \text{Tid} \uplus \{\perp\}$ is a thread identifier (\perp is used for initialization events) and $sn \in \mathbb{N}$ is a serial numbers. The functions tid and sn return the thread identifier and serial number of an event. We denote the set of all events by Event . Given a set $E \subseteq \text{Event}$ and $\tau \in \text{Tid}$, we write E^τ for $\{e \in E \mid \text{tid}(e) = \tau\}$.

Definition 3.2. Let E be a set of events and $\text{lab} : E \rightarrow \text{Lab}$ be a labeling function (Lab is defined in [Def. 2.2](#)).

- A relation dp is a *dependency relation* for E and lab if $\text{dp} \subseteq E \times E$ and $\langle e_1, e_2 \rangle \in \text{dp}$ implies $\text{tid}(e_1) = \text{tid}(e_2)$, $\text{sn}(e_1) < \text{sn}(e_2)$, and $\text{typ}(\text{lab}(e_1)) \in \{\text{R}, \text{RMW}, \text{R}^*\}$.
- A relation rf is a *reads-from* relation for E and lab if the following hold:
 - If $\langle w, r \rangle \in \text{rf}$, then $w, r \in E$, $\text{typ}(\text{lab}(w)) \in \{\text{W}, \text{RMW}\}$, $\text{typ}(\text{lab}(r)) \in \{\text{R}, \text{RMW}, \text{R}^*\}$, $\text{loc}(\text{lab}(w)) = \text{loc}(\text{lab}(r))$, and $\text{val}_w(\text{lab}(w)) = \text{val}_r(\text{lab}(r))$.
 - If $\langle w_1, r \rangle, \langle w_2, r \rangle \in \text{rf}$, then $w_1 = w_2$ (that is, rf^{-1} is functional).
 - For every $r \in \{e \in E \mid \text{typ}(\text{lab}(e)) \in \{\text{R}, \text{RMW}, \text{R}^*\}\}$, we have $\langle w, r \rangle \in \text{rf}$ for some $w \in E$.
- A relation mo is a *modification order* for E and lab if it is a disjoint union of relations $\{\text{mo}_x\}_{x \in \text{Loc}}$, where each mo_x is a strict total order on

$$\{w \in E \mid \text{typ}(\text{lab}(w)) \in \{\text{W}, \text{RMW}\} \wedge \text{loc}(\text{lab}(w)) = x\}.$$

Definition 3.3. An *execution graph* is a tuple $G = \langle E, \text{lab}, \text{dp}, \text{rf}, \text{mo} \rangle$ where E is a finite set of events, $\text{lab} : E \rightarrow \text{Lab}$ is a labeling function, and dp , rf , and mo are respectively a dependency relation, a reads-from relation, and a modification order for E and lab . We denote by EGraph the set of all execution graphs. The components of G are denoted by $G.E$, $G.\text{lab}$, $G.\text{dp}$, $G.\text{rf}$, and $G.\text{mo}$. We write $G.\text{typ}$ for $\lambda e \in G.E. \text{typ}(G.\text{lab}(e))$ and use similar notations for loc , val_R , val_W , and mod . We use the following notations for restricted sets of events:

$$G.R \triangleq \{e \in G.E \mid G.\text{typ}(e) \in \{\text{R}, \text{RMW}, \text{R}^*\}\} \quad G.W \triangleq \{e \in G.E \mid G.\text{typ}(e) \in \{\text{W}, \text{RMW}\}\}$$

$$\text{For } \alpha \in \{\text{RMW}, \text{F}, \text{R}^*, \text{CTRL}\} : \quad G.\alpha \triangleq \{e \in G.E \mid G.\text{typ}(e) = \alpha\}$$

We employ subscripts and superscripts to restrict sets of events to certain properties (e.g., $G.W_x^{\text{rel1}} = \{w \in G.W \mid G.\text{loc}(w) = x \wedge G.\text{mod}(w) \supseteq \text{rel1}\}$).

We assume that execution graphs are always initialized, as defined next.

Definition 3.4. The set of Init initialization events consists of an event w_x^{Init} for every $x \in \text{Loc}$ with $\text{tid}(w_x^{\text{Init}}) = \perp$ and $\text{sn}(w_x^{\text{Init}}) = 0$. The initial labeling function lab^{Init} is the function assigning the label $W(\text{rlx}, x, 0)$ to each every initialization event $w_x^{\text{Init}} \in \text{Init}$. The initial execution graph G_\emptyset is given by $G_\emptyset \triangleq \langle \text{Init}, \text{lab}^{\text{Init}}, \emptyset, \emptyset, \emptyset \rangle$. An execution graph G is *initialized* if it extends G_\emptyset (that is: $\text{Init} \subseteq G.E$ and $\text{lab}^{\text{Init}} \subseteq G.\text{lab}$).

Our representation of events induces a *program order*, in which events of the same thread are ordered according to their serial numbers and initialization events precede all other events (i.e., $\langle \tau_1, sn_1 \rangle < \langle \tau_2, sn_2 \rangle$ if $\tau_1 = \tau_2 \neq \perp$ and $sn_1 < sn_2$, or $\tau_1 = \perp$ and $\tau_2 \neq \perp$). We denote by $G.\text{po}$ the restriction of this order to $G.E$ (i.e., $G.\text{po} \triangleq \{\langle e_1, e_2 \rangle \in G.E \times G.E \mid e_1 < e_2\}$).

To define RC20-consistent execution graphs, we need several additional derived relations. First, we define the *synchronizes with* relation and the *happens-before* relation:

$$G.sw \triangleq [G.E^{\exists rel^1}] ; ([G.F] ; G.po)^? ; G.rf^+ ; (G.po ; [G.F])^? ; [G.E^{\exists acq}] \quad (\textit{synchronizes with})$$

$$G.hb \triangleq (G.po \cup G.sw)^+ \quad (\textit{happens-before})$$

Happens-before consists of all paths formed of program order edges (*po*) and synchronization edges (*sw*). The latter contains *rf*-edges between release writes and acquire reads. Synchronization is also induced by:

- (1) *rf*-edges from relaxed writes that follow release fences;
- (2) *rf*-edges to relaxed reads that are followed by acquire fences; and
- (3) *release sequences* (using C11's terminology)—sequences of *rf*-edges, where all internal events in the sequence are RMWs (necessarily, as they both read and write).

We note that (R)C11's synchronization also includes $[G.W^{\exists rel^1}] ; (G.po \cap G.mo) ; G.rf^+ ; (G.po ; [G.F])^? ; [G.E^{\exists acq}]$. To simplify the model, observing that this extension is rarely used, it was recently discarded from the C/C++ memory model [Boehm et al. 2018]. Accordingly, we use here the simpler version of *sw*.

Second, we use the standard *from-read* relation (a.k.a. *reads-before*), which relates reads with subsequent writes to the same location (write events that are *mo*-later than the write event that the read event reads its value from):

$$G.fr \triangleq (G.rf^{-1} ; G.mo) \setminus [G.E] \quad (\textit{from-read})$$


Now, following (R)C11, RC20-consistency is defined as the conjunction of four constraints:

Definition 3.5. An execution graph G is RC20-consistent if the following hold:

- $G.mo ; G.rf^? ; G.hb^?$ is irreflexive. (WRITE COHERENCE)
- $G.fr ; G.rf^? ; G.hb$ is irreflexive. (READ COHERENCE)
- $G.fr ; G.mo$ is irreflexive. (ATOMICITY)
- $G.po \cup G.rf$ is acyclic. (PO-RF)

The first two constraints ensure that the per-location order ($mo \cup fr$) ; $rf^?$ agrees with *hb*. The third constraint guarantees the atomicity of RMWs (in particular, it implies that two RMWs cannot read from the same event). The last constraint is a conservative strengthening of the original C11 model employed in the RC11 model [Lahav et al. 2017] (following [Boehm and Demsky 2014]) for avoiding the out-of-thin-air problem. To the best of our knowledge, this acyclicity condition is assumed in all C11 verification methods and tools e.g., [Doherty et al. 2019; Kokologianakis et al. 2017; Vafeiadis and Narayan 2013].

Example 3.6. Consider the following execution graph of the standard message passing test (MP):



Due to the *rf*-edge from $W(r1x, y, 1)$ to $R(r1x, y, 1)$ we obtain synchronization (*sw*) from $F(rel)$ to

$F(\text{acq})$, and thus hb from $W(r1x, x, 1)$ to $R(r1x, x, 0)$. Since fr goes in the opposite direction, this execution violates **READ COHERENCE**, and so it is not RC20-consistent. Omitting one (or both) of the fences will remove the synchronization, making this graph RC20-consistent.

There are several differences between RC20 and RC11:

- For simplicity, RC20 does not include SC accesses, whose semantics was rectified in the RC11 model [Lahav et al. 2017]. We are aware of only a few (practical) algorithms that actually employ SC accesses and become wrong when release/acquire accesses are used instead. In these cases, SC-fences can be used instead of SC accesses, although this may incur a certain cost when targeting ARMv8 that has special support for SC accesses.
- C11's SC-fences, whose semantics was rectified in the RC11 model [Lahav et al. 2017], are not included in RC20 as primitives. Instead, in RC20 SC-fences are syntactic sugar for sequences of instructions: $\text{fence}(\text{acq}); _ := f.\text{FADD}(0, \text{acqrel}); \text{fence}(\text{rel})$ where f is a distinguished otherwise-unused location. The two fences in the sequence ensure that SC-fences will behave both as an acquire fence and as a release fence for inducing synchronization using the sw relation (w.r.t. later writes and prior reads). The RMW instruction (FADD) ensures that hb will totally order the SC-fences in the execution graph, which together with **WRITE COHERENCE** and **READ COHERENCE**, provides us with the required guarantees.
- For the simplicity of the presentation, our formulation so far did not discuss C11's non-atomic accesses that have "catch-fire" semantics. These can be easily added (see §6).

Remark 1. Interestingly, the above encoding of SC fences results in a semantics that is strictly stronger (i.e., allowing less behaviors) than RC11's. Our model can be shown to be equivalent to a model with SC fences as primitives that includes $[F^{\text{sc}}]; (\text{po} \cup \text{rf})^+; [F^{\text{sc}}]$ and $[F^{\text{sc}}]; (\text{po} \cup \text{sw})^+; (\text{rf} \cup \text{mo}; \text{rf}^? \cup \text{fr}; \text{rf}^?); (\text{po} \cup \text{sw})^+; [F^{\text{sc}}]$ as components of hb (together with po and sw). In turn, RC11 only requires that $[F^{\text{sc}}]; (\text{hb} \cup \text{hb}; (\text{rf} \cup \text{mo}; \text{rf}^? \cup \text{fr}; \text{rf}^?); \text{hb}); [F^{\text{sc}}]$ is acyclic. We conjecture that the existing mapping of RC11 to hardware, as well as the compiler transformations that are sound under RC11, can be validated for our semantics. This will provide: (i) stronger guarantees to the programmers with no additional cost; and (ii) a more principled and parsimonious approach for the semantics of SC fences (they are precisely captured by $\text{fence}(\text{acq}); _ := f.\text{FADD}(0, \text{acqrel}); \text{fence}(\text{rel})$) instead of their rather ad-hoc semantics in RC11. The latter is in line with other models (e.g., TSO and the fragment of RC11 consisting of release/acquire atomics and non-atomics) in which an RMW to a distinguished otherwise-unused location is *equivalent* to an SC-fence. We note that previous work on verification under RC11 assumed a similar stronger simplified semantics for SC fences [Dang et al. 2019].

We end this section by recalling the definition of sequential consistency in the declarative framework. This formulation is used in our robustness notions (in §4) that are based on comparing the set of RC20-consistent execution graphs of a given program to the set of SC-consistent execution graphs of that program. We follow the standard definition (see, e.g., [Alglave et al. 2014]) using the derived relation:

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

Definition 3.7. An execution graph G is *SC-consistent* if $G.\text{hb}_{\text{SC}}$ is irreflexive.

It is easy to see that SC-consistency implies RC20-consistency.

3.2 From Concurrent Programs to Execution Graphs

In this section we present the definitions needed for formally relating execution graphs with concurrent programs, and use them to define what it means that a given program may fail under

RC20. For example, we will have that the execution graph in Ex. 3.6 is indeed generated by the program in that example and that this program may not fail under RC20. To so do, we define when an execution graph G is *generated by a program Pr with a final state \bar{q}* by using a special memory system, called FG (for *free graphs*). This system sets no restrictions whatsoever in its transitions (every sequence of pairs in $\text{Tid} \times \text{TLab}$ is a trace of FG), but rather maintains in its states the current execution graph. Then, the execution graphs generated by a program Pr with final state \bar{q} are defined as the execution graphs reachable by FG when FG executes a trace of Pr that reaches \bar{q} .

To define FG, we use the following notation for adding a fresh event to an execution graph. It “receives as input” the current execution graph G , the thread identifier τ of the new event, the label l of the new event, the set D of events on which the new event depends, and the write event w that the new event reads from if it is a read, or placed as the **mo**-successor of if it is a write. When the new event is an RMW ($\text{typ}(l) = \text{RMW}$), its **rf**-source and **mo**-predecessor must coincide (this is in line with **WRITE COHERENCE** and **ATOMICITY**). When the new event is a fence or a control event, we do not need the event w .

NOTATION 3.8. For $G \in \text{EGraph}$ and $\tau \in \text{Tid}$, $\text{next}(G, \tau)$ denotes the “next available event” for thread τ in G , given by $\text{next}(G, \tau) \triangleq \langle \tau, \max\{\text{sn}(e) \mid e \in G.E^\tau\} + 1 \rangle$. For $G \in \text{EGraph}$, $\tau \in \text{Tid}$, $l \in \text{Lab}$, $D \subseteq \text{Event}$, and $w \in \text{Event} \cup \{\perp\}$, $\text{add}(G, \tau, l, D, w)$ denotes the tuple $\langle E', \text{lab}', \text{dp}', \text{rf}', \text{mo}' \rangle$ defined as follows, where $e = \text{next}(G, \tau)$:

$$\begin{aligned} E' &= G.E \cup \{e\} \\ \text{lab}' &= G.\text{lab} \cup \{e \mapsto l\} \\ \text{dp}' &= G.\text{dp} \cup (D \times \{e\}) \\ \text{rf}' &= \begin{cases} G.\text{rf} \cup \{\langle w, e \rangle\} & \text{typ}(l) \in \{\text{R}, \text{RMW}, \text{R}^*\} \\ G.\text{rf} & \text{otherwise} \end{cases} \\ \text{mo}' &= \begin{cases} G.\text{mo} \cup \{\langle w', e \rangle \mid \langle w', w \rangle \in G.\text{mo}^?\} \cup \{\langle e, w' \rangle \mid \langle w, w' \rangle \in G.\text{mo}\} & \text{typ}(l) \in \{\text{W}, \text{RMW}\} \\ G.\text{mo} & \text{otherwise} \end{cases} \end{aligned}$$

To properly set the dependency relation between the events of the graph, together with the current execution graph, FG carries a *dependency store* in its state, which records for every register r the graph events that correspond to the reads on which the value of r depends.

Definition 3.9. A *dependency store* is a function $\psi : \text{Reg} \rightarrow \mathcal{P}(\text{Event})$. For $R \subseteq \text{Reg}$, we write $\psi(R)$ for the set $\bigcup_{r \in R} \psi(r)$, and $\psi[R \mapsto E]$ for the dependency store ψ' that is identical to ψ (possibly) except for $\psi'(r) = E$ for every $r \in R$.

Having $e \in \psi(r)$ means that the value of r depends on the event e . For example, after executing $a := x.\text{load}(r1x)$ we will have $\psi(a) = \{e\}$, where e is the (fresh) read event generated for this load instruction. For $a := x.\text{load}(r1x); b := x.\text{load}(r1x); c := a + b$, we will have *two* events in $\psi(c)$. This allows us to generate appropriate dependency edges, when, say, $y.\text{store}(c, r1x)$ comes next. (Our assumption that different threads use disjoint registers simplifies the definition of ψ .)

Now, we define the memory system FG.

Definition 3.10. The memory system FG is given by:

- Its states are pairs $\langle G, \psi \rangle$, where G is an execution graph and ψ is a dependency store.
- Its initial state is $\langle G_\emptyset, \psi_\emptyset \rangle$, where G_\emptyset is the initial execution graph (as defined in Def. 3.4) and $\psi_\emptyset \triangleq \lambda r. \emptyset$. \emptyset is the initial dependency store.
- Its transitions are given by:

$$\begin{array}{c}
\text{typ}(L) \in \{R, W, \text{RMW}, R^*\} \\
\frac{\text{typ}(L) \in \{R, \text{RMW}, R^*\} \implies G.\text{val}_W(w) = \text{val}_R(L) \quad G' = \text{add}(G, \tau, \text{lab}(L), \psi(R_{\text{in}}(L)), w) \quad \psi' = \psi[R_{\text{out}}(L) \mapsto \{\text{next}(G, \tau)\}]}{\langle G, \psi \rangle \xrightarrow{\tau, L}_{\text{FG}} \langle G', \psi' \rangle}
\end{array}
\qquad
\begin{array}{c}
\text{typ}(L) \in \{F, \text{CTRL}\} \\
\frac{G' = \text{add}(G, \tau, \text{lab}(L), \psi(R_{\text{in}}(L)), \perp)}{\langle G, \psi \rangle \xrightarrow{\tau, L}_{\text{FG}} \langle G', \psi \rangle} \\
\frac{\text{lab}(L) = \varepsilon \quad \psi' = \psi[R_{\text{out}}(L) \mapsto \psi(R_{\text{in}}(L))]}{\langle G, \psi \rangle \xrightarrow{\tau, L}_{\text{FG}} \langle G, \psi' \rangle}
\end{array}$$

We use FG to formally associate execution graphs with programs.

Definition 3.11. Let Pr be a program. An execution graph G is:

- generated by Pr with final state \bar{q} and dependency store ψ if $\langle \bar{q}, G, \psi \rangle$ is reachable in $Pr \parallel \text{FG}$.
- generated by Pr with final state \bar{q} if $\langle \bar{q}, G, \psi \rangle$ is reachable in $Pr \parallel \text{FG}$ for some ψ .
- generated by Pr if $\langle \bar{q}, G, \psi \rangle$ is reachable in $Pr \parallel \text{FG}$ for some \bar{q} and ψ .

Using the last definition, the semantics to a program Pr under the RC20 model is taken to be the set of RC20-consistent execution graphs that are generated by Pr . For safety under RC20, we use the following definition, which captures when an assertion violation might arise under RC20.

Definition 3.12. A program Pr may fail under RC20 if some RC20-consistent execution graph G is generated by Pr with a failing final state \bar{q} (see Def. 2.9).

In addition, we can now state the (well known) equivalence of the SC memory system (Def. 2.8) and its declarative formulation (Def. 3.7) in terms of program safety.

PROPOSITION 3.13. A program Pr may fail under SC (Def. 2.10) iff some SC-consistent execution graph G is generated by Pr with a failing final state \bar{q} .

Finally, we will use certain receptiveness properties ensuring that the set of generated execution graphs for a given program is closed under changes of the values read in read events with no outgoing dependency edges. Intuitively, if nothing depends on the difference between two graphs, then this difference cannot be observed by the program.

Definition 3.14. An event r is a G -irrelevant read if $G.\text{typ}(e) = R$ and $r \notin \text{dom}(G.\text{dp})$. We denote by $G.\text{IR}$ the set of all G -irrelevant read events.

Definition 3.15. Two execution graphs G and G' are equal up to a set T of events, denoted $G \sim_T G'$, if the following hold:

- $G.E = G'.E$, $G.\text{dp} = G'.\text{dp}$, and $G.\text{mo} = G'.\text{mo}$.
- $G.\text{typ} = G'.\text{typ}$, $G.\text{loc} = G'.\text{loc}$, $G.\text{val}_W = G'.\text{val}_W$, and $G.\text{mod} = G'.\text{mod}$.
- $G.\text{val}_R(r) = G'.\text{val}_R(r)$ for every $r \notin T$.
- $G.\text{rf}; [\text{Event} \setminus T] = G'.\text{rf}; [\text{Event} \setminus T]$.

We say that G and G' are observationally equivalent if $G \sim_{G.\text{IR}} G'$.

The following property ensures that the set of execution graphs generated by a program is closed under observational equivalence.

Definition 3.16. Two program states $\bar{q}, \bar{q}' \in Pr.Q$ are equivalent up to a set $R \subseteq \text{Reg}$, denoted by $\bar{q} \sim_R \bar{q}'$, if $\bar{q}(\tau) = \langle pc, \phi \rangle$ implies that $\bar{q}'(\tau) = \langle pc, \phi' \rangle$ for some local store ϕ' satisfying $\phi'(r) = \phi(r)$ for every $r \notin R$.

LEMMA 3.17. Suppose that $G \sim_T G'$ for some $T \subseteq G.\text{IR}$. If $\langle \bar{q}, G, \psi \rangle$ is reachable in $Pr \parallel \text{FG}$, then $\langle \bar{q}', G', \psi \rangle$ is reachable in $Pr \parallel \text{FG}$ for some program state $\bar{q}' \in Pr.Q$ such that $\bar{q} \sim_R \bar{q}'$ for $R = \{r \in \text{Reg} \mid \psi(r) \cap T \neq \emptyset\}$.


4 ROBUSTNESS AGAINST RC20

In this section we introduce the two robustness notions studied in this paper: robustness against RC20 and observational robustness against RC20. We show how they can be reduced to certain properties of *SC-consistent* execution graphs that are generated by the given program. Since we only consider robustness against RC20, we will refer to these properties simply as *robustness* and *observational robustness*.

Robustness requires that RC20-consistency and SC-consistency coincide for execution graphs generated by the program:

Definition 4.1. A program Pr is *robust* if all RC20-consistent execution graphs generated by Pr are also SC-consistent.

Example 4.2. It is easy to see that the MP in Ex. 3.6 is robust. The same program without fences, however, is not robust as the RC20-consistent but SC-inconsistent execution graph on the middle demonstrates (the graph on the right is discussed in Ex. 4.15):



The above robustness notion (w.r.t. different weak memory models) is also known as *trace-robustness* [Bouajjani et al. 2013] or *execution-graph-robustness* [Lahav and Margalit 2019]. It implies *state-robustness*: if some RC20-consistent execution graph G is generated with a final program state \bar{q} , then $\langle \bar{q}, M \rangle$ is reachable in $Pr \parallel SC$ for some $M : \text{Loc} \rightarrow \text{Val}$. In particular, the following property ensures that safety verification of robust programs may assume SC semantics.

PROPOSITION 4.3. *A robust program may fail under RC20 (Def. 3.12) iff it may fail under SC (Def. 2.10).*

Our first main theorem formulates a technical property of execution graphs that holds for all SC-consistent graphs of a given program *iff* that program is robust. It allows us to avoid the universal quantification over RC20-consistent execution graphs of Def. 4.1 and paves the way for deciding robustness via reachability analysis under SC. (The next section clarifies this reduction.)

NOTATION 4.4. $G.w_x^{\max}$ denotes the $G.mo$ -maximal write to x in G (i.e., $G.w_x^{\max} \triangleq \max_{G.mo} G.W_x$).

Definition 4.5. A *non-robustness witness* for a program Pr is a tuple $\langle \bar{q}, G, \tau, L \rangle$ such that the following hold:

- G is generated by Pr with final state \bar{q} .
- G is SC-consistent.
- \bar{q} enables the transition $\langle \tau, L \rangle$.
- $\text{typ}(L) \in \{R, W, RMW, R^*\}$.
- $G.w_{1oc(L)}^{\max} \in \text{dom}(G.hb_{SC}^? ; [E^\tau])$.
- The following hold for some $w \in G.W_{1oc(L)}$:
 - $w \neq G.w_{1oc(L)}^{\max}$.
 - $w \notin \text{dom}(G.mo ; G.rf^? ; G.hb^? ; [E^\tau])$.
 - If $\text{typ}(L) \in \{W, RMW\}$, then $w \notin \text{dom}(G.rf ; [G.RMW])$.
 - If $\text{typ}(L) \in \{RMW, R^*\}$, then $G.\text{val}_w(w) = \text{val}_R(L)$.

THEOREM 4.6. *A program Pr is not robust iff there exists a non-robustness witness for Pr .*

A non-robustness witness forms a “borderline execution” of a given program—an execution graph G of the program that is still SC-consistent, but one of the possible program steps after generating G results in an RC20-consistent but SC-inconsistent graph. Roughly speaking, this scenario happens when (under SC) some read/write/RMW step (labeled with L) accessing location x by thread τ *must* be executed after the maximal write to x in G (namely, there is a $G.\text{hb}_{\text{SC}}$ -path from $G.w_x^{\max}$ to some event of thread τ), but some other write w to x in G has no $G.\text{mo}; G.\text{rf}^?; G.\text{hb}^?$ -path to thread τ . The latter condition ensures that thread τ is not “aware” of a write that is later than w , implying that under RC20 we may extend G to the execution graph $G' = \text{add}(G, \tau, \text{lab}(L), _, w)$. Since $G.w_{\text{loc}(L)}^{\max} \in \text{dom}(G.\text{hb}_{\text{SC}}^?; [E^\tau])$, we will have that G' is SC-inconsistent. More care has to be taken because of the **ATOMICITY**-condition: if the enabled transition is a write (or an RMW) then RC20 forbids placing the new generated event $\text{next}(G, \tau)$ as the mo -successor of writes that are already read by RMWs, and thus we require $w \notin \text{dom}(G.\text{rf}; [G.\text{RMW}])$ in this case. Furthermore, for G' to be an execution of the program, we also require that $G.\text{val}_w(w) = \text{val}_R(L)$ in case that $\text{typ}(L) \in \{\text{RMW}, \text{R}^*\}$. The latter condition ensures that the program indeed allows a step leading from G to G' .

Example 4.7. A non-robustness witness for the (non-fenced) MP program in Ex. 4.2 is given by $\langle \bar{q}, G, \tau, L \rangle$, where \bar{q} is the state of the program after the first thread is fully executed and the second thread is just before the load of x (with $a = 1$); G is the execution graph obtained from the one presented in the example above by removing the last read event of the second thread; τ is the identifier of the second thread; and L is any transition label with $\text{typ}(L) = \text{R}$, $\text{loc}(L) = x$, $\text{mod}(L) = \text{rlx}$, $\text{R}_{\text{in}}(L) = \emptyset$, and $\text{R}_{\text{out}}(L) = \{b\}$. In particular, we have $G.w_x^{\max} \in \text{dom}(G.\text{hb}_{\text{SC}}^?; [E^\tau])$ ($G.w_x^{\max}$ is the event labeled with $W(\text{rlx}, x, 1)$), and for w being the initial write event to x , we have $w \notin \text{dom}(G.\text{mo}; G.\text{rf}^?; G.\text{hb}^?; [E^\tau])$.

We note that the standard DRF-theorem ensuring that there are not any weak behaviors for programs that are race free under SC (e.g., by using locks implemented with BCAS instructions) is a simple corollary of the above precise characterization. A non-robustness witness necessarily implies a race under SC between $G.w_{\text{loc}(L)}^{\max}$ and an event that can be added to G in the next step of thread τ .

4.1 Observational Robustness

The above robustness notion is too strict for supporting an important use case of relaxed accesses in which the stale values observed by relaxed accesses (and cannot be observed under SC) are in *speculative reads* that are never used by the program. This use case was discussed in [Sinclair et al. 2017] and demonstrated with the *sequence lock* implementation (seqlock for short), a special locking mechanism used in Linux that avoids writer starvation [Boehm 2012]. The next example shows the challenge seqlock poses to robustness verification.

Example 4.8. Consider the litmus test in Fig. 3, extracted from an efficient implementation of seqlock with a single writer and a single reader. The writer executes two iterations in which it increases the version number (z) to be odd, writes the protected data (x and y), and publishes the data by increasing the version number to be even again. The reader waits in a loop until it reads an even version number, then (speculatively) reads the data, and reads again the version number to check that it has not changed. If the latter read yields a different version number, the reader never uses the data that was read and retries. Finally, if the reader reads the same version number, it can safely use the data (modeled here as `assert(a = b)`). Access modes and fences were carefully picked to ensure correctness while optimizing performance [Boehm 2012]. In particular, note that the read of the protected data is relaxed.

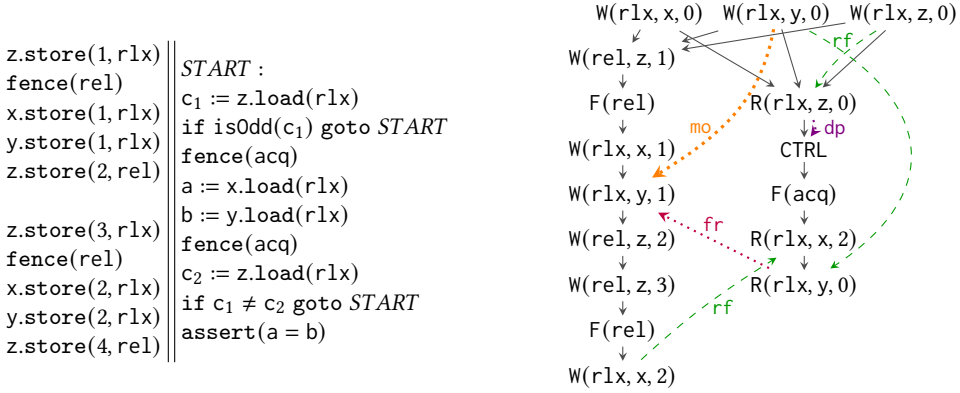


Fig. 3. Seqlock litmus test.

A robustness violation—an execution graph of the program that is RC20-consistent but SC-inconsistent—is presented on the right in Fig. 3 (to avoid clutter we omit some *mo* and *fr* edges). It is caused by the reader reading an updated value from x but an old (the initial) value from y . After reading an updated value from x , the next acquire fence that the reader will perform (before using the values) will synchronize with the last release fence of the writer (creating an *sw*-edge), so that reading $z = 0$ again is forbidden (the reader can only read $z \geq 3$), and the (taint) data will be necessarily discarded. We note that standard *state* robustness does not hold here as well: a program state in which the reader has $a = 2$ and $b = 0$ is reachable under RC20 but not under SC.

To support this use case, we define a weaker robustness notion, which we refer to as *observational robustness*. This notion allows non-SC execution graphs of the program, provided that nothing depends on events that read stale values. To formally define this notion, we adapt the declarative formulation of SC to ignore *fr*-edges induced by read events with no outgoing dependencies. For that we use two additional derived relations:

$$G.\mathit{fr}^{\mathit{dp}} \triangleq [\mathit{dom}(G.\mathit{dp}) \cup G.R^*]; G.\mathit{fr} \quad G.\mathit{hb}_{\mathit{SC}}^{\mathit{dp}} \triangleq (G.\mathit{po} \cup G.\mathit{rf} \cup G.\mathit{mo} \cup G.\mathit{fr}^{\mathit{dp}})^+$$

While in $\mathit{hb}_{\mathit{SC}}$, via *fr*, a read r is always placed before every write w' written *mo*-after the write w that r reads from, the $\mathit{hb}_{\mathit{SC}}^{\mathit{dp}}$ relation ignores such orders when (i) $r \notin G.R^*$ (that is, r does not arise from a wait instruction or a failed CAS) and (ii) no event depends on r . We note that, **ATOMICITY** ensures that $[G.\mathit{RMW}]; G.\mathit{fr} \subseteq G.\mathit{mo}$, so $G.\mathit{hb}_{\mathit{SC}}^{\mathit{dp}}$ always contains *fr*-edges from RMW events.

Using $\mathit{hb}_{\mathit{SC}}^{\mathit{dp}}$, observational robustness is naturally defined as follows.

Definition 4.9. A program Pr is *observationally robust* if $G.\mathit{hb}_{\mathit{SC}}^{\mathit{dp}}$ is irreflexive in every RC20-consistent execution graph G that is generated by Pr .

Example 4.10. Revisiting the examples above, note that the seqlock test (Ex. 4.8) is *observationally robust*. In particular, the execution graph in Fig. 3 has no $\mathit{hb}_{\mathit{SC}}^{\mathit{dp}}$ -cycle. A dependency edge from a $R(\mathit{rlx}, y, v)$ -event will be only added if the previous read of x is reading the same value v . The (non-fenced) MP program (Ex. 4.7) is not observationally robust. Without the final `assert(b ≠ 0)` instruction, that program is observationally robust (but not robust).

To see that observational robustness is not overly weak, we establish the analogue of Prop. 4.3, ensuring that verification of observationally robust programs may safely assume SC-semantics.

This easily follows from the fact that the SC-inconsistent execution graphs allowed by observational robustness are always observationally equivalent to SC-consistent ones (as defined in Def. 3.15), which we prove next.

LEMMA 4.11. *Let G be an RC20-consistent execution graph with $G.\text{hb}_{\text{SC}}^{\text{dp}}$ being irreflexive. Then, G is observationally equivalent to some SC-consistent execution graph G' .*

Lemma 4.11 and the receptiveness property (Lemma 3.17) easily entail the analogue of Prop. 4.3.

PROPOSITION 4.12. *An observationally robust program may fail under RC20 (Def. 3.12) iff it may fail under SC (Def. 2.10).*

Remark 2. Using Lemma 4.11, we have that observational robustness of a program Pr implies that every RC20-consistent execution graph generated by Pr is observationally equivalent to some SC-consistent execution graph. We note that the converse does not hold. For example, consider the following program:

$$\begin{array}{l} x.\text{store}(1, \text{rel}) \\ y.\text{store}(1, \text{rel}) \end{array} \left\| \begin{array}{l} a := y.\text{load}(\text{acq}) // 1 \\ b := z.\text{load}(\text{acq}) // 0 \\ \text{assert}(b = b) \end{array} \right\| \begin{array}{l} z.\text{store}(1, \text{rel}) \\ c := x.\text{load}(\text{acq}) // 0 \\ \text{assert}(c = c) \end{array}$$

The execution graph G that corresponds to the annotated outcome has a cycle in $G.\text{hb}_{\text{SC}}^{\text{dp}}$ (using two $G.\text{fr}^{\text{dp}}$ -edges: from the read of z and from the read of x). Nevertheless, by changing the (irrelevant) read of y to read 0, we obtain an SC-consistent graph that is observationally equivalent to G .

Next, to be able to check for observational robustness of a given program, we present *non-observational-robustness witnesses* which, as before, allow us to verify observational robustness by considering only SC-executions of the program. For this definition, we use *tainted* versions of several relations w.r.t. a set $T \subseteq G.\text{IR}$ of G -irrelevant reads (see Def. 3.14), defined as follows:

Definition 4.13. Given $T \subseteq G.\text{IR}$, the T -tainted versions of rf , sw , and hb are defined as follows:⁴

$$\begin{aligned} G.\text{rf}_{\bar{T}} &\triangleq G.\text{rf} ; [\text{Event} \setminus T] \\ G.\text{sw}_{\bar{T}} &\triangleq [G.E^{\exists \text{rel}}] ; ([G.F] ; G.\text{po})^? ; G.\text{rf}_{\bar{T}}^+ ; (G.\text{po} ; [G.F])^? ; [G.E^{\exists \text{acq}}] \\ G.\text{hb}_{\bar{T}} &\triangleq (G.\text{po} \cup G.\text{sw}_{\bar{T}})^+ \end{aligned}$$

Definition 4.14. A *non-observational-robustness witness* for a program Pr is a tuple $\langle \bar{q}, G, \tau, L, \psi, T \rangle$ such that the following hold:

- G is generated by Pr with final state \bar{q} and dependency store ψ .
- G is SC-consistent.
- \bar{q} enables the transition $\langle \tau, L \rangle$.
- $\text{typ}(L) \in \{\text{W}, \text{RMW}, \text{R}^*, \text{CTRL}\}$.
- $T \subseteq G.\text{IR}$.
- $G.\text{rf} ; [T] \subseteq G.\text{hb}_{\text{SC}} ; G.\text{po} \setminus G.\text{rf}_{\bar{T}}^? ; G.\text{hb}_{\bar{T}}$.
- Either $T \cap \psi(\text{R}_{\text{in}}(L)) \neq \emptyset$ or the following hold:
 - $\text{typ}(L) \in \{\text{W}, \text{RMW}, \text{R}^*\}$.
 - $G.w_{\text{loc}(L)}^{\text{max}} \in \text{dom}(G.\text{hb}_{\text{SC}}^? ; [E^r])$.
 - The following hold for some $w \in G.W_{\text{loc}(L)}$:
 - * $w \neq G.w_{\text{loc}(L)}^{\text{max}}$.
 - * $w \notin \text{dom}(G.\text{mo} ; G.\text{rf}_{\bar{T}}^? ; G.\text{hb}_{\bar{T}}^? ; [E^r])$.
 - * If $\text{typ}(L) \in \{\text{W}, \text{RMW}\}$, then $w \notin \text{dom}(G.\text{rf} ; [G.\text{RMW}])$.
 - * If $\text{typ}(L) \in \{\text{RMW}, \text{R}^*\}$, then $G.\text{val}_w(w) = \text{val}_R(L)$.

Non-observational-robustness witnesses are similar to non-robustness witnesses (Def. 4.5), but they also contain a dependency store ψ and a set T of “tainted” read events in G . The events in T are the read events in G that could have read from an mo -earlier write event (w.r.t. their current rf -source) under RC20 but not under SC (i.e., $G.\text{rf} ; [T] \subseteq G.\text{hb}_{\text{SC}} ; G.\text{po} \setminus G.\text{rf}_{\bar{T}}^? ; G.\text{hb}_{\bar{T}}$.)

⁴By definition, T cannot contain events in $G.\text{RMW}$, so we always have $G.\text{rf}_{\bar{T}}^+ = G.\text{rf}^* ; G.\text{rf}_{\bar{T}}$.

The existence of such a read event would entail a robustness violation, but *not* an *observational* robustness violation, which may be violated only if some later event (in particular, a CTRL event) depends on an event in T (i.e., $T \cap \psi(R_{\text{in}}(L)) \neq \emptyset$). When $\text{typ}(L) \in \{\text{W}, \text{RMW}, \text{R}^*\}$, we have the same conditions as in [Def. 4.5](#), but crucially, the read-from edges in G to tainted read events should not limit the writes that threads can observe under RC20. Hence, we use the tainted versions of rf and hb in the $w \notin \text{dom}(G.\text{mo}; G.\text{rf}_T^?; G.\text{hb}_T^?; [E^r])$ condition.

Example 4.15. A non-observational-robustness witness for the (non-fenced) MP program in [Ex. 4.2](#) consists of the graph depicted on the right in that example, where the set T includes the highlighted event reading x . The read event in T could have read from the initial write under RC20 but not under SC; and the program enables a transition that depends on the event in T (executing the `assert(b ≠ 0)`).

Next, we establish that non-observational-robustness witnesses must exist for non-observationally-robust programs. The proof is rather intricate DRF-style argument—RC20 can deviate from SC early on, and to construct the witness for a non-robust program we need to globally “fix” these deviations.

THEOREM 4.16. *If Pr is not observationally robust, then there exists a non-observational-robustness witness for Pr .*

The following lemma is a key lemma in the proof of [Thm. 4.16](#).

LEMMA 4.17. *Let G be an RC20-consistent execution graph. Suppose that $G.\text{hb}_{\text{SC}}^{\text{dp}}$ is irreflexive. Then, there exist an SC-consistent execution graph G' and a set $T \subseteq G.\text{IR}$ such that $G \sim_T G'$ and $G'.\text{rf}; [T] \subseteq G'.\text{hb}_{\text{SC}}; G'.\text{po} \setminus G'.\text{rf}_T^?; G'.\text{hb}_T^-$.*

PROOF (OUTLINE). Let S be a maximal acyclic relation such that

$$G.\text{po} \cup G.\text{rf} \cup G.\text{mo} \cup [\text{dom}(G.\text{dp}) \cup G.\text{R}^*]; G.\text{fr} \subseteq S \subseteq G.\text{po} \cup G.\text{rf} \cup G.\text{mo} \cup G.\text{fr}.$$

Define the following:

- For every $r \in G.\text{R}$, let $W_r = \{w \mid \langle r, w \rangle \in G.\text{fr} \setminus S^+\}$.
- Let $T = \{r \in G.\text{R} \mid W_r \neq \emptyset\}$.
- For every $r \in T$, let $w_r = \max_{G.\text{mo}} W_r$.

Then, one shows that the following properties hold:

- $T \subseteq G.\text{IR}$.
- If $\langle w, r \rangle \in G.\text{rf}; [T]$, then $\langle w, w_r \rangle \in G.\text{mo}$.
- $\langle w_r, r \rangle \in S^+; G.\text{po}$ for every $r \in T$.
- $\langle r, w \rangle \in S^+$ for every $r \in T$ and w such that $\langle w_r, w \rangle \in G.\text{mo}$.

Now, we let G' be the execution graph obtained from G by modifying each event $r \in T$ to read from w_r (and adapting read values accordingly). Our construction ensures that $G \sim_T G'$, $G.\text{rf} \subseteq (G'.\text{mo})^?; G'.\text{rf}; G.\text{fr} \cap S \subseteq G'.\text{fr}$, and $S \subseteq G'.\text{hb}_{\text{SC}}$. With these properties, we can show that the two required conditions hold. \square

Next, we outline the proof of [Thm. 4.16](#).

PROOF OF [Thm. 4.16](#) (OUTLINE). Let \mathcal{G} be the set consisting of all execution graphs G such that (i) G is RC20-consistent; (ii) G is generated by Pr ; and (iii) $G.\text{hb}_{\text{SC}}^{\text{dp}}$ is not irreflexive. Since Pr is not observationally robust, \mathcal{G} is not empty. Let G_1 be a minimal element in \mathcal{G} , in the sense that every proper prefix of G_1 is not in \mathcal{G} . Let $\bar{q}_1 \in Pr.Q$ and $\psi_1 : \text{Reg} \rightarrow \mathcal{P}(\text{Event})$ such that $\langle \bar{q}_1, G_1, \psi_1 \rangle$ is reachable in $Pr \parallel \text{FG}$. Let $G_2 \in \text{EGraph}$, $\bar{q}_2 \in Pr.Q$, $\psi_2 : \text{Reg} \rightarrow \mathcal{P}(\text{Event})$, $\tau \in \text{Tid}$, and $L = \langle l, R_{\text{in}}, R_{\text{out}} \rangle \in \text{TLab}$ such that:

- $l \in \text{Lab}$.
- $\langle \bar{q}_2, G_2, \psi_2 \rangle$ is reachable in $Pr \parallel \text{FG}$.
- $\bar{q}_2 \xrightarrow{\tau, L}_{Pr} \bar{q}$ for some $\bar{q} \in Pr.Q$ such that $\bar{q} \xrightarrow{\langle \epsilon, \dots \rangle^*}_{Pr} \bar{q}_1$.
- $\langle G_2, \psi_2 \rangle \xrightarrow{\tau, L}_{\text{FG}} \langle G_1, \psi_1 \rangle$.

Then, by definition, there exists some $w_c \in \text{Event} \cup \{\perp\}$ such that the following hold:


- $G_1 = \text{add}(G_2, \tau, l, \psi_2(R_{\text{in}}), w_c)$.
- If $\text{typ}(l) \in \{\text{R}, \text{RMW}, \text{R}^*\}$, then $G_2.\text{val}_W(w_c) = \text{val}_R(l)$.
- If $\text{typ}(l) \notin \{\text{F}, \text{CTRL}\}$, then $w_c \in G_2.W_{\text{loc}}(l)$.
- If $\text{typ}(l) \in \{\text{F}, \text{CTRL}\}$, then $w_c = \perp$.

Let $c = \text{next}(G_2, \tau)$ (i.e., the event added to G_2 to obtain G_1). Then, $G_2 = G_1 \setminus \{c\}$ is RC20-consistent, and the minimality of G_1 ensures that $G_2.\text{hb}_{\text{SC}}^{\text{dp}}$ is irreflexive.

By Lemma 4.17 there exist an SC-consistent execution graph G_3 and a set $T \subseteq G_2.\text{IR}$ such that $G_2 \sim_T G_3$ and $G_3.\text{rf}; [T] \subseteq G_3.\text{hb}_{\text{SC}}; G_3.\text{po} \setminus G_3.\text{rf}_{\bar{T}}^?; G_3.\text{hb}_{\bar{T}}$. Since $T \subseteq G_2.\text{IR}$, Lemma 3.17 implies that $\langle \bar{q}_3, G_3, \psi_2 \rangle$ is reachable in $Pr \parallel \text{FG}$ for some state $\bar{q}_3 \in Pr.Q$ such that $\bar{q}_2 \sim_R \bar{q}_3$ for $R = \{r \in \text{Reg} \mid \psi_2(r) \cap T \neq \emptyset\}$. We also have that \bar{q}_3 enables the transition $\langle \tau, L \rangle$ if $\psi_2(R_{\text{in}}) \cap T = \emptyset$; and that \bar{q}_3 enables $\langle \tau, (l_3, R_{\text{in}}, R_{\text{out}}) \rangle$ for some event label $l_3 \in \text{Lab}$.

Now, if $\psi_2(R_{\text{in}}) \cap T \neq \emptyset$, then $\langle \bar{q}_3, G_3, \tau, (l_3, R_{\text{in}}, R_{\text{out}}), \psi_2, T \rangle$ is a non-observational-robustness witness for Pr . Otherwise, we can show that $\text{typ}(l) \in \{\text{W}, \text{RMW}, \text{R}^*\}$ and $\langle c, c \rangle \in (G_1.\text{mo} \cup G_1.\text{fr})$; $G_1.\text{hb}_{\text{SC}}^{\text{dp}}; G_1.\text{po}$. Let $E_4 = \text{dom}(G_3.\text{hb}_{\text{SC}}; [E^\tau])$, $G_4 = G_3 \cap E_4$ and $T' = T \cap E_4$. Since G_4 is a prefix of G_3 and $G_3.E^\tau \subseteq G_4.E$, there exist $\bar{q}_4 \in Pr.Q$ and dependency store ψ_4 , such that $\langle \bar{q}_4, G_4, \psi_4 \rangle$ is reachable in $Pr \parallel \text{FG}$ and \bar{q}_4 enables $\langle \tau, L \rangle$. As a prefix of G_3 , G_4 is SC-consistent, and since $G_3.\text{rf}; [T] \subseteq G_3.\text{hb}_{\text{SC}}; G_3.\text{po} \setminus G_3.\text{rf}_{\bar{T}}^?; G_3.\text{hb}_{\bar{T}}$, our construction ensures that $G_4.\text{rf}; [T'] \subseteq G_4.\text{hb}_{\text{SC}}; G_4.\text{po} \setminus G_4.\text{rf}_{\bar{T}'}^?; G_4.\text{hb}_{\bar{T}'}$. It follows that $\langle \bar{q}_4, G_4, \tau, L, \psi_4, T' \rangle$ is a non-observational-robustness witness for Pr . \square


Example 4.18. An important detail in the definition of a non-observational-robustness witness is the fact that we use $G.\text{rf}_{\bar{T}}^?; G.\text{hb}_{\bar{T}}^?$ rather than $G.\text{rf}^?; G.\text{hb}^?$ in the condition checking that thread τ is not “aware” of a later write than w . Intuitively speaking, since the execution graph in the witness is SC-consistent, the tainted reads in this graph also read the most recent values, but these reads must not make the thread aware of more writes, since they could have read from an **mo**-earlier write under RC20. The following example of a “nested” SB litmus test demonstrates this issue. Under RC20 (or any other weak memory model), we may get the non-SC outcome $b = d = 0$, which, because of the conditionals, implies that the program is not observationally robust. The graph component of a possible non-observational-robustness witness is presented on the right (all initialization event summarized in one event and **mo**-edges are omitted) where the highlighted events are the tainted read events in the set T .



However, if we use the non-tainted versions of **rf** and **hb** in the definition of a non-observational-robustness witness, we will not be able to include the read from x in T (since we will wrongly

have that the thread is “aware” of $W(\text{rel}, x, 1)$ and cannot read the initial value 0). In fact, with $G.\text{rf}^? ; G.\text{hb}^?$ instead of $G.\text{rf}_T^? ; G.\text{hb}_T^?$, we will not have *any* non-observational-robustness witness for this non-observationally-robust example.

Example 4.19. The converse of the claim in [Thm. 4.16](#) does not hold. The following simple adaptation of the SB litmus test is an observationally robust program that has an non-observational-robustness witness (whose execution graph component is presented on the right, where the highlighted event is the only event in T). Developing an exact (and verifiable) characterization of observational-robustness is challenging and is left for future work.



Finally, by adding (vacuous) assert instructions we can reduce robustness verification to the existence of a non-observational-robustness witness.

LEMMA 4.20. *Let Pr be a program in which every read instruction $r := x.\text{load}(o_R)$ is followed by a (vacuous) assert instruction $\text{assert}(r = r)$. Then, there exists a non-robustness witness for Pr iff there exists a non-observational-robustness witness for Pr .*

Using [Lemma 4.20](#), we obtain a reduction from robustness verification to the problem of verifying whether a non-observational-robustness witness exists or not. In the rest of this paper we focus on the latter problem, and obtain a decision procedure for robustness using this reduction.

5 VERIFYING OBSERVATIONAL ROBUSTNESS

In this section we show how to automatically check for the existence of a non-observational-robustness witness by exploring the reachable states in an instrumented SC semantics. We present the instrumented semantics in two stages. First, [§5.1](#) presents an “operational” version of [Thm. 4.16](#) in the form of an (infinite) memory system that we call SC_T . This system extends the standard operational construction of SC-consistent execution graphs with instrumentation and monitoring for non-observational-robustness witnesses. Then, in [§5.2](#) we present SC_M , a finite memory system instrumenting the standard SC memory system, and show that it simulates SC_T . Finally, we obtain that non-observational-robustness witness exists iff certain error state is reachable under SC_M .

5.1 The Memory System SC_T

The memory system SC_T provides an “operational” version of [Thm. 4.16](#). It is based on an operationalized version of the declarative SC semantics (see [Def. 3.7](#)), which restricts the steps of the graph generation system FG ([Def. 3.10](#)) so that reads only read from the last executed write (to the relevant location) and that writes are always placed in the end of the (per-location) mo -order. In addition, SC_T includes instrumentation and monitoring that (i) tracks a “taint set” T of read events that read stale values during the execution; and (ii) enters an error state, denoted by \perp , when a non-observational-robustness witness is detected.

Definition 5.1. The memory system SC_T is given by:

- Its states consists of the error state \perp , and all tuples of the form $\langle G, \psi, T \rangle$, where G is an execution graph, ψ is a dependency store, and $T \subseteq \{e \in G.E \mid G.\text{typ}(e) = R\}$.
- Its initial state is $\langle G_\emptyset, \psi_\emptyset, \emptyset \rangle$ (where G_\emptyset and ψ_\emptyset are defined in [Def. 3.10](#)).
- Its transitions are given in [Fig. 4](#).

$$\begin{array}{c}
\text{SILENT/FENCE/CONTROL} \\
\frac{\langle G, \psi \rangle \xrightarrow{\tau, L}_{\text{FG}} \langle G', \psi' \rangle \quad \text{lab}(L) = \varepsilon \vee \text{typ}(L) \in \{\text{F}, \text{CTRL}\}}{\langle G, \psi, T \rangle \xrightarrow{\tau, L}_{\text{SC}_T} \langle G', \psi', T \rangle} \\
\\
\text{NON-READ} \\
\frac{\langle G, \psi \rangle \xrightarrow{\tau, L}_{\text{FG}} \langle G', \psi' \rangle \quad \text{typ}(L) \in \{\text{W}, \text{RMW}, \text{R}^*\} \quad G' = \text{add}(G, \tau, \dashv, \dashv, G.w_{\text{loc}(L)}^{\max})}{\langle G, \psi, T \rangle \xrightarrow{\tau, L}_{\text{SC}_T} \langle G', \psi', T \rangle} \\
\\
\text{READ} \\
\frac{\langle G, \psi \rangle \xrightarrow{\tau, L}_{\text{FG}} \langle G', \psi' \rangle \quad \text{typ}(L) = \text{R} \quad G' = \text{add}(G, \tau, \dashv, \dashv, G.w_{\text{loc}(L)}^{\max}) \quad G.w_{\text{loc}(L)}^{\max} \in \text{dom}(G.\text{hb}_{\text{SC}}; [E^\tau]) \implies G.w_{\text{loc}(L)}^{\max} \in \text{dom}(G.\text{rf}_T^?; G.\text{hb}_T^?; [E^\tau])}{\langle G, \psi, T \rangle \xrightarrow{\tau, L}_{\text{SC}_T} \langle G', \psi', T \rangle} \quad \frac{\text{ERR-DEP} \quad \text{typ}(L) \in \{\text{W}, \text{RMW}, \text{R}^*, \text{CTRL}\} \quad \psi(\text{R}_{\text{in}}(L)) \cap T \neq \emptyset}{\langle G, \psi, T \rangle \xrightarrow{\tau, L}_{\text{SC}_T} \perp} \\
\\
\text{TAINT-READ} \\
\frac{\langle G, \psi \rangle \xrightarrow{\tau, L}_{\text{FG}} \langle G', \psi' \rangle \quad \text{typ}(L) = \text{R} \quad G' = \text{add}(G, \tau, \dashv, \dashv, G.w_{\text{loc}(L)}^{\max}) \quad T' = T \cup \{\text{next}(G, \tau)\} \quad G.w_{\text{loc}(L)}^{\max} \in \text{dom}(G.\text{hb}_{\text{SC}}; [E^\tau]) \quad G.w_{\text{loc}(L)}^{\max} \notin \text{dom}(G.\text{rf}_T^?; G.\text{hb}_T^?; [E^\tau])}{\langle G, \psi, T \rangle \xrightarrow{\tau, L}_{\text{SC}_T} \langle G', \psi', T' \rangle} \\
\\
\text{ERR-SC} \\
\frac{\text{typ}(L) \in \{\text{W}, \text{RMW}, \text{R}^*\} \quad G.w_{\text{loc}(L)}^{\max} \in \text{dom}(G.\text{hb}_{\text{SC}}; [E^\tau]) \quad w \neq G.w_{\text{loc}(L)}^{\max} \quad w \in G.W_{\text{loc}(L)} \quad w \notin \text{dom}(G.\text{mo}; G.\text{rf}_T^?; G.\text{hb}_T^?; [E^\tau]) \quad \text{typ}(L) \in \{\text{W}, \text{RMW}\} \implies w \notin \text{dom}(G.\text{rf}; [G.\text{RMW}]) \quad \text{typ}(L) \in \{\text{RMW}, \text{R}^*\} \implies G.\text{val}_w(w) = \text{val}_R(L)}{\langle G, \psi, T \rangle \xrightarrow{\tau, L}_{\text{SC}_T} \perp}
\end{array}$$

Fig. 4. SC_T transitions.

Each of the (non-error) steps of SC_T extends the current graph with one event (like the FG system in Def. 3.10) while recording in the set T all the read events added to the graph that read tainted values, maintaining the set T to consist of all the tainted read events in the current graph. Exactly as in the definition of non-observational-robustness witness, the latter formally means that there is a hb_{SC} -path, but not an $\text{rf}_T^?$; $\text{hb}_T^?$ -path from the mo -maximal write to the relevant location to the thread executing the read. In addition, two steps lead the system to the error state when a non-observational-robustness witness is detected. The soundness and completeness of SC_T is established in the following theorem.

THEOREM 5.2. *There exists a non-observational-robustness witness for a program Pr iff some state of the form $\langle \bar{q}, \perp \rangle$ is reachable in $Pr \parallel \text{SC}_T$.*

5.2 The Memory System SC_M

We show how to automatically check whether a state of the form $\langle _, \perp \rangle$ is reachable in $Pr \parallel \text{SC}_T$. Since execution graphs (for programs with loops) may grow unboundedly, SC_T is an infinite state transition system whose traces cannot be naively explored. To address this, we show that maintaining the whole execution graph that was generated so far, as done in SC_T , is unnecessary. Instead, it is possible to summarize SC_T 's state and record only the properties of $\langle G, \psi, T \rangle$ that are needed for checking which transitions are enabled in each state and whether a transition to the error state is enabled or not. To do so, we define a *finite* (but precise) abstraction of SC_T , a memory system which we call SC_M (for SC with Monitors), that simulates SC_T , so that the traces of SC_T that reach the error state coincide with those of SC_M that reach the error state.

Next, we present SC_M 's states and the transitions between them. We write $q_{\text{SC}_M}(G, \psi, T)$ for the SC_M state that abstracts the SC_T state $\langle G, \psi, T \rangle$. States of SC_M are tuples of the form $q_{\text{SC}_M} =$

$\langle M, R_{\text{taint}}, V_{\text{SC}}, V \rangle$. We use $q_{\text{SC}_M} \cdot M$, $q_{\text{SC}_M} \cdot R_{\text{taint}}$, $q_{\text{SC}_M} \cdot V_{\text{SC}}$, and $q_{\text{SC}_M} \cdot V$ to refer to the different components. Below, we gradually present the structure and meaning of each component and describe how it is maintained in the system's transitions.

Memory (M). Since SC_T monitors *SC-executions*, for checking what transitions are enabled (whether $\langle G, \psi, T \rangle$ enables a transition $\langle \tau, L \rangle$ in SC_T), we only need to track the last value written to every location. This $\text{Loc} \rightarrow \text{Val}$ mapping is stored in the M component. Formally, we have

$$q_{\text{SC}_M}(G, \psi, T) \cdot M = \lambda x. G \cdot \text{val}_W(G \cdot w_x^{\text{max}}).$$

Then, $\langle \tau, L \rangle$ is enabled iff $\text{typ}(L) \in \{W, F, \text{CTRL}\}$ (these are always enabled by the memory) or $\text{typ}(L) \in \{R, \text{RMW}, R^*\}$ and $\text{val}_R(L) = q_{\text{SC}_M} \cdot M(\text{loc}(L))$. (The rest of the system only instruments and monitors the runs, so it does not affect the transitions between non-error states.) Initially, we have $M_\emptyset = \lambda x. 0$, since the initial execution graph contains initialization to 0 for all locations. The maintenance of M is trivial: write and RMW steps to x update the value of x , while other transitions keep the memory intact.

Tainted Registers (R_{taint}). Instead of maintaining the (unbounded) set T of tainted events, SC_M keeps a (bounded) set $R_{\text{taint}} \subseteq \text{Reg}$ of *tainted registers*. A register is *tainted* if its value depends on some event $e \in T$. Formally, we have

$$q_{\text{SC}_M}(G, \psi, T) \cdot R_{\text{taint}} = \{r \in \text{Reg} \mid \psi(r) \cap T \neq \emptyset\}.$$

Then, we replace the condition $\psi(R_{\text{in}}(L)) \cap T \neq \emptyset$ in SC_T 's *ERR-DEP* step with $R_{\text{in}}(L) \cap R_{\text{taint}} \neq \emptyset$. Initially, we have $R_{\text{taint}\emptyset} = \emptyset$. Registers are added to R_{taint} when a taint value is loaded to them (in a read instruction), or if any expression that includes tainted registers is assigned to them. Registers are removed from R_{taint} when they are overwritten with a non-taint value (e.g., by another read instruction). Whether a read transition reads a taint value or not (i.e., whether RC20 allows to read from some other event in the graph) is detected using the components described below.

hb_{SC}-tracker (V_{SC}). The V_{SC} component in SC_M states is used to track whether a thread is “*hb_{SC}-aware*” of the maximal write to some location. This piece of information is needed for checking the $G \cdot w_{\text{loc}(L)}^{\text{max}} \in \text{dom}(G \cdot \text{hb}_{\text{SC}}; [E^\tau])$ condition that appears in SC_T 's *READ*, *TAINT-READ*, and *ERR-SC* steps. The V_{SC} component is defined and maintained exactly as in [Lahav and Margalit 2019]. To save space, we omit the details here and simply assume transitions of the form $V_{\text{SC}} \xrightarrow{\tau, l} V'_{\text{SC}}$ (where $\tau \in \text{Tid}$ and $l \in \text{Lab}$) for maintaining this component, and a query of the form $x \in V_{\text{SC}}(\tau)$ that “returns true” iff $G \cdot w_x^{\text{max}} \in \text{dom}(G \cdot \text{hb}_{\text{SC}}^?; [E^\tau])$. Initially, we start with $V_{\text{SC}\emptyset}$, where $x \in V_{\text{SC}\emptyset}(\tau)$ always “returns true”. We refer the reader to [Lahav and Margalit 2019, §5] for details.

RC20 tracker (V). The V component in SC_M states is the most complex component in SC_M 's states. It provides the additional information required for detecting robustness violations, as well as for differentiating between normal and taint read steps (for maintaining $q_{\text{SC}_M} \cdot R_{\text{taint}}$). First, we need to track for every thread τ and location x , whether, besides the *mo*-maximal write, there exists a write w to x in the current execution graph that does not have an-*mo*; $\text{rf}_{\tau}^?$; $\text{hb}_{\tau}^?$ -path to some event of thread τ (which means that w can be observed by thread τ under RC20). Indeed, this information is needed for checking the $G \cdot w_{\text{loc}(L)}^{\text{max}} \in \text{dom}(G \cdot \text{rf}_{\tau}^?; G \cdot \text{hb}_{\tau}^?; [E^\tau])$ condition in SC_T 's *READ* and *TAINT-READ* steps, as well as the $w \notin \text{dom}(G \cdot \text{mo}; G \cdot \text{rf}_{\tau}^?; G \cdot \text{hb}_{\tau}^?; [E^\tau])$ condition in SC_T 's *ERR-SC* step. Furthermore, for SC_T 's *ERR-SC* step, we also need to know:

- whether w is already read by some RMW event in the graph ($w \notin \text{dom}(\text{rf}; [\text{RMW}])$), in which case we cannot add a write placed as the *mo*-successor of w without violating *ATOMICITY*; and

- whether w 's written value is a value that can be read by the program when adding a certain label l , since for CAS and blocking instructions (BCAS and wait), the program only allows to add R^* or RMW events that read certain values.

We achieve this by maintaining a tuple $V = \langle V^a, V^c, V^r, V^w, V_{RMW}^a, V_{RMW}^c, V_{RMW}^r, V_{RMW}^w \rangle$, where: $V^\alpha, V_{RMW}^\alpha : \text{Tid} \times \text{Loc} \rightarrow \mathcal{P}(\text{Val})$ for $\alpha \in \{a, c, r\}$, and $V^w, V_{RMW}^w : \text{Loc} \times \text{Loc} \rightarrow \mathcal{P}(\text{Val})$. The V^c and V_{RMW}^c components (standing for *current view* and *current RMW-view*) are the ones actually used for tracking the information mentioned above: we have $v \in V^c(\tau, x)$ iff the value v is written by some write w to x in the current execution graph such that $w \neq w_x^{\text{max}}$ and $w \notin \text{dom}(\text{mo}; \text{rf}_T^?; \text{hb}_T^?; [E^\tau])$; and we have $v \in V_{RMW}^c(\tau, x)$ iff $w \notin \text{dom}(\text{rf}; [\text{RMW}])$ in addition to the above conditions. The other six components provide the additional instrumentation required for a faithful maintenance of V^c and V_{RMW}^c . Roughly speaking:

- V^a and V_{RMW}^a (*acquire view* and *acquire RMW-view*) are needed to handle acquire fences.
- V^r and V_{RMW}^r (*release view* and *release RMW-view*) are needed to handle release fences.
- V^w and V_{RMW}^w (*location view* and *location RMW-view*) are needed to handle reads and RMWs, which incorporate the view of the write event they read from in the thread view. (Crucially, since we are running under SC, we only need to record this data for **mo**-maximal write events.)

We use $q_{SC_M}.V^\alpha$ and $q_{SC_M}.V_{RMW}^\alpha$ (for $\alpha \in \{a, c, r, w\}$) to directly accesses these views. Their formal meaning is given by ($G.\text{val}_w$ is lifted to sets of events in the obvious way):

$$\begin{aligned}
q_{SC_M}(G, \psi, T).V^c &= \lambda\tau, x. G.\text{val}_w[G.W_x^{\neq\text{max}} \setminus \text{dom}(R; [E^\tau])] \\
q_{SC_M}(G, \psi, T).V^r &= \lambda\tau, x. G.\text{val}_w[G.W_x^{\neq\text{max}} \setminus \text{dom}(R; [E^\tau \cap G.F^{\exists\text{rel}}])] \\
q_{SC_M}(G, \psi, T).V^a &= \lambda\tau, x. G.\text{val}_w[G.W_x^{\neq\text{max}} \setminus \text{dom}(R; ([G.E^{\exists\text{rel}}]; ([G.F]; G.\text{po})^?; G.\text{rf}_T^+)^?; [E^\tau])] \\
q_{SC_M}(G, \psi, T).V^w &= \lambda y, x. G.\text{val}_w[G.W_x^{\neq\text{max}} \setminus \text{dom}(R; [G.E^{\exists\text{rel}}]; ([G.F]; G.\text{po})^?; G.\text{rf}_T^*; [G.w_y^{\text{max}}])] \\
q_{SC_M}(G, \psi, T).V_{RMW}^c &= \lambda\tau, x. G.\text{val}_w[G.W_x^{\neq\text{max}} \setminus \text{dom}(R; [E^\tau] \cup R_{RMW})] \\
q_{SC_M}(G, \psi, T).V_{RMW}^r &= \lambda\tau, x. G.\text{val}_w[G.W_x^{\neq\text{max}} \setminus \text{dom}(R; [E^\tau \cap G.F^{\exists\text{rel}}] \cup R_{RMW})] \\
q_{SC_M}(G, \psi, T).V_{RMW}^a &= \lambda\tau, x. G.\text{val}_w[G.W_x^{\neq\text{max}} \setminus \text{dom}(R; ([G.E^{\exists\text{rel}}]; ([G.F]; G.\text{po})^?; G.\text{rf}_T^+)^?; [E^\tau] \cup R_{RMW})] \\
q_{SC_M}(G, \psi, T).V_{RMW}^w &= \lambda y, x. G.\text{val}_w[G.W_x^{\neq\text{max}} \setminus \text{dom}(R; [G.E^{\exists\text{rel}}]; ([G.F]; G.\text{po})^?; G.\text{rf}_T^*; [G.w_y^{\text{max}}] \cup R_{RMW})]
\end{aligned}$$

where $G.W_x^{\neq\text{max}} = G.W_x \setminus \{G.w_x^{\text{max}}\}$, $R \triangleq G.\text{mo}; G.\text{rf}_T^?; G.\text{hb}_T^?$, and $R_{RMW} \triangleq G.\text{rf}; [G.\text{RMW}]$.

To explain these equations (which serve as our simulation invariants), we call a value v of location x *non-robustly readable* for thread τ if thread τ can read v from x under RC20 from some write event that is not the **mo**-maximal write to x in the current graph. Thus, a value v of x is *not* non-robustly readable for thread τ if τ is already aware of some **mo**-later write to x (“aware” here means an $\text{rf}_T^?; \text{hb}_T^?$ -path). Then, $V^c(\tau, x)$ consists of all non-robustly readable values of x for thread τ ; $V^a(\tau, x)$ consists of all non-robustly readable values of x for thread τ that will remain non-robustly readable for τ even after it executes an acquire fence; and $V^r(\tau, x)$ consists of all values of x that may remain non-robustly readable for any other thread even when it synchronizes (via the **sw** relation) with the last release fence of thread τ . In turn, $V^w(y, x)$ consists of all values of x that may remain non-robustly readable for any thread after it reads (the most recent value) from $y \neq x$. The V_{RMW}^α variants are similar, but they refer to non-robustly *writable* values for thread τ , that is: values of non-**mo**-maximal write events that thread τ may overwrite by putting the new write as their **mo**-successors.

Algorithm 1 Thread views maintenance for a step of thread τ labeled with $l \in \text{Lab}$ with $\text{typ}(l) \in \{\text{R}, \text{W}, \text{RMW}, \text{R}^*\}$ and $\text{loc}(l) = x$

- 1: $\forall \alpha \in \{\text{a}, \text{c}\} : V^\alpha(\tau, x) := \emptyset$
 - 2: **if** $\text{typ}(l) \in \{\text{W}, \text{RMW}\}$ **then**
 - 3: $V^r(\tau, x) := \{M(x)\}$
 - 4: **if** $\text{typ}(l) \in \{\text{R}, \text{RMW}, \text{R}^*\}$ **then**
 - 5: $\forall y \neq x : V^a(\tau, y) := V^w(x, y)$
 - 6: **if** $\text{mod}(l) \supseteq \text{acq}$ **then**
 - 7: $\forall y \neq x : V^c(\tau, y) := V^w(x, y)$
 - 8: **if** $\text{typ}(l) \in \{\text{W}, \text{RMW}\}$ **then**
 - 9: $\forall \alpha \in \{\text{a}, \text{c}, \text{r}\}, \pi \neq \tau : V^\alpha(\pi, x) := \{M(x)\}$
-

Algorithm 2 Thread RMW-views maintenance for a step of thread τ labeled with $l \in \text{Lab}$ with $\text{typ}(l) \in \{\text{R}, \text{W}, \text{RMW}, \text{R}^*\}$ and $\text{loc}(l) = x$

- 1: $\forall \alpha \in \{\text{a}, \text{c}\} : V_{\text{RMW}}^\alpha(\tau, x) := \emptyset$
 - 2: **if** $\text{typ}(l) = \text{W}$ **then**
 - 3: $V_{\text{RMW}}^r(\tau, x) := \{M(x)\}$
 - 4: **if** $\text{typ}(l) \in \{\text{R}, \text{RMW}, \text{R}^*\}$ **then**
 - 5: $\forall y \neq x : V_{\text{RMW}}^a(\tau, y) := V_{\text{RMW}}^w(x, y)$
 - 6: **if** $\text{mod}(l) \supseteq \text{acq}$ **then**
 - 7: $\forall y \neq x : V_{\text{RMW}}^c(\tau, y) := V_{\text{RMW}}^w(x, y)$
 - 8: **if** $\text{typ}(l) = \text{W}$ **then**
 - 9: $\forall \alpha \in \{\text{a}, \text{c}, \text{r}\}, \pi \neq \tau : V_{\text{RMW}}^\alpha(\pi, x) := \{M(x)\}$
-

Algorithm 3 Thread views maintenance for a step of thread τ labeled with $F(o_f)$

- | | |
|--|--|
| <ol style="list-style-type: none"> 1: if $o_f \supseteq \text{acq}$ then 2: $\forall x : V^c(\tau, x) := V^a(\tau, x)$ 3: $\forall x : V_{\text{RMW}}^c(\tau, x) := V_{\text{RMW}}^a(\tau, x)$ | <ol style="list-style-type: none"> 4: if $o_f \supseteq \text{rel}$ then 5: $\forall x : V^r(\tau, x) := V^c(\tau, x)$ 6: $\forall x : V_{\text{RMW}}^r(\tau, x) := V_{\text{RMW}}^c(\tau, x)$ |
|--|--|
-

Algorithm 4

Location view maintenance for $l = \text{W}(o, x, _)$

- 1: $\forall y \neq x : V^w(y, x) := \{M(x)\}$
 - 2: $\forall y \neq x : V_{\text{RMW}}^w(y, x) := \{M(x)\}$
 - 3: **if** $o \supseteq \text{rel}$ **then**
 - 4: $\forall y \neq x : V^w(x, y) := V^c(\tau, y)$
 - 5: $\forall y \neq x : V_{\text{RMW}}^w(x, y) := V_{\text{RMW}}^c(\tau, y)$
 - 6: **else**
 - 7: $\forall y \neq x : V^w(x, y) := V^r(\tau, y)$
 - 8: $\forall y \neq x : V_{\text{RMW}}^w(x, y) := V_{\text{RMW}}^r(\tau, y)$
-

Algorithm 5

Location view maintenance for $l = \text{RMW}(o, x, _)$

- 1: $\forall y \neq x : V^w(y, x) := \{M(x)\}$
 - 2:
 - 3: **if** $o \supseteq \text{rel}$ **then**
 - 4: $\forall y \neq x : V^w(x, y) := V^c(\tau, y)$
 - 5: $\forall y \neq x : V_{\text{RMW}}^w(x, y) := V_{\text{RMW}}^c(\tau, y)$
 - 6: **else**
 - 7: $\forall y \neq x : V^w(x, y) := V^r(\tau, y)$
 - 8: $\forall y \neq x : V_{\text{RMW}}^w(x, y) := V_{\text{RMW}}^r(\tau, y)$
-

We note that the following invariants always hold:

$$\begin{aligned} V^a(\tau, x) \subseteq V^c(\tau, x) \subseteq V^r(\tau, x) & \quad V_{\text{RMW}}^a(\tau, x) \subseteq V_{\text{RMW}}^c(\tau, x) \subseteq V_{\text{RMW}}^r(\tau, x) \\ \forall \alpha \in \{\text{a}, \text{c}, \text{r}\} : V_{\text{RMW}}^\alpha(\tau, x) \subseteq V^\alpha(\tau, x) & \quad V_{\text{RMW}}^w(y, x) \subseteq V^w(y, x) \end{aligned}$$

Initially, since each location has only one write in the initial graph, we start with V_\emptyset where all components always return the empty set of values. To maintain V , we define a 5-ary relation $\langle M, V \rangle \xrightarrow{\tau, l} V'$ which ascribes how to update V when thread τ takes a step labeled with l starting from memory M . We define this relation using “pseudo-code update algorithms”. For every transition, we first “execute” [Algorithm 1](#) and [Algorithm 2](#) to update V^α and V_{RMW}^α (for $\alpha \in \{\text{a}, \text{c}, \text{r}\}$) or [Algorithm 3](#) for fence steps. Then, V^w and V_{RMW}^w are updated using [Algorithm 4](#) for write steps or [Algorithm 5](#) for RMW steps. In these algorithms we write $X := Y$ for $X := X \cup Y$ and $X := Y$ for $X := X \cap Y$.

[Algorithm 1](#) updates V^c , V^a , and V^r . To understand its steps, consider an access to location x by thread τ . Since we are running under SC, every such access will make τ aware of the most recent write to x , so that there are not any non-robustly readable values of x for thread τ . Accordingly, [Line 1](#) makes $V^c(\tau, x)$ and $V^a(\tau, x)$ empty. If τ writes to x , then the previous value of x in the memory may be non-robustly readable for every other thread that will synchronize with the last

$$\begin{array}{c}
\text{SILENT} \\
\frac{L = \langle \varepsilon, R_{\text{in}}, R_{\text{out}} \rangle \quad R_{\text{in}} \cap R_{\text{taint}} = \emptyset \implies R'_{\text{taint}} = R_{\text{taint}} \setminus R_{\text{out}} \quad R_{\text{in}} \cap R_{\text{taint}} \neq \emptyset \implies R'_{\text{taint}} = R_{\text{taint}} \cup R_{\text{out}}}{\langle M, R_{\text{taint}}, V_{\text{SC}}, V \rangle \xrightarrow{\tau, L}_{\text{SC}_M} \langle M, R'_{\text{taint}}, V_{\text{SC}}, V \rangle} \\
\\
\text{FENCE} \\
\frac{\text{typ}(L) = F \quad \langle M, V \rangle \xrightarrow{\tau, \text{lab}(L)} V'}{\langle M, R_{\text{taint}}, V_{\text{SC}}, V \rangle \xrightarrow{\tau, L}_{\text{SC}_M} \langle M, R_{\text{taint}}, V_{\text{SC}}, V' \rangle} \\
\\
\text{READ} \\
\frac{\text{typ}(L) = R \quad x = \text{loc}(L) \quad \text{val}_R(L) = M(x) \quad x \notin V_{\text{SC}}(\tau) \vee V^c(\tau, x) = \emptyset \quad R'_{\text{taint}} = R_{\text{taint}} \setminus R_{\text{out}}(L) \quad V_{\text{SC}} \xrightarrow{\tau, \text{lab}(L)} V'_{\text{SC}} \quad \langle M, V \rangle \xrightarrow{\tau, \text{lab}(L)} V'}{\langle M, R_{\text{taint}}, V_{\text{SC}}, V \rangle \xrightarrow{\tau, L}_{\text{SC}_M} \langle M, R'_{\text{taint}}, V'_{\text{SC}}, V' \rangle} \\
\\
\text{ERR-SC} \\
\frac{\text{typ}(L) \in \{W, \text{RMW}, R^*\} \quad x = \text{loc}(L) \quad x \in V_{\text{SC}}(\tau) \quad \text{typ}(L) = W \implies V^c_{\text{RMW}}(\tau, x) \neq \emptyset \quad \text{typ}(L) = \text{RMW} \implies \text{val}_R(L) \in V^c_{\text{RMW}}(\tau, x) \quad \text{typ}(L) = R^* \implies \text{val}_R(L) \in V^c(\tau, x)}{\langle M, R_{\text{taint}}, V_{\text{SC}}, V \rangle \xrightarrow{\tau, L}_{\text{SC}_M} \perp} \\
\\
\text{NON-READ} \\
\frac{\text{typ}(L) \in \{W, \text{RMW}, R^*\} \quad x = \text{loc}(L) \quad \text{typ}(L) \in \{\text{RMW}, R^*\} \implies \text{val}_R(L) = M(x) \quad \text{typ}(L) \in \{W, \text{RMW}\} \implies M' = M[x \mapsto \text{val}_W(L)] \quad \text{typ}(L) = R^* \implies M' = M \quad R'_{\text{taint}} = R_{\text{taint}} \setminus R_{\text{out}}(L) \quad V_{\text{SC}} \xrightarrow{\tau, \text{lab}(L)} V'_{\text{SC}} \quad \langle M, V \rangle \xrightarrow{\tau, \text{lab}(L)} V'}{\langle M, R_{\text{taint}}, V_{\text{SC}}, V \rangle \xrightarrow{\tau, L}_{\text{SC}_M} \langle M', R'_{\text{taint}}, V'_{\text{SC}}, V' \rangle} \\
\\
\text{CONTROL} \\
\frac{\text{typ}(L) = \text{CTRL}}{\langle M, R_{\text{taint}}, V_{\text{SC}}, V \rangle \xrightarrow{\tau, L}_{\text{SC}_M} \langle M, R_{\text{taint}}, V_{\text{SC}}, V \rangle} \\
\\
\text{TAINT-READ} \\
\frac{\text{typ}(L) = R \quad x = \text{loc}(L) \quad x \in V_{\text{SC}}(\tau) \quad V^c(\tau, x) \neq \emptyset \quad R'_{\text{taint}} = R_{\text{taint}} \cup R_{\text{out}}(L) \quad V_{\text{SC}} \xrightarrow{\tau, \text{lab}(L)} V'_{\text{SC}}}{\langle M, R_{\text{taint}}, V_{\text{SC}}, V \rangle \xrightarrow{\tau, L}_{\text{SC}_M} \langle M, R'_{\text{taint}}, V'_{\text{SC}}, V \rangle} \\
\\
\text{ERR-DEP} \\
\frac{\text{typ}(L) \in \{W, \text{RMW}, R^*, \text{CTRL}\} \quad R_{\text{in}}(L) \cap R_{\text{taint}} \neq \emptyset}{\langle M, R_{\text{taint}}, V_{\text{SC}}, V \rangle \xrightarrow{\tau, L}_{\text{SC}_M} \perp}
\end{array}$$

Fig. 5. SC_M transitions.

release fence of thread τ (Line 3). For other locations, if τ reads from x it is now confined as the write event view of x ascribes (Lines 5 and 7). A relaxed read will only affect other locations once an acquire fence is placed (only $V^a(\tau, y)$ is restricted). Finally, if τ writes to x , then the previous (overwritten) value of x in the memory is non-robustly readable by all other threads (Line 9).

Algorithm 2 updates V^c_{RMW} , V^a_{RMW} , and V^r_{RMW} . Its steps are exactly the same as Algorithm 1, where the only exceptions are in Line 3 and Line 9, which should not be performed if the access of thread τ is an RMW. In this case, the event writing the current value $M(x)$ will be read by the RMW event added to graph in thread τ , and so, to satisfy **ATOMICITY**, it should not affect the values that are non-robustly writable by other threads.

Algorithm 3 updates the thread views when fences are performed. For an acquire fence, this requires to propagate the acquire views into the current views (Lines 2 and 3); and for a release fence, we propagate the current views into the release views (Lines 5 and 6).

Algorithm 4 updates the location views when non-RMW writes are performed. Lines 1 and 2 reflect the fact that after writing to x by thread τ , the previous value of x in the memory may be non-robustly readable/writable for every other thread that will read from some other location y . In addition, after reading from x , other threads will be confined by what thread τ “releases” in its write: the current view for release writes or the release view for relaxed writes. Algorithm 5 is similar. To support release sequences, Lines 4, 5, 7 and 8, “absorb” the thread view in the location view instead of overwriting it.


Using these algorithms, we can define the (finite) memory system SC_M .

Definition 5.3. The memory system SC_M is given by:

- Its states consists of the error state \perp , and all tuples of the form $\langle M, R_{\text{taint}}, V_{SC}, V \rangle$ (where each of these components is of the form described above).
- Its initial state is $\langle M_\emptyset, R_{\text{taint}\emptyset}, V_{SC\emptyset}, V_\emptyset \rangle$ (all as defined above).
- Its transitions are given in Fig. 5.

The transitions of SC_M in Fig. 5 follow the descriptions above, using the predefined $V_{SC} \rightarrow V'_{SC}$ and $\langle M, V \rangle \rightarrow V'$ relations. In particular, a register is marked as taint in a read step reading from x by thread τ iff some value of x may be non-robustly read by thread τ ($V^c(\tau, x) \neq \emptyset$) while the maximal write to x must be executed before τ 's read step in SC-executions ($x \in V_{SC}(\tau)$). The ERR-SC and ERR-DEP steps detect robustness violations: either by performing a non-robust non-read transition, or by depending on a taint read.

Example 5.4. We demonstrate how SC_M simulates SC_T for establishing the non-observational-robustness of the (non-fenced) MP program in Ex. 4.2 by reaching the witness discussed in Ex. 4.15. The following illustration depicts three program states (top part) during a run, the corresponding execution graphs in the run of SC_T (middle part; the set T consists of the highlighted events), and the corresponding states in the run of SC_M (bottom part). To simplify the presentation, we only present the part of SC_M 's states that is relevant for detecting the robustness violation (which is relatively small in this example).



To reach the first (leftmost) state, we run the first thread until it completes its execution. At this point we have $x = y = 1$ in the memory, no tainted registers ($R_{\text{taint}} = \emptyset$), the second thread is not “ hb_{SC} -aware” of the maximal write to x ($x \notin V_{SC}.V(T_2)$), and the second thread may non-robustly read (i.e., can read under RC20 but not under SC if we continue from this state) the value 0 from x ($V^c(T_2, x) = \{0\}$). Then, since we are running under (instrumented) SC, the second thread has to read 1 from y . This does not change the memory or the set of tainted registers, but it does make the thread “ hb_{SC} -aware” of the maximal write to x (via the added rf -edge on y). Hence, we have $x \in V_{SC}.V(T_2)$ in the second state. Next, when the second thread reads x , it marks b as a tainted register—the thread is “ hb_{SC} -aware” of the maximal write to x but can non-robustly read some value of x ($V^c(T_2, x) = \{0\}$). Finally, when we use the value of the tainted register b (by executing the assertion) a robustness violation is detected via SC_M 's ERR-DEP transition.

If we include a release fence before $y.\text{store}(1, r1x)$ in the first thread and an acquire fence before $b := x.\text{load}(r1x)$ in the second thread, the register b will not become tainted, and the program will be (correctly) deemed as robust:

- After the first thread executes the release fence, we will have $V^r(\tau_1, x) = \emptyset$;
- Then, after the first thread executes $y.\text{store}(1, r1x)$, we will have $V^w(y, x) = \emptyset$;
- Then, after the second thread reads 1 from y , we will have $V^a(\tau_2, x) = \{0\} \cap \emptyset = \emptyset$;
- Then, after the second thread executes the acquire fence, we will have $V^c(\tau_2, x) = \emptyset$;
- Then, when the second thread reads x , it cannot non-robustly read any value of x , so the register b remains untainted.

Next, we establish the equivalence of SC_T and its finite abstraction SC_M , and derive our final results as simple corollaries of [Theorems 4.16, 5.2 and 5.5](#) and [Corollary 5.6](#).

THEOREM 5.5. *A state $\langle \bar{q}, \perp \rangle$ is reachable in $Pr \parallel \text{SC}_T$ iff it is reachable in $Pr \parallel \text{SC}_M$.*

COROLLARY 5.6. *There exists a non-observational-robustness witness for Pr iff $\langle \bar{q}, \perp \rangle$ is reachable in $Pr \parallel \text{SC}_M$ for some $\bar{q} \in Pr.Q$.*

COROLLARY 5.7. *If $\langle \bar{q}, \perp \rangle$ is not reachable in $Pr \parallel \text{SC}_M$ for every $\bar{q} \in Pr.Q$, then Pr is observationally robust.*

Finally, a decision procedure for robustness follows from [Thm. 4.6](#), [Lemma 4.20](#), and [Corollary 5.6](#).

COROLLARY 5.8. *A program Pr is robust iff $\langle \bar{q}, \perp \rangle$ is not reachable in $Pr' \parallel \text{SC}_M$ for every $\bar{q} \in Pr'.Q$, where Pr' is the program obtained from Pr by adding a (vacuous) assert instruction $\text{assert}(r = r)$ after every read instruction $r := x.\text{load}(o_R)$.*

It also follows that robustness verification against RC20 is PSPACE-complete. The upper bound follows by reduction to reachability in an instrumented SC semantics; whereas the lower bound directly follows from [\[Lahav and Margalit 2019\]](#), as the model studied there is a fragment of RC20.

6 EXTENSION WITH NON-ATOMICS

The extension of our results to cover C11-style non-atomics with “catch-fire” semantics is straightforward: we can identify data-races on non-atomics while checking for robustness. Next, we describe this extension. We refer to the extended model as RC20+NA.

First, we include “na” in the set Mod of access modes, where na is included both in Mod_R (read access modes) and in Mod_W (write access modes). We assume a set Loc_{na} of *non-atomic locations* disjoint from the set Loc of *atomic locations*, and require that all accesses to the locations in Loc_{na} are na-accesses, and conversely, that all na-accesses are to locations in Loc_{na} . We further assume that non-atomic accesses are used in plain loads and stores, but not in RMWs or wait instructions.

To give semantics to programs with non-atomic accesses, we define *data race* on non-atomics, and consider consistent execution graphs that have such races as program failures.

Definition 6.1. Two events a and b are called *conflicting* in an execution graph G if $a, b \in G.E$, $G.\text{loc}(a) = G.\text{loc}(b)$, and $W \in \{G.\text{typ}(a), G.\text{typ}(b)\}$. A pair $\langle a, b \rangle$ is called a *race* in G if a and b are conflicting events in G and $\langle a, b \rangle \notin G.\text{hb} \cup G.\text{hb}^{-1}$. An execution graph G is called *racy* if there is some race $\langle a, b \rangle$ in G with $na \in \{G.\text{mod}(a), G.\text{mod}(b)\}$.

Definition 6.2. A program Pr may fail under RC20+NA if either (i) it may fail under RC20 ([Def. 3.12](#)) or (ii) some RC20-consistent execution graph G generated by Pr is racy.

Our goal is to extend SC_T , and, in turn, SC_M , so that these systems can be used in order to establish the fact that a given program may fail under RC20+NA iff it may fail under SC. We note

that the definition of failure under SC is unchanged: for the SC memory system non-atomic accesses are not at all distinguished from atomic ones. To do so, we apply a simple program transformation:

- (1) We split every $x \in \text{Loc}_{n_a}$ into $|\text{Tid}|$ variables, $x_1, \dots, x_{|\text{Tid}|}$, one for each thread.
- (2) We replace every (non-atomic) write to x with a sequence of $|\text{Tid}|$ (non-atomic) writes writing to $x_1, \dots, x_{|\text{Tid}|}$ the value that was written to x .
- (3) We replace every (non-atomic) read operation from x of thread T_i with a read from x_i followed by a write (of the read value that was read) to x_i .

We denote by $\text{trans}(Pr)$ the program obtained by applying this transformation on Pr .

Example 6.3. Applying the above transformation to the program on the left generates the program on the right:

$$\begin{array}{c}
 a := x.\text{load}(na) \\
 y.\text{store}(1, r1x)
 \end{array}
 \left\| \begin{array}{c}
 \text{wait}(y = 1, r1x) \\
 b := x.\text{load}(na) \\
 y.\text{store}(2, rel)
 \end{array} \right\|
 \left\| \begin{array}{c}
 \text{wait}(y = 2, acq) \\
 x.\text{store}(1, na)
 \end{array} \right\|
 \rightsquigarrow
 \begin{array}{c}
 a := x_1.\text{load}(na) \\
 x_1.\text{store}(a, na) \\
 y.\text{store}(1, r1x)
 \end{array}
 \left\| \begin{array}{c}
 \text{wait}(y = 1, r1x) \\
 b := x_2.\text{load}(na) \\
 x_2.\text{store}(b, na) \\
 y.\text{store}(2, rel)
 \end{array} \right\|
 \left\| \begin{array}{c}
 \text{wait}(y = 2, acq) \\
 x_1.\text{store}(1, na) \\
 x_2.\text{store}(1, na) \\
 x_3.\text{store}(1, na)
 \end{array} \right\|$$

This program transformation essentially ensures that all races are expressed as write-write races. We use per-thread variables to make sure that two concurrent reads to the same location in the original program will not induce a race in the transformed program. Applying this transformation allows us to detect races by checking that whenever a thread executes a non-atomic access it is already aware of the latest write to the same (non-atomic) location. Since our instrumented semantics is designed to track the set of locations of which every is aware to the latest write, this approach requires only minimal changes in SC_T and SC_M . Formally, we have the following:

LEMMA 6.4. *The following are equivalent for every concurrent program Pr :*

- (i) *Some RC20-consistent graph generated by Pr is racy.*
- (ii) *There exists some RC20-consistent execution graph G that is generated by $\text{trans}(Pr)$ and a $(G.\text{po} \cup G.\text{rf})$ -maximal event $e \in G.E$ such that $G.\text{mod}(e) = na$ and $G'.w_{G'.\text{loc}(e)}^{\max} \notin \text{dom}(G'.\text{rf}^?; G'.\text{hb}; [E^{\tau}])$ for $G' = G \setminus \{e\}$.*

Condition (ii) above is similar to the condition that SC_T (and, in turn, SC_M) monitors, so that by using the transformation above we do not need to change the instrumentation in SC_T (and SC_M). Indeed, it suffices to add the following transition to SC_T (we refer to the extend system as SC_T^{na}):

$$\frac{\text{mod}(L) = na \quad G.w_{\text{loc}(L)}^{\max} \notin \text{dom}(G.\text{rf}_T^?; G.\text{hb}_T; [E^{\tau}])}{\langle G, \psi, T \rangle \xrightarrow{\tau, L}_{\text{SC}_T^{\text{na}}} \perp}$$

This additional transition allows the system to move to \perp when a race is detected.

THEOREM 6.5. *If $\langle \bar{q}, \perp \rangle$ is not reachable in $\text{trans}(Pr) \parallel \text{SC}_T^{\text{na}}$ for every $\bar{q} \in \text{trans}(Pr).Q$, then Pr may fail under RC20+NA iff Pr may fail under SC.*

Finally, we extend SC_M with the following step that precisely matches the step added to SC_T (we refer to the extend system as SC_M^{na}):

$$\frac{\text{mod}(L) = na \quad V^c(\tau, \text{loc}(L)) \neq \emptyset}{\langle M, R_{\text{taint}}, V_{\text{SC}}, V \rangle \xrightarrow{\tau, L}_{\text{SC}_M^{\text{na}}} \perp}$$

The simulation argument relating SC_T and SC_M is easily extended to relate SC_T^{na} and SC_M^{na} , which provides us with our final result:

COROLLARY 6.6. *If $\langle \bar{q}, \perp \rangle$ is not reachable in $\text{trans}(Pr) \parallel \text{SC}_M^{\text{na}}$ for every $\bar{q} \in \text{trans}(Pr).Q$, then Pr may fail under RC20+NA iff Pr may fail under SC.*

7 IMPLEMENTATION AND EVALUATION

We have implemented our algorithm as an extension of “Rocker” [Lahav and Margalit 2019]—a robustness checking tool against the RA model (a fragment of RC20).⁵ The implementation transforms an input program in a minimal C-like language to an instrumented Promela program that can be verified by the SPIN model checker [Holzmann 1997]. The instrumentation follows the description above, but instead of checking reachability in an instrumented SC semantics, we instrument the program itself, and equivalently check for reachability of a state representing an observational robustness violation under the usual SC interleaving semantics.

Remark 3. Similarly to what was done in [Lahav and Margalit 2019], as an implementation optimization, it is possible to group together different values, making the range of the different views in the instrumented semantics to be the powerset of some fixed partition of Val rather than the powerset of Val (which corresponds to the powerset of the trivial partition $\{\{v\} \mid v \in \text{Val}\}$). The selected partition should be sufficiently expressive for checking whether enabled successful CAS/BCAS/wait instructions may read a stale value. In our implementation, for every location x , we use a partition of the form $\{\{v\} \mid v \in \text{Crit}_x\} \cup \{\text{Val} \setminus \text{Crit}_x\}$, where Crit_x are all values that successful CAS/BCAS/wait instructions of location x may read. (We require the user to annotate the input program with this information.) In particular, for programs that have no CAS/BCAS/wait instructions at all, all views are boolean predicates (rather than functions to $\mathcal{P}(\text{Val})$).

We have performed a series of experiments on litmus tests, examples from [Lahav and Margalit 2019], and additional algorithms that have been shown to be challenging-to-get-right. Figure 6 summarizes our evaluation results when run on an Intel®Core™i5-6300U CPU @2.40GHz GNU/Linux machine. In several cases, for verifying robustness, we need to replace busy-loops with wait instructions or use BCAS instead of CAS. Clearly, such transformations preserve the program semantics. The columns “Program” and “#T” show the benchmark name and its number of threads. The column “Rob” indicates whether the program is robust (“R”) and observationally robust (“O”). The “Time” column presents the time that was needed to automatically verify robustness for each method (i.e., detecting non-robustness witnesses vs. detecting non-observational-robustness witnesses), as well as the running time of SPIN on the given program without any instrumentation (in the column named “SC”). The latter may shed light on the size of the each verification problem and the overhead required for robustness verification w.r.t. standard explicit model checking. The number in parentheses indicates the percentage of the time used for compiling the verifier produced by SPIN (using gcc -O2), which, in some cases, dominates the overall search time. We note that the time needed for Rocker to create the instrumented program and for SPIN to create the verifier are negligible (< 0.3s).

We interpret these results as showing that the time difference between merely checking robustness vs. observational robustness is rather small. Naturally, for non-robust but observationally robust programs, we have a significant gap.

For comparison with existing work, we have run our tool on programs that employ only release/acquire accesses from [Lahav and Margalit 2019]. Due to the more fine-grained instrumentation required to handle the full RC20 model, we obtain slightly increased runtime compared to the RA-only robustness verifier. This can be seen in Fig. 7 (all programs there are robust against RA).

Next, we describe some of the verified examples and our findings:

arc: An implementation of Rust’s *Atomic Reference Counter* [Doko and Vafeiadis 2017; Rust 2020] used by three threads repeatedly going online and offline. It relies on the synchronization induced by relaxed RMWs inside release sequences, and our tool verified the robustness of this mechanism.

⁵Our implementation and the examples it was tested on are available in the artifact accompanying this paper.

Program	#T	Rob		Time (sec)		
		R	O	R	O	SC
arc	3	✓	✓	67.1 (8%)	71.2 (8%)	3.9 (26%)
peterson	2	✗	✗	1.1 (100%)	1.1 (100%)	0.9 (100%)
peterson-fix1	2	✓	✓	1.1 (100%)	1.1 (100%)	0.9 (100%)
peterson-fix2	2	✓	✓	1.1 (100%)	1.1 (100%)	0.9 (100%)
singleton	4	✓	✓	2.8 (96%)	3.0 (97%)	1.1 (100%)
singleton-rlx	4	✓	✓	2.8 (96%)	3.1 (97%)	1.1 (100%)
wait-free-ring	2	✓	✓	0.9 (100%)	1.0 (100%)	0.9 (100%)
dekker-rlx	2	✓	✓	1.5 (100%)	1.6 (100%)	0.9 (100%)
seqlock rdmw	3	✗	✓	7.1 (100%)	59.7 (12%)	30.9 (4%)
seqlock fence	3	✗	✓	8.0 (100%)	81.8 (10%)	39.9 (3%)
seqlock rdmw-rw	3	✗	✓	6.4 (100%)	53.5 (13%)	31.9 (3%)
seqlock fence-rw	3	✗	✓	6.7 (100%)	74.6 (9%)	41.5 (3%)
chase-lev	3	✓	✓	57.7 (2%)	58.8 (3%)	27.1 (4%)
lock exchange	2	✓	✓	1.0 (100%)	1.0 (100%)	0.9 (100%)
spinlock4-rlx	4	✓	✓	2.5 (52%)	2.9 (48%)	1.3 (77%)

Fig. 6. Evaluation results

Program	#T	Time (sec)		
		R	O	Rocker
peterson-ra	2	1.1 (100%)	1.1 (100%)	1.0 (100%)
dekker-ra	2	1.4 (100%)	1.4 (100%)	1.1 (100%)
seqlock-ra	3	62.2 (10%)	59.1 (10%)	58.3 (8%)
chase-lev-ra	3	59.8 (3%)	55.7 (3%)	39.9 (3%)
spinlock4-ra	4	2.4 (63%)	2.5 (60%)	2.3 (61%)
lamport2-3	3	102.3 (9%)	95.5 (10%)	69.9 (8%)
rcu	4	54.7 (5%)	46.1 (6%)	52.0 (4%)
rcu-offline	3	80.3 (13%)	77.5 (14%)	61.4 (12%)
ticketlock4	4	12.1 (12%)	11.8 (14%)	11.7 (13%)

Fig. 7. Comparison with *Rocker* [Lahav and Margalit 2019] on robust RA programs

peterson: An implementation of Peterson’s lock for C11 by V’jukov [2008]. We (automatically) identified a robustness violation when one of the threads repeatedly tries to enter the critical section. Previous verification efforts (including an informal argument for correctness [Williams 2008], testing [V’jukov 2013], as well as a formal proof in a dedicated program logic [Dalvandi et al. 2020; Doherty et al. 2019]) focused on just a single attempt to enter the critical section by each thread, and, thus, did not detect this issue. We also note that the fact that both threads cannot simultaneously be in the critical section (which previous work established for a single critical section by each thread) does not suffice in C11: we need locked regions to be ordered by *hb*. We propose (and we verified) two fixes to the implementation: [fix1] Promoting a certain relaxed read to acquire; and [fix2] Placing an additional acquire fence before entering the critical section (and then we can demote an existing acquire read to be relaxed).

singleton: A double-checked locking pattern from the “boost” library [Bahmann and Blechmann 2012]. We found that robustness (and, consequently, correctness) still hold even if we demote the read inside the lock to be relaxed (thus suggesting a possible performance improvement).

wait-free-ring: A wait-free ring buffer with a single producer and a single consumer taken from the “boost” library [Bahmann and Blechmann 2012].

dekker: An optimized (rather tricky) implementation by Williams [2010] of the classical Dekker’s mutual exclusion algorithm that employs relaxed accesses and SC fences.

seqlocks: Optimized implementations of Seqlock (see §4.1) based on [Boehm 2012] used by three threads (which non-deterministically write and read from the protected data structure). These programs are observationally robust, but not robust. We verified the two implementations of Boehm [2012]: one using “read-don’t-modify-write” instructions (“seqlock rdmw”), and one using acquire fences (“seqlock fence”). We also tried a writer that uses efficient relaxed writes with appropriate release fences (no “rw” suffix), as well as a less efficient writer using release writes (“rw” suffix).

chase-lev: An optimized work stealing double-ended queue that uses relaxed accesses [Chase and Lev 2005; Kang 2018]. The queue owner inserts and removes from the bottom of the queue. Stealers steal from the top of the queue. In this case, to prove robustness, we had to make the write to “bot” (the bottom index) in the owner pop function to be release. We also had to add SC-fences in the owner push function between the relaxed read of “bot” and the read acquire of “top”, as well as before the stealer starts stealing (to allow stealing loops). In this case, these strengthenings are due to limitations of the robustness correctness criterion, rather than bugs in the implementation.

8 RELATED WORK AND CONCLUSIONS

We provided a sound and precise robustness verification method against RC20, a slight variant of the RC11 memory model, which includes release/acquire and relaxed reads, writes, and RMWs, as well as release/acquire fences. To support speculative relaxed reads as used in seqlock algorithms, we introduced an observational robustness notion, and showed how it can be soundly verified. We implemented our method and evaluated it on multiple examples, demonstrating the effectiveness of this approach for detecting subtle bugs during the development of concurrent algorithms (or migrating algorithms written for SC) for a rather complex weak memory model.

Previous work studied robustness against hardware models, e.g., [Alglave et al. 2017; Alglave and Maranget 2011; Burckhardt et al. 2007; Derevenetc and Meyer 2014], and mostly against the x86-TSO model, e.g., [Abdulla et al. 2015b,a; Bouajjani et al. 2013, 2018, 2011; Burckhardt and Musuvathi 2008; Burnim et al. 2011; Gotsman et al. 2012; Linden and Wolper 2011, 2013; Liu et al. 2012; Owens 2010]. From a complexity perspective, our results show that verification of robustness against RC20 is similar to verification of robustness against TSO—they are both PSPACE-complete. Nevertheless, compared to TSO, the technical challenge here stems from the lack of (traditional) operational semantics, delicate synchronization definition, and rather weak RMW operations.

A (very imprecise) robustness criteria against a *programming language* memory model emerges from the well-known DRF guarantee [Adve and Hill 1990; Gharachorloo et al. 1992], and our “non-robustness witnesses” can be viewed as a precise DRF condition. More recently, as mentioned above, Lahav and Margalit [2019] provided a precise robustness analysis for release/acquire semantics. Our work builds and extends this result to include relaxed accesses (with acyclic($po \cup rf$)) and release/acquire fences. Observational robustness is, to the best of our knowledge, a novel notion, arising when analyzing the use case of relaxed accesses as in the seqlock algorithm.

For shared memory distributed systems, robustness of transactional programs against serializability was studied in [Beillahi et al. 2019a,b; Bernardi and Gotsman 2016; Brutschy et al. 2018; Fekete et al. 2005; Nagar and Jagannathan 2018]. Except for [Beillahi et al. 2019a,b], these papers do not provide precise verification methods, but rather practical over-approximations.

ACKNOWLEDGMENTS

We thank the POPL’21 reviewers for their helpful feedback. This research was supported by the Israel Science Foundation (grant number 5166651). The second author was also supported by the Alon Young Faculty Fellowship.

REFERENCES

- Parosh Aziz Abdulla, Jatin Arora, Mohamed Faouzi Atig, and Shankaranarayanan Krishna. 2019. Verification of programs under the release-acquire semantics. In *PLDI*. ACM, New York, NY, USA, 1117–1132. <https://doi.org/10.1145/3314221.3314649>
- Parosh Aziz Abdulla, Mohamed Faouzi Atig, Magnus Lång, and Tuan Phong Ngo. 2015b. Precise and sound automatic fence insertion procedure under PSO. In *NETYS*. Springer International Publishing, Cham, 32–47.
- Parosh Aziz Abdulla, Mohamed Faouzi Atig, and Tuan-Phong Ngo. 2015a. 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
- 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>
- 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>
- 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>
- 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>
- 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>
- 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
- Helge Bahmann and Tim Blechmann. 2012. Atomic usage examples. Retrieved June 08, 2020 from https://www.boost.org/doc/libs/1_55_0/doc/html/atomic/usage_examples.html
- 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>
- Sidi Mohamed Beillahi, Ahmed Bouajjani, and Constantin Enea. 2019a. Checking robustness against snapshot isolation. In *Computer Aided Verification*. Springer International Publishing, Cham, 286–304.
- Sidi Mohamed Beillahi, Ahmed Bouajjani, and Constantin Enea. 2019b. Robustness against transactional causal consistency. In *CONCUR 2019*, Vol. 140. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 30:1–30:18. <https://doi.org/10.4230/LIPIcs.CONCUR.2019.30>
- 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/LIPIcs.CONCUR.2016.7>
- 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>
- Hans-J. Boehm and Brian Demsky. 2014. Outlawing ghosts: avoiding out-of-thin-air results. In *MSPC*. ACM, New York, NY, USA, Article 7, 6 pages. <https://doi.org/10.1145/2618128.2618134>
- Hans-J Boehm, Olivier Giroux, and Viktor Vafeiades. 2018. P0982R1: Weaken release sequences. <http://www.openstd.org/jtc1/sc22/wg21/docs/papers/2018/p0982r1.html>. Accessed: 2020-05-27.
- 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
- 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.
- Ahmed Bouajjani, Roland Meyer, and Eike Möhlmann. 2011. Deciding robustness against total store ordering. In *ICALP*. Springer, Berlin, Heidelberg, 428–440.
- 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>
- 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>
- 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
- Jabob 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.
- David Chase and Yossi Lev. 2005. Dynamic circular work-stealing deque. In *SPAA*. ACM, New York, NY, USA, 21–28. <https://doi.org/10.1145/1073970.1073974>

- Sadegh Dalvandi, Simon Doherty, Brijesh Dongol, and Heike Wehrheim. 2020. Owicki-Gries Reasoning for C11 RAR. In *ECOOP*, Vol. 166. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 11:1–11:26. <https://doi.org/10.4230/LIPIcs.ECOOP.2020.11>
- Hoang-Hai Dang, Jacques-Henri Jourdan, Jan-Oliver Kaiser, and Derek Dreyer. 2019. RustBelt meets relaxed memory. *Proc. ACM Program. Lang.* 4, POPL, Article 34 (Dec. 2019), 29 pages. <https://doi.org/10.1145/3371102>
- 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>
- Egor Derevenetc and Roland Meyer. 2014. Robustness against Power is PSpace-complete. In *ICALP*. Springer, Berlin, Heidelberg, 158–170.
- 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>
- Marko Doko and Viktor Vafeiadis. 2017. Tackling real-life relaxed concurrency with FSL++. In *ESOP*. Springer Berlin Heidelberg, Berlin, Heidelberg, 448–475.
- 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>
- Shaked Flur, Kathryn E. Gray, Christopher Pulte, Susmit Sarkar, Ali Sezgin, Luc Maranget, Will Deacon, and Peter Sewell. 2016. Modelling the ARMv8 architecture, operationally: concurrency and ISA. In *POPL*. ACM, New York, 608–621. <https://doi.org/10.1145/2837614.2837615>
- 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\(92\)90052-0](https://doi.org/10.1016/0743-7315(92)90052-0)
- 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
- Gerard J. Holzmann. 1997. The model checker SPIN. *IEEE Transactions on software engineering* 23, 5 (1997), 279–295.
- Jeehoon Kang. 2018. Deque proof. Retrieved June 11, 2020 from <https://github.com/jeehoonkang/crossbeam-rfcs/blob/deque-proof/text/2018-01-07-deque-proof.md>
- 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>
- 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>
- Ori Lahav and Roy Margalit. 2019. Robustness against release/acquire semantics. In *PLDI*. ACM, New York, NY, USA, 126–141. <https://doi.org/10.1145/3314221.3314604>
- 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>
- 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>
- 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/10.1007/978-3-642-36742-7_24
- Feng Liu, Nayden Nedev, Nedyalko Prisadnikov, 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>
- Roy Margalit and Ori Lahav. 2020. Verifying observational robustness against a C11-style memory model. https://www.cs.tau.ac.il/~orilahav/papers/popl21_robustness_full.pdf (full version of this paper).
- Kartik Nagar and Suresh Jagannathan. 2018. Automated detection of serializability violations under weak consistency. In *CONCUR*, Vol. 118. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 41:1–41:18. <https://doi.org/10.4230/LIPIcs.CONCUR.2018.41>
- Scott Owens. 2010. Reasoning about the implementation of concurrency abstractions on x86-TSO. In *ECOOP*. Springer-Verlag, Berlin, Heidelberg, 478–503.
- Anton Podkopaev, Ori Lahav, and Viktor Vafeiadis. 2019. Bridging the Gap Between Programming Languages and Hardware Weak Memory Models. *PACMPL* 3, POPL, Article 69 (Jan. 2019), 31 pages. <https://doi.org/10.1145/3290382>
- Rust. 2020. std::sync::Arc. Retrieved June 11, 2020 from <https://doc.rust-lang.org/src/alloc/sync.rs.html#213-216>
- Matthew D. Sinclair, Johnathan Alsop, and Sarita V. Adve. 2017. Chasing away RAs: semantics and evaluation for relaxed atomics on heterogeneous systems. In *ISCA*. ACM, New York, NY, USA, 161–174. <https://doi.org/10.1145/3079856.3080206>
- 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>
- Dmitriy V’jukov. 2008. C++ atomics and memory ordering. Retrieved June 23, 2020 from <https://bartoszmilewski.com/2008/12/01/c-atomics-and-memory-ordering/#comment-111>

- Dmitriy V'jukov. 2013. Relacy race detector. Retrieved July 06, 2020 from <http://www.1024cores.net/home/relacy-race-detector>
- 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
- Anthony Williams. 2010. Implementing Dekker's algorithm with fences. Retrieved June 08, 2020 from https://www.justsoftwaresolutions.co.uk/threading/implementing_dekkers_algorithm_with_fences.html