



Automated Robustness Verification of Concurrent Data Structure Libraries against Relaxed Memory Models

KARTIK NAGAR, IIT Madras, India

ANMOL SAHOO, Purdue University, USA

ROMIT ROY CHOWDHURY, Chennai Mathematical Institute, India

SURESH JAGANNATHAN, Purdue University, USA

Clients reason about the behavior of concurrent data structure libraries such as sets, queues, or stacks using specifications that capture well-understood correctness conditions, such as linearizability. The implementation of these libraries, however, focused as they are on performance, may additionally exploit relaxed memory behavior allowed by the language or underlying hardware that weaken the strong ordering and visibility constraints on shared-memory accesses that would otherwise be imposed by a sequentially consistent (SC) memory model. As an alternative to developing new specification and verification mechanisms for reasoning about libraries under relaxed memory model, we instead consider the orthogonal problem of *library robustness*, a property that holds when all possible behaviors of a library implementation under relaxed memory model are also possible under SC. In this paper, we develop a new *automated* technique for verifying robustness of library implementations in the context of a C11-style memory model. This task is challenging because a most-general client may invoke an unbounded number of concurrently executing library operations that can manipulate an unbounded number of shared locations. We establish a novel inductive technique for verifying library robustness that leverages prior work on the robustness problem for the C11 memory model based on the search for a non-robustness witness under SC executions. We crucially rely on the fact that this search is carried out over SC executions, and use high-level SC specifications (including linearizability) of the library to verify the absence of a non-robustness witness. Our technique is compositional - we show how we can safely preserve robustness of multiple interacting library implementations and clients using additional SC fences to guarantee robustness of entire executions. Experimental results on a number of complex realistic library implementations demonstrate the feasibility of our approach.

CCS Concepts: • **Theory of computation** → **Program verification**; *Concurrency*.

Additional Key Words and Phrases: Relaxed Memory Models, Concurrent Library implementations, Robustness

ACM Reference Format:

Kartik Nagar, Anmol Sahoo, Romit Roy Chowdhury, and Suresh Jagannathan. 2024. Automated Robustness Verification of Concurrent Data Structure Libraries against Relaxed Memory Models. *Proc. ACM Program. Lang.* 8, OOPSLA2, Article 362 (October 2024), 28 pages. <https://doi.org/10.1145/3689802>

1 Introduction

Unlike a sequentially consistent (SC) [41] memory model where concurrently executing threads all witness the same view of memory, relaxed memory models such as those found in C11 [7, 8, 15] allow each thread to have its own logical view of memory, resulting in behaviors like store buffering that are not expressible under SC [37]. A program written to take advantage of the performance benefits

Authors' Contact Information: [Kartik Nagar](#), IIT Madras, Chennai, India, nagark@cse.iitm.ac.in; [Anmol Sahoo](#), Purdue University, West Lafayette, USA, sahoo9@purdue.edu; [Romit Roy Chowdhury](#), Chennai Mathematical Institute, Chennai, India, romit@cmi.ac.in; [Suresh Jagannathan](#), Purdue University, West Lafayette, USA, suresh@cs.purdue.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2024 Copyright held by the owner/author(s).

ACM 2475-1421/2024/10-ART362

<https://doi.org/10.1145/3689802>

afforded by executing under a relaxed memory model M , but whose behaviors can nonetheless be fully explained in terms of SC (interleaved) executions, is said to be *robust* against M . Intuitively, a robust program is one in which every relaxed memory execution is also an SC one.

In this paper, we consider the problem of *automatically* verifying the robustness of library implementations against RC20 [43], a useful variant of the C11 memory model. C11 allows the expression of high-performance concurrent C/C++ code through the use of annotations affixed to memory accesses. Release (rel) and acquire (acq) annotations enable a form of message-passing between “release” writes and “acquire” reads [38], yielding a causally-consistent ordering between the operations that precede the “release” and those that follow the “acquire”; and, relaxed (rlx) annotations are used to compile accesses to single hardware loads and stores with no additional synchronization, other than basic store coherence guarantees [33] provided by the hardware. RC20 provides an improved semantics for atomic accesses [40], prevents undesirable executions involving out-of-thin-air reads, and simplifies the semantics of synchronization for release-acquire pairs.

Robustness under the RC20 memory model is a well-studied problem [38, 43] for which algorithms to check robustness of finite-state programs operating over a bounded number of threads have been developed. By showing a program to be robust, one can then reason about its correctness w.r.t. a high-level specification under the SC memory model, which is a significantly simpler problem. Consequently, any correctness guarantees proven under SC also continue to hold under RC20. Thus, proving library robustness can also enable a similar pathway to reasoning about correctness of programs using library implementations under RC20. Prior efforts addressing the correctness of library implementations under relaxed memory models have mostly focused on developing newer forms of specification [18, 44, 49, 52] that allows reasonable non-SC behaviors and also exposes the internals of the underlying memory model. While these specifications can admit highly efficient library implementations, proving their correctness requires significant manual effort because of the need to closely correlate abstract state-based specifications with the event-based relaxed memory model semantics. Furthermore, using these relaxed specifications to prove correctness of client programs using the library implementation is also highly non-trivial. On the other hand, by proving a library implementation to be robust in the context of the most general client, we can extend this robustness guarantee to any client program using the library, thus allowing one to reason about its correctness under SC, which is a well-studied problem.

Following the classical modular verification paradigm, we would like to separately verify the robustness of library implementations and their clients and then compose these guarantees together to establish robustness of the overall execution. We consider the following strategy: first verify a library implementation to be robust in the presence of a *most general client*, which can call the library methods an arbitrary number of times across an arbitrary number of threads, and then use this robustness guarantee to prove the robustness of any client program potentially calling multiple libraries. Clearly, if a library is not robust on its own, any client program using the library (and which does not restrict the concurrent behaviors of the library) would also be non-robust.

Unfortunately, we find that naïvely composing individually robust libraries may not lead to an overall robust execution. To illustrate this, consider the two simple register libraries \mathcal{L}_1 and \mathcal{L}_2 given in Fig. 1. It is well known that under the RC20 memory model, atomic accesses to a single location follow SC-per-location semantics, and thus both the libraries (which access single distinct locations ℓ_1 and ℓ_2) will be individually robust for the most general client. However,




Fig. 1. Example showing non-compositionality of robustness

we can construct a standard store-buffering execution formed out of interaction with both libraries, as shown by the sessions τ_1 and τ_2 in the figure (assume that the initial value at locations ℓ_1 and ℓ_2 is 0). This execution is not possible under SC, because in any interleaving, at least one of set1 or set2 (and hence the underlying store) operations will happen first, so that at least one of the get1 or get2 operations would return 1. To recover overall robustness, we can add an SC-fence when crossing libraries within every session. In the RC20 memory model, an SC-fence is defined using a composition of three instructions: `fence(acq); fadd(f , 0, acqrel); fence(rel)` where f is an identifier for a distinguished location not accessed by any other instruction [43]. The `fadd` instruction guarantees synchronization between any two SC fences.

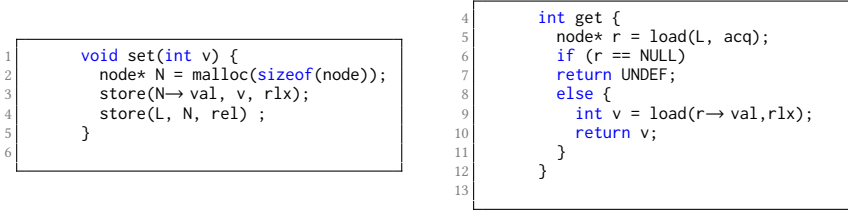
In the execution above, suppose we place an SC fence between the two invocations in each session. Then, the total ordering among the fences would prevent both the loads from ignoring the stores under RC20, thus prohibiting the store buffering anomaly. On the other hand, if even one of the SC-fences is removed, store buffering is still possible. We call an execution *well-fenced* if there is an SC-fence between every adjacent pair of invocations to different libraries in every session.

Having well-fenced executions then gives us a way to decompose the overall robustness problem into robustness of individual libraries. In this work, we leverage this observation and make significant progress on the problem of verifying robustness of programs using library implementations: (i) we propose an automated, sound (but not complete) verification procedure for proving robustness of individual library implementations against the most general client, (ii) we formally prove that a well-fenced execution comprising of calls to individually robust libraries is guaranteed to be robust.

Verifying individual library robustness conceptually requires us to compare the set of all SC and RC20 executions of the library implementation. This is a non-trivial problem because the most general client can invoke library methods an arbitrary number of times concurrently across an arbitrary number of threads. Previous efforts on robustness verification of relaxed memory programs [11, 38, 43] are only applicable for finite state programs involving a bounded number of threads. Our primary contribution in this regard is a fully automated verification strategy that is not constrained by these restrictions.

Our verification strategy is based on searching SC executions of a library implementation for potential *non-robustness witnesses* [38, 43]. SC executions are constructed by simply treating all memory accesses as being sequentially consistent, ignoring the actual annotations affixed to these accesses in the program. A non-robustness witness is a prefix of an SC execution that contains a location whose latest write w would be witnessed by the next operation (say a read r) in that execution, but which may not necessarily be witnessed in an RC20 execution, i.e. informally, $\text{hb}_{\text{SC}}(w, r)$ holds but $\text{hb}_{\text{RC20}}(w, r)$ does not, where hb_{SC} and hb_{RC20} define a happens-before ordering among read and write actions in SC and RC20 executions, resp. Verifying the absence of such a witness is tantamount to showing robustness in a RC20 setting [38, 43]. We use an inductive strategy to cover the infinite set of executions, and crucially take advantage of axioms and constraints provided by the library's SC specification to help discharge the verification conditions.

We note that unlike prior approaches to robustness verification [38, 43] which are both sound and complete, our approach sacrifices completeness for the sake of automation by devising sufficient but not necessary conditions for proving absence of a non-robustness witness. However, this makes our verification strategy more amenable to automation, and also better suited to library implementations. We also crucially rely on expressive SC specifications to rule out infeasible non-robustness witnesses, with the caveat that spurious non-robustness witnesses can arise if the SC specification is not expressive enough. However, our extensive experimental evaluation suggests that the SC specification of the standard data structures is rich enough to verify a number of complex implementations.



```

1 void set(int v) {
2     node* N = malloc(sizeof(node));
3     store(N->val, v, rlx);
4     store(L, N, rel);
5 }
6
4 int get {
5     node* r = load(L, acq);
6     if (r == NULL)
7         return UNDEF;
8     else {
9         int v = load(r->val, rlx);
10        return v;
11    }
12 }
13

```

Fig. 2. A stylized C11 implementation of a Register library. The contents of variables N and L can be accessed by multiple threads; variable r , on the other hand, is thread-local.

This paper makes the following contributions:

- (1) We provide a verification methodology for proving robustness of a concurrent data structure library executing under the RC20 memory model. To enable automated verification, our proof methodology exploits various constraints derived from the library’s SC specification. We provide a common SMT-based framework to systematically relate these high-level constraints with the low-level event-based guarantees required for robustness proofs.
- (2) We introduce a new notion of robustness, called *induced subgraph robustness*, tailored for reasoning about robustness of library implementations that may have benign non-robustness as part of their speculative computations. This weaker definition permits executions to have library-internal non-robust operations, as long as their effects do not manifest in a method’s return value.
- (3) We extend our methodology for proving robustness of an individual library to compose multiple robust libraries using SC fences to guarantee overall execution robustness. Our composition guarantees also extend to induced subgraph robustness. Our initial experiments indicate that the performance penalty of putting additional SC fences is not significant, and potentially masked by the synchronization already performed by the individual libraries to maintain robustness.
- (4) We have applied our approach on a number of realistic concurrent C11 library implementations, including real-world benchmarks such as Meta’s Folly lock-free queue implementation that make sophisticated use of relaxed atomics [26]. Our results establish, often for the first time, robustness against RC20 of well-known library implementations.

The remainder of the paper is structured as follows. The next section presents a motivating example to illustrate the core ideas underlying our approach for proving robustness of a library. §3 introduces a representative language, along with an axiomatic formalization of the memory model. The derivation of the inductive invariant used for our robustness proof is given in §4. We formalize the definition of induced subgraph robustness in §5. We provide our methodology to compose robust libraries in §6. Details about the implementation and SMT encoding are given in §7. Experimental results are presented in §8. Related work and conclusions are given in §9.

2 Motivating Example

Consider an implementation of a concurrent register data structure shown in Figure 2. Although not very efficient, it nonetheless exhibits access patterns commonly found in a number of real-world library implementations, and it allows us to provide a simple, concise demonstration of the core ideas of our verification strategy. The register data structure has two methods ($\text{set}(v)$ and get), with the intended semantics that get returns the most recent value that was assigned to the register by a set operation, or else UNDEF if no set operations have yet executed. In Figure 2, the set method is implemented by allocating a new node containing the value (in the field val), storing

a reference to the node in a shared variable L . The `get` method reads the variable L and then dereferences the `val` field (if L is not `NULL`) to return the value. Notice that each memory access is annotated with an access mode, which has an effect on its behavior, as well as the behavior of subsequent accesses.

If we ignore the annotations defining relaxed access modes and assume all reads and writes to shared-memory are SC, then every concurrent execution can be thought of as an interleaving of statements from different threads, all accessing a single shared memory. Under SC, it is straightforward to prove that the above implementation is correct, which means that all the invocations in an execution follow the intended register semantics.

2.1 RC20 Memory Model and Robustness

In order to describe executions under the RC20 relaxed memory model, we consider the read/write events generated during an execution, as well as dependency relations between these events. In the following, we give a brief informal description of these relations; they are formally defined in the next section. The relations include (i) *reads-from* (rf) that relates a write event to the read event which reads from it, (ii) *modification-order* (mo), a total ordering on all write events to the same location, and (iii) *from-read* (fr), a relation that relates a read event to write events that occur later than the write event it reads from according to the mo order. Events in the same thread are all totally ordered according to a session order (so) relation, a generalization of program order adapted to a concurrent library setting (a session is defined as a sequence of library method invocations performed by the same thread).

The access modes associated with read/write events are used to determine the *synchronizes-with* relation (sw): write and read events that are related by rf are also related by sw if the write is annotated with `rel` access mode and the read is annotated with `acq` mode (there are also other ways to establish sw , more details in §3.2). Finally, these primitive relations are used to define two derived relations: (i) hb_{SC} , the transitive closure of rf , fr , mo , so , and (ii) hb , the transitive closure of sw and so . Intuitively, hb_{SC} will be used to define the behavior of SC executions, while hb will define the behavior of RC20 executions.

For the purpose of this section, we are mainly interested in how the RC20 memory model uses hb to constrain the behavior of read events by forcing them to not ignore write events that are in hb order to them. This means that a read event e_r cannot read from a write event e_w if there is another write event e'_w such that $e_w \xrightarrow{mo} e'_w$ and $e'_w \xrightarrow{hb} e_r$ (more precisely, e'_w must be in hb -order before some previous event in the session containing e_r). In this case, e_r must read from either e'_w or an event mo -after e'_w .

To illustrate this, consider an execution of the register library consisting of two invocations, `set(v)` and `get`. The events generated during this execution are shown in Figure 3, with the two events on the left side of the figure generated by `set(v)` and the two events on the right generated by `get`. Assuming that the `malloc` statement in `set(v)` allocates a new node N , the rf relation from the store of L in `set` to the load of L in `get` induces a sw relation (because of the `rel` and `acq` annotations associated with the store operation in the `set` method (line 4) and the load operation in the `get` method (line 5)). Due to this, the write of $N \rightarrow val$ in `set` comes in hb -order before the read of $N \rightarrow val$ in `get`. As a result, this load must read the value v . However, if the access mode of either the load or store to L are changed to `rlx`, then the store to $N \rightarrow val$ would no longer be in hb order, allowing the load to $N \rightarrow val$ to ignore it. The execution




Fig. 3. An execution of the register library




Fig. 4. Establishing the inductive invariant for the location L in the register library implementation

in Figure 3 corresponds to classical message-passing behavior, and is possible under both SC and RC20.

The robustness problem then asks whether any arbitrary RC20 execution of the library implementation is also possible under SC. In order to solve the robustness problem, we consider an alternative event-based characterization of SC executions: SC constrains the behavior of read events by forcing them to not ignore write events which are in hb_{SC} order before them. This characterization is almost identical to the RC20 memory model, replacing the hb relation with hb_{SC} , and this observation leads to the following sufficient condition for robustness[43]: *At every step during an SC execution, if the next event is a read event e_r to location ℓ , and if the most recent write event e_w to ℓ (according to the mo ordering) is in hb_{SC} order to the session τ containing e_r , then e_w is also in hb ordering to τ .* This means that, under RC20, the read event e_r cannot ignore the write event e_w , and hence it must have the same behavior as under SC. Notice that although we are trying to establish robustness of RC20 executions, the robustness condition itself needs to be checked over SC executions only. This is because we are essentially searching for the minimal or first robustness violation, and hence the execution leading up to the violation must be itself robust and hence possible under SC.

Coming back to the execution in Figure 3, just before the event $R\langle N \rightarrow val, v, rlx \rangle$, the most recent write event to the location $N \rightarrow val$, $W\langle N \rightarrow val, v, rlx \rangle$, is in both hb_{SC} and hb order, thus ensuring that the read event has the same behavior under both SC and RC20. On the other hand, if the access modes of either the load or store to L were to be changed to rlx , robustness would be violated by the read event to $N \rightarrow val$, resulting in a manifestation of a non-robustness witness.

2.2 Verifying Robustness

Our verification strategy aims to prove the absence of non-robustness witnesses, which are nothing but violations of the sufficient condition for robustness: an SC execution of the library implementation that contains a write event on a location ℓ that is in hb_{SC} but not hb order to a session where the next event is a read event on ℓ . Since a library implementation can be invoked an arbitrary number of times by clients across arbitrary number of threads, we have to reason about robustness of an unbounded number of executions. To make this problem tractable, we define an inductive proof strategy that essentially uses the absence of non-robustness witnesses as an inductive invariant. We prove the correctness of the inductive invariant systematically on a per-location basis, and use information about the SC behavior of the library implementation (such as its high-level specification under SC) to facilitate automated verification.

To illustrate, we describe how to verify the robustness of the register implementation of Figure 2. Executions of this implementation access two classes of locations: (i) L and (ii) $N \rightarrow val$ for different values of N. We reason about the absence of non-robustness witnesses involving these location

classes separately. Intuitively, the `rel/acq` annotations to the store/load of location L should ensure robust behavior of loads to location $N \rightarrow \text{val}$. However, it is not clear how the implementation ensures robustness for accesses to location L . We prove that it does so by establishing the following inductive invariant over all SC executions: *if the most recent write to the location L (according to mo) comes hb_{SC} before another event, then it also comes hb before it*. This would then ensure that a subsequent load to L would have the same behavior under SC and RC20. We use induction on the number of events in an SC execution, and consider adding events in an order that obeys the hb_{SC} relation. Figure 4 depicts some of the more interesting inductive cases; for each case, we consider an arbitrary SC execution G of the library (depicted by the rectangle) and a new event μ to be added to the execution. The most recent write to location L in the execution G is labelled μ_w . Inductively, we assume that if μ_w comes hb_{SC} before another event in G , it also comes hb before it.

The various cases exhaustively consider new additions to hb_{SC} due to the inclusion of μ . We highlight the different ways in which the inductive invariant is maintained. Figure 4(a) considers the case where μ is a read event to L , which can cause a new hb_{SC} edge to be established through the `rf` relation. Now, since all writes to L use the access mode `rel` while all reads to L use access mode `acq`, the `rf` relation induces a `sw` relation, and thus a `hb` relation so that the required condition continues to hold. Figure 4(b) depicts an execution in which μ is a write to L that results in an incoming `fr` edge from an existing read (μ_r) in G . Note that the dotted arrow labeled with hb_{SC} indicates that the hb_{SC} relation is established through some sequence of events. Now, since μ_w occurs hb_{SC} -before μ_r , it also occurs hb_{SC} -before μ . However, in the new execution after addition of the event μ , notice that μ now becomes the most recent write to L , and it cannot occur hb_{SC} before any event in G (because we are adding events in hb_{SC} order). Hence, the required invariant holds trivially, and we do not have to consider any hb_{SC} paths from μ_w .

A more interesting case is shown in Figure 4(c). Here μ is a write event to $N \rightarrow \text{val}$ and its addition to the execution establishes a new hb_{SC} path from μ_w through a `fr` relation on μ_r . On the surface, it would seem that this depicts a non-`hb` (but hb_{SC}) path from μ_w to μ . However, in an actual execution, the `fr` relation involving $N \rightarrow \text{val}$ can never arise. In the following, we show how we can infer the infeasibility of this execution by using the fact that we are only considering SC executions.

Notice that in a SC execution, there is a unique write to any location of the class $N \rightarrow \text{val}$, because it only happens inside the `set` method, and `malloc` will always return a fresh location. Since μ_r does not read from this write event, and since there is no other write event to $N \rightarrow \text{val}$, μ_r must read the initial value of the location (say 0), which will also be returned by `get` invocation containing μ_r . However, the register implementation is agnostic to the actual value passed as an argument to `set` (i.e., `set`'s control-flow is unaffected by v) and thus we can assume data independence of its arguments [1]. In particular, we prohibit executions where `set` is passed the initialization value 0. But the `get` invocation containing μ_r returns 0, and the register specification says that for every `get` invocation which return a non-UNDEF value, there must be a `set` invocation whose argument value is the same as the return value of the `get` invocation. We thus derive a contradiction, establishing infeasibility of the execution in Fig. 4(c). Notice how we are able to leverage program structure constraints (uniqueness of write to $N \rightarrow \text{val}$ enforced by `malloc`) and the register specification for this reasoning.

The cases depicted in Figures 4(d) and 4(e) can also be handled in a fashion similar to the handling of cases 4(a) - 4(c). In particular, the execution shown in Figure 4(d) is not possible due to uniqueness of writes to $N \rightarrow \text{val}$. Figure 4(e) depicts a feasible execution. In this case, there will already exist a `hb` ordering between the event $W(N \rightarrow \text{val}, _, _)$ and μ (as shown in Figure 3) via a path comprising `sw` and `so` edges. Using the inductive hypothesis that μ_w also occurs `hb`-before $W(N \rightarrow \text{val}, _, _)$ and the transitivity of `hb`, we can infer that μ_w occurs `hb`-before μ . Collectively, our case analysis shows robustness for the location L , guaranteeing that there can be no non-robustness witnesses caused

$c \in \text{Constant}$	$v \in \text{Value} ::= c \mid \ell$
$\ell \in \text{Location}$	$e \in \text{Expr} ::= v \mid x \mid e_1 == e_2 \mid \dots$
$x, y \in \text{Var}$	$I \in \text{Method} ::= \text{method } m(x) = s$
$\tau \in \text{SessionId}$	$s \in \text{Stmt} ::= \text{skip} \mid x = e \mid x = \text{malloc}(c)$
$m \in \text{MethodName}$	$\mid \text{store}(x, e, o_W) \mid y = \text{load}(x, o_R) \mid y = \text{cas}(x, e_1, e_2, o_U, o_R) \mid \text{fence}(o_F)$
$\mu \in \text{MemEvt}$	$\mid y = \text{fadd}(x, e, o_U) \mid s_1; s_2 \mid \text{if } e \text{ then } s_1 \text{ else } s_2 \mid \text{while } e \text{ do } s \mid \text{return } e$
$\alpha_X \in \text{Mod}_X \text{ where } X \in \{R, W, U, F\}$	$\alpha \in \text{Action} ::= \epsilon \mid R\langle \ell, v, o_R \rangle \mid W\langle \ell, v, o_W \rangle \mid U\langle \ell, v_1, v_2, o_U \rangle \mid F\langle o_F \rangle$
	$\gamma \in \text{InvkEvt} ::= \text{Invk}\langle i, m, v_a, v_r, \tau \rangle$
	$\omega \in \text{Evt} ::= \gamma \mid \text{Sil}\langle \gamma \rangle \mid \text{Mem}\langle i, \tau, \alpha \rangle$

Fig. 5. Domains and language syntax.

by accesses to this location. In each case, either we cannot establish the hb_{SC} relation, or if we can, then hb relation is also established. A similar analysis can be carried out for the location $N \rightarrow \text{val}$.

The following sections formalize these intuitions, generalizes the approach to deal with implementations that use loops and synchronization operations like `cas` and `fence`, and builds an automated procedure to perform the above reasoning.

3 Preliminaries

3.1 Syntax and Program Semantics

To describe our approach, we consider a C-style imperative language, with standard control-flow operators (sequencing `(;)`, `while`, `if-then-else`, etc.), and operations to access thread-local and shared state (Figure 5). A library implementation of a data structure is given by a set of methods written in this language, with each method corresponding to an operation of the data structure; we prohibit methods from invoking other methods.

Library implementations have access to a shared heap i.e., an unbounded collection of locations (Location), shared among all sessions; additionally, library methods can bind and read local variables whose values are accessible to the method only for the lifetime of the invocation. We assume that every instance of the `malloc(c)` command returns a unique location on the heap, allocating c units of memory. We also require that a library implementation owns its accessed locations, which means they cannot be accessed by clients (or other library implementations) [49]. A method body is a sequence of statements, at least one of which must be a return.

The most interesting part of the language are statements that access shared memory. Specifically, statements `load(x, oR)` and `store(x, e, oW)` are used to load values from, and store values to, locations in shared memory whose address is contained in variable x , resp. Following the RC20 memory model, each shared memory access in our language is also associated with an *access mode*, taken from the set $\text{Mod} \triangleq \{\text{rlx}, \text{acq}, \text{rel}, \text{acqrel}\}$. These modes define the consistency level of the memory access. Note that the language does not allow non-atomics, as every access is identified with a consistency level. We define specific subsets of access modes for the different types of memory accesses:

$$\begin{aligned} \text{Mod}_R &= \{\text{rlx}, \text{acq}\} & \text{Mod}_U &= \{\text{rlx}, \text{acq}, \text{rel}, \text{acqrel}\} \\ \text{Mod}_W &= \{\text{rlx}, \text{rel}\} & \text{Mod}_F &= \{\text{acq}, \text{rel}, \text{acqrel}\} \end{aligned}$$

A partial order \sqsubseteq is also defined over these access modes, which intuitively orders them according to their consistency level: $\text{rlx} \sqsubseteq \text{acq} \sqsubseteq \text{acqrel}$ and $\text{rlx} \sqsubseteq \text{rel} \sqsubseteq \text{acqrel}$. Additionally, the `cas(x, e1, e2, oU, oR)` operation atomically updates the location ℓ bound to x to the evaluation of e_2 if e_1 evaluates to the value stored in ℓ , returning true in this case, and returns false otherwise.

$$\begin{aligned}
& \mathcal{I} \in \text{Invocations} = \text{SessionId} \rightarrow (\text{InvkEvt} \cup \{\perp\}) \times (\text{Stmt} \cup \{\perp\}) \\
& \varrho \in \text{Environments} = \text{SessionId} \rightarrow \text{EnvLocal} \\
& \mathcal{L} \in \text{Library} \in \mathcal{P}(\text{Method}) \\
& \rho \in \text{EnvLocal} = \text{Var} \rightarrow \text{Value}
\end{aligned}$$

$$\begin{array}{c}
\text{INVK} \\
\frac{
\begin{array}{l}
\mathcal{I}(\tau) = (_, \perp) \quad m = \lambda x. s \in \mathcal{L} \\
\varrho' = \varrho[\tau \mapsto \varrho(\tau)[x \mapsto v_a]] \\
i \text{ unique} \quad \gamma = \text{Invk}\langle i, m, v_a, v_r, \tau \rangle \\
\mathcal{I}' = \mathcal{I}[\tau \mapsto \langle \gamma, s \rangle]
\end{array}
}{
\langle \varrho, \mathcal{I} \rangle \xRightarrow{\gamma} \langle \varrho', \mathcal{I}' \rangle
}
\end{array}
\quad
\begin{array}{c}
\text{RET} \\
\frac{
\begin{array}{l}
\mathcal{I}(\tau) = \langle \gamma, \text{return}(e); s \rangle \\
\gamma = \text{Invk}\langle _, _, v_r, \tau \rangle \quad e \rightsquigarrow_{\varrho(\tau)} v_r \\
\mathcal{I}' = \mathcal{I}[\tau \mapsto \langle \gamma, \perp \rangle] \quad \varrho' = \varrho[\tau \mapsto \rho_\perp]
\end{array}
}{
\langle \varrho, \mathcal{I} \rangle \xRightarrow{e} \langle \varrho', \mathcal{I}' \rangle
}
\end{array}$$

$$\begin{array}{c}
\text{STEPLocal} \\
\frac{
\begin{array}{l}
\mathcal{I}(\tau) = \langle \gamma, s \rangle \quad \langle \varrho(\tau), s \rangle \xrightarrow{\epsilon}_\tau \langle \rho', s' \rangle \\
\mathcal{I}' = \mathcal{I}[\tau \mapsto \langle \gamma, s' \rangle] \quad \varrho' = \varrho[\varrho(\tau) \mapsto \rho']
\end{array}
}{
\langle \varrho, \mathcal{I} \rangle \xRightarrow{\text{Sil}(\gamma)} \langle \varrho', \mathcal{I}' \rangle
}
\end{array}
\quad
\begin{array}{c}
\text{STEPMEM} \\
\frac{
\begin{array}{l}
\mathcal{I}(\tau) = \langle \gamma, s \rangle \quad \langle \varrho(\tau), s \rangle \xrightarrow{\alpha}_\tau \langle \rho', s' \rangle \\
\mathcal{I}' = \mathcal{I}[\tau \mapsto \langle \gamma, s' \rangle] \\
\varrho' = \varrho[\varrho(\tau) \mapsto \rho'] \quad i \text{ unique} \quad \mu = \text{Mem}\langle i, \tau, \alpha \rangle
\end{array}
}{
\langle \varrho, \mathcal{I} \rangle \xRightarrow{\mu} \langle \varrho', \mathcal{I}' \rangle
}
\end{array}$$

Fig. 6. Global Reduction Rules

The semantics of library implementations are defined in terms of a most general client which can call arbitrary methods in any arbitrary session. Formally, for a library $\mathcal{L} \in \mathcal{P}(\text{Method})$, we define a labeled transition system (LTS), $\Omega_{\mathcal{L}} = (\Phi, \text{Evt}, \Rightarrow)$, where Φ denotes a set of states, Evt denotes a set of events (also used as labels on transitions) and $\Rightarrow \subseteq \Phi \times \text{Evt} \times \Phi$ defines a transition relation over states and events. We write $\xRightarrow{\omega}$ for the relation $\{\langle s_1, s_2 \rangle \mid \langle s_1, \omega, s_2 \rangle \in \Rightarrow\}$.

A *state* in Φ is a tuple, comprising a set of thread-local environments (ϱ) indexed by session ids and a set of current invocations (\mathcal{I} , also indexed by session ids) that records the state of each session along with the invocation event of the method currently executing in that session. In an invocation event, $\text{Invk}\langle i, m, v_a, v_r, \tau \rangle$, i represents an invocation number (a unique identifier for a particular invocation), m is a method name, v_a and v_r represent the argument supplied to and value returned by the method, and τ is a session identifier. The prophesied return value v_r must match the value yielded by e in any $\text{return } e$ statement executed during the execution of m with invocation number i . This follows the convention introduced by [49] to describe method invocations in a relaxed memory setting using a single event, instead of the traditional invocation and response events used in the SC setting.

We note that SessionId is an unbounded set of sessions. The *initial* state s_\perp of Ω_L is $\langle \varrho_\perp, \mathcal{I}_\perp \rangle$. \mathcal{I}_\perp maps each session in SessionId to $\langle \perp, \perp \rangle$, while ϱ_\perp maps each variable to an initial value.

Figure 6 defines the semantics of clients. It models the behavior of a most-general client. A client can invoke a new method using the **INVK** rule, which also establishes an initial local environment. Notice that $\mathcal{I}(\tau) = (_, \perp)$ ensures that τ is not an active session. The **RET** rule corresponds to the return of a method. We use the notation $e \rightsquigarrow_{\varrho(\tau)} v_r$ to indicate that the expression e evaluate to v_r under the local environment $\varrho(\tau)$. Notice that the **RET** rule sets the second component of $\mathcal{I}(\tau)$ to \perp , thus allowing future calls in τ through **INVK**. The **STEP** rules allow a method invocation active in a session to take either a silent step (**STEPLocal**) or a step that affects memory (**STEPMEM**)¹.

Steps taken by an invocation that do not access shared memory are recorded using a silent event ($\text{Sil}(\gamma)$). Steps that involve either reading or writing from shared memory are recorded using memory events that are of the form: $\text{Mem}\langle i, \tau, \alpha \rangle$, where i is a unique event identifier, τ is the session containing the event, and α defines the particular operation performed on the memory

¹The local reduction rules defining $\xrightarrow{\alpha}_\tau$ are straightforward and provided in the supplemental material, §2.

(read (R), write (W), read-modify-update (U), or fence (F)). We assume the language modeled by the LTS is receptive [51] and does not constrain the specific values read from memory. This assumption is consistent with prior work [43] that also separates out the memory system from the program semantics, by allowing load operations to read arbitrary values from memory. These values are then constrained by the memory system as defined in §3.2.

3.2 Memory System

We use declarative specifications of the RC20 memory model to construct the memory system. These specifications constrain *execution graphs*, which consist of the memory events that we defined in the previous subsection, along with a number of binary relations among these events. Formally, an execution graph $G = \langle M, \text{rf}, \text{mo}, \text{so} \rangle$ consists of a set of memory events M along with the binary relations *reads-from* (rf), *modification-order* (mo) and *session-order* (so). Our development closely follows prior work in this area [33, 38, 43].

Notation. For a relation R , we write $R^?$ for the reflexive and R^+ for the transitive closure of R , R^{-1} denotes its inverse, while $\text{dom}(R)$ denotes its domain. $R_1; R_2$ denotes the composition of the two relations R_1 and R_2 . Given a set of memory events M , $[M]$ denotes the identity relation on M . Hence, $[M_1]; R; [M_2]$ denotes $R \cap (M_1 \times M_2)$. Given an execution graph $G = \langle M, \text{rf}, \text{mo}, \text{so} \rangle$, we use the notation $G.X$ to denote the subcomponent X of G . Given a subset of events $M' \subseteq G.M$ and binary relation $G.R$ over $G.M$, we write $G.R|_{M'}$ to denote the projection of R over the events in M' . We also use the notation $G.M_\tau$ to denote the set of events in $G.M$ of session τ , and use $G.M_\top$ to denote set of events in $G.M$ of memory event type \top ($\top \in \{R, W, U, F\}$). We use $G.M_{W+U}$ to indicate the set of all write and update events. $G.M_{\sqsubseteq \chi}$ is used to denote the set of all events in $G.M$ whose access mode is related by the partial order \sqsubseteq , with χ . $G.M_\ell$ for location ℓ denotes the set of memory events accessing ℓ . We can also combine these notations to denote subsets of events satisfying multiple properties, for example, $G.M_{W, \ell}$ denotes write events to location ℓ . We use $G.w_\ell^{\max}$ to denote the most recent write (or update) event in $G.M$ to ℓ according to the $G.\text{mo}$ order. Given a memory event $\mu = \text{Mem}(i, \tau, \alpha)$, we use various projection functions $\text{sess}(\mu)$, $\text{type}(\mu)$, $\text{loc}(\mu)$, $\text{rval}(\mu)$, $\text{wval}(\mu)$, $\text{mod}(\mu)$ to denote the session, memory event type, memory location, read value, write value, access mode respectively. Similarly, for an invocation event γ , we use functions $\text{method}(\gamma)$, $\text{arg}(\gamma)$, $\text{ret}(\gamma)$ to denote the method name, argument, and return value resp. Given a relation R between memory events, we use the notation $\mu_1 \xrightarrow{R} \mu_2$ to denote $R(\mu_1, \mu_2)$ for memory events μ_1, μ_2 .

Execution Graphs. A valid execution graph $G = \langle M, \text{rf}, \text{mo}, \text{so} \rangle$ obeys certain well-formedness constraints, irrespective of the memory model. These well-formedness constraints follow directly from the definition of the various relations and include the following: (i) $\text{mo} \subseteq M_{W+U} \times M_{W+U}$ is a total order between write/update events to the same location, (ii) so is a total order between events in the same session, (iii) $\text{rf} \subseteq M_{W+U} \times M_{U+R}$ is total on its co-domain and only relates events to the same location and (iv) if $\mu_1 \xrightarrow{\text{rf}} \mu_2$, then $\text{wval}(\mu_1) = \text{rval}(\mu_2)$. We denote these constraints by χ_{base} . In order to understand how execution graphs are constrained by memory models, we first describe the specification for SC. Given an execution graph $G = \langle M, \text{rf}, \text{mo}, \text{so} \rangle$, we define the following derived relations:

$$\begin{aligned} G.\text{fr} &= (G.\text{rf}^{-1}; G.\text{mo}) \setminus [G.M] & G.\text{sw} &= [G.M_{\text{re1}}]; ([G.M_F]; G.\text{so})^?; G.\text{rf}^+; (G.\text{so}; [G.M_F])^?; [G.M_{\text{acq}}] \\ G.\text{hb}_{\text{SC}} &= (G.\text{rf} \cup G.\text{mo} \cup G.\text{fr} \cup G.\text{so})^+ & G.\text{hb} &= (G.\text{sw} \cup G.\text{so})^+ \end{aligned}$$

The fr relation relates a read event with write events that it does not witness. Execution graph G is said to be *SC-consistent* if hb_{SC} is irreflexive. The hb_{SC} ordering is an interleaving of events across sessions; making it irreflexive ensures that each read event reads from the most recent write event, formalizing our intuitive understanding of SC executions. We denote this constraint by $\chi_{\text{sc}}(G)$.

The *sw* relation relates *rel* writes with *acq* reads that read from them, as well as *rel* and *acq* fences that have *rf*-related write and read events between them; *hb* then transitively expands *sw* while also taking into account *so*. An execution graph G is said to be *RC20-consistent* if the following conditions hold: (1) $G.\text{fr}; (G.\text{rf})^?; G.\text{hb}$ is irreflexive; (2) $G.\text{mo}; (G.\text{rf})^?; (G.\text{hb})^?$ is irreflexive; (3) $G.\text{fr}; G.\text{mo}$ is irreflexive; and, (4) $G.\text{rf} \cup G.\text{so}$ is acyclic.

We denote this set of constraints by $\chi_{\text{RC20}}(G)$. Intuitively, the first two irreflexivity constraints guarantee that read and write operations obey *hb* ordering, i.e. a read event cannot overlook a write event that *happens-before* the read, and the modification order between two write events must agree with the *happens-before* order. The third irreflexivity constraint guarantees CAS semantics, i.e., two distinct update events cannot read from the same write event. The last constraint is required to prohibit out-of-thin-air reads [9]. Notice that $\chi_{\text{SC}}(G) \Rightarrow \chi_{\text{RC20}}(G)$, and hence, if an execution graph is SC-consistent, it is also RC20-consistent. The above specification of RC20-consistency matches with the development used in previous work [43].

Traces. We now present the memory systems for the SC and RC20 memory models. We define a common parameterized system that can be instantiated with either χ_{SC} or χ_{RC20} . The labeled transition system for memory model X is given by $\mathcal{MS}_X = \langle \mathcal{G}, \text{MemEvt}, \rightarrow_X \rangle$. Here, \mathcal{G} is the set of all execution graphs, MemEvt are memory events as defined earlier, and $\rightarrow_X \subseteq \mathcal{G} \times \text{MemEvt} \times \mathcal{G}$ are labeled transitions. We define an initial execution graph $G_\perp = \langle M_\perp, \emptyset, \emptyset, \emptyset \rangle$, where M_\perp contains an initial write event $W\langle \ell, v_\ell^\perp, \text{rlx} \rangle$ to every location ℓ .

Each transition adds a new memory event, corresponding to an arbitrary action α in arbitrary session τ , as long as the new execution graph obeys the well-formedness constraints and the consistency constraints of the memory model. We now take the product of the program semantics (Ω_L) and memory system (\mathcal{MS}_X) to describe executions of a library implementation \mathcal{L} under memory model X . In the product, the transition of the program semantics $\xRightarrow{\mu}$ and the transition of the memory system \xrightarrow{X} must agree on the memory event. For method invocations, returns, or silent transitions of the program semantics, there will be no transitions in the memory system. A *trace* of the combined transition system begins from the initial state $\langle s_\perp, G_\perp \rangle$ and contains a finite number of transitions. A *complete trace* must end in a final state $\langle s_\perp, G_\perp \rangle \xRightarrow{\omega_1} \dots \xRightarrow{\omega_n} \langle s, G \rangle$, where in the state s , all invocations have completed their executions, i.e. there are no pending return statements in any session.

In the following, we consider an execution of a library implementation L under memory model X to be a tuple $E = \langle t, \Gamma, G \rangle$ where t is a trace of the combined transition system $\Omega_L \times \mathcal{MS}_X$, Γ is the set of all method invocation events and G is the execution graph in the final state of the trace. We also define an invocation session order so_{inv} , which relates invocation events belonging to the same session in the order in which they appear in the trace t . We use the notation \mathcal{E}_X^L to denote the set of all such executions. Similarly, a complete execution corresponds to a complete trace, and we use \mathcal{CE}_X^L to denote the set of all complete executions. We also use the notation $E.G$ to denote the execution graph of E , and $E.M, E.\text{rf}, E.\text{mo}, E.\text{so}$ to denote the various components of the execution graph G .

3.3 Robustness and Library Correctness under SC

The robustness problem, in general, asks whether every RC20 execution of a program is also possible under SC. In the context of library implementations, we re-define the notion of execution-graph robustness, originally proposed in [38], as follows:

Definition 3.1. An execution $E \in \mathcal{E}_{\text{RC20}}^L$ of a library implementation L is *execution-graph robust*, if $E \in \mathcal{E}_{\text{SC}}^L$. A library implementation L is execution-graph robust if all of its complete executions $E \in \mathcal{CE}_{\text{RC20}}^L$ are execution-graph robust.

Our verification strategy (described in §4) uses the SC specification of a library implementation to discharge the verification conditions. The correctness of a library implementation under SC is typically specified using the notion of linearizability, which tries to establish a simulation between concurrent executions of the implementation and sequential executions of a reference implementation of the same data structure. However, previous works ([23, 25]) have also proposed alternative declarative specifications which are equivalent to linearizability and which directly constrain the argument/return value of invocations using a collection of axioms. Such a declarative specification is more amenable to SMT encoding, and hence we have used such specifications in our work. The specifications also establish a *happens-before* ordering among invocations (hb_{inv}) which is required to be a total order.

For example, to provide a declarative specification of the register library of Fig. 2, we first define a binary predicate *match* over invocation events in an execution, which is used in the following three specification predicates. These take as input the invocation events Γ in an execution E of the library implementation:

$$\begin{aligned} \text{match}(\gamma_1, \gamma_2) &\triangleq \text{method}(\gamma_1) = \text{set} \wedge \text{method}(\gamma_2) = \text{get} \wedge \text{arg}(\gamma_1) = \text{ret}(\gamma_2) \\ \chi_{\text{GETSET}}(\Gamma) &\triangleq \forall \gamma \in \Gamma. \text{method}(\gamma) = \text{get} \wedge \text{ret}(\gamma) \neq \text{UNDEF} \Rightarrow \exists \gamma' \in \Gamma. \text{match}(\gamma', \gamma) \wedge \\ &\quad \gamma' \xrightarrow{\text{hb}_{\text{inv}}} \gamma \\ \chi_{\text{GETFROM}}(\Gamma) &\triangleq \forall \gamma_1, \gamma_2, \gamma_3 \in \Gamma. \neg(\text{match}(\gamma_1, \gamma_2) \wedge \text{method}(\gamma_3) = \text{set} \wedge \gamma_1 \xrightarrow{\text{hb}_{\text{inv}}} \gamma_3 \wedge \gamma_3 \xrightarrow{\text{hb}_{\text{inv}}} \gamma_2) \\ \chi_{\text{GETUNDEF}}(\Gamma) &\triangleq \forall \gamma_1, \gamma_2 \in \Gamma. \text{method}(\gamma_1) = \text{set} \wedge \text{method}(\gamma_2) = \text{get} \wedge \gamma_1 \xrightarrow{\text{hb}_{\text{inv}}} \gamma_2 \\ &\quad \Rightarrow \text{ret}(\gamma_2) \neq \text{UNDEF} \\ \chi_{\text{LIN}} &\triangleq (\text{so}_{\text{inv}} \Rightarrow \text{hb}_{\text{inv}}) \wedge \text{hb}_{\text{inv}} \text{ is a total order} \end{aligned}$$

χ_{GETSET} constrains all *get* invocations that return non-UNDEF values to match their return values with the argument value of some *set* invocation (and also defines hb_{inv} relation between them), χ_{GETFROM} disallows scenarios in which a *get* invocation returns the value of an older (not most-recent) *set* operation, while χ_{GETUNDEF} disallows a *get* operation to return UNDEF if there is *set* operation before it. χ_{LIN} ensures that hb_{inv} is a total order and obeys the session order among invocations. The set $\chi_{\text{REG}} = \{\chi_{\text{GETSET}}, \chi_{\text{GETFROM}}, \chi_{\text{GETUNDEF}}, \chi_{\text{LIN}}\}$ defines a specification of the register library equivalent to linearizability[25]. Under SC, it is easy to see that any execution of the implementation of Fig. 2 satisfies χ_{REG} .

In general, given a library implementation \mathcal{L} of data structure \mathcal{D} with specification $\chi_{\mathcal{D}}$, we assume library correctness under SC, which means that every complete execution $E = \langle t, \Gamma, G \rangle \in \mathcal{CE}_{\text{SC}}^L$ satisfies the specification, i.e. $\mathcal{A}_i(\Gamma)$ holds for all $\mathcal{A}_i \in \chi_{\mathcal{D}}$. Declarative specifications equivalent to linearizability have been defined for all common data structures such as stack, queue, set, etc. in previous works [25]. These specifications can be easily encoded as FOL formulae over the domain of invocation events, constraining the method names, arguments, return value and the session order relation. Note that specifications do not directly constrain any internal event of the library implementation, since these would not be directly observable to a client, and the specification should be agnostic of the implementation.

4 Induction for Robustness

In this section, we adapt an existing approach [43] to checking robustness of programs under the RC20 memory model to the library setting, and in particular derive an inductive strategy for establishing robustness. Our strategy is based on deriving sufficient conditions for robustness, such

that if these conditions are maintained at every step in every SC execution of the implementation, then it is guaranteed to be execution-graph robust. In order to describe this robustness condition, we first define the notion of prefix of executions.

Definition 4.1. Given a complete execution $E = \langle t, \Gamma, G \rangle \in \mathcal{CE}_{SC}^L$ of implementation L under SC, a prefix $E' = \langle t', \Gamma', G' \rangle \in \mathcal{E}_{SC}^L$ is an execution obtained from a prefix t' of the trace t . That is, there exists t'' such that $t = t' \circ t''$ where \circ composes two traces if the final state in t' is the same as the initial state of t'' .

Note that for every relation R defined in the execution graph G , the relation R' in the execution graph G' will be $R|_{G'.M}$. That is because for events $e_1, e_2 \in G'.M$, if $e_1 \xrightarrow{R} e_2$, then since both events also occur in the trace t' , we would also have $e_1 \xrightarrow{R'} e_2$. The prefix E' also obeys the following property:

$$\forall \mu' \in E'.M. \forall \mu \in E.M. \mu \xrightarrow{E.so} \mu' \Rightarrow \mu \in E'.M$$

That is, if an event μ' is in the prefix, then every event μ before it in the same session must also be in the prefix. We now define the next event of a prefix:

Definition 4.2. Given a prefix $E' = \langle t', \Gamma', G' \rangle$ of E , μ is called the next event of E' if $t' \xRightarrow{\omega_1} C_1 \xRightarrow{\omega_2} C_2 \Rightarrow \dots \xRightarrow{\omega_{n-1}} C_{n-1} \xRightarrow{\mu} C_n$ is also a prefix of t , where $\omega_1, \dots, \omega_{n-1}$ are either invocation events or local events and μ is a memory event.

The next event of prefix E' is denoted by $\text{next}(E', E)$. For prefix E' and $\mu = \text{next}(E', E)$, $E' + \mu$ is used to denote the execution corresponding to the prefix of E ending in event μ . The following definition characterizes a non-robustness witness[43], i.e. a SC execution which leads to a non-robust RC20 execution.

Definition 4.3. Given a complete SC execution $E \in \mathcal{CE}_{SC}^L$ of implementation L , let E' be a prefix of E and event $\mu = \text{next}(E', E)$ such that $\text{loc}(\mu) = \ell$ and $\text{sess}(\mu) = \tau$. Then E' is a non-robustness witness if the following conditions are true:²

- (1) $E'.w_\ell^{\max} \in \text{dom}(E'.\text{hb}_{SC}^?; [E'.M_\tau])$
- (2) There exists $\mu_w \in E'.M_{\ell, \ell} \cup E'.M_{\cup, \ell}$ such that $\mu_w \neq E'.w_\ell^{\max}$,
 $\mu_w \notin \text{dom}(E'.\text{mo}; E'.\text{rf}^?; E'.\text{hb}^?; [E'.M_\tau])$ and if $\text{act}(\mu) \neq R$ then $\mu_w \notin \text{dom}(E'.\text{rf}; [E'.M_\cup])$
 and

The non-robustness witness definition formalizes our intuitive understanding of a non-robust execution: for a prefix E' , the next event μ can be non-robust (i.e. can have a different behavior under RC20 and SC memory models) if the most recent write event w_ℓ^{\max} to ℓ occurs hb_{SC}-before the session containing μ , so under SC, if μ is a read event, then it must read from w_ℓ^{\max} . But w_ℓ^{\max} does not occur hb-before the session containing μ (because of the presence of the earlier write event μ_w), so that in a RC20 execution, μ can read from μ_w . Note that if μ is not a read event, then the definition further constrains the earlier write event μ_w to not be related by rf to an update event, since the CAS semantics even under the C11 memory model would not allow another write (in this case μ) to come between the write (μ_w) from which an update events reads from. In essence, referring back to the memory system defined in Sec. 3.2, $E'.G \xrightarrow{\mu}_{RC20} E'.G + \mu$ would be a valid transition in the RC20 memory model, but this would not be a valid transition in the SC system.

Further, the execution leading up to such a non-robustness witness must obey SC semantics, thus characterizing the first (or minimal) violation of robustness. If we can ensure such a non-robustness

²Adapted from Definition 4.5 in [43]

witness does not arise during any SC execution, the library implementation is guaranteed to be robust.

THEOREM 4.4. *Given a library implementation L , if every prefix E' of every complete execution $E \in \mathcal{CE}_{SC}^L$ is not a non-robustness witness, then L is execution-graph robust³.*

The above theorem is a direct application of Theorem 4.6 in [43] but applied to executions of a library implementation in the presence of the most general client. In order to illustrate this idea, we revisit the execution in Figure 3. Notice that just before the event $R(N \rightarrow \text{val}, v, rlx)$, the most recent write event to the location $N \rightarrow \text{val}$, $W(N \rightarrow \text{val}, v, rlx)$, is in both hb_{SC} and hb order, thus ensuring that the read event has the same behavior under both SC and RC20. On the other hand, suppose we modified the library implementation and changed the access modes of either the load or store of L to be rlx . Then, we would get a non-robustness witness E' consisting of the two events of the set invocation (on the left) and the first event of the get invocation ($R(L, N, rlx)$), with the next event being the read of $N \rightarrow \text{val}$. This witness can be obtained through an SC execution, in the sense that it does not violate χ_{SC} , and it obeys both the conditions of Definition 4.3. In particular, with the next event to E' being a read to location $N \rightarrow \text{val}$, the most recent write event $w_{N \rightarrow \text{val}}^{\max}$ is $W(N \rightarrow \text{val}, v, rlx)$ and it is in hb_{SC} order to the session on the right (thus obeying condition-1), but it is not in hb order to it. As a result, the initializing write event to $N \rightarrow \text{val}$ (not depicted in the figure) would take the role of the event μ_w in condition-2 of Definition 4.3. Hence, relaxing the access mode of either of the accesses to the location L renders the library implementation non-robust.

While Definition 4.3 simplifies our task of proving robustness of executions by reducing it to a search problem over SC executions, this is still a tall order, as it requires maintaining information about relations between low-level events across an unbounded number of arbitrarily long executions involving an unbounded number of heap locations. To address this issue, our verification strategy uses induction on the prefixes of SC executions, and then shows the absence of a non-robustness witness at every step of every SC execution. We can try to use the non-existence of a non-robustness witness itself as an inductive invariant. That is, for an SC execution E , we consider a prefix E' and assume inductively that for $\mu = \text{next}(E', E)$, E' and μ do not form a non-robustness witness, and then prove that for $\mu' = \text{next}(E' + \mu, E)$, $E' + \mu$ and μ' do not form a non-robustness witness. Notice from Def. 4.3 that E' and μ not forming a non-robustness witness would mean $\neg((1) \wedge (2)) \equiv (1) \Rightarrow \neg(2)$, which essentially means that if $E'.w_\ell^{\max}$ is in hb_{SC} order before session τ , then it should also be in hb order before τ .

Unfortunately, we find that this does not constitute an appropriate inductive invariant, since it only constrains the events that access location ℓ , but the next event of the new prefix $E' + \mu$ may access a different location after this event is added. Hence, the inductive hypothesis involving ℓ is not useful. We could try to use as inductive invariant the required property for all locations: i.e. $\forall \ell$, if $E'.w_\ell^{\max}$ is in hb_{SC} order before some session τ , it is in hb order before it. However, this condition is too strong and not necessary for establishing non-robustness. This is because even if this condition does not hold for some location (say ℓ'), which means that $E'.w_{\ell'}^{\max}$ is in hb_{SC} order before the session τ but not in hb order before it, the next event in τ must access ℓ' to form the actual non-robustness witness. The presence of such an event $E'.w_{\ell'}^{\max}$ is a necessary condition for forming an actual non-robustness witness, which we formalize below as a potential witness:

³All proofs can be found in supplemental material, §3

Definition 4.5. Given an execution $E \in \mathcal{CE}_{SC}^L$, a prefix E' of E , $\mu \in E'.M$ and a write event $\mu' \in E'.M_W^{\max} \cup E'.M_U^{\max}$, E' , μ' and μ form a potential non-robustness witness if $\mu' \xrightarrow{E'.hb_{SC}} \mu$ and $\neg(\mu' \xrightarrow{E'.hb} \mu)$.

$E.M_W^{\max}$ indicates the set of maximal events in E according to the mo ordering. Note that to obtain an actual non-robustness witness, we would have to instantiate the above definition with $\mu' = w_\ell^{\max}$ for some location ℓ , and the next event in the session containing μ must also access ℓ . Figure 7 shows a potential non-robustness witness for the register library implementation. It shows an execution with two set invocations which happen in different sessions. Because of the mo relation between the two write events to location L, the write event $W(N_1 \rightarrow val, v_1, rlx)$ comes in hb_{SC} order to the last event of session τ_2 ($W(L, N_2, rel)$), but it is not in hb order before it. Using the notation of the above definition, E' would be the execution depicted in Figure 7, μ' would be the write event to $N_1 \rightarrow val$, and μ would be the write event to L in session τ_2 .

A potential non-robustness witness, consisting of prefix E' and the write event μ' can become an actual non-robustness witness if E' can be expanded into another prefix E'' such that the next event μ'' of E'' accesses the location of μ' , μ' continues to remain the most recent write to its location, and the non-hb hb_{SC} relation continues to exist between μ_w and μ'' . For the potential non-robustness witness of Fig. 7, a subsequent immediate access to location $N_1 \rightarrow val$ would result in an actual non-robustness witness. While every potential witness does not necessarily become an actual witness, an actual witness must have been at some stage a potential witness. For example, the potential non-robustness witness in our example will never result in an actual witness, since there will be no further access to the location $N_1 \rightarrow val$ in an SC execution. Intuitively, this is because of the register semantics, as the effect of a set invocation γ_2 “overwrites” the effects of a previous set invocation γ_1 through the write to location L, and the $\chi_{GETFROM}$ specification prohibits a get invocation to read from an older set invocation.




Fig. 7. Potential non-robustness witness in the register library implementation.

While every potential witness does not necessarily become an actual witness, an actual witness must have been at some stage a potential witness. For example, the potential non-robustness witness in our example will never result in an actual witness, since there will be no further access to the location $N_1 \rightarrow val$ in an SC execution. Intuitively, this is because of the register semantics, as the effect of a set invocation γ_2 “overwrites” the effects of a previous set invocation γ_1 through the write to location L, and the $\chi_{GETFROM}$ specification prohibits a get invocation to read from an older set invocation.

Table 1. Sufficient conditions for establishing robustness inductively.

$\Phi_{src}(E, \ell)$	$\forall \mu_m \in E.M_W^{\max} \cup E.M_U^{\max}. \forall \mu_1, \mu_2 \in E.M. \text{loc}(\mu_m) = \ell \wedge$ $\wedge \mu_m \xrightarrow{E.hb^?} \mu_1 \wedge \mu_1 \xrightarrow{E.fr \vee E.mo \vee E.rf} \mu_2 \Rightarrow \exists \mu_3 \in E.M. \mu_m \xrightarrow{E.hb} \mu_3 \wedge \mu_3 \xrightarrow{E.hb^?} \mu_2$
$\Phi_{dst}(E, \ell)$	$\forall \mu_m \in E.M_W^{\max} \cup E.M_U^{\max}. \forall \mu_1, \mu_2 \in E.M. \text{loc}(\mu_m) = \ell \wedge \text{loc}(\mu_2) = \ell$ $\wedge \mu_m \xrightarrow{E.hb_{SC}} \mu_1 \wedge \mu_1 \xrightarrow{E.so} \mu_2 \Rightarrow \exists \mu_3 \in E.M. \mu_m \xrightarrow{E.hb} \mu_3 \wedge \mu_3 \xrightarrow{E.so} \mu_2$

In order to verify robustness, we then have to ensure that either a potential non-robustness witness is not formed at all (thus ensuring robustness at the source), or if a potential robustness witness does form, it does not manifest into an actual non-robustness witness (thus ensuring robustness at the destination). As shown in Table 1, we present two conditions (Φ_{src} and Φ_{dst}) that ensure the absence of a non-robustness witness by disallowing it at the source and destination, resp. These conditions are parametric on executions and locations.

We first focus on Φ_{src} , which inductively tries to maintain the non-existence of a potential non-robustness witness. In order to understand $\Phi_{src}(E, \ell)$, we note that the only difference between the definitions of hb and hb_{SC} is that the rf relation between relaxed accesses as well as fr and mo

relations between any two accesses induce an hb_{SC} relation, but they do not induce an hb relation. Hence, whenever hb_{SC} is established due to these relations, we should also be able to establish hb . Φ_{src} considers a maximal write event μ_w to location ℓ in E . If there exists a hb_{SC} (and hence inductively hb) relation from μ_w to some event μ_1 , and if this hb_{SC} relation were to be extended to μ_2 by a fr , mo or rf relation between μ_1 and μ_2 , then there should also exist a hb relation from μ_w to μ_2 , through some event μ_3 . To illustrate, we refer back to Figure 4, which essentially shows the different cases to be handled in order to establish $\Phi_{\text{src}}(E, L)$ for every SC execution E of the register implementation. In particular, for the case depicted in Figure 4(a), using the notation of Φ_{src} , $\mu_w = \mu_1 = W\langle L, _, \text{rel} \rangle$ and $\mu_2 = \mu_3 = R\langle L, _, \text{acq} \rangle$. Cases depicted in Figure 4(b)-(d) do not satisfy the antecedent of Φ_{src} (because there does not exist events μ_1, μ_2 with the required relation between them), while for the case depicted in Figure 4(e), $\mu_1 = W\langle N \rightarrow \text{val}, _, _ \rangle$, $\mu_2 = R\langle N \rightarrow \text{val}, _, _ \rangle$, and $\mu_3 = R\langle L, N, \text{acq} \rangle$; μ_3 is not depicted in the figure, but it would occur before μ_2 in the same get invocation. Note that the antecedent of Φ_{src} simply means that $\mu_w \xrightarrow{\text{hb}} \mu_2$, but explicitly maintaining the event μ_3 which establishes this hb relation helps us in the SMT encoding of Φ_{src} , as we will explain in the next section.

Now, if Φ_{src} cannot be established, then a potential non-robustness witness may be formed, as illustrated in Fig. 7. Notice that this means that $\Phi_{\text{src}}(E, N_1 \rightarrow \text{val})$ does not hold for the execution E of Fig. 7. Then, we use the condition Φ_{dst} to prevent such a potential non-robustness witness from turning into an actual witness. In words, this condition considers a scenario when there is a maximal write event μ_w^{max} to location ℓ which occurs hb_{SC} -before event μ_1 , with a later event μ_2 in the same session accessing the location ℓ . This could lead to an actual non-robustness witness, with the prefix being the execution corresponding to the trace leading upto the state just before μ_2 . However, in such a scenario, Φ_{dst} enforces the existence of another event μ_3 before μ_2 (μ_3 could be the same as μ_1) in the same session, such that μ_w^{max} occurs hb -before μ_3 . This would prevent the formation of an actual non-robustness witness involving μ_w^{max} and μ_2 .

To illustrate, we have already established through Fig. 7 that Φ_{src} does not hold for the location $N_1 \rightarrow \text{val}$. To show Φ_{dst} , we let μ_w be $W\langle N_1 \rightarrow \text{val}, v_1, \text{rlx} \rangle$, while μ_1 would be some event which occurs in so-order before μ_2 , which has to be another access to $N_1 \rightarrow \text{val}$. As per the program structure of the register library, μ_2 can only be a read event, which must occur in a get invocation. The scenario exactly corresponds to Fig. 3, where, using the notation of Φ_{dst} , $\mu_w = W\langle N \rightarrow \text{val}, v, \text{rlx} \rangle$, $\mu_1 = R\langle L, N, \text{acq} \rangle$ and $\mu_2 = R\langle N \rightarrow \text{val}, v, \text{rlx} \rangle$. However, in this case, μ_1 must read from $W\langle L, N, \text{acq} \rangle$, because there is no another write to L which writes N (guaranteed by the malloc semantics). Hence, we can instantiate μ_3 with μ_1 itself in the consequent of Φ_{dst} .

Notice that if $\Phi_{\text{src}}(E, \ell)$ holds, then $\Phi_{\text{dst}}(E, \ell)$ also holds, because hb_{SC} would imply hb from write events to ℓ , in which case we can take $\mu_3 = \mu_1$. However, in general, Φ_{dst} does not imply Φ_{src} (as we saw for the location $N_1 \rightarrow \text{val}$). The reason we separate out the two conditions is because for some locations, it is easier to establish $\Phi_{\text{src}}(E, \ell)$, since the condition just requires one access to the location ℓ , as opposed to $\Phi_{\text{dst}}(E, \ell)$ which requires two accesses. The distinction will become more clear when we go through our SMT encoding in §7.

To summarize, $\Phi_{\text{src}}(E, \ell)$ prevents the formation of a potential non-robustness witness involving location ℓ , while $\Phi_{\text{dst}}(E, \ell)$ prevents a potential non-robustness witness involving ℓ from turning into an actual non-robustness witness. If we can show that $\forall \ell \in \text{Location}. \forall E \in \mathcal{CE}_{\text{SC}}^L. \Phi_{\text{dst}}(E, \ell) \vee \Phi_{\text{src}}(E, \ell)$, then we can inductively show execution-graph robustness.

THEOREM 4.6. *Given a library implementation L , if $\forall \ell \in \text{Location}. \forall E \in \mathcal{CE}_{\text{SC}}^L. \Phi_{\text{dst}}(E, \ell) \vee \Phi_{\text{src}}(E, \ell)$, then L is execution-graph robust.*

We note that Theorem 4.6 does not hold in the other direction, i.e. Φ_{dst} and Φ_{src} are not necessary for ensuring execution-graph robustness. In particular, there is a gap between the definition of

a non-robustness witness (Def. 4.3) and the conditions Φ_{src} and Φ_{dst} . These conditions merely ensure that if the maximal write μ_m to a location appears hb_{SC} -before the next access to it, it also appear hb -before it. They do not consider the non-maximal write μ_w from Def. 4.3 at all, which is actually the root cause of the non-robust behavior. In particular, consider the scenario where a non-maximal write μ_w may be in rf relation to an update event, and hence cannot cause a robustness violation if the next event (say μ) to be added is a write/update event to the same location as μ_w , and the maximal write μ_m is in hb_{SC} order to the session containing μ . In this case, $act(\mu) \neq R$, $\mu_w \in dom(E'.rf; [E'.M]_U)$ (here E' is the set of events in the execution, see point(2) in Def. 4.3). Hence, we do not require μ_w and consequently μ_m to be hb -before μ .

Fig. 8 concretely demonstrates this scenario.

Assume that all the events in the execution come from some library invocations. The update operation to ℓ_1 in τ_2 reads from the write to ℓ_1 in τ_1 . The read in τ_3 to location ℓ_2 reads from the write in τ_2 , hence the maximal write to ℓ_1 (which is the update in τ_2) is in hb_{SC} order, but not in hb order to τ_3 , since the accesses to ℓ_2 are relaxed. Excluding the write to ℓ_1 in




Fig. 8. Execution illustrating incompleteness of Φ_{src} and Φ_{dst} .

τ_3 , the execution is a potential non-robustness witness. Let E' be this execution, with the next event $\mu = W(\ell_1, 3, rlx)$. Following the notation used in Def. 4.3, we have the maximal write to ℓ_1 , $\mu_m = U(\ell_1, 1, 2, rlx, rlx)$, and a non-maximal write $\mu_w = W(\ell_1, 1, rlx)$. Now, $E' + \mu$ is not an actual non-robustness witness, because even though μ_m is hb_{SC} order and μ_w is not in hb order to the session τ_3 , μ_w is in rf order to an update event, and hence μ_w does not satisfy point(2) of Def. 4.3.

We note that except scenarios like the one given above, Φ_{src} and Φ_{dst} precisely model the absence of non-robustness witnesses, i.e. if these conditions are violated, it would imply the presence of an actual non-robustness witness. A distinguishing property of executions like Fig. 8 is that they involve at least one update event (generated through CAS operation) to a location. We have observed that in library implementations, either all writes to a location happen through CAS operations, or none do (i.e. it doesn't happen that there is both a CAS and a normal write operation to the same location). If all accesses to a location are through CASes, our automated verification strategy (given in §7) uses additional CAS constraints, which helps us avoid any false positives (like the execution in Fig. 8) arising out of the imprecision of Φ_{src} and Φ_{dst} .

5 Induced Subgraph Robustness

We now discuss the challenges that arise in using our verification strategy on real-world library implementations, and the mechanisms we employ to overcome them.

We observe that the notion of execution-graph robustness is often too strong for real-world library implementations that use low-level synchronization primitives such as compare-and-swap (CAS). A typical CAS-based synchronization pattern is as follows: a method invocation performs a number of speculative reads without any synchronization, followed by some computation based on these speculative reads, and then finally stores the result of this computation to the global state if no other concurrent method invocation has made conflicting changes. The pattern uses CAS operations to determine if a conflict exists. If another concurrent invocation has indeed made changes to the global state, the speculative computation is ignored (since the corresponding CAS would fail), and the pattern is retried. These speculative computations are often performed using relaxed accesses, and they may generate an arbitrary number of read/write events. Any non-robust behavior of such events would result in a failing CAS, in effect restarting the entire invocation. It is clearly ineffective to establish the robustness of events that are generated by such failed

computation. A concrete example of a library implementation demonstrating this pattern is given in the supplemental material, §4.

In particular, an RC20 execution of a library implementation may be still be effectively robust against SC even if it exhibits potentially non-robust actions if all events that actually affect the return value of any method invocation obey SC semantics. To capture this distinction, we define a new notion of robustness called *induced subgraph robustness* that only focuses on the robustness of events that actually effect the client observable behavior of an invocation. Given an execution graph $G = \langle M, \text{rf}, \text{mo}, \text{so} \rangle$, $G' = \langle M', \text{rf}', \text{mo}', \text{so}' \rangle$ is called an induced subgraph of G if $M' \subseteq M$ and $\text{rf}' = \text{rf}|_{M'}$, $\text{mo}' = \text{mo}|_{M'}$, $\text{so}' = \text{so}|_{M'}$.

Definition 5.1. A complete execution $E = \langle t, \Gamma, G \rangle \in \mathcal{CE}_{\text{RC20}}^L$ of a library implementation L is *induced subgraph robust* if there exists a complete execution $E' = \langle t', \Gamma', G' \rangle \in \mathcal{CE}_{\text{SC}}^L$ such that (i) $\Gamma' = \Gamma$, (ii) For all invocation events $\gamma_1, \gamma_2 \in \Gamma$, γ_1 occurs before γ_2 in trace $t \Leftrightarrow \gamma_1$ occurs before γ_2 in trace t' and (iii) G' is an induced subgraph of G . L is *induced subgraph robust* if all of its complete executions $E \in \mathcal{CE}_{\text{RC20}}^L$ are induced subgraph robust.

Notice that induced subgraph robustness only considers complete executions of the library implementation. For every RC20 execution, there must exist another complete execution under SC such that the behavior of the method invocations (i.e. argument and return values) and the memory events which are responsible for this behavior remains the same as the original RC20 execution. The latter is captured by requiring the execution graph G' of the SC execution to be an induced subgraph of the RC20 execution. Referring back to the CAS based pattern used in libraries, the motivation behind requiring G' to be an induced subgraph of G is that G' would presumably only contain the events of G corresponding to the computation which resulted in a successful CAS, while the remaining events in G corresponding to failing speculative computation can be ignored. This is ensured by the fact that the behavior of the method invocation events remains the same. The client of the library need not care about the non-robustness of events which do not affect the observable behavior.

While induced subgraph robustness is a more useful correctness criterion for real-world library implementations, our verification strategy of the previous section aims to show execution-graph robustness. In order to connect these two notions, we consider a pre-processing step that transforms the implementation so that it does not generate those events which have no impact on an invocation's return value. Our goal is to ensure that all events remaining in the transformed program will have an effect on a method's return values. Induced subgraph robustness of the original implementation is then reduced to checking execution-graph robustness of the new implementation. Formally, we define a robustness-preserving transformation as follows:

Definition 5.2. A *robustness-preserving* program transformation is a function ρ that takes as input a library implementation L and outputs another library implementation $\rho(L)$ that obeys the following two conditions:

- (1) For every complete RC20 execution $E = \langle t, \Gamma, G \rangle \in \mathcal{CE}_{\text{RC20}}^L$ of L , there exists a complete RC20 execution $E' = \langle t', \Gamma', G' \rangle \in \mathcal{CE}_{\text{RC20}}^{\rho(L)}$ of $\rho(L)$ such that (i) $\Gamma' = \Gamma$, (ii) the order of invocation events is the same in both the traces t and t' and (iii) G' is an induced subgraph of G .
- (2) $\mathcal{CE}_{\text{SC}}^{\rho(L)} \subseteq \mathcal{CE}_{\text{SC}}^L$.

The first condition in the above definition allows executions of $\rho(L)$ to not contain all the events in executions of L , as long as the observable behavior of invocation events remains the same. The second condition ensures that the transformation does not add any new behaviors. If we can show

that the transformed implementation $\rho(L)$ is execution graph robust, then this would imply that the original implementation is induced subgraph robust. We can directly apply our induction strategy of Section 4 on $\rho(L)$ to determine its execution-graph robustness.

THEOREM 5.3. *Given a library implementation L and a robustness-preserving transformation ρ , if $\rho(L)$ is execution-graph robust, then L is induced subgraph robust.*

While checking whether a transformation is robustness preserving would be hard in general—since it requires comparing all RC20 executions of two implementations—applying the definition to the CAS based synchronization pattern used in libraries is quite straightforward. In particular, we consider the transformation ρ_{CAS} , which focuses on loops whose loop condition is based on the success of a cas instruction, and which only generates read events in iterations where the CAS fails. In such cases, only the last iteration of the loop where the cas succeeds actually matters, and all previous iterations generate benign read events whose robustness can be ignored. The ρ_{CAS} transformation is implemented as a simple syntactic analysis in the following manner: For every while loop whose loop constraint is a cas operation, we first check that only read operations are performed in any iteration. Following that, we unroll the loop once in the original implementation, and check for any dependencies (through local variables) from read operations in the first iteration to write operations, return value of the method, or any operation in the second iteration. If there are no such dependencies, we conclude that the loop iterations are independent, and only the events in the last iteration need to be preserved. In this case, we remove the while loop, and replace the cas operation with a bcas. In such cases, it eliminates the loop, and replaces the cas with a blocking cas (bcas) with the same parameters [38]. The bcas operation blocks until its compare is successful, ensuring that only the events in the last iteration of the loop in the original implementation will be generated.

LEMMA 5.4. *The transformation ρ_{CAS} is a robustness-preserving transformation.*

Intuitively, for any RC20 execution E of the original implementation L , we can construct a RC20 execution E' of $\rho_{CAS}(L)$ because all the write operations in E can be directly replicated in E' , since the only events that will occur in E but not E' will be read events who do not have any dependences on the write events. Even though these read events can potentially induce more hb relations in E , this only restricts the behavior of other events which are preserved from E to E' , and hence this behavior can be replicated in E' where $E'.hb \subseteq E.hb$.

Finally, we note that the notion of induced subgraph robustness is closely related to the previously proposed notion of observational robustness [43], which also allows benign non-robust events without any outgoing dependencies. The major difference is that we explicitly maintain the induced subgraph property, and the fact that invocation events retain the same behavior. Further, in heap manipulating programs, it often happens that the location to be read by a subsequent read operation depends on the value read from a previous read (as in the get implementation in the register library). In such a scenario, observational robustness would then enforce robustness of the earlier read (because there is an outgoing dependency), but in our case, induced subgraph robustness may allow both reads to be non-robust if there is no outgoing dependency to the return value/global state.

6 Compositionality

As discussed in the introduction, individual library robustness as defined in Def. 3.1 or Def. 5.1 is not sufficient for establishing robustness of executions involving multiple libraries. The issue is that robustness of a library itself does not provide enough synchronization guarantees that would be required to establish whole execution robustness. To compensate for this, we consider adding an SC-fence when crossing libraries within every thread. The store buffering example involving the

INVK	INVKF
$ \begin{array}{l} I(\tau) = \langle y', \perp \rangle \quad m = \lambda x.s \in \mathcal{L} \\ (y' = \text{Invk}(_, m', _, _, \tau) \wedge m' \in \mathcal{L}) \vee y' = \perp \\ \varrho' = \varrho[\tau \mapsto \varrho(\tau)[x \mapsto v_a]] \\ i \text{ unique} \quad \gamma = \text{Invk}(i, m, v_a, v_r, \tau) \\ I' = I[\tau \mapsto \langle y, s \rangle] \end{array} $	$ \begin{array}{l} I(\tau) = \langle y', \perp \rangle \quad y' = \text{Invk}(_, m'', _, _, \tau) \quad m'' \notin \mathcal{L} \\ m = \lambda x.s \in \mathcal{L} \quad m' = \lambda x.\text{fence}(\text{sc}); s \\ \varrho' = \varrho[\tau \mapsto \varrho(\tau)[x \mapsto v_a]] \\ i \text{ unique} \quad \gamma = \text{Invk}(i, m, v_a, v_r, \tau) \\ I' = I[\tau \mapsto \langle y, s \rangle] \end{array} $
$\langle \varrho, I \rangle \xRightarrow{y} \langle \varrho', I' \rangle$	$\langle \varrho, I \rangle \xRightarrow{y} \langle \varrho', I' \rangle$

Fig. 9. New Global Reduction Rules of $\Omega_{\mathbb{L}}$

two register libraries clearly demonstrates that we need SC-fences when crossing libraries within every thread, but the question is whether this is sufficient in general for any execution involving any robust library implementations? We answer this in the positive, and formally prove that for executions composed of calls to multiple robust libraries, if there is an SC-fence within each thread when crossing different libraries, then the overall execution is guaranteed to be robust.

In the following, we assume each pair of libraries \mathcal{L}_1 and \mathcal{L}_2 are disjoint, i.e. there is no common method that belongs to both libraries. We also require that the sets of memory locations that may be accessed by each library are also disjoint from each other, which is enforced by the ownership assumption mentioned at the beginning of §3. We denote an SC fence (i.e. the three instruction program `fence(acq); fadd(f , 0, acqrel); fence(rel)`) by `fence(sc)`. Note that the `hb` relation will be total among all the update events that are generated by the `fadd` instructions, due to the `acqrel` annotation, while the `acq` and `rel` fences are required to synchronize with any corresponding fences (if present) in the library implementations.

Given a set of libraries $\mathbb{L} = \{\mathcal{L}_1, \dots, \mathcal{L}_n\}$, we define a LTS $\Omega_{\mathbb{L}}$, which is almost exactly the same as the LTS $\Omega_{\mathcal{L}}$ that was defined in §3 for a single library \mathcal{L} . Similar to $\Omega_{\mathcal{L}}$, each state of $\Omega_{\mathbb{L}}$ also consists of a thread-local environment per active session, and the current invocations indexed by session id, except that now, these invocations can come from any of the libraries in \mathbb{L} . The transition rules also remain the same, except the INVK rule in Fig. 6, which is replaced by the two new rules given in Fig. 9.

The new INVK rule is used to invoke a method m of library $\mathcal{L} \in \mathbb{L}$ if the method preceding it in its session belonged to the same library; otherwise, in order to invoke a method following one from a different library, the INVKF rule must be used. In this case the method call is preceded with an SC-fence statement but the rest of the rule is the same as INVK.

Because of the altered INVK rules, executions of $\Omega_{\mathbb{L}} \times \mathcal{MS}_X$, for $X = \text{SC}$ or RC20 , will always have an SC-fence at the boundary when crossing from a method of one library to a method of a different library in session order. Notice that there is no SC-fence between invocations of methods from the same library. Intuitively, the executions generated by $\Omega_{\mathbb{L}} \times \mathcal{MS}_X$ maintain a *well-fencedness* property: in any (complete or partial) execution, if there is an event e_1 generated by a method belonging to \mathcal{L}_1 preceding an event e_2 generated by a method belonging to a different library \mathcal{L}_2 in session order, then there must be an SC-fence in between them. Let $\mathcal{CE}_X^{\mathbb{L}}$ denote the set all complete executions generated by $\Omega_{\mathbb{L}} \times \mathcal{MS}_X$ for the memory system X .

THEOREM 6.1. *Given a set of libraries \mathbb{L} , if each library in \mathbb{L} is execution-graph robust, then all executions in $\mathcal{CE}_{\text{RC20}}^{\mathbb{L}}$ are also execution graph robust.*

We prove the above theorem by contradiction, starting with a non-robust execution in $\mathcal{CE}_{\text{RC20}}^{\mathbb{L}}$ and inspecting the shape of a minimal hb_{SC} cycle that must be present in the execution graph. We construct an execution comprising solely of events within a single library $\mathcal{L} \in \mathbb{L}$ in which the hb_{SC} cycle is preserved, thus providing a contradiction to the guarantee of robustness of \mathcal{L} . In the case where the original cycle involves events from multiple libraries, we use the presence of SC

fences between libraries to infer hb synchronization, and we construct an RC20-execution of \mathcal{L} by converting these hb relations to so by changing the assignment of invocations to sessions.

We can also extend the above result for induced subgraph robustness, using the same definition as given in Def. 5.1 for complete executions in $\mathcal{CE}_{\text{RC20}}^{\mathcal{L}}$.

THEOREM 6.2. *Given a set of libraries \mathbb{L} , if each library in \mathbb{L} is induced subgraph robust, then all executions in $\mathcal{CE}_{\text{RC20}}^{\mathcal{L}}$ are also induced subgraph robust.*

Intuitively, there may be benign non-robustness in executions of individual libraries, which does not affect its observable behavior. In such a scenario, for multi-library executions, this non-robustness would still exist in the overall execution. Hence, we may not be able to show the overall execution is execution-graph robust, but instead, we can safely remove the non-robustness of each individual library to show induced subgraph robustness of the entire execution. This preserves the observable behavior of each invocation.

Our formulation thus far still does not allow a client to also perform atomic RC20 accesses. However, it is possible to view the client program as a library itself, with every program segment in every session between two library calls encapsulated as a separate method in this artificially generated client library. Then, Theorems 6.1 and 6.2 allow us to decompose the overall robustness problem into proving robustness for the actual libraries and the (synthetic) client library. In particular, as a special case, we can consider a client program which does not access any shared variable (i.e., leaves all the shared memory operations to the libraries), or a client program which only accesses shared memory inside locks. Such client programs will be robust according to our definition (i.e., the artificially generated client library will be robust), and hence composing them with robust libraries guarantee overall robustness.

We note that the usage of the SC fences is as important as the robustness of the individual libraries to guarantee overall robustness of the composite execution. If a library \mathcal{L} can generate a (non-benign) non-robust execution on its own, it is obvious that we can orchestrate a composite execution involving multiple libraries along with \mathcal{L} which would also be non-robust. A reader may wonder whether we can perhaps place more SC fences to perhaps not require individual library robustness, e.g. before and after every library invocation, regardless of whether the surrounding invocations belonged to the same library or not. However, this would also not work, because we can consider a library which has two methods corresponding to the two sessions involved in the store buffering anomaly, i.e. each method has two memory events, a store followed by load to different location, with opposite order of locations in the two methods. Essentially, each method performs the memory operations of one of the sessions involved in the store buffering anomaly. Then, in an execution where each of the methods is called in a separate session, even with the presence of SC fences before and after the invocations, store buffering is still possible.

7 Automated Verification

Our automated verification strategy relies on discovering violations of Φ_{src} and Φ_{dst} , which essentially corresponds to discovering potential and actual non-robustness witnesses. Towards this end, we model memory events using FOL domains, executions as relations between events and then we generate FOL queries that instantiate memory events involved in establishing Φ_{src} and Φ_{dst} . For example, referring back to Table 1, violations of Φ_{src} would require us to instantiate events μ_w, μ_1, μ_2 satisfying the antecedent of Φ_{src} , while the negation of the consequent can be simplified as $\neg(\mu_w \xrightarrow{\text{hb}} \mu_2)$.

In addition to memory events which are directly involved in Φ_{src} and Φ_{dst} , for the library method containing these events, the FOL query also instantiates an event for every program statement present in the implementation of the method. We call the entire set of instantiated events as a *partial*

execution, which is essentially the non-robust core of an actual execution. If such a non-robust core cannot be instantiated, then no potential or actual non-robustness witnesses can exist. Thus, if the generated FOL formulae are not satisfiable, we can conclude that Φ_{src} or Φ_{dst} hold, implying robustness of the library implementation.

Next, we require that the partial execution must be a part of a valid SC execution. To ensure this, we encode the constraints on the executions in terms of FOL formula obtained from the χ_{base} and χ_{SC} constraints of §3.2. For each relation $R \in \{rf, mo, fr, sw, so, hb, hb_{SC}\}$, we encode constraints that ensure or prohibit their presence between pairs of events. For instance, all events belonging to the same invocation must be related by so, two events writing to the same location must be related by mo, there must be a rf between a unique write event to a location and a read event returning the same value, etc. We also encode how derived relations depend on the base relations.

Note that Φ_{src} and Φ_{dst} need to be checked individually for every location used in any execution of the implementation. While an implementation can allocate and access an unbounded number of locations on the heap during executions, to generate our encoding, we use a fixed, finite number of *location classes*. We define a location class for every shared global variable and every field (for record types allocated on the heap). For example, the location classes for the register implementation of Fig. 2 are the global variable `L` and the field `val`.

Instantiating partial executions with a fixed, finite number of events may result in a number of false positives, i.e. partial executions that would not be a part of any complete execution. To prune these false positives, we introduce an analysis phase before checking Φ_{src} and Φ_{dst} that derives useful constraints obeyed by any SC execution of the library implementation. These constraints are then expressed as universally quantified FOL formulae and added to the encoding. The derived constraints can be broadly classified into two classes: (1) *program structure* constraints that are derived by a static analysis of the implementation and (2) *specification* constraints that are directly obtained from the specification of the implemented data structure (as described in §3.3). Since we assume that the library implementation is correct under SC, we can directly incorporate these specification constraints.⁴ We find these derived constraints to be effective in helping to establish robustness conditions over low-level memory events.

For each library implementation, we generate (i) a set of access constraints from the provided implementation code and the library specification that e.g., identify locations written to at most once in an execution, are always written to within the same invocation, etc., or (ii) CAS constraints that identify locations exclusively modified using `acqrel` CAS operations and thus are totally ordered under the `hb` relation, etc. We additionally extract (iii) specification constraints from the axiomatic, declarative specification of the data structure to generate additional `hb` edges. For example, for the register implementation shown in Figure 1 we add the specifications χ_{GETSET} and $\chi_{GETUNDEF}$, which state that a `get` should return the value from another set and if a `get` returns `UNDEF`, then no set method should have been executed before it.

```

Input: Library Implementation  $L$ , Specification
        axioms  $\Psi_D$ 
 $L' \leftarrow \rho_{CAS}(L)$ 
 $\Psi_{Analysis} \leftarrow \text{ConstraintAnalysis}(L', \Psi_D)$ 
 $\Psi \leftarrow \Psi_{base} \wedge \Psi_{SC} \wedge \Psi_{Analysis}$ 
foreach  $\ell \in \text{LocClass}(L')$  do
    if  $\neg \text{Check-}\Phi_{src}(\Psi, L', \ell)$  then
        if  $\neg \text{Check-}\Phi_{dst}(\Psi, L', \ell)$  then
            return  $L$  may be non-robust
        end
    end
end
return  $L$  is robust

```

Fig. 10. Main Algorithm

⁴Complete details about the SMT encoding and the derived constraints are provided in the supplemental material, §5.

Figure 10 depicts the algorithm for checking robustness, combining the various components discussed above. The algorithm takes as input a library implementation L , along with the specification axioms ($\Psi_{\mathcal{D}}$) of the implemented data structure \mathcal{D} . We first perform a robustness-preserving CAS-to-BCAS transformation ρ_{CAS} (shown in 5.4) to obtain the modified implementation L' . We then populate the constraint set $\Psi_{Analysis}$ as described above. We then iterate through all the location classes, and check if either Φ_{src} or Φ_{dst} hold for each location class. Each check for Φ_{src} or Φ_{dst} will check the feasibility of instantiating a partial execution involving a violation to the corresponding condition, in the presence of the constraints Ψ . This feasibility check is reduced to checking the satisfiability of a FOL formula.

8 Evaluation

We have implemented a tool called ROBOCOP to test our methodology on real-world benchmarks. ROBOCOP takes a library of methods written in C11 and produces a result that indicates if the library is robust under RC20. ROBOCOP directly implements the algorithm of Figure 10. ROBOCOP parses the provided inputs, performs the ρ_{CAS} robustness transformation (5.4), and generates the necessary SMTLIB encoding needed to specify constraints (§7), which is then discharged by Z3.

Our evaluation considers a number of real-world benchmarks adapted from the literature and open-source repositories. All these implementations have SC specifications which have been proven to be correct, but their robustness under RC20 has not been established automatically prior to this work. All experiments were executed on an Intel® Core™ i5-7200U CPU @ 2.50 GHz, Ubuntu 18.04 machine using Z3 4.8.10. Table 2 summarizes key results. Column Loc denotes the number of location classes considered for verification. Columns RLX, RA, and Tot denote the number of relaxed accesses, release-acquire operations (accesses and fences), and total operations performed, which includes accesses to location classes and fence operations in the library source code. As these numbers indicate, most of the benchmarks make meaningful use of relaxed and release-acquire operations. Our benchmarks cover the following commonly occurring access patterns in real-world libraries - static locations, dynamically allocated locations and locations accessed by an offset into an array.

Next, we describe situations where the derived constraints aid the automated verification procedure. We note that removing any of these constraints either produces a robustness violation that is a false-positive or the solver loops until timeout, trying to unroll and instantiate memory events.

Access Constraints. The Treiber stack, Lockfree queue, MPMC Unbounded queue and Non-blocking set follow the pattern of creating a new node on a push/add operation and then linking it to the main data structure. Thus, the updates to the fields of these new nodes can only be performed inside these methods and removing this constraint generates false-positive memory events where two writes may update this field and lead to a non-robustness witness, if observed by a read.

In the case of the SPSC bounded queue, MPMC bounded and unbounded queues and Chase-Lev deque, the threads access an underlying array, using integer indices. The access constraint-generated states that writes to different indices are made by different threads, ensuring that a read cannot potentially read from two different writes, leading to a non-robustness witness.

CAS Constraints: Since all the data structures are written to be lock-free, they make use of a CAS update to a designated location to update the data structure and retrying if the CAS fails. In the Treiber stack, this is the head of the stack and in the Lockfree queue, these are the head and tail of the queues. Similarly, the indices signalling where to write into the array for the SPSC bounded queue, MPMC bounded queue and MPMC unbounded queue are updated using fetch-and-add instructions, ensuring that a unique thread gets access to the index element. The CAS constraints ensure that two events are not instantiated that read from different orderings to these locations, since the CAS's impose a total order.

Table 2. Results of applying ROBOCOP to RC20 concurrent data structure libraries. All benchmarks were verified to be robust.

Benchmark	Time (s)	Locs	RLX	RA	Tot
Atomic Reference Counter [20]	10.21	3	2	4	6
Singleton [5]	8.43	4	0	6	6
Read Copy Update [39]	9.20	4	0	8	8
Spinlock [47]	12.54	2	2	4	6
Seqlock [10]	10.82	4	4	4	8
Ticketlock [47]	4.54	2	1	3	4
Lamport Mutex [42]	9.78	5	0	12	12
Peterson Lock [43]	6.54	3	3	3	6
Dekkers Mutex [54]	8.22	3	8	2	10
Treiber Stack [53]	13.67	3	5	3	8
Herlihy-Wing Queue [30]	12.39	2	0	4	4
TwoLock Queue [45]	18.12	6	7	7	14
Lockfree Queue [45]	24.55	4	5	7	12
SPSC Queue [50]	11.05	3	4	4	8
MPMC Bounded Queue [44]	18.45	4	2	8	10
MPMC Unbounded Queue [26]	25.74	6	4	12	16
Non-blocking Set [29]	19.42	3	6	6	12
Work-stealing Queue [17]	28.40	3	10	4	14

Specification constraints: For each benchmark, we add the specification constraint and map it to the internal program statements of the benchmarks. For example, the AddRem constraint for stacks, queues and sets ensures that in the respective implementations, there is an hb edge between a push and a pop that have matching argument and return values. This is translated to an hb edge between the program points in the method invocations where the method executes, in the SC order.

In the case of the libraries that create and link nodes, such as Treiber stack and Lockfree queue, this ensures that a pop operation always reads from a unique push operation (thus ruling out a non-robust witness, where a pop may read from two different push operations). Similarly, for implementations which use indexing, this also ensures that there is a single push operation at a certain index, that the pop operation reads from.

To demonstrate that our approach can also detect robustness violations, we systematically relaxed memory access statements in these benchmarks to create a non-robust implementation. ROBOCOP was successfully able to provide counter-examples in the form of partial executions that serve as a witness to the violation for all modified benchmarks.⁵

Finally, we performed an experiment to understand the impact of the SC-fence insertion overhead in the presence of multiple instances of a stack or queue library. We create a benchmark scenario where there are N producers, N consumers and N intermediaries. The producers add messages to a queue and the intermediaries remove messages from this queue. Then, the intermediaries push these messages to another queue and finally the consumers remove messages from the second queue. This benchmark models a message bus or pipeline, observed in real-world software. To ensure compositionality, the intermediaries need to insert a fence after the pop from the first queue and before the push to the second queue. We measure how long it takes to complete a fixed number

⁵Additional details are provided in the supplemental material.

of operations, from producers to consumers. We run the experiments on a `c7g.metal` AWS instance, that has 64 ARMv8.4 cores. We run the benchmarks for 1M operations by each thread, over $N = 16$ consumers ($N = 1$ for SPSC queue), producers and intermediaries. We run the benchmark for 3 cases, where in each benchmark the intermediary data structures are set to be the Boost 1.74 libraries for SPSC queue, MPMC queue and MPMC stack and observe 4.79%, 7.39% and 3.69% increase in runtime when SC fences (DMB ISH) are inserted between pop and push operations.

9 Related Work and Conclusion

There has been substantial prior work on determining robustness against hardware models [3, 4, 12, 19], with the x86-TSO memory model being particularly well-studied [11, 48]. Guaranteeing data-race freedom (DRF-SC) [2, 21] is a well-known instantiation of robustness applicable to a language's concurrency semantics. The notion of execution graph robustness closely resembles the DRF-SC property, and indeed prior work [43] has formally proved a correspondence between the two. However, in this work, our main contribution is a fully automated approach for verifying this property in the context of library implementations in the presence of the most general client. As discussed earlier, [38] presents a method to verify execution-graph robustness against the C11 release-acquire concurrency semantics, implementing their procedure using a model-checker that operates over programs with a finite data domain. [43] extends this result to additionally support relaxed accesses and release/acquire fences. Their approach also takes into account speculative actions as part of their robustness formulation, similar to our definition of effect robustness, but as described earlier, our notion of robustness is weaker and better suited for library implementations. Other works [34–36] have also proposed model-checking based techniques for verifying programs under weak memory that cannot, however, be soundly used for verifying library implementations in the presence of a most-general client.

A number of prior works have also considered the specification and verification problem of libraries in a relaxed memory setting [6, 13, 22, 24, 31, 49, 52]; these efforts, however, do not consider automated verification tooling or robustness arguments in their proof methodology. [14, 16, 27, 28, 32] also propose correctness notions that are weaker than linearizability - we consider the incorporation of these ideas as a topic for future research.

Recent work [18, 44, 49] has developed new proof techniques to modularly reason about clients that interact with libraries which often have weaker specifications that expose relaxed memory behavior. We focus on an orthogonal problem in this paper - establishing the robustness of library implementations that internally use relaxed memory primitives. As we have shown here, libraries that internally use relaxed accesses may have enough synchronization to make them robust, enabling a pathway to automated verification. Having said that, we also believe that synchronization specifications of libraries as proposed by these other efforts would be useful in addition to robustness guarantees, since they would allow an optimal SC fence placement strategy for guaranteeing robustness of executions involving multiple libraries.

This paper presents a modular verification strategy for verifying robustness of programs using library implementations. Our verification strategy adapts the notion of execution-graph robustness to the library setting, and exploits specification axioms of the library under SC, to generate constraints sufficient to imply a suitable inductive robustness invariant. We also show how to effectively compose robustness guarantees of multiple libraries. We have successfully demonstrated our technique on a number of challenging real-world concurrent data structure implementations that meaningfully exploit sophisticated weak memory behavior. Our results suggest that automated robustness proofs can be effectively applied to ascertain whether concurrent libraries, specified assuming sequential consistency, can be safely refined to exploit relaxed atomics.

Data-Availability Statement

The software that supports Section. 8 is available on Zenodo [46].

Acknowledgments

This material is based in part upon work supported by the National Science Foundation under grant CCF-FMiTF 2019263. The first author is also supported by a grant from the Tezos Foundation.

References

- [1] Parosh Aziz Abdulla, Frédéric Haziza, Lukás Holík, Bengt Jonsson, and Ahmed Rezzine. 2013. An Integrated Specification and Verification Technique for Highly Concurrent Data Structures. In *TACAS*.
- [2] Sarita V. Adve and Mark D. Hill. 1990. Weak Ordering—a New Definition. *SIGARCH Comput. Archit. News* 18, 2SI (may 1990), 2–14. <https://doi.org/10.1145/325096.325100>
- [3] 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 (2017), 6:1–6:38. <https://doi.org/10.1145/2994593>
- [4] Jade Alglave and Luc Maranget. 2011. Stability in Weak Memory Models. In *Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings (Lecture Notes in Computer Science, Vol. 6806)*, Ganesh Gopalakrishnan and Shaz Qadeer (Eds.). Springer, 50–66. https://doi.org/10.1007/978-3-642-22110-1_6
- [5] Helge Bahmann and Tim Blechmann. 2012. https://www.boost.org/doc/libs/1_85_0/libs/atomic/doc/html/atomic/usage_examples.html#boost_atomic.usage_examples.singleton. Accessed 31-07-2024.
- [6] Mark Batty, Mike Dodds, and Alexey Gotsman. 2013. Library Abstraction for C/C++ Concurrency. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. Association for Computing Machinery, 235–248.
- [7] Mark Batty, Kayvan Memarian, Scott Owens, Susmit Sarkar, and Peter Sewell. 2012. Clarifying and Compiling C/C++ Concurrency: from C++11 to POWER. In *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, (POPL 2012), Philadelphia, Pennsylvania, USA, January 22-28, 2012*. 509–520. <https://doi.org/10.1145/2103656.2103717>
- [8] Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber. 2011. Mathematizing C++ concurrency. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, Thomas Ball and Mooly Sagiv (Eds.). ACM, 55–66. <https://doi.org/10.1145/1926385.1926394>
- [9] Hans-Juergen Boehm and Brian Demsky. 2014. Outlawing Ghosts: Avoiding Out-Of-Thin-Air Results. In *Proceedings of the workshop on Memory Systems Performance and Correctness, MSPC '14, Edinburgh, United Kingdom, June 13, 2014*, Jeremy Singer, Milind Kulkarni, and Tim Harris (Eds.). ACM, 7:1–7:6. <https://doi.org/10.1145/2618128.2618134>
- [10] Hans-J. Boehm. 2012. Can seqlocks get along with programming language memory models?. In *Proceedings of the 2012 ACM SIGPLAN Workshop on Memory Systems Performance and Correctness (Beijing, China) (MSPC '12)*. Association for Computing Machinery, New York, NY, USA, 12–20. <https://doi.org/10.1145/2247684.2247688>
- [11] Ahmed Bouajjani, Egor Derevenetc, and Roland Meyer. 2013. Checking and Enforcing Robustness against TSO. In *Programming Languages and Systems - 22nd European Symposium on Programming, ESOP 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings (Lecture Notes in Computer Science, Vol. 7792)*, Matthias Felleisen and Philippa Gardner (Eds.). Springer, 533–553. https://doi.org/10.1007/978-3-642-37036-6_29
- [12] Sebastian Burckhardt, Rajeev Alur, and Milo M. K. Martin. 2007. CheckFence: Checking Consistency of Concurrent Data Types on Relaxed Memory Models. In *Proceedings of the ACM SIGPLAN 2007 Conference on Programming Language Design and Implementation, San Diego, California, USA, June 10-13, 2007*, Jeanne Ferrante and Kathryn S. McKinley (Eds.). ACM, 12–21. <https://doi.org/10.1145/1250734.1250737>
- [13] Sebastian Burckhardt, Alexey Gotsman, Madanlal Musuvathi, and Hongseok Yang. 2012. Concurrent Library Correctness on the TSO Memory Model. In *European Symposium on Programming*, Helmut Seidl (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 87–107.
- [14] Sebastian Burckhardt, Alexey Gotsman, Hongseok Yang, and Marek Zawirski. 2014. Replicated data types: specification, verification, optimality. In *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20-21, 2014*. 271–284. <https://doi.org/10.1145/2535838.2535848>
- [15] C11 Last Accessed Nov 9 2022. C11 Memory Model (std::memory_order). https://en.cppreference.com/w/cpp/atomic/memory_order.

- [16] Armando Castañeda, Sergio Rajsbaum, and Michel Raynal. 2015. Specifying Concurrent Problems: Beyond Linearizability and up to Tasks. In *Distributed Computing*, Yoram Moses (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 420–435.
- [17] David Chase and Yossi Lev. 2005. Dynamic Circular Work-Stealing Deque. In *Proceedings of the Seventeenth Annual ACM Symposium on Parallelism in Algorithms and Architectures* (Las Vegas, Nevada, USA) (SPAA '05). Association for Computing Machinery, New York, NY, USA, 21–28. <https://doi.org/10.1145/1073970.1073974>
- [18] Hoang-Hai Dang, Jaehwang Jung, Jaemin Choi, Duc-Thuan Nguyen, William Mansky, Jeehoon Kang, and Derek Dreyer. 2022. Compass: Strong and Compositional Library Specifications in Relaxed Memory Separation Logic. In *PLDI '22: 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation, San Diego, CA, USA, June 13 - 17, 2022*, Ranjit Jhala and Isil Dillig (Eds.). ACM, 792–808. <https://doi.org/10.1145/3519939.3523451>
- [19] Egor Derevenetc and Roland Meyer. 2014. Robustness Against POWER is PSpace-Complete. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 8573)*, Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias (Eds.). Springer, 158–170. https://doi.org/10.1007/978-3-662-43951-7_14
- [20] Marko Doko and Viktor Vafeiadis. 2017. Tackling Real-Life Relaxed Concurrency with FSL++. In *Programming Languages and Systems*, Hongseok Yang (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 448–475.
- [21] Stephen Dolan, K. C. Sivaramakrishnan, and Anil Madhavapeddy. 2018. Bounding Data Races in Space and Time. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2018, Philadelphia, PA, USA, June 18-22, 2018*, Jeffrey S. Foster and Dan Grossman (Eds.). ACM, 242–255. <https://doi.org/10.1145/3192366.3192421>
- [22] Brijesh Dongol, Radha Jagadeesan, James Riely, and Alasdair Armstrong. 2018. On Abstraction and Compositionality for Weak-Memory Linearisability. In *Verification, Model Checking, and Abstract Interpretation (VMCAI)*. Springer International Publishing, 183–204.
- [23] Michael Emmi and Constantin Enea. 2018. Sound, Complete, and Tractable Linearizability Monitoring for Concurrent Collections. *Proc. ACM Program. Lang.* 2, POPL (2018), 25:1–25:27. <https://doi.org/10.1145/3158113>
- [24] Michael Emmi and Constantin Enea. 2019. Weak-Consistency Specification via Visibility Relaxation. *Proc. ACM Program. Lang.* 3, POPL, Article 60 (Jan 2019), 28 pages. <https://doi.org/10.1145/3290373>
- [25] Michael Emmi, Constantin Enea, and Jad Hamza. 2015. Monitoring Refinement via Symbolic Reasoning. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation, Portland, OR, USA, June 15-17, 2015*, David Grove and Stephen M. Blackburn (Eds.). ACM, 260–269. <https://doi.org/10.1145/2737924.2737983>
- [26] Folly Last Accessed Nov 9 2022. Meta Folly MPMC Queue. <https://github.com/facebook/folly/blob/main/folly/MPMCQueue.h>.
- [27] Andreas Haas, Thomas A. Henzinger, Andreas Holzer, Christoph M. Kirsch, Michael Lippautz, Hannes Payer, Ali Sezgin, Ana Sokolova, and Helmut Veith. 2016. Local Linearizability for Concurrent Container-Type Data Structures. In *27th International Conference on Concurrency Theory (CONCUR 2016) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 59)*, José Desharnais and Radha Jagadeesan (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 6:1–6:15. <https://doi.org/10.4230/LIPIcs.CONCUR.2016.6>
- [28] Nir Hemed, Noam Rinetzk, and Viktor Vafeiadis. 2015. Modular Verification of Concurrency-Aware Linearizability. In *Distributed Computing*, Yoram Moses (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 371–387.
- [29] Maurice Herlihy and Nir Shavit. 2012. *The Art of Multiprocessor Programming, Revised Reprint* (1st ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- [30] M Herlihy and J Wing. 1990. Linearizability: A Correctness Condition for Concurrent Objects. In *ACM TOPLAS*, 12(3).
- [31] Radha Jagadeesan, Gustavo Petri, Corin Pitcher, and James Riely. 2013. Quarantining Weakness. In *European Symposium on Programming*, Matthias Felleisen and Philippa Gardner (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 492–511.
- [32] Radha Jagadeesan and James Riely. 2014. Between Linearizability and Quiescent Consistency. In *Automata, Languages, and Programming*, Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 220–231.
- [33] Jeehoon Kang, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, and Derek Dreyer. 2017. A Promising Semantics for Relaxed-Memory Concurrency. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages* (Paris, France) (POPL 2017). Association for Computing Machinery, New York, NY, USA, 175–189. <https://doi.org/10.1145/3009837.3009850>
- [34] Michalis Kokologiannakis, Ori Lahav, Konstantinos Sagonas, and Viktor Vafeiadis. 2018. Effective stateless model checking for C/C++ concurrency. *Proc. ACM Program. Lang.* 2, POPL (2018), 17:1–17:32. <https://doi.org/10.1145/3158105>
- [35] Michalis Kokologiannakis, Iason Marmanis, Vladimir Gladstein, and Viktor Vafeiadis. 2022. Truly Stateless, Optimal Dynamic Partial Order Reduction. *Proc. ACM Program. Lang.* 6, POPL (2022), 1–28. <https://doi.org/10.1145/3498711>

- [36] Michalis Kokologiannakis, Azalea Raad, and Viktor Vafeiadis. 2019. Model Checking for Weakly Consistent Libraries. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2019, Phoenix, AZ, USA, June 22-26, 2019*, Kathryn S. McKinley and Kathleen Fisher (Eds.). ACM, 96–110. <https://doi.org/10.1145/3314221.3314609>
- [37] Ori Lahav, Nick Giannarakis, and Viktor Vafeiadis. 2016. Taming Release-Acquire Consistency. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (St. Petersburg, FL, USA) (POPL '16). Association for Computing Machinery, New York, NY, USA, 649–662. <https://doi.org/10.1145/2837614.2837643>
- [38] Ori Lahav and Roy Margalit. 2019. Robustness Against Release/Acquire Semantics. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2019, Phoenix, AZ, USA, June 22-26, 2019*, Kathryn S. McKinley and Kathleen Fisher (Eds.). ACM, 126–141. <https://doi.org/10.1145/3314221.3314604>
- [39] Ori Lahav and Viktor Vafeiadis. 2015. Owicki-Gries Reasoning for Weak Memory Models. In *Automata, Languages, and Programming*, Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 311–323.
- [40] Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, and Derek Dreyer. 2017. Repairing Sequential Consistency in C/C++11. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2017, Barcelona, Spain, June 18-23, 2017*, Albert Cohen and Martin T. Vechev (Eds.). ACM, 618–632. <https://doi.org/10.1145/3062341.3062352>
- [41] Leslie Lamport. 1979. How to Make a Multiprocessor Computer That Correctly Executes Multiprocess Programs. *IEEE Trans. Computers* 28, 9 (1979), 690–691. <https://doi.org/10.1109/TC.1979.1675439>
- [42] Leslie Lamport. 1987. A fast mutual exclusion algorithm. *ACM Trans. Comput. Syst.* 5, 1 (jan 1987), 1–11. <https://doi.org/10.1145/7351.7352>
- [43] Roy Margalit and Ori Lahav. 2021. Verifying Observational Robustness Against a C11-style Memory Model. *Proc. ACM Program. Lang.* 5, POPL (2021), 1–33. <https://doi.org/10.1145/3434285>
- [44] Glen Mével and Jacques-Henri Jourdan. 2021. Formal Verification of a Concurrent Bounded Queue in a Weak Memory Model. *Proc. ACM Program. Lang.* 5, ICFP, Article 66 (aug 2021), 29 pages. <https://doi.org/10.1145/3473571>
- [45] Maged M. Michael and Michael L. Scott. 1996. Simple, Fast, and Practical Non-Blocking and Blocking Concurrent Queue Algorithms. In *Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing, Philadelphia, Pennsylvania, USA, May 23-26, 1996*, James E. Burns and Yoram Moses (Eds.). ACM, 267–275. <https://doi.org/10.1145/248052.248106>
- [46] Kartik Nagar, Anmol Sahoo, Romit Roy Chowdhury, and Suresh Jagannathan. 2024. *Artifact - Automated Robustness Verification of Concurrent Data Structure Libraries Against Relaxed Memory Models*. <https://doi.org/10.5281/zenodo.13626195>
- [47] Scott Owens. 2010. Reasoning about the Implementation of Concurrency Abstractions on x86-TSO. In *European Conference on Object-Oriented Programming*. <https://api.semanticscholar.org/CorpusID:28735505>
- [48] Scott Owens. 2010. Reasoning about the Implementation of Concurrency Abstractions on x86-TSO. In *ECOOP 2010 - Object-Oriented Programming, 24th European Conference, Maribor, Slovenia, June 21-25, 2010. Proceedings (Lecture Notes in Computer Science, Vol. 6183)*, Theo D'Hondt (Ed.). Springer, 478–503. https://doi.org/10.1007/978-3-642-14107-2_23
- [49] Azalea Raad, Marko Doko, Lovro Rozic, Ori Lahav, and Viktor Vafeiadis. 2019. On Library Correctness Under Weak Memory Consistency: Specifying and Verifying Concurrent Libraries Under Declarative Consistency Models. *PACMPL* 3, POPL (2019), 68:1–68:31. <https://doi.org/10.1145/3290381>
- [50] Boris Schling. 2011. *The Boost C++ Libraries*. XML Press.
- [51] Jaroslav Sevcik, Viktor Vafeiadis, Francesco Zappa Nardelli, Suresh Jagannathan, and Peter Sewell. 2013. CompCertTSO: A Verified Compiler for Relaxed-Memory Concurrency. *J. ACM* 60, 3 (2013), 22:1–22:50. <https://doi.org/10.1145/2487241.2487248>
- [52] Abhishek Kr Singh and Ori Lahav. 2023. An Operational Approach to Library Abstraction under Relaxed Memory Concurrency. *Proc. ACM Program. Lang.* 7, POPL (2023), 1542–1572. <https://doi.org/10.1145/3571246>
- [53] R.K. Treiber. 1986. *Systems programming: Coping with Parallelism. Technical Report RJ-5118*. Technical Report. IBM Almaden Research Center.
- [54] Anthony Williams. Last Accessed Nov 9, 2022. Implementing Dekker's algorithm with fences. https://www.justsoftwaresolutions.co.uk/threading/implementing_dekkers_algorithm_with_fences.html

Received 2024-04-06; accepted 2024-08-18