

An Operational Approach to Library Abstraction under Relaxed Memory Concurrency

ABHISHEK KR SINGH, Tel Aviv University, Israel

ORI LAHAV, Tel Aviv University, Israel

Concurrent data structures and synchronization mechanisms implemented by expert developers are indispensable for modular software development. In this paper, we address the fundamental problem of library abstraction under weak memory concurrency, and identify a general library correctness condition allowing clients of the library to reason about program behaviors using the specification code, which is often much simpler than the concrete implementation. We target (a fragment of) the RC11 memory model, and develop an equivalent operational presentation that exposes knowledge propagation between threads, and is sufficiently expressive to capture library behaviors as totally ordered operational execution traces. We further introduce novel access modes to the language that allow intricate specifications accounting for library internal synchronization that is not exposed to the client, as well as the library's demands on external synchronization by the client. We illustrate applications of our approach in several examples of different natures.

CCS Concepts: • **Theory of computation** → **Concurrency**; **Operational semantics**; **Program verification**; • **Software and its engineering** → **Semantics**; **Software verification**; **Abstraction, modeling and modularity**.

Additional Key Words and Phrases: Relaxed memory consistency, Concurrent objects, Linearizability, Library abstraction

ACM Reference Format:

Abhishek Kr Singh and Ori Lahav. 2023. An Operational Approach to Library Abstraction under Relaxed Memory Concurrency. *Proc. ACM Program. Lang.* 7, POPL, Article 53 (January 2023), 31 pages. <https://doi.org/10.1145/3571246>

1 INTRODUCTION

Library abstraction constitutes a powerful means to achieve modularity in software development. It allows expert library developers to write optimized implementations of common programming tasks, and “once and for all” establish that these implementations admit their corresponding specifications. In turn, users of these implementations, called *clients* of the library, may reason about program behaviors assuming only the libraries' specifications, with no understanding, or even no access, to the implementations. From a formal standpoint, a prerequisite for applying library abstraction is to identify a condition that provably allows library developers to abstract away from a particular client program when verifying their implementations, while ensuring the soundness of client reasoning using the specifications in any (valid) client program.

In the case of sequential programs, library abstraction is straightforward: specifications may be given using pre- and post-conditions, and it is easy to establish the soundness of client reasoning that is based on them. Concurrent programs, however, require more attention. Assuming

Authors' addresses: [Abhishek Kr Singh](#), Tel Aviv University, Israel, abhishek.uor@gmail.com; [Ori Lahav](#), Tel Aviv University, Israel, orilahav@tau.ac.il.



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

© 2023 Copyright held by the owner/author(s).

2475-1421/2023/1-ART53

<https://doi.org/10.1145/3571246>

an underlying sequentially consistent (SC) memory system [Lamport 1979], classical linearizability (w.r.t. some sequential specification) [Herlihy and Wing 1990] ensures refinement w.r.t. a sequential object [Filipović et al. 2010], and can be seen as a library abstraction condition. More recent approaches employ “code as specification” (see, e.g., [Gotsman and Yang 2011]), where implementations are specified using (simple) code, and library abstraction amounts to contextual refinement between two pieces of code. Linearizability can be seen as a particular instance of this approach, where specifications are obtained by wrapping a sequential implementation within a global lock [Bouajjani et al. 2015]. Furthermore, to make such specification formalism more useful, one often enriches the given programming language with special specification constructs, that are employed only in the specification code and are relatively easy to reason about (e.g., atomic blocks in concurrent programs).

The current paper studies library abstraction under relaxed (a.k.a. weak) shared-memory concurrency. In particular, we consider a C11-style weak memory model with several kinds of memory-access modes (a.k.a. memory orderings), placed differently on the spectrum between efficient implementation (with an optimizing compiler targeting a modern multicore hardware) and strong consistency guarantees. The infamous complexity of programming under such a model makes library abstraction indispensable. In particular, library abstraction allows clients to use the programming guarantees supplied by the model, e.g., clients should be able to rely on the synchronization induced by a library (such as a lock library) for ensuring data-race freedom in their programs, and then apply the model’s guarantee that data-race free programs have strong semantics.

Remark 1. Like many formal verification frameworks for C11, we are unable to work with the original model, which allows unrestricted cycles in the union of ‘program order’ and ‘reads-from’, and thus exhibits “out-of-thin-air” behaviors and fails to provide the most basic data-race-freedom guarantee. We follow [Boehm and Demsky 2014] and its formalization in the RC11 model in [Lahav et al. 2017] to conservatively disallow *all* such cycles. This entails a certain performance penalty for maintaining the load-store order between relaxed accesses [Ou and Demsky 2018]. We also note that the fragment of the RC11 handled in this paper lacks release/acquire fences and sequentially consistent accesses, which are left to future work.

Our main contribution is a correctness condition for libraries that provably ensures (contextual) refinement between implementations and their respective specifications under the weak memory model, with several distinctive properties:

- Our proposed condition is based on *totally ordered execution traces* of the library in question (a decision motivated below). Accordingly, we present a novel equivalent operational version of the (originally) declarative memory model, and use it in the correctness condition.
- To account for various restrictions that libraries may impose on their clients (e.g., refrain from data races between two specific methods), our result supports rich library calling policies, and refinement is conditioned on the client’s adherence to the calling policy. Importantly, whether a program adheres to the policy or not is checked against the library specification, which allows the application of the abstraction theorem by the client without any knowledge of the implementation.
- Our specification language allows simple *lock-based* specifications of libraries that do not provide synchronization guarantees for their clients, and dually of libraries that rely on the client for performing the synchronization between library calls. For that matter, we introduce a novel access mode to the programming language (and memory model).

We illustrate the application of our approach for an RCU synchronization mechanism which we specify by relying solely on locks (§8.1), as well as for a relaxed concurrent queue object that does not expose its internal synchronization to its clients (§8.2). We also derive a (local) data-race-freedom guarantee for the memory model as an instance of library abstraction (§8.3).

Outline. This paper is organized as follows. In §2 we present an informal overview of the challenges we address and our solutions. In §3 we define concurrent programs and their semantics independent of a memory model. In §4 we present the memory model semantics, which we call dRC11 (d for ‘declarative’), as a fragment of RC11 extended with novel constraints for the new access modes. In §5 we present an operational version of the memory model, which we call pRC11 (p for ‘propagation’), that is needed for defining the library correctness condition. In §6 we introduce the notions required to formulate the library abstraction theorem. In §7 we state and prove the main theorem. In §8 we demonstrate applications of the abstraction theorem (which can assist in understanding the crux of the abstraction theorem even before reading the more technical material in Section 3 to 7). We discuss the relation to other work in §9 and conclude in §10. The *full version* of this paper available in [Singh and Lahav 2022] contains additional detailed proofs.

2 KEY CHALLENGES AND IDEAS

We outline the main challenges and the key ideas in our solutions. We keep the discussion and examples informal, leaving the formal development to later sections.

2.1 Library Correctness Criterion

The main challenge lies in establishing a library correctness condition that will ensure contextual refinement between a library specification $L^\#$ (given as code) and its implementation L under weak memory semantics. Naturally, such condition should consider each object in isolation and avoid quantification over all possible clients. Moreover, we opt for an *operational* condition that, like standard linearizability, is based on totally ordered histories generated by the object in question. The latter desideratum is in contrast with previous work on library specifications in (R)C11 [Batty et al. 2013; Raad et al. 2019] and adaptations of linearizability to weak memory concurrency (see, e.g., [Dongol et al. 2018]), which employ multiple partial orders in their correctness criteria. We believe that an operational approach, based on *one* total timeline, may be more intuitive for users (also considering informal arguments), and will allow easier adoption of standard techniques and tools that were applied before to verify refinement between transition systems (see e.g., our initial experiment with the FDR refinement checker in §8.1).

While having an operational correctness criterion may seem contradicting to the fact that standard weak memory formulations, e.g., (R)C11, are *declarative* (a.k.a. axiomatic memory models), unlike previous work (e.g., [Batty et al. 2013]), we consider this to be only a superficial matter. Indeed, declarative specifications of models that forbid “reads from the future” (i.e., impose acyclicity of the union of the ‘program order’ and the ‘read-from’ relations; called an ‘in-order’ semantics in [Cho et al. 2021]) can be equally characterized as operational memory models: the states consist of the current execution trace ordered by multiple partial order relations as needed, and transitions are only between *consistent* states as specified by the consistency condition of the declarative model (see, e.g., the RAG transition system for the release/acquire fragment of C11 in [Lahav and Margalit 2019]). More concise (and possibly more intuitive) formulations of such semantics may use timestamps and views, and the equivalence to the “operationalized” declarative model is witnessed by a standard forward simulation relation (see, e.g., [Kaiser et al. 2017; Kang et al. 2017]).

Now, operational semantics (in particular, operationalized versions of declarative models), naturally give rise to a set of *library histories*, and allow the application of standard linearizability. For that matter, one considers a “most general client” program that invokes the library methods in the most general way expected by the library. Then, the set of library histories consists of all sequences of invocations (with the values of the arguments) and responses (with the returned values) of the library methods by different threads that are obtained in operational traces of that client. Under SC,

inclusion of the sets of histories between two libraries may serve as a library abstraction condition (i.e., if every history of a library L is also a history of $L^\#$, then L refines $L^\#$).

Aiming to apply a similar condition, following previous work on abstraction under TSO [Burckhardt et al. 2012] and under non-volatile memory [Khyzha and Lahav 2022], we observe that under weak memory semantics, inclusion of sets of histories is generally *unsound* as a library correctness condition.¹ Next, we provide a simple (contrived) example of this issue.

Example 2.1. Consider a library L with two methods `foo` and `bar` and the following specification (here and henceforth we assume that all variables are initialized to 0):

```
foo(): store(x, 1, rel); return();
```

```
bar(): a := load(x, acq); return(a);
```

Suppose that the library enforces a policy on its clients: (1) `foo` and `bar` must be called *exactly once* in *different* threads; and (2) `bar` must be called *after* `foo` in the execution order. For example, the following program adheres to this call policy:

```
foo();
store(y, 1, rlx);
||
b := load(y, rlx);
if b = 1 then
c := bar();
```

The accesses to `y` ensure that in the generated traces the call to `bar` indeed appears only after `foo` returns. Nevertheless, speaking in C11 terminology, since these accesses are relaxed (as annotated by `rlx` mode), they do *not* induce a *happens-before* order between the calls, and thus, following the specification, the read from `x` inside `bar` may return 0 (the initial value) or 1. Accordingly, a “naive” library developer may attempt to efficiently implement L as follows (where \oplus denotes a non-deterministic choice):

```
foo(): return();
```

```
bar(): return(0)  $\oplus$  return(1);
```

Is this implementation correct? In histories of the most general client of L that respects the library’s calling policy (i.e., in histories generated by executing the above program) one cannot observe a difference between the implementation and the specification. But, nevertheless, contextual refinement does *not* hold (so the implementation cannot be considered correct) for two different reasons illustrated by the following client programs:

```
store(z, 1, rlx);
foo();
store(y, 1, rlx);
||
b := load(y, rlx); // 1
if b = 1 then
c := bar(); // 1
d := load(z, rlx); // 0
```

```
foo();
store(y, 1, rel);
||
b := load(y, acq); // 1
if b = 1 then
c := bar(); // 0
```

The annotated behaviors in both examples are allowed when the suggested implementation of `foo` and `bar` is used, but disallowed when the specification is used. Indeed, in the example on the left, since the specification ensures release/acquire synchronization when `bar` returns 1 (via the `rel` and `acq` annotations in the accesses to `x`), the write to `z` happens-before the read from `z`, and in this case the model disallows the read from `z` to observe the overwritten initial value. In turn, in the example on the right, since the *client* uses release/acquire accesses, the write in `foo` happens-before the read in `bar`, which similarly ensures that the read from `x` cannot observe the initial value.

To address this challenge, we need to make library histories more expressive, so we can avoid the (SC) pitfall that identifies the execution order and the synchronization order. Concretely, in the

¹Interestingly, it follows from our result that for a language that employs *only* RC11-style relaxed accesses, inclusion of sets of histories is a sound condition (like it is for SC). Such a language is, however, too weak to be considered useful.

example above, we need to expose the facts that: (1) `bar` returning 1 *entails* a happens-before relation from *the call of `foo` to the return of `bar`*; and (2) `bar` returning 0 *forbids* a happens-before relation from *the return of `foo` to the call of `bar`*. What does that mean in operational terms without talking about the happens-before partial order between call and return actions? Our solution consists of the following:

- We introduce a novel operational semantics of the memory model and show that it is equivalent to the original declarative model. The operational semantics makes *knowledge propagation* between threads explicit. Roughly speaking, an access that is executed by one thread is first unknown to the other threads, and later non-deterministically propagates to each other threads. Access modes impose certain constraints on the propagation order and the ability to read from (locally) unknown events. In the simplest case—for programs with only release/acquire accesses—the propagation order has to follow the program order and reads can only read from known writes.
- We include propagation of method invocations and responses in the memory trace and, in turn, in library histories. These steps have no effect on the outcomes of the operational semantics (i.e., what values can be read and when), but they serve as crucial “markers” in library history, making histories expressive enough for validating contextual refinement given the inclusion of the sets of histories. In other words, we keep the inclusion of the sets of library histories as the correctness criterion, and recover its soundness by including appropriate markers (call and return propagation among different threads) in the histories.

In particular, revisiting the example above, the proposed implementation will have a history in which the propagation of the call marker of `foo` is after the return of `bar` but `bar` returns 1; as well as a history in which the propagation of the return marker of `foo` is before the call of `bar` but `bar` returns 0. Both histories are impossible for the specification, so library correctness does not hold.

2.2 Specification under Relaxed Memory Concurrency

A second challenge that we address is related to specification of libraries that expose weak behaviors. A straightforward approach to specifying concurrent data structures is to take a simple sequential implementation of the data structure, and wrap every operation inside a per-object lock. In fact, correctness according to the classical linearizability criterion [Herlihy and Wing 1990] is equivalent to refinement with respect to such lock-based specifications [Bouajjani et al. 2015; Filipović et al. 2010]. Since a lock entails a total order on all operations, the resulting specification is typically easy to understand and enables reasoning similar to reasoning about sequential programs.

This approach, however, has major shortcomings under relaxed memory concurrency, as it identifies the execution order with the synchronization order, which is only justified assuming SC. For instance, consider the following program that uses a concurrent queue (inspired by [Mével and Jourdan 2021]):

```
x := 1;           || a := q.dequeue; // 1
q.enqueue(1);    || b := x; // 0
```

Assuming a lock-based specification of the queue, the client may easily conclude that the annotated behavior ($a = 1$ and $b = 0$, where 0 is the initial value of x and the queue is initially empty) is impossible. Indeed, the fact that the dequeue operation returned the enqueued value entails that the lock of the dequeue operation must have been acquired after the lock of the enqueue operation was released. In turn, since locks necessarily provide release/acquire synchronization, the write to x happens-before the read from x , and thus the read cannot observe the overwritten initial value. Now, while the implementation of a queue may provide such synchronization guarantees for its client, it is also possible that the queue is implemented using efficient C11 relaxed accesses,

and thus comes with weaker guarantees that do not allow the client to rely on library-induced synchronization (so the accesses to x above are racy and $a = 1$ but $b = 0$ is possible). In that case the lock-based specification would be too strong.

We note that such weak guarantees do not imply that the queue is not “linearizable”: there may still exist a total order on all queue operations that agrees with the execution order (which can be any total order extending the union of the program order and the reads-from relation), and the queue exhibits FIFO behavior w.r.t. that order. A weak queue may be useful transferring elements of base types, whereas, in order to be used for transferring ownership via pointers, the client is responsible to appropriately place fences before and after invoking the weak queue methods.

Remark 2. While release/acquire synchronization may seem natural between enqueue and dequeue of the same element, the lock-based specification implies synchronization also in other cases, which may or may not be supplied (the queue studied in [Mével and Jourdan 2021] provides a case in point). For example, a client relying on a lock-based specification can deduce that the following annotated outcomes are all disallowed (where \perp denotes an empty queue):

$x := 1;$	\parallel	$q.\text{enqueue}(2);$	\parallel	$b := q.\text{dequeue};$	$// 1$
$q.\text{enqueue}(1);$	\parallel	$a := x;$	$// 0$	$c := q.\text{dequeue};$	$// 2$

$x := 1;$	\parallel	$q.\text{enqueue}(1);$
$a := q.\text{dequeue};$	$// \perp$	$b := x;$
		$// 0$

Accordingly, the challenge lies in identifying specification constructs that can be used instead of standard locks and be sufficiently flexible to account for different synchronization guarantees for the client. Ideally, like locks, they should allow straightforward reasoning about the library behaviors. We observe that existing concurrency models, C11 in particular, lack such constructs.

To address this challenge, we propose to extend the language with specialized instructions. Concretely, we identify that there is a range of possible guarantees that lie between C11’s release/acquire accesses (which provide synchronization) and relaxed accesses (which do not), and introduce a novel type of memory accesses that lies on this spectrum. Roughly speaking, these memory accesses, which we call *partial release/acquire*, ensure synchronization only for certain variables (see §4 for the formal definition).² C11’s release/acquire accesses can be seen as partial release/acquire accesses that ensure synchronization for *all* variables, whereas C11’s relaxed accesses can be seen as partial release/acquire accesses that do *not* ensure synchronization for *any* variable.

Partial release/acquire accesses can be used to construct locks that provide synchronization *only* for library variables, and these special locks can be used to specify libraries with weak synchronization guarantees. In particular, in the queue example above, locks implemented by *partial* release/acquire accesses would behave just like standard locks from library perspective, but will not allow the clients to rely on their induced synchronization. More generally, these accesses allow a library-internal release/acquire synchronization, which is needed for correct behaviors of the library, but is considered invisible from the point of view of the client.

3 CONCURRENT PROGRAMS: SYNTAX AND MEMORY-INDEPENDENT SEMANTICS

In this section we begin to present the formal preliminaries for our results. As standard in memory models, it is convenient to break the semantics into: a *program* semantics (a.k.a. thread subsystem) and a *memory* semantics. Next, we focus on the program part, presenting syntax (§3.1), semantics (§3.2), and the synchronization with a (parametric) memory system (§3.3).

²We generally require libraries and clients to operate on different “address spaces” and never access the same shared variable. Lifting this simplifying assumption (e.g., as was done for SC in [Gotsman and Yang 2013]) is beyond our current scope.

3.1 Program Syntax

We employ a simple programming language, which supports the distinction between clients and libraries variable spaces.³ We use the following domains (and metavariables ranging over them):

(variable) spaces	$X, Y \in \text{Space} \triangleq \{X, Y, \dots\}$	variables	$x, y \in \text{Loc} \triangleq \{x, y, \dots\}$
thread identifiers	$\tau, \pi \in \text{Tid} = \{T_1, T_2, \dots, T_N\}$	registers	$r \in \text{Reg} = \{a, b, \dots\}$
read modes	$o_R \in \text{Mod}_R \triangleq \{\text{na}, \text{rlx}, \text{pacq}, \text{acq}\}$	values	$v \in \text{Val} \triangleq \{0, 1, 2, \dots\}$
write modes	$o_W \in \text{Mod}_W \triangleq \{\text{na}, \text{rlx}, \text{prel}, \text{rel}\}$	method names	$f \in F \quad \text{main} \notin F$

In particular, as we will see below, every memory access is to a particular variable in a particular variable space (like address spaces in operating systems); registers are thread-local variables; and the name `main` is reserved for non-library operations. Memory instructions have access modes (a.k.a. memory orderings), which determine their “strength” (`na` for non-atomics, `rlx` for relaxed, `prel/pacq` for partial release/acquire, and `rel/acq` for release/acquire). We also assume a relation \sqsubset that orders modes according to their strength (following the left-to-right enumeration of the sets Mod_R and Mod_W above).

The language provides the following constructs, inspired by C11 atomics:

- Expressions are constructed with arithmetic and boolean operations over registers and values. We use e to range over expressions, and leave the exact expression grammar parametric.
- Thread local instructions (which do not interact with the shared-memory system):
 - Assignments of the form $r := e$, used for storing an expression e in a register r .
 - Conditionals of the form `if e goto $n_1 \mid \dots \mid n_m$` (where $n_1, \dots, n_m \in \mathbb{N}$), used to non-deterministically jump to some program counter among $\{n_1, \dots, n_m\}$ when e evaluates to non-zero or, otherwise, skipping.
 - Additional local instructions (standard loop constructs, unconditional non-deterministic choice, etc.). Below, we will use such instructions in our examples, and their semantics should be clear.
- Shared-memory single accesses:
 - Write (a.k.a. store) instructions to memory of the form `store(X, x, e, o_W)` for storing into a shared variable x in space X the value that e evaluates to with access mode o_W .
 - Read (a.k.a. load) instructions from memory of the form $r := \text{load}(X, x, o_R)$ for loading the value from a shared variable x in space X into a register r with access mode o_R .
- Read-modify-write (RMW) instructions:
 - Fetch-and-add (FAA) instructions of the form $r := \text{FADD}(X, x, e, o_R, o_W)$ for atomically incrementing a variable x in space X by the value of e with read mode o_R and write mode o_W .
 - Compare-and-swap (CAS) instructions of the form $r := \text{CAS}(X, x, e_R, e_W, o_R, o_W, o_R^{\text{fail}})$. This instruction atomically loads the value from x in space X into r , compares it to the value e_R , and overwrites it by the value of e_W in case if the loaded value coincides with the value e_R . The load part will have mode o_R if comparison succeeds and o_R^{fail} otherwise; and the store part (if it happens) has mode o_W .
- Library interaction: `call(f)` for calling a method f and `return` for returning to the caller. For simplicity, we do not provide any argument passing mechanism and we will use the full register store for that matter. (If needed, each component may store the values it needs in the memory, and reload them later on.)

To construct programs we introduce three syntactic categories, each of which builds on the previous one:

³The partition to variable spaces is only needed to support the “partial release/acquire” access mode in library specifications, and can be completely ignored for specifications without such accesses.

- *Instruction sequences* represent the (sequential) implementation of each method (including `main`). Formally, an instruction sequence I is a function from a non-empty finite domain of the form $\{0, \dots, n\}$ (representing the possible program counters) to the set of instructions. We say that an instruction sequence is *flat* if it does not include `call($_$)` instructions.
- *Sequential programs* consist of a “main” method accompanied with implementations of every method $f \in F$. Formally, a sequential program sPr is a function assigning an instruction sequence to every $f \in \{\text{main}\} \cup F$. To avoid modeling a call stack and simplify the framework, we require that $sPr(f)$ is a flat instruction sequence for every $f \in F$.
- *Concurrent programs*, which we often call *programs*, are top-level parallel compositions of sequential programs, all accompanied by the same method implementations. Formally, a (concurrent) program Pr is a mapping assigning a sequential program to every $\tau \in \text{Tid}$, with $Pr(\tau)(f) = Pr(\pi)(f)$ for every $\tau, \pi \in \text{Tid}$ and $f \in F$.

In our examples, we often write instruction sequences as sequences of instructions delimited by “;”, and concurrent programs using ‘||’ between the main method of each thread. We also refer to the program threads as τ_1, τ_2, \dots following their left-to-right order in the program listing.

3.2 Program Semantics

We give semantics to the syntactic objects above using labeled transition systems.

Definition 3.1. A *labeled transition system* (LTS) is a tuple $A = \langle \Sigma, Q, q_0, T \rangle$, where Σ is a set of *transition labels*, Q is a set of *states*, $q_0 \in Q$ is the *initial state*, and $T \subseteq Q \times \Sigma \times Q$ is a set of *transitions*. We denote by $A.\Sigma$, $A.Q$, $A.q_0$, and $A.T$ the components of an LTS A . We write $q \xrightarrow{\sigma} q'$ to denote a transition $\langle q, \sigma, q' \rangle \in T$, and $\xrightarrow{\sigma}_A$ for the relation $\{\langle q, q' \rangle \mid q \xrightarrow{\sigma} q' \in A.T\}$, and \rightarrow_A for $\bigcup_{\sigma \in \Sigma} \xrightarrow{\sigma}_A$. For a sequence $t \in A.\Sigma^*$, we write \xrightarrow{t}_A for the composition $\xrightarrow{t(1)}_A ; \dots ; \xrightarrow{t(|t|)}_A$. A sequence $t \in A.\Sigma^*$ such that $A.q_0 \xrightarrow{t}_A q$ for some $q \in A.Q$ is called a *trace* of A . We denote by $\text{traces}(A)$ the set of all traces of A . A state $q \in A.Q$ is called *reachable* in A if $A.q_0 \xrightarrow{t}_A q$ for some $t \in \text{traces}(A)$. For a trace t and a set $\Theta \subseteq \Sigma$ of transition labels, we write $t|_{\Theta}$ for the longest subsequence of t over Θ .

Next, we define the LTSs induced by instruction sequences, sequential programs, and concurrent programs. We often identify the syntactic objects with the LTS they induce (e.g., when writing expressions like $sPr.Q$ for a sequential program sPr). The transition labels of these LTSs feature *action labels*, which represent the interactions that a program may have with the memory.

Definition 3.2. An *action label* l takes one of the following forms: a read $R(X, x, v_R, o_R)$, a write $W(X, x, v_W, o_W)$, a read-modify-write $\text{RMW}(X, x, v_R, v_W, o_R, o_W)$, a call $\text{CALL}(f, \phi)$, and a return $\text{RET}(\phi)$, where $X \in \text{Space}$, $x \in \text{Loc}$, $v_R, v_W \in \text{Val}$, $o_R \in \text{Mod}_R$, $o_W \in \text{Mod}_W$, $f \in F$, and $\phi : \text{Reg} \rightarrow \text{Val}$. We denote by Lab the set of all action labels. The functions typ , sp , loc , val_R , val_W , mod_R , mod_W , callee , and store respectively retrieve (when applicable) the type ($R/W/\dots$), space (X), variable (x), read value (v_R), written value (v_W), read mode (o_R), write mode (o_W), callee method name (f), and store (ϕ) of an action label.

Next, we define the LTS induced by an instruction sequence.

Definition 3.3. An *instruction sequence state* is a pair $\langle pc, \phi \rangle$, where $pc \in \mathbb{N}$, called *program counter*, stores the current instruction pointer inside the sequence, and $\phi : \text{Reg} \rightarrow \text{Val}$, called *local store*, records the values of the registers. Local stores are extended to apply on expressions in the standard way. The *LTS induced by an instruction sequence* I is an LTS over instruction sequence states with: $\text{Lab}_{\epsilon} \triangleq \text{Lab} \cup \{\epsilon\}$ as the set of transition labels (that is, the set of all action labels extended with ϵ for silent transitions); $\langle 0, \phi_{\text{init}} \rangle$ where $\phi_{\text{init}} \triangleq \lambda r. 0$ as the initial state; and the transitions as given in Fig. 1 (additional thread local instructions can be standardly added).

$$\begin{array}{c}
\frac{I(pc) = r := e \quad \phi' = \phi[r \mapsto \phi(e)]}{\langle pc, \phi \rangle \xrightarrow{\epsilon}_I \langle pc + 1, \phi' \rangle} \\
\\
\frac{I(pc) = \mathbf{if} \ e \ \mathbf{goto} \ n_1 \mid \dots \mid n_m \quad \begin{array}{l} \phi(e) \neq 0 \implies pc' \in \{n_1, \dots, n_m\} \\ \phi(e) = 0 \implies pc' = pc + 1 \end{array}}{\langle pc, \phi \rangle \xrightarrow{\epsilon}_I \langle pc', \phi \rangle} \\
\\
\frac{I(pc) = r := \mathbf{load}(X, x, o_R) \quad \begin{array}{l} l = R(X, x, v, o_R) \\ \phi' = \phi[r \mapsto v] \end{array}}{\langle pc, \phi \rangle \xrightarrow{l}_I \langle pc + 1, \phi' \rangle} \\
\\
\frac{I(pc) = \mathbf{store}(X, x, e, o_W) \quad l = W(X, x, \phi(e), o_W)}{\langle pc, \phi \rangle \xrightarrow{l}_I \langle pc + 1, \phi \rangle} \\
\\
\frac{I(pc) = r := \mathbf{FADD}(X, x, e, o_R, o_W) \quad \begin{array}{l} l = \mathbf{RMW}(X, x, v, v + \phi(e), o_R, o_W) \\ \phi' = \phi[r \mapsto v] \end{array}}{\langle pc, \phi \rangle \xrightarrow{l}_I \langle pc + 1, \phi' \rangle} \\
\\
\frac{I(pc) = r := \mathbf{CAS}(X, x, e_R, e_W, o_R, o_W, o_R^{\text{fail}}) \quad \begin{array}{l} v = \phi(e_R) \implies l = \mathbf{RMW}(X, x, v, \phi(e_W), o_R, o_W) \\ v \neq \phi(e_R) \implies l = R(X, x, v, o_R^{\text{fail}}) \\ \phi' = \phi[r \mapsto v] \end{array}}{\langle pc, \phi \rangle \xrightarrow{l}_I \langle pc + 1, \phi' \rangle}
\end{array}$$

Fig. 1. Transitions of LTS induced by an instruction sequence

Recall that program semantics is separate from memory semantics, which is why the read and RMW transitions in Fig. 1 can observe any value. It is only important that each transition that interacts with the memory announces itself in the transition label. The `call`(`_`) and `return` instructions are not handled at the level of instruction sequences, but receive special semantics at the level of sequential programs, as defined next.

Definition 3.4. A *sequential program state* is a tuple $q = \langle pc, \phi, pc_s, f \rangle$, where: $\langle pc, \phi \rangle$ is an instruction sequence state storing the state of the sequence currently running; $pc_s \in \mathbb{N} \cup \{\perp\}$, called *the stored program counter*, is used to remember the program position to jump to when the current instruction sequence returns ($pc_s = \perp$ means that the main method is currently running); and $f \in F \cup \{\mathbf{main}\}$, called *the active method*, tracks the method that is currently running. We denote by $q.pc$, $q.\phi$, $q.pc_s$, and $q.f$ the components of a sequential program state q .

Definition 3.5. The *LTS induced by a sequential program* sPr is given by:

- The set of transition labels is $\text{Lab}_\epsilon \times (F \cup \{\mathbf{main}\})$. The functions `lab` and `method` respectively retrieve the action label (or ϵ) and method name of a transition label. All functions on action labels (`typ`, `sp`, `loc`, ...) are lifted to sequential program transition labels in the obvious way.
- The states are sequential program states, as defined in Def. 3.4.
- The initial state is $\langle 0, \phi_{\text{init}}, \perp, \mathbf{main} \rangle$.
- The transitions are given by:

$$\frac{f \in \{\mathbf{main}\} \cup F \quad \langle pc, \phi \rangle \xrightarrow{l_\epsilon}_{sPr(f)} \langle pc', \phi' \rangle}{\langle pc, \phi, pc_s, f \rangle \xrightarrow{l_{\epsilon, f}}_{sPr} \langle pc', \phi', pc_s, f \rangle}$$

$$\frac{\begin{array}{l} sPr(\mathbf{main})(pc) = \mathbf{call}(f) \\ l = \mathbf{CALL}(f, \phi) \end{array}}{\langle pc, \phi, \perp, \mathbf{main} \rangle \xrightarrow{l_{\mathbf{main}}}_{sPr} \langle 0, \phi, pc + 1, f \rangle} \quad \frac{\begin{array}{l} sPr(f)(pc) = \mathbf{return} \\ l = \mathbf{RET}(\phi) \end{array}}{\langle pc, \phi, pc_s, f \rangle \xrightarrow{l_f}_{sPr} \langle pc_s, \phi, \perp, \mathbf{main} \rangle}$$

The first transition, which applies for any method (`main` or other), lifts the instruction-sequence transition to the level of sequential programs. The second transition passes control from the main method to some other method, jumping the program counter to the first instruction and storing the return point ($pc + 1$). Finally, the third transition passes control back using the stored return point. (We do not need to record a call stack since we assume that $sPr(f)$ is flat for every $f \in F$.)

Finally, the LTS induced by a concurrent program interleaves the thread transitions.

Definition 3.6. A (concurrent) program state \bar{p} is a mapping assigning a sequential program state to every $\tau \in \text{Tid}$. The LTS induced by a program Pr is an LTS over program states, with: $\text{ProgLab} \triangleq \text{Tid} \times \text{Lab}_\epsilon \times (\text{F} \cup \{\text{main}\})$ as the set of transition labels; $\bar{p}_{\text{init}} \triangleq \lambda\tau. \langle 0, \phi_{\text{init}}, \perp, \text{main} \rangle$ as the initial state; and the following transitions:

$$\frac{\bar{p}(\tau) \xrightarrow{l_{\epsilon,f}}_{Pr(\tau)} q'}{\bar{p} \xrightarrow{\tau, l_{\epsilon,f}}_{Pr} \bar{p}[\tau \mapsto q']}$$

Below, for a program transition label $\alpha \in \text{ProgLab}$, the functions `tid`, `lab`, and `method` respectively retrieve the thread identifier (τ), the action label (or ϵ) (l_ϵ), and the method name (f) of α . Functions on action labels (`typ`, `sp`, `loc`, ...) are lifted to program transition labels in the obvious way.

3.3 Synchronizing Programs and Memories

To give semantics to programs running under a particular memory model, we synchronize the transitions of a program Pr with a memory system. For now, we leave the memory system parametric, and assume it is represented by some LTS \mathcal{M} whose set of transition labels consists of non-silent program transition labels (elements of $\text{Tid} \times \text{Lab} \times (\text{F} \cup \{\text{main}\})$) as well as a (disjoint) set $\mathcal{M}.\Theta$ of internal memory actions (which we use later for non-deterministic propagation of knowledge between threads).

Definition 3.7. The composition of a program Pr and a memory system \mathcal{M} , denoted by $Pr \bowtie \mathcal{M}$, is the LTS whose transition labels are the elements of $\text{ProgLab} \cup \mathcal{M}.\Theta$; states are pairs $\langle \bar{p}, M \rangle \in Pr.Q \times \mathcal{M}.Q$; initial state is $\langle \bar{p}_{\text{init}}, \mathcal{M}.q_\emptyset \rangle$; and transitions are given by:

$$\frac{\alpha \in \text{Tid} \times \text{Lab} \times (\text{F} \cup \{\text{main}\}) \quad \bar{p} \xrightarrow{\alpha}_{Pr} \bar{p}' \quad M \xrightarrow{\alpha}_{\mathcal{M}} M'}{\langle \bar{p}, M \rangle \xrightarrow{\alpha}_{Pr \bowtie \mathcal{M}} \langle \bar{p}', M' \rangle} \quad \frac{\alpha \in \text{Tid} \times \{\epsilon\} \times (\text{F} \cup \{\text{main}\}) \quad \bar{p} \xrightarrow{\alpha}_{Pr} \bar{p}'}{\langle \bar{p}, M \rangle \xrightarrow{\alpha}_{Pr \bowtie \mathcal{M}} \langle \bar{p}', M \rangle} \quad \frac{\alpha \in \mathcal{M}.\Theta \quad M \xrightarrow{\alpha}_{\mathcal{M}} M'}{\langle \bar{p}, M \rangle \xrightarrow{\alpha}_{Pr \bowtie \mathcal{M}} \langle \bar{p}, M' \rangle}$$

The above transitions are “synchronized transitions” of Pr and \mathcal{M} , using the labels to decide what to synchronize on. Both the program and the memory take the same step for transition labels that are common to both LTSs, only the program steps for transition labels that are program internal (i.e., with $\text{lab}(\alpha) = \epsilon$) and only the memory steps for transition labels that are memory internal.

Example 3.8. The most well-known memory system is the one of sequential consistency, denoted here by SC. This memory system simply tracks the most recent value written to each variable, and has no internal transitions ($\text{SC}.\Theta = \emptyset$). Formally, it is defined by $\text{SC}.Q \triangleq (\text{Space} \times \text{Loc}) \rightarrow \text{Val}$, $\text{SC}.q_\emptyset \triangleq \lambda X, x. 0$, and \rightarrow_{SC} is given by:

$$\frac{l = \text{R}(X, x, v_R, _) \quad m(X, x) = v_R}{m \xrightarrow{\tau, l, f}_{\text{SC}} m} \quad \frac{l = \text{W}(X, x, v_W, _) \quad m' = m[(X, x) \mapsto v_W]}{m \xrightarrow{\tau, l, f}_{\text{SC}} m'} \quad \frac{l = \text{RMW}(X, x, v_R, v_W, _) \quad m(X, x) = v_R \quad m' = m[(X, x) \mapsto v_W]}{m \xrightarrow{\tau, l, f}_{\text{SC}} m'} \quad \frac{\text{typ}(l) \in \{\text{CALL}, \text{RET}\}}{m \xrightarrow{\tau, l, f}_{\text{SC}} m}$$

4 THE dRC11 MEMORY MODEL

In this section we introduce the weak memory model that we assume in this paper. This model, which we call dRC11, is a declarative (a.k.a. axiomatic) model forming an extension of (a fragment of) the RC11 model [Lahav et al. 2017] with specialized semantics for the novel `prel`/`pacq` accesses.

NOTATION 4.1 (RELATIONAL NOTATIONS). Given a (binary) relation R , $\text{dom}(R)$ and $\text{codom}(R)$ denote its domain and codomain, and R^2 , R^+ , and R^* denote its reflexive, transitive, and reflexive-transitive closures. The inverse of a relation R is denoted by R^{-1} , and the (left) composition of two

relations R_1 and R_2 is denoted by $R_1 ; R_2$. We denote by $[A]$ the identity relation on a set A . In particular, $[A] ; R ; [B] = R \cap (A \times B)$. When A is finite, we write $[a_1, \dots, a_n]$ instead of $[\{a_1, \dots, a_n\}]$.

We start by defining *execution graphs*. Their nodes, called *events*, represent memory accesses, and their directed edges are of different kinds: *program order* represents the order imposed by the program; *reads-from* mapping maps each read event to the write event it obtains its value from; and *modification order* (a.k.a. *coherence order*) provides a total order on the writes to every variable. The precise definitions are given next.

Definition 4.2. An event e is a tuple $\langle \tau, s, l, f \rangle$, where $\tau \in \text{Tid} \uplus \{\perp\}$, called the event's *thread identifier* (\perp is used for initialization events); $s \in \mathbb{N}$, called the event's *serial identifier*; $l \in \text{Lab}$, called the event's *label* (as defined in Def. 3.2), and $f \in F \cup \{\text{main}\}$, called the event's *method*. The functions `tid`, `sn`, `lab`, and `method` return the thread identifier (τ), identifier (s), action label (l), and method of an event (f). All functions on action labels (`typ`, `sp`, `loc`, ...) are lifted to events in the obvious way. We denote by E the set of all events, and define the following subsets:

$$\begin{aligned} R &\triangleq \{e \in E \mid \text{typ}(e) \in \{\text{R}, \text{RMW}\}\} & W &\triangleq \{e \in E \mid \text{typ}(e) \in \{\text{W}, \text{RMW}\}\} & \text{RMW} &\triangleq R \cap W \\ \text{CALL} &\triangleq \{e \in E \mid \text{typ}(e) = \text{CALL}\} & \text{RET} &\triangleq \{e \in E \mid \text{typ}(e) = \text{RET}\} & \text{CR} &\triangleq \text{CALL} \cup \text{RET} \end{aligned}$$

We employ subscripts and superscripts to restrict sets of events to certain properties, e.g., $W_X = \{w \in W \mid \text{sp}(w) = X\}$, $R_{X,x}^{\text{pacq}} = \{r \in R \mid \text{sp}(r) = X \wedge \text{loc}(r) = x \wedge \text{mod}_W(r) \sqsupseteq \text{pacq}\}$, $\text{CALL}_F = \{e \in \text{CALL} \mid \text{callee}(e) \in F\}$, $\text{RET}_F = \{e \in \text{RET} \mid \text{method}(e) \in F\}$, $\text{CR}_F = \text{CALL}_F \cup \text{RET}_F$, and $E^\tau = \{e \in E \mid \text{tid}(e) = \tau\}$ for any $E \subseteq E$. The set `Init` of *initialization events* is given by:

$$\text{Init} \triangleq \{\langle \perp, 0, W(X, x, 0, \text{rlx}), \text{main} \rangle \mid X \in \text{Space}, x \in \text{Loc}\}.$$

Definition 4.3. An *execution graph* G is a tuple $\langle E, po, rf, mo \rangle$, where:

- E is a finite set of events, such that $\text{Init} \subseteq E$ and $\text{tid}(e) \neq \perp$ for every $e \in E \setminus \text{Init}$.
- po is a *program order* for E , that is: $po = (\bigsqcup_{\tau \in \text{Tid}} po_\tau) \uplus (\text{Init} \times (E \setminus \text{Init}))$, for some relations po_τ , such that each po_τ is a strict total order on E^τ .
- rf is a *reads-from relation* for E , that is a relation on E satisfying:
 - If $\langle w, r \rangle \in rf$, then $w \in W$ and $r \in R$.
 - If $\langle w, r \rangle \in rf$, then $\text{sp}(w) = \text{sp}(r)$, $\text{loc}(w) = \text{loc}(r)$, and $\text{val}_W(w) = \text{val}_R(r)$.
 - $w_1 = w_2$ whenever $\langle w_1, r \rangle, \langle w_2, r \rangle \in rf$ (each read reads from at most one write).
 - $E \cap R \subseteq \text{codom}(rf)$ (each read reads from some write).
- mo is a *modification order* for E , that is: $mo = \bigsqcup_{X \in \text{Space}, x \in \text{Loc}} mo_{X,x}$, for some relations $mo_{X,x}$, such that each $mo_{X,x}$ is a strict *total* order on $E \cap W_{X,x}$.

We denote the components of G by $G.E$, $G.po$, $G.rf$, and $G.mo$. For any set $E' \subseteq E$, we write $G.E'$ for $G.E \cap E'$ (e.g., $G.W = G.E \cap W$).

To formally associate execution graphs with programs we use a memory system called FG (for “free graphs”), whose states are execution graphs. This system allows *all* possible program transitions, while recording them in its state, which is the current execution graph with (almost) arbitrary reads-from and modification order relations.

Definition 4.4. The memory system FG is the LTS whose transition labels are program transition labels (i.e., there are not any internal memory actions); states are execution graphs; initial state is G_{Init} , defined by $G_{\text{Init}}.E \triangleq \text{Init}$ and $G_{\text{Init}}.X \triangleq \emptyset$ for every other component of G ; and transitions are given in Fig. 2.

The transitions of FG are based on helper notations used to extend an execution graph G with a fresh event e at the end of the executing thread. For memory accesses, it requires to pick a write

$$\begin{array}{c}
\text{WRITE/READ/RMW} \\
\text{typ}(l) \in \{W, R, \text{RMW}\} \quad e = \text{NextEvent}(G.E, \tau, l, f) \\
w \in G.W_{\text{sp}(e), \text{loc}(e)} \quad e \in R \implies \text{val}_W(w) = \text{val}_R(e) \\
G' = \text{Add}(G, e, w) \\
\hline
G \xrightarrow{\tau, l, f}_{\text{FG}} G'
\end{array}
\qquad
\begin{array}{c}
\text{CALL/RETURN} \\
\text{typ}(l) \in \{\text{CALL}, \text{RET}\} \\
e = \text{NextEvent}(G.E, \tau, l, f) \\
G' = \text{Add}(G, e) \\
\hline
G \xrightarrow{\tau, l, f}_{\text{FG}} G'
\end{array}$$

Fig. 2. Transitions of FG system for generating all candidate execution graphs

event w in G , called the *write-predecessor* of e , that is: (1) the **rf**-source of e if e is a read; (2) the **mo**-immediate predecessor of e if e is a write; and (3) both the **rf**-source and the **mo**-immediate predecessor of e if e is an RMW. For that matter, we employ the following notations:

NOTATION 4.5. Given a set E of events, $\tau \in \text{Tid}$, $l \in \text{Lab}$, and $f \in F$, $\text{NextEvent}(E, \tau, l, f)$ denotes the event with thread identifier τ , label l , method f , and a minimal fresh serial identifier w.r.t. E , that is: $\text{NextEvent}(E, \tau, l, f) \triangleq \langle \tau, s, l, f \rangle$, where $s = \min\{n \in \mathbb{N} \mid \langle \tau, n, l, f \rangle \notin E\}$.

NOTATION 4.6. For an execution graph G and events e and w , $\text{Add}(G, e, w)$ denotes the tuple $\langle E', \text{po}', \text{rf}', \text{mo}' \rangle$, where:

$$\begin{array}{l}
E' = G.E \cup \{e\} \\
\text{rf}' = \begin{cases} G.\text{rf} \cup \{\langle w, e \rangle\} & e \in R \\ G.\text{rf} & \text{otherwise} \end{cases} \\
\text{po}' \triangleq G.\text{po} \cup ((G.E^{\text{tid}(e)} \cup \text{Init}) \times \{e\}) \\
\text{mo}' = \begin{cases} G.\text{mo} \cup \text{dom}(G.\text{mo}^?; [w]) \times \{e\} \\ \cup \{e\} \times \text{codom}([w]; G.\text{mo}) & e \in W \\ G.\text{mo} & \text{otherwise} \end{cases}
\end{array}$$

Similarly, for an execution graph G and event e , $\text{Add}(G, e)$ denotes the tuple $\langle E', \text{po}', G.\text{rf}, G.\text{mo} \rangle$, where E' and po' are defined as above.

In the sequel, it will be useful to note that all execution graphs generated by FG when synchronized with a given program satisfy the following well-formedness property:

Definition 4.7. An execution graph G is *well-formed* if the following hold for every $f \in F$:

- $\{\{e \in E \mid \text{method}(e) \neq f\}; G.\text{po}\}; \{\{e \in E \mid \text{method}(e) = f\}\} \subseteq G.\text{po}^?; [\text{CALL}_{\{f\}}]; G.\text{po}$.
- $\{\{e \in E \mid \text{method}(e) = f\}; G.\text{po}\}; \{\{e \in E \mid \text{method}(e) \neq f\}\} \subseteq G.\text{po}^?; [\text{RET}_{\{f\}}]; G.\text{po}$.

PROPOSITION 4.8. If $\langle \bar{p}, G \rangle$ is reachable in $\text{Pr} \bowtie \text{FG}$ for some program Pr , then G is well-formed.

Now, to filter for *consistent* graphs among all candidate execution graphs generated by FG for a given program, we define several derived relations (some parametrized by $X \in \text{Space}$):

$$\begin{array}{l}
G.\text{fr} \triangleq (G.\text{rf}^{-1}; G.\text{mo}) \setminus [E] \quad (\text{from-read, a.k.a. read-before}) \\
G.\text{sw}_{\text{base}} \triangleq [W^{\text{rel}}]; G.\text{rf}^+; [R^{\text{acq}}] \quad (\text{global synchronization}) \\
G.\text{sw}_X \triangleq [W_X^{\text{prel}}]; G.\text{rf}^+; [R_X^{\text{pacq}}] \quad (\text{per-space synchronization}) \\
G.\text{hb}_X \triangleq (G.\text{po} \cup G.\text{sw}_{\text{base}} \cup G.\text{sw}_X)^+ \quad (\text{per-space happens-before})
\end{array}$$

The **fr** relation is standard in weak memory models [Alglave et al. 2014], relating every read (or RMW) to subsequent writes (or RMWs) as dictated by **rf** and **mo** (every write w that is **mo**-after the **rf**-source of a read r is **fr**-after r). The **sw_{base}** and **sw_X** relations formally capture synchronization: **sw_{base}** is for “global” synchronization which is formed by reads-from edges between **rel** and **acq** accesses (just like in C11 [Lahav et al. 2017]), whereas **sw_X** is for “space-internal” synchronization that affects only accesses to space X and is formed by reads-from edges *in space* X between **prel**




Fig. 3. Illustration of forbidden patterns in dRC11

and `pacq`. The use of rf^+ (rather than just rf) is for supporting *release sequences* as in C11, which ensures the synchronization between a release write w and an acquire read r also if there is chain of reads-from edges between them ($\langle w, u_1 \rangle, \langle u_1, u_2 \rangle, \dots, \langle u_{n-1}, u_n \rangle, \langle u_n, r \rangle \in rf$ where u_1, \dots, u_n are RMWs, which can be relaxed).⁴ Finally, paths composed of the program order (po) and the sw_{base} and sw_X relations form the per-space *happens-before* relation, which again refines the C11 happens-before relation (defined as $(G.po \cup G.sw_{base})^+$).

Using the above definitions, the consistency of an execution graph is defined as follows.

Definition 4.9. An execution graph G is *dRC11-consistent* if the following hold:


- For every $X \in \text{Space}$, $[R_X]; G.fr; G.rf^?; G.hb_X$ is irreflexive. (READ COHERENCE)
- For every $X \in \text{Space}$, $[W_X]; G.mo; G.rf^?; G.hb_X$ is irreflexive. (WRITE COHERENCE)
- $G.mo; G.rf$ is irreflexive. (RMW1)
- $G.fr; G.mo$ is irreflexive. (RMW2)
- $G.po \cup G.rf$ is acyclic. (PO-RF)

These constraints, depicted in Fig. 3, are variants of the ones in (R)C11, where the only essential difference is in the coherence constraints that here use $G.hb_X$ for restricting accesses to X .⁵ The **RMW1** and **RMW2** constraints are needed for ensuring the right behavior of RMWs including their atomicity. The **PO-RF** constraint is an addition of RC11 on top of C11, which is a conservative solution to the “out-of-thin-air” problem that arises if $po \cup rf$ -cycles are allowed [Batty et al. 2015; Kang et al. 2017].

The next example demonstrates the role of the coherence constraints:

Example 4.10. Consider the following standard “message-passing” litmus test (parametric in the space Y and the access modes o_W, o_R):

```
store(X, x, 1, rlx); || a := load(Y, y, o_R); // 1
store(Y, y, 1, o_W); || b := load(X, x, rlx); // 0
```



The annotated behavior is disallowed only if the synchronization on y (in space Y) is visible for the read of x (in space X). This will be the case only if (1) $o_W = \text{rel}$ and $o_R = \text{acq}$; or (2) $o_W \sqsupseteq \text{prel}$, $o_R \sqsupseteq \text{pacq}$, and $Y = X$ (i.e., the four accesses are to the same space). (In particular, if $Y \neq X$, then only $o_W = \text{rel}$ and $o_R = \text{acq}$ would forbid the annotated behavior.) To see how this follows from

⁴Another component of C11’s release sequence, which allows forming synchronization using a relaxed write *po*-after a release write to the same location, was omitted from the standard in C++20. (See https://en.cppreference.com/w/cpp/atomic/memory_order [Accessed July 2022].)

⁵There are some presentational differences w.r.t. the model in [Lahav et al. 2017]: (1) RC11 uses two events related by an *rmw*-edge to represent RMWs, so **RMW1** is not needed and **RMW2** has a different formulation (called **ATOMICITY** in [Lahav et al. 2017]); and (2) RC11 uses the “extended coherence order” (**eco**) which allows it to merge both coherence constraints.

the coherence constraints, note that these conditions are needed to ensure $\langle b, c \rangle \in \text{sw}_{\text{base}} \cup \text{sw}_X$, which in turn implies $\langle a, d \rangle \in \text{hb}_X$. Now, since $\langle \text{init}, a \rangle \in \text{po} \subseteq \text{hb}_X$, **WRITE COHERENCE** ensures that $\langle \text{init}, a \rangle \in \text{mo}$, and so $\langle d, a \rangle \in \text{fr}$. In turn, **READ COHERENCE** disallows hb_X from a to d .

By connecting the definition of candidate execution graphs for a given program (using the FG system) and dRC11-consistency, we define the allowed program behaviors under dRC11 (more precisely, possible reachable program states).

Definition 4.11. A program state \bar{p} is *reachable for a program* Pr under dRC11 if $\langle \bar{p}, G \rangle$ is reachable in $Pr \bowtie \text{FG}$ for some dRC11-consistent execution graph G .

Finally, again following (R)C11, data races on non-atomics are considered as programming errors (so non-atomic accesses can be heavily optimized by the hardware and the compiler). Accordingly, we define racy execution graphs and racy programs.

Definition 4.12. Two events e_1 and e_2 form a *race* in an execution graph G , if $e_1, e_2 \in G.W \cup G.R$, $e_1 \neq e_2$, $\text{sp}(e_1) = \text{sp}(e_2)$, $\text{loc}(e_1) = \text{loc}(e_2)$, $|\{e_1, e_2\} \cap W| \geq 1$, $|\{e_1, e_2\} \cap E^{\text{na}}| \geq 1$, and $\langle e_1, e_2 \rangle \notin G.\text{hb}_{\text{sp}(e_1)} \cup G.\text{hb}_{\text{sp}(e_1)}^{-1}$. An execution graph G is *racy* if some two events form a race in G .

Definition 4.13. A program Pr is *racy* under dRC11 if $\langle \bar{p}, G \rangle$ is reachable in $Pr \bowtie \text{FG}$ for some program state \bar{p} and racy dRC11-consistent execution graph G .

5 THE OPERATIONAL MEMORY SYSTEM: pRC11

In this section we introduce an operational version of the dRC11 memory model, called pRC11, which exposes knowledge propagation steps as internal memory steps, and is particularly suitable for library abstraction.

Since dRC11-consistency is *prefix-closed* w.r.t. $\text{po} \cup \text{rf}$ [Kokologiannakis et al. 2017], dRC11 can be easily made operational by adapting the LTS FG (Def. 4.4) to require dRC11-consistency after each step of the execution graph generation (rather than one time at the end). However, following §2.1, we opt for a more elaborate model with explicit point-to-point *propagation transitions* marking the steps in which some event of thread τ becomes visible to another thread π . In particular, this allows us to observe the propagation of the call/return events in memory traces, which is essential in our definition of the library correctness condition in §7. We expose propagation transitions in traces of pRC11 as internal memory steps labeled with *propagation labels* as defined next.

Definition 5.1. A *propagation label* is a triple, denoted by $p = \text{EP}(e, \tau, X)$, where $e \in E$ (propagated event), $\tau \in \text{Tid} \setminus \{\text{tid}(e)\}$ (destination thread identifier), and $X \in \text{Space}$ (destination space). We use E , ptid and psp to retrieve the components (e , τ , and X , respectively) of a propagation label p . All functions on events and action labels (tid , $\text{typ} \dots$) are lifted to propagation labels in the obvious way. We denote by PLab the set of all propagation labels.

Then, the pRC11 memory system is defined as follows.

Definition 5.2. The pRC11 memory system is an LTS whose set of transition labels is $\text{ProgLab} \cup \text{PLab}$ (i.e., $\text{pRC11}.\Theta = \text{PLab}$); states are pairs of the form $M = \langle G, K \rangle$, where G is an execution graph and K is a *knowledge mapping* for G , that is a function in $\text{Tid} \rightarrow \text{Space} \rightarrow \mathcal{P}(G.E)$; initial state is $M_{\text{init}} \triangleq \langle G_{\text{init}}, K_{\text{init}} \rangle$, where $K_{\text{init}} \triangleq \lambda \tau. \lambda X. \text{Init}(G_{\text{init}}$ is defined in Def. 4.4); and transitions are as given in Fig. 4.

In addition to the current execution graph G , pRC11's states record a *knowledge mapping* K that records for each thread τ and space X all events that have already propagated to τ for space X (as well as all events associated with actions executed by τ itself). We refer to the set $K(\tau)(X)$ as the X -knowledge of thread τ .

$$\begin{array}{c}
\text{WRITE/READ/RMW} \\
\text{typ}(l) \in \{W, R, \text{RMW}\} \quad e = \text{NextEvent}(G.E, \tau, l, f) \quad X = \text{sp}(e) \\
w \in G.W_{X, \text{loc}(e)} \quad e \in R \implies \text{val}_W(w) = \text{val}_R(e) \\
w \notin \text{dom}(G.\text{mo}; G.\text{rf}^?; [K(\tau)(X)]) \quad e \in W \implies w \notin \text{dom}(G.\text{rf}; [\text{RMW}]) \\
e \in R^{\text{pacq}} \implies \text{dom}(G.\text{rf}^*; [w]) \subseteq K(\tau)(X) \quad e \in R^{\text{acq}} \implies \forall Y. \text{dom}(G.\text{rf}^*; [w]) \subseteq K(\tau)(Y) \\
G' = \text{Add}(G, e, w) \quad k' = \lambda Y. K(\tau)(Y) \cup \{e\} \quad K' = K[\tau \mapsto k'] \\
\hline
\langle G, K \rangle \xrightarrow{\tau, l, f}_{\text{pRC11}} \langle G', K' \rangle \\
\\
\text{CALL/RETURN} \\
\text{typ}(l) \in \{\text{CALL}, \text{RET}\} \quad e = \text{NextEvent}(G.E, \tau, l, f) \\
G' = \text{Add}(G, e) \quad k' = \lambda Y. K(\tau)(Y) \cup \{e\} \quad K' = K[\tau \mapsto k'] \\
\hline
\langle G, K \rangle \xrightarrow{\tau, l, f}_{\text{pRC11}} \langle G', K' \rangle \\
\\
\text{PROPAGATE} \\
e \in G.E \setminus K(\tau)(X) \\
e \in W_X^{\text{prel}} \cup W^{\text{rel}} \cup \text{CR} \implies (E_X \cup \text{CR}) \cap \text{dom}(G.\text{hb}_X; [e]) \subseteq K(\tau)(X) \\
k' = K(\tau)[X \mapsto K(\tau)(X) \cup \{e\}] \quad K' = K[\tau \mapsto k'] \\
\hline
\langle G, K \rangle \xrightarrow{\text{EP}(e, \tau, X)}_{\text{pRC11}} \langle G, K' \rangle
\end{array}$$

Fig. 4. Transitions of pRC11.

The WRITE/READ/RMW transition in Fig. 4 executes a memory access by thread τ , by adding a corresponding event e to the current execution graph while imposing certain conditions on w , the write-predecessor of e . Intuitively, the main imposed condition, $w \notin \text{dom}(G.\text{mo}; G.\text{rf}^?; [K(\tau)(X)])$, requires that w is not overwritten by any other write that τ is already aware of for the space of e . More precisely, if e is an access in space X , then τ should not have in its X -knowledge any write that is **mo**-later than w or any read that reads from a write that is **mo**-later than w . In addition:

- If e is a write (or RMW), then w should not be already read by an RMW. This condition is needed to ensure the atomicity of RMWs (corresponds to **rmw2** in Def. 4.9).
- If e is a read (or RMW) with **acq** or **pacq** mode, then w (the write that e reads from) has to be already present in the thread’s knowledge: in its X -knowledge if e is **pacq**, or in Y -knowledge for all Y if e is **acq**. (Moreover, if w is an RMW, to account for release sequences, this should hold not only for w , but also for every event on the **rf**-chain entering w .)

Finally, every thread certainly knows about its own actions, so in addition to extending G , this step also extends the knowledge of τ by adding the event e to all spaces.

The CALL/RETURN transition is simple: it adds a corresponding event e to the current execution graph and extends the knowledge of the executing thread to include e .

The PROPAGATE transition is a non-deterministic internal memory step that extends the threads’ knowledge. It picks some event e that is not in τ ’s X -knowledge and adds e to $K(\tau)(X)$. When e is a **prel**-write to space X , a **rel**-write (to any space), or a call or return marker, then the propagation of e to the X -knowledge of thread τ can be done only after all accesses to X , as well as all call and return markers, that are $G.\text{hb}_X$ -before e have propagated to the X -knowledge of thread τ . What is *not* constrained is equally important: for instance, relaxed writes can propagate “out-of-order”.

Example 5.3. pRC11’s transitions are best understood via the “message-passing” litmus test presented in Example 4.10. Consider first the case that $o_W = \text{rel}$ and $o_R = \text{acq}$. Then, due to the constraints on acquire reads and release writes: (1) since the read of y is acquire, it has to read

from a write that is in T_2 's X -knowledge hold for every space X : and (2) since the write of y is release, this write can only propagate to the X -knowledge of T_2 (for every X) after the po-earlier write to x has propagated there. This means that if the read of y retrieves 1, then T_2 already has the write to x in its X -knowledge, so when later reading x it cannot read from the overwritten initial value. Similarly, if $o_W = \text{prel}$, $o_R = \text{pacq}$, and $Y = X$, then, (i) and (ii) apply for $X = X$, and the same reasoning holds. Nevertheless, in every other case, T_2 can read the overwritten initial value: because either (1) the *read* from y is too weak and it allows to read from an event that has not yet propagated to the thread's X -knowledge; or (2) the *write* to y is too weak and it can propagate to the thread's X -knowledge before the po-earlier write to X has propagated.

Example 5.4. Being equivalent to dRC11, pRC11 provides “per-location-SC” (a.k.a. coherence). As a concrete example, $a = 1 \wedge b = 0$ in the following program on the right is disallowed.

To see this in pRC11, note that read events are also included in the threads' knowledge. Concretely, after executing the first read, the second thread has its own read in its X -knowledge. Then, reading later from the (implicit) initialization write is forbidden by pRC11 since that write and the first read will be have $G.\text{mo}$; $G.\text{rf}$ between them.

```
store(X, x, 1, rlx); || a := load(X, x, rlx);
                    || b := load(X, x, rlx);
```

Remark 3. It is also instructive to consider a simplified fragment with only one variable space and without na/rlx accesses. In this fragment, we can simply talk about the knowledge of each thread (instead of the X -knowledge for each X), and observe that: reads by thread τ can only read from writes that thread τ already knows about; and the propagation order respects **hb** (in particular, it follows the program order). This actually makes the reads deterministic: they have to read-from the **mo**-latest write among all known writes to the relevant variable. The resulting model is similar to message passing models for causal consistency [Beillahi et al. 2021].

Based on the pRC11 memory system, we define reachable program states and racy programs.

Definition 5.5. A program state \bar{p} is *reachable* for a program Pr under pRC11 if $\langle \bar{p}, \langle G, K \rangle \rangle$ is reachable in $Pr \triangleright \text{pRC11}$ for some $\langle G, K \rangle \in \text{pRC11.Q}$. A program Pr is *racy* under pRC11 if $\langle \bar{p}, \langle G, K \rangle \rangle$ is reachable in $Pr \triangleright \text{pRC11}$ for some program state \bar{p} and $\langle G, K \rangle \in \text{pRC11.Q}$ such that G is a racy execution graph (see Def. 4.12).

5.1 Equivalence of pRC11 and dRC11

We state our equivalence result between pRC11 and dRC11, relating Def. 5.5 to Definitions 4.11 and 4.13.

THEOREM 5.6 (EQUIVALENCE OF THE MODELS). *A program state \bar{p} is reachable for a program Pr under dRC11 iff it is reachable for Pr under pRC11. Furthermore, Pr is racy under dRC11 iff it is racy under pRC11.*

Next, we describe the main steps in the proof (the full proof is given in [Singh and Lahav 2022]). First, for the right-to-left directions, it suffices to establish the following invariants on reachable pRC11-states.

Definition 5.7. A knowledge mapping K is *well-formed* for an execution graph G if the following hold for every $\tau \in \text{Tid}$ and $X \in \text{Space}$:

- (1) $G.E^\tau \subseteq K(\tau)(X)$.
- (2) $(E_X \cup \text{CR}) \cap \text{dom}(G.\text{hb}_X; [(W_X^{\text{prel}} \cup W^{\text{rel}}) \cap K(\tau)(X)]) \subseteq K(\tau)(X)$.
- (3) $(E_X \cup \text{CR}) \cap \text{dom}(G.\text{hb}_X; [E^\tau]) \subseteq K(\tau)(X)$.

PROPOSITION 5.8. *If $\langle G, K \rangle$ is reachable in pRC11, then K is well-formed for G .*

LEMMA 5.9. *If $\langle G, K \rangle$ is reachable in pRC11, then G is dRC11-consistent.*

For the converse, one starts with an dRC11-consistent execution graph G , and has to traverse its events (following the program order, so it can be synchronized with the program), and intersperse propagation actions to make it a valid trace of pRC11. To construct this traversal, we define the following relations, where $P \triangleq \{p \in \text{PLab} \mid E(p) \in G.E\}$:

$$R_{\text{prop}} \triangleq \{\langle e, p \rangle \in G.E \times P \mid E(p) = e\}$$

$$T \triangleq \{\langle p, p' \rangle \in P \times P \mid \text{ptid}(p) = \text{ptid}(p') \wedge \text{psp}(p) = \text{psp}(p') \wedge \langle E(p), E(p') \rangle \in [E_{\text{psp}(p)} \cup \text{CR}] ; G.\text{hb}_{\text{psp}(p)} ; [W_{\text{psp}(p)}^{\text{prel}} \cup W^{\text{rel}} \cup \text{CR}]\}$$

$$R_{\text{rfp}} \triangleq \{\langle p, e \rangle \in P \times G.E \mid \langle E(p), e \rangle \in G.\text{rf}^+ \wedge \text{ptid}(p) = \text{tid}(e) \wedge e \in R_{\text{psp}(p)}^{\text{pacq}}\}$$

$$R_{\text{rf}} \triangleq \{\langle p, e \rangle \in P \times G.E \mid \langle E(p), e \rangle \in G.\text{rf}^+ \wedge \text{ptid}(p) = \text{tid}(e) \wedge e \in R^{\text{acq}}\}$$

$$R_{\text{fr}} \triangleq \{\langle r, p \rangle \in G.E \times P \mid \langle r, E(p) \rangle \in G.\text{fr} ; G.\text{rf}^2 \wedge \text{ptid}(p) = \text{tid}(r) \wedge \text{psp}(p) = \text{sp}(r)\}$$

$$R_{\text{mo}} \triangleq \{\langle w, p \rangle \in G.E \times P \mid \langle w, E(p) \rangle \in G.\text{mo} ; G.\text{rf}^2 \wedge \text{ptid}(p) = \text{tid}(w) \wedge \text{psp}(p) = \text{sp}(w)\}$$

$$R \triangleq G.\text{po} \cup G.\text{rf} \cup R_{\text{prop}} \cup T \cup R_{\text{rfp}} \cup R_{\text{rf}} \cup R_{\text{fr}} \cup R_{\text{mo}}$$

Then, the proof proceeds by showing that R is acyclic, and that every total order of $G.E \cup P$ extending R induces a trace of pRC11. Indeed, the relations above are in one-to-one correspondence with the conditions in the steps of pRC11.

6 LIBRARIES AND THEIR CLIENTS

In this section we list the necessary definitions for the library abstraction theorem, and state the key properties that are used in its proof.

Client-library composition. A library L is a function mapping a set $\text{dom}(L) \subseteq F$ of method names to flat instruction sequences representing the method bodies. We only consider the case where libraries and their clients never access the same variable space. To formally define this syntactic restriction, we use the following notations for spaces used by libraries and their clients:

- $\text{Space}(I)$ denotes the set of variable spaces mentioned in an instruction sequence I .
- For a library L , $\text{Space}(L) \triangleq \bigcup_{f \in \text{dom}(L)} \text{Space}(L(f))$.
- For a program Pr and a set $F \subseteq F$, $\text{Space}(Pr \setminus F) \triangleq \bigcup_{\tau \in \text{Tid}, f \in (F \cup \{\text{main}\}) \setminus F} \text{Space}(Pr(\tau)(f))$.

Then, client-library composition is defined as follows.

Definition 6.1. A library L is *safe* for a program Pr if $\text{Space}(L) \cap \text{Space}(Pr \setminus \text{dom}(L)) = \emptyset$. When L is safe for Pr , we write $Pr[L]$ for the program obtained from Pr by setting $Pr(\tau)(f) = L(f)$ for every $\tau \in \text{Tid}$ and $f \in \text{dom}(L)$.

Next, we observe that the safety condition above ensures that execution graphs generated by $Pr[L]$ satisfy the following conditions relating the location spaces and the invoked methods of the graph events.

PROPOSITION 6.2. *Let L be a library that is safe for a program Pr . Suppose that $\langle \bar{p}, G \rangle$ is reachable in $Pr[L] \bowtie FG$. Then, the following hold for every $e \in G.R \cup G.W$:*

- If $\text{sp}(e) \in \text{Space}(L)$, then $\text{method}(e) \in \text{dom}(L)$.
- If $\text{sp}(e) \in \text{Space}(Pr \setminus \text{dom}(L))$, then $\text{method}(e) \notin \text{dom}(L)$.

Definition 6.3. A set $F \subseteq F$ is *encapsulated in an execution graph G* if for every $e_1, e_2 \in G.R \cup G.W$ with $\text{method}(e_1) \in F$ and $\text{method}(e_2) \notin F$, we have $\text{sp}(e_1) \neq \text{sp}(e_2)$.

From [Prop. 6.2](#), we obtain the following.

PROPOSITION 6.4. *Let L be a library that is safe for a program Pr . If $\langle \bar{p}, G \rangle$ is reachable in $Pr[L] \bowtie FG$, then $\text{dom}(L)$ is encapsulated in G .*

The notions of well-formedness and encapsulated set of methods are lifted to memory states in the obvious way:

Definition 6.5. Let $M = \langle G, K \rangle$ be a memory state.

- $M = \langle G, K \rangle$ is *well-formed* if G is well-formed ([Def. 4.7](#)) and K is well-formed for G ([Def. 5.7](#)).
- F is *encapsulated in M* if it is encapsulated in G ([Def. 6.3](#)).

Since every reachable state in $Pr[L] \bowtie \text{pRC11}$ is also reachable in $Pr[L] \bowtie FG$, the following is an immediate consequence of [Propositions 4.8, 5.8](#) and [6.4](#).

PROPOSITION 6.6. *Let L be a library that is safe for a program Pr . If $\langle \bar{p}, M \rangle$ is reachable in $Pr[L] \bowtie \text{pRC11}$, then M is well-formed and $\text{dom}(L)$ is encapsulated in M .*

Client-library program states. We define the composition of a program state \bar{p}_{cl} representing a client state and a library state \bar{p}_{lib} representing a library state as follows.

Definition 6.7. The *composition of two program states \bar{p}_{cl} and \bar{p}_{lib} w.r.t. a set $F \subseteq \mathbb{F}$* , denoted by $\bar{p}_{\text{cl}}[F \mapsto \bar{p}_{\text{lib}}]$, is given by:

$$\bar{p}_{\text{cl}}[F \mapsto \bar{p}_{\text{lib}}] \triangleq \lambda \tau. \begin{cases} \langle \bar{p}_{\text{lib}}(\tau) \cdot \text{pc}, \bar{p}_{\text{lib}}(\tau) \cdot \phi, \bar{p}_{\text{cl}}(\tau) \cdot \text{pc}_s, \bar{p}_{\text{cl}}(\tau) \cdot f \rangle & \bar{p}_{\text{cl}}(\tau) \cdot f \in F \\ \bar{p}_{\text{cl}}(\tau) & \text{otherwise} \end{cases}$$

This definition uses \bar{p}_{cl} for threads that are not currently inside a method in F , and \bar{p}_{lib} , but with the stored program counter and active method of \bar{p}_{cl} , for threads that are inside a method in F .

Histories. Histories record the interactions between libraries and clients. Formally, a *history h* of a library L is a sequence of transition labels representing a call to a method of L , a return from a method of L , or propagation of these call and return events. To define the history induced by a program, we employ the following notations (for every $F \subseteq \mathbb{F}$):

$$\begin{aligned} \text{Call}_F &\triangleq \{\alpha \in \text{ProgLab} \mid \text{typ}(\alpha) = \text{CALL} \wedge \text{callee}(\alpha) \in F\} & \text{CP}_F &\triangleq \{p \in \text{PLab} \mid E(p) \in \text{CALL}_F\} \\ \text{Ret}_F &\triangleq \{\alpha \in \text{ProgLab} \mid \text{typ}(\alpha) = \text{RET} \wedge \text{method}(\alpha) \in F\} & \text{RP}_F &\triangleq \{p \in \text{PLab} \mid E(p) \in \text{RET}_F\} \\ \text{HLab}_F &\triangleq \text{Call}_F \cup \text{Ret}_F \cup \text{CP}_F \cup \text{RP}_F \end{aligned}$$

Definition 6.8. Let $F \subseteq \mathbb{F}$. The *F -history* induced by a trace t of $Pr \bowtie \text{pRC11}$ for some program Pr , denoted by $H_F(t)$, is given by $H_F(t) \triangleq t|_{\text{HLab}_F}$. This notion is extended to sets of traces in the obvious way. The set of *F -histories of a program Pr* , denoted by $H_F(Pr)$, is given by $H_F(Pr) \triangleq H_F(\text{traces}(Pr \bowtie \text{pRC11}))$.

Client-library trace restrictions. We extract library and client transitions from a given trace as follows.

Definition 6.9. For $F \subseteq \mathbb{F}$, the *F -restriction* and the *\bar{F} -restriction* of a trace t of $Pr \bowtie \text{pRC11}$, denoted by $t|_F$ and $t|_{\bar{F}}$ (respectively), are given by:

$$t|_F \triangleq t|_{\{\alpha \in \text{ProgLab} \cup \text{PLab} \mid \text{method}(\alpha) \in F\} \cup \text{Call}_F \cup \text{CP}_F} \quad t|_{\bar{F}} \triangleq t|_{\{\alpha \in \text{ProgLab} \cup \text{PLab} \mid \text{method}(\alpha) \notin F\} \cup \text{Ret}_F \cup \text{RP}_F}$$

Note that both the F -restriction and the \bar{F} -restriction of a trace t contain the F -history induced by t as a subsequence.

Restricting memory states. Similarly, we will need to restrict memory states to client/library, as defined next.

Definition 6.10. The restriction of an execution graph G w.r.t. a set $E \subseteq \mathbb{E}$, denoted by $G|_E$, is defined by: $G|_{E.E} \triangleq E \cup \text{Init}$ and $G|_{E.X} \triangleq [G|_{E.E}] ; G.X ; [G|_{E.E}]$ for every other component (i.e., $X \in \{\text{po}, \text{rf}, \text{mo}\}$).

Definition 6.11. The restriction of a memory state $M = \langle G, K \rangle$ w.r.t. a set $E \subseteq \mathbb{E}$, denoted by $M|_E$, is given by $M|_E \triangleq \langle G|_E, K|_E \rangle$, where $K|_E \triangleq \lambda \tau. \lambda X. K(\tau)(X) \cap E$.

Definition 6.12. Let $F \subseteq \mathbb{F}$. The F -events and the \bar{F} -events denoted by E_F and $E_{\bar{F}}$ (respectively), are given by:

$$E_F \triangleq \{e \in \mathbb{E} \mid \text{method}(e) \in F\} \cup \text{CALL}_F \quad E_{\bar{F}} \triangleq \{e \in \mathbb{E} \mid \text{method}(e) \notin F\} \cup \text{RET}_F$$

The F -restriction and the \bar{F} -restriction of a memory state $M = \langle G, K \rangle$, denoted by $M|_F$ and $M|_{\bar{F}}$ (respectively), are given by $M|_F \triangleq M|_{E_F}$ and $M|_{\bar{F}} \triangleq M|_{E_{\bar{F}}}$.

Again, we note that both the F -restriction and the \bar{F} -restriction of a memory M contain the call events invoking methods in F and the return events that complete these invocations.

Restriction and merge properties. The following lemmas summarize the critical properties of the pRC11 memory system that allow us to compose memory traces of clients and libraries. In the proof of the abstraction theorem below we rely only on these properties of pRC11.

LEMMA 6.13 (RESTRICTION-1). *Suppose that $M \xrightarrow{\alpha}_{\text{pRC11}} M'$ and let $F \subseteq \mathbb{F}$ that is encapsulated in M' . Then, the following hold:*

- (1) *If $\text{method}(\alpha) \in F$ or $\alpha \in \text{HLab}_F$, then $M|_F \xrightarrow{\alpha}_{\text{pRC11}} M'|_F$.*
- (2) *If $\text{method}(\alpha) \notin F$ or $\alpha \in \text{HLab}_F$, then $M|_{\bar{F}} \xrightarrow{\alpha}_{\text{pRC11}} M'|_{\bar{F}}$.*

LEMMA 6.14 (RESTRICTION-2). *Suppose that $M \xrightarrow{\alpha}_{\text{pRC11}} M'$ and let $F \subseteq \mathbb{F}$. Then, the following hold:*

- (1) *If $\text{method}(\alpha) \notin F$ and $\alpha \notin \text{HLab}_F$, then $M|_F = M'|_F$.*
- (2) *If $\text{method}(\alpha) \in F$ and $\alpha \notin \text{HLab}_F$, then $M|_{\bar{F}} = M'|_{\bar{F}}$.*

LEMMA 6.15 (MERGE). *Suppose that $F \subseteq \mathbb{F}$ is encapsulated in a well-formed memory state $M = \langle G, K \rangle$. Then, the following hold:*

- (1) *Let α be an pRC11 transition label with $\text{method}(\alpha) \in F$. Suppose that if $\alpha \in \text{ProgLab}$ and $\text{typ}(\alpha) \in \{\text{W}, \text{R}, \text{RMW}\}$, then $\text{method}(e) \in F$ for every $e \in G.E$ with $\text{sp}(e) = \text{sp}(\alpha)$. Then, $M|_F \xrightarrow{\alpha}_{\text{pRC11}} M'_F$ implies that $M \xrightarrow{\alpha}_{\text{pRC11}} M'$ for some M' such that $M'|_F = M'_F$.*
- (2) *Let α be an pRC11 transition label with $\text{method}(\alpha) \notin F$. Suppose that if $\alpha \in \text{ProgLab}$ and $\text{typ}(\alpha) \in \{\text{W}, \text{R}, \text{RMW}\}$, then $\text{method}(e) \notin F$ for every $e \in G.E$ with $\text{sp}(e) = \text{sp}(\alpha)$. Then, $M|_{\bar{F}} \xrightarrow{\alpha}_{\text{pRC11}} M'_{\bar{F}}$ implies that $M \xrightarrow{\alpha}_{\text{pRC11}} M'$ for some M' such that $M'|_{\bar{F}} = M'_{\bar{F}}$.*

7 THE LIBRARY ABSTRACTION THEOREM

In this section we state and prove the library abstraction theorem. This theorem assumes two libraries implementing the same set of methods, an *implementation* L and a *specification* $L^\#$ with $F = \text{dom}(L) = \text{dom}(L^\#)$, and two programs, a *client* Pr and a *most general client* MGC . The latter is representing the library's calling policy.

Example 7.1. For a library with no restrictions whatsoever on its clients (beyond the separation of variable spaces) one can use a most general client that repeatedly invokes arbitrary library methods with arbitrary stores. We denote this client by *MGC*_{free}. On the right we present the code of the main method in each thread τ in *MGC*_{free} for $dom(L) = \{f_1, \dots, f_n\}$. We use **havoc** for arbitrarily modifying all registers.

```
BEGIN : havoc; goto f1 | ... | fn;
f1 : call(f1); goto BEGIN;
...
fn : call(fn); goto BEGIN;
```

Example 7.2. Alternatively, a policy that requires to call the library methods in a race-free fashion is captured by a most general client that uses for each thread the code on the right. Here, we hold a lock L while executing every method. We assume a standard lock implementation using release/acquire accesses (see §8).

```
BEGIN : havoc; goto f1 | ... | fn;
f1 : acquire(L); call(f1);
      release(L); goto BEGIN;
...
fn : acquire(L); call(fn);
      release(L); goto BEGIN;
```

Beyond syntactic separation of variable spaces between libraries and their clients (see Def. 6.1), the library abstraction theorem has two conditions:

- L should *refine* $L^\#$ w.r.t. *MGC*, denoted by $L \sqsubseteq_{MGC} L^\#$, and formally defined by:

$$L \sqsubseteq_{MGC} L^\# \iff \Delta \quad H_F(MGC[L]) \subseteq H_F(MGC[L^\#]) \wedge MGC[L] \text{ is not racy under pRC11}$$

- Pr should *adhere to MGC* w.r.t. $L^\#$, denoted by $Pr \sqsubseteq_{L^\#} MGC$, and formally defined by:

$$Pr \sqsubseteq_{L^\#} MGC \iff \Delta \quad H_F(Pr[L^\#]) \subseteq H_F(MGC[L^\#]) \wedge Pr[L^\#] \text{ is not racy under pRC11}$$

The first condition is an obligation of the library developer. It requires that all histories generated for the most general client using the implementation are also generated with the specification. Additionally, the implementation should not have data races when run by the most general client. Importantly, both conditions do not mention Pr : library developers have to be able to verify their implementations “once and for all” without access to a particular client program.

Dually, the second condition is an obligation of the client. It requires that the client adheres to the library policy (otherwise, the blame is on the client), which means that all histories generated by the client program should be also generated by the most general client, which expresses that policy. Additionally, the client program should not have data races. Importantly, both conditions do not mention L : clients should be able to apply the abstraction theorem without access to the library implementation. In fact, it suffices to assume that Pr uses $L^\#$ both when checking for adherence to the library’s calling policy and for checking for data-race freedom (this can be important for both checks if the calling policy relies on the values returned by the library).

Now, what should the theorem guarantee? Intuitively, we want all client behaviors observable when using L to be observable when using $L^\#$. Thus, for every trace t generated by $Pr[L]$ that reaches a program state \bar{p} and a memory state M , there should exist a corresponding trace $t^\#$ generated by $Pr[L^\#]$ that reaches a program state $\bar{p}^\#$ and a memory state $M^\#$, and a client should not be able to observe the difference between t and $t^\#$, \bar{p} and $\bar{p}^\#$, and M and $M^\#$. This, however, does not mean that there is not any difference between these objects: L and $L^\#$ may perform different operations (leading to different traces), use different variables (leading to different memories), and internally use different registers (leading to different program states). We capture “client-equivalence” by requiring that: (1) the \bar{F} -parts of t and $t^\#$ coincide (using Def. 6.9); (2) the state $\bar{p}^\#$ is obtained from \bar{p} by only changing the instruction sequence state for threads that are currently inside a method in F (using Def. 6.7); and (3) the \bar{F} -restriction of M and $M^\#$ coincide (using Def. 6.12). Finally, we also want to ensure that $Pr[L]$ is not racy.

All in all, we reach the following statement of the abstraction theorem.

THEOREM 7.3 (LIBRARY ABSTRACTION). *Let L and $L^\#$ be libraries implementing the same set F of methods. Let MGC and Pr be programs, such that both L and $L^\#$ are safe for both MGC and Pr . Suppose that $L \sqsubseteq_{MGC} L^\#$ and $Pr \sqsubseteq_{L^\#} MGC$. Then, the following hold:*

- *If $\langle \bar{p}_{\text{Init}}, M_{\text{Init}} \rangle \xrightarrow{t}_{Pr[L] \bowtie pRC11} \langle \bar{p}, M \rangle$, then there exist $t^\#$ and $\langle \bar{p}^\#, M^\# \rangle$ such that the following hold:*
 - (1) $\langle \bar{p}_{\text{Init}}, M_{\text{Init}} \rangle \xrightarrow{t^\#}_{Pr[L^\#] \bowtie pRC11} \langle \bar{p}^\#, M^\# \rangle$; (2) $t^\#|_{\bar{F}} = t|_{\bar{F}}$; (3) $\bar{p}^\# = \bar{p}[F \mapsto \bar{p}_{\text{lib}}]$ for some \bar{p}_{lib} (in particular, $\bar{p}^\#(\tau) = \bar{p}(\tau)$ whenever $\bar{p}(\tau).f \notin F$); and (4) $M^\#|_{\bar{F}} = M|_{\bar{F}}$.
- *$Pr[L]$ is not racy under $pRC11$.*

In particular, the conditions of [Thm. 7.3](#) ensure that $L \sqsubseteq_{Pr} L^\#$. The following property, which allows compositional verification of a library consisting of several (non-interacting) libraries, is obtained as a corollary of the abstraction theorem as in [\[Khyzha and Lahav 2021\]](#).

COROLLARY 7.4 (COMPOSITIONALITY). *Let L_1, \dots, L_n be libraries implementing pairwise disjoint sets of methods, such that $\text{Space}(L_1), \dots, \text{Space}(L_n), \text{Space}(L_1^\#), \dots, \text{Space}(L_n^\#)$, and $\text{Space}(MGC \setminus \text{dom}(L_1 \uplus \dots \uplus L_n))$ are pairwise disjoint. Suppose that for every $1 \leq i \leq n$, we have $L_i \sqsubseteq_{MGC_i} L_i^\#$ for $MGC_i = MGC[L_1^\# \uplus \dots \uplus L_{i-1}^\# \uplus L_{i+1}^\# \uplus \dots \uplus L_n^\#]$. Then, $L_1 \uplus \dots \uplus L_n \sqsubseteq_{MGC} L_1^\# \uplus \dots \uplus L_n^\#$.*

The rest of this section is devoted to sketch the main steps in the proof of the abstraction theorem. First, we prove a general ‘‘composition lemma’’ which allows us to take a client’s portion from one trace and glue it together with a library’s portion from another trace, provided that the two traces induce the same history. The proof of this lemma is by induction on the length of the traces, where [Lemmas 6.13](#) to [6.15](#) provide the main tools for the induction step.

LEMMA 7.5 (COMPOSITION). *Let L and L' be libraries implementing the same set F of methods such that both are safe for a program Pr , and L is also safe for a program Pr' . Suppose that $\langle \bar{p}_{\text{Init}}, M_{\text{Init}} \rangle \xrightarrow{t_{\text{cl}}}_{Pr[L'] \bowtie pRC11} \langle \bar{p}_{\text{cl}}, M_{\text{cl}} \rangle$ and $\langle \bar{p}_{\text{Init}}, M_{\text{Init}} \rangle \xrightarrow{t_{\text{lib}}}_{Pr'[L] \bowtie pRC11} \langle \bar{p}_{\text{lib}}, M_{\text{lib}} \rangle$, with $H_F(t_{\text{cl}}) = H_F(t_{\text{lib}})$. Then, $\langle \bar{p}_{\text{Init}}, M_{\text{Init}} \rangle \xrightarrow{t}_{Pr[L] \bowtie pRC11} \langle \bar{p}_{\text{cl}}[F \mapsto \bar{p}_{\text{lib}}], M \rangle$ for some trace t and memory state M such that $t|_{\bar{F}} = t_{\text{cl}}|_{\bar{F}}$, $t|_F = t_{\text{lib}}|_F$, $M|_{\bar{F}} = M_{\text{cl}}|_{\bar{F}}$, and $M|_F = M_{\text{lib}}|_F$.*

Now, using the composition lemma, we are able to show the following key property.

LEMMA 7.6. *Under the conditions of [Thm. 7.3](#), $H_F(Pr[L]) \subseteq H_F(MGC[L^\#])$.*

PROOF (OUTLINE). Assume otherwise, and let h be a shortest history in $H_F(Pr[L]) \setminus H_F(MGC[L^\#])$. Let t be a shortest trace of $Pr[L] \bowtie pRC11$ with $H_F(t) = h$. Since ε (the empty history) is clearly in $H_F(MGC[L^\#])$, we know that t is non-empty. Consider the last transition label α in t , and let t' such that $t = t' \cdot \alpha$. The minimality of t ensures that α must be an element of HLab_F (i.e., call, return, call propagation or return propagation of some method in F). Let $h' = H_F(t')$. The minimality of h further ensures that $h' \in H_F(MGC[L^\#])$. Let $t'_\#$ and $\langle \bar{p}'_\#, M'_\# \rangle$ such that $H_F(t'_\#) = h'$ and $\langle \bar{p}_{\text{Init}}, M_{\text{Init}} \rangle \xrightarrow{t'_\#}_{MGC[L^\#] \bowtie pRC11} \langle \bar{p}'_\#, M'_\# \rangle$. Let $\langle \bar{p}', M' \rangle$ be such that $\langle \bar{p}_{\text{Init}}, M_{\text{Init}} \rangle \xrightarrow{t'}_{Pr[L] \bowtie pRC11} \langle \bar{p}', M' \rangle \xrightarrow{\alpha}_{Pr[L] \bowtie pRC11} \langle \bar{p}, M \rangle$. We consider the following cases:

- (1) $\alpha \in \text{Call}_F \cup \text{CP}_F$: Using [Lemma 7.5](#) (applied with $L := L^\#$, $L' := L$, $Pr := Pr$, and $Pr' := MGC$), there exist $t''_\#$ and $M''_\#$ such that $\langle \bar{p}_{\text{Init}}, M_{\text{Init}} \rangle \xrightarrow{t''_\#}_{Pr[L^\#] \bowtie pRC11} \langle \bar{p}'[F \mapsto \bar{p}'_\#], M''_\# \rangle$, $H_F(t''_\#) = h'$, and $M''_\#|_{\bar{F}} = M'_\#|_{\bar{F}}$. Then, using [Lemmas 6.13](#) and [6.15](#), it is straightforward to show that α is enabled in $\langle \bar{p}'[F \mapsto \bar{p}'_\#], M''_\# \rangle$, and it follows that $h = h' \cdot \alpha = H_F(t''_\#) \cdot \alpha \in H_F(Pr[L^\#])$. Since $Pr \sqsubseteq_{L^\#} MGC$, we obtain $h \in H_F(MGC[L^\#])$, which contradicts our assumption.

- (2) $\alpha \in \text{Ret}_F \cup \text{RP}_F$: Using [Lemma 7.5](#) (applied with $L := L$, $L' := L^\#$, $Pr := \text{MGC}$, and $Pr' := Pr$), there exist t'' and M'' such that $\langle \bar{p}_{\text{Init}}, M_{\text{Init}} \rangle \xrightarrow{t''}_{\text{MGC}[L] \bowtie \text{pRC11}} \langle \bar{p}'_F[F \mapsto \bar{p}'], M'' \rangle$, $H_F(t'') = h'$, and $M''|_F = M'|_F$. Then, using [Lemmas 6.13](#) and [6.15](#), it is straightforward to show that α is enabled in $\langle \bar{p}'_F[F \mapsto \bar{p}'], M'' \rangle$, and it follows that $h = h' \cdot \alpha = H_F(t'') \cdot \alpha \in H_F(\text{MGC}[L])$. Since $L \sqsubseteq_{\text{MGC}} L^\#$, we obtain $h \in H_F(\text{MGC}[L^\#])$, which contradicts our assumption. \square

Then, to prove the abstraction theorem we need one additional lemma.

LEMMA 7.7. *Suppose that $M_{\text{Init}} \xrightarrow{t}_{\text{pRC11}} \langle G, K \rangle$. Then, there exist a trace t' and a knowledge mapping K' for G such that the following hold:*

- $M_{\text{Init}} \xrightarrow{t'}_{\text{pRC11}} \langle G, K' \rangle$.
- $t'|_{\text{ProgLab}} = t|_{\text{ProgLab}}$.
- $K'(\tau, X) \cap (W_X^{\text{prel}} \cup W_X^{\text{rel}} \cup \text{CR}) \subseteq \text{dom}(G.\text{hb}_X^?; [E^\tau])$ for every $\tau \in \text{Tid}$ and $X \in \text{Space}$.

The abstraction theorem is proved as follows.

PROOF OF THM. 7.3 (SKETCH). By [Lemma 7.6](#), we have $H_F(Pr[L]) \subseteq H_F(\text{MGC}[L^\#])$. Then, the first part of the claim directly follows using [Lemma 7.5](#) by letting $L := L^\#$, $L' := L$, $Pr := Pr$, and $Pr' := \text{MGC}$. For the second part, suppose that $Pr[L]$ is racy under pRC11. Let $\langle \bar{p}, \langle G, K \rangle \rangle$ be a reachable state in $Pr[L] \bowtie \text{pRC11}$ such that G is racy, and let e_1, e_2 be two events that form a race in G . Let $E = \text{dom}((G.\text{po} \cup G.\text{rf})^*; [e_1, e_2])$ and $G' = G|_E$. Clearly, e_1 and e_2 form a race in G' . It is also easy to see that there exist \bar{p}' and K' such that $\langle \bar{p}', \langle G', K' \rangle \rangle$ is reachable in $Pr[L] \bowtie \text{pRC11}$. By [Lemma 7.7](#), there exists a knowledge mapping K'' for G' such that $\langle \bar{p}', \langle G', K'' \rangle \rangle$ is reachable in $Pr[L] \bowtie \text{pRC11}$ as well and $K''(\tau, X) \cap \text{CR} \subseteq \text{dom}(G'.\text{hb}_X^?; [E^\tau])$ for every $\tau \in \text{Tid}$ and $X \in \text{Space}$. By [Prop. 6.6](#), F is encapsulated in $\langle G', K'' \rangle$. Thus, since $\text{sp}(e_1) = \text{sp}(e_2)$, either (i) $e_1 \in E_{\bar{F}} \wedge e_2 \in E_{\bar{F}}$ or (ii) $e_1 \in E_F \wedge e_2 \in E_F$. If (i) holds, we obtain a contradiction by obtaining a reachable state in $Pr[L^\#] \bowtie \text{pRC11}$ in which e_1 and e_2 form a race. If (ii) holds, we obtain a contradiction by obtaining a reachable state in $\text{MGC}[L] \bowtie \text{pRC11}$ in which e_1 and e_2 form a race. \square

8 ILLUSTRATIVE APPLICATIONS

In this section we present three examples of applications of library abstraction of different kinds. To simplify the presentation, we use an extended programming language with standard loops and conditionals, mechanisms for passing arguments in method calls/returns, arrays, etc. While these are not directly supported in our development above, the intention should be clear and the translation to programs as defined in [§3](#) is straightforward (but tedious). Each of the libraries below has its own private variable space denoted by X_L . For simplicity, we use macros for lock operations:

```

acquire(x)  $\triangleq$  LOCK : l := CAS( $X_L, x, 0, 1, \text{acq}, \text{rlx}, \text{rlx}$ ); if l  $\neq$  0 goto LOCK
release(x)  $\triangleq$  store( $X_L, x, 0, \text{rel}$ )
pacquire(x)  $\triangleq$  LOCK : l := CAS( $X_L, x, 0, 1, \text{pacq}, \text{rlx}, \text{rlx}$ ); if l  $\neq$  0 goto LOCK
prelease(x)  $\triangleq$  store( $X_L, x, 0, \text{prel}$ )

```

8.1 Read-Copy-Update Synchronization

Read-copy-update (RCU) is a synchronization mechanism, heavily used in the Linux kernel, that allows a (single) writer to safely manipulate a data structure while multiple readers are concurrently accessing it [[Mckenny 2004](#)]. For that matter, it provides *RCU critical sections* for the readers (delimited by `rcu_read_lock` and `rcu_read_unlock`) and a synchronization method (`synchronize_rcu`) for the writer that waits for all readers currently in critical sections to exit

them. Readers should access the structure inside a critical section and the writer should not perform destructive updates before calling the synchronization method [Desnoyers et al. 2012].

Intuitively, every invocation of `synchronize_rcu` begins another ‘grace period’, and the implementation ensures that “read-side critical sections cannot span grace periods”.⁶ Following various long-lasting informal discussions among developers, Alglave et al. [2018] formalized this guarantee in the context of their proposed Linux Kernel memory model. They provided specialized ad hoc semantics to the RCU primitives in the form of declarative consistency constraints.

Next, we use our framework to demonstrate a simple specification of RCU under weak memory that is based on *standard locks*. We believe that our specification has the advantage of being more parsimonious and amenable to formal verification of client programs using RCU: it can be understood by relying solely on the semantics of locks, and allows verification using techniques and tools that already support reasoning about locks.

Our specification is given below. It assumes an MGC in which a particular thread (the writer) is repeatedly calling `synchronize_rcu`, and each other thread is a reader that interleaves invocations of `rcu_read_lock` and `rcu_read_unlock` with its own thread identifier, thus acquiring and releasing a per-reader lock $l[\tau]$. The set *Readers* consist of all thread identifiers except for the writer’s identifier, and the `foreach` loop iterates over this set in an arbitrary order (which may vary between invocations).

```
rcu_read_lock( $\tau$ ) :
  acquire( $l[\tau]$ );
  return();
```

```
rcu_read_unlock( $\tau$ ) :
  release( $l[\tau]$ );
  return();
```

```
synchronize_rcu :
  foreach  $\tau \in \text{Readers}$ 
    acquire( $l[\tau]$ );
    release( $l[\tau]$ );
  return();
```

Each RCU critical section is protected by a per-reader lock. For synchronization, the writer acquires and immediately releases each of the reader locks. This ensures that RCU critical sections do not span grace periods. For example, using the specification, it is easy to conclude that the following behavior of the client program is disallowed (as usual, 0 is the initial value of all variables):

<pre>store(X, x, 1, rlx); synchronize_rcu(); store(X, y, 1, rlx);</pre>	<pre>rcu_read_lock(T_2); a := load(X, x, rlx); // 0 b := load(X, y, rlx); // 1 rcu_read_unlock(T_2);</pre>
---	--

Indeed, roughly speaking, since the read of x returns 0, we know that the writer must have acquired the reader’s lock after the reader has released it. In turn, the value 1 has been written to y only after the reader exited its critical section.

Remark 4. Gotsman et al. [2013] introduced a verification technique *assuming SC* that is based on a simpler RCU specification. In their specification each reader thread τ sets a flag `racs[τ]` to 1 when entering the critical section, and resets it to 0 when exiting the critical section. The writer waits for all reader flags to be 0, which directly means that a critical section cannot span over more than one grace period. Using release/acquire accesses, their specification is as follows:

⁶See <https://lwn.net/Articles/573497/> [Accessed July 2022].

```
rcu_read_lock( $\tau$ ) :
  store( $X_L$ , rcs[ $\tau$ ], 1, rel);
  return();
```

```
rcu_read_unlock( $\tau$ ) :
  store( $X_L$ , rcs[ $\tau$ ], 0, rel);
  return();
```

```
synchronize_rcu :
  foreach  $\tau \in$  Readers
    repeat
      a := load( $X_L$ , rcs[ $\tau$ ], acq)
      until a = 0;
  return();
```

Interestingly, this specification generates the same invocation-response sequences (i.e., same histories as employed in standard linearizability) as the one we propose. Nevertheless, it is *too weak* under weak memory concurrency, and, in fact, it allows the annotated behavior of the client above.

Following [Alglave et al. 2018; Desnoyers et al. 2012], we present a simplified implementation of the RCU library with weak memory constructs. It uses a global and per-reader ‘phase bits’, phase and rphase[τ], as well as per-reader flags, rcs[τ]. For correctness, it employs barriers (`smp_mb` in Linux or `atomic_thread_fence(memory_order_seq_cst)` in C11). While these are not included in the model we consider, they can be implemented here by acquire-release RMWs to an otherwise unused variable [Lahav et al. 2016]. Thus, `fence()` below is syntactic sugar for $r := \text{FADD}(X_L, f, 0, \text{acq}, \text{rel})$.

```
rcu_read_lock( $\tau$ ) :
  a := load( $X_L$ , phase, acq);
  store( $X_L$ , rphase[ $\tau$ ], a, rel);
  store( $X_L$ , rcs[ $\tau$ ], 1, rel);
  fence();
  return();
```

```
rcu_read_unlock( $\tau$ ) :
  store( $X_L$ , rcs[ $\tau$ ], 0, rel);
  return();
```

```
synchronize_rcu :
  fence();
  store( $X_L$ , phase, 1, rel);
  foreach  $\tau \in$  Readers
    repeat
      a := load( $X_L$ , rcs[ $\tau$ ], acq);
      b := load( $X_L$ , rphase[ $\tau$ ], acq)
      until a = 0  $\vee$  b = 1;
  store( $X_L$ , phase, 0, rel);
  foreach  $\tau \in$  Readers
    repeat
      a := load( $X_L$ , rcs[ $\tau$ ], acq);
      b := load( $X_L$ , rphase[ $\tau$ ], acq)
      until a = 0  $\vee$  b = 0;
  return();
```

In this implementation, for entering the critical section each reader τ stores phase in rphase[τ] and announces it enters the critical section by setting rcs[τ] to 1. When exiting the critical section, the reader resets rcs[τ] to 0. For the writer synchronization, the writer switches phase from 0 to 1, waits for each reader until it sees rcs[τ] = 0 or rphase[τ] = 1, switches phase back to 0, and waits again for each reader until it sees rcs[τ] = 0 or rphase[τ] = 0.

To verify refinement (i.e., to show $H_F(MGC[L]) \subseteq H_F(MGC[L^\#])$) we used the FDR refinement checker [Gibson-Robinson et al. 2014], which has been used before for automatic linearizability checking [Lowe 2017], and gives us guarantees up to a certain bound. Concretely, we managed to complete the automatic check for two readers with at most four unpropagated events between every two threads and for three readers with at most one unpropagated event. Nevertheless, we note that the length of the traces being checked is unbounded. A standard limitation here is the modeling gap between the paper definitions and FDR’s encoding, which uses communicating sequential processes (CSP). Using the fact that the library specification and implementation employ solely `rel/acq` accesses, which implies that propagation has to follow the program order and reads only read from propagated writes (see Remark 3), we can model pRC11’s memory state using FIFO buffers, which are supported by FDR. (But, using Thm. 7.3, refinement is guaranteed also when

non-rel/acq accesses are used by clients as in the example above.) We leave a systematic FDR encoding, and possibly the verification of more efficient variants of the above implementation that use rlx accesses to future work.

Finally, we note that while every history of the implementation is also a history of the specification, the implementation has important advantages over the specification. For instance, consider a scenario where a reader is repeatedly entering critical sections and the writer invokes `synchronize_rcu`. With the specification, to complete its invocation, the writer has to be “lucky” and catch a free lock between two reader’s sections. In turn, with the implementation, the writer is able to complete its invocation using the phase bit, without any particular “lucky” scheduling.

8.2 A Relaxed Concurrent Queue

Our second example continues the discussion in §2.2 and presents a specification and an implementation of a weak concurrent queue. The queue object is specified as follows:

```
enqueue(v) :
  pacquire(LT);
  tp := load(XL, T, rlx);
  store(XL, q[tp], v, rlx);
  store(XL, T, tp + 1, rlx);
  prelease(LT);
  return();
```

```
dequeue :
  pacquire(LH);
  hp := load(XL, H, rlx);
  hc := load(XL, q[hp], rlx);
  if hc ≠ ⊥ then
    store(XL, H, hp + 1, rlx);
  prelease(LH);
  return(hc);
```

The specification code uses an unbounded array q (represented by infinitely many variables $q[0], q[1], \dots$) for the queue elements that initially contains only \perp . The variables H and T store the index of the current head and tail of the queue. Then, the specification wraps a sequential implementation inside locks, making three choices (aiming to demonstrate the framework’s flexibility):

- Except for the locks, all accesses are *relaxed*, so synchronization is only induced by the locks.
- The locks are accessed using *partial* acquire/release. This means that the queue does not expose its internal synchronization to the client, thus allowing all behaviors demonstrated in §2.2.
- The `enqueue` and `dequeue` methods use *two different locks*, which allows for certain “non-linearizable” behaviors. For example, after `enqueue` returns in one thread, a `dequeue` by another thread can retrieve \perp (that signifies an empty queue). Indeed, a relaxed read from $q[k]$ may return either the value that was written to $q[k]$ or the initial value. To avoid these behaviors, clients, if they want, may form their own synchronization between enqueues and dequeues (and analyze possible behaviors of the queue using the above specification).

A possible implementation that avoids the locks and uses only relaxed accesses is shown next.

```
enqueue(v) :
  BEGIN : tp := load(XL, T, rlx);
  if tp ≠ 0 then
    tc := load(XL, q[tp - 1], rlx);
    if tc = ⊥ goto BEGIN;
  tp' := CAS(XL, T, tp, tp + 1, rlx, rlx, rlx);
  if tp' ≠ tp goto BEGIN;
  store(XL, q[tp'], v, rlx);
  return();
```

```
dequeue :
  BEGIN : hp := load(XL, H, rlx);
  hc := load(XL, q[hp], rlx);
  if hc = ⊥ then
    return(⊥);
  else
    hp' := CAS(XL, H, hp, hp + 1, rlx, rlx, rlx);
    if hp' ≠ hp goto BEGIN;
  return(hc);
```

This implementation uses CAS instructions to update H and T . It exploits the freedom allowed by the specification: no synchronization is ever induced between queue operations, and certain “non-linearizable” behaviors are allowed (when the read of $q[hp]$ returns the overwritten \perp value).

We note that it is important that `enqueue` waits for the previous cell being filled, which corresponds to all `enqueue`'s being totally ordered in the specification. In [Singh and Lahav 2022] we provide a sketch of the refinement proof, which generally follows a standard simulation argument albeit assuming non-SC memory and considering call/return propagation as observable transitions.

8.3 Local Data-Race-Freedom as an Instance of Library Abstraction

Our third application relies on the library policy component (*MGC*) of the abstraction theorem to derive a *local data-race-freedom guarantee* (LDRF, for short) for the underlying memory model (a fragment of RC11).⁷ Local data-race-freedom guarantees, introduced in [Dolan et al. 2018] and further developed in [Cho et al. 2021], generalize the more well-known (global) data-race-freedom guarantee, by ensuring strong semantics for locations accessed by non-racy executions. Crucially, the premise of these guarantees (i.e., what locations are racy) is checked assuming the strong semantics for the specified locations, which makes LDRF particularly useful for modular reasoning [Cho et al. 2021].

In our case, in the role of the strong semantics we have release/acquire (RA, for short) semantics (thus we establish what is called LDRF-RA in [Cho et al. 2021]), and we want to show that if a program avoids races on a set S of variables, then it is safe to assume that the accesses to S have RA semantics (as if they have `rel/acq` access modes). Moreover, it suffices to establish the premise, namely the avoidance of races on S , assuming the accesses to S have RA semantics.

To formulate this intricate property in terms of refinement, consider a library implementing methods `writex` and `readx` for every variable $x \in S$. The library specification ($L^\#$) is as follows:

```
writex(v) : store(ℳL, x, v, rel); return();
```

```
readx : a := load(ℳL, x, acq); return(a);
```

The implementation (L) employs non-atomic accesses (it could use any access mode):

```
writex(v) : store(ℳL, x, v, na); return();
```

```
readx : a := load(ℳL, x, na); return(a);
```

Then, the assurance provided by LDRF is precisely (conditioned) contextual refinement between L and $L^\#$. Now, to express the data-race-freedom premise we take *MGCdrf* to be a program that repeatedly and non-deterministically calls the library methods with arbitrary arguments, but avoids races by properly using standard per-variable readers-writer lock (which allows concurrent read operations but write operations require exclusive access). That is, before calling `writex` (respectively, `readx`) *MGCdrf* takes a lock on x with a write-mode (respectively, read-mode).

Adherence to the calling policy as specified by *MGCdrf* (i.e., $Pr \sqsubseteq_{L^\#} MGCdrf$) is equivalent to the LDRF premise. In particular, $H_{dom(L)}(Pr[L^\#]) \subseteq H_{dom(L)}(MGCdrf[L^\#])$ means that in histories in $H_{dom(L)}(Pr[L^\#])$ for every variable $x \in S$: (1) before calling `writex` by thread τ , the return markers of all previous invocations of `writex` and `readx` have propagated to thread τ ; and (2) before calling `readx` by thread τ , the return markers of all previous invocations of `writex` have propagated to thread τ . In addition, adherence to the policy in our abstraction theorem assumes $L^\#$ (rather than L) just like LDRF's premise assuming RA semantics for accesses to S .

Accordingly, by applying the abstraction theorem we can establish LDRF-RA by proving our library correctness criterion $L \sqsubseteq_{MGCdrf} L^\#$ for the libraries above. This is straightforward: the conditions on propagation events of calls and returns that arise by $L^\#$ are already covered by *MGCdrf*. For instance, if an invocation of `readx` reads from an invocation of `writex` in an execution of *MGCdrf*[$L^\#$], then due to the `rel/acq` accesses by $L^\#$, the call `writex` has to propagate before `readx` returns, but this already holds in traces of *MGCdrf*[$L^\#$] due to *MGCdrf*.

⁷To the best of our knowledge, we are the first to observe that (local) data-race-freedom guarantee is an instance of library abstraction.

9 RELATED WORK

We have adopted several key instruments proposed in earlier work to achieve library abstraction. Concretely, [Burckhardt et al. \[2012\]](#) addressed this challenge for the x86-TSO architecture [[Owens et al. 2009](#)], and achieved library abstraction by including call/return markers in the thread-local store buffers, and exposing the propagation of these markers to the main memory in library histories. We adapt this idea to allow point-to-point communication and apply it in a much weaker memory model. (Moreover, [Burckhardt et al. 2012](#) does not consider libraries that rely on synchronization by the client or not meant to expose their internal synchronization to the client, since these issues do not arise on x86-TSO.) In turn, from [Khyzha and Lahav 2022](#)], which studied library abstraction under non-volatile memory, we adopt the general formalism, the proof strategy of the abstraction theorem by relying on a general composition lemma ([Lemma 7.5](#)), and the modeling of a library’s calling policy as a client program (*MGC*).

Another closely related work is [Batty et al. 2013](#)], which provides the first library abstraction result for C11. For having simple specifications (which they want, like us, to be pieces of code) [Batty et al. \[2013\]](#) extended (a fragment of) C11 with specialized “atomic block” constructs, whose semantics is similar to the semantics of software transactions. In contrast, we use locks with “partial release/acquire” semantics, and have no need in including transactional features, which may be harder for client reasoning. Other crucial differences is that [Batty et al. 2013](#)] does not support library’s calling policies, and their correctness condition is based on partially ordered execution traces (with ‘guarantee’ and ‘deny’ relations), while we opt for considering only totally ordered histories, which, we believe, are easier to grasp. We also note that the relaxed atomics employed in [Batty et al. 2013](#)] have the original C11 semantics (*without* the *PO-RF* constraint, and thus *with* “out-of-thin-air” behaviors), which, as they show, does not allow fully compositional reasoning.

Library abstraction under weak memory was also studied from a *declarative* point of view in [[Dongol et al. 2018](#); [Raad et al. 2019](#)], using partial orders for exposing the externally visible synchronization induced by concurrent objects. Concretely, in the framework of [[Raad et al. 2019](#)], libraries are specified as collections of execution graphs, and the abstraction condition relates the graphs produced by the implementation to those in the specification. Their specification framework does not depend on the language constructs, and it is sufficiently expressive for having direct specifications of non-standard non-linearizable objects (e.g., queues with certain non-FIFO behaviors). We believe that there is a price for the generality: rich declarative specifications are often hard to understand and informally apply, and require non-standard methods for client reasoning. In addition, [[Raad et al. 2019](#)]’s ‘well-formedness’ requirement (which corresponds to our “policy adherence”) is defined in terms of the implementation library (unlike in our abstraction theorem).

[Smith et al. \[2020\]](#) studied linearizability on a general hardware memory model, and related it to a certain refinement notion, which they call “object refinement”. In contrast, we are aiming for standard contextual refinement between an implementation and its specification code. Other works developed correctness conditions for concurrent objects under weakly consistent memory, but were not formally related to a refinement notion (e.g., [[Doherty et al. 2018](#)]).

Weakly consistent objects were also studied by [Emmi and Enea \[2019\]](#) who proposed a specification framework based on extending standard histories with a *visibility* relation that is subject to varying constraints. Their approach was applied to non-linearizable Java concurrent objects, and allows verification by a generalization of the concept of linearization points [[Krishna et al. 2020](#)].

Other recent works propose *logical* approaches to library specification and verification under weak memory [[Dalvandi and Dongol 2021a,b](#); [Dang et al. 2022](#); [Mével and Jourdan 2021](#)]. In contrast to our work, their specifications are given as Hoare triples and refinement is understood from a program logic perspective. Specifically, for a fragment of the model we study (without non-atomics),

Dalvandi and Dongol [2021a,b] propose an Owicki-Gries-style Hoare logic and specify libraries with view-based object semantics. Their approach is, however, limited to clients who synchronize only through the library operations. In turn, Mével and Jourdan [2021] and Dang et al. [2022] are based on the idea of *logical atomicity* [Birkedal et al. 2021; Jung 2019], and provide mechanized verification of certain library implementations in the Iris framework [Jung et al. 2015]. The weak memory model in [Dang et al. 2022] is similar to the one we study (RC11), while the model in [Mével and Jourdan 2021] is Multicore OCaml, which provides much stronger accesses than RC11. The library studied in [Mével and Jourdan 2021] is a queue that is stronger than our example in §8.2. It exposes synchronization to the client between from an enqueue of some element to the dequeue of that specific element, but still, unlike a lock-based queue, does not expose synchronization in all other cases (e.g., between two different enqueues).

Finally, the idea of having explicit steps for propagating knowledge between threads appeared before in semantics of relaxed memory concurrency, particularly in an operational model for the POWER architecture [Sarkar et al. 2011], but we are not aware of a similar existing model for RC11.

10 CONCLUSION

We established a library abstraction for (a fragment of) the RC11 model extended with specialized partial release/acquire accesses. The latter are used to have simple lock-based specifications of libraries that do not expose internal synchronization to their clients or libraries that rely on synchronization by the client. The condition for library abstraction is based on inclusion of sets of totally ordered library histories. To achieve this, we developed an operational version of the model with explicit steps for propagation of knowledge between threads, and included the propagation of method invocation and responses in library histories.

Our work has several limitations, which are interesting to lift and address in future work. In particular: (1) The fragment of RC11 that we consider excludes fences and SC-atomics; (2) We only consider partial correctness and finite histories ignoring liveness issues and thus our refinement notion is not termination preserving (see [Gotsman and Yang 2011] for a possible approach to lift this limitation); (3) We disallow nested method calls (in particular, recursion), as common in works relating linearizability to refinement (also under SC); and (4) We disallow libraries and their clients to transfer ownership of data structures among themselves or to run in a shared address space (see [Gotsman and Yang 2013] for a possible approach to lift this limitation).

Verification techniques for both library development and client reasoning are beyond the scope of the current work. In particular, it is interesting to see how our specification constructs can be supported by a program logic, which will pave the way to foundational mechanized refinement proofs as in [Dang et al. 2022], as well as by model checking tools, such as GenMC [Kokologiannakis et al. 2019] and C11Tester [Luo and Demsky 2021], that can be used for client reasoning.

Finally, we are interested in developing a view-based operational model (see, e.g., [Kaiser et al. 2017; Kang et al. 2017]) with explicit propagation steps that will be equivalent to pRC11, but possibly more intuitive. We also believe that pRC11 (in particular, its `rel/acq` fragment, see Remark 3) may be of independent interest for developing verification techniques, by relying on the analogy between pRC11's knowledge propagation and well-studied x86-TSO store propagation steps (see e.g., [Bouajjani et al. 2018; Enea and Farzan 2016]).

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their valuable feedback. This work was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 851811) and the Israel Science Foundation (grant numbers 1566/18 and 814/22).

REFERENCES

- Jade Alglave, Luc Maranget, Paul E. McKenney, Andrea Parri, and Alan Stern. 2018. Frightening Small Children and Disconcerting Grown-ups: Concurrency in the Linux Kernel. In *ASPLOS*. ACM, New York, NY, USA, 405–418. <https://doi.org/10.1145/3173162.3177156>
- Jade Alglave, Luc Maranget, and Michael Tautschnig. 2014. Herding Cats: Modelling, Simulation, Testing, and Data Mining for Weak Memory. *ACM Trans. Program. Lang. Syst.* 36, 2, Article 7 (July 2014), 74 pages. <https://doi.org/10.1145/2627752>
- Mark Batty, Mike Dodds, and Alexey Gotsman. 2013. Library Abstraction for C/C++ Concurrency. In *POPL*. ACM, New York, NY, USA, 235–248. <https://doi.org/10.1145/2429069.2429099>
- Mark Batty, Kayvan Memarian, Kyndylan Nienhuis, Jean Pichon-Pharabod, and Peter Sewell. 2015. The Problem of Programming Language Concurrency Semantics. In *ESOP*. Springer, Berlin, Heidelberg, 283–307. http://dx.doi.org/10.1007/978-3-662-46669-8_12
- Sidi Mohamed Beillahi, Ahmed Bouajjani, and Constantin Enea. 2021. Robustness Against Transactional Causal Consistency. *Logical Methods in Computer Science* Volume 17, Issue 1 (Feb. 2021). [https://doi.org/10.23638/LMCS-17\(1:12\)2021](https://doi.org/10.23638/LMCS-17(1:12)2021)
- Lars Birkedal, Thomas Dinsdale-Young, Armaël Guéneau, Guilhem Jaber, Kasper Svendsen, and Nikos Tzevelekos. 2021. Theorems for Free from Separation Logic Specifications. *Proc. ACM Program. Lang.* 5, ICFP, Article 81 (Aug. 2021), 29 pages. <https://doi.org/10.1145/3473586>
- Hans-Juergen Boehm and Brian Demsky. 2014. Outlawing Ghosts: Avoiding Out-of-thin-air Results. In *MSPC*. ACM, New York, NY, USA, Article 7, 6 pages. <https://doi.org/10.1145/2618128.2618134>
- Ahmed Bouajjani, Michael Emmi, Constantin Enea, and Jad Hamza. 2015. Tractable Refinement Checking for Concurrent Objects. In *POPL*. ACM, New York, NY, USA, 651–662. <https://doi.org/10.1145/2676726.2677002>
- Ahmed Bouajjani, Constantin Enea, Suha Orhun Mutluergil, and Serdar Tasiran. 2018. Reasoning About TSO Programs Using Reduction and Abstraction. In *CAV*. Springer International Publishing, Cham, 336–353. https://doi.org/10.1007/978-3-319-96142-2_21
- Sebastian Burckhardt, Alexey Gotsman, Madanlal Musuvathi, and Hongseok Yang. 2012. Concurrent Library Correctness on the TSO Memory Model. In *ESOP*. Springer, Berlin, Heidelberg, 87–107. https://doi.org/10.1007/978-3-642-28869-2_5
- Minki Cho, Sung-Hwan Lee, Chung-Kil Hur, and Ori Lahav. 2021. Modular Data-Race-Freedom Guarantees in the Promising Semantics. In *PLDI*. ACM, New York, NY, USA, 867–882. <https://doi.org/10.1145/3453483.3454082>
- Sadeh Dalvandi and Brijesh Dongol. 2021a. Verifying C11-Style Weak Memory Libraries. In *PPoPP*. ACM, New York, NY, USA, 451–453. <https://doi.org/10.1145/3437801.3441619>
- Sadeh Dalvandi and Brijesh Dongol. 2021b. Verifying C11-Style Weak Memory Libraries via Refinement. *CoRR* abs/2108.06944 (2021). arXiv:2108.06944 <https://arxiv.org/abs/2108.06944>
- 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*. ACM, New York, NY, USA, 792–808. <https://doi.org/10.1145/3519939.3523451>
- Mathieu Desnoyers, Paul E. McKenney, Alan S. Stern, Michel R. Dagenais, and Jonathan Walpole. 2012. User-Level Implementations of Read-Copy Update. *IEEE Trans. Parallel Distrib. Syst.* 23, 2 (2012), 375–382. <https://doi.org/10.1109/TPDS.2011.159>
- Simon Doherty, Brijesh Dongol, Heike Wehrheim, and John Derrick. 2018. Making Linearizability Compositional for Partially Ordered Executions. In *iFM*. Springer International Publishing, Cham, 110–129. https://doi.org/10.1007/978-3-319-98938-9_7
- Stephen Dolan, KC Sivaramakrishnan, and Anil Madhavapeddy. 2018. Bounding Data Races in Space and Time. In *PLDI*. ACM, New York, NY, USA, 242–255. <https://doi.org/10.1145/3192366.3192421>
- Brijesh Dongol, Radha Jagadeesan, James Riely, and Alasdair Armstrong. 2018. On Abstraction and Compositionality for Weak-Memory Linearizability. In *VMCAI*. Springer International Publishing, Cham, 183–204. https://doi.org/10.1007/978-3-319-73721-8_9
- 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>
- Constantin Enea and Azadeh Farzan. 2016. On Atomicity in Presence of Non-atomic Writes. In *TACAS*. Springer, Berlin, Heidelberg, 497–514. https://doi.org/10.1007/978-3-662-49674-9_29
- Ivana Filipović, Peter O’Hearn, Noam Rinetzky, and Hongseok Yang. 2010. Abstraction for concurrent objects. *Theoretical Computer Science* 411, 51 (2010), 4379–4398. <https://www.sciencedirect.com/science/article/pii/S0304397510005001>
- Thomas Gibson-Robinson, Philip Armstrong, Alexandre Boulgakov, and Andrew W. Roscoe. 2014. FDR3 – A Modern Refinement Checker for CSP. In *TACAS*. Springer, Berlin, Heidelberg, 187–201. https://doi.org/10.1007/978-3-642-54862-8_13
- Alexey Gotsman, Noam Rinetzky, and Hongseok Yang. 2013. Verifying Concurrent Memory Reclamation Algorithms with Grace. In *ESOP*. Springer, Berlin, Heidelberg, 249–269. https://doi.org/10.1007/978-3-642-37036-6_15

- Alexey Gotsman and Hongseok Yang. 2011. Liveness-Preserving Atomicity Abstraction. In *ICALP*. Springer, Berlin, Heidelberg, 453–465. https://doi.org/10.1007/978-3-642-22012-8_36
- Alexey Gotsman and Hongseok Yang. 2013. Linearizability with Ownership Transfer. *Logical Methods in Computer Science* Volume 9, Issue 3 (Sept. 2013). <https://lmcs.episciences.org/931>
- Maurice P. Herlihy and Jeannette M. Wing. 1990. Linearizability: A Correctness Condition for Concurrent Objects. *ACM Trans. Program. Lang. Syst.* 12, 3 (July 1990), 463–492. <https://doi.org/10.1145/78969.78972>
- Ralf Jung. 2019. Logical atomicity in Iris: The good, the bad, and the ugly. In *Iris Workshop*. <https://people.mpi-sws.org/~jung/iris/talk-iris2019.pdf>
- Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *POPL*. ACM, New York, NY, USA, 637–650. <https://doi.org/10.1145/2676726.2676980>
- Jan-Oliver Kaiser, Hoang-Hai Dang, Derek Dreyer, Ori Lahav, and Viktor Vafeiadis. 2017. Strong Logic for Weak Memory: Reasoning About Release-Acquire Consistency in Iris. In *ECOOP*. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 17:1–17:29. <https://doi.org/10.4230/LIPIcs.ECOOP.2017.17>
- Jeehoon Kang, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, and Derek Dreyer. 2017. A Promising Semantics for Relaxed-Memory Concurrency. In *POPL*. ACM, New York, NY, USA, 175–189. <https://doi.org/10.1145/3009837.3009850>
- Artem Khyzha and Ori Lahav. 2021. Taming x86-TSO Persistency. *Proc. ACM Program. Lang.* 5, POPL, Article 47 (Jan. 2021), 29 pages. <https://doi.org/10.1145/3434328>
- Artem Khyzha and Ori Lahav. 2022. Abstraction for Crash-Resilient Objects. In *ESOP*. Springer International Publishing, Cham, 262–289. https://doi.org/10.1007/978-3-030-99336-8_10
- Michalis Kokologiannakis, Ori Lahav, Konstantinos Sagonas, and Viktor Vafeiadis. 2017. Effective Stateless Model Checking for C/C++ Concurrency. *Proc. ACM Program. Lang.* 2, POPL, Article 17 (Dec. 2017), 32 pages. <https://doi.org/10.1145/3158105>
- Michalis Kokologiannakis, Azalea Raad, and Viktor Vafeiadis. 2019. Model Checking for Weakly Consistent Libraries. In *PLDI*. ACM, New York, NY, USA, 96–110. <https://doi.org/10.1145/3314221.3314609>
- Siddharth Krishna, Michael Emmi, Constantin Enea, and Dejan Jovanović. 2020. Verifying Visibility-Based Weak Consistency. In *ESOP*. Springer International Publishing, Cham, 280–307. https://doi.org/10.1007/978-3-030-44914-8_11
- Ori Lahav, Nick Giannarakis, and Viktor Vafeiadis. 2016. Taming Release-Acquire Consistency. In *POPL*. ACM, New York, NY, USA, 649–662. <https://doi.org/10.1145/2837614.2837643>
- Ori Lahav and Roy Margalit. 2019. Robustness Against Release/Acquire Semantics. In *PLDI*. ACM, New York, NY, USA, 126–141. <https://doi.org/10.1145/3314221.3314604>
- Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, and Derek Dreyer. 2017. Repairing Sequential Consistency in C/C++11. In *PLDI*. ACM, New York, NY, USA, 618–632. <https://doi.org/10.1145/3062341.3062352>
- Leslie Lamport. 1979. How to Make a Multiprocessor Computer That Correctly Executes Multiprocess Programs. *IEEE Trans. Computers* 28, 9 (Sept. 1979), 690–691. <https://doi.org/10.1109/TC.1979.1675439>
- Gavin Lowe. 2017. Analysing Lock-Free Linearizable Datatypes Using CSP. In *Concurrency, Security, and Puzzles - Essays Dedicated to Andrew William Roscoe on the Occasion of His 60th Birthday*. Springer, 162–184. https://doi.org/10.1007/978-3-319-51046-0_9
- Weiyu Luo and Brian Demsky. 2021. C11Tester: A Race Detector for C/C++ Atomics. In *ASPLOS*. ACM, New York, NY, USA, 630–646. <https://doi.org/10.1145/3445814.3446711>
- Paul E. Mckenney. 2004. *Exploiting Deferred Destruction: An Analysis of Read-Copy-Update Techniques in Operating System Kernels*. Ph. D. Dissertation. Oregon Health & Science University.
- 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>
- Peizhao Ou and Brian Demsky. 2018. Towards Understanding the Costs of Avoiding Out-of-Thin-Air Results. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 136 (Oct. 2018), 29 pages. <https://doi.org/10.1145/3276506>
- Scott Owens, Susmit Sarkar, and Peter Sewell. 2009. A Better x86 Memory Model: x86-TSO. In *TPHOLS*. Springer, Berlin, Heidelberg, 391–407. https://doi.org/10.1007/978-3-642-03359-9_27
- Azalea Raad, Marko Doko, Lovro Rožić, Ori Lahav, and Viktor Vafeiadis. 2019. On Library Correctness under Weak Memory Consistency: Specifying and Verifying Concurrent Libraries under Declarative Consistency Models. *Proc. ACM Program. Lang.* 3, POPL, Article 68 (Jan. 2019), 31 pages. <https://doi.org/10.1145/3290381>
- Susmit Sarkar, Peter Sewell, Jade Alglave, Luc Maranget, and Derek Williams. 2011. Understanding POWER Multiprocessors. In *PLDI*. ACM, New York, NY, USA, 175–186. <https://doi.org/10.1145/1993498.1993520>
- Abhishek Kr Singh and Ori Lahav. 2022. An Operational Approach to Library Abstraction under Relaxed Memory Concurrency (Extended Version). https://www.cs.tau.ac.il/~orilahav/papers/popl23_lib_full.pdf
- Graeme Smith, Kirsten Winter, and Robert J. Colvin. 2020. Linearizability on Hardware Weak Memory Models. *Form. Asp. Comput.* 32, 1 (Feb. 2020), 1–32. <https://doi.org/10.1007/s00165-019-00499-8>

Received 2022-07-07; accepted 2022-11-07