Library Abstraction for C/C++ Concurrency

— extended version —

Mark Batty

University of Cambridge

Mike Dodds

University of York

Alexey Gotsman

IMDEA Software Institute

Abstract

When constructing complex concurrent systems, abstraction is vital: programmers should be able to reason about concurrent libraries in terms of abstract specifications that hide the implementation details. Relaxed memory models present substantial challenges in this respect, as libraries need not provide sequentially consistent abstractions: to avoid unnecessary synchronisation, they may allow clients to observe relaxed memory effects, and library specifications must capture these.

In this paper, we propose a criterion for sound library abstraction in the new C11 and C++11 memory model, generalising the standard sequentially consistent notion of linearizability. We prove that our criterion soundly captures all client-library interactions, both through call and return values, and through the subtle synchronisation effects arising from the memory model. To illustrate our approach, we verify implementations against specifications for the lock-free Treiber stack and a producer-consumer queue. Ours is the first approach to compositional reasoning for concurrent C11/C++11 programs.

Categories and Subject Descriptors D.2.4 [*Software Engineering*]: Software/Program Verification; F.3.1 [*Logics and Meanings of Programs*]: Specifying and Verifying and Reasoning about Programs

General Terms Languages, Theory, Verification

Keywords Verification, Concurrency, Modularity, C, C++

1. Introduction

Software developers often encapsulate functionality in libraries, and construct complex libraries from simpler ones. The advantage of this is information hiding: the developer need not understand each library's implementation, but only its more abstract *specification*. On a sequential system, a library's internal actions cannot be observed by its client, so its specification can simply be a relation from initial to final states of every library invocation. This does not suffice on a concurrent system, where the invocations can overlap and interact with each other. Hence, a concurrent library's specification is often given as just another library, but with a simpler (e.g., atomic) implementation; the two libraries are called *concrete* and *abstract*, respectively. Validating a specification means showing that the simpler implementation *abstracts* the more complex one, i.e., reproduces all its client-observable behaviours. Library abstraction has to take into account a variety of ways in which a client and library can interact, including values passed at library calls and returns, the contents of shared data structures and, in this paper, the *memory model*.

The memory model of a concurrent system governs what values can be returned when the system reads from shared memory. In a traditional *sequentially consistent (SC)* system, the memory model is straightforward: there is a total order over reads and writes, and each read returns the value of the most recent write to the location being accessed [14]. However, modern processors and programming languages provide *relaxed memory models*, where there is no total order of memory actions, and the order of actions observed by a thread may not agree with program order, or with that observed by other threads.

In this paper, we propose a criterion for library abstraction on the relaxed memory model defined by the new ISO C11 [11] and C++11 [12] standards (henceforth, the 'C11 model'). We handle the core of the C11 memory model, leaving more esoteric features, such as release-consume atomics and fences, as future work (see §9). The C11 model is designed to support common compiler optimisations and efficient compilation to architectures such as x86, Power, ARM and Itanium, which themselves do not guarantee SC. It gives the programmer fine-grained control of relaxed behaviour for individual reads and writes, and is defined by a set of axiomatic constraints, rather than operationally. Both of these properties produce subtle interactions between the client and the library that must be accounted for in abstraction.

Our criterion is an evolution of *linearizability* [4, 6, 9, 10], a widely-used abstraction criterion for non-relaxed systems. Like linearizability, our approach satisfies the *Abstraction Theorem*: if one library (a specification) abstracts another (an implementation), then the behaviours of any client using the implementation are contained in the behaviours of the client using the specification. This result allows complex library code to be replaced by simpler specifications, for verification or informal reasoning. Hence, it can be viewed as giving a proof technique for contextual refinement that avoids considering all possible clients. Our criterion is compositional, meaning that a library consisting of several smaller non-interacting libraries can be abstracted by considering each sub-library separately. When restricted to the SC fragment of C11, our criterion implies classical linearizability (but not vice versa).

The proposed criterion for library abstraction gives the first sound technique for specifying C11 and C++11 concurrent libraries. To justify its practicality, we have applied it two typical concurrent algorithms: a non-blocking stack and an array-based queue. To do this, we have adapted the standard linearization point technique to the axiomatic structure of the C11 model. These case studies represent the first step towards verified concurrent libraries for C11 and C++11.

Technical challenges. Apart from managing the mere complexity of the C11 model, defining a criterion for library abstraction requires us to deal with several challenges that have not been considered in prior work.

First, the C11 memory model is defined axiomatically, whereas existing techniques for library abstraction, such as linearizability, have focused on operational trace-based models. To deal with this, we propose a novel notion of a *history*, which records all interac-

tions between a client and a library. Histories in our work consist of several partial orders on call and return actions. This is in contrast to variants of linearizability, where histories are linear sequences (for this reason, in the following we avoid the term 'linearizability'). We define an abstraction relation on histories as inclusion over partial orders, and lift this relation to give our abstraction criterion for libraries: one library abstracts another if any history of the former can be reproduced in abstracted form by the latter.

Second, C11 offers the programmer a range of options for concurrently accessing memory, each with different trade-offs between consistency and performance. These choices can subtly affect other memory accesses across the client-library boundary—a particular choice of consistency level inside the library might force or forbid reading certain values in the client, and vice versa. This is an intended feature: it allows C11 libraries to define synchronisation constructs that offer different levels of consistency to clients. We propose a method for constructing histories that captures such client-library interactions uniformly. The Abstraction Theorem certifies that our histories indeed soundly represent *all* possible interactions.

Finally, some aspects of the C11 model conflict with abstraction. Most problematically, the model permits *satisfaction cycles*. In satisfaction cycles, the effect of actions executed down a conditional branch is what causes the branch to be taken in the first place. This breaks the straightforward assumption that faults are confined to either client or library code: a misbehaving client can cause misbehaviour in a library, which can in turn cause the original client misbehaviour! For these reasons, we actually define *two* distinct library abstraction criteria: one for general C11, and one for a language without the feature leading to satisfaction cycles. The former requires an a priori check that the client and the library do not access each others' internal memory locations, which hinders compositionality. The latter lifts this restriction (albeit for a C11 model modified to admit incomplete program runs) and thus provides evidence that satisfaction cycles are to blame for non-compositional behaviour. Our results thus illuminate corner cases in C11 that undermine abstraction, and may inform future revisions of the model.

As we argue in §9, many of the techniques we developed to address the above challenges should be applicable to other models similar to C11.

Structure. In the first part of the paper, we describe informally how algorithms can be expressed and specified in the C11 memory model (§2), and our abstraction criteria (§3). We then present the model formally (§4 and §5), followed by the criteria (§6) and a method for establishing their requirements (§7). Proofs are given in §C.

2. C11 Concurrency and Library Specification

In this section we explain the form of our specifications for C11 concurrent libraries, together with a brief introduction to programming in the C11 concurrency model itself. As a running example, we use a version of the non-blocking Treiber stack algorithm [21] implemented using the concurrency primitives in the subset of C11 that we consider. Figure 1a shows its specification, and Figure 1b its implementation, which we have proved to correspond (§7 and §E). For readability, we present examples in a pseudocode instead of the actual C/C++ syntax. Several important features are highlighted in red—these are explained below.

Stack specification. As noted in §1, specifications are just alternative library implementations that have the advantage of simplicity, in exchange for inefficiency or nondeterminism. The specification in Figure 1a represents the stack as a sequence abstract data type and provides the three expected methods: init, push and pop. A correct stack implementation should provide the illu-

SPECIFICATION:

```
atomic Seq S;



void init() {
  store_REL(&S,empty);
}

void push(int v) {
  Seq s, s2;
  if (nondet()) while(1);
  atom_sec {
    s = load_RLX(&S);
    s2 = append(s,v);
    CAS_RLX,REL(&S,s,s2);
  }
}



int pop() {
  Seq s;
  if (nondet()) while(1);
  atom_sec {
    s = load_ACQ(&S);
    if (s == empty)
      return EMPTY;
    CAS_RLX,RLX(&S,s,tail(s));
    return head(s);
  }
}
```

(a)

IMPLEMENTATION:

```
struct Node {
  int data;
  Node *next;
};
atomic Node *T;

void init() {
  store_REL(&T,NULL);
}

void push(int v) {
  Node *x, *t;
  x = new Node();
  x->data = v;
  do {
    t = load_RLX(&T);
    x->next = t;
  } while
    (!CAS_RLX,REL(&T,t,x));
}

int pop() {
  Node *t, *x;
  do {
    t = load_ACQ(&T);
    if (t == NULL)
      return EMPTY;
    x = t->next;
  } while
    (!CAS_RLX,RLX(&T,t,x));
  return t->data;
}
```

(b)

Figure 1. The Treiber stack. For simplicity, we let pop leak memory. The CASes in the specification always succeed.

sion of atomicity of operations to concurrent threads. We specify this by wrapping the bodies of push and pop in *atomic sections*, denoted by atom_sec. Atomic sections are not part of the standard C11 model—for specification purposes, we have extended the language with a prototype semantics for atomic section (§5). Both push and pop may non-deterministically diverge, as common stack implementations allow some operations to starve (in concurrency parlance, they are lock-free, but not wait-free). All these are the expected features of a specification on an SC memory model. We now explain the features specific to C11.

The sequence S holding the abstract state is declared atomic. In C11, programs must not have data races on normal variables; any location where races can occur must be explicitly identified as atomic and accessed using the special commands load, store, and CAS (*compare-and-swap*). The latter combines a load and a store into a single operation executed atomically. A CAS takes three arguments: a memory address, an expected value and a new value. The command atomically reads the memory address and, if it contains the expected value, updates it with the new one. Due to our use of atomic sections, the CASes in the specification always succeed. We use CASes here instead of just stores, because, for subtle technical reasons, the latter have a stronger semantics in C11 than our atomic sections (see release sequences in §A).

The load and store commands are annotated with a *memory order* that determines the trade-off between consistency and performance for the memory access; CASes are annotated with two memory orders, as they perform both a load and a store. The choice of memory orders inside a library method can indirectly affect its clients, and thus, a library specification must include them. In the

stack specification, several memory operations have the ***release-acquire*** memory orders, denoted by the subscripts REL (for stores) and ACQ (for loads). To explain its effect, consider the following client using the stack according to a typical message-passing idiom:

```
            int a, b, x=0;
                      ┃  do {
x=1;                  ┃    a = pop();
push(&x);             ┃  } while (a==EMPTY);
                      ┃  b=*a;
```

The first thread writes 1 into x and calls `push(&x)`; the second thread pops the address of x from the stack and then reads its contents. In general, a relaxed memory model may allow the second thread to read 0 instead of 1, e.g., because the compiler reorders x=1 and push(&x). The release-acquire annotations guarantee that this is not the case: when the ACQ load of S in pop reads the value written by the REL store to S in push, the two commands ***synchronise***. We define this notion more precisely later, but informally, it means that the ordering between the REL store and ACQ load constrains the values fetched by reads from other locations, such as the read *a in the client.

To enable this message-passing idiom, the specification only needs to synchronise from pushes to pops; it need not synchronise from pops to pushes, or from pops to pops. To avoid unnecessary synchronisation, the specification uses the ***relaxed*** memory order (RLX). This order is weaker than release-acquire, meaning that the set of values a relaxed load can read from memory is less constrained; additionally, relaxed loads and stores do not synchronise with each other. However, relaxed operations are very cheap, since they compile to basic loads and stores without any additional hardware barrier instructions. Hence, the specification allows implementations that are efficient, yet support the intended use of the stack for message passing. On the other hand, as we show below, it intentionally allows non-SC stack behaviours.

Stack implementation. Figure 1b gives our implementation of the Treiber stack. The stack is represented by a heap-allocated linked list of nodes, accessed through a top-of-stack pointer T. Only the latter needs to be `atomic`, as it is the only point of contention among threads. The `push` function repeatedly reads from the top pointer T, initialises a newly created node x to point to the value read, and tries to swing T to point to x using a CAS; `pop` is implemented similarly. For simplicity, we let `pop` leak memory.

Like the specification, the implementation avoids unnecessary hardware synchronisation by using the relaxed memory order RLX. However, the load of T in pop is annotated ACQ, since the command x = t->next accesses memory based on the value read, and hence, requires it to be up to date.

What does it mean for the implementation in Figure 1b to meet the specification in Figure 1a? As well as returning the right values, it must also faithfully implement the correct synchronisation. To understand how this can be formalised, we must therefore explain how synchronisation works in C11's semantics.

C11 model structure. The C11 memory model is defined axiomatically. An ***execution*** of a program consists of a set of ***actions*** and several partial orders on it. An action describes a memory operation, including the information about the thread that performed it, the address accessed and the values written and/or read. The semantics of a program is given by the set of executions consistent with the program code and satisfying the ***axioms*** of the memory model (see Figure 4 for a flavour of these). Here is a program with one of its executions, whose outcome we explain below:

STORE BUFFERING (SB):

```
store_RLX(&x,0); store_RLX(&y,0)

store_RLX(&x,1) ┃ store_RLX(&y,1)
load_RLX(&y)    ┃ load_RLX(&x)
```


Note that, in diagrams representing executions, we omit thread identifiers from actions. Several of the most important relations in an execution are:

- ***sequenced-before*** (sb), a transitive and irreflexive relation ordering actions by the same thread according to their program order.
- ***initialised-before*** (ib), ordering initial writes to memory locations before all other actions in the execution[1]. Above we have shown ib by a dotted line dividing the two kinds of actions.
- ***reads-from*** (rf), relating reads r to the writes w from which they take their values: $w \xrightarrow{\text{rf}} r$.
- ***happens-before*** (hb), showing the precedence of actions in the execution. In the fragment of C11 that we consider, it is transitive and irreflexive.

Happens-before is the key relation, and is the closest the C11 model has to the notion of a global time in an SC model: a read must not read any write to the same location related to it in hb other than its immediate predecessor. Thus, for writes w_1 and w_2 and a read r accessing the same location, the following shapes are forbidden:

$$w_1 \xrightarrow{\text{hb}} w_2 \xrightarrow{\text{hb}} r \qquad\qquad r \overset{\text{hb}}{\underset{\text{rf}}{\rightleftarrows}} w_1 \qquad \text{(RD)}$$

However, in contrast to an SC model, hb is partial in C11, and some reads can read from hb-unrelated writes: we might have $w \xrightarrow{\text{rf}} r$, but not $w \xrightarrow{\text{hb}} r$.

Memory orders. By default, memory reads and writes in C11 are ***non-atomic*** (NA). The memory model guarantees that data-race free programs with only non-atomic memory accesses have SC behaviour. A data race occurs when two actions on the same memory location, at least one of which is a write, and at least one of which is a non-atomic access, are unrelated in happens-before, and thus, intuitively, can take place 'at the same time'. Hence, non-expert programmers who write code that is free from both data races and atomic accesses need not understand the details of the relaxed memory model. Data races are considered faults, resulting in undefined behaviour for the whole program.

The three main atomic memory orders, from least to most restrictive, are relaxed, release-acquire and sequentially consistent[2]. We have already seen the first two in the stack example above. The third, ***sequentially consistent*** (SC), does not allow relaxed behaviour: if all actions in a race-free program are either non-atomic or SC, the program exhibits only sequentially-consistent behaviour [1, 3]. However, the SC memory order is more expensive.

The weakest memory order, ***relaxed***, exhibits a number of relaxations, as the C11 model places very few restrictions on which write a relaxed read might read from. For example, consider the (SB) example above. The outcome shown there is allowed by C11, but cannot be produced by any interleaving of the threads' actions. C11 disallows it if all memory accesses are annotated as SC.

---

[1] This is a specialisation of the *additional-synchronises-with* relation from the C11 model [1] to programs without dynamic thread creation, to which we restrict ourselves in this paper (see §4).

[2] Release-consume atomics and fences [1] are left for future work (see §9).

The ***release-acquire*** memory orders allow more relaxed behaviour than SC, while still providing some guarantees. Consider the following execution of the client of the stack in Figure 1a or 1b we have seen above:

MESSAGE PASSING (MP):

```
                                   store_NA(&x,0)
                                       ib |
    int a, b, x=0;                  store_NA(&x,1)
                                       sb |          ⋮
x=1;        ‖ do {a=pop();}       call push(&x)  rf,hb  load_ACQ(&S,j)
push(&x);   ‖ while                    sb |            sb |
            ‖    (a==EMPTY);     rmw_RLX,REL(&S,i,j)    ret pop(&x)
            ‖ b=*a;                     |        rf      sb |
                                        ⋮             load_NA(&x,1)
```

Here rmw (read-modify-write) is a combined load-store action produced by a CAS. In this case, the ACQ load in pop synchronises with the REL store part of the CAS in push that it reads from. This informal notion of synchronisation we mentioned above is formalised in the memory model by including the corresponding rf edge into hb. Then, since sb ∪ ib ⊆ hb and hb is transitive, both writes to x happen before the read. Hence, by (RD), the read from x by the second thread is forced to read from the most recent write, i.e., 1.

If all the memory order annotations in the Treiber stack were relaxed, the second thread could read 0 from x instead. Furthermore, without release-acquire synchronisation, there would be a data race between the non-atomic write of x in the first thread and the non-atomic read of x in the second.

The release-acquire memory orders only synchronise between pairs of reads and writes, but do not impose a total order over memory accesses, and therefore allow non-SC behaviour. For example, if we annotate writes in (SB) with REL, and reads with ACQ, then the outcome shown there will still be allowed: each load can read from the initialisation without generating a cycle in hb or violating (RD). We can also get this outcome if we use push and pop operations on two instances of the stack from Figure 1a or 1b instead of load and store. Thus, both the implementation and the specification of the stack allow it to have non-SC behaviour.

To summarise, very roughly, release-acquire allows writes to be delayed but not reordered, while relaxed allows both. Relaxed actions produce even stranger behaviour, including what we call ***satisfaction cycles***:

SATISFACTION CYCLE (SCL):

```
a=load_RLX(&x) ‖ b=load_RLX(&y)    load_RLX(&x,8)   load_RLX(&y,8)
if (a==8)      ‖ if (b==8)            sb |    rf  rf   | sb
  store_RLX(&y,a) ‖ store_RLX(&x,8)  store_RLX(&y,8)  store_RLX(&x,8)
```

Here, each conditional satisfies its guard from a later write in the other thread. This is possible because relaxed reads and writes do not create any happens-before ordering, and thus neither read is constrained by (RD). Unlike relaxed, release-acquire does not allow satisfaction cycles. If the loads and stores in the example were annotated release-acquire, then both rf edges would also be hb edges. This would produce an hb cycle, which is prohibited by the memory model. Satisfaction cycles are known to be a problematic aspect of the C11 model; as we show in this paper, they also create difficulties for library abstraction.

3. Library Abstraction Informally

Histories. Our approach to abstraction is based on the notion of a ***history***, which concisely records all interactions between the client and the library in a given execution. Clients and libraries can affect

---

We also omit some C11 subtleties that are orthogonal to abstraction (see §4).

one another in several ways in C11. Most straightforwardly, the library can observe the parameters passed by the client at calls, and the client can observe the library's return values. Therefore, a history includes the set of all call and return actions in the execution. However, clients can also observe synchronisation and other memory-model effects inside a method. These more subtle interactions are recorded by two kinds of partial order: *guarantees* and *denies*.

Synchronisation internal to the library can affect the client by forcing reads to read from particular writes. For example, in (MP) from §2, the client is forced to read 1 from x because the push and pop methods synchronise internally in a way that generates an hb ordering between the call to push and the return from pop. If the methods did not hb-synchronise, the client could read from either of the writes to x. The client can thus observe the difference between two library implementations with different internal synchronisation, even if all call and return values are identical. To account for this, the ***guarantee*** relation in a history of an execution records hb edges between library call and return actions.

Even non-synchronising behaviour inside the library can sometimes be observed by the client. For example, the C11 model requires the existence of a total order mo over all atomic writes to a given location. This order cannot go against hb, but is not included into it, as this would make the model much stronger, and would hinder efficient compilation onto very weak architectures, such as Power and ARM [20]. Now, consider the following:

DENY (DN):

```
                                    ⋮            rf,hb (forbidden!)
                              load_ACQ(&x,0) ←- - - - - - - -
                                    sb |                      ⋮
    store_REL(&x,1);            call lib()                    ⋮
                              - - - - -|      mo            library
load_ACQ(&x); ‖ lib();          library |    ↗  |  - - - - -
lib();        ‖ store_REL(&x,0);        |       ↓    ret lib()
                                        ↓          sb |      ⋮
                                             store_REL(&x,0)
```

In this execution, a write internal to the invocation of lib in the second thread is mo-ordered after a write internal to the invocation of lib in the first thread. This forbids the client from reading 0 from x. To see this, suppose the contrary holds. Then the ACQ load synchronises with the REL store of 0, yielding an hb edge. By transitivity with the client sb edges, which are included in hb, we get an hb edge from ret lib in the second thread to call lib in the first. Together with the library's sb edges, this yields an hb edge going against the library-internal mo one, which is prohibited by the memory model. To account for such interactions, the ***deny*** relations in a history of an execution record hb or other kinds of edges between return and call actions that the client *cannot* enforce due to the structure of the library, e.g., the hb-edge from ret lib to call lib above.

Abstraction in the presence of relaxed atomics. As we noted in §1, we actually propose two library abstraction criteria: one for the full memory model described in §2, and one for programs without relaxed atomics. We discuss the former first.

Two library executions with the same history are observationally equivalent to clients, even if the executions are produced by different library implementations. By defining a sound abstraction relation over histories, we can therefore establish abstraction between libraries. To this end, we need to compare the histories of libraries under every client context. Fortunately, we need not examine every possible client: it suffices to consider behaviour under a ***most general client***, whose threads repeatedly invoke library methods in any order and with any parameters. Executions under this client generate all possible histories of the library, and thus represent all client-library interactions (with an important caveat,

discussed below). We write $[\![\mathcal{L}]\!]I$ for the set of executions of the library $\mathcal{L}$ under the most general client starting from an ***initial state*** I. Initial states are defined formally in §4, but, informally, record initialisation actions such as the ones shown in (SB). The set $[\![\mathcal{L}]\!]I$ gives the denotation of the library considered in isolation from its clients, and in this sense, defines a ***library-local semantics*** of $\mathcal{L}$.

This library-local semantics allows us to define library abstraction. We now quote its definition and formulate the corresponding Abstraction Theorem, introducing some of the concepts used in them only informally. This lets us highlight their most important features that can be discussed independently of the formalities. We fill in the missing details in §6.1, after we have presented the C11 model more fully.

For the memory model with relaxed atomics, a history contains one guarantee and one deny relation.

DEFINITION 1. *A **history** is a triple $H = (A, G, D)$, where A is a set of call and return actions, and $G, D \subseteq A \times A$.*

Library abstraction is defined on pairs of libraries $\mathcal{L}$ and sets of their initial states $\mathcal{I}$. It relies on a function $\mathsf{history}(\cdot)$ extracting the history from a given execution of a library, which we lift to sets of executions pointwise. The notation $[\![\mathcal{L}, R]\!]$ and the notion of safety are explained below.

DEFINITION 2. *For histories (A_1, G_1, D_1) and (A_2, G_2, D_2), we let $(A_1, G_1, D_1) \sqsubseteq (A_2, G_2, D_2)$ if $A_1 = A_2$, $G_1 = G_2$ and $D_2 \subseteq D_1$.*

*For safe $(\mathcal{L}_1, \mathcal{I}_1)$ and $(\mathcal{L}_2, \mathcal{I}_2)$, $(\mathcal{L}_1, \mathcal{I}_1)$ **is abstracted by** $(\mathcal{L}_2, \mathcal{I}_2)$, written $(\mathcal{L}_1, \mathcal{I}_1) \sqsubseteq (\mathcal{L}_2, \mathcal{I}_2)$, if for any relation R containing only edges from return actions to call actions, we have*

$$\forall I_1 \in \mathcal{I}_1, H_1 \in \mathsf{history}([\![\mathcal{L}_1, R]\!]I_1).$$
$$\exists I_2 \in \mathcal{I}_2, H_2 \in \mathsf{history}([\![\mathcal{L}_2, R]\!]I_2). \, H_1 \sqsubseteq H_2.$$

The overall shape of the definition is similar to that of linearizability on SC memory models [10]: any behaviour of the concrete library $\mathcal{L}_1$ relevant to the client has to be reproducible by the abstract library $\mathcal{L}_2$. However, there are several things to note.

First, we allow the execution of the abstract library to deny less than the concrete one, but require it to provide the same guarantee. Intuitively, we can strengthen the deny, because this only allows more client behaviours.

Second, we do not consider raw executions of the most general client of $\mathcal{L}_1$ and $\mathcal{L}_2$, but those whose happens-before relation can be ***extended*** with an arbitrary set R of edges between return and call actions without contradicting the axioms of the memory model; $[\![\mathcal{L}_1, R]\!]I_1$ and $[\![\mathcal{L}_2, R]\!]I_2$ denote the sets of all such extensions. The set R represents the happens-before edges that can be enforced by the client: such happens-before edges are not generated by the most general client and, in the presence of relaxed atomics, have to be considered explicitly (this is the caveat to its generality referred to above). We consider only return-to-call edges, as these are the ones that represent synchronisation inside the client (similarly to how call-to-return edges in the guarantee represent synchronisation inside the library; cf. (MP)). The definition requires that, if an extension of the concrete library is consistent with R, then so must be the matching execution of the abstract one.

Finally, the abstraction relation is defined only between ***safe*** libraries that do not access locations internal to the client and do not have faults, such as data races.

As we show in §7, the specification of the Treiber stack in Figure 1a abstracts its implementation in Figure 1b.

Abstraction theorem. We now formulate a theorem that states the correctness of our library abstraction criterion. We consider programs $\mathcal{C}(\mathcal{L})$ with a single client $\mathcal{C}$ and a library $\mathcal{L}$ (the case of multiple libraries is considered in §6.1). The Abstraction Theorem

states that, if we replace the (implementation) library $\mathcal{L}_1$ in a program $\mathcal{C}(\mathcal{L}_1)$ with another (specification) library $\mathcal{L}_2$ abstracting $\mathcal{L}_1$, then the set of client behaviours can only increase. Hence, when reasoning about $\mathcal{C}(\mathcal{L}_1)$, we can soundly replace $\mathcal{L}_1$ with $\mathcal{L}_2$ to simplify reasoning.

In the theorem, $[\![\mathcal{C}(\mathcal{L})]\!]\mathcal{I}$ gives the set of executions of $\mathcal{C}(\mathcal{L})$ from initial states in a set $\mathcal{I}$, $\uplus$ combines the initial states of a client and a library, and $\mathsf{client}(\cdot)$ selects the parts of executions generated by client commands. We call $(\mathcal{C}(\mathcal{L}), \mathcal{I})$ ***non-interfering***, if $\mathcal{C}$ and $\mathcal{L}$ do not access each others' internal memory locations in executions of $\mathcal{C}(\mathcal{L})$ from initial states in $\mathcal{I}$. The notion of safety for $\mathcal{C}(\mathcal{L})$ is analogous to the one for libraries.

THEOREM 3 (Abstraction). *Assume that $(\mathcal{L}_1, \mathcal{I}_1)$, $(\mathcal{L}_2, \mathcal{I}_2)$, $(\mathcal{C}(\mathcal{L}_2), \mathcal{I} \uplus \mathcal{I}_2)$ are safe, $(\mathcal{C}(\mathcal{L}_1), \mathcal{I} \uplus \mathcal{I}_1)$ is non-interfering and $(\mathcal{L}_1, \mathcal{I}_1) \sqsubseteq (\mathcal{L}_2, \mathcal{I}_2)$. Then $(\mathcal{C}(\mathcal{L}_1), \mathcal{I} \uplus \mathcal{I}_1)$ is safe and*

$$\mathsf{client}([\![\mathcal{C}(\mathcal{L}_1)]\!](\mathcal{I} \uplus \mathcal{I}_1)) \subseteq \mathsf{client}([\![\mathcal{C}(\mathcal{L}_2)]\!](\mathcal{I} \uplus \mathcal{I}_2)).$$

The requirement of non-interference is crucial, because it ensures that clients can only observe library behaviour through return values and memory-model effects, rather than by 'opening the box' and observing internal states. The drawback of Theorem 3 is that it requires us to establish the non-interference between the client $\mathcal{C}$ and the *concrete* library $\mathcal{L}_1$, e.g., via a type system or a program logic proof. As we show below, we cannot weaken this condition to allow checking non-interference on the client $\mathcal{C}$ using the *abstract* library $\mathcal{L}_2$, as is standard in data refinement on SC memory models [7]. This makes the reasoning principle given by the theorem less compositional, since establishing non-interference requires considering the composed behaviour of the client and the concrete library—precisely what library abstraction is intended to avoid! However, this does not kill compositional reasoning completely, as non-interference is often simple to check even globally. We can also soundly check other aspects of safety, such as data-race freedom, on $\mathcal{C}(\mathcal{L}_2)$. Furthermore, as we show in §6.1, the notion of library abstraction given by Definition 2 is compositional for non-interfering libraries. As we now explain, we can get the desired theorem allowing us to check non-interference on $\mathcal{C}(\mathcal{L}_2)$ for the fragment of the language excluding relaxed atomics.

Abstraction without relaxed atomics. Restricting ourselves to programs without the relaxed memory order (and augmenting the axiomatic memory model to allow incomplete program runs, as described in §4) allows strengthening our result in three ways:

1. We no longer need to quantify over client happens-before edges R, like in Definition 2. Instead, we enrich histories with an additional deny relation, which is easier to deal with in practice than the quantification. Hence, without relaxed atomics, the caveat to the generality of the most general client does not apply.

2. Abstraction on histories can be defined by inclusion on guarantees, rather than by equality.

3. We no longer need to show that the unabstracted program $\mathcal{C}(\mathcal{L}_1)$ is non-interfering. Rather, non-interference is a consequence of the safety of the abstracted program $\mathcal{C}(\mathcal{L}_2)$.

The first two differences make proofs of library abstraction slightly easier, but are largely incidental otherwise. In particular, quantification over client happens-before edges in Definition 2, although unpleasant, does not make library abstraction proofs drastically more complicated. Requiring the guarantees of the concrete and abstract executions to be equal in this definition just results in more verbose specifications in certain cases. In contrast, the last difference is substantial.

The price of satisfaction cycles. For each of the three above differences we have a counterexample showing that Theorem 3 will

not hold if we change the corresponding condition to the one required in the case without relaxed atomics. All of these counterexamples involve satisfaction cycles, which can only be produced by relaxed atomics. Our results show that this language feature makes the reasoning principles for C11 programs less compositional. Due to space constraints, here we present only the counterexample for point 3 above; the others are given in §D. In §6.3, we identify the corresponding place in the proof of the Abstraction Theorem for the language without relaxed atomics where we rely on the absence of satisfaction cycles.

Consider the following pair of libraries $\mathcal{L}_1$ and $\mathcal{L}_2$:

```
𝓛₁:  atomic int x;          𝓛₂:  atomic int x;
     int m() {                   int m() {
       store_RLX(&x,42);           return 42;
       return load_RLX(&x);      }
     }
```

Here x is a library-internal location. We have $\mathcal{L}_1 \sqsubseteq \mathcal{L}_2$, since both method implementations behave exactly the same, assuming that the internal location x is not modified outside the library. Unsafe clients can distinguish between $\mathcal{L}_1$ and $\mathcal{L}_2$. For example, the client

$$\text{print m(); } \| \text{ store}_{\text{RLX}}(\&x,0);$$

can print 0 when using $\mathcal{L}_1$, but not $\mathcal{L}_2$. However, any non-trivial library behaves erratically when clients corrupt its private data structures, and thus, it is reasonable for abstraction to take into account only well-behaved clients that do not do this. We therefore contend that $\mathcal{L}_1$ *should* be abstracted by $\mathcal{L}_2$ according to any sensible abstraction criterion.

The above misbehaved client violates non-interference when using either $\mathcal{L}_1$ or $\mathcal{L}_2$. However, we can define a more complicated client $\mathcal{C}$ such that $\mathcal{C}(\mathcal{L}_2)$ is non-interfering, but $\mathcal{C}(\mathcal{L}_1)$ is not:

```
a=m()          ‖ b=load_RLX(&y)
if (a==0)      ‖ if (b==0)
 store_RLX(&y,0)‖  store_RLX(&x,0)
```


The execution of $\mathcal{C}(\mathcal{L}_1)$ given on the right violates non-interference due to a satisfaction cycle: a fault in the client causes the library to misbehave by returning 0 instead of 42, and the effect of this misbehaviour is what causes the client fault in the first place! Since the abstract library $\mathcal{L}_2$ is completely resilient to client interference, its method will always return 42, and thus, the satisfaction cycle will not appear and the client will not access the variable x. Note that this counterexample is not specific to our notion of library abstraction: *any* such notion for C11 considering $\mathcal{L}_2$ to be a specification for $\mathcal{L}_1$ cannot allow checking non-interference using $\mathcal{L}_2$.

For expository reasons, we have given a very simple counterexample. This program would be easy to detect and eliminate, e.g., using a simple type system: one *syntactic* path in the client is guaranteed to result in the forbidden access to the library's internal state. However, the same kind of behaviour can occur with dynamically-computed addresses: the client stepping out of bounds of an array overwrites the library state, causing it to misbehave, which in turn causes the original client misbehaviour. For this kind of example, proving non-interference becomes non-trivial.

It is unclear to us whether satisfaction cycles are observable in practice. They are disallowed by even the weakest C11 target architectures, such as Power and ARM [20], because these architectures respect read-to-write dependencies created by control-flow. It is also clear that the C11 language designers would like to for-

bid satisfaction cycles: the C++11 standard [11, Section 29.3, Paragraph 11] states that, although (SCL) from §2 is permitted, "implementations should not allow such behaviour". This apparent contradiction is because certain compiler optimisations, such as common subexpression elimination and hoisting from loops, can potentially create satisfaction cycles (see [15] for discussion). Since avoiding them would require compilers to perform additional analysis and/or limit optimisations, the standard does not disallow satisfaction cycles outright. Our results provide an extra argument that allowing satisfaction cycles is undesirable.

4. C11: Language and Thread-Local Semantics

To define the C11 model, we use a lightly modified version of its formalisation proposed by Batty et al. [1]. In the interests of simplicity, we consider a simple core language, instead of the full C/C++, and omit some of the features of the memory model. We do not handle two categories of features: those that are orthogonal to abstraction, and more esoteric features that would complicate our results. In the first category we have dynamic memory allocation, dynamic thread creation, blocked CASes and non-atomic initialisation of atomic locations (we also do not present the treatment of locks here, although we handle a bounded number of them in our formal development; see §A). In the second category we have memory fences and release-consume atomics, discussed in §9.

The semantics of a C11 program is given by a set of executions, generated in two stages. The first stage, described in this section, generates a set of *action structures* using a sequential thread-local semantics which takes into account only the structure of every thread's statements, not the semantics of memory operations. In particular, the values of reads are chosen arbitrarily, without regard for writes that have taken place. The second stage, described in §5, filters out the action structures that are inconsistent with the C11 memory model. It does this by constructing additional relations and checking the resulting executions against the axioms of the model.

Programming language. We assume that the memory consists of locations from Loc containing values in Val. We assume a function sort : Loc $\rightarrow$ {ATOM, NA}, showing whether memory accesses to a given location must be atomic (ATOM) or non-atomic (NA); see §2. The program syntax is as follows:

$$
\begin{aligned}
C &::= \text{skip} \mid c \mid m \mid C; C \mid \text{if}(x)\, \{C\}\, \text{else}\, \{C\} \mid \\
&\quad \text{while}(x)\, \{C\} \mid \text{atom\_sec}\, \{c\} \\
\mathcal{L} &::= \{m = C_m \mid m \in M\} \\
\mathcal{C}(\mathcal{L}) &::= \text{let}\, \mathcal{L}\, \text{in}\, C_1 \parallel \ldots \parallel C_n
\end{aligned}
$$

A program consists of a **library** $\mathcal{L}$ implementing methods $m \in$ Method and its **client** $\mathcal{C}$, given by a parallel composition of threads $C_1, \ldots, C_n$. The commands include skip, an arbitrary set of base commands $c \in$ BComm (e.g., atomic and non-atomic loads and stores, and CASes), method calls $m \in$ Method, sequential composition, branching on the value of a location $x \in$ Loc and loops. Our language also includes atomic sections, ensuring that a base command executes atomically. Atomic sections are not part of C/C++, but are used here to express library specifications. We assume that every method called by the client is defined in the library, and we disallow nested method calls.

We assume that every method accepts a single parameter and returns a single value. Parameters and return values are passed by every thread via distinguished locations in memory, denoted $\text{param}_t, \text{retval}_t \in$ Loc for $t = 1..n$, such that $\text{sort}(\text{param}_t) = \text{sort}(\text{retval}_t) = \text{NA}$. The rest of memory locations are partitioned into those owned by the client (CLoc) and the library (LLoc):

$$\text{Loc} = \text{CLoc} \uplus \text{LLoc} \uplus \{\text{param}_t, \text{retval}_t \mid t = 1..n\}.$$

The property of non-interference introduced in §3 then requires that a library or a client access only the memory locations belonging

to them (except the ones used for passing parameters and return values). We provide pointers on how we can relax the requirement of static address space partitioning in §9.

Actions. Executions are composed of ***actions***, defined as follows:

$$\lambda, \mu \in \mathsf{MemOrd} ::= \mathsf{NA} \mid \mathsf{SC} \mid \mathsf{ACQ} \mid \mathsf{REL} \mid \mathsf{RLX}$$
$$\varphi \in \mathsf{Effect} ::= \mathsf{store}_\lambda(x, a) \mid \mathsf{load}_\lambda(x, a)$$
$$\mid \mathsf{rmw}_{\lambda,\mu}(x, a, b) \mid \mathsf{call}\ m(a) \mid \mathsf{ret}\ m(a)$$
$$u, v, w, q, r \in \mathsf{Act} ::= (e, g, t, \varphi)$$

Here $e \in \mathsf{AId}$ is a unique ***action identifier***, and λ, μ are memory orders (§2) of memory accesses. Every instance of an atomic section occurring in an execution has a unique identifier $g \in \mathsf{SectId}$. Atomic sections only have force when multiple actions have the same section identifier, so actions outside any section are simply assigned a unique identifier each. The domains of the rest of the variables are as follows: $t \in \{0, \ldots, n\}$, $x \in \mathsf{Loc}$, $a, b \in \mathsf{Val}$. We allow actions by a dummy thread 0 to denote memory initialisation. We only consider actions whose memory orders respect location sorts given by sort, and we do not allow rmw actions of sort NA.

Loading or storing a value a at a location x generates the obvious action. A read-modify-write action $(e, g, t, \mathsf{rmw}_{\lambda,\mu}(x, a, b))$ arises from a successful compare-and-swap command. It corresponds to reading the value a from the location x and atomically overwriting it with the value b; λ and μ give the memory orders of the read and the write, respectively, and have to be different from NA. The value a in $(e, g, t, \mathsf{call}\ m(a))$ or $(e, g, t, \mathsf{ret}\ m(a))$ records the parameter param_t or the return value retval_t passed between the library method and its client. We refer to call and return actions as ***interface actions***.

For an action u we write $\mathsf{sec}(u)$ for its atomic section identifier, and we denote the set of all countable sets of actions by $\mathcal{P}(\mathsf{Act})$. We omit e, g and λ annotations from actions when they are irrelevant. We also write _ for an expression whose value is irrelevant. We use $(t, \mathsf{read}_\lambda(x, a))$ to mean any of the following:

$$(t, \mathsf{load}_\lambda(x, a)); \quad (t, \mathsf{rmw}_{\lambda,\_}(x, a, \_));$$
$$(t, \mathsf{call}\ \_(a)), \text{ if } x = \mathsf{param}_t \text{ and } \lambda = \mathsf{NA};$$
$$(t, \mathsf{ret}\ \_(a)), \text{ if } x = \mathsf{retval}_t \text{ and } \lambda = \mathsf{NA}.$$

We use $(t, \mathsf{write}_\lambda(x, a))$ to mean $(t, \mathsf{store}_\lambda(x, a))$ or $(t, \mathsf{rmw}_{\_,\lambda}(x, \_, a))$. We call the two classes of actions ***read actions*** and ***write actions***, respectively.

Thread-local semantics. The thread-local semantics generates a set of ***action structures***—triples $(A, \mathsf{sb}, \mathsf{ib})$, where $A \in \mathcal{P}(\mathsf{Act})$, and $\mathsf{sb}, \mathsf{ib} \subseteq A \times A$ are the sequenced-before and initialised-before relations introduced in §2. We assume that sb is transitive and irreflexive and relates actions by the same thread; ib relates initialisation actions with thread identifier 0 to the others. We do not require sb to be a total order: in C/C++, the order of executing certain program constructs is unspecified.

For a base command $c \in \mathsf{BComm}$, we assume a set $\langle c \rangle_t \in \mathcal{P}(\mathcal{P}(A) \times \mathcal{P}(A \times A))$ of all pairs of action sets and sb relations that c produces when executed by a thread t (the ib relations are missing, as they are relevant only for a whole program). Note that base commands may include conditionals and loops, and thus can give rise to an arbitrary number of actions; we give a separate explicit semantics to conditionals and loops only because they are used in the most general client in §6.1. Definitions of $\langle c \rangle_t$ for sample base commands c are given in Figure 2. Note that, in the thread-local semantics, a read from memory, such as $\mathsf{load}_\mu(y)$ in the figure, yields an arbitrary value. A CAS command generates an rmw action, if successful, and a load, otherwise.

We define an ***initial state*** of a program $\mathcal{C}(\mathcal{L})$ by a function

$$I \in (\mathsf{LLoc} \uplus \mathsf{CLoc}) \rightharpoonup_{fin} (\mathsf{Val} \times \mathsf{MemOrd}),$$

giving the initial values of memory locations, together with the memory orders of initial writes to them. We define the set of action structures $\langle\langle \mathcal{C}(\mathcal{L}) \rangle\rangle I$ of a program $\mathcal{C}(\mathcal{L})$ in Figure 3. Note that this set of action structures corresponds to complete runs of $\mathcal{C}(\mathcal{L})$. The clause for atom_sec $\{c\}$ assigns the same atomic section identifier to all actions generated by c. The clause for a call to a method m brackets structures generated by its implementation C_m with call and return actions. We have omitted the clause for loops (§B). For simplicity, we assume that the variable in the condition of a branch is always non-atomic.

Assumptions. We make several straightforward assumptions about the structures in $\langle c \rangle_t$ for $c \in \mathsf{BComm}$:

- Structures in $\langle c \rangle_t$ are finite and contain only load, store and read-modify-write actions by t with unique action identifiers.

- For any $(A, \mathsf{sb}) \in \langle c \rangle_t$, sb is transitive and irreflexive.

- Structures in $\langle c \rangle_t$ are insensitive to the choice of action and atomic section identifiers: applying any bijection to such identifiers from a structure in $\langle c \rangle_t$ produces another structure in $\langle c \rangle_t$.

- Atomic sections in every $(A, \mathsf{sb}) \in \langle c \rangle_t$ are contiguous in sb:

$$\forall u, v, q.\ \mathsf{sec}(u) = \mathsf{sec}(v) \wedge u \xrightarrow{\mathsf{sb}} q \xrightarrow{\mathsf{sb}} v \implies \mathsf{sec}(q) = \mathsf{sec}(u).$$

So as not to obfuscate presentation, we only consider programs $\mathcal{C}(\mathcal{L})$ that use param_t and retval_t correctly. We assume that in any action structure of a program, only library actions by t read param_t and write to retval_t, and only client actions by t read retval_t and write to param_t. We also require that param_t and retval_t be initialised before an access: in any structure $(A, \mathsf{sb}, \mathsf{ib})$ of $\mathcal{C}(\mathcal{L})$,

$$(\forall u = (t, \mathsf{call}\ \_) \in A.\ \exists w = (t, \mathsf{write}(\mathsf{param}_t, \_)).\ w \xrightarrow{\mathsf{sb}} u) \wedge$$
$$(\forall u = (t, \mathsf{ret}\ \_) \in A.\ \exists w = (t, \mathsf{write}(\mathsf{retval}_t, \_)).\ w \xrightarrow{\mathsf{sb}} u).$$

Additional assumptions without relaxed atomics. For our result without relaxed atomics (§6.2), we currently require the following additional assumptions:

- The structure set $\langle c \rangle_t$ accounts for c fetching any values from the memory locations it reads (see §B for formalisation).

- Any structure of a base command c inside an atomic section accesses at most one atomic location. This is sufficient for our purposes, since a library specification usually accesses a single such location containing the abstract state of the library.

- We modify the standard C11 model by requiring that a program's semantics include structures corresponding to execution prefixes. In the standard C11 model, all executions are complete (although possibly infinite). We define $\langle C_i \rangle_t^p$ for a thread C_i by

$$\langle C_i \rangle_t^p = \{(A_p, \mathsf{sb}_p) \mid \exists (A, \mathsf{sb}) \in \langle C_i \rangle_t.$$
$$A_p \subseteq A \wedge \forall u \in A, v \in A_p.\ u \xrightarrow{\mathsf{sb}} v \implies u \in A_p\}.$$

This is necessary due to the interaction between our prototype atomic section semantics and the C11 model. It weakens the notion of atomicity: atomic sections at the end of a prefix may be partially executed, and therefore more weakly ordered than their completed counterparts. Eliminating it will require a deeper understanding of the relationship between atomicity and the notion of incomplete program runs. See §6.3 for its use in the proof.

5. C11: Axiomatic Memory Model

The axiomatic portion of the C11 model takes a set of action structures of the program, generated by the thread-local semantics in §4, and filters out those which are inconsistent with the memory model. To formulate it, we enrich action structures with extra relations.

Executions. The semantics of a program consists of a set of *executions*, $X = (A, \mathsf{sb}, \mathsf{ib}, \mathsf{rf}, \mathsf{sc}, \mathsf{mo}, \mathsf{sw}, \mathsf{hb})$, where $A \in \mathcal{P}(\mathsf{Act})$ and the rest are relations on A:

- sb, ib, rf and hb introduced in §4; rf is such that its reverse is a partial function from read to write actions on the same location and with the same value.

- *sequentially consistent order* (sc), ordering SC reads from and writes to the same location. The projection of sc to each atomic location is a transitive, irreflexive total order, while writes to distinct locations are unrelated[3].

- *modification order* (mo), ordering writes (but not reads!) to the same atomic (i.e., of the sort ATOM) location. The projection of mo to each atomic location is a transitive, irreflexive total order.

- *synchronises-with* (sw), defining synchronisation.

We write $\mathsf{r}(X)$ for the component r of the execution X.

We can now define the denotation $[\![\mathcal{C}(\mathcal{L})]\!]$ of a program and the notion of safety and non-interference introduced informally in §3. An execution X is *valid*, if it satisfies the validity axioms shown in Figure 4; it is *safe*, if it satisfies the safety axioms in Figure 5, and it is *non-interfering* if it satisfies the NONINTERF axiom from Figure 5. We explain the axioms shown in the rest of this section. Intuitively, validity axioms correspond to properties that are enforced by the runtime, while safety axioms correspond to properties that the programmer must ensure to avoid faults. To simplify the following explanations, Figures 4 and 5 do not show axioms dealing with CASes and locks. To keep the presentation tractable, we have also omitted some corner cases from SWDEF and SCREADS in Figure 4. The missing axioms and cases are given in §A, and our results are established for the memory model including these (the correctness of the stack in Figure 1 actually relies on a corner case in SWDEF).

For a program $\mathcal{C}(\mathcal{L})$ and an initial state I, we let $[\![\mathcal{C}(\mathcal{L})]\!]I$ be the set of valid executions X, whether safe or not, such $(A, \mathsf{sb}(X), \mathsf{ib}(X)) \in \langle\mathcal{C}(\mathcal{L})\rangle I$. We write $[\![\mathcal{C}(\mathcal{L})]\!]\mathcal{I}$ to stand for its obvious lifting to sets $\mathcal{I}$ of initial states. A program $\mathcal{C}(\mathcal{L})$ is *safe* when run from I if every one of its valid executions is (and similarly for non-interference and sets of initial states). An unsafe program has undefined behaviour.

The validity axioms define sw and hb directly in terms of the other relations (SWDEF and HBDEF). The hb relation is constructed from the sb, ib and sw, and as follows from ACYCLICITY, has to be irreflexive. The $/\!\sim$ operator in the definition of hb is needed to handle atomic sections; for now the reader should ignore it. The sw relation is derived from sc and rf. The rf, sc and mo relations are only constrained by the axioms, not defined directly. We explain the validity axioms by first considering a language fragment with non-atomic memory accesses only and then gradually expanding it to include the other memory orders.

Non-atomic memory accesses. The values read by non-atomic reads are constrained by DETREAD and RFNONATOMIC. DETREAD requires every read to have an associated rf edge when the location read was previously initialised, i.e., when there is a write to it that happened before the read. Executions with reads from uninitialised locations are valid, but, as we explain below, unsafe. RFNONATOMIC requires that a read only reads from the write to the same location immediately preceding it in hb; cf. (RD). In the absence of other synchronisation, this means that a thread can read only from its own previous writes or initial values, since by HBDEF, $\mathsf{sb} \cup \mathsf{ib} \subseteq \mathsf{hb}$. Threads can establish the necessary syn-

chronisation using atomic operations (which we explain now) or locks (which we elide here; see §A).

SC atomics. The strong semantics of SC actions is enforced by organising all SC reads and writes over a *single* location into a total order sc, which cannot form a cycle with hb (ACYCLICITY). According to SCREADS, an SC read can only read from the closest sc-preceding write. Thus, if all memory accesses are annotated as SC in (SB) from §2, the result shown there is forbidden. Indeed, by SCREADS and ACYCLICITY, the store of 1 to y has to follow the load of 0 from y in sc, and similarly for x. This yields a cycle in $\mathsf{sb} \cup \mathsf{sc}$, contradicting ACYCLICITY. Note that the model requires the existence of sc, but does not include all of it into hb. As a consequence, one cannot use the ordering of, say, two SC reads in sc to constrain the values that can be read according to RFNONATOMIC.

By SWDEF, an rf edge between an SC write and an SC read generates an sw edge, which is then included into hb by HBDEF. Release-acquire atomics have the same effect, as we now explain.

Release-acquire atomics. By SWDEF, an ACQ read synchronises with the REL write it reads from. For example, if in (SCL) we annotated all writes with REL and all reads with ACQ, then the rf edges would be included into hb and the execution would be prohibited by ACYCLICITY.

For atomics weaker than SC, there is no total order on all operations over a given location analogous to sc; this is why (SB) is allowed. Instead, they satisfy a weaker property of *coherence*: all writes (but not reads) to a single atomic location are organised into a total modification order mo, which has to be consistent with hb (HBvsMO). SC writes to the location are also included into mo, and in such cases the latter has to be consistent with sc (MOvsSC).

Since reads are not included into mo, we do not have an analogue of SCREADS, and thus, a read has more freedom to choose which write it reads from. The only constraints on atomic accesses weaker than SC are given by coherence axioms—COWR, CORW and CORR. For example, COWR says that a read r that happened after a write w_2 cannot read from a write w_1 earlier in mo.

Relaxed atomics. Like release-acquire atomics, relaxed atomics respect coherence, given by the mo order and the axioms COWR, CORW and CORR. However, rf edges involving them do not generate synchronisation edges sw. The only additional constraint on relaxed reads is given by RFATOMIC, which prohibits reads 'from the future', i.e., from writes later in hb; cf. (RD). This and the fact that coherence axioms enforce no constraints on actions over distinct locations allows (SCL). If all the loads and stores in (SCL) were to the same location, it would be forbidden by CORW.

Safety axioms. The safety axioms in Figure 5 define the conditions under which a program is faulty. DRF constrains pairs of actions over the same location, with at least one write. It requires that such pairs on distinct threads, one of which is a non-atomic access, are related by hb, and on the same thread, by sb (recall that in C/C++, the order of executing certain program constructs is unspecified, and thus, sb is partial). SAFEREAD prohibits reads from uninitialised locations.

The NONINTERF axiom is not part of the C11 memory model, but formalises the property of non-interference required for our results to hold (§3); it is technically convenient for us to consider it together with the other safety axioms. NONINTERF requires that the library and the client only read from and write to the locations they own, except param_t and retval_t used for communication (§4). The axiom classifies an action as performed by the library or the client depending on its position in sb with respect to calls and returns.

Atomic sections. Atomic sections are a widespread idiom for defining library specifications. In an SC memory model, we can

---

[3] In the original C11 model [1], sc is a total order on SC operations over all locations. The formulation here is is equivalent to the original one (§C), but more convenient for defining library abstraction.

$$\langle \mathsf{store}_\lambda(x, \mathsf{load}_\mu(y)) \rangle_t = \{(\{u, v\}, \{(u, v)\}) \mid \exists a', e_1, e_2, g_1, g_2.\, e_1 \neq e_2 \wedge g_1 \neq g_2 \wedge$$
$$u = (e_1, g_1, t, \mathsf{load}_\mu(y, a')) \wedge v = (e_2, g_2, t, \mathsf{store}_\lambda(x, a'))\}$$

$$\langle *y = \mathsf{CAS}_{\lambda,\mu}(x, a, b) \rangle_t = \{(\{u, v\}, \{(u, v)\}) \mid \exists e_1, e_2, g_1, g_2, a'.\, e_1 \neq e_2 \wedge g_1 \neq g_2 \wedge a' \neq a \wedge$$
$$(u = (e_1, g_1, t, \mathsf{rmw}_{\lambda,\mu}(x, a, b)) \wedge v = (e_2, g_2, t, \mathsf{store}_{\mathsf{NA}}(y, 1))) \vee (u = (e_1, g_1, t, \mathsf{load}_\lambda(x, a')) \wedge v = (e_2, g_2, t, \mathsf{store}_{\mathsf{NA}}(y, 0)))\}$$

Figure 2. Definitions of $\langle c \rangle_t$ for sample base commands. Here $x, y \in \mathsf{Loc}$ and $a, b \in \mathsf{Val}$ are constants.

$$\langle \mathsf{skip} \rangle_t = \{(\emptyset, \emptyset)\}$$
$$\langle C_1; C_2 \rangle_t = \{(A_1 \uplus A_2, \mathsf{sb}_1 \cup \mathsf{sb}_2 \cup \{(u, v) \mid u \in A_1 \wedge v \in A_2\}) \mid (A_1, \mathsf{sb}_1) \in \langle C_1 \rangle_t \wedge (A_2, \mathsf{sb}_2) \in \langle C_2 \rangle_t\}$$
$$\langle \mathsf{if}(x)\,\{C_1\}\,\mathsf{else}\,\{C_2\} \rangle_t = \{(\{u\} \uplus A, \mathsf{sb} \cup \{(u, v) \mid v \in A\}) \mid \exists a.\, (A, \mathsf{sb}) \in \langle C_1 \rangle_t \wedge u = (\_, \_, t, \mathsf{load}_{\mathsf{NA}}(x, a)) \wedge a \neq 0\} \cup$$
$$\{(\{u\} \uplus A, \mathsf{sb} \cup \{(u, v) \mid v \in A\}) \mid (A, \mathsf{sb}) \in \langle C_2 \rangle_t \wedge u = (\_, \_, t, \mathsf{load}_{\mathsf{NA}}(x, 0))\}$$
$$\langle \mathsf{atom\_sec}\,\{C\} \rangle_t = \{(\{(e, g, t, \varphi) \mid (e, \_, t, \varphi) \in A\}, \{((e_1, g, t, \varphi_1), (e_2, g, t, \varphi_2)) \mid$$
$$((e_1, \_, t, \varphi_1), (e_2, \_, t, \varphi_2)) \in \mathsf{sb}\}) \mid (A, \mathsf{sb}) \in \langle C \rangle_t \wedge g \in \mathsf{SectId}\}$$
$$\langle m \rangle_t = \{(A \cup \{u\} \cup \{v\}, \mathsf{sb} \cup \{(u, v)\} \cup \{(u, q), (q, v) \mid q \in A\}) \mid (A, \mathsf{sb}) \in \langle C_m \rangle_t \wedge u = (\_, \_, t, \mathsf{call}\, m(\_)) \wedge v = (\_, \_, t, \mathsf{ret}\, m(\_))\}$$
$$\langle \mathsf{let}\,\{m = C_m \mid m \in M\}\,\mathsf{in}\,C_1 \parallel \ldots \parallel C_n \rangle I = \{(A_0 \uplus (\bigcup_{t=1}^n A_t), \bigcup_{t=1}^n \mathsf{sb}_t, (A_0 \times (\bigcup_{t=1}^n A_t))) \mid (\forall t = 1..n.\, (A_t, \mathsf{sb}_t) \in \langle C_t \rangle_t) \wedge$$
$$(\forall t = 1..n.\, \forall u.\, \exists\, \mathsf{finitely\ many}\, v.\, (v, u) \in \mathsf{sb}_t) \wedge (A_0 = \bigcup\{(e, g, 0, \mathsf{store}_\lambda(x, a)) \mid I(x) = (a, \lambda) \wedge e \in \mathsf{AId} \wedge g \in \mathsf{SectId}\})\}$$

Figure 3. Thread-local semantics. $A \uplus B$ is the union of the sets of actions A and B with disjoint sets of action and atomic section identifiers.

HBDEF. $\mathsf{hb} = ((\mathsf{sb} \cup \mathsf{ib} \cup \mathsf{sw})/\sim)^+$, where
$$R/\sim = R \cup \{(u, v) \mid \mathsf{sec}(u) \neq \mathsf{sec}(v) \wedge \exists u', v'.\, \mathsf{sec}(u) = \mathsf{sec}(u') \wedge \mathsf{sec}(v) = \mathsf{sec}(v') \wedge u' \xrightarrow{R} v'\}$$

SWDEF$^\approx$. $\forall w, r.\, w \xrightarrow{\mathsf{sw}} r \iff \begin{pmatrix} \exists t_1, t_2, \lambda, \mu, x.\, t_1 \neq t_2 \wedge \lambda \in \{\mathsf{SC}, \mathsf{REL}\} \wedge \mu \in \{\mathsf{SC}, \mathsf{ACQ}\} \\ \wedge\, w = (t_1, \mathsf{write}_\lambda(x, \_)) \wedge r = (t_2, \mathsf{read}_\mu(x, \_)) \wedge w \xrightarrow{\mathsf{rf}} r \end{pmatrix}$ ACYCLICITY. $\mathsf{hb} \cup \mathsf{sc}$ is acyclic

DETREAD. $\forall r.\, (\exists x, w'.\, w' \xrightarrow{\mathsf{hb}} r \wedge w' = (\_, \mathsf{write}(x, \_)) \wedge r = (\_, \mathsf{read}(x, \_))) \iff (\exists w.\, w \xrightarrow{\mathsf{rf}} r)$ RFATOMIC.

RFNONATOMIC. $\forall w, r, x.\, w \xrightarrow{\mathsf{rf}} r \wedge w = (\_, \mathsf{write}(x, \_)) \wedge r = (\_, \mathsf{read}(x, \_)) \wedge \mathsf{sort}(x) = \mathsf{NA}$
$$\implies w \xrightarrow{\mathsf{hb}} r \wedge \neg\exists w'.\, w' = (\_, \mathsf{write}(x, \_)) \wedge w \xrightarrow{\mathsf{hb}} w' \xrightarrow{\mathsf{hb}} r \qquad \neg\exists r, w.\quad r \underset{\mathsf{rf}}{\overset{\mathsf{hb}}{\rightleftarrows}} w$$

SCREADS$^\approx$. $\forall w, r.\, w \xrightarrow{\mathsf{rf}} r \wedge r = (\_, \mathsf{read}_{\mathsf{SC}}(x, \_)) \wedge w = (\_, \mathsf{write}_{\mathsf{SC}}(x, \_)) \implies w \xrightarrow{\mathsf{sc}} r \wedge \neg\exists w'.\, w' = (\_, \mathsf{write}(x, \_)) \wedge w \xrightarrow{\mathsf{sc}} w' \xrightarrow{\mathsf{sc}} r$

HBVSMO. $\neg\exists w_1, w_2.$ MOVSSC. $\neg\exists w_1, w_2.$ COWR. $\neg\exists w_1, w_2.$ CORW. $\neg\exists r, w_1, w_2.$ CORR. $\neg\exists r_1, r_2, w_1, w_2.$

$$w_1 \underset{\mathsf{mo}}{\overset{\mathsf{hb}}{\rightleftarrows}} w_2 \qquad w_1 \underset{\mathsf{sc}}{\overset{\mathsf{mo}}{\rightleftarrows}} w_2 \qquad \begin{matrix} w_1 \xrightarrow{\mathsf{mo}} w_2 \\ \;\;\searrow_{\mathsf{rf}} \quad \downarrow{\mathsf{hb}} \\ r \end{matrix} \qquad \begin{matrix} w_2 \xrightarrow{\mathsf{rf}} r \\ \nwarrow \quad \downarrow{\mathsf{hb}} \\ \mathsf{mo}\;\; w_1 \end{matrix} \qquad \begin{matrix} w_1 \xrightarrow{\mathsf{rf}} r_1 \\ \mathsf{mo}\uparrow \qquad \downarrow{\mathsf{hb}} \\ w_2 \xrightarrow{\mathsf{rf}} r_2 \end{matrix}$$

ASMO. $\forall u, v.\, u \xrightarrow{\mathsf{mo}} v \wedge \mathsf{sec}(u) = \mathsf{sec}(v) \implies \neg\exists q.\, u \xrightarrow{\mathsf{mo}} q \xrightarrow{\mathsf{mo}} v \wedge \mathsf{sec}(u) \neq \mathsf{sec}(q)$
ASSC. $\forall u, v.\, u \xrightarrow{\mathsf{sc}} v \wedge \mathsf{sec}(u) = \mathsf{sec}(v) \implies \neg\exists q.\, u \xrightarrow{\mathsf{sc}} q \xrightarrow{\mathsf{sc}} v \wedge \mathsf{sec}(u) \neq \mathsf{sec}(q)$

Figure 4. Selected validity axioms of the C11 memory model. Axioms simplified for the purposes of presentation are marked by $\approx$.

DRF. $\forall u, v, x, t_1, t_2.\, (u, v \in A \wedge u \neq v \wedge u = (t_1, \_(x, \_)) \wedge v = (t_2, \_(x, \_)) \wedge (u = (t_1, \mathsf{write}(x, \_)) \vee v = (t_2, \mathsf{write}(x, \_)))) \implies$
$$((t_1 \neq t_2 \implies (u \xrightarrow{\mathsf{hb}} v \vee v \xrightarrow{\mathsf{hb}} u \vee \mathsf{sort}(x) = \mathsf{ATOM})) \wedge (t_1 = t_2 \implies (u \xrightarrow{\mathsf{sb}} v \vee v \xrightarrow{\mathsf{sb}} u)))$$

SAFEREAD. $\qquad\qquad\qquad\qquad \forall r.\, r \in A \wedge r = (\_, \mathsf{read}(\_, \_)) \implies \exists w.\, w \xrightarrow{\mathsf{rf}} r$

NONINTERF. $\qquad\qquad \forall u, x, t.\, (u \in A \wedge t \neq 0 \wedge u = (t, \_(x, \_)) \wedge x \notin \{\mathsf{param}_t, \mathsf{retval}_t \mid t = 1..n\}) \implies$
$$((\exists v.\, v = (\_, \mathsf{call}\, \_) \wedge v \xrightarrow{\mathsf{sb}} u \wedge \neg\exists q.\, q = (\_, \mathsf{ret}\, \_) \wedge v \xrightarrow{\mathsf{sb}} q \xrightarrow{\mathsf{sb}} u) \iff (x \in \mathsf{LLoc}))$$

Figure 5. Selected safety axioms of the C11 memory model

DETREAD$_l$. $\qquad \forall r, t.\, ((\exists x.\, r = (\_, \mathsf{read}(x, \_)) \wedge x \neq \mathsf{param}_t \wedge \exists w.\, w \xrightarrow{\mathsf{hb}} r \wedge w = (\_, \mathsf{write}(x, \_))) \iff \exists w'.\, w' \xrightarrow{\mathsf{rf}} r) \wedge$
$$(\forall a, b.\, r = (t, \mathsf{read}(\mathsf{param}_t, a)) \in A \wedge u = (\_, \mathsf{call}\, \_(b)) \wedge u \xrightarrow{\mathsf{sb}} r \wedge (\neg\exists v.\, v = (\_, \mathsf{ret}\, \_) \wedge u \xrightarrow{\mathsf{sb}} v \xrightarrow{\mathsf{sb}} r) \implies a = b)$$

SAFEREAD$_l$. $\qquad\qquad\qquad \forall r, x, t.\, r \in A \wedge r = (t, \mathsf{read}(x, \_)) \wedge x \neq \mathsf{param}_t \implies \exists w.\, w \xrightarrow{\mathsf{rf}} r$

DETREAD$_c$. $\qquad \forall r, t.\, ((\exists x.\, r = (\_, \mathsf{read}(x, \_)) \wedge x \neq \mathsf{retval}_t \wedge \exists w.\, w \xrightarrow{\mathsf{hb}} r \wedge w = (\_, \mathsf{write}(x, \_))) \iff \exists w'.\, w' \xrightarrow{\mathsf{rf}} r) \wedge$
$$(\forall a, b.\, r = (t, \mathsf{read}(\mathsf{retval}_t, a)) \wedge u = (\_, \mathsf{ret}\, \_(b)) \wedge u \xrightarrow{\mathsf{sb}} r \wedge (\neg\exists v.\, v = (\_, \mathsf{call}\, \_) \wedge u \xrightarrow{\mathsf{sb}} v \xrightarrow{\mathsf{sb}} r) \implies a = b)$$

SAFEREAD$_c$. $\qquad\qquad\qquad \forall r, x, t.\, (r \in A \wedge r = (t, \mathsf{read}(x, \_)) \wedge x \neq \mathsf{retval}_t \implies \exists w.\, w \xrightarrow{\mathsf{rf}} r) \wedge$
$$(r \in A \wedge r = (t, \mathsf{read}(\mathsf{retval}_t, \_)) \wedge r \neq (\_, \mathsf{ret}\, \_) \implies \exists u.\, u \xrightarrow{\mathsf{hb}} r \wedge u = (t, \mathsf{ret}\, \_))$$

Figure 6. Axioms for library (§6.1) and client (§6.3) executions

define their semantics by simply requiring that no other events are interleaved with actions inside an atomic section. Unfortunately, the relaxed memory model of C11 does not admit such a simple definition. The straightforward solution of imposing a total order on all instances of atomic sections would rule out relaxed specifications that we would like to give, such as the Treiber specification from §2. Hence, we have extended C11 with a prototype notion of atomic sections suitable for its relaxed-memory setting (inspired by the semantics of transactions in [5]). This notion represents only the first step towards a natural specification language for relaxed C11, which is an interesting problem in itself.

The axioms defining the semantics of atomic sections are HB-DEF, ASMO and ASSC in Figure 4 and ATOMAS in Figure 7 (deferred to §A for brevity). They capture the expected properties of atomicity. Thus, in HBDEF, we factor $sb \cup ib \cup sw$ over atomic sections using $/\sim$: e.g., if an action u happens-before another action v, then u also happens-before any other action from the same atomic section as v. ASMO and ASSC require that actions from the same atomic section be contiguous in mo and sc. ATOMAS constrains relaxed actions, which do not generate hb edges. ASMO, ASSC and ATOMAS are trivially satisfied when every action has a unique atomic section identifier. Additionally, in this case HBDEF simplifies to $hb = (sb \cup ib \cup sw)^+$, which is how it is defined in standard C11 [1]. Thus, if every action executes in a separate atomic section, our augmented model coincides with standard C11.

6. Library Abstraction in Detail

We first define formally the concepts used in the definition of library abstraction (Definition 2) and the Abstraction Theorem (Theorem 3) from §3 for the memory model with relaxed atomics. We then show how the Abstraction Theorem can be strengthened for a fragment of the language excluding them (§6.2) and give the proof outlines for both theorems (§6.3).

6.1 Library Abstraction in the Presence of Relaxed Atomics

History definition. We formally define the history function, which selects a history in the sense of Definition 2 from a library execution. For an execution X, we let

$$\text{history}(X) = (\text{interf}(X), \text{hbL}(X), \text{scL}(X))$$

and lift history to sets of executions pointwise. Here $\text{interf}(X)$ is the projection of $A(X)$ to interface (call and return) actions. The hbL selector computes the guarantee part of the history. We let $\text{hbL}(X)$ be the projection of $\text{hb}(X)$ to pairs of actions of the form $((\_, \text{call } \_), (\_, \text{ret } \_))$ and pairs of calls and returns (in any order) by the same thread. We record only edges of the above form, since it can be shown that any happens-before edge between interface actions in a library execution under its most general client can be obtained as a transitive closure of such edges. Intuitively, call-to-return edges are the ones that represent the synchronisation between library method invocations, as illustrated by (MP) in §2.

The scL selector computes the deny part of the history. We let $\text{scL}(X)$ be the projection of $((\text{hb}(X) \cup \text{sc}(X))^+)^{-1}$ to pairs of actions of the form $((\_, \text{ret } \_), (\_, \text{call } \_))$. This component is needed, since the ACYCLICITY axiom (Figure 4) mandates that sc cannot form a cycle with hb, but does not include sc into hb. Thus, when a library relates a call action u to a return action v with hb and sc, the client cannot relate v to u with the same relations, as this would invalidate ACYCLICITY. We add only return-to-call edges into $\text{scL}(X)$, as these are the edges that represent synchronisation inside the client (similarly to how call-to-return edges represent synchronisation inside the library in $\text{hbL}(X)$). One might think that the deny component of the history should have included edges recording potential violations of other similar axioms, e.g., HB-VSMO, as suggested by (DN). However, in the case of the model

with relaxed atomics, we are forced to quantify over client happens-before edges R in Definition 2. As it happens, this makes it unnecessary to consider axioms other than ACYCLICITY (see the proof of the Theorem 3 in §C). These axioms, however, have to be taken into account in the case without relaxed atomics (§6.2).

Library-local semantics. We define the most general client as follows. Take $n \geq 1$ and let $\{m_1, \ldots, m_l\}$ be the methods implemented by a library $\mathcal{L}$. We let

$$\text{MGC}_n(\mathcal{L}) = (\text{let } \mathcal{L} \text{ in } C_1^{\text{mgc}} \parallel \ldots \parallel C_n^{\text{mgc}}),$$

where C_t^{mgc} is

$$\text{while}(\text{nondet}()) \{ \quad \text{if}(\text{nondet}()) \{m_1\}$$
$$\text{else if}(\text{nondet}()) \{m_2\} \ldots \text{else} \{m_l\} \}$$

Here we use the obvious generalisation of loops and conditionals to branch expressions that yield a non-deterministic value. To allow parameters of methods to be chosen arbitrarily, we replace the axioms DETREAD from Figure 4 and SAFEREAD from Figure 5 by DETREAD_l and SAFEREAD_l from Figure 6, which mandate that reads from param_t lack associated rf edges, while nonetheless yielding identical values within a single method call. As part of the proof of the Theorem 3 (§6.3), we show that this client is indeed most general in a certain formal sense (with the caveat concerning the need to extend its executions with client happens-before edges mentioned in §3).

We note that some libraries require their clients to pass only certain combinations of parameters or issue only certain sequences of method calls. Such contracts could be accommodated in our framework by restricting the most general client appropriately; we do not handle them here so as not to complicate the presentation.

For an initial library state $I \in \text{LLoc} \rightharpoonup_{fin} \text{Val} \times \text{MemOrd}$, a *library execution* of $\mathcal{L}$ from I is an execution from $[\![\text{MGC}_n(\mathcal{L})]\!]I$ for some $n \geq 1$. A library execution is *valid* if it satisfies the validity axioms with DETREAD_l instead of DETREAD; it is *safe* if it satisfies the safety axioms with SAFEREAD_l instead of SAFEREAD. We let $[\![\mathcal{L}]\!]I$ be the set of all valid library executions of $\mathcal{L}$ from I and lift $[\![\mathcal{L}]\!]$ to sets of initial states pointwise. We say that a library $\mathcal{L}$ is safe when run from I if so is every execution in $[\![\mathcal{L}]\!]I$; for a set $\mathcal{I}$ of initial states, $(\mathcal{L}, \mathcal{I})$ is safe if $\mathcal{L}$ is safe when run from any $I \in \mathcal{I}$. The notion of a *non-interfering* library is defined similarly.

Extended executions. In Definition 2, we use library executions whose happens-before relation is extended with extra edges recording constraints enforced by the client. Consider an execution X and a relation R over interface actions from $A(X)$. The *extension* of X with R is an execution that has the same components as X, except the happens-before relation is replaced by $(\text{hb}(X) \cup R)^+$. An extension of a library execution with R is *admissible* when it satisfies the corresponding validity axioms, but with HBDEF replaced by

EXTHBDEF. $\text{hb} = ((\text{sb} \cup \text{ib} \cup \text{sw} \cup R)/\sim)^+$.

For an initial library state I, we let $[\![\mathcal{L}, R]\!]I$ be the set of admissible executions of $\mathcal{L}$ from I extended with R. This completes the definition of components used in Definition 2.

Execution projections. Finally, we define the client function used in Theorem 3. Consider a valid execution X of $\mathcal{C}(\mathcal{L})$. An action $u \in A$ is a *library action*, if it is a call or return action, an action of the form $(0, \text{write}(x, \_))$ for $x \in \text{LLoc}$, or if

$$\exists v. v = (\_, \text{call } \_) \wedge v \xrightarrow{\text{sb}(X)} u \wedge$$
$$\neg \exists q. q = (\_, \text{ret } \_) \wedge v \xrightarrow{\text{sb}(X)} q \xrightarrow{\text{sb}(X)} u.$$

An action $u \in A$ is a *client action*, if it is a call or return action, it is an action of the form $(0, \text{write}(x, \_))$ for $x \in \text{CLoc}$, or the negation of the above property holds. We define the execution $\text{lib}(X)$ by

restricting the action set to library actions and projecting all the relations in X accordingly. We use a similar projection $\text{client}(X)$ to client actions and lift client and lib to sets of executions pointwise.

Properties of library abstraction. For the fragment of the language with an SC semantics—i.e., allowing only non-atomic memory accesses and SC atomics—Definition 2 implies classical linearizability. This follows from Theorem 3 and the fact that classical linearizability is equivalent to observational abstraction on the SC memory model [6]. However, the converse is not true: since our notion of library abstraction validates Theorem 3 for clients in the full C11, it distinguishes between SC libraries that classical linearizability would consider equivalent (see §D).

Using Theorem 3, we can obtain the expected property that, like the classical notion of linearizability, our notion of library abstraction is compositional (with a caveat that non-interference among libraries has to be checked globally). Formally, consider libraries $\mathcal{L}_1, \ldots, \mathcal{L}_k$ with disjoint sets of declared methods and assume the splitting of the library address space into regions belonging to each library: $\text{LLoc} = \text{LLoc}_1 \uplus \ldots \uplus \text{LLoc}_k$. Consider sets of initial states $\mathcal{I}_1, \ldots, \mathcal{I}_k$ such that $\forall j = 1..k. \forall I \in \mathcal{I}_j. \text{dom}(I) \subseteq \text{LLoc}_j$. We adjust the notion of library safety so that NONINTERF for $\mathcal{L}_j$ is checked with respect to locations in LLoc_j. Let $(\mathcal{L}'_1, \mathcal{I}'_1), \ldots, (\mathcal{L}'_k, \mathcal{I}'_k)$ be corresponding library specifications. We define $\mathcal{L}$, respectively, $\mathcal{L}'$ as the library implementing all methods of $\mathcal{L}_1, \ldots, \mathcal{L}_k$, respectively, $\mathcal{L}'_1, \ldots, \mathcal{L}'_k$ and having the set of initial states $\mathcal{I}_1 \uplus \ldots \uplus \mathcal{I}_k$, respectively, $\mathcal{I}'_1 \uplus \ldots \uplus \mathcal{I}'_k$. We assume that any combination of implementations or specifications of different libraries is non-interfering. The following theorem is shown by abstracting $\mathcal{L}_j$ to $\mathcal{L}'_j$ one by one using Theorem 3.

THEOREM 4. *If $(\mathcal{L}_j, \mathcal{I}_j) \sqsubseteq (\mathcal{L}'_j, \mathcal{I}'_j)$ for $j = 1..k$, then $\mathcal{L} \sqsubseteq \mathcal{L}'$.*

6.2 Library Abstraction without Relaxed Atomics

We call an action with at least one RLX annotation ***relaxed***. In this section, we restrict ourselves to programs whose action structures do not have any relaxed actions, and we augment the C11 thread-local semantics as described in the assumptions of §4. Among other things, these changes allow us to remove the quantification over client happens-before edges R from Definition 2, at the expense of including an additional deny relation into the history.

DEFINITION 5. *An **extended history** is a quadruple (A, G, D, D'), where A is a set of interface actions, and $G, D, D' \subseteq A \times A$.*

For an execution $X = (A, \text{sb}, \text{ib}, \text{rf}, \text{sc}, \text{mo}, \text{sw}, \text{hb})$, we let

$$\text{Ehistory}(X) = (\text{interf}(X), \text{hbL}(X), \text{scL}(X), \text{denyL}(X)),$$

which we lift to sets of executions pointwise. The selectors interf, hbL and scL are defined as in §6.1.

The relation $\text{denyL}(X)$ is defined similarly to scL, but whereas the latter records client hb and sc edges that can violate ACYCLICITY, the former includes the client hb edges that can violate the axioms HBVSMO, COWR and SCREADS (Figure 4). We do not have to consider other similar axioms, such as RFATOMIC, CORW and CORR since in any valid execution X without relaxed actions, we have $\text{rf}(X) \subseteq \text{hb}(X)$. Due to this, the hb client edges violating, HBVSMO and COWR, CORW and CORR can be covered by the same relations. For the case of HBVSMO and COWR, $\text{denyL}(X)$ includes the dashed edges of all possible instantiations of the following diagrams:

where $u = (\_, \text{ret } \_)$ and $v = (\_, \text{call } \_)$. This records client hb edges that violate HBVSMO or COWR (cf. (DN) in §3). The diagrams are thus constructed systematically by 'breaking' the hb edge. For the case of SCREADS, $\text{denyL}(X)$ includes edges corresponding to a corner case in this axiom omitted from Figure 4. The full version of the axiom and the corresponding deny diagram are given in §A.

DEFINITION 6. *We let $(A_1, G_1, D_1, D'_1) \preceq (A_2, G_2, D_2, D'_2)$, if $A_1 = A_2$, $G_2 \subseteq G_1$, $D_2 \subseteq D_1$ and $D'_2 \subseteq D'_1$.*

*For safe $(\mathcal{L}_1, \mathcal{I}_1)$ and $(\mathcal{L}_2, \mathcal{I}_2)$, $(\mathcal{L}_1, \mathcal{I}_1)$ **is abstracted by** $(\mathcal{L}_2, \mathcal{I}_2)$, written $(\mathcal{L}_1, \mathcal{I}_1) \preceq (\mathcal{L}_2, \mathcal{I}_2)$, if*

$$\forall I_1 \in \mathcal{I}_1, H_1 \in \text{Ehistory}(\llbracket \mathcal{L}_1 \rrbracket I_1).$$
$$\exists I_2 \in \mathcal{I}_2, H_2 \in \text{Ehistory}(\llbracket \mathcal{L}_2 \rrbracket I_2). H_1 \preceq H_2.$$

Unlike in Definition 2, here an abstracted history can guarantee *fewer* happens-before edges to the client: without relaxed atomics, removing edges from happens-before can only permit more client behaviours or make the client unsafe. We note that checking the inclusion between the components of the history given by denyL, required by Definition 6, is simpler in practice than quantifying over client happens-before edges R in Definition 2.

For executions X and Y, we write $X \preceq Y$ when all their components except hb are equal, and $\text{hb}(Y) \subseteq \text{hb}(X)$.

THEOREM 7 (Abstraction without relaxed atomics). *Assume that $(\mathcal{L}_1, \mathcal{I}_1)$, $(\mathcal{L}_2, \mathcal{I}_2)$ and $(\mathcal{C}(\mathcal{L}_2), \mathcal{I} \uplus \mathcal{I}_2)$ are safe and $(\mathcal{L}_1, \mathcal{I}_1) \preceq (\mathcal{L}_2, \mathcal{I}_2)$. Then $(\mathcal{C}(\mathcal{L}_1), \mathcal{I} \uplus \mathcal{I}_1)$ is safe and*

$$\forall X \in \text{client}(\llbracket \mathcal{C}(\mathcal{L}_1) \rrbracket (\mathcal{I} \uplus \mathcal{I}_1)).$$
$$\exists Y \in \text{client}(\llbracket \mathcal{C}(\mathcal{L}_2) \rrbracket (\mathcal{I} \uplus \mathcal{I}_2)). X \preceq Y.$$

Unlike Theorem 3, this one allows the programmer to check non-interference on $\mathcal{C}(\mathcal{L}_2)$—i.e., with respect to a library specification—and to conclude that $\mathcal{C}(\mathcal{L}_1)$ is non-interfering. Since the abstract library can have a smaller guarantee than the concrete one, the happens-before of the execution $\mathcal{C}(\mathcal{L}_2)$ may also be smaller than that of the execution $\mathcal{C}(\mathcal{L}_1)$.

Like the notion of library abstraction from §6.1, the one proposed here implies classical linearizability for the SC fragment of the language (but not vice versa) and is compositional (§C). However, here the latter property does not require us to check non-interference globally: a composition of several non-interfering libraries is non-interfering.

6.3 Proof Outlines

Client-local semantics. We start by defining the ***client-local semantics*** of a client $\mathcal{C} = C_1 \parallel \ldots \parallel C_n$, which is a counterpart of the library-local semantics defined in §6.1. Let M be the set of methods that can be called by $\mathcal{C}$. Consider the program

$$\mathcal{C}(\cdot) = (\text{let } \{m = \text{skip} \mid m \in M\} \text{ in } C_1 \parallel \ldots \parallel C_n),$$

where every method is implemented by a stub that returns immediately after having been called. Moreover, we allow methods to return arbitrary values by replacing the axioms DETREAD from Figure 4 and SAFEREAD from Figure 5 by DETREAD$_c$ and SAFEREAD$_c$ from Figure 6. We call executions of the above program ***client executions*** of $\mathcal{C}$. A client execution is ***valid***, if it satisfies the validity axioms with DETREAD$_c$ instead of DETREAD; it is ***safe***, if it satisfies the safety axioms with SAFEREAD$_c$ instead of SAFEREAD. For an initial client state $I \in \text{CLoc} \rightharpoonup_{fin} (\text{Val} \times \text{MemOrd})$, let $\llbracket \mathcal{C} \rrbracket I$ be the set of all valid client executions of $\mathcal{C}$ from I; for a set $\mathcal{I}$ of initial states we define $\llbracket \mathcal{C} \rrbracket \mathcal{I}$ as expected. The notion of an extended client execution is similar to the one of an extended library execution from §6.1. For an extended execution

X we let $\mathsf{core}(X)$ be the execution obtained from X by recomputing $\mathsf{hb}(X)$ from $\mathsf{sb}(X)$, $\mathsf{ib}(X)$ and $\mathsf{sw}(X)$ according to HBDEF.

Client-side history selectors. Consider a client execution X. We define $\mathsf{hbC}(X), \mathsf{scC}(X) \subseteq \mathsf{interf}(X) \times \mathsf{interf}(X)$, which are analogous to hbL and scL from §6.1, but select the information about the client part of the execution that is relevant to the library. We let $\mathsf{hbC}(X)$ be the projection of $\mathsf{hb}(X)$ to pairs of actions of the form $((\_, \mathsf{ret}\ \_), (\_, \mathsf{call}\ \_))$ and pairs of calls and returns (in any order) by the same thread. We select edges of this form as they are the ones that record synchronisation enforced by the client. We let $\mathsf{scC}(X)$ be the projection of $((\mathsf{hb}(X) \cup \mathsf{sc}(X))^+)^{-1}$ to pairs of actions of the form $((\_, \mathsf{call}\ \_), (\_, \mathsf{ret}\ \_))$.

Proof outline for Theorem 3. Consider an execution X of $\mathcal{C}(\mathcal{L}_1)$ from the initial state $I \uplus I_1$, where $I \in \mathcal{I}$ and $I_1 \in \mathcal{I}_1$. We start by decomposing X into a client execution $\mathsf{client}(X)$ and a library execution $\mathsf{lib}(X)$ and showing that

$$\mathsf{client}(X) \in [\![\mathcal{C}, \mathsf{hbL}(\mathsf{core}(\mathsf{lib}(X)))]\!]I;$$
$$\mathsf{lib}(X) \in [\![\mathcal{L}_1, \mathsf{hbC}(\mathsf{core}(\mathsf{client}(X)))]\!]I_1.$$

The second inclusion justifies that the most general client of the library defined in §6.1 is indeed most general, as it can reproduce the behaviour of $\mathcal{L}_1$ under any client $\mathcal{C}$. This comes with the caveat that an execution of the most general client of $\mathcal{L}_1$ has to be extended with $\mathsf{hbL}(\mathsf{core}(\mathsf{lib}(X)))$ to obtain $\mathsf{lib}(X)$, as the library-local semantics does not generate happens-before edges enforced by client synchronisation (and similarly for $\mathsf{client}(X)$ in the first inclusion). The above decomposition step relies on X being non-interfering. Using the fact that $(\mathcal{L}_1, \mathcal{I}_1) \sqsubseteq (\mathcal{L}_2, \mathcal{I}_2)$, we prove that there exist $I_2 \in \mathcal{I}_2$ and $Y \in [\![\mathcal{L}_2, \mathsf{hbC}(\mathsf{core}(\mathsf{client}(X)))]\!]I_2$ such that $\mathsf{history}(\mathsf{lib}(X)) \sqsubseteq \mathsf{history}(Y)$. Here we use the quantification of client happens-before edges R in Definition 2 to handle the extension of the library execution with $\mathsf{hbC}(\mathsf{core}(\mathsf{client}(X)))$. We then compose the executions X and Y into the desired execution of $\mathcal{C}(\mathcal{L}_2)$. This step uses the fact that the deny component of $\mathsf{history}(Y)$ is smaller than that of $\mathsf{history}(\mathsf{lib}(X))$. $\qquad\square$

Proof outline for Theorem 7. Unlike in Theorem 3, here we need to consider the case when an execution X of $\mathcal{C}(\mathcal{L}_1)$ has actions violating non-interference. In this case, we identify an "earliest" faulting action u and construct a valid execution that is a prefix of X ending just before u and is thus non-interfering. This is only possible because, without relaxed atomics, we do not have satisfaction cycles, and because the theorem is stated over an augmented thread-local semantics (§4), with prefix executions added to the semantics. We convert the resulting execution into one of $\mathcal{C}(\mathcal{L}_2)$ as in the proof of Theorem 3 and conjoin the action the action u to it. This yields an execution of $\mathcal{C}(\mathcal{L}_2)$ violating non-interference and contradicting the assumption about its safety.

In the case when X is non-interfering, the proof is similar to that of Theorem 3. Some additional work is needed to deal with the fact that Definition 6 does not have a quantification over client happens-before edges and allows the abstract guarantee to be smaller than the concrete one. $\qquad\square$

7. Establishing Library Abstraction

In this section, we discuss the proof process for establishing abstraction between libraries in the sense of Definitions 2 and 6. To reason about programs on the C11 memory model, we use axiomatisations of the action structures generated by them, which give a simple mathematical interface to the program semantics. To prove library safety, required by Definitions 2 and 6, we consider all the execution shapes of the most general client, and check these against the C11 safety axioms. We now explain how we prove the correspondence between the executions of concrete and abstract libraries, using Definition 2 for illustration.

Effect points. Consider an execution $X_1 \in [\![\mathcal{L}_1, R]\!]I_1$. We construct an execution $X_2 \in [\![\mathcal{L}_2, R]\!]I_2$ whose history witnesses the existential in Definition 2 using an adapted version of the *linearization point* method for proving linearizability. The method constructs the abstract execution by calling a method specification at a fixed linearization point in every method invocation of the concrete execution; intuitively, it is at this point that the concrete method 'takes effect'. In our adaptation, we construct X_2 by substituting calls to library method implementations in X_1 for the corresponding calls to specifications, and choosing appropriate values for reads and different orders between actions. These values and orders are chosen based on the orders over actions we call **effect points**, picked for each concrete method invocation in X_1. Thus, the various partial orders over the effect points in the concrete execution dictate the order of precedence for the effects of method invocations in the abstract execution. In contrast with the original linearization point method, the latter order does not have to be linear. We now explain this technique in more detail on an example.

Example: Treiber stack. We have proved the correctness of the stack in Figure 1b with respect to its specification in Figure 1a according to Definition 2 (full details are given in §E). As effect points in X_1, we pick the rmw actions modifying the stack's top pointer T, which correspond to successful CASes (this is the same choice as the that of linearization points when proving the linearizability of Treiber's stack on an SC memory model). The order $\mathsf{mo}(X_1)$ over these actions defines a total order over successful push and pop invocations. When substituting invocations of method implementations for invocations of specifications, we use $\mathsf{mo}(X_1)$ to decide $\mathsf{rf}(X_2)$ and $\mathsf{mo}(X_2)$ for the abstract location S. Namely, suppose we have two method invocations U and V in X_1, such that the rmw of U is the immediate predecessor to the rmw of V in $\mathsf{mo}(X_1)$. When substituting corresponding invocations of specification methods U' and V', we set up $\mathsf{rf}(X_2)$ and $\mathsf{mo}(X_2)$ so that the load from S in V' reads from the rmw in U', and $\mathsf{mo}(X_2)$ orders the rmw on S in U' right before rmw on it in V'. Fixing $\mathsf{rf}(X_2)$ and $\mathsf{mo}(X_2)$ in the abstract execution immediately fixes the values of reads, which finishes the construction of X_2.

Example: producer-consumer queue. We have similarly verified a non-blocking producer-consumer queue according to Definition 6 (see §E). The queue is intended for communication between a single producer thread and a single consumer thread and provides three methods: init, enq and deq. The implementation stores values in a finite cyclic array, while the specification stores them using the abstract data type of a sequence. To ensure that the consumer calling deq observes up-to-date values, it must synchronise with the producer calling enq using release-acquire atomics.

8. Related Work

Relaxed-memory behaviour has become widespread in real concurrent systems. As a result, some algorithm designers have begun to publish algorithms with memory-model annotations such as fences [2, 16, 17]. However, the corresponding formulations of correctness properties and their proofs are generally informal. While there has been some work on formally verifying programs on weak memory models [13, 19, 22], none has proposed a compositional reasoning method, like we do. Ours is the first approach that formulates the notion of a correct library specification and provides a method for establishing it on C11, or any similarly relaxed model.

Our work is an evolution of *linearizability* [10], a correctness criterion that has been widely adopted in the concurrent algorithms community. In the SC fragment of C11, our definition of library abstraction implies classical linearizability. History abstraction in classical linearizability is defined by linearization of the partial order over non-overlapping methods invocations, and the guarantee

portion in our histories can be seen as lifting this to the C11 setting. Classical linearizability has no equivalent to our deny relation; the conflicts between relations that are captured by deny do not occur in an SC setting where all events are totally ordered.

Recent work has formalised the intuition that linearizability corresponds to observational abstraction [6] and has extended it to handle liveness [8], resource-transferring programs [9] and the x86 memory model [4]. The latter work is the closest to this paper; in particular, we borrow the decompose-compose approach in the proof of the Abstraction Theorem (§6.3) from it. However, while its objective is the same as ours—abstraction for relaxed-memory concurrent libraries—the technical challenges and the machinery developed to address them are very different. The x86 memory model can be defined by a small-step operational semantics, which has an underlying total order on abstract machine memory events [18]. Linearizability for x86 is therefore a relatively mild extension to classical linearizability, which simply represents some of these events in (linear) histories. In contrast, C11 constrains relaxed behaviour through partial ordering, controlled by an axiomatic semantics. There are no abstract machine events to linearize, which motivates our novel definition of a history as a set of partial orders and history abstraction as inclusion over them. Furthermore, x86 is a substantially stronger model than C11, with (SB) from §2 being the only significant relaxation [18]. Our approach is the first technique for specifying client-visible effects of relaxations inside libraries on weaker memory models.

9. Conclusion and Future Work

We have proposed the first sound criterion for library abstraction suitable for the C11 memory model and demonstrated its practicality on two small, but typical, relaxed libraries. Our criterion is certainly complex, but much of this complexity arises from the real-world intricacies of the C11 memory model. In turn, the complexity of the model arises from a multitude of target platforms of C and C++—two of the world's most widely-used programming languages. Despite this complexity, the criterion allows developers to establish that C11 libraries satisfy simple, reusable specifications that precisely describe the level of consistency guaranteed. This is an essential ingredient for supporting modular development of complex software on relaxed memory models.

In addition, our approach is the first compositional reasoning technique for an axiomatically defined relaxed memory model, and highlights general principles for abstraction on such models. We have good reason to believe that our techniques can be reused for other memory models, as the conditions for library abstraction fall out naturally from obligations arising when trying to prove the Abstraction Theorem according to the approach in §6.3. In particular, the histories in our approach are constructed uniformly, with deny relations obtained straightforwardly from axioms by 'breaking' hb edges (§6.2). In particular, our preliminary investigations show that these techniques can be used to define a notion of library abstraction for an axiomatic formulation of the x86 memory model [18].

Our specifications describe precisely the level of synchronisation provided by the library, although in some cases this makes them more verbose. This is motivated by the fact that libraries on C11 can offer relaxed interfaces to clients, without either giving up all synchronisation guarantees or enforcing sequential consistency. If information about the internal synchronisation used to ensure library correctness were not described in these interfaces, clients would have to duplicate it, thereby decreasing performance. On the other hand, requiring library interfaces to be SC would rule out libraries that use weaker memory orders to achieve efficiency while preserving basic correctness properties, again decreasing performance. Our prototype atomic section semantics represents a first attempt at a syntactic specification idiom for relaxed algorithms,

albeit with limitations as described in §4. We leave a more comprehensive treatment of relaxed atomic sections to future work.

Our two formulations of library abstraction (Definitions 2 and 6) and the Abstraction Theorem (Theorems 3 and 7) identify the feature of the current C/C++ memory model that does not allow fully compositional reasoning about libraries. As we argued in §3, this deficiency is not specific to our definition of library abstraction, but would be inherent to *any* sensible one. We hope that these insights will inform future revisions of the C/C++ memory model.

Our development omits memory fences and release-consume atomics, which are the more advanced features of the C11 memory model. A memory fence is a synchronisation construct that affects many memory actions, rather than just one. Release-consume is a special-purpose memory order which compiles more efficiently to Power and ARM processors. We conjecture that our methods can be used to handle these features. As both of them generate more possible client-library interactions, this will require adding additional relations to histories.

To concentrate on the core challenges of library abstraction in C11, we assumed that the data structures of the client and its libraries are completely disjoint (§4). We hope to lift this restriction by combining our results with a recent generalisation of classical linearizability allowing transfers of memory ownership [9]. Similarly, a previous generalisation of linearizability to handle liveness properties [8] could be used to strengthen specifications of the kind shown in Figure 1a to guarantee properties such as lock-freedom.

References

[1] M. Batty, S. Owens, S. Sarkar, P. Sewell, and T. Weber. Mathematizing C++ concurrency. In *POPL*, 2011.

[2] H.-J. Boehm. Can seqlocks get along with programming language memory models? In *MSPC*, 2012.

[3] H.-J. Boehm and S. V. Adve. Foundations of the C++ concurrency memory model. In *PLDI*, 2008.

[4] S. Burckhardt, A. Gotsman, M. Musuvathi, and H. Yang. Concurrent library correctness on the TSO memory model. In *ESOP*, 2012.

[5] S. Burckhardt, D. Leijen, M. Fähndrich, and M. Sagiv. Eventually consistent transactions. In *ESOP*, 2012.

[6] I. Filipović, P. O'Hearn, N. Rinetzky, and H. Yang. Abstraction for concurrent objects. In *ESOP*, 2009.

[7] I. Filipović, P. O'Hearn, N. Torp-Smith, and H. Yang. Blaiming the client: On data refinement in the presence of pointers. *FAC*, 22, 2010.

[8] A. Gotsman and H. Yang. Liveness-preserving atomicity abstraction. In *ICALP*, 2011.

[9] A. Gotsman and H. Yang. Linearizability with ownership transfer. In *CONCUR*, 2012.

[10] M. P. Herlihy and J. M. Wing. Linearizability: a correctness condition for concurrent objects. *TOPLAS*, 12, 1990.

[11] ISO/IEC. *Programming Languages – C++, 14882:2011*.

[12] ISO/IEC. *Programming Languages – C, 9899:2011*.

[13] M. Kuperstein, M. T. Vechev, and E. Yahav. Partial-coherence abstractions for relaxed memory models. In *PLDI*, 2011.

[14] L. Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Comp.*, 28, 1979.

[15] J. Manson. The Java memory model. PhD Thesis. Department of Computer Science, University of Maryland, 2004.

[16] M. M. Michael. Scalable lock-free dynamic memory allocation. In *PLDI*, 2004.

Action structures for locks

$\langle \mathsf{lock}(\ell) \rangle_t = \{ (\{(e, g, t, \mathsf{lock}(\ell))\}, \emptyset), (\{(e, g, t, \mathsf{block}(\ell))\}, \emptyset) \mid e \in \mathsf{AId} \land g \in \mathsf{SectId} \}$

$\langle \mathsf{unlock}(\ell) \rangle_t = \{ (\{(e, g, t, \mathsf{unlock}(\ell))\}, \emptyset) \mid e \in \mathsf{AId} \land g \in \mathsf{SectId} \}$

$\langle \mathsf{let}\ \{ m = C_m \mid m \in M \}\ \mathsf{in}\ C_1 \parallel \dots \parallel C_n \rangle I = \{ (A_0 \cup (\bigcup_{t=1}^n A_t), \bigcup_{t=1}^n \mathsf{sb}_t, (A_0 \times (\bigcup_{t=1}^n A_t))) \mid$
$\qquad (\forall t = 1..n.\ (A_t, \mathsf{sb}_t) \in \langle C_t \rangle_t) \land (\forall t = 1..n.\ \forall u.\ \exists \text{ finitely many } v.\ (v, u) \in \mathsf{sb}_t) \land$
$\qquad (\neg \exists u, v, t.\ (u, v) \in \mathsf{sb}_t \land u = (\_, \_, \_, \mathsf{block}(\_))) \land (A_0 = \bigcup \{ (e, g, 0, \mathsf{store}_x(x, a)) \mid I(x) = (a, \lambda) \land e \in \mathsf{AId} \land g \in \mathsf{SectId} \}) \}$

Additional validity axioms

ATOMRMW. $\forall w, u.\ w \xrightarrow{\mathsf{rf}} u \land u = (\_, \mathsf{rmw}(\_)) \implies w \xrightarrow{\mathsf{mo}} u \land \neg \exists w'.\ w \xrightarrow{\mathsf{mo}} w' \xrightarrow{\mathsf{mo}} u$

ATOMAS. $\forall w, w', r, u, v, x.\ w \xrightarrow{\mathsf{rf}} r \land \mathsf{sec}(u) = \mathsf{sec}(r) \neq \mathsf{sec}(w) = \mathsf{sec}(v) \land r = (\_, \mathsf{read}(x, \_)) \land \mathsf{sort}(x) = \mathsf{ATOM} \implies \neg (w \xrightarrow{\mathsf{mo}} v)$
$\qquad \land (u = (\_, \mathsf{write}(x, \_)) \implies w \xrightarrow{\mathsf{mo}} u \land \neg \exists w''.\ \mathsf{sec}(w'') \neq \mathsf{sec}(u) \land w \xrightarrow{\mathsf{mo}} w'' \xrightarrow{\mathsf{mo}} u)$
$\qquad \land (u = (\_, \mathsf{read}(x, \_)) \land w' \xrightarrow{\mathsf{rf}} r \land \mathsf{sec}(w') \neq \mathsf{sec}(r) \implies w'' = w)$

LOCKS. $\forall u, v.\ u = (\_, \mathsf{lock}(\ell)) \land v = (\_, \mathsf{lock}(\ell)) \land u \xrightarrow{\mathsf{sc}} v \implies \exists q.\ q = (\_, \mathsf{unlock}(\ell)) \land u \xrightarrow{\mathsf{sc}} q \xrightarrow{\mathsf{sc}} v$

SWDEF. $\forall w, r.\ w \xrightarrow{\mathsf{sw}} r \iff (\exists \ell.\ w \xrightarrow{\mathsf{sc}} r \land w = (\_, \mathsf{unlock}(\ell)) \land r = (\_, \mathsf{lock}(\ell))) \lor$
$(\exists t_1, t_2, \lambda, \mu, x, w'.\ t_1 \neq t_2 \land \lambda \in \{\mathsf{SC}, \mathsf{REL}\} \land \mu \in \{\mathsf{SC}, \mathsf{ACQ}\} \land w = (t_1, \mathsf{write}_\lambda(x, \_)) \land r = (t_2, \mathsf{read}_\mu(x, \_)) \land w \xrightarrow{\mathsf{rs}}_{t_1} w' \xrightarrow{\mathsf{rf}} r),$
where $w \xrightarrow{\mathsf{rs}}_t w' \overset{\mathsf{def}}{\iff} \exists w_1.\ w \xrightarrow{\mathsf{mo}}^* w_1 \land (\forall w_2.\ w \xrightarrow{\mathsf{mo}}^* w_2 \xrightarrow{\mathsf{mo}}^* w' \implies (w_2 = (t, \_) \lor w_2 = (\_, \mathsf{rmw}(\_))))$

SCREADS.
$\forall w, r, x.\ w \xrightarrow{\mathsf{rf}} r \land r = (\_, \mathsf{read}_{\mathsf{SC}}(x, \_)) \implies$
$\quad ((w = (\_, \mathsf{write}_{\mathsf{SC}}(x, \_)) \land w \xrightarrow{\mathsf{sc}} r \land \neg \exists w'.\ w' = (\_, \mathsf{write}(x, \_)) \land w \xrightarrow{\mathsf{sc}} w' \xrightarrow{\mathsf{sc}} r) \lor$
$\quad (\exists \lambda.\ w = (\_, \mathsf{write}_\lambda(x, \_)) \land \lambda \neq \mathsf{SC} \land \forall w_1.\ (w_1 = (\_, \mathsf{write}(x, \_)) \land w_1 \xrightarrow{\mathsf{sc}} r \land \neg \exists w_2.\ w_2 = (\_, \mathsf{write}(x, \_)) \land w_1 \xrightarrow{\mathsf{sc}} w_2 \xrightarrow{\mathsf{sc}} r) \implies$
$\qquad\qquad\qquad\qquad \neg (w \xrightarrow{\mathsf{hb}} w_1)))$

Additional safety axiom

SAFELOCK. $(\forall v, t, \ell.\ v \in A \land v = (t, \mathsf{unlock}(\ell)) \implies \exists u.\ u = (t, \mathsf{lock}(\ell)) \land u \xrightarrow{\mathsf{sb}} v \land \neg \exists q.\ q = (\_, \_(\ell)) \land u \xrightarrow{\mathsf{sc}} q \xrightarrow{\mathsf{sc}} v) \land$
$(\neg \exists u, v, t, \ell.\ u = (t, \mathsf{lock}(\ell)) \land v = (t, \mathsf{block}(\ell)) \land u \xrightarrow{\mathsf{sb}} v \land \neg \exists q.\ q = (\_, \_(\ell)) \land u \xrightarrow{\mathsf{sc}} q \xrightarrow{\mathsf{sc}} v)$

Figure 7. Action structures for locks and additional C11 memory model axioms

---

[17] M. M. Michael, M. T. Vechev, and V. A. Saraswat. Idempotent work stealing. In *PPOPP*, 2009.

[18] S. Owens, S. Sarkar, and P. Sewell. A better x86 memory model: x86-TSO. In *TPHOLs*, 2009.

[19] T. Ridge. A rely-guarantee proof system for x86-TSO. In *VSTTE*, 2010.

[20] S. Sarkar, P. Sewell, J. Alglave, L. Maranget, and D. Williams. Understanding POWER multiprocessors. In *PLDI*, 2011.

[21] R. K. Treiber. Systems programming: Coping with parallelism. Technical Report RJ 5118, IBM Almaden Research Center, 1986.

[22] I. Wehrman and J. Berdine. A proposal for weak-memory local reasoning. In *LOLA*, 2011.

## A. Additional Definitions for the C11 Model

In §4–5, we omitted the treatment of locks and CASes from the description of the memory model and simplified some of the axioms, even though the proofs of our theorems do not make such simplifications. Here we provide the missing definitions.

We handle programs with bounded numbers of locks $\ell \in \mathsf{Lock}$, acquired and released using commands $\mathsf{lock}(\ell)$ and $\mathsf{unlock}(\ell)$, respectively. We thus extend the set of actions as follows: $\varphi ::= \dots \mid \mathsf{lock}(\ell) \mid \mathsf{unlock}(\ell) \mid \mathsf{block}(\ell)$, where $\ell \in \mathsf{Lock}$. An action $(e, g, t, \mathsf{block}(\ell))$ represents a deadlocked attempt to acquire a lock ℓ. We split the set of locks into client and library ones ($\mathsf{Lock} = \mathsf{CLock} \uplus \mathsf{LLock}$) and consider only programs where the client and the library use locks from CLock and LLock, respectively.

In Figure 7, we give the actions structures for lock operations, omitted from Figure 3, and the C11 axioms omitted from Figures 4

and 5. Note that the action structures of a whole program do not include those where a thread executes actions after blocking. To adjust the axioms to the model with locks, we require that the sc relation totally orders actions of the form $(\_, \mathsf{lock}(\_))$, $(\_, \mathsf{block}(\_))$ or $(\_, \mathsf{unlock}(\_))$ for each lock, in addition to SC actions. The axioms ATOMRMW and LOCKS define the behaviour of CAS commands and locks. ATOMAS is an additional axiom for atomic sections, in the spirit of ATOMRMW. SWDEF and SCREADS are full versions of the axioms presented in a simplified form in Figure 4. The former adds synchronisation via *release sequences* [1] (defined by the rs relation), and the latter allows SC reads from non-SC writes. SAFELOCK is an additional safety axiom, flagging a double unlock or a double lock in the same thread as a fault.

All the theorems stated in the paper stay valid for this extension of the model. To account for the full version of the SCREADS axiom, we add additional edges to $\mathsf{denyL}(X)$. For $X = (A, \mathsf{sb}, \mathsf{ib}, \mathsf{rf}, \mathsf{sc}, \mathsf{mo}, \mathsf{sw}, \mathsf{hb})$, $\mathsf{denyL}(X)$ includes the dashed edges of all possible instantiations of the following diagram:

$$\begin{pmatrix} w = (\_, \mathsf{write}_\lambda(x, \_)) \land \lambda \neq \mathsf{SC} \land \\ r = (\_, \mathsf{read}_{\mathsf{SC}}(x, \_)) \land \\ w_1 = (\_, \mathsf{write}(x, \_)) \land \\ (\neg \exists w_2.\ w_2 = (\_, \mathsf{write}(x, \_)) \land \\ w_1 \xrightarrow{\mathsf{sc}} w_2 \xrightarrow{\mathsf{sc}} r) \end{pmatrix} \land \begin{array}{c} r \xleftarrow{\mathsf{rf}} w \\ \downarrow \mathsf{hb} \\ u \\ \mathsf{sc} \quad \downarrow \\ v \\ \downarrow \mathsf{hb} \\ w_1 \end{array}$$

where u is a return, and v is a call. Its counterpart $\mathsf{denyC}(X)$ contains the dashed edges assuming u is a call, and v is a return.

B. Auxiliary Definitions

Notation. In the following we use interf and other selectors not only on executions, but also on sets of actions and relations over them. We write $r|_A$ for the projection of a relation r to a set of actions A. We write $X|_A$ for the obvious lifting to executions X.

The formalisation of a restriction on base commands. We formalise the requirement that the set $\langle c \rangle_t$ account for c fetching any values from the memory locations it reads (§4) as follows:

$$\forall c, t. \forall (A, \mathsf{sb}) \in \langle c \rangle_t. \forall A'. \forall u = (e, g, t, \mathsf{read}_\lambda(x, a)) \in A'.$$
$$A' \subseteq A \wedge (\neg \exists v \in A'. u \xrightarrow{\mathsf{sb}} v) \wedge$$
$$(\forall q \in A, v \in A'. q \xrightarrow{\mathsf{sb}} v \implies q \in A') \wedge$$
$$\implies$$
$$\exists (A'', \mathsf{sb}'') \in \langle c \rangle_t. \exists u' = (e, g, t, \mathsf{read}_\lambda(x, b)).$$
$$A'' \supseteq (A' \setminus \{u\}) \cup \{u'\} \wedge$$
$$\mathsf{sb}''|_{(A' \setminus \{u\}) \cup \{u'\}} =$$
$$\mathsf{sb}|_{A' \setminus \{u\}} \cup \{(q, u') \mid q \in A' \wedge (q, u) \in \mathsf{sb}\} \wedge$$
$$(\forall q \in A'', v \in ((A' \setminus \{u\}) \cup \{u'\}). q \xrightarrow{\mathsf{sb}} v \implies$$
$$q \in ((A' \setminus \{u\}) \cup \{u'\})). \tag{1}$$

Semantics of loops. The actions structures for loops, omitted from Figure 3, are given in Figure 8.

C11 well-formedness axioms. In Figure 9, we summarise the well-formedness constraints on various relations that we stated in plain text earlier, to make references to them easier in our proofs.

Relationship to the original C11 formalisation. Consider the memory model without atomic sections, i.e., without the axioms ASMO and ASSC and with HBDEF formulated as follows:

$$\mathsf{hb} = (\mathsf{sb} \cup \mathsf{ib} \cup \mathsf{sw})^+.$$

In the formalisation of this model by Batty et al. [1], the sc relation totally orders SC operations on *all* locations and locks. Namely, the axioms SCWF and ACYCLICITY are replaced by the following:

- SCWF′. sc is a transitive, irreflexive, total order on all actions of the form $(\_, \mathsf{lock}(\_)), (\_, \mathsf{block}(\_)), (\_, \mathsf{unlock}(\_)), (\_, \_\mathsf{sc} \_)$.
- HBVSSC. $\neg \exists u, v. u \xrightarrow{\mathsf{sc}} v \wedge v \xrightarrow{\mathsf{hb}} u$.
- ACYCLICITY′. hb is irreflexive.

Let M_1 be the C11 model with these axioms, and M_2, the model adopted in this paper. The following proposition, proved in §C.13, shows that the two are equivalent.

PROPOSITION 8. *1. Every valid execution X in M_1 has a corresponding valid execution X' in M_2, with $\mathsf{sc}(X') \subseteq \mathsf{sc}(X)$, and all other components of the execution the same.*
2. Every valid execution X' in M_2 has a corresponding valid execution X in M_1, with $\mathsf{sc}(X') \subseteq \mathsf{sc}(X)$, and all other components of the execution the same.
3. Consider a valid execution X in M_1 and a valid execution X' in M_2, such that $\mathsf{sc}(X') \subseteq \mathsf{sc}(X)$, and all other components of the executions are the same. Then X is safe if and only if X' is safe.

C. Proofs of Theorems

C.1 Proof of Theorem 3 (Abstraction)

PROPOSITION 9. *If an execution X is safe, then so is its extension with any R.*

In particular, extensions of library executions used in Definition 2 are safe.

Since action structures of base commands do not include calls and returns, by the construction in Figure 3, calls and returns in the same thread are totally related by sb. This implies

PROPOSITION 10. *For any valid execution X,*

$$\{(u, v) \in \mathsf{interf}(\mathsf{hb}(X)) \mid$$
$$\exists t. u = (t, \_) \wedge v = (t, \_)\} = \mathsf{interf}(\mathsf{sb}(X)).$$

The following lemma states that an execution of a complete program generates two executions in the client-local and library-local semantics.

LEMMA 11 (Decomposition). *For any $X \in [\![\mathcal{C}(\mathcal{L}_1)]\!](I \uplus I_1)$ satisfying NONINTERF,*

$$\mathsf{client}(X) \in [\![\mathcal{C}, \mathsf{hbL}(\mathsf{core}(\mathsf{lib}(X)))]\!]I; \tag{2}$$
$$\mathsf{lib}(X) \in [\![\mathcal{L}_1, \mathsf{hbC}(\mathsf{core}(\mathsf{client}(X)))]\!]I_1. \tag{3}$$

Furthermore,

- $\mathsf{client}(X)$ *and* $\mathsf{lib}(X)$ *satisfy* NONINTERF;
- *if X is unsafe, then so is either* $\mathsf{client}(X)$ *or* $\mathsf{lib}(X)$*; and*
- $\mathsf{scC}(\mathsf{client}(X)) \cup \mathsf{scL}(\mathsf{lib}(X))$ *is acyclic.*

The following lemma states that any pair of appropriately extended client and library executions with an acyclic union of deny edges can be combined into a valid execution of the complete program.

LEMMA 12 (Composition). *Consider*

$$X \in [\![\mathcal{C}, \mathsf{hbL}(\mathsf{core}(Y))]\!]I;$$
$$Y \in [\![\mathcal{L}_2, \mathsf{hbC}(\mathsf{core}(X))]\!]I_2,$$

such that

- $A(X) \cap A(Y) = \mathsf{interf}(X) = \mathsf{interf}(Y)$ *and otherwise action and atomic section identifiers in $A(X)$ and $A(Y)$ are different;*
- $\mathsf{interf}(\mathsf{sb}(X)) = \mathsf{interf}(\mathsf{sb}(Y))$;
- $\mathsf{scC}(X) \cup \mathsf{scL}(Y)$ *is acyclic; and*
- X *and* Y *satisfy* NONINTERF.

Then for some $Z \in [\![\mathcal{C}(\mathcal{L}_2)]\!](I \uplus I_2)$ we have $X = \mathsf{client}(Z)$ and $Y = \mathsf{lib}(Z)$. Furthermore, if X is unsafe, then so is Z.

Proofs of the above lemmas are given in the following sections.

PROOF OF THEOREM 3. Consider $X \in [\![\mathcal{C}(\mathcal{L}_1)]\!](I \uplus I_1)$ for $I \in \mathcal{I}$ and $I_1 \in \mathcal{I}_1$. Since X satisfies NONINTERF, by Lemma 11, we have that (2) and the other conclusions of this lemma hold. Note that, since $(\mathcal{L}_1, \mathcal{I}_1)$ is safe by Proposition 9, $\mathsf{lib}(X)$ is safe. Hence, if X is unsafe, then so is $\mathsf{client}(X)$.

Let R be the projection of $\mathsf{hbC}(\mathsf{core}(\mathsf{client}(X)))$ to pairs of actions of the form $((\_, \mathsf{ret}\ \_), (\_, \mathsf{call}\ \_))$. Then by Proposition 10,

$$\mathsf{hbC}(\mathsf{core}(\mathsf{client}(X))) \setminus R \subseteq$$
$$\mathsf{interf}(\mathsf{sb}(\mathsf{client}(X))) \subseteq \mathsf{interf}(\mathsf{sb}(\mathsf{lib}(X))). \tag{4}$$

Thus, from (3), we obtain $\mathsf{lib}(X) \in [\![\mathcal{L}_1, R]\!]I_1$. Since $(\mathcal{L}_1, \mathcal{I}_1) \sqsubseteq (\mathcal{L}_2, \mathcal{I}_2)$, by Definition 2 for some $I_2 \in \mathcal{I}_2$ and $Y \in [\![\mathcal{L}_2, R]\!]I_2$, we have

$$\mathsf{history}(\mathsf{lib}(X)) \sqsubseteq \mathsf{history}(Y).$$

Hence,

$$\mathsf{interf}(Y) = \mathsf{interf}(\mathsf{lib}(X));$$
$$\mathsf{hbL}(Y) = \mathsf{hbL}(\mathsf{lib}(X)); \quad \mathsf{scL}(Y) = \mathsf{scL}(\mathsf{lib}(X)),$$

Then

$$\mathsf{scC}(\mathsf{client}(X)) \cup \mathsf{scL}(Y) \subseteq \mathsf{scC}(\mathsf{client}(X)) \cup \mathsf{scL}(\mathsf{lib}(X)),$$

$$\langle \text{while}(x)\ \{C\}\rangle_t = \{(\bigcup_{i=1}^{k}(\{u_i\} \uplus A_i) \cup \{u_{k+1}\}, \bigcup_{i=1}^{k}\text{sb}_i \cup \{(u,v) \mid u \in A_i \uplus \{u_i\} \wedge v \in A_j \uplus \{u_j\} \wedge 1 \leq i < j \leq k\} \cup$$
$$\{(u_i,v) \mid v \in A_i\} \cup \{(u,u_{k+1}) \mid u \in \bigcup_{i=1}^{k}(\{u_i\} \uplus A_i) \cup \{u_{k+1}\}\}) \mid$$
$$(A_i,\text{sb}_i) \in \langle C\rangle_t \wedge k \geq 0 \wedge (\forall i = 1..k.\ \exists a.\ u_i = (\_,\_,t,\text{load}_{\mathsf{NA}}(x,a)) \wedge a \neq 0) \wedge u_{k+1} = (\_,\_,t,\text{load}_{\mathsf{NA}}(x,0))\}$$
$$\cup$$
$$\{(\bigcup_{i=1}^{\infty}(\{u_i\} \uplus A_i), \bigcup_{i=1}^{\infty}\text{sb}_i \cup \{(u,v) \mid u \in A_i \uplus \{u_i\} \wedge v \in A_j \uplus \{u_j\} \wedge i < j\} \cup \{(u_i,v) \mid v \in A_i\} \mid$$
$$(A_i,\text{sb}_i) \in \langle C\rangle_t \wedge (\forall i.\ \exists a.\ u_i = (\_,\_,t,\text{load}_{\mathsf{NA}}(x,a)) \wedge a \neq 0)\}$$

Figure 8. Action structures for loops

SBWF. sb is transitive, irreflexive and relates only actions by the same thread

RFWF. $(\forall w_1, w_2, r.\ w_1 \xrightarrow{\text{rf}} r \wedge w_2 \xrightarrow{\text{rf}} r \implies w_1 = w_2) \wedge (\forall w, r.\ w \xrightarrow{\text{rf}} r \implies \exists x, a.\ w = (\_,\text{write}(x,a)) \wedge r = (\_,\text{read}(x,a)))$

SCWF. sc is a union of transitive, irreflexive and total orders on SC actions for each atomic location, and on actions of the form $(\_,\text{lock}(\_))$, $(\_,\text{block}(\_))$ or $(\_,\text{unlock}(\_))$ for each lock

MOWF. mo is the union of transitive, irreflexive and total orders on writes to each atomic location

Figure 9. C11 well-formedness axioms

---

and we have obtained from Lemma 11 that this is acyclic. By Proposition 10, we also have

$$\text{interf}(\text{sb}(Y)) = \text{interf}(\text{sb}(\text{lib}(X))).$$

Then from $Y \in [\![\mathcal{L}_2, R]\!]I_2$ and (4), we get

$$Y \in [\![\mathcal{L}_2, \text{hbC}(\text{core}(\text{client}(X)))]\!]I_2. \tag{5}$$

We have:

$$\text{interf}(Y) = \text{interf}(\text{lib}(X)) = \text{interf}(\text{client}(X)).$$

Since action structures of base commands are insensitive to the choice of action and atomic section identifiers (§4), we can assume that $A(\text{client}(X)) \cap A(Y) = \text{interf}(Y)$, and otherwise action and atomic section identifiers in $A(\text{client}(X))$ and $A(Y)$ are different. Then from (2) and (5), by Lemma 12 we get that for some $Z \in [\![\mathcal{C}(\mathcal{L}_2)]\!](I \uplus I_2)$ we have $\text{client}(Z) = \text{client}(X)$. Furthermore, if X is unsafe, then so is $\text{client}(X)$, and then by Lemma 12, so is Z. Thus the safety of $\mathcal{C}(\mathcal{L}_2)$ implies that of $\mathcal{C}(\mathcal{L}_1)$. □

C.2 Proof of Theorem 7 (Abstraction without relaxed atomics)

Recall that here we restrict ourselves to programs whose action structures do not have any relaxed actions, and we prefix-close the set of action structures of every program. We use the following variation on the notion of an execution.

DEFINITION 13. *We call an execution X **almost-valid**, if:*

- X *satisfies all the validity axioms except* RFNONATOMIC *and possibly* DETREAD*;*
- X *violates* RFNONATOMIC *solely because of the existence of some number of pairs of actions $u, v \in A(X)$ accessing locations other than* param_t *or* retval_t *such that*

$$u \xrightarrow{\text{rf}(X)} v \wedge \neg(u \xrightarrow{\text{hb}(X)} v);$$

- X *may violate* DETREAD *only because* DETREAD *rules out some of the above edges in* $\text{rf}(X)$.
- $\text{hb} \cup \text{rf} \cup \text{sc}$ *is acyclic.*

Let $(\![\mathcal{C}(\mathcal{L})]\!)I$ be the set of almost-valid executions of $\mathcal{C}(\mathcal{L})$ from I. We define the notions execution extensions, library and client executions derived from the above as expected.

PROPOSITION 14 (Absence of causal cycles). *For a valid execution X without relaxed actions, $\text{rf}(X) \subseteq \text{hb}(X)$, and hence,*

$\text{hb}(X) \cup \text{rf}(X) \cup \text{sc}(X)$ *is acyclic. When X has an action structure generated as in Figure 3, the inverse of this relation is well-founded; this is also true when X is almost-valid.*

Note that the acyclicity guarantee is not valid for almost-valid executions, as the following example shows:

LEMMA 15 (Prefix construction). *Consider $X \in [\![\mathcal{C}(\mathcal{L})]\!]I$ and $A' \subseteq A(X)$ such that*

$$\forall u \in A(X).\ u = (0, \_) \implies u \in A' \wedge$$
$$\forall v \in A'.\ \forall u \in A(X).\ u \xrightarrow{(\text{hb}(X) \cup \text{rf}(X) \cup \text{sc}(X))^+} v \implies u \in A'.$$

Then $X|_{A'} \in [\![\mathcal{C}(\mathcal{L})]\!]I$. The above also holds for $\mathcal{L}$ instead of $\mathcal{C}(\mathcal{L})$. It also holds for almost-valid, rather than valid, executions, provided at least one action violating RFNONATOMIC *is in A'.*

PROPOSITION 16. *For any library execution X,*

$$\text{interf}(\text{hb}(X)) = (\text{hbL}(X))^+.$$

For any client execution X,

$$\text{interf}(\text{hb}(X)) = (\text{hbC}(X))^+.$$

PROPOSITION 17 (Execution reduction). *Assume $R' \subseteq R$ and the extension of a client or library execution X with R is admissible. Then its extension with R' is either admissible or almost-admissible. Furthermore, if the extension of X with R is unsafe, then so is its extension with R'.*

LEMMA 18 (Read completion). *If $Y \in (\![\mathcal{C}(\mathcal{L})]\!)I$, then there is an execution $Z \in [\![\mathcal{C}(\mathcal{L})]\!]I$ that violates DRF. The above also holds for $\mathcal{L}$ instead of $\mathcal{C}(\mathcal{L})$.*

LEMMA 19 (Appending an action). *Let $X \in [\![\mathcal{C}(\mathcal{L})]\!]I$ and $(A', \text{sb}', \text{ib}') \in \langle\mathcal{C}(\mathcal{L})\rangle I$ be such that X satisfies* NONINTERF*,*

$$A' = A(X) \uplus \{u\}, \quad u = (\_, \mathsf{a}(x, \_)), \quad \mathsf{a} \in \{\text{read}, \text{write}\},$$
$$\text{sb}(X) \subseteq \text{sb}', \quad \text{sb}' \setminus \text{sb}(X) \subseteq (A(X) \times \{u\}),$$

and u violates NONINTERF *in $(A', \text{sb}', \text{ib}')$. Then there exists $X' \in [\![\mathcal{C}(\mathcal{L})]\!]I$ such that for some $v = (\_, \mathsf{a}(x, \_))$ we have*

$$A(X') = A(X) \uplus \{v\},$$
$$\text{sb}(X') = \text{sb}(X) \cup \{(q, v) \mid (q, u) \in \text{sb}'\}.$$

The above also holds for $\mathcal{L}$ instead of $\mathcal{C}(\mathcal{L})$.

LEMMA 20 (Execution extension). *If the extension of a safe execution $X \in [\![\mathcal{L}]\!]I$ with R is admissible and safe, and $Y \in [\![\mathcal{L}']\!]I'$ is a valid and safe execution such that*

$$\mathsf{Ehistory}(X) \preceq \mathsf{Ehistory}(Y),$$

then its extension with R is also admissible and safe.

For an execution X, we let $\mathsf{EscC}(X)$ be the projection of $(\mathsf{hb}(X) \cup \mathsf{sc}(X) \cup \mathsf{rf}(X))^+$ to pairs of actions of the form $((\_, \mathsf{ret}\ \_), (\_, \mathsf{call}\ \_))$.

PROPOSITION 21 (Composition generalisation). *Lemma 12 holds when X and Z are almost-admissible, provided $\mathsf{EscC}(X) \cup \mathsf{scL}(Y)$.*

Proofs of the above lemmas are given in the following sections.

PROOF OF THEOREM 7. Let n be the number of threads in $\mathcal{C}(\mathcal{L})$. Consider $X' \in [\![\mathcal{C}(\mathcal{L}_1)]\!](I \uplus I_1)$ for $I \in \mathcal{I}$ and $I_1 \in \mathcal{I}_1$.

If X' violates NONINTERF, by Proposition 14, there exists an action $v \in A(X')$ that violates NONINTERF such that there does not exist another action $u \in A(X')$ violating NONINTERF and satisfying

$$u \xrightarrow{(\mathsf{hb}(X') \cup \mathsf{rf}(X') \cup \mathsf{sc}(X'))^+} v.$$

Note that v is not an interface action. Let

$$A' = \{u \mid u = (0, \_) \in A(X')\} \cup$$
$$\{u \mid u \xrightarrow{(\mathsf{hb}(X') \cup \mathsf{rf}(X') \cup \mathsf{sc}(X'))^+} v\}$$

and $X = X'|_{A'}$. Then by Lemma 15, $X \in [\![\mathcal{C}(\mathcal{L}_1)]\!](I \uplus I_1)$. Furthermore, by the choice of action v, X satisfies NONINTERF.

Let $\mathsf{sb}' = \mathsf{sb}(X')|_{A(X) \uplus \{v\}}$. Then by Proposition 14,

$$(A(X) \uplus \{v\}, \mathsf{sb}', \mathsf{ib}|_{A(X) \uplus \{v\}}) \in \langle \mathcal{C}(\mathcal{L}_1) \rangle (I \uplus I_1),$$
$$\mathsf{sb}(X) \subseteq \mathsf{sb}', \quad \mathsf{sb}' \setminus \mathsf{sb}(X) \subseteq (A(X) \times \{v\}).$$

If X' satisfies NONINTERF, we let $X = X'$.
By Lemma 11,

$$\mathsf{client}(X) \in [\![\mathcal{C}, \mathsf{hbL}(\mathsf{core}(\mathsf{lib}(X)))]\!]I; \tag{6}$$
$$\mathsf{lib}(X) \in [\![\mathcal{L}_1, \mathsf{hbC}(\mathsf{core}(\mathsf{client}(X)))]\!]I_1. \tag{7}$$

and the other conclusions of Lemma 11 hold. Note that, by Proposition 9, $\mathsf{lib}(X)$ is safe. Hence, if X is unsafe, then so is $\mathsf{client}(X)$.

From (7) and Proposition 17, for some Y_1, $\mathsf{lib}(X)$ is the admissible extension of Y_1 with $\mathsf{hbC}(\mathsf{core}(\mathsf{client}(X)))$, where either $Y_1 \in [\![\mathcal{L}_1]\!]I_1$, or $Y_1 \in (\![\mathcal{L}_1]\!)I_1$. However, in the latter case, by Lemma 18, $\mathcal{L}_1$ is unsafe when run from I_1, contradicting the assumptions of the theorem. Thus, $Y_1 \in [\![\mathcal{L}_1]\!]I_1$.

Consider the case when $X \ne X'$ and the action v is a library action. Let $A'' = A(Y_1) \uplus \{v\}$. From $\mathsf{sb}(Y_1) = \mathsf{sb}(\mathsf{lib}(X))$ and sb', we can easily construct sb'', $\mathsf{ib}'' \in A' \times A'$ such that

$$(A', \mathsf{sb}'', \mathsf{ib}'') \in \langle \mathsf{MGC}_n(Y_1) \rangle I_1,$$
$$\mathsf{sb}(Y_1) \subseteq \mathsf{sb}'', \quad \mathsf{sb}'' \setminus \mathsf{sb}(Y_1) \subseteq (A(Y_1) \times \{v\}).$$

By Lemma 19, we can then construct an execution $Y_1' \in [\![\mathcal{L}_1]\!]I_1$ that contains v. But then this execution violates NONINTERF, which contradicts the safety of $\mathcal{L}_1$. Hence, the case considered is impossible.

Since $(\mathcal{L}_1, \mathcal{I}_1) \sqsubseteq (\mathcal{L}_2, \mathcal{I}_2)$, for some $I_2 \in \mathcal{I}_2$ and $Y_2 \in [\![\mathcal{L}_2]\!]I_2$ we have $\mathsf{Ehistory}(Y_1) \preceq \mathsf{Ehistory}(Y_2)$. Note that by Proposition 10, this implies

$$\mathsf{interf}(\mathsf{sb}(Y_1)) = \mathsf{interf}(\mathsf{sb}(Y_2)). \tag{8}$$

By (7) and Lemma 20, the execution Z obtained by extending Y_2 with $\mathsf{hbC}(\mathsf{core}(\mathsf{client}(X)))$ is admissible and safe. Then

$$Z \in [\![\mathcal{L}_2, \mathsf{hbC}(\mathsf{core}(\mathsf{client}(X)))]\!]I_2. \tag{9}$$

The relations $\mathsf{scL}(\mathsf{lib}(X))$ and $\mathsf{scL}(Z)$ can be obtained by an appropriate projection from

$$(\mathsf{hbL}(Y_1) \cup \mathsf{hbC}(\mathsf{core}(\mathsf{client}(X))) \cup \mathsf{scL}(Y_1))^+,$$

respectively,

$$(\mathsf{hbL}(Y_2) \cup \mathsf{hbC}(\mathsf{core}(\mathsf{client}(X))) \cup \mathsf{scL}(Y_2))^+.$$

Since $\mathsf{hbL}(Y_2) \subseteq \mathsf{hbL}(Y_1)$ and $\mathsf{scL}(Y_2) \subseteq \mathsf{scL}(Y_1)$, this implies $\mathsf{scL}(Z) \subseteq \mathsf{scL}(\mathsf{lib}(X))$. Hence,

$$\mathsf{scC}(\mathsf{client}(X)) \cup \mathsf{scL}(Z) \subseteq \mathsf{scC}(\mathsf{client}(X)) \cup \mathsf{scL}(\mathsf{lib}(X)),$$

and we have obtained from Lemma 11 that this is acyclic.
We have:

$$\mathsf{hbL}(\mathsf{core}(Z)) = \mathsf{hbL}(Y_2) \subseteq \mathsf{hbL}(Y_1) = \mathsf{hbL}(\mathsf{core}(\mathsf{lib}(X))).$$

From (6), by Proposition 17, there exists W such that $\mathsf{client}(X) \preceq W$,

$$W \in [\![\mathcal{C}, \mathsf{hbL}(\mathsf{core}(Z))]\!]I \vee W \in (\![\mathcal{C}, \mathsf{hbL}(\mathsf{core}(Z))]\!)I, \tag{10}$$

and if $\mathsf{client}(X)$ is unsafe, then so is W.
We have:

$$\mathsf{interf}(Z) = \mathsf{interf}(Y_2) = \mathsf{interf}(Y_1) =$$
$$\mathsf{interf}(\mathsf{lib}(X)) = \mathsf{interf}(\mathsf{client}(X)) = \mathsf{interf}(W).$$

Since action structures for base commands are insensitive to the choice of action and atomic section identifiers (§4), we can assume that $A(Z) \cap A(W) = \mathsf{interf}(Z)$, and otherwise the action and atomic section identifiers from $A(Z)$ and $A(W)$ are different. By (8), we also have:

$$\mathsf{interf}(\mathsf{sb}(Z)) = \mathsf{interf}(\mathsf{sb}(Y_2)) = \mathsf{interf}(\mathsf{sb}(Y_1)) =$$
$$\mathsf{interf}(\mathsf{sb}(\mathsf{lib}(X))) = \mathsf{interf}(\mathsf{client}(X)) = \mathsf{interf}(W).$$

Since

$$\mathsf{EscC}(W) \subseteq \mathsf{scC}(\mathsf{client}(X))$$

we get that

$$\mathsf{EscC}(W) \cup \mathsf{scL}(Z)$$

is acyclic.
Since, $\mathsf{hbC}(\mathsf{core}(W)) = \mathsf{hbC}(\mathsf{core}(\mathsf{client}(X)))$, from (9) we get

$$Z \in [\![\mathcal{L}_2, \mathsf{hbC}(\mathsf{core}(W))]\!]I_2.$$

From this and (10), by Proposition 21, for some V we have $\mathsf{client}(X) \preceq W = \mathsf{client}(V)$,

$$V \in [\![\mathcal{C}(\mathcal{L}_2)]\!](I \uplus I_2) \vee V \in (\![\mathcal{C}(\mathcal{L}_2)]\!)(I \uplus I_2).$$

Furthermore, if W is unsafe, then so is V. However, by Lemma 18, the second disjunct contradicts the safety of $\mathcal{C}(\mathcal{L}_2)$. Hence, $V \in [\![\mathcal{C}(\mathcal{L}_2)]\!](I \uplus I_2)$. Furthermore, if X is unsafe, then so is V.

If $X = X'$, this completes the proof of the theorem.

Let us now return to the case when X was obtained from X' by applying Lemma 15 with respect to the action v. Then v is a client action (the case of a library action has already been covered above). Note that $\mathsf{sb}(\mathsf{client}(V)) = \mathsf{sb}(\mathsf{client}(X))$. Analogously to how it was done before, we can append v to V using Lemma 19, getting an execution $V' \in [\![\mathcal{C}(\mathcal{L}_2)]\!](I \uplus I_2)$ violating NONINTERF and contradicting the safety of $\mathcal{C}(\mathcal{L}_2)$. □

C.3 Proof of Lemma 11 (Decomposition)

PROPOSITION 22. *Consider an execution X a set of actions $B \subseteq A(X)$. Let $X' = X|_B$. For every one of the following axioms, if X satisfies the axiom, then X' does:* SBWF, RFWF, SCWF, MOWF, ATOMRMW, ACYCLICITY, RFNONATOMIC, RFATOMIC, HB-vsMO, MOvsSC, CORR, COWR, CORW, ASMO, ASSC, ATOMAS.

It is easy to show that

$$(A(\mathsf{client}(X)), \mathsf{sb}(\mathsf{client}(X)), \mathsf{ib}(\mathsf{client}(X))) \in \langle \mathcal{C}(\cdot) \rangle I;$$
$$(A(\mathsf{lib}(X)), \mathsf{sb}(\mathsf{lib}(X)), \mathsf{ib}(\mathsf{lib}(X))) \in \langle \mathsf{MGC}_n(\mathcal{L}_1) \rangle I_1,$$

where n is the number of threads in $\mathcal{C}(\mathcal{L}_1)$. It is also easy to check that NONINTERF holds for $\mathsf{client}(X)$ and $\mathsf{lib}(X)$. We now show that $\mathsf{client}(X)$ and $\mathsf{lib}(X)$ are admissible. By Proposition 22, we only need to check the validity of LOCKS, SWDEF, SCREADS, EXTHBDEF and DETREAD$_c$ or DETREAD$_l$. We have:

- LOCKS is valid, because the client and the library access disjoint sets of locks (§A).

- Since X satisfies NONINTERF, $\mathsf{mo}(X)$ edges do not relate actions by different components; and $\mathsf{rf}(X)$ edges do not relate actions by different components, except accesses to param_t and retval_t. Besides, clients and the libraries access disjoint sets of locks. Thus, SWDEF and SCREADS are valid.

- For similar reasons, DETREAD$_c$ is valid for $\mathsf{client}(X)$ and DETREAD$_l$, for $\mathsf{lib}(X)$.

We now show that EXTHBDEF holds of $\mathsf{client}(X)$ (the case of $\mathsf{lib}(X)$ is analogous). Let $R = \mathsf{hbL}(\mathsf{core}(\mathsf{lib}(X)))$. We need to show that

$$\mathsf{hb}(\mathsf{client}(X)) =$$
$$((\mathsf{ib}(\mathsf{client}(X)) \cup \mathsf{sb}(\mathsf{client}(X)) \cup \mathsf{sw}(\mathsf{client}(X)) \cup R)/\sim)^+.$$

We have:

$$((\mathsf{ib}(\mathsf{client}(X)) \cup \mathsf{sb}(\mathsf{client}(X)) \cup \mathsf{sw}(\mathsf{client}(X)) \cup R)/\sim)^+ \subseteq$$
$$((\mathsf{ib}(X) \cup \mathsf{sb}(X) \cup \mathsf{sw}(X) \cup \mathsf{hb}(X))/\sim)^+ \subseteq \mathsf{hb}(X).$$

Since $\mathsf{ib}(\mathsf{client}(X))$, $\mathsf{sb}(\mathsf{client}(X))$, $\mathsf{sw}(\mathsf{client}(X))$ and R relate client actions only, this implies

$$((\mathsf{ib}(\mathsf{client}(X)) \cup \mathsf{sb}(\mathsf{client}(X)) \cup \mathsf{sw}(\mathsf{client}(X)) \cup R)/\sim)^+ \subseteq$$
$$\mathsf{hb}(\mathsf{client}(X)).$$

Let us now show that

$$\mathsf{hb}(\mathsf{client}(X)) \subseteq$$
$$((\mathsf{ib}(\mathsf{client}(X)) \cup \mathsf{sb}(\mathsf{client}(X)) \cup \mathsf{sw}(\mathsf{client}(X)) \cup R)/\sim)^+.$$

Consider $(u, v) \in \mathsf{hb}(\mathsf{client}(X))$. Then $u, v \in A(\mathsf{client}(X))$ and there exists a path from u to v consisting of edges in $\mathsf{ib}(X)/\sim$, $\mathsf{sw}(X)/\sim$ and $\mathsf{sb}(X)/\sim$. We only consider then case when the path does not have any edges in $\mathsf{ib}(X)/\sim$; the other case is handled trivially. By the construction in Figure 3, we can split $\mathsf{sb}(X)/\sim$ edges to make all intermediate interface actions in the program order explicit. Since X satisfies NONINTERF, edges in $\mathsf{sw}(X)$ do not relate internal actions by different components. Hence, any segment of library actions on the path starts with a call action u' and ends with a return action v'. Then $(u', v') \in R$. Thus, every such segment can be replaced by an edge in R, yielding a path from u to v with edges from $\mathsf{sb}(\mathsf{client}(X))/\sim$, $\mathsf{sw}(\mathsf{client}(X))/\sim$ and R. Therefore,

$$(u, v) \in$$
$$((\mathsf{ib}(\mathsf{client}(X)) \cup \mathsf{sb}(\mathsf{client}(X)) \cup \mathsf{sw}(\mathsf{client}(X)) \cup R)/\sim)^+,$$

as required.

Thus, we have shown that

$$\mathsf{client}(X) \in [\![\mathcal{C}, \mathsf{hbC}(\mathsf{core}(\mathsf{lib}(X)))]\!] I;$$
$$\mathsf{lib}(X) \in [\![\mathcal{L}_1, \mathsf{hbL}(\mathsf{core}(\mathsf{client}(X)))]\!] I_1.$$

Assume that

$$\mathsf{scC}(\mathsf{client}(X)) \cup \mathsf{scL}(\mathsf{lib}(X))$$

has a cycle. But then this cycle can be converted into one using only edges in $\mathsf{hb}(X)$ and $\mathsf{sc}(X)$, contradicting ACYCLICITY. Hence, the above relation is acyclic.

Assume now that X is unsafe, i.e., DRF, SAFEREAD or SAFE-LOCK is violated. Consider first the case when DRF does not hold for X. Since X satisfies NONINTERF, both actions u and v violating DRF have to come from the same component. Hence, either $\mathsf{client}(X)$ or $\mathsf{lib}(X)$ is unsafe, as required. The case of SAFELOCK is handled similarly, using the fact that the client and the library access disjoint sets of locks. If SAFEREAD is violated in X, then for some action $r = (t, \mathsf{read}(x, \_)) \in A(X)$, there is no action w such that $w \xrightarrow{\mathsf{rf}(X)} r$. Consider the case when r is a client action (the case of a library action is handled similarly). Then there is no action w such that $w \xrightarrow{\mathsf{rf}(\mathsf{client}(X))} r$. If $x \neq \mathsf{retval}_t$, then this violates SAFEREAD$_c$. If $x = \mathsf{retval}_t$, then there is no action $u = (t, \mathsf{ret} \_)$ such that $u \xrightarrow{\mathsf{hb}(X)} r$, for otherwise by DETREAD for X, we would have $w \xrightarrow{\mathsf{rf}(X)} r$ for some w. But then SAFEREAD$_c$ is violated in $\mathsf{client}(X)$. $\qquad\square$

C.4 Proof of Lemma 12 (Composition)

PROPOSITION 23. *If one of the following axioms is violated in an execution X satisfying NONINTERF and with the action structure generated as in Figure 3, then all the actions involved into the violating configuration come from* $\mathsf{client}(X)$ *or* $\mathsf{lib}(X)$: ATOMRMW, RFATOMIC, SCREADS, HBVSMO, MOVSSC, COWR, CORW, CORR, ASMO, ASSC, ATOMAS.

We show how to construct the required

$$Z = (A, \mathsf{sb}, \mathsf{ib}, \mathsf{rf}, \mathsf{sc}, \mathsf{mo}, \mathsf{sw}, \mathsf{hb}).$$

Let

$$A = A(X) \cup A(Y);$$
$$\mathsf{ib} = \{(u, v) \mid \exists t.\, u = (0, \_) \in A \wedge v = (t, \_) \in A \wedge t \neq 0\};$$
$$\mathsf{sb} = (\mathsf{sb}(X) \cup \mathsf{sb}(Y))^+; \quad \mathsf{mo} = \mathsf{mo}(X) \cup \mathsf{mo}(Y).$$

By the assumptions of the lemma,

$$A(X) \cap A(Y) = \mathsf{interf}(A(X)) = \mathsf{interf}(A(Y))$$

and otherwise action and atomic section identifiers in $A(X)$ and $A(Y)$ are different, and

$$\mathsf{interf}(\mathsf{sb}(X)) = \mathsf{interf}(\mathsf{sb}(Y)).$$

Then $(A, \mathsf{sb}, \mathsf{ib}) \in \langle \mathcal{C}(\mathcal{L}_2) \rangle (I \uplus I_2)$, SBWF holds, and since X and Y satisfy NONINTERF, MOWF holds as well. Furthermore, the projections of A, sb, ib and mo to client, respectively, library actions are $A(X)$, $\mathsf{sb}(X)$, $\mathsf{ib}(X)$, $\mathsf{mo}(X)$, respectively, $A(Y)$, $\mathsf{sb}(Y)$, $\mathsf{ib}(Y)$, $\mathsf{mo}(Y)$.

We let $\mathsf{rf} = \mathsf{rf}(X) \cup \mathsf{rf}(Y) \cup R_l \cup R_c$, where R_l contains edges relating library reads from param_t to the corresponding client writes, and R_c, edges relating client reads from retval_t to the corresponding library writes. We now define R_l formally (R_c can be defined in the same way). For every action $r = (t, \mathsf{read}(\mathsf{param}_t, a)) \in A(Y)$, R_l contains an edge defined as follows. By DETREAD$_l$ for Y, for some action $u \in A(Y)$ we have

$$u = (t, \mathsf{call} \_(a)) \wedge u \xrightarrow{\mathsf{sb}(Y)} r \wedge$$
$$\neg \exists v.\, v = (t, \mathsf{ret} \_) \wedge u \xrightarrow{\mathsf{sb}(Y)} v \xrightarrow{\mathsf{sb}(Y)} r.$$

Since $\mathsf{interf}(A(X)) = \mathsf{interf}(A(Y))$, we have $u \in A(X)$. By the restrictions on accesses to param_t and retval_t imposed in §4 and DETREAD$_c$ for X, there exists an action $w = (t, \mathsf{write}(\mathsf{param}_t, a)) \in X$ such that $w \xrightarrow{\mathsf{rf}(X)} u$. Then the corresponding edge in R_l is (w, r). RFWF holds for rf defined in this way, and its projections to client and library actions are $\mathsf{rf}(X)$ and $\mathsf{rf}(Y)$.

Let $\mathsf{sw} = \mathsf{sw}(X) \cup \mathsf{sw}(Y)$ and $\mathsf{hb} = ((\mathsf{ib} \cup \mathsf{sb} \cup \mathsf{sw})/\sim)^+$, so that HBDEF holds by construction. Since X and Y satisfy NONINTERF,

the projections of sw to client and library actions are $\mathsf{sw}(X)$ and $\mathsf{sw}(Y)$. Let hb_c, respectively, hb_l be the projections of hb to client, respectively, library actions. We now show that $\mathsf{hb}_c = \mathsf{hb}(X)$ (we can show $\mathsf{hb}_l = \mathsf{hb}(Y)$ analogously). Let

$$R = \mathsf{hbL}(((\mathsf{ib}(Y) \cup \mathsf{sb}(Y) \cup \mathsf{sw}(Y))/{\sim})^+) =$$
$$\mathsf{hbL}(((\mathsf{sb}(Y) \cup \mathsf{sw}(Y))/{\sim})^+).$$

By EXTHBDEF, we need to show that

$$\mathsf{hb}_c = ((\mathsf{ib}(X) \cup \mathsf{sb}(X) \cup \mathsf{sw}(X) \cup R)/{\sim})^+.$$

We have:

$$((\mathsf{ib}(X) \cup \mathsf{sb}(X) \cup \mathsf{sw}(X) \cup R)/{\sim})^+ \subseteq$$
$$((\mathsf{ib}(X) \cup \mathsf{sb}(X) \cup \mathsf{sb}(Y) \cup \mathsf{sw}(X) \cup \mathsf{sw}(Y))/{\sim})^+ \subseteq \mathsf{hb}.$$

Since $\mathsf{ib}(X)$, $\mathsf{sb}(X)$, $\mathsf{sw}(X)$ and R relate only client actions in A, this implies

$$((\mathsf{ib}(X) \cup \mathsf{sb}(X) \cup \mathsf{sw}(X) \cup R)/{\sim})^+ \subseteq \mathsf{hb}_c.$$

We now show

$$\mathsf{hb}_c \subseteq ((\mathsf{ib}(X) \cup \mathsf{sb}(X) \cup \mathsf{sw}(X) \cup R)/{\sim})^+.$$

Consider a path in

$$((\mathsf{ib}(X) \cup \mathsf{sb}(X) \cup \mathsf{sb}(Y) \cup \mathsf{sw}(X) \cup \mathsf{sw}(Y))/{\sim})^+$$

between two client actions. Like in the proof of Lemma 11 the case when a path has an edge from $\mathsf{ib}(X)$ is trivial, and otherwise, any segment of library actions on the path starts with a call and ends with a return. Hence, we can replace every such segment with an edge in R, obtaining a path in $(\mathsf{ib}(X) \cup \mathsf{sb}(X) \cup \mathsf{sw}(X) \cup R)^+$, as required.

Let $\mathsf{sc} = \mathsf{sc}(X) \cup \mathsf{sc}(Y)$. Let us now show that ACYCLICITY holds of Z. Assume now that there is a cycle in $\mathsf{hb} \cup \mathsf{sc}$ and consider the corresponding cycle with edges from

$$(\mathsf{sb}(X)/{\sim}) \cup (\mathsf{sb}(Y)/{\sim}) \cup (\mathsf{sw}(X)/{\sim}) \cup$$
$$(\mathsf{sw}(Y)/{\sim}) \cup \mathsf{sc}(X) \cup \mathsf{sc}(Y)$$

(it is easy to see that ib edges cannot participate in the cycle). As before, we can split $\mathsf{sb}/{\sim}$ edges to make all intermediate interface actions in the program order explicit. If all actions on the cycle come from only X or only Y, we get a contradiction with ACYCLICITY. Otherwise, the cycle contains at least two interface actions. Then, every segment of library actions on the cycle starts with a call and ends with a return; conversely, every segment of client actions starts with a return and ends with a call. We can thus convert every one of these segments into an edge from $(\mathsf{scC}(X))^{-1}$ or $(\mathsf{scL}(Y))^{-1}$ contradicting the assumptions of the lemma. Thus, ACYCLICITY holds. The sc relation constructed above also validates SCWF, and the projections of sc to client and library actions are $\mathsf{sc}(X)$ and $\mathsf{sc}(Y)$. This and the fact that the client and the library use disjoint sets of locks implies that LOCKS holds for Z.

All of the above shows that $\mathsf{client}(Z) = X$ and $\mathsf{lib}(Z) = Y$.

We now show that SWDEF holds of Z. First, take w and r such that $w \xrightarrow{\mathsf{sw}} r$. Then either $w \xrightarrow{\mathsf{sw}(X)} r$ or $w \xrightarrow{\mathsf{sw}(Y)} r$. Since the client and library projections of sc, mo and rf coincide with the corresponding relations in X and Y, we obtain the judgement on the right-hand side of the equivalence in SWDEF for Z. Assume now that the judgement holds for Z. Then by NONINTERF for X and Y, all the actions involved come from the same component. Hence, again by SWDEF for X and Y, we get $w \xrightarrow{\mathsf{sw}} r$.

If RFNONATOMIC is violated in Z, then either all actions involved into the violating configuration come from the same component, or r is a read from param_t or retval_t. In the former case, the violating configuration can be reproduced in X or Y, contradicting

their admissibility. In the latter case, it is easy to check that the axiom holds of Z by construction of rf. The validity of DETREAD for Z is justified similarly.

By Proposition 23, any violation of the rest of the axioms has actions coming from a single component. The edges from the relations involved exist in X or Y, so we get a violation in one of these executions. Hence, the rest of the axioms are valid for Z.

Thus, we have shown that $Z \in [\![\mathcal{C}(\mathcal{L}_2)]\!](I \uplus I_2)$. Assume now that X is unsafe, i.e., it violates DRF, $\mathrm{SAFEREAD}_c$ or SAFELOCK. Since $\mathsf{client}(Z) = X$ and $\mathsf{lib}(Z) = Y$, it is easy to show that in this case Z violates DRF, SAFEREAD or SAFELOCK (in the case of SAFELOCK this relies on the fact that the client and the library access disjoint sets of locks). $\qquad\square$

C.5 Proof of Proposition 14 (Absence of causal cycles)

Consider $w \xrightarrow{\mathsf{rf}(X)} r$. If the location accessed is non-atomic, then RFNONATOMIC implies that $w \xrightarrow{\mathsf{hb}(X)} r$. If the location is atomic, then $w \xrightarrow{\mathsf{hb}(X)} r$ according to SWDEF. Then $\mathsf{hb}(X) \cup \mathsf{rf}(X) \cup \mathsf{sc}(X) = \mathsf{hb}(X) \cup \mathsf{sc}(X)$. This is acyclic by ACYCLICITY.

We now show that $(\mathsf{hb}(X) \cup \mathsf{rf}(X) \cup \mathsf{sc}(X))^{-1}$ is well-founded (when X is valid or almost valid). Assume there is an infinite chain in this relation. Since a program has finitely many threads, there are infinitely many actions $u_1, u_2, \ldots$ by the same thread in the chain. Note that for i and j such that $i < j$, we cannot have $u_i \xrightarrow{\mathsf{sb}} u_j$, for this would contradict the acyclicity of $\mathsf{hb}(X) \cup \mathsf{rf}(X) \cup \mathsf{sc}(X)$. Since action structures of base commands are finite, by the definition in Figure 3, the above chain contains infinitely many actions related by sb one way or another; according to our previous observation, this has to be a chain in sb^{-1}. However, this relation is well-founded, which yields a contradiction. $\qquad\square$

C.6 Proof of Lemma 15 (Prefix construction)

We only consider the case when X is valid and is an execution of $\mathcal{C}(\mathcal{L})$. The cases of an almost-valid execution and an execution of $\mathcal{L}$ are handled similarly.

Let $X_p = X|_{A'}$. Since $A(X_p)$ is closed under $(\mathsf{sb}(X))^{-1}$, we know that $(A(X_p), \mathsf{sb}(X_p), \mathsf{ib}(X_p)) \in \langle \mathcal{C}(\mathcal{L})\rangle I$.

By Proposition 22, it is sufficient to check the axioms SWDEF, SCREADS, LOCKS, HBDEF and DETREAD. We have:

- If $r \in A(X_p)$ and $w \xrightarrow{\mathsf{rf}(X)} r$, then $w \in A(X_p)$. Furthermore, if $w \xrightarrow{\mathsf{rs}}_{t_1} w'$ in X and $w, w' \in A(X_p)$, then $w \xrightarrow{\mathsf{rs}}_{t_1} w'$ in X_p. Hence, SWDEF holds.

- If $\neg(w \xrightarrow{\mathsf{hb}(X)} w_1)$, then $\neg(w \xrightarrow{\mathsf{hb}(X_p)} w_1)$. Also, if $w_1 \xrightarrow{\mathsf{sc}} w_2 \xrightarrow{\mathsf{sc}} r$ and $r \in A(X_p)$, then $w_2, w_1 \in A(X_p)$. Hence, SCREADS holds.

- LOCKS is valid, since $A(X_p)$ is closed under $(\mathsf{sc}(X))^{-1}$.

- If $u \xrightarrow{\mathsf{hb}(X_p)} v$ and $v \in A(X_p)$, then any action on a corresponding path of $(\mathsf{ib}(X) \cup \mathsf{sb}(X) \cup \mathsf{sw}(X))/{\sim}$ edges from u to v is in $A(X_p)$. Hence, HBDEF holds.

- DETREAD holds, because $A(X_p)$ is closed under $(\mathsf{hb}(X))^{-1}$ and $(\mathsf{rf}(X))^{-1}$. $\qquad\square$

C.7 Proof of Proposition 17 (Execution reduction)

Let Y be the the extension of X with R'. It is easy to check that, if a validity axiom other than RFNONATOMIC, $\mathrm{DETREAD}_c$ or $\mathrm{DETREAD}_l$ is true of the extension of X with R, then it is also true of Y, and if a safety axiom is violated in the extension of X with R, then so it is in Y.

If $\textsc{DetRead}_c$ or $\textsc{DetRead}_l$ (depending on whether X is a client or a library execution) is violated by a read action, and RFNonAtomic is not violated by the same action, then we can satisfy $\textsc{DetRead}_c$ or $\textsc{DetRead}_l$ by dropping the offending edge from $\mathsf{rf}(Y)$. We then get either an admissible or an almost-admissible execution. □

C.8 Proof of Lemma 18 (Read completion)

We consider only the case when Y is an execution of $\mathcal{L}$; the case of $\mathcal{C}(\mathcal{L})$ is handled similarly.

There exist actions $w, r \in A(Y)$, such that $w \xrightarrow{\mathsf{rf}(Y)} r$, but it is not the case that $w \xrightarrow{\mathsf{hb}(Y)} r$. Let us chose the read action $r \in A(Y)$ violating RFNonAtomic in this way such that for any other such action $r' \in A(Y)$ we do not have $r' \xrightarrow{(\mathsf{hb}(Y) \cup \mathsf{rf}(Y) \cup \mathsf{sc}(Y))^+} r$. Such an action r exists, since $\mathsf{hb}(Y) \cup \mathsf{rf}(Y) \cup \mathsf{sc}(Y)$ is acyclic. Let w be the corresponding action such that $w \xrightarrow{\mathsf{rf}(Y)} r$.

Let
$$A' = \{u \mid u \xrightarrow{(\mathsf{hb}(Y) \cup \mathsf{rf}(Y) \cup \mathsf{sc}(Y))^*} r\},$$
and let $Y' = Y|_{A'}$. Note that $w, r \in A'$ and no other read action violating RFNonAtomic or $\textsc{DetRead}_l$ is in A'. Also, by Proposition 14, $\neg \exists u. r \xrightarrow{\mathsf{sb}(Y')} u$. By Lemma 15, $(A(Y'), \mathsf{sb}(Y'), \mathsf{ib}(Y')) \in \langle \mathcal{L} \rangle I$ and Y' satisfies all the validity axioms except r violates RFNonAtomic and possibly $\textsc{DetRead}_l$ by reading from w. Let $r = (\_, \mathsf{read}_{\mathsf{NA}}(x, a))$, then x is not param_t or retval_t.

If there is no action $w' = (\_, \mathsf{write}(x, \_)) \in A(Y')$ such that $w' \xrightarrow{\mathsf{hb}(Y')} r$, then we construct the execution Z from Y' by dropping the offending rf edge $w \xrightarrow{\mathsf{rf}(Y)} r$. The execution Z satisfies $\textsc{DetRead}_l$ and RFNonAtomic, and thus $Z \in [\![\mathcal{L}]\!]I$. We also have $\neg(w \xrightarrow{\mathsf{hb}(Z)} r)$ and $\neg(r \xrightarrow{\mathsf{hb}(Z)} w)$, which means that Z violates DRF.

Otherwise, there exists an action $w' = (\_, \mathsf{write}(x, b)) \in A(Y')$ such that for any other action $w'' = (\_, \mathsf{write}(x, \_)) \in A(Z)$ we do not have $w' \xrightarrow{\mathsf{hb}(Y')} w'' \xrightarrow{\mathsf{hb}(Y')} r$. In this case, we construct Z by replacing the action r with an action r' that reads b instead of a and keeping all relations isomorphic to the original ones, except the rf edge $w \xrightarrow{\mathsf{rf}(Y)} r$ is replaced by $w' \xrightarrow{\mathsf{rf}(Y)} r'$. The resulting execution satisfies $\textsc{DetRead}_l$ and RFNonAtomic. It is easy to check that the validity of the other axioms is not affected, and thus, the execution is valid. According to (1) in §B, $(A(Z), \mathsf{sb}(Z), \mathsf{ib}(Z)) \in (\![\mathcal{L}]\!)I$. We also have $\neg(w \xrightarrow{\mathsf{hb}(Z)} r')$ and $\neg(r' \xrightarrow{\mathsf{hb}(Z)} w)$, which means that Z violates DRF. □

C.9 Proof of Lemma 19 (Appending an action)

Let
$$X = (A, \mathsf{sb}, \mathsf{ib}, \mathsf{rf}, \mathsf{sc}, \mathsf{mo}, \mathsf{sw}, \mathsf{hb}).$$
We only consider the case when X is an execution of $\mathcal{C}(\mathcal{L})$.

If needed, we put the action u at the end of the corresponding suborders of mo and sc, obtaining mo' and sc'. Let
$$\mathsf{hb}' = (\mathsf{sb}' \cup \mathsf{ib}' \cup \mathsf{sw})^+$$
and
$$Y = (A', \mathsf{sb}', \mathsf{ib}', \mathsf{rf}, \mathsf{sc}', \mathsf{mo}', \mathsf{sw}, \mathsf{hb}').$$
Then SCwf, MOwf, MOvsSC and HBDEF hold of Y.

Since there are no outgoing sb' edges out of u, the only edges in $\mathsf{hb}' \setminus \mathsf{hb}$ are of the form $(\_, u)$. Hence, HBvsMO and ACYCLICITY hold.

Since there are no outgoing sc, mo, hb edges out of u, and no rf edges involving u, the following axioms hold of

Y: SWDEF, ATOMRMW, CoRR, CoWR, CoRW, RFAtomic, RFNonAtomic, SCReads. Furthermore, since X satisfies NonInterf and u violates it, there is no action v in X such that $\mathsf{sec}(u) = \mathsf{sec}(v)$ and v accesses the same location as u. Hence, ASmo, ASsc and ATOMAS hold.

Locks holds for Y, since it holds for X. Finally, SBwf and RFwf hold trivially.

Thus, Y satisfies all the validity axioms, except possibly DETREAD. Besides, $\mathsf{hb}(Y) \cup \mathsf{rf}(Y) \cup \mathsf{sc}(Y)$ is acyclic.

If u is a store, then DETREAD is satisfied, and Y is the desired execution X'. If u accesses a non-atomic location, then Y can be completed to a valid execution like in the proof of Lemma 18.

Consider now the case when $u = (e, g, t, \mathsf{read}_\lambda(x, a))$ and $\mathsf{sort}(x) = \mathsf{ATOM}$. If there is no action $w = (\_, \mathsf{write}(x, \_))$ such that $w \xrightarrow{\mathsf{hb}'} u$, then Y satifies DETREAD and is the desired X'. Otherwise, take the action $w = (\_, \mathsf{write}(x, b))$ that is the last write to x in mo', if u is a load, and the penultimate one, if u is a read-modify-write. If w is annotated with SC, by MOvsSC, it is also the last write to x in sc', except possibly u.

Let r be identical to u, except it reads b instead of a and $A'' = A' \setminus \{u\} \uplus \{r\}$. Let
$$X' = (A'', \mathsf{sb}'', \mathsf{ib}'', \mathsf{rf}'', \mathsf{sc}'', \mathsf{mo}'', \mathsf{sw}'', \mathsf{hb}''),$$
where sb'', ib'', sc'', mo'' are relations isomorphic to sb, ib, sc, mo, but with r instead of u, $\mathsf{rf}'' = \mathsf{rf} \cup \{(w, r)\}$, and sw'' and hb'' are computed from the other relations according to SWDEF and HBDEF. We have that $\mathsf{sw}'' \setminus \mathsf{sw} \subseteq A'' \times \{r\}$, and hence, $\mathsf{hb}'' \setminus \mathsf{hb}' \subseteq A'' \times \{r\}$ and ACYCLICITY and HBvsMO hold. By (1) in §B, $(A'', \mathsf{sb}'', \mathsf{ib}'') \in \langle \mathcal{C}(\mathcal{L}) \rangle I$. Furthermore, DETREAD holds of X'.

Since there are no outgoing edges from r in hb'', sc'' and mo'', RFAtomic, SCReads and CoRW hold of X'. By the choice of w, we also get that ATOMRMW, CoWR and CoRR hold.

Since X satisfies NONINTERF and u violates it, there is no action v in X' such that $\mathsf{sec}(r) = \mathsf{sec}(v)$ and v accesses the same location as r. Hence, ASmo, ASsc and ATOMAS hold. It is easy to see that the other validity axioms hold of X', and thus, it is the desired execution. □

C.10 Proof of Lemma 20 (Execution extension)

Let Z be the extension of Y with R. Then ExTHBDEF holds of Z.

Assume that ACYCLICITY does not hold of Z, i.e., $\mathsf{hb}(Z) \cup \mathsf{sc}(Y)$ has a cycle. Consider the corresponding cycle of edges in
$$(\mathsf{sb}(Y)/{\sim}) \cup (\mathsf{sw}(Y)/{\sim}) \cup R \cup \mathsf{sc}(Y)$$
($\mathsf{ib}(Y)$ edges cannot participate in it). If the cycle does not involve any edges in R, then it also exists in Y, contradicting its validity. As usual, we can split $\mathsf{sb}(Y)/{\sim}$ edges to make all intermediate interface actions in the program order explicit. We can also assume that the cycle contains at least two interface actions. We can then convert the cycle into one with edges from $(\mathsf{scL}(Y))^{-1} \cup R \subseteq (\mathsf{scL}(X))^{-1} \cup R$, which contradicts the admissibility of X extended with R.

For executions without relaxed actions, RFAtomic follows from ACYCLICITY, and so is valid.

Assume that RFNonAtomic does not hold of Z, i.e., for some w, r and x, we have
$$\mathsf{sort}(x) = \mathsf{NA} \wedge w \xrightarrow{\mathsf{rf}(Y)} r \wedge$$
$$w = (\_, \mathsf{write}(x, \_)) \wedge r = (\_, \mathsf{read}(x, \_))$$
and, additionally, for some w' we have
$$w' = (\_, \mathsf{write}(x, \_)) \wedge w \xrightarrow{\mathsf{hb}(Z)} w' \xrightarrow{\mathsf{hb}(Z)} r.$$

Then w and w', w' and r, if unordered by $\mathsf{hb}(Y)$ would form a race in Y, contradicting its safety. Hence, $w \xrightarrow{\mathsf{hb}(Y)} w'$ or $w' \xrightarrow{\mathsf{hb}(Y)} r$, $w' \xrightarrow{\mathsf{hb}(Y)} r$ or $r \xrightarrow{\mathsf{hb}(Y)} w'$. We cannot have $w' \xrightarrow{\mathsf{hb}(Y)} w$ or $r \xrightarrow{\mathsf{hb}(Y)} w'$, as this would give a rise to a cycle in $\mathsf{hb}(Z)$. Hence,

$$w \xrightarrow{\mathsf{hb}(Y)} w' \xrightarrow{\mathsf{hb}(Y)} r.$$

But this contradicts the validity of RFNonAtomic for Y.

Since Y is safe, by SafeRead$_l$, for any read action $r \in A(Y)$ accessing a location other than param_t, we have $\exists w.\, w \xrightarrow{\mathsf{rf}(Y)} r$. Hence, DetRead$_l$ holds for Z.

The only other validity axioms that can be affected by extending hb with R are SCReads, HBvsMO, CoRR, CoWR, CoRW. We show that these axioms are valid of Z using the fact that $\mathsf{denyL}(Y) \subseteq \mathsf{denyL}(X)$. We consider only the case of HBvsMO. The cases of CoWR and SCReads are considered similarly, the case of CoRR is covered by that of CoWR, and the case of CoRW, by that of HBvsMO.

Assume that HBvsMO does not hold for Z. Then for some w_1 and w_2, we have $w_1 \xrightarrow{\mathsf{hb}(Z)} w_2$ and $w_2 \xrightarrow{\mathsf{mo}(Y)} w_1$. Consider the path of edges from

$$(\mathsf{ib}(Y)/{\sim}) \cup (\mathsf{sb}(Y)/{\sim}) \cup (\mathsf{sw}(Y)/{\sim}) \cup R$$

yielding $w_1 \xrightarrow{\mathsf{hb}(Z)} w_2$. Let us split $\mathsf{sb}(Z)/{\sim}$ edges to make all intermediate interface actions in the program order explicit. The path must include edges from R, for otherwise HBvsMO would be violated in Y. Thus, the path includes at least two interface actions. Let u and v be the first, respectively, last interface action on it. Then u is a return, v is a call, and

$$r \xrightarrow{\mathsf{hb}(Y)} u \xrightarrow{\mathsf{hb}(Z)} v \xrightarrow{\mathsf{hb}(Y)} w$$
$$\underset{\mathsf{mo}(Y)}{\longleftarrow}$$

Hence, $(u, v) \in \mathsf{denyL}(Y) \subseteq \mathsf{denyL}(X)$, so that

$$r \xrightarrow{\mathsf{hb}(X)} u \qquad v \xrightarrow{\mathsf{hb}(X)} w$$
$$\underset{\mathsf{mo}(X)}{\longleftarrow}$$

As usual, using the transitivity of $\mathsf{hb}(Z)$, we can eliminate non-interface actions on the path connecting u and v in $\mathsf{hb}(Z)$, showing that

$$(u, v) \in (\mathsf{interf}(\mathsf{hb}(Y)) \cup R)^+.$$

Then by Propositions 16,

$$(u, v) \in (\mathsf{hbL}(Y) \cup R)^+ \subseteq (\mathsf{hbL}(X) \cup R)^+.$$

But this contradicts the admissibility of X extended with R.

Finally, the extension of Y with R is safe by Proposition 9. □

C.11 Proof of Proposition 21 (Composition generalisation)

The only non-trivial obligation to show is the acyclicity of $\mathsf{hb}(Z) \cup \mathsf{rf}(Z) \cup \mathsf{sc}(Z)$ in the case of almost-valid executions in Lemma 12. Consider the construction of the corresponding relations in this lemma and assume that $\mathsf{hb}(Z) \cup \mathsf{rf}(X) \cup \mathsf{sc}(Z)$ has a cycle. It is easy to show that any such cycle can be converted into one in

$$\mathsf{hb}(X) \cup \mathsf{rf}(X) \cup \mathsf{sc}(X)$$

or

$$\mathsf{EscC}(X) \cup \mathsf{scL}(Y)$$

contradicting the conditions of the proposition. Thus, $\mathsf{hb}(Z) \cup \mathsf{rf}(X) \cup \mathsf{sc}(Z)$ is acyclic. By Proposition 14, this implies the acyclicity of $\mathsf{hb}(Z) \cup \mathsf{rf}(Z) \cup \mathsf{sc}(Z)$. □

C.12 Proof of Theorem 4 (Compositionality)

Consider libraries $\mathcal{L}_1$ and $\mathcal{L}_2$ with the corresponding address spaces LLoc_1 and LLoc_2, such that $\mathsf{LLoc} = \mathsf{LLoc}_1 \uplus \mathsf{LLoc}_2$, and initial states $\mathcal{I}_1$ and $\mathcal{I}_2$, such that $\forall j = 1..k.\, \forall I \in \mathcal{I}_j.\, \mathsf{dom}(I) \subseteq \mathsf{LLoc}_j$. We write $\mathcal{L}_1 \uplus \mathcal{L}_2$ for the library obtained by concatenating their method declarations.

To simplify presentation, in our development we have assumed that any method called in a program by the client belongs to the library. Our setting and the proof of Theorems 3 and 7 can be simply generalised to allow the client to have its private methods (as before, nested method calls are disallowed).

We first show Theorem 4 for the case of programs without relaxed atomics.

LEMMA 24. *If* $(\mathcal{L}_1, \mathcal{I}_1)$ *and* $(\mathcal{L}_2, \mathcal{I}_2)$ *are safe, then so is* $(\mathcal{L}_1 \uplus \mathcal{L}_2, \mathcal{I}_1 \uplus \mathcal{I}_2)$.

This is easily shown using the techniques from the proof of Theorem 7.

LEMMA 25. *If* $(\mathcal{L}_2, \mathcal{I}_2) \sqsubseteq (\mathcal{L}_2', \mathcal{I}_2')$, *and* $(\mathcal{L}_1 \uplus \mathcal{L}_2, \mathcal{I}_1 \uplus \mathcal{I}_2)$ *and* $(\mathcal{L}_1 \uplus \mathcal{L}_2', \mathcal{I}_1 \uplus \mathcal{I}_2')$ *are safe, then*

$$(\mathcal{L}_1 \uplus \mathcal{L}_2, \mathcal{I}_1 \uplus \mathcal{I}_2) \preceq (\mathcal{L}_1 \uplus \mathcal{L}_2', \mathcal{I}_1 \uplus \mathcal{I}_2').$$

Theorem 25 is obtained by applying Lemma 25 repeatedly to abstract one library after another. Lemma 24 guarantees the safety of composed libraries arising during this construction.

Consider an execution X of $\mathcal{L}_1 \uplus \mathcal{L}_2$. We let $\mathsf{lib}_i(X)$ be the projection of X to actions of $\mathcal{L}_i$; this excludes calls to and returns from the other library. We let $\mathsf{client}_i(X)$ be the projection of X to actions excluding the internal actions by $\mathcal{L}_i$; this includes calls to and returns from both libraries.

PROOF OF LEMMA 25. Viewing $\mathcal{L}_1$ as a client and $\mathcal{L}_2$ as a library in $\mathcal{L}_1 \uplus \mathcal{L}_2$, from the above generalisation of Theorem 7, we obtain

$$\mathsf{client}_2(\llbracket \mathcal{L}_1 \uplus \mathcal{L}_2 \rrbracket(\mathcal{I}_1 \uplus \mathcal{I}_2)) \preceq \mathsf{client}_2(\llbracket \mathcal{L}_1 \uplus \mathcal{L}_2' \rrbracket(\mathcal{I}_1 \uplus \mathcal{I}_2')).$$

Consider $X_1 \in \llbracket \mathcal{L}_1 \uplus \mathcal{L}_2 \rrbracket(\mathcal{I}_1 \uplus \mathcal{I}_2)$ and a corresponding $X_2 \in \llbracket \mathcal{L}_1 \uplus \mathcal{L}_2' \rrbracket(\mathcal{I}_1 \uplus \mathcal{I}_2')$ such that $\mathsf{client}_2(X_1) \preceq \mathsf{client}_2(X_2)$. Then

$$\mathsf{interf}(X_1) = \mathsf{interf}(X_2), \quad \mathsf{hbL}(X_1) \subseteq \mathsf{hbL}(X_2),$$
$$\mathsf{scL}(\mathsf{core}(\mathsf{lib}_1(X_2))) \subseteq \mathsf{scL}(\mathsf{core}(\mathsf{lib}_1(X_1))),$$
$$\mathsf{denyL}(\mathsf{core}(\mathsf{lib}_1(X_2))) \subseteq \mathsf{denyL}(\mathsf{core}(\mathsf{lib}_1(X_1))).$$

We need to show that

$$\mathsf{scL}(X_2) \subseteq \mathsf{scL}(X_1), \quad \mathsf{denyL}(X_2) \subseteq \mathsf{denyL}(X_1).$$

We only show the former inclusion; the latter is handled similarly.

By SCwf, it is easy to check that for any valid and safe execution X,

$$\mathsf{scL}(X) = (\mathsf{scL}(\mathsf{core}(\mathsf{lib}_1(X))) \cup$$
$$\mathsf{scL}(\mathsf{core}(\mathsf{lib}_2(X))) \cup \mathsf{interf}(\mathsf{sb}(X)))^+.$$

Hence, we only need to show that

$$\mathsf{scL}(\mathsf{core}(\mathsf{lib}_2(X_2))) \subseteq \mathsf{scL}(\mathsf{core}(\mathsf{lib}_2(X_1))).$$

Consider $(u, v) \in \mathsf{scL}(\mathsf{core}(\mathsf{lib}_2(X_2)))$. We now trace the construction of X_2 from X_1 in the proof of Theorem 7. In this proof, $\mathsf{lib}_2(X_2)$ is constructed by extending an execution $Y_2 \in \llbracket \mathcal{L}_2' \rrbracket I_2'$. Thus,

$$(u, v) \in \mathsf{scL}(\mathsf{core}(\mathsf{lib}_2(X_2))) = \mathsf{scL}(\mathsf{lib}_2(Y_2)).$$

The execution Y_2 is obtained from an execution $Y_1 \in \llbracket \mathcal{L}_2 \rrbracket I_2$ such that $\mathsf{Ehistory}(Y_1) \preceq \mathsf{Ehistory}(Y_2)$. Thus,

$$(u, v) \in \mathsf{scL}(\mathsf{lib}_2(Y_2)) \subseteq \mathsf{scL}(\mathsf{lib}_2(Y_1)).$$

Finally, the execution $\mathsf{lib}_2(X_1)$ is obtained by extending Y_1. Hence,

$$(u, v) \in \mathsf{scL}(\mathsf{lib}_2(Y_1)) = \mathsf{scL}(\mathsf{core}(\mathsf{lib}_2(X_1))),$$

as required.   □

We now turn to the case of programs with relaxed atomics. For this case, we assume that the libraries considered in Theorem 4 and all their possible combinations satisfy NONINTERF (unlike above, where this could be checked modularly using Lemma 24).

PROOF OF LEMMA 25 IN THE PRESENCE OF RELAXED ATOMICS. Let R be a relation containing only edges from return actions to call actions. An appropriate generalisation of Theorem 3 yields

$$\mathsf{client}_2([\![\mathcal{L}_1 \uplus \mathcal{L}_2, R]\!](\mathcal{I}_1 \uplus \mathcal{I}_2)) \subseteq \mathsf{client}_2([\![\mathcal{L}_1 \uplus \mathcal{L}'_2, R]\!](\mathcal{I}_1 \uplus \mathcal{I}'_2)).$$

Consider $X_1 \in [\![\mathcal{L}_1 \uplus \mathcal{L}_2, R]\!](\mathcal{I}_1 \uplus \mathcal{I}_2)$ and a corresponding $X_2 \in [\![\mathcal{L}_1 \uplus \mathcal{L}'_2, R]\!](\mathcal{I}_1 \uplus \mathcal{I}'_2)$ such that $\mathsf{client}_2(X_1) = \mathsf{client}_2(X_2)$. Then

$$\mathsf{interf}(X_1) = \mathsf{interf}(X_2), \quad \mathsf{hbL}(X_1) = \mathsf{hbL}(X_2),$$
$$\mathsf{scL}(\mathsf{lib}_1(X_2)) = \mathsf{scL}(\mathsf{lib}_1(X_1)).$$

From the proof of Theorem 3, we can see that

$$\mathsf{scL}(\mathsf{lib}_2(X_2)) \subseteq \mathsf{scL}(\mathsf{lib}_2(X_1)).$$

All together, these imply

$$\mathsf{scL}(X_2) \subseteq \mathsf{scL}(X_1),$$

as required.   □

C.13   Proof of Proposition 8 (Equivalence of Different Model Formulations)

1. Let us construct X' from X by projecting $\mathsf{sc}(X)$ per location or lock. Suppose X' is invalid. We consider each of the axioms sensitive to $\mathsf{sc}(X')$:

- If SCWF$'$ is false, then SCWF could not have been true over X, violating our assumptions.

- If ACYCLICITY is false, then there is a cycle in $\mathsf{sc}(X') \cup \mathsf{hb}(X') \subseteq \mathsf{sc}(X) \cup \mathsf{hb}(X)$. By HBVSSC for X, the $\mathsf{hb}(X)$ edges in this cycle must all be in $\mathsf{sc}(X)$, and thus, we have a cycle in $\mathsf{sc}(X)$, violating SCWF for X.

- If SCREADS, LOCKS or MOVSSC or SWDEF is false in X', then it is false in X, because the validity of these axioms depends only on actions on the same location.

2. Noting that $\mathsf{sc}(X') \cup \mathsf{hb}(X')$ is acyclic, let X be the execution with $\mathsf{sc}(X)$ be a total order on SC and lock actions including $(\mathsf{sc}(X') \cup \mathsf{hb}(X'))^+$. Then SCWF and HBVSSC hold by construction. The only other axioms that are sensitive to sc are SCREADS, LOCKS, MOVSSC and SWDEF. However, if any of these axioms is false in X, then it is false in X', because its validity only depends on actions on the same location.

3. The only safety axiom sensitive to sc is SAFELOCK and it only depends on sc edges between actions on the same location or lock.   □

D.   Additional counterexamples

We now give counterexamples showing that we cannot strengthen Theorem 3 to remove differences 2 and 3 with Theorem 7 noted in §3, and that the classical linearizability does not imply our notion of library abstraction for the SC fragment of the language.

2. Weakening the guarantee in abstraction is unsound. Consider the following two libraries:

```
𝓛₁:                              𝓛₂:

  atomic int a;                    atomic int a;
  storeREL(&a,0);                  storeREL(&a,0);

  void wait() {                    void wait() {
    while (!loadACQ(&a)) ;           while(nondet()) ;
  }                                }

  void signal() {                  void signal() {
    storeREL(&a,1);                  ;
  }                                }
```

The library $\mathcal{L}_1$ is abstracted by $\mathcal{L}_2$ according to Definition 6. However, the following client $\mathcal{C}$ can print 1 when using $\mathcal{L}_1$, but not $\mathcal{L}_2$:

```
          atomic int x, y, z;
          int r1 = 0, r2 = 0, r3 = 0;
      storeSC(&x,0); storeSC(&y,0); storeSC(&z,0);

r3 = loadRLX(&z);  ‖ wait();              ‖ x = 1;
if (r3)            ‖ r1 = loadRLX(&y);    ‖ signal();
  storeRLX(&y,r3); ‖ if (r1)              ‖
                   ‖   r2 = x;            ‖
                   ‖ if (r2)              ‖
                   ‖   storeRLX(&z,1);    ‖
                   ‖ print r2;            ‖
```

The corresponding execution of $\mathcal{C}(\mathcal{L}_1)$ is:

Note that this includes a satisfaction cycle. Now, in the program $\mathcal{C}(\mathcal{L}_2)$, the read from x in the second thread happens only when r1 is assigned to 1. By tracing the sequence of assignments, we can check that this will only happen if we read 1 from x: in the presence of satisfaction cycles, we can set things up so that a read is only executed if it reads a given value! However, the access to x in the first thread cannot read 1 in $\mathcal{C}(\mathcal{L}_2)$, since the executions of the program have no hb edge from the write to x in the third thread to its read in the second. Hence, $\mathcal{C}(\mathcal{L}_2)$ is safe and cannot print 1.

3. Removing quantification over client happens-before edges R (even at the expense of adding the denies from §6.2 and §A) into histories is unsound. Consider the following library $\mathcal{L}_1$ with three methods:

```
          atomic int x, y, z;
          int r1 = 0, r2 = 0, r3 = 0;
      storeSC(&x,0); storeSC(&y,0); storeSC(&z,0);

int m1() {            void m2() {             void m3() {
  r3 = loadRLX(&z);     r1 = loadRLX(&y);       x = 1;
  if (r3)              if (r1)                 }
    storeRLX(&y,r3);     r2 = x;
}                      if (r2)
                        storeRLX(&z,1);
                      return r2;
                    }
```

We assume that a client of the library must invoke every library method at most once and in a separate thread.

Consider the following client, which is reproduces the key behaviours of the most general client of the library:

$$\texttt{m1} \parallel \texttt{m2} \parallel \texttt{m3}$$

This program is safe, despite non-atomic accesses to x in both m1 and m2. Like in the example above, the read from x in m1 happens only when r1 is assigned to 1, which will happen only if we read 1 from x. However, in the above program, the access to x in m1 cannot read 1, since the executions of the program have no hb edge from the write to x in m3 to its read in m1. For the same reason, m1 always returns 0. Hence, $\mathcal{L}_1$ is safe. It is easy to check that it is abstracted by the following specification $\mathcal{L}_2$, according to Definition 6:

```
         atomic int x, y, z;
         int r1 = 0, r2 = 0, r3 = 0;
   storeSC(&x,0); storeSC(&y,0); storeSC(&z,0);

  int m1() {       void m2() {        void m3() {
    ;                return 0;          ;
  }                }                  }
```
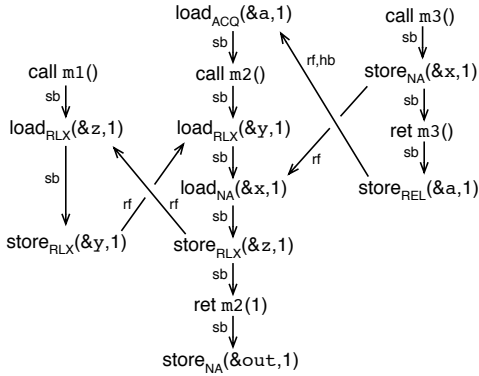
Now consider the following client $\mathcal{C}$ that makes m3 happen before m1:

```
            atomic int a;
            int out = 0;
            storeSC(&a,0);

m1();  ||   while (!loadACQ(&a));  ||   m3();
       ||   if (m2())              ||   storeREL(&a,1);
       ||     out = 1;            ||
       ||   print out;           ||
```

The program $\mathcal{C}(\mathcal{L}_2)$ cannot print 1, but $\mathcal{C}(\mathcal{L}_1)$ can:

Hence, in the presence of satisfaction cycles, the most general client is not most general: removing an hb edge enforced by the client can cause the library not to become unsafe, but just to return different values.

Classical linearizability does not imply library abstraction. Consider the following pair of libraries $\mathcal{L}_1$ and $\mathcal{L}_2$:

```
$\mathcal{L}_1$:  atomic int x;    $\mathcal{L}_2$:  atomic int x;
      int m() {                int m() {
        ;                        storeSC(&x,0);
      }                        }
```

$\mathcal{L}_1$ is linearized by $\mathcal{L}_2$ according to the classical definition. However, $\mathcal{L}_1$ is not abstracted by $\mathcal{L}_2$, either according to Definition 2 or Definition 6. This is because, in a history with more than one call to m, the component given by scL is empty for $\mathcal{L}_1$, but non-empty for $\mathcal{L}_2$. Our notion of abstraction distinguishes between the two libraries, since the presence of SC operations inside a method can prohibit certain relaxations for its clients using weaker memory operations.

E. Proofs of Example Algorithms

In this section, we give proofs of correctness for two algorithms: the lock-free Treiber stack, and a producer-consumer queue.

E.1 Treiber's stack

The Treiber stack specification and implementation are given in §2.

E.1.1 Execution Axiomatisation

We start by giving axiomatisations of the stack specification and implementation. These provide simple mathematical interfaces to the underlying semantics, and make proof simpler. We give the axiomatisations as action structure shapes. These shapes are parameterised with their start and end actions, and their important values, allowing them to be substituted for one another in place. Substitution of specification executions for implementation executions plays an important role in the proof of abstraction below.

Specification A successful push has the following shape:

$$S^{a,j}_{\mathbf{pushS},t,l,v_t} \;\triangleq\; (\{a,\dots b\},\; a \xrightarrow{\mathsf{sb}} \dots \xrightarrow{\mathsf{sb}} b)$$

where

$$
\begin{aligned}
a:\quad & \mathsf{call}(\mathrm{push}, v_t)\\
& \mathsf{store_{NA}}(b, \mathrm{nondet})\\
& \mathsf{load_{NA}}(b, 0)\\
& \mathsf{load_{RLX}}(\&S, l)\\
& \mathsf{store_{NA}}(s, l)\\
& \mathsf{load_{NA}}(s, s_t)\\
& \mathsf{load_{NA}}(v, v_t)\\
& \mathsf{store_{REL}}(s2, \mathrm{append}(s_t, v_t))\\
& \mathsf{rmw_{RLX,REL}}(\&S, s, s2)\\
b:\quad & \mathsf{ret}(\mathrm{push})
\end{aligned}
$$

A blocking failed push has the following shape:

$$S^{a}_{\mathbf{pushF},t,v_i} \;\triangleq\; (\{a,\dots b\},\; a \xrightarrow{\mathsf{sb}} \dots \xrightarrow{\mathsf{sb}} b)$$

where

$$
\begin{aligned}
a:\quad & \mathsf{call}(\mathrm{push}, v)\\
& \mathsf{store_{NA}}(b, \mathrm{nondet})\\
& \mathsf{load_{NA}}(b, i) \text{ such that } i \neq 0\\
b:\quad & \mathrm{block}
\end{aligned}
$$

A succussful pop has the following shape:

$$S^{a,b}_{\mathbf{popS},t,l} \;\triangleq\; (\{a,\dots b\},\; a \xrightarrow{\mathsf{sb}} \dots \xrightarrow{\mathsf{sb}} b)$$

where

$$
\begin{aligned}
a:\quad & \mathsf{call}(\mathrm{pop})\\
& \mathsf{store_{NA}}(b, \mathrm{nondet})\\
& \mathsf{load_{NA}}(b, 0)\\
& \mathsf{load_{ACQ}}(\&S, l)\\
& \mathsf{store_{NA}}(s, l)\\
& \mathsf{load_{NA}}(s, l')\\
& \mathsf{store_{NA}}(s2, \mathrm{tail}(l'))\\
& \mathsf{rmw_{REL}}(\&S, s, s2)\\
b:\quad & \mathsf{ret}(\mathrm{pop}, \mathrm{head}(s))
\end{aligned}
$$

A blocking failed pop has the following shape:

$$S^{a}_{\mathbf{popF},t} \;\triangleq\; (\{a,\dots b\},\; a \xrightarrow{\mathsf{sb}} \dots \xrightarrow{\mathsf{sb}} b)$$

where

$$
\begin{aligned}
a:\quad & \mathsf{call}(\mathrm{pop})\\
& \mathsf{store_{NA}}(b, \mathrm{nondet})\\
& \mathsf{load_{NA}}(b, 0)\\
b:\quad & \mathrm{block}
\end{aligned}
$$

An empty failed pop has the following shape

$$S^{a,b}_{\mathbf{popE},t,l} \;\triangleq\; (\{a\dots b\},\; a \xrightarrow{\mathsf{sb}} \dots \xrightarrow{\mathsf{sb}} b)$$

where

$$
\begin{aligned}
a:\ & \mathsf{call}(\mathsf{pop}) \\
& \mathsf{store}_{\mathsf{NA}}(b, \mathrm{nondet}) \\
& \mathsf{load}_{\mathsf{NA}}(b, 0) \\
& \mathsf{load}_{\mathsf{ACQ}}(\&S, l) \\
& \mathsf{store}_{\mathsf{NA}}(s, l) \\
& \mathsf{load}_{\mathsf{NA}}(s, empty) \\
b:\ & \mathsf{ret}(\mathsf{pop}, EMPTY)
\end{aligned}
$$

Implementation   We first define the following action structure corresponding to a single failed iteration of the central loop of push:

$$
L^{a,b} \triangleq \quad
\begin{aligned}
a:\ & \mathsf{load}_{\mathsf{RLX}}(s, s_t) \\
& \mathsf{store}_{\mathsf{NA}}(t, s_t) \\
& \mathsf{load}_{\mathsf{NA}}(t, t_t) \\
& \mathsf{load}_{\mathsf{NA}}(x, x_t) \\
& \mathsf{store}_{\mathsf{NA}}(x_t \to \mathrm{next}, t_t) \\
& \mathsf{load}_{\mathsf{NA}}(t, t_t') \\
& \mathsf{load}_{\mathsf{RLX}}(s, s_t') \\
& \mathsf{store}_{\mathsf{NA}}(b, 0) \\
b:\ & \mathsf{load}_{\mathsf{NA}}(b, 0)
\end{aligned}
$$

We write $L_n^{a,z}$ for an action structure consisting of n iterations of L, connected end to end (ie, $L^{a,b} \cup L^{b,c} \cup L^{c,d} \ldots \cup L^{y,z}$). We write L_*^a for an infinite chain.

We define a successful execution as

$$
I_{\mathbf{pushS},s,t,l,v_t}^{a,} \triangleq (\{a, \ldots, s\} \cup L_n^{x,y}, a \xrightarrow{\mathsf{sb}}{}^* b \xrightarrow{\mathsf{sb}} x \xrightarrow{\mathsf{sb}}{}^* y \xrightarrow{\mathsf{sb}} c \xrightarrow{\mathsf{sb}}{}^* d)
$$

where the initial prefix of the action structure is defined as follows:

$$
\begin{aligned}
a:\ & \mathsf{call}(\mathsf{push}, v_t) \\
& \mathsf{store}_{\mathsf{NA}}(n_t \to \mathrm{data}, \_) \\
& \mathsf{store}_{\mathsf{NA}}(n_t \to \mathrm{next}, \_) \\
& \mathsf{store}_{\mathsf{NA}}(x, n_t) \\
& \mathsf{load}_{\mathsf{NA}}(v, v_t) \\
& \mathsf{load}_{\mathsf{NA}}(x, x_t) \\
b:\ & \mathsf{store}_{\mathsf{NA}}(x_t \to \mathrm{data}, v_t)
\end{aligned}
$$

and the final, successful iteration of the loop, is defined as follows:

$$
\begin{aligned}
c:\ & \mathsf{load}_{\mathsf{RLX}}(s, l) \\
& \mathsf{store}_{\mathsf{NA}}(t, l) \\
& \mathsf{load}_{\mathsf{NA}}(t, t_t) \\
& \mathsf{load}_{\mathsf{NA}}(x, x_t) \\
& \mathsf{store}_{\mathsf{NA}}(x_t \to \mathrm{next}, t_t) \\
& \mathsf{load}_{\mathsf{NA}}(t, t_t') \\
& \mathsf{load}_{\mathsf{NA}}(x, x_t') \\
& \mathsf{rmw}_{\mathsf{ACQ},\mathsf{REL}}(s, t_t') \\
& \mathsf{store}_{\mathsf{NA}}(b, 1) \\
& \mathsf{load}_{\mathsf{NA}}(b, 1) \\
d:\ & \mathsf{ret}(\mathsf{push})
\end{aligned}
$$

We define an action structure for a failed push as follows:

$$
I_{\mathbf{pushF},s,t}^{a} \triangleq (\{a, \ldots, \} \cup L_*^c, a \xrightarrow{\mathsf{sb}}{}^* b \xrightarrow{\mathsf{sb}} c \xrightarrow{\mathsf{sb}}{}^*)
$$

where the initial prefix of the action structure is defined as follows (identically to a successful push):

$$
\begin{aligned}
a:\ & \mathsf{call}(\mathsf{push}, v_t) \\
& \mathsf{store}_{\mathsf{NA}}(n_t \to \mathrm{data}, \_) \\
& \mathsf{store}_{\mathsf{NA}}(n_t \to \mathrm{next}, \_) \\
& \mathsf{store}_{\mathsf{NA}}(x, n_t) \\
& \mathsf{load}_{\mathsf{NA}}(v, v_t) \\
& \mathsf{load}_{\mathsf{NA}}(x, x_t) \\
b:\ & \mathsf{store}_{\mathsf{NA}}(x_t \to \mathrm{data}, v_t)
\end{aligned}
$$

For pop a failed loop iteration has the following shape :

$$
M^{a,b} \triangleq \quad
\begin{aligned}
a:\ & \mathsf{load}_{\mathsf{ACQ}}(s, s_t) \\
& \mathsf{store}_{\mathsf{NA}}(t, s_t) \\
& \mathsf{load}_{\mathsf{NA}}(t, t_t) \text{ such that } t_t \neq \mathrm{NULL} \\
& \mathsf{load}_{\mathsf{NA}}(t, t_t') \\
& \mathsf{load}_{\mathsf{NA}}(t \to \mathrm{next}, t_t'') \\
& \mathsf{store}_{\mathsf{NA}}(x, t_t'') \\
& \mathsf{load}_{\mathsf{RLX}}(s, s_t') \\
& \mathsf{load}_{\mathsf{NA}}(t, t_t''') \\
& \mathsf{store}_{\mathsf{NA}}(b, 0) \\
b:\ & \mathsf{load}_{\mathsf{NA}}(b, 0)
\end{aligned}
$$

We write $M_n^{a,z}$ for an action structure consisting of n iterations of M (ie, $M^{a,b} \cup M^{b,c} \cup M^{c,d} \ldots \cup M^{y,z}$). We write M_*^a for an infinite chain.

For a successful pop, we have the following action structure:

$$
I_{\mathbf{popS},t,s,v}^{a,p} \triangleq (\{a \ldots c\} \cup M_n^{q,r}, \ a \xrightarrow{\mathsf{sb}} q \xrightarrow{\mathsf{sb}}{}^* r \xrightarrow{\mathsf{sb}} b \xrightarrow{\mathsf{sb}}{}^* c)
$$

The initial action is just the call to pop.

$$
a:\ \mathsf{call}(\mathsf{pop})
$$

The suffix corresponding to the successful iteration of the main loop is defined as follows:

$$
\begin{aligned}
b:\ & \mathsf{load}_{\mathsf{ACQ}}(s, s_t) \\
& \mathsf{store}_{\mathsf{NA}}(t, s_t) \\
& \mathsf{load}_{\mathsf{NA}}(t, t_t) \text{ such that } t_t \neq NULL \\
& \mathsf{load}_{\mathsf{NA}}(t, t_t') \\
& \mathsf{load}_{\mathsf{NA}}(t \to \mathrm{next}, t_t'') \\
& \mathsf{store}_{\mathsf{NA}}(x, t_t'') \\
& \mathsf{load}_{\mathsf{NA}}(t, t_t''') \\
& \mathsf{load}_{\mathsf{NA}}(x, x_t) \\
& \mathsf{rmw}_{\mathsf{RLX},\mathsf{REL}}(s, t_t''', x_t) \\
& \mathsf{store}_{\mathsf{NA}}(b, 1) \\
& \mathsf{load}_{\mathsf{NA}}(b, 1) \\
& \mathsf{load}(t, t_t'''') \\
& \mathsf{load}(t_t'''' \to \mathrm{data}, r) \\
c:\ & \mathsf{ret}(\mathsf{pop}, r)
\end{aligned}
$$

For a blocking pop, we have the following action structure:

$$
I_{\mathbf{popF},t}^{a,b} \triangleq (\{a\} \cup M_*^b, \ a \xrightarrow{\mathsf{sb}} b \xrightarrow{\mathsf{sb}}{}^*)
$$

For a failed empty pop we have

$$
I_{\mathbf{popE},t,s}^{a,p} \triangleq (\{a \ldots c\} \cup M_n^{q,r}, \ a \xrightarrow{\mathsf{sb}} q \xrightarrow{\mathsf{sb}}{}^* r \xrightarrow{\mathsf{sb}} b \xrightarrow{\mathsf{sb}}{}^* c)
$$

The suffix corresponding to a failed iteration of the central loop is defined as follows:

$$
\begin{aligned}
b:\ & \mathsf{load}_{\mathsf{ACQ}}(s, s_t) \\
& \mathsf{store}_{\mathsf{NA}}(t, s_t) \\
& \mathsf{load}_{\mathsf{NA}}(t, \mathrm{NULL}) \\
c:\ & \mathsf{ret}(\mathsf{pop}, EMPTY)
\end{aligned}
$$

E.1.2   Safety

We first define a total order $<$ over calls. We choose this total order using modification order over the final rmw event in each successful push and pop. That is $I_1 < I_2$ if for the rmw events $a_1 \in I_1$ and $a_2 \in I_2$, $a_1 \xrightarrow{\mathsf{mo}} a_2$. For failing empty pop calls, we use the rforder between rmw and acquire-reads. That is, we add such operations at the earliest point in the order such that it reads from no later operations in the order. Blocking push and pop operations have no rmw events, and so can be included anywhere in the order.

Let X_{imp} be an arbitrary but fixed execution of the library implementation. Define the following predicate to help with proving safety:

$Q(c) \triangleq$ Either $c = \text{null}$, or

- c is a location.

- $c \to \text{next}$ and to $c \to \text{val}$ are in LLoc.

- There is exactly one write in the execution to $c \to \text{next}$ and to $c \to \text{val}$.

- The write to $c \to \text{next}$ has value n for which the predicate $Q(n)$ holds.

LEMMA 26. *Let v be the value observed by any atomic read in X_{imp}. Then $Q(v)$ holds in X_{imp}.*

Proof. By induction over $<$. $\square$

LEMMA 27. *Any read from T in an implementation pop is hb ordered with respect to any ¡-preceding write to T in push.*

Proof. All writes are CASes. Writes in push are annotated release, while reads in pop are annotated acquire. The result follows by appeal to the semantics of release sequences. $\square$

LEMMA 28. *The Treiber implementation $(L_1, \mathcal{I}_1)$ is safe over the MGC.*

Consider the safety axioms that could be violated.

NONINTERF. By Lemma 26, all locations pointed to in the stack at any point in the algorithm are in LLoc. Other locations are method-local, and so striaightforward to prove.

SAFELOCK. There are no locks in the implementation.

SAFEREAD. Follows immediately from the fact that initialisation hb-precedes all library calls.

DRF. All reads and writes are either atomic or local to the call, aside from reads and writes from the next and val fields of nodes. Writes to these locations are succeeded by a write to T, the head of the list, while reads from this location are preceded by reads from T. All such reads-writes to T are ordered in hb, by Lemma 27. Consequently, any pair consisting of a read and write to the next or val field of a node are also hb related. Each node is written by exactly one action, as nodes are never reused. Therefore, we need not worry about write-write races. $\square$

LEMMA 29. *The Treiber specification $(L_1, \mathcal{I}_1)$ is safe over the MGC.*

Proof. NONINTERF is satisfied as the algorithm accesses a fixed set of locations. There are no locks, so SAFELOCK does not apply. SAFEREAD follows by the same argument as for the implementation. DRF is satisfied trivially as there are no non-atomic reads or writes. $\square$

E.1.3 Abstraction

Constructing a Specification Execution

DEFINITION 30. *Let R be some arbitrary return-to-call relation. Let $X_{\text{imp}} \in [\![L_1, R]\!]\mathcal{I}_1$ be an arbitrary, fixed execution of the Treiber stack implementation under the MGC. We construct a corresponding execution X_{spec} of the specification as follows.*

We first construct an action structure C_{spec}. Let

$$X_{\text{imp}} = (A_i, \text{sb}_i, \text{ib}_i, \text{rf}_i, \text{sc}_i, \text{mo}_i, \text{sw}_i, \text{hb}_i)$$

and let $C_{\text{imp}} = (A_i, \text{sb}_i, \text{ib}_i)$ be the corresponding action structure. We construct C_{spec} by substituting successful implementation push operations with successful specification push operations, blocking implementation pushs with blocking specification pushs, and so on.

We use $<$ to construct a function k on actions:

- if $I^{a,b}$ is an init call, $k(a) = [\,]$;

- if $I^{a,b}$ is a successful push of v, and $I^{a',b'}$ is the preceding call, then $k(a) = \text{append}(k(a'), v)$.

- if $I^{a,b}$ is a successful pop, and $I^{a',b'}$ is the preceding call, then $k(a) = \text{tail}(v)$.

- otherwise if $I^{a,b}$ is a call and $I^{a',b'}$ is it's predecessor, $k(a) = k(a')$.

Starting with C_{imp}, we construct the action structure as in the previous proof, by subsituting calls. For each successful push, we substitute a specification for an implementation:

$$C' = C[S^{x,y}_{\text{pushS},t,v,k(x)}/I^{x,y}_{\text{pushS},t,v}]$$

We do the same things with successful pop and blocking push and pop. The result is an action structure C_{spec}.

LEMMA 31. *C_{spec} is an action structure for the Treiber stack specification. That is $C_{\text{spec}} \in \langle L_2 \rangle \mathcal{I}_2$.*

Proof. Straightforward from the structure of the MGC. $\square$

We now construct the execution X_{spec}. We take the action set and hb and ib relations from C_{spec}, and construct the rest of the relations as follows:

- rf over location S is dictated by $<$; that is, the call reads from the write in its most recent successful predecessor. All other locations are thread-local and dictated by sb.

- mo over location S is dictated by $<$.

- sc is empty, as there are no SC actions.

- sw is determined by rf and the SWDEF axiom. Note that release sequences will generate synchronisation.

- hb is determined by the HBDEF axiom.

To relate the values observed in the specification execution to the implementation executions X_{imp}, we define a predicate $P: \text{Act} \times \text{Loc} \times \text{List}$ as the least fixed point of the following recursive definition:

$P(c, xs) \triangleq$ Either $c = \text{null}$, or

- c is a location, and xs is a list of values.

- There exists exactly one write to $c \to \text{next}$ and to $c \to \text{val}$.

- The write to $c \to \text{val}$ has value $\text{head}(xs)$, and the write to $c \to \text{next}$ has value n for which the predicate $P(s, n, \text{tail}(xs))$ holds.

LEMMA 32. *Let $I^{a,b}$ be a successful push or pop operation, and $S^{a,b}$ the corresponding specification operation. In $I^{a,b}$ pick the final rmw operation (call this r) and in $S^{a,b}$ similarly pick the rmw operation. Let c be the value read in the implementation and let xs be the value read in the specification. The two values read in the executions are related by the predicate $P(c, xs)$.*

Proof. Prove the result by induction down the order $>$. The first operation in the order reads $xs = [\,]$ and $c = \text{null}$ from the initialisation, so the property holds trivially.

Now consider an arbitrary call. As $<$ has been constructed to follow mo, the rmw r must read from it's most recent successful $<$-predecessor. By construction, the corresponding specification call must also read from it's $<$-predecessor. Now case-split on the type of the $<$-preceding successful call: either push or pop.

Suppose the preceding successful call is a push. The implementation call pushes a new node onto the stack, while the specification appends the same value onto the list. We only need to show that the

new value written by the rmw satisfies the predicate. The call performs no writes between the initialisation of the new node being pushed onto the stack and the rmw (and so no other inter-thread hb edges). It is easy to see from this and the structure of the thread that the new node satisfies the predicate for the value being pushed onto the stack.

Suppose instead the preceding successful call is a pop. The implementation call simply redirects the top pointer to the next node in the stack. The predicate therefore holds immediately by its inductive structure. □

LEMMA 33. X_{spec} *is valid.*

PROOF.

SBWF. Trivial by the structure of the substituion.

RFWF. This is satisfied by construction for all reads and writes aside from the return value for a successful pop. For this, we need to show that the value written to the return value is the same as the one read in the client. Lemma 32 shows that the sequence of values readable from the stack in any call corresponds to the values in the specification list. We can therefore conclude that the return value is set correctly.

SCWF and SCREADS. Irrelevant, as there are no SCactions.

MOWF. By construction only atomic write actions are ordered in mo.

ATOMRMW. By construction, mofollows $<$, which is dictated by rf.

LOCKS. Irrelevant, as there are no locks.

BLOCK. By construction, block actions have no sc-successors.

RFATOMIC. By construction, hb over the internal actions of library calls must agree with rf: intra-thread rf edges are ordered by sb, while intra-thread hb-ordering must follow $<$. For rf to disagree with hb, there must be a client return-to-call order in R between two calls which disagrees with $<$. But any such order that could violate RFATOMIC in the specification execution would also violate RFATOMIC in X_{imp}. The result thus follows by contradiction.

ACYCLICITY and HBVSMO. By the same argument as RFATOMIC: violating the axiom in the constructed specification execution would mean that the axiom was violated in the implementation execution, contradicting the assumption.

CORW, COWR and CORR. By construction, reads follow mo, satisfying CORW. Any execution violating COWR or CORR would violate the same axioms in the implementation execution.

RFNONATOMIC. Trivial as all non-atomic stores are local to a single call.

DETREAD. Satisfied trivially by the construction, and by the presence of the initialisation operation hb-preceding any of the calls.

ASMO and ASSC. There are no single atomic sections with multiple mo-related actions, and there are no sc-related actions at all.

ATOMAS. For the first consequent, both rf and mo are constructed to follow $<$, which ensures that they conform to the axiom. The second consequent requires that all the reads in an atomic section to the same location read from the same action. This holds by construction for our execution, because all reads follow $<$. The third consequent is satisfied trivially because each atomic section contains at most one write. □

LEMMA 34. $X_{\text{spec}} \in [\![L_2]\!]\mathcal{I}_2$.

PROOF. Appeal to Lemmas 31 and 33. □

Inclusion over histories

LEMMA 35. $\text{history}(X_{\text{imp}}) \sqsubseteq \text{history}(X_{\text{spec}})$

Proof. We need not consider any deny edges, as these are only induced by sc orders, which are absent in both X_{imp} and X_{spec}.

For guarantee edges, we must show that the projections of hb to inter-thread call-to-return and intrathread interface actions are the same for X_{imp} and X_{spec}.

Intrathread hb is trivial: by construction of the X_{spec} we have the same sb over calls and returns.

For inter-thread hb, we need not consider blocking calls, as they do not write and do not have a return action. In successful calls, hb edges are only generated by either reads from S or client dependencies in R. Client dependencies are by construction also in X_{imp}. All successful calls are ordered by mo over their final write to S. By the semantics of release sequences, there exists an sw edge from call A to return B if and only iff A is a successful push and B is a successful pop such that $A < B$. The same holds for the implementation.

As failed-on-empty-stack implementation pop calls do not write, they generate no outgoing sw edges, and thus no non-client inter-thread hb edges. The same holds for the specification. They generate an incoming rfedge from the $<$-preceding successful push or pop operation, and by the semantics of release sequences, a sw edge from all $<$-preceding successful push operations. (Other sw edges generated by failing reads are subsumed into the sw edges generated by the release sequence). The same is true of the specification pop, by the semantics of release sequences. □

E.2 Producer-Consumer

In this section we present a producer-consumer buffer. The library is to be used by a single producer thread to push items into a buffer to be processed by a single consumer thread. The interface in this library has two method calls: enq() and deq(). We provide a specification of the queue, a C/C++11 implementation, and a proof that the implementation linearizes the specification.

E.2.1 Algorithm

The producer-consumer queue is intended for communication between a single producer thread and a single consumer thread. The library has three method calls: init(), enq() and deq(). Figure 10 gives a specification and implementation.

The specification and implementation are written in the subset of the language without relaxed atomics. This allows us to use the stronger abstraction theorem. In this setting, values cannot be communicated between threads without synchronising or creating a race. In the producer-consumer example, we do not want the specification to require the creation of happens-before edges from the dequeueing thread to the enqueueing thread. In order to avoid communication in that direction, we make the enqueueing thread unaware of the progress of the dequeueing thread, and allow spurious failure instead.

Specification. To help with specification, we assume that locations can store arbitrarily large abstract datatypes – in this case, an unbounded list representing all previously enqueued values. We assume functions empty_list(), append(L,m), length(L) and index(L,i) for manipulating lists. Such unbounded datatypes would of course be impractical to implement, but serve very well as a specification idiom.

To ensure that the consumer thread observes the most recent version of enqueued objects, the specification generates happens-before edges from the enq() that adds a particular item, to the

SPECIFICATION:

```
atomic list l;
int c;


int init() {
  l = empty_list();
  c = 0;
}

int enq(T* k) {
  if (nondet()) return 0;
  else {
    list lt = load_ACQ(l);
    lt = append(lt, k);
    store_REL(l, lt);
  }
}



int deq() {
  list lt = load_ACQ(l);
  if (length(lt) <= c)
    return 0;
  else {
    T* k = index(lt, c);
    ++c;
    return k;
  }
}
```

IMPLEMENTATION:

```
int s = N;
T a[s];
atomic int b, f;

int init() {
  b = 0;
  f = 0;
}

int enq(T* k) {
  int ft = load_ACQ(f);
  int bt = load_ACQ(b);
  if (ft <= bt-s) return 0;
  else {
    a[bt % s] = k;
    ++bt;
    store_REL(b, bt);
  }
}

int deq() {
  int ft = load_ACQ(f);
  int bt = load_ACQ(b);
  if (bt <= ft) return 0;
  else {
    T* k = a[ft % s];
    ++ft;
    store_REL(f, ft);
    return k;
  }
}
```

Figure 10. The producer-consumer queue.

deq() that removes it. It does this using release-acquire synchronisation on the location l holding the list of enqueued values. The specification avoids unwanted synchronisation from deq() to enq() by not mutating l in deq(); instead it maintains a counter c in deq() marking the current position in the list of enqueued items.

The specification models the possibility of failure due to the queue being full or empty. Emptiness is modelled by checking the counter c, but fullness is modelled by simple nondeterministic failure to avoid committing to a particular buffer size.

Implementation. The implementation uses a fixed-size circular array to store enqueued items. A *back* counter b tracks the next cell to be filled, and a *front* counter f tracks the next cell to be read. Both enq() and deq() read both counters before they proceed to check for fullness / emptyness. Note that this means the implementation synchronises from deq() to enq(), unlike the specification.

This implementation is more complex than the specification in two ways. First, each element of the queue is stored in a separate location, with its own synchronisation; in the specification, all elements of the array are part of the same location. Second, the array is finite and cyclic, and operations must avoid reading stale queue items or overwriting current items; in the specification, the list of enqueued items is infinite.

Note that this implementation is not sequentially consistently, and as a result it does not necessarily make progress. The locations f and b can produce a variant on the store-buffering example from §2. In the extreme, both calls can repeatedly fail, with enq() reading a value for f that makes the array appear full, and deq() reading a value for b that makes it appear empty.

E.2.2 Protocol

The client must obey a protocol for the specification to be valid: all initialisation must be called from a master thread that happens-before all other threads, calls to enq must be on the same thread and all calls to deq must be on one other thread.

Let $lm(a)$ be a mapping from an action a to the name of the library call that produced them, or the client: in this case: **enq**, **deq** or **client**. Define the protocol as a filter on a set of executions X, given the mapping lm:

$$sl\_protocol(X, lm) =$$
$$\{A \mid A \in X \land$$
$$\exists t_1, t_2. \forall a \in A.actions.$$
$$(lm(a) = \textbf{enq}) \Rightarrow (a = (t_1, \_)) \land$$
$$(lm(a) = \textbf{deq}) \Rightarrow (a = (t_2, \_)) \}$$

Applying this restriction to the most general client gives us a *protocol restricted* most general client:

$$\mathsf{MGC}_{pr}(L) = (\mathsf{let}\ L\ \mathsf{in}\ C^{\textbf{enq}} \parallel C^{\textbf{deq}})$$

where

$$C^{\textbf{enq}} = (\textbf{enq})^* \qquad \text{and} \qquad C^{\textbf{deq}} = (\textbf{deq})^*.$$

We use the stronger form of the abstraction theorem that applies only in the absence of relaxed atomics 6, replacing the MGC with the protocol-restricted most-general-client MGC_{pr}.

THEOREM 36 (Abstraction). *If: $(L_1, \mathcal{I}_1)$, $(L_2, \mathcal{I}_2)$ are safe under MGC_{pr}, and $(L_1, \mathcal{I}_1) \preceq (L_2, \mathcal{I}_2)$ under MGC_{pr}, and C satisfies the protocol pr, and $(C(L_2), \mathcal{I} \uplus \mathcal{I}_2)$, then $(C(L_1), \mathcal{I} \uplus \mathcal{I}_1)$ is safe and*

$$\forall X \in \mathsf{client}(\llbracket C(L_1) \rrbracket (\mathcal{I} \uplus \mathcal{I}_1)).$$
$$\exists Y \in \mathsf{client}(\llbracket C(L_2) \rrbracket (\mathcal{I} \uplus \mathcal{I}_2)). X \sqsubseteq Y.$$

E.2.3 Axiomatisation of action structures

Specification. The following action structures correspond to the actions generated by calls to the specification functions. In the specification, we have four sorts of calls to the library: enqueues, both succeeding and failing, and dequeues, both succeeding and failing. Each of the four sorts of calls generates a different execution fragment. A call to init generates the following:

$$X_{\textbf{init},t,d}^{a,d} =$$
$$(\{a, b, c, d\}), \{a \xrightarrow{\text{sb}} b \xrightarrow{\text{sb}} c \xrightarrow{\text{sb}} d\})$$

where

$$a = (\_, \_, t, \mathsf{call}\ \textbf{init}(d))$$
$$b = (\_, \_, t, \mathsf{store}_{NA}(l, \mathtt{empty\_list}()))$$
$$c = (\_, \_, t, \mathsf{store}_{NA}(c, 0))$$
$$d = (\_, \_, t, \mathsf{ret}\ \textbf{init}(1))$$

A successful enq generates the following:

$$X_{enq,t,k,l_t}^{a,f} =$$
$$(\{a, \dots f\}), \{a \xrightarrow{\text{sb}} b \xrightarrow{\text{sb}} c \xrightarrow{\text{sb}} d \xrightarrow{\text{sb}} f\})$$

where

$$a = (\_, \_, t, \mathsf{call}\ \textbf{enq}(k))$$
$$b = (\_, \_, t, \mathsf{store}_{\mathsf{NA}}(b, \mathtt{nondet}))$$
$$c = (\_, \_, t, \mathsf{load}_{\mathsf{NA}}(b, 0))$$
$$d = (\_, \_, t, \mathsf{load}_{ACQ}(l, l_t))$$
$$e = (\_, \_, t, \mathsf{store}_{REL}(l, \mathtt{append}(l_t, k)))$$
$$f = (\_, \_, t, \mathsf{ret}\ \textbf{enq}(1))$$

A failing enq generates the following:

$$X^{a,d}_{enq,t,k} =$$
$$(\{a,b,c,d)\}, \{a \xrightarrow{sb} b \xrightarrow{sb} c \xrightarrow{sb} d\})$$

where

$a = (\_, \_, t, \mathsf{call}\ \mathbf{enq}(k))$
$b = (\_, \_, t, \mathsf{store_{NA}}(b, \mathtt{nondet}))$
$c = (\_, \_, t, \mathsf{load_{NA}}(b, i))$ such that $i \neq 0$
$d = (\_, \_, t, \mathsf{ret}\ \mathbf{enq}(0))$

A successful deq generates the following:

$$X^{a,f}_{deq,t,d,l_t,v} =$$
$$(\{a,b,c,d,e,f)\}, \{a \xrightarrow{sb} b \xrightarrow{sb} c \xrightarrow{sb} d \xrightarrow{sb} e \xrightarrow{sb} f\})$$

where

$a = (\_, \_, t, \mathsf{call}\ \mathbf{deq}(d))$
$b = (\_, \_, t, \mathsf{load}_{ACQ}(l, l_t))$
$c = (\_, \_, t, \mathsf{load}_{NA}(c, v))$ such that $v < \mathtt{length}(l_t)$
$d = (\_, \_, t, \mathsf{store}_{NA}(c, v+1))$
$e = (\_, \_, t, \mathsf{ret}\ \mathbf{deq}(\mathtt{index}(l_t, c)))$

A failing deq generates the following:

$$X^{a,e}_{deq,t,d,l_t,v} =$$
$$(\{a,b,c,d,e)\}, \{a \xrightarrow{sb} b \xrightarrow{sb} c \xrightarrow{sb} d \xrightarrow{sb} e\})$$

where

$a = (\_, \_, t, \mathsf{call}\ \mathbf{deq}(d))$
$b = (\_, \_, t, \mathsf{load}_{ACQ}(l, l_t))$
$c = (\_, \_, t, \mathsf{load}_{NA}(c, v))$ such that $\mathtt{length}(l_t) \leq v$
$e = (\_, \_, t, \mathsf{ret}\ \mathbf{deq}(0))$

Implementation. The following action structures correspond to the actions generated by calls to the implementation functions, starting with **init**:

$$X^{a,e}_{\mathbf{init},t,n} =$$
$$(\{a,b,c,d,e,f,g)\}, \{a \xrightarrow{sb} b \xrightarrow{sb} c \xrightarrow{sb} d \xrightarrow{sb} e\})$$

where

$a = (\_, \_, t, \mathsf{call}\ \mathbf{init}(n))$
$b = (\_, \_, t, \mathsf{store}_{NA}(s, n))$
$c = (\_, \_, t, \mathsf{store}_{NA}(b, 0))$
$d = (\_, \_, t, \mathsf{store}_{NA}(f, 0))$
$e = (\_, \_, t, \mathsf{ret}\ \mathbf{init}(1))$

A successful **enq** gives rise to:

$$X^{a,g}_{\mathbf{enq},t,k,f_t,b_t,k} =$$
$$(\{a,b,c,d,e,f,g)\}, \{a \xrightarrow{sb} b \xrightarrow{sb} c \xrightarrow{sb} d \xrightarrow{sb} e \xrightarrow{sb} f \xrightarrow{sb} g\})$$

where

$a = (\_, \_, t, \mathsf{call}\ \mathbf{enq}(k))$
$b = (\_, \_, t, \mathsf{load}_{ACQ}(f, f_t))$
$c = (\_, \_, t, \mathsf{load}_{ACQ}(b, b_t))$ such that $b_t - s < f_t$
$e = (\_, \_, t, \mathsf{store}_{NA}(a + (b_t \% s), k))$
$f = (\_, \_, t, \mathsf{store}_{REL}(b, b_t + 1))$
$g = (\_, \_, t, \mathsf{ret}\ \mathbf{enq}(1))$

or the **enq** fails and we have:

$$X^{a,e}_{\mathbf{enq},t,k,f_t,b_t,k} =$$
$$(\{a,b,c,d,e)\}, \{a \xrightarrow{sb} b \xrightarrow{sb} c \xrightarrow{sb} d \xrightarrow{sb} e\})$$

where

$a = (\_, \_, t, \mathsf{call}\ \mathbf{enq}(k))$
$b = (\_, \_, t, \mathsf{load}_{ACQ}(f, f_t))$
$c = (\_, \_, t, \mathsf{load}_{ACQ}(b, b_t))$ such that $f_t \leq b_t - s$
$e = (\_, \_, t, \mathsf{ret}\ \mathbf{enq}(0))$

Similarly, either **deq** is successful, and we have:

$$X^{a,h}_{\mathbf{deq},t,d,f_t,b_t,k} =$$
$$(\{a,b,c,d,e,f,g,h)\}, \{a \xrightarrow{sb} b \xrightarrow{sb} c \xrightarrow{sb} d \xrightarrow{sb} e \xrightarrow{sb} f \xrightarrow{sb} g \xrightarrow{sb} h\})$$

where

$a = (\_, \_, t, \mathsf{call}\ \mathbf{deq}(d))$
$b = (\_, \_, t, \mathsf{load}_{ACQ}(f, f_t))$
$c = (\_, \_, t, \mathsf{load}_{ACQ}(b, b_t))$ such that $f_t < b_t$
$e = (\_, \_, t, \mathsf{load}_{NA}(a + (f_t \% s), v))$
$f = (\_, \_, t, \mathsf{store}_{NA}(k, v))$
$g = (\_, \_, t, \mathsf{store}_{REL}(f, f_t + 1))$
$h = (\_, \_, t, \mathsf{ret}\ \mathbf{deq}(k))$

or the **deq** fails and we have:

$$X^{a,e}_{\mathbf{deq},t,d,f_t,b_t,k} =$$
$$(\{a,b,c,d,e)\}, \{a \xrightarrow{sb} b \xrightarrow{sb} c \xrightarrow{sb} d \xrightarrow{sb} e\})$$

where

$a = (\_, \_, t, \mathsf{call}\ \mathbf{deq}(d))$
$b = (\_, \_, t, \mathsf{load}_{ACQ}(f, f_t))$
$c = (\_, \_, t, \mathsf{load}_{ACQ}(b, b_t))$ such that $b_t \leq f_t$
$e = (\_, \_, t, \mathsf{ret}\ \mathbf{deq}(0))$

E.2.4 Implementation and Specification Safety

Let $(L_2, \mathcal{I}_2)$ be the producer-consumer specification, and $(L_1, \mathcal{I}_1)$ be the implementation. To apply the abstraction theorem, we must prove the following:

LEMMA 37 (Specification Safety). $(L_1, \mathcal{I}_1)$ and $(L_2, \mathcal{I}_2)$ are safe over the MGC_{pr}.

Proof: Assume one of them is not. Then there exists an execution, $X_2 \in \mathsf{MGC}_{pr}(L_2, \mathcal{I}_2)$ or $X_1 \in \mathsf{MGC}_{pr}(L_1, \mathcal{I}_1)$ such that X_2 or X_1 has a safety violation. Consider each type of safety violation in turn:

Location disjointness   Recall NONINTERF:

$$\forall a, x, t. (a \in A \wedge (a = (t, \mathsf{write}(x, \_)) \vee a = (t, \mathsf{read}(x, \_))) \wedge$$
$$x \notin \{\mathsf{param}_t, \mathsf{retval}_t \mid t = 1..n\}) \implies$$
$$((\exists b. b = (\_, \mathsf{call}\ \_) \wedge b \xrightarrow{sb} a \wedge$$
$$\neg \exists c. c = (\_, \mathsf{ret}\ \_) \wedge b \xrightarrow{sb} c \xrightarrow{sb} a) \iff$$
$$(x \in \mathsf{LLoc})).$$

First note that there are no client actions in the MGC_{pr}, so the read or write, a, must originate from a library call. Consequently, there is a sequenced before preceding call action without an intervening return action, so we must show that $x \in \mathsf{LLoc}$ for a. Observe that in L_1 and L_2, all accesses are at static global locations, except accesses to statically allocated elements of the task array in the implementation, L_1. In the axiomatisation, these accesses, action e in **enq** and e in **deq** use a dynamically calculated address. Both accesses of the array have indices that are calculated modulo

the size of the array, so no accesses can escape its bounds, preserving disjointness.

Mutex safety Recall SAFELOCK:

$$(\forall b, t, \ell.\, b = (t, \mathsf{unlock}(\ell)) \implies$$
$$\quad \exists a.\, a = (t, \mathsf{lock}(\ell)) \wedge a \xrightarrow{\mathsf{sb}} b \wedge a \xrightarrow{\mathsf{sc}} b \wedge$$
$$\quad \neg \exists c.\, c = (\_, \_(\ell)) \wedge a \xrightarrow{\mathsf{sc}} c \xrightarrow{\mathsf{sc}} b) \wedge$$
$$(\neg \exists a, b, t, \ell.\, a = (t, \mathsf{lock}(\ell)) \wedge b = (t, \mathsf{block}(\ell)) \wedge a \xrightarrow{\mathsf{sb}} b \wedge a \xrightarrow{\mathsf{sc}} b \wedge$$
$$\quad \neg \exists c.\, c = (\_, \_(\ell)) \wedge a \xrightarrow{\mathsf{sc}} c \xrightarrow{\mathsf{sc}} b).$$

There are no mutex actions in the specification or implementation, so both conjuncts of SAFELOCK are true.

Indeterminate read Recall SAFEREAD:

$$\forall a, x.\, (a \in A \wedge a = (\_, \mathsf{read}(x, \_)) \implies \exists b.\, b \xrightarrow{\mathsf{rf}} a \wedge b = (\_, \mathsf{write}(x, \_)).$$

The client performs initialisation of all global variables, and that happens-before all library calls. That combined with DETREAD, ensures there is always a reads-from edge as required.

Data races Recall DRF:

$$\forall a, b, x, \mathsf{a}, \mathsf{b}, t_1, t_2.$$
$$\quad (a, b \in A \wedge a \neq b \wedge a = (t_1, \mathsf{a}(x, \_)) \wedge$$
$$\quad b = (t_2, \mathsf{b}(x, \_)) \wedge (\mathsf{a} = \mathsf{write} \vee \mathsf{b} = \mathsf{write})) \implies$$
$$\quad (t_1 \neq t_2 \implies (a \xrightarrow{\mathsf{hb}} b \vee b \xrightarrow{\mathsf{hb}} a \vee \mathsf{sort}(x) = \mathsf{ATOM})) \wedge$$
$$\quad (t_1 = t_2 \implies (a \xrightarrow{\mathsf{sb}} b \vee b \xrightarrow{\mathsf{sb}} a)).$$

Suppose there is a violation of the second conjunct in the implication, then there must be a write and another access of the same location, on the same thread, unrelated by sequenced-before. The race cannot be within one call in either the implementation or the specification, because inspection of the axiomatisation shows all actions are totally ordered by sequenced before. There can be no inter-call unsequenced actions on the same thread because the MGC_{pr} totally orders calls by sequenced before. Consequently, the second conjunct of the implication is not violated.

Now suppose there is a violation of the first conjunct of the implication. Then there must be a write and another access of the same non-atomic location, on different threads, unrelated by happens-before. The specification accesses one non-local non-atomic location, c. It is only accessed by the dequeueing thread, so all accesses are related by sequenced-before, and as a result, happens-before.

The only non-local non-atomic locations accessed by the implementation are reads of s and accesses of the array, a. The only write of s is in the initialisation, and that happens-before all of the reads.

Suppose there were a race between accesses to the array, there must be an enqueue on one thread and a dequeue on the other such that the value of b read by **enq** modulo s is the same as the value of f read by **deq** modulo s1.

Let $W_{a+(b_{enq}\%s)}$ be the write of the array and $R_{a+(f_{deq}\%s)}$ be the read of the array that might race, let f_{enq} and b_{enq} be the values of f and b read by **enq**, and f_{deq} and b_{deq} be the values read by **deq**. Noting that the conditional was passed in both calls, brings our knowledge to the following:

1. $b_{enq}\%s = f_{deq}\%s$
2. $b_{enq} - s < f_{enq}$
3. $f_{deq} < b_{deq}$

LEMMA 38. $b_{enq} \leq f_{deq} \Rightarrow W_{a+(b_{enq}\%s)} \xrightarrow{\mathsf{hb}} R_{a+(f_{deq}\%s)}$

We have both: $b_{enq} \leq f_{deq}$ and $f_{deq} < b_{deq}$, so $b_{enq} < b_{deq}$. The value of b is modified only by the increment in each call of **enq**, therefore, the read of b in **deq** reads from and synchronises with the write of b in the **enq** call we are considering or some later call

in sequenced-before. Either way, through transitivity of happens-before, we have $W_{a+(b_{enq}\%s)} \xrightarrow{\mathsf{hb}} R_{a+(f_{deq}\%s)}$. □

LEMMA 39. $f_{deq} < b_{enq} \Rightarrow R_{a+(f_{deq}\%s)} \xrightarrow{\mathsf{hb}} W_{a+(b_{enq}\%s)}$

We have both: $f_{deq} < b_{enq}$ and $b_{enq}\%s = f_{deq}\%s$, so $f_{deq} \leq b_{enq} - s$. Furthermore, we have $b_{enq} - s < f_{enq}$, so $f_{deq} < f_{enq}$.

The value of f is modified only by the increment in each call of **deq**, therefore, the read of f in **enq** reads from and synchronises with the write of f in the **deq** call we are considering or some later call in sequenced-before. Either way, through transitivity of happens-before, we have $R_{a+(f_{deq}\%s)} \xrightarrow{\mathsf{hb}} W_{a+(b_{enq}\%s)}$. □

The two lemmas cover all possible read values, so we always have a happens-before edge between potentially racy accesses of the array, and there are no data races.

E.2.5 Proving Abstraction

Now we show that $(L_1, \mathcal{I}_1) \preceq (L_2, \mathcal{I}_2)$.

Constructing a Specification Execution

DEFINITION 40. *Let $X_{\mathrm{imp}} \in [\![L_1]\!]\mathcal{I}_1$ be an arbitrary but fixed execution of the producer-consumer implementation. We construct a corresponding execution X_{spec} of the specification as follows:*

We first construct an action structure C_{spec}. Let

$$X_{\mathrm{imp}} = (A_i, \mathsf{sb}_i, \mathsf{ib}_i, \mathsf{rf}_i, \mathsf{sc}_i, \mathsf{mo}_i, \mathsf{sw}_i, \mathsf{hb}_i)$$

and let $C_{\mathrm{imp}} = (A_i, \mathsf{sb}_i, \mathsf{ib}_i)$ be the corresponding action structure. We construct C_{spec} by substituting successful implementation enq operations with successful specification enq operations, failing implementation enqs with failing specification enqs, and so on.

We first define a total order $<$ over calls. We choose this total order such that for any pair of calls $I_1, I_2 \in X_{\mathrm{client}}$, $I_1 < I_2$ if I_2 reads from a store in I_1 to either f or b. By construction, the first element in the order must be the single initial call to init.

This order $<$ is acyclic because reads and writes between threads are are release-acquire, meaning that $<$ must be acyclic by the acyclicity of hb.

We use the order $<$ to construct a function j on actions:

1. if $X^{a,b}$ is an init call, $j(a) = []$;

2. if $X^{a,b}$ is a successful enq for value v, and $X^{a',b'}$ is the immediately preceding operation in $<$, then $j(a) = \mathsf{append}(v, j(a'))$.

3. otherwise, for all other operations $X^{a,b}$ with immediately preceding operation $X^{a',b'}$, $j(a) = j(a')$.

The function j maps from call actions in the implementation, to the contents of the l variable in the specification. Intuitively, it builds the abstract state of the queue.

We now construct C_{spec}. We begin with C_{imp}. For every successful enq operation in the action structure, we substitute the implementation sub-execution for a specification sub-execution:

$$C' = C[S^{x,y}_{\mathbf{enqS},t,k,b,j(x)} / I^{x,y}_{\mathbf{enqS},t,k,f,b}]$$

Note that here the back pointer b from the implementation gives the value of c in the specification, while the function j is gives a value to l in the specification.

We do the same to failing enq operations, successful and failing deq, and init. As these substitutions commute, we can do them in any order. The result is an action structure C_{spec}.

LEMMA 41. *C_{spec} is an action structure for the producer-consumer specification. That is, $C_{\mathrm{spec}} \in \langle L_2 \rangle \mathcal{I}_2$.*

PROOF. We have performed straightforward surgery on the action struture C_{imp}, replacing enq with enq, etc. It is trivial to see

that the resulting execution is a action structure for the protocol-restricted MGC. Values have been chosen deliberately to ensure consistency between the resulting executions, but are irrelevant to the correctness of the action structure. □

We now construct the execution

$$X_{\text{spec}} = (A_s, \text{sb}_s, \text{ib}_s, \text{rf}_s, \text{sc}_s, \text{mo}_s, \text{sw}_s, \text{hb}_s).$$

We take the action set, happens-before and initialised-before relations from $C_{\text{spec}} = (A_s, \text{sb}_s, \text{ib}_s)$. Other relations are constructed as follows:

- All reads and writes for location c are on the same thread, or in the initialisation, and so rf for c simply follows sb and ib.

  For location l, reads may be on either thread, but writes are all on the enqueuing thread. For reads on the same thread, rf is follows sb. For reads between threads, we use the $<$ order as an oracle. For any $I^{a,b}_{\textbf{deqS},t,k}$, we find its immediate $<$-predecessor $I^{a',b'}_{\textbf{enqS},t',k'}$, and generate an rf from the write to l in the corresponding $S^{a',b'}_{\textbf{enqS},t',k'}$ to the read in $S^{a,b}_{\textbf{deqS},t,k}$. If $I^{a,b}_{\textbf{deqS},t,k}$ has no $<$-preceding enqueue, we direct the read for l to init.

- mo is determined by sb, because the locations c and l are each modified on a single thread only.

- sc is empty, as there are no SC actions.

- sw is determined by rf and the SWDEF axiom.

- hb is determined by the HBDEF axiom.

The following lemma relates the values stored in X_{imp} to values in X_{spec}.

LEMMA 42. *Let $I^{x,y}$ be an* enq *or* deq *call in X_{imp} and let $S^{x,y}$ be the corresponding call in X_{spec}. Identify in X_{imp} the action for the read from b, and in X_{spec} the action for the read from l. These actions are related as follows:*

- length$(lt) = bt$.
- *for $i \in [bt..(bt - ft)]$, the value of the* hb-*preceding write at location $a[(i-1) \bmod s]$ is same as $lt[i]$.*

PROOF. By induction on the order $<$. Pick a call $I^{x,y}$ in X_{imp} and let $S^{x,y}$ be the corresponding call in X_{spec}. By construction, the current call must read its value of l from the $<$-preceding call. Case-split on the type of the $<$-preceding call:

enqueue. To see that length$(\text{lt}) = b$, observe that b is incremented in this call, and one element is added to the list, incrementing its length, so this part of the invariant is preserved.

The successful **enq** writes the parameter value to element $b - 1$ in the implementation, and the specification appends it to the queue as required. The second part of the invariant is preserved.

dequeue. In the **deq** case, the implementation increments f which reduces the number of index values that must be considered. Neither b nor the list in the specification are changed, so the invariant is preserved.

init. Trivial.

□

LEMMA 43. X_{spec} *is valid.*

PROOF. We consider each condition in turn:

SBWF   sb is only modified by substituting linear sequences of actions for other linear sequences. T

RFWF   For client actions, the actions and relations are the same, and the predicate remains true if reads-from edges are removed. For the library, we have chosen the execution to be consistent.

It remains to show that the values in the specification calls are correctly read from associated writes. By the structure of the original execution X_{imp} we know that b is incremented in each successful call to deq. In X_{spec}, c takes it's value from the increment in the hb-preceding call to deq. By simple induction, all reads of c have values that satisfy the axiom.

For intra-call non-atomic reads, we examine the given values to ensure that they are consistent. For the call value passed to enq, we know by construction that the appropriate value will be added to the list. For the correctness of return value for deq, we appeal to Lemma 42 to ensure that the value in l read by deq corresponds to the value read from the array in the concrete implementation. Once again by Lemma 42, we know that deq will fail only if the corresponding implementation call could fail.

SCWF, SCREADS, HBVSSC **and** MOVSSC   There are no SC actions in the implementation or specification.

MOWF   By construction, all writes to the same location are mo-related. In X_{spec}, mois determined sbso the rest of the axiom follows from SBWF.

ATOMRMW   There are no read-modify-write actions in the implementation or specification.

ACYCLICITY   By assumption, hb_i in the execution X_{imp} does not contain cycles. The constructed execution X_{spec} removes all hb-ordering from the dequeueing to enqueueing thread, and adds only hb-ordering from the enqueueing to dequeuing thread. By construction, this cannot break irreflexivity.

LOCKS   There are no locks in the implementation or specification.

RFATOMIC   All atomic reads of atomic writes automatically satisfy the predicate because in reading, they synchronise. For reads of the stores in init, note that the stores in init precede all other calls in happens-before, according to the protocol.

HBVSMO   mo is determined by sb, and hb is similarly determined by sb and HBDEF.

CORR, COWR **and** CORW   The library actions on the **enq** and **deq** threads are totally ordered by sequenced before, precluding intra thread library coherence violations. The only inter-thread edge is the read-from edge at l, from calls of **enq** to calls of **deq**. This reads-from edge synchronises, creating a happens-before edge, making coherence violations impossible.

RFNONATOMIC **and** DETREAD   For DETREAD, l is the only atomic location. All reads in the library are hb-preceded by writes in the initialisation, which satisfies one direction of the iff. For the other direction, we know by construction that every read of l is rf-related to an appropriate write.

For RFNONATOMIC, the only inter-call non-atomic reads in the library are those of c in calls to **deq**. These are totally ordered by sequenced before, and are hb-related only to actions on the same thread.

BLOCK   There are no blocking actions in the library.

□

LEMMA 44. $X_{\text{spec}} \in \llbracket L_2 \rrbracket \mathcal{I}_2$.

PROOF. Appeal to Lemmas 41 and 43. □

Inclusion over histories

LEMMA 45. $\text{hbL}(X_{\text{spec}}) \subseteq \text{hbL}(X_{\text{imp}})$.

PROOF. All intra-thread happens-before edges over the call and return actions are the same. The only inter-thread happens-before edges created by the specification are a result of **deq** calls reading l from writes of l in **enq** calls. By construction, these edges correspond to reads of b which synchronise. Therefore all happens-before edges in X_{spec} are covered by X_{imp}, as required. □

LEMMA 46. $\mathsf{denyL}(X_{\mathrm{spec}}) \subseteq \mathsf{denyL}(X_{\mathrm{imp}})$.

PROOF. Consider each deny part of the history in turn, and show that the edges created by the specification are a subset of the edges created by the implementation:

HBVSMO. The only modification order edges in the specification are over the writes to l. All calls that write l are totally ordered by sequenced-before in both the specification and the implementation, so this part of the deny is covered.

COWR. By the structure of the specification, any read from a write must hb-preceed that writes mo-successor. Thus the specification generates no deny.

SCREADS. There are no sc actions in the specification, so this component of the deny is empty.

□

LEMMA 47.

$$\forall H_1 \in \mathsf{Ehistory}(X_{\mathrm{imp}}). \, \exists H_2 \in \mathsf{Ehistory}(X_{\mathrm{spec}}). \, H_2 \subseteq H_1.$$

Proof. By appeal to the previous two lemmas. The sc component of the history has no effect, as both the implementation and specification include no sc orders. □