

Relaxed Separation Logic: A Program Logic for C11 Concurrency

Viktor Vafeiadis

Max Planck Institute for Software Systems
(MPI-SWS)
viktor@mpi-sws.org

Chinmay Narayan

Indian Institute of Technology, Delhi
chinmay@cse.iitd.ac.in

Abstract

We introduce *relaxed separation logic* (RSL), the first program logic for reasoning about concurrent programs running under the C11 relaxed memory model. From a user’s perspective, RSL is an extension of concurrent separation logic (CSL) with proof rules for the various kinds of C11 atomic accesses. As in CSL, individual threads are allowed to access non-atomically only the memory that they own, thus preventing data races. Ownership can, however, be transferred via certain atomic accesses. For SC-atomic accesses, we permit arbitrary ownership transfer; for acquire/release atomic accesses, we allow ownership transfer only in one direction; whereas for relaxed atomic accesses, we rule out ownership transfer completely. We illustrate RSL with a few simple examples and prove its soundness directly over the axiomatic C11 weak memory model.

Categories and Subject Descriptors D.3.1 [Programming Languages]: Formal Definitions and Theory; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs

Keywords Concurrency; Weak memory models; C/C++; Proof system; Separation logic

1. Introduction

Wanting to enable many hardware and software optimizations, modern programming language definitions provide rather weak guarantees on the semantics of concurrent memory accesses allowing, for example, different threads to observe shared operations happening in different orders. One such case is the concurrency model adopted by the 2011 revisions of the C and C++ standards (ISO/IEC 9899:2011;

ISO/IEC 14882:2011), which we will study in this paper and refer to as the C11 model.

C11 provides several kinds of memory accesses—non-atomic, relaxed atomic, acquire atomic, release atomic, and sequentially consistent (SC) atomic—each providing different consistency guarantees. On the one end of the spectrum, races on non-atomic accesses result in completely undefined behaviour (they are treated as programming errors); on the other end, SC-atomic accesses are globally synchronized. The guarantees provided by relaxed, acquire, and release accesses lie somewhere in between: different threads can observe them happening in different orders.

The reason for having all these kinds of accesses is that they map differently to the various common architectures, and have very different implementation costs. Non-atomic and relaxed atomic accesses are generally rather cheap as they correspond to vanilla machine loads and stores, and may be reordered by the compiler and/or by an out-of-order execution unit. At the other end of the spectrum, SC accesses are very expensive because their implementation involves a full memory barrier. The cost of acquire and release accesses depends a lot on the architecture. On x86, they are compiled down to plain reads and writes (Batty et al. 2011) and are therefore cheap. On PowerPC and ARM, the cost is somewhat higher as they induce a memory barrier, but of a weaker kind than full memory barriers (Sarkar et al. 2012).

Our goal is to help C11 programmers by providing them with sound reasoning principles for concurrent programs. We show that C11 concurrency supports resource reasoning in the style of separation logic (O’Hearn 2007); in particular, ownership can be transferred along acquire/release atomic memory accesses and does not require SC-accesses.

We develop *relaxed separation logic* (RSL), a program logic that follows the resourceful reading of separation logic triples. When we assert the Hoare triple $\{P\} Cmd \{Q\}$, we say that the command Cmd will not access any memory other than that given by its precondition, P , or subsequently acquired during its execution. We thus support the parallel composition rule of separation logic,

$$\frac{\{P_1\} Cmd_1 \{Q_1\} \quad \{P_2\} Cmd_2 \{Q_2\}}{\{P_1 * P_2\} Cmd_1 || Cmd_2 \{Q_1 * Q_2\}} \quad (\text{PAR})$$

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

OOPSLA ’13, October 29–31, 2013, Indianapolis, Indiana, USA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-2374-1/13/10...\$15.00.
<http://dx.doi.org/10.1145/2509136.2509532>

which ensures that the two threads do not have any races on non-atomic memory accesses, a condition required by C11.

To handle acquire/release atomics, we introduce two new assertion kinds, $\text{Rel}(\ell, Q_1)$ and $\text{Acq}(\ell, Q_2)$. These denote respectively the permissions to perform a release-write of some value v at location ℓ and give away ownership of the resource described by $Q_1(v)$, or an acquire-read and gain ownership of $Q_2(v)$. With these assertion forms we provide simple proof rules for release writes and acquire reads, similar to those for releasing and acquiring mutual exclusion locks in concurrent separation logic.

Besides RSL itself, the main contribution of this work was to define the meaning of Hoare triples in a relaxed memory model setting, so as to prove the soundness of RSL. This was rather challenging for three main reasons.

No global state/time: Traditionally, $\{P\} \text{Cmd} \{Q\}$ asserts that if we execute Cmd in an initial state satisfying P and it terminates, then the final state will satisfy Q . In C11 concurrency, however, the terms “initial state” and “final state” are ill-defined, because there exist no global notions of time or state.

To interpret triples, we thus resort to *logical* local notions of time and state. We define a logical notion of local state at each event of a program execution, and thread the logical state through C11 “happens-before” edges.

Assertions in heaps: Our assertions for dealing with acquire and release atomics require that the logical heaps used to interpret them contain assertions. This results in a circularity in the model of assertions, which for simplicity we resolve by storing syntactic assertions.

No operational semantics: Concurrent program logics are typically proved sound over an operational or a trace semantics. In either case, the meaning of Hoare triples can be defined in terms of an auxiliary predicate by induction over the length of an execution trace. These definitions cannot directly be extended to the C11 model as there is no obvious total order for the induction. Our solution is to order the C11 events according to the total number of events that happen before them.

As a secondary contribution, we observed that the semantics of relaxed atomic memory accesses in C11 is too weak to permit even the most basic reasoning principles about them, which in turn renders basic compiler optimizations unsound. In order to allow such reasoning principles, we proposed a crude fix to C11, which we discuss in Section 6.

In the remainder of this paper, we define a minimal concurrent programming language (§2), review the C11 concurrency model (§3), describe the assertions and proofs rules of RSL (§4), verify a few examples using RSL (§5), explain the problems caused by relaxed accesses and their resolution (§6), present the semantics of assertions and Hoare triples and sketch the main parts of the soundness proof (§7). We conclude with a discussion of related and future work (§8).

A Coq formalization of the soundness proof of RSL can be found at the following URL.

<http://www.mpi-sws.org/~viktor/rsl/>

2. Programming Language

In order to focus on the concurrency aspects of C11 and to avoid the inherent complexity of a large language like C, we introduce a minimal concurrent programming language featuring the various kinds of memory accesses supported by C11. Following Batty et al. (2012), we omit consume reads from the model, because they are relevant only for a few architectures (PowerPC and ARM) and substantially complicate the model. For simplicity, we also omit memory fences.

To make the local sequential execution order explicit, we present the grammar of expressions in A-normal form (cf. Flanagan et al. 1993). Atomic expressions, $e \in \text{AExp}$, consist of variables and values (locations and numbers). Program expressions, $E \in \text{Exp}$, consist of atomic expressions, let-bound computations, conditionals, loops, parallel composition, memory allocation, loads, stores, and atomic compare-and-swap (**CAS**) instructions.

$$\begin{aligned} v \in \text{Val} &::= \ell \mid n && \text{where } \ell \in \text{Loc}, n \in \mathbb{N} \\ e \in \text{AExp} &::= x \mid v && \text{where } x \in \text{Var} \\ E \in \text{Exp} &::= e \mid \text{let } x = E \text{ in } E' \mid \text{if } e \text{ then } E \text{ else } E' \\ & \mid \text{repeat } E \text{ end} \mid E_1 \parallel E_2 \mid \text{alloc}() \\ & \mid [e]_X \mid [e]_Y := e' \mid \text{CAS}_{Z,W}(e, e', e'') \end{aligned}$$

where $X \in \{\text{sc}, \text{acq}, \text{rlx}, \text{na}\}$, $Y \in \{\text{sc}, \text{rel}, \text{rlx}, \text{na}\}$,
 $Z \in \{\text{sc}, \text{rel_acq}, \text{acq}, \text{rel}, \text{rlx}\}$, $W \in \{\text{sc}, \text{acq}, \text{rlx}\}$

As in C, in conditional expressions we treat zero as false and non-zero values as true. The construct **repeat** E **end** executes E repeatedly until it returns a non-zero value.

Memory accesses are annotated by their mode: sequentially consistent (sc), acquire (acq), release (rel), combined release-acquire (rel_acq), relaxed (rlx), or non-atomic (na). According to the C standard, not all modes are available for all accesses: reads cannot be releases, writes cannot be acquires, **CAS**s cannot be non-atomic. These restrictions are to avoid redundancy in the language. For example, an acquire write, if such a thing were allowed, would behave exactly the same as a relaxed write.

CAS is an atomic operation used to heavily in lock-free concurrent algorithms. It takes a location, ℓ , and two values, v' and v'' , as arguments. It atomically checks if the value at location is v' or not. If the value is same as v' , then **CAS** succeeds: it atomically writes v'' to ℓ and returns the old value. If the value is different, **CAS** fails: it returns that value and does not modify the location. **CAS** is annotated with two access modes: one to be used for the successful case, and one for the unsuccessful case.

For conciseness in examples, we will often write expressions such as $[E]_{\text{na}}$ instead of **let** $x = E$ **in** $[x]_{\text{na}}$. We also write $E_1; E_2$ instead of **let** $x = E_1$ **in** E_2 when $x \notin \text{fv}(E_2)$.

```

new_lock() = let  $x = \mathbf{alloc}()$  in  $[x]_{\text{rel}} := 1; x$ 
spin( $x$ ) = repeat  $[x]_{\text{rlx}}$  end
lock( $x$ ) = repeat spin( $x$ ); CASacq,rlx( $x, 1, 0$ ) end
unlock( $x$ ) =  $[x]_{\text{rel}} := 1$ 

```

Figure 1. Simple spinlock implementation.

Spinlock Example There are two important uses of acquire/release accesses: in implementing locks, and in message passing. Relaxed accesses are useful in cases of optimistic reads, where the value read, if it is of interest to the algorithm, will be read again by an acquire read, or an acquire fence will be issued. For example, in the simple CAS-based spinlock implementation shown in Figure 1, $lock(x)$ performs an acquire-on-success CAS and $unlock(x)$ does a release write. The optimistic $spin(x)$ loop that waits for the lock to become free, in contrast, does relaxed reads. The combined release-acquire CAS is supposed to be used for operations that atomically release one lock and acquire another—this is possible, for example, if the locks are represented as different bits of the same word. Further examples can be found in McKenney and Garst (2011).

3. The C11 Memory Model

The C11 memory model is defined axiomatically in terms of program executions. A program *execution* consists of a set of actions and several binary relations over them. Actions describe the memory operations performed by the program, and are labelled with information about the memory order of the operation, the address accessed and the values read and/or written.

$$\text{Act} ::= \text{skip} \mid \text{W}_{(\text{sc}|\text{rel}|\text{rlx}|\text{na})}(\ell, v) \mid \text{R}_{(\text{sc}|\text{acq}|\text{rlx}|\text{na})}(\ell, v) \\ \mid \text{RMW}_{(\text{sc}|\text{rel}|\text{acq}|\text{acq}|\text{rel}|\text{rlx})}(\ell, v, v') \mid \text{A}(\ell)$$

In summary, we have a no-op action; SC, release, relaxed and non-atomic writes; SC, acquire, relaxed, and non-atomic reads; atomic read-modify-write actions; and allocations. The no-op (skip) action represents local computations, thread forks and joins.

For the subset of C11 we consider, an execution contains the following relations:¹

- *Sequenced-before* (sb) relates actions of the same thread that follow one another in control flow. We have $\text{sb}(a, b)$ if a and b belong to the same thread and a immediately precedes b in the thread’s control flow, or a is a fork action and b the first action of the forked thread, or b is a join action and a the last action of the joined thread.
- The *reads-from* map (rf) maps every read action r to the write action w that wrote the value read by r .

¹The full model includes two additional relations, dd (data dependency) and dob (dependency ordered before), used to define the happens-before relation for consume reads.

- The *memory-order* relation (mo) is a total order on the store actions writing to the same atomic location.
- The *sequential-consistency* order (sc) is a total order over all SC-atomic actions.

Formally, let AName be a countably infinite of action names. Then, an execution, \mathcal{X} , is represented as a tuple, $\langle \mathcal{A}, \text{lab}, \text{sb}, \text{rf}, \text{mo}, \text{sc} \rangle$, where $\mathcal{A} \subseteq_{\text{fin}} \text{AName}$ is the set of action names included in the execution, $\text{lab} \in \mathcal{A} \rightarrow \text{Act}$ maps every action identifier to its label, $\text{rf} \in \mathcal{A} \rightarrow \mathcal{A}$ is the reads-from map, and $\text{sb}, \text{mo}, \text{sc} \in \mathcal{P}(\mathcal{A} \times \mathcal{A})$ are the sequenced-before relation, the memory order, and the sequential consistency order respectively.

From these relations, C11 defines a number of derived relations, the most important of which are: the *synchronizes-with* relation and the *happens-before* order.

- *Synchronizes-with* (sw) relates acquire reads with the release writes that precede in mo order the write whose value was read by the acquire read provided that all the writes between these two writes belong to the same thread or are RMW operations.
- *Happens-before* (hb) is a partial order on actions formalizing the intuition that one action was completed before the other. In the C11 subset we consider, $\text{hb} = (\text{sb} \cup \text{sw})^+$.

The semantics of a program is given by the set of *consistent* executions. An execution is said to be consistent if it satisfies the axioms of the memory model, which will be presented shortly. If, however, any of these consistent executions contains a data race on non-atomic actions, i.e. events generated from two conflicting operations on the same non-atomic location not ordered by hb in either direction, then the program is deemed to have arbitrary semantics. Thus, any sound program logic for C11 concurrency must ensure its specifications imply race-freedom for non-atomic actions.

Expression Semantics Let CExp denote closed expressions (i.e., ones with no free variables). The semantics of such closed expressions, $\llbracket E \rrbracket$, is given in Figure 2 as a set of tuples $\langle \text{res}, \mathcal{A}, \text{lab}, \text{sb}, \text{fst}, \text{lst} \rangle$. These tuples represent finite complete executions as well as finite incomplete execution prefixes (used to model infinite executions), where:

- (1) res is the result of evaluating the expression or \perp if the execution is incomplete;
- (2) \mathcal{A} is the set of all actions contained in the execution;
- (3) lab labels the actions with the corresponding operations;
- (4) sb represents the sequenced-before relation; and
- (5) fst and lst are the first and last actions in the sb-order.

For uniformity, we record the last action even in incomplete executions. In the parallel composition case, the auxiliary function $\text{combine}(\text{res}_1, \text{res}_2)$ returns res_1 if $\text{res}_2 \neq \perp$ and \perp otherwise. In the $\llbracket E \rrbracket$ semantics, allocations can return an arbitrary new location, and reads can read an arbitrary value. These will later be constrained by the consistency axioms.

$$\begin{aligned}
\llbracket - \rrbracket &: \text{CExp} \rightarrow \mathbb{P}(\langle \text{res} : \text{Val} \cup \{\perp\}, \mathcal{A} : \mathbb{P}(\text{AName}), \text{lab} : \mathcal{A} \rightarrow \text{Act}, \text{sb} : \mathbb{P}(\mathcal{A} \times \mathcal{A}), \text{fst} : \mathcal{A}, \text{lst} : \mathcal{A} \rangle) \\
\llbracket v \rrbracket &\stackrel{\text{def}}{=} \{\langle v, \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge \text{lab}(a) = \text{skip}\} \\
\llbracket \text{alloc}() \rrbracket &\stackrel{\text{def}}{=} \{\langle \ell, \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge \ell \in \text{Loc} \wedge \text{lab}(a) = \text{A}(\ell)\} \\
\llbracket [v]_Z := v' \rrbracket &\stackrel{\text{def}}{=} \{\langle v', \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge \text{lab}(a) = \text{W}_Z(v, v')\} \\
\llbracket [v]_Z \rrbracket &\stackrel{\text{def}}{=} \{\langle v', \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge v' \in \text{Val} \wedge \text{lab}(a) = \text{R}_Z(v, v')\} \\
\llbracket \text{CAS}_{X,Y}(v, v_o, v_n) \rrbracket &\stackrel{\text{def}}{=} \{\langle v', \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge v' \in \text{Val} \wedge v' \neq v_o \wedge \text{lab}(a) = \text{R}_Y(v, v')\} \\
&\cup \{\langle v_o, \{a\}, \text{lab}, \emptyset, a, a \rangle \mid a \in \text{AName} \wedge \text{lab}(a) = \text{RMW}_X(v, v_o, v_n)\} \\
\llbracket \text{let } x = E_1 \text{ in } E_2 \rrbracket &\stackrel{\text{def}}{=} \{\langle \perp, \mathcal{A}_1, \text{lab}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1 \rangle \mid \langle \perp, \mathcal{A}_1, \text{lab}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1 \rangle \in \llbracket E_1 \rrbracket\} \\
&\cup \{\langle \text{res}_2, \mathcal{A}_1 \uplus \mathcal{A}_2, \text{lab}_1 \cup \text{lab}_2, \text{sb}_1 \cup \text{sb}_2 \cup \{(lst_1, \text{fst}_2)\}, \text{fst}_1, \text{lst}_2 \rangle \mid \\
&\quad \langle v_1, \mathcal{A}_1, \text{lab}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1 \rangle \in \llbracket E_1 \rrbracket \wedge \langle \text{res}_2, \mathcal{A}_2, \text{lab}_2, \text{sb}_2, \text{fst}_2, \text{lst}_2 \rangle \in \llbracket E_2[v_1/x] \rrbracket\} \\
\llbracket \text{repeat } E \text{ end} \rrbracket &\stackrel{\text{def}}{=} \{\langle \text{res}_N, \biguplus_{i \in [1..N]} \mathcal{A}_i, \bigcup_{i \in [1..N]} \text{lab}_i, \bigcup_{i \in [1..N]} \text{sb}_i \cup \{(lst_1, \text{fst}_2), \dots, (lst_{N-1}, \text{fst}_N)\}, \text{fst}_1, \text{lst}_N \rangle \mid \\
&\quad \forall i. \langle \text{res}_i, \mathcal{A}_i, \text{lab}_i, \text{sb}_i, \text{fst}_i, \text{lst}_i \rangle \in \llbracket E \rrbracket \wedge (i \neq N \implies \text{res}_i = 0) \wedge \text{res}_N \neq 0\} \\
\llbracket E_1 \parallel E_2 \rrbracket &\stackrel{\text{def}}{=} \{\langle \text{combine}(\text{res}_1, \text{res}_2), \mathcal{A}_1 \uplus \mathcal{A}_2 \uplus \{a_{\text{fork}}, a_{\text{join}}\}, \text{lab}_1 \cup \text{lab}_2 \cup \{a_{\text{fork}} \mapsto \text{skip}, a_{\text{join}} \mapsto \text{skip}\}, \\
&\quad \text{sb}_1 \cup \text{sb}_2 \cup \{(a_{\text{fork}}, \text{fst}_1), (a_{\text{fork}}, \text{fst}_2), (lst_1, a_{\text{join}}), (lst_2, a_{\text{join}})\}, a_{\text{fork}}, a_{\text{join}} \rangle \mid \\
&\quad \langle \text{res}_1, \mathcal{A}_1, \text{sb}_1, \text{fst}_1, \text{lst}_1 \rangle \in \llbracket E_1 \rrbracket \wedge \langle \text{res}_2, \mathcal{A}_2, \text{sb}_2, \text{fst}_2, \text{lst}_2 \rangle \in \llbracket E_2 \rrbracket \wedge a_{\text{fork}}, a_{\text{join}} \in \text{AName}\}
\end{aligned}$$

Figure 2. Semantics of closed program expressions.

$$\begin{aligned}
&\nexists x. \text{hb}(x, x) && \text{(IrreflexiveHB)} \\
&\forall \ell. \text{totalorder}(\{a \in \mathcal{A} \mid \text{iswrite}_\ell(a)\}, \text{mo}) \wedge \text{hb}_\ell \subseteq \text{mo} && \text{(ConsistentMO)} \\
&\text{totalorder}(\{a \in \mathcal{A} \mid \text{isSeqCst}(a)\}, \text{sc}) \wedge \text{hb}_{\text{SeqCst}} \subseteq \text{sc} \wedge \text{mo}_{\text{SeqCst}} \subseteq \text{sc} && \text{(ConsistentSC)} \\
&\forall b. \text{rf}(b) \neq \perp \iff \exists \ell, a. \text{iswrite}_\ell(a) \wedge \text{isread}_\ell(b) \wedge \text{hb}(a, b) && \text{(ConsistentRFdom)} \\
&\forall a, b. \text{rf}(b) = a \implies \exists \ell, v. \text{iswrite}_{\ell, v}(a) \wedge \text{isread}_{\ell, v}(b) \wedge \neg \text{hb}(b, a) && \text{(ConsistentRF)} \\
&\forall a, b. \text{rf}(b) = a \wedge (\text{mode}(a) = \text{na} \vee \text{mode}(b) = \text{na}) \implies \text{hb}(a, b) && \text{(ConsistentRFna)} \\
&\forall a, b. \text{rf}(b) = a \wedge \text{isSeqCst}(b) \implies \text{isc}(a, b) \vee \neg \text{isSeqCst}(a) \wedge (\forall x. \text{isc}(x, b) \implies \neg \text{hb}(a, x)) && \text{(RestrSCReads)} \\
&\nexists a, b. \text{hb}(a, b) \wedge \text{mo}(\text{rf}(b), \text{rf}(a)) \wedge \text{locs}(a) = \text{locs}(b) && \text{(CoherentRR)} \\
&\nexists a, b. \text{hb}(a, b) \wedge \text{mo}(\text{rf}(b), a) \wedge \text{iswrite}(a) \wedge \text{locs}(a) = \text{locs}(b) && \text{(CoherentWR)} \\
&\nexists a, b. \text{hb}(a, b) \wedge \text{mo}(b, \text{rf}(a)) \wedge \text{iswrite}(b) \wedge \text{locs}(a) = \text{locs}(b) && \text{(CoherentRW)} \\
&\forall a. \text{isrmw}(a) \wedge \text{rf}(a) \neq \perp \implies \text{mo}(\text{rf}(a), a) \wedge \nexists c. \text{mo}(\text{rf}(a), c) \wedge \text{mo}(c, a) && \text{(AtomicRMW)} \\
&\forall a, b, \ell. \text{lab}(a) = \text{lab}(b) = \text{A}(\ell) \implies a = b && \text{(ConsistentAlloc)}
\end{aligned}$$

where $\text{iswrite}_{\ell, v}(a) \stackrel{\text{def}}{=} \exists X, v_{\text{old}}. \text{lab}(a) \in \{\text{W}_X(\ell, v), \text{RMW}_X(\ell, v_{\text{old}}, v)\}$ $\text{iswrite}_\ell(a) \stackrel{\text{def}}{=} \exists v. \text{iswrite}_{\ell, v}(a)$

$\text{isread}_{\ell, v}(a) \stackrel{\text{def}}{=} \exists X, v_{\text{new}}. \text{lab}(a) \in \{\text{R}_X(\ell, v), \text{RMW}_X(\ell, v, v_{\text{new}})\}$ etc.

$\text{rsElem}(a, b) \stackrel{\text{def}}{=} \text{sameThread}(a, b) \vee \text{isrmw}(b)$

$\text{rseq}(a) \stackrel{\text{def}}{=} \{a\} \cup \{b \mid \text{rsElem}(a, b) \wedge \text{mo}(a, b) \wedge (\forall c. \text{mo}(a, c) \wedge \text{mo}(c, b) \implies \text{rsElem}(a, c))\}$

$\text{sw} \stackrel{\text{def}}{=} \{(a, b) \mid \text{mode}(a) \in \{\text{rel}, \text{rel_acq}, \text{sc}\} \wedge \text{mode}(b) \in \{\text{acq}, \text{rel_acq}, \text{sc}\} \wedge \text{rf}(b) \in \text{rseq}(a)\}$

$\text{hb} \stackrel{\text{def}}{=} (\text{sb} \cup \text{sw})^+$

$\text{hb}_\ell \stackrel{\text{def}}{=} \{(a, b) \in \text{hb} \mid \text{iswrite}_\ell(a) \wedge \text{iswrite}_\ell(b)\}$

$X_{\text{SeqCst}} \stackrel{\text{def}}{=} \{(a, b) \in X \mid \text{isSeqCst}(a) \wedge \text{isSeqCst}(b)\}$

$\text{isc}(a, b) \stackrel{\text{def}}{=} \text{iswrite}_{\text{locs}(b)}(a) \wedge \text{sc}(a, b) \wedge \nexists c. \text{sc}(a, c) \wedge \text{sc}(c, b) \wedge \text{iswrite}_{\text{locs}(b)}(c)$

Figure 3. Axioms satisfied by consistent C11 executions, $\text{Consistent}(\mathcal{A}, \text{lab}, \text{sb}, \text{rf}, \text{mo}, \text{sc})$.




Figure 4. Sample executions violating coherency conditions (Batty et al. 2011).

Consistent Executions According to the C11 model, an execution is consistent, $\text{Consistent}(\mathcal{A}, \text{lab}, \text{sb}, \text{rf}, \text{mo}, \text{sc})$, if all of the properties shown in Figure 3 hold.

- (IrreflexiveHB) The happens-before order, hb , must be irreflexive: an action cannot happen before itself.
- (ConsistentMO) All write actions on an atomic location ℓ must be totally ordered by mo , and be consistently ordered by hb (restricted to the location ℓ).
- (ConsistentSC) The sc relation must be a total order and include both hb and mo restricted to SC actions. This in effect means that SC actions are globally synchronized.
- (ConsistentRFdom) The reads-from map, rf , is defined for those read (or RMW) actions for which the execution contains an earlier write (or RMW) to the same location.
- (ConsistentRF) Each entry in the reads-from map, rf , should map a read to an earlier or concurrent write to the same location and with the same value.
- (ConsistentRFna) Further, if a read reads from a write and either the read or write are non-atomic, then the write must have happened before the read. Batty et al. (2011) also require the write to be *visible*: i.e. not to have been overwritten by another write that happened before the read. This extra condition is unnecessary, as it follows from CoherentWR.
- (RestrSCreads) SC reads are further restricted to read only from the immediately preceding SC write to the same location in sc order or from a non-SC write that has not happened before that immediately preceding SC write.
- (CoherentRR, CoherentWR, CoherentRW) Next, we have three per-location coherence properties relating mo , hb , and rf . These properties require that mo never contradicts hb or the observed read order, and that rf never reads values that have been overwritten by more recent actions that happened before the read. These coherence properties are depicted in Figure 4.
- (AtomicRMW) Each read-modify-write action should execute atomically: it should read from the immediately preceding write in mo .
- (ConsistentAlloc) Finally, the same location cannot be allocated twice by different allocation actions.²

Remark Our model differs in a few minor ways from that of Batty et al. (2011, 2012). First, we have incorporated the C standard’s “additional synchronized with” (asw) relation in sb rather than in sw , because it describes synchronization induced by control flow rather than by data flow.

Second, our sw -relation also relates acquire reads with release writes (whenever the read returns a value written by or after the release write), even if the two actions belong

²This axiom suffices, because we do not support deallocation. Had we included deallocation, we would instead require there to be a deallocation actions between any two allocation actions of the same location. The formalized C11 model by Batty et al. (2011, 2012) does not model allocation.

```

let  $a = \text{alloc}()$  in
let  $c = \text{alloc}()$  in
   $[c]_{\text{rlx}} := 0;$ 
   $\left( \begin{array}{l} [a]_{\text{na}} := 7; \\ [c]_{\text{rel}} := 1 \end{array} \parallel \text{repeat } [c]_{\text{acq}} \text{ end}; \right)$ 

```

Figure 5. Message passing example showing transfer of ownership of the non-atomic location a .

to the same thread (and are thus sb -related), whereas Batty et al. (2012) do not add any sb -related actions to the sw -relation. Since relating such actions also by sw does not affect execution consistency, we do so for uniformity, which eases the definition of validity of Hoare triples in §7.

Finally, in the standard, the sb and sw relations are taken to be strict partial orders, corresponding to the transitive closure of our relations. Conversely, our sb relation can be defined in terms of the sb order from the C and C++ standards as follows, $\text{sb}_{\text{our}} = \{(a, b) \in \text{sb}_{\text{std}} \cup \text{asw}_{\text{std}} \mid \nexists c. (a, c) \in \text{sb}_{\text{std}} \cup \text{asw}_{\text{std}} \wedge (c, b) \in \text{sb}_{\text{std}} \cup \text{asw}_{\text{std}}\}$. Again, we found the non-transitive versions slightly more convenient when defining the meaning of Hoare triples.

4. Relaxed Separation Logic

To motivate RSL, consider the message passing program shown in Figure 5. The thread on the left updates some data structure using non-atomic memory accesses (here, the location a), and then signals to other threads that the data structure has been updated by performing a release write to c . The thread on the right repeatedly performs acquire reads until it notices that $[c] \neq 0$. Then, it can conclude that the thread on the left has finished its work, and so may safely access the data structure without interfering with it.

This message passing idiom is correct (i.e., race-free) because whenever an acquire read sees the value written by a release write, the write “synchronizes with” the acquire read. Thus, as hb is transitive, any event that happened before the write (e.g., by being sequenced before it), also happens before the read. This, in turn, justifies the ownership transfer from the writing thread to the reading thread.

To model such ownership transfers, RSL extends the grammar of separation logic assertions, P , with three new assertion forms, $\text{Rel}(e, \mathcal{Q})$, $\text{Acq}(e, \mathcal{Q})$, and $\text{RMWAcq}(e, \mathcal{Q})$, where \mathcal{Q} ranges over functions from values to assertions. Formally, RSL assertions are given by following grammar:

$$\begin{aligned}
 P, Q ::= & \text{false} \mid P \Rightarrow Q \mid \forall x. P \mid \text{emp} \mid e \stackrel{k}{\mapsto} e' \mid P * Q \\
 & \mid \text{Rel}(e, \mathcal{Q}) \mid \text{Acq}(e, \mathcal{Q}) \mid \text{RMWAcq}(e, \mathcal{Q}) \\
 & \mid \text{Init}(e) \mid \text{Uninit}(e)
 \end{aligned}$$

where k ranges over fractional permissions ($\text{Perm} = (0, 1]$, see Boyland 2003). We have the usual classical first order logic constructs (the three primitive ones and the derived: true , \wedge , \vee , \neg , \exists), the three assertions forms pertinent to separation logic (empty heap, a single memory cell with

Note that this rule does not assert anything about the value read. A more useful rule is the following, which asserts that the value read must be one that may have been written.

$$\left\{ \begin{array}{l} \text{Init}(\ell) * \\ \text{Acq}(\ell, \mathcal{Q}) \end{array} \right\} [\ell]_{\text{rlx}} \left\{ \begin{array}{l} v. \text{Acq}(\ell, \mathcal{Q}) \wedge \\ (\mathcal{Q}(v) \neq \text{false}) \end{array} \right\} \quad (\text{R-RLX}^*)$$

Similar to W-RLX^* , this latter rule is not sound in C11, but is so in the strengthened model of Section 6.

In RSL, we permit multiple readers to read the value written by a single release write. Concretely, consider the scenario where thread A initializes two data structures and signals by a release write that it has finished its work. Then thread B can do an acquire read and notice that A has finished its initialization and then access the first data structure non-atomically. Likewise, thread C can do an acquire read and access the second data structure non-atomically. Such an execution does not have data races and should therefore be permitted. In terms of our program logic, this means that acquire read permissions should be splittable and joinable as follows:

$$\begin{array}{l} \text{Acq}(\ell, \mathcal{Q}_1) * \text{Acq}(\ell, \mathcal{Q}_2) \\ \iff \text{Acq}(\ell, \lambda v. \mathcal{Q}_1(v) * \mathcal{Q}_2(v)) \end{array} \quad (\text{ACQ-SPLIT})$$

Read-Modify-Write Instructions The next new assertion form, $\text{RMWAcq}(\ell, \mathcal{Q})$, is used in the following proof rule for atomic compare-and-swaps.

$$\frac{\begin{array}{l} P \Rightarrow \text{Init}(\ell) * \text{RMWAcq}(\ell, \mathcal{Q}) * \text{true} \\ P * \mathcal{Q}(v) \Rightarrow \text{Rel}(\ell, \mathcal{Q}') * \mathcal{Q}'(v') * R[v/y] \\ X \in \{\text{rel}, \text{rlx}\} \Rightarrow \mathcal{Q}(v) = \text{emp} \\ X \in \{\text{acq}, \text{rlx}\} \Rightarrow \mathcal{Q}'(v') = \text{emp} \\ \{P\} [\ell]_Y \{y. y \neq v \Rightarrow R\} \end{array}}{\{P\} \text{CAS}_{X,Y}(\ell, v, v') \{y. R\}} \quad (\text{CAS}^*)$$

The rule has five premises. First, the precondition must ensure that we have permission to do a RMW-read from ℓ and acquire ownership of $\mathcal{Q}(v)$. Second, we require the update performed by the successful CAS to be valid: that is, to have the necessary release permission, to satisfy $\mathcal{Q}'(v')$, the assertion that is to be transferred away, and to separately also satisfy the postcondition. As a precondition for this update, we get to assume not only that the initial precondition holds, but also that we have access to the state acquired by ownership transfer, $\mathcal{Q}(v)$.

The next two premises take the access modes into account, suitably restricting the ownership that can be acquired or released. If the successful CAS is of release or relaxed kind, then it does not synchronize with the write whose value it read, so it should not acquire any ownership. This is ensured by demanding that $\mathcal{Q}(v) = \text{emp}$. Symmetrically, if the successful CAS is of acquire or relaxed kind, it does not synchronize with the reads seeing the value it produced, so it should not release any ownership. This is ensured by demanding that $\mathcal{Q}'(v') = \text{emp}$.

Finally, we require that failed CASs also satisfy the postcondition, R , under the assumption that a value different from the expected one was read.

In its general form, the CAS^* rule is sound in the strengthened model of Section 6. In the standard model, it is sound only when $X \in \{\text{rel_acq}, \text{sc}\}$.

Unlike multiple normal reads, multiple successful CAS instructions cannot all read from (and therefore potentially synchronize with) the same write. This follows from the AtomicRMW axiom, which requires RMW actions to read from the immediately preceding write in mo-order. Therefore, it is sound to duplicate the RMW-acquire permission,

$$\begin{array}{l} \text{RMWAcq}(\ell, \mathcal{Q}) \\ \iff \text{RMWAcq}(\ell, \mathcal{Q}) * \text{RMWAcq}(\ell, \mathcal{Q}) \end{array} \quad (\text{RMW-SPLIT})$$

because the semantics ensures that at most one process will effectively be able to use this permission at any given instant.

In order to be able to prove the last premise of the CAS^* rule, we also support the following rule, allowing us to carve out a plain acquire permission from an RMW-acquire one.

$$\frac{\forall v. \mathcal{Q}'(v) = \text{emp} \vee \mathcal{Q}(v) = \mathcal{Q}'(v) = \text{false}}{\text{RMWAcq}(\ell, \mathcal{Q}) \iff \text{RMWAcq}(\ell, \mathcal{Q}) * \text{Acq}(\ell, \mathcal{Q}')} \quad (\text{RMW-ACQ-SPLIT})$$

The premise of RMW-ACQ-SPLIT ensures that the assertion that we have carved out for plain reads is empty, except perhaps for the values where $\mathcal{Q}(v)$ is false, in which case $\mathcal{Q}'(v)$ may also be false.

Allocation of Atomic Locations Whenever a new atomic location is allocated, the verifier is free to choose a suitable ownership assertion \mathcal{Q} and attach it to the newly allocated location, and moreover to choose whether the ownership of \mathcal{Q} will be acquired using plain reads or using successful CASs. We thus have the following two rules.

$$\begin{array}{l} \{\text{emp}\} \text{alloc}() \{ \ell. \text{Rel}(\ell, \mathcal{Q}) * \text{Acq}(\ell, \mathcal{Q}) \} \quad (\text{A-R}) \\ \{\text{emp}\} \text{alloc}() \{ \ell. \text{Rel}(\ell, \mathcal{Q}) * \text{RMWAcq}(\ell, \mathcal{Q}) \} \quad (\text{A-M}) \end{array}$$

Following the C standard, newly allocated locations are not initialized, and thus do not generate the $\text{Init}(\ell)$ permission required for reading them. To enable reading from these locations, the programmer must first initialize them by performing a plain write as we have already seen.

Non-Atomic Locations Finally, the rules for non-atomic accesses are exactly as in concurrent separation logic. Allocation returns an uninitialized new cell with full permission; writing requires full permission of a location (whether initialized or not), whereas reading works also with partial permission but requires the location to be initialized.

$$\begin{array}{l} \{\text{emp}\} \text{alloc}() \{x. \text{Uninit}(x)\} \quad (\text{A-NA}) \\ \left\{ \ell \xrightarrow{1} _ \vee \text{Uninit}(\ell) \right\} [\ell]_{\text{na}} := v \left\{ \ell \xrightarrow{1} v \right\} \quad (\text{W-NA}) \\ \left\{ \ell \xrightarrow{k} v \right\} [\ell]_{\text{na}} \left\{ x. x = v \wedge \ell \xrightarrow{k} v \right\} \quad (\text{R-NA}) \end{array}$$

Let $\mathcal{Q}_J(v) \stackrel{\text{def}}{=} (v = 0 \wedge \text{emp}) \vee (v = 1 \wedge J)$
 $\text{Lock}(x, J) \stackrel{\text{def}}{=} \text{Rel}(x, \mathcal{Q}_J) * \text{RMWA}c\text{q}(x, \mathcal{Q}_J) * \text{Init}(x)$

| | |
|---|--|
| $\text{new_lock}() \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \{J\} \\ \text{let } x = \text{alloc}() \text{ in} \\ \left\{ \begin{array}{l} J * \text{Rel}(x, \mathcal{Q}_J) * \\ \text{RMWA}c\text{q}(x, \mathcal{Q}_J) \end{array} \right\} \\ [x]_{\text{rel}} := 1 \\ \{ \text{Lock}(x, J) \} \end{array} \right.$ | $\text{lock}(x) \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \{ \text{Lock}(x, J) \} \\ \text{repeat} \\ \left\{ \begin{array}{l} \text{Lock}(x, J) \\ \text{spin}(x); \\ \text{Lock}(x, J) \end{array} \right\} \\ \text{CAS}_{\text{acq,rlx}}(x, 1, 0) \\ \left\{ \begin{array}{l} y. \text{Lock}(x, J) * \\ (y = 1 \wedge J \vee \\ y = 0 \wedge \text{emp}) \end{array} \right\} \end{array} \right.$ |
| $\text{unlock}(x) \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \{ J * \text{Lock}(x, J) \} \\ [x]_{\text{rel}} := 1 \\ \{ \text{Init}(x) * \text{Lock}(x, J) \} \\ \{ \text{Lock}(x, J) \} \end{array} \right.$ | $\text{end} \\ \{ J * \text{Lock}(x, J) \}$ |

Figure 7. Verification of the lock module.

Let $\mathcal{Q}(x) \stackrel{\text{def}}{=} \text{if } x = 0 \text{ then emp else } a \mapsto 7$
 $\{ \text{emp} \}$
 $\text{let } a = \text{alloc}() \text{ in}$
 $\{ \text{Uninit}(a) \}$
 $\text{let } c = \text{alloc}() \text{ in}$
 $\{ \text{Uninit}(a) * \text{Rel}(c, \mathcal{Q}) * \text{Acq}(c, \mathcal{Q}) \}$
 $[c]_{\text{rlx}} := 0;$
 $\left\{ \begin{array}{l} \{ \text{Uninit}(a) * \text{Rel}(c, \mathcal{Q}) * \text{Acq}(c, \mathcal{Q}) * \text{Init}(c) \} \\ \{ \text{Uninit}(a) * \text{Rel}(c, \mathcal{Q}) \} \end{array} \right\} \parallel \left\{ \begin{array}{l} \{ \text{Acq}(c, \mathcal{Q}) * \text{Init}(c) \} \\ \text{repeat } [c]_{\text{acq}} \text{ end} \\ \{ \text{true} * a \mapsto 7 \} \\ [a]_{\text{na}} := [a]_{\text{na}} + 1 \\ \{ \text{true} * a \mapsto 8 \} \end{array} \right.$
 $[a]_{\text{na}} := 7$
 $\{ a \mapsto 7 * \text{Rel}(c, \mathcal{Q}) \}$
 $[c]_{\text{rel}} := 1$
 $\{ \text{Rel}(c, \mathcal{Q}) \}$
 $\{ a \mapsto 8 * \text{true} \}$

Figure 8. Verification of the message passing example.

These rules ensure that all accessed location have been allocated and there are no races on non-atomic memory locations, and moreover that only initialized locations are read.

5. Examples

We now illustrate RSL by proving simple race-free programs involving release-acquire synchronization patterns. Ownership transfer along those release-acquire synchronizations is necessary to prove them correct, that is, to show that they are memory safe and do not contain data races. To make the proof outlines more concise, we define the following shorthand notations.

$\text{Emp} \stackrel{\text{def}}{=} \lambda v. \text{emp}$
 $\text{IAc}q(\ell, v, P) \stackrel{\text{def}}{=} \text{Init}(\ell) * \text{Ac}q(\ell, \text{Emp}[v := P])$
 $\text{IRMWA}c\text{q}(\ell, v, P) \stackrel{\text{def}}{=} \text{Init}(\ell) * \text{RMWA}c\text{q}(\ell, \text{Emp}[v := P])$

Figure 7: Lock Module As our first example, we consider the lock module introduced in Figure 1. Here we show that

any invariant J may be attached to a lock so that we get the same specifications as in concurrent separation logic:

$$\begin{aligned} & \{J\} \text{new_lock}() \{x. \text{Lock}(x, J)\} \\ & \{ \text{Lock}(x, J) \} \text{lock}(x) \{ J * \text{Lock}(x, J) \} \\ & \{ J * \text{Lock}(x, J) \} \text{unlock}(x) \{ \text{Lock}(x, J) \} \\ & \text{Lock}(x, J) \iff \text{Lock}(x, J) * \text{Lock}(x, J) \end{aligned}$$

As expected, creating a lock requires the invariant J to hold initially and returns a token confirming that the lock exists and protects the invariant J . Acquiring the lock requires this token and obtains ownership of the invariant. Conversely, releasing the lock requires the invariant to hold and transfers it away. Finally, the token saying that x is a lock protecting resource J can be freely duplicated.

To derive this specification, we define the predicates:

$\mathcal{Q}_J(v) \stackrel{\text{def}}{=} (v = 0 \wedge \text{emp}) \vee (v = 1 \wedge J)$
 $\text{Lock}(x, J) \stackrel{\text{def}}{=} \text{Rel}(x, \mathcal{Q}_J) * \text{RMWA}c\text{q}(x, \mathcal{Q}_J) * \text{Init}(x)$

The $\mathcal{Q}_J(v)$ predicate describes the invariant associated with the location x implementing the lock. It assigns empty ownership when the lock is held ($v = 0 \wedge \text{emp}$), and ownership of the invariant, J , when the lock is free ($v = 1 \wedge J$). The $\text{Lock}(x, J)$ predicate contains permissions to access the lock by performing release-writes and acquire-RMWs. It also contains the knowledge that the lock is initialized.

In $\text{new_lock}()$, we use the A-M and W-REL rules to initialize the location and transfer away the ownership of J . Similarly, in $\text{unlock}(x)$, we use the W-REL to transfer away the ownership of J and then the INIT-SPLIT rule to remove the duplicate $\text{Init}(x)$ fact. In $\text{lock}(x)$, we use the R-RLX rule for the relaxed optimistic read in the $\text{spin}(x)$ loop, and then the CAS* and the R-RLX* rules to deal with the CAS. Finally, the fact that the $\text{Lock}(x, J)$ predicate can be freely duplicated follows immediately from REL-SPLIT, RMW-SPLIT, and INIT-SPLIT.

Figure 8: Message Passing As our second example, we consider the message passing idiom of Figure 5. Here, by constructing a proof, we conclude that the program has no data races and moreover, when both threads terminate, we have $[a] = 8$. The proof illustrates the use of the $\text{Ac}q(-, -)$ predicate, and the rules A-NA, A-R, W-RLX*, W-NA, W-REL, R-ACQ, and R-NA.

Figure 9: Partial Ownership Transfer Our next example is a variant of the message passing program we have just seen, where after the synchronization between the two threads, both threads read from a . This is valid because two concurrent read accesses do not count as a data race.

In order to verify this program, we use fractional permissions and transfer the partial ownership of the non-atomic location a from the first to the second thread. The first thread writes to a , and then performs a release write to x , giving

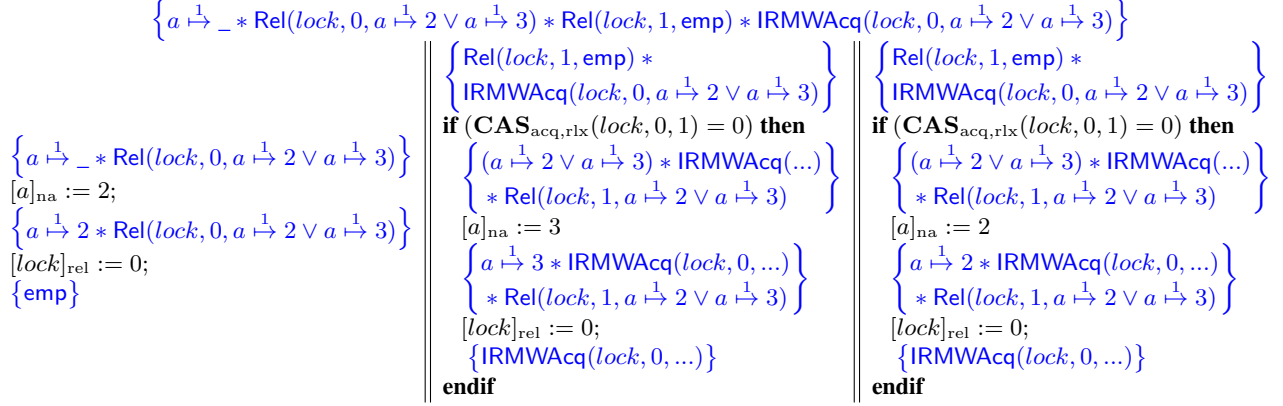


Figure 10. Example illustrating the use of CAS to implement a lock.

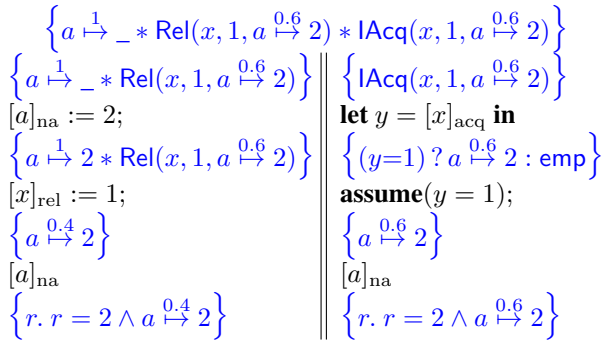


Figure 9. Example illustrating fractional ownership and transfer thereof.

away the partial permission $a \mapsto^{0.6} 2$ (using the W-REL rule). With its remaining $a \mapsto^{0.4} 2$ permission, it then reads a using the R-NA rule. The second thread synchronizes with the write to x and gets the $a \mapsto^{0.6} 2$ permission (using the R-ACQ rule), after which it reads a and also gets the value 2 (using the R-NA rule).

Figure 10: Transfer of Permission in Both Directions

Our next example demonstrates the use of CAS directly to implement a simple mutual exclusion lock. (We could of course use the lock module verified previously, but we include this example in order to demonstrate the CAS* rule again.) Here, for a change, we implement a non-blocking “tryLock” command using a conditional, rather than a blocking locking command using a loop.

The lock is implemented by a single location, lock , storing 0 if the lock is free and 1 if it is held. (This is opposite to the covention of the earlier lock module.) The lock protects a resource invariant describing the memory cell a . Initially, the first thread starts with the lock acquired and owning a : it establishes the resource invariant and releases the lock. The other two threads start without knowing that the lock was initially held; they both have the permission to write the value 1 to the lock without releasing any owner-




Figure 11. Program with a possible dependency cycle.




Figure 12. Execution exhibiting the dependency cycle.

ship, $\text{Rel}(\text{lock}, 1, \text{emp})$. By itself, this permission is pretty useless: the threads can do a blind relaxed-atomic write to lock setting its value to 1 (acquired) but without gaining any information. What makes this permission useful, is its combination with the other permission they have, namely to read the state of an unacquired lock with a CAS and get ownership of the resource invariant. Successfully performing the CAS enables them to later release the lock with the same resource invariant.

6. Dealing with Relaxed Memory Accesses

A serious deficiency of the C11 memory model is that it allows “out of thin air” reads, as illustrated by the program in Figure 11, adapted from Batty et al. (2013).

In this program, two locations are initialized with the value 0, and then two threads are forked, each writing 1 to the one location provided the other already has the value 1. Intuitively, one would expect that the writes would never be executed, but actually the C11 concurrency model permits this outcome. The questionable execution, depicted in Figure 12, is consistent as the two reads can get the value 1 by reading from the corresponding conditional stores.

This counterintuitive behaviour is extremely problematic for formal reasoning as it inhibits even the simplest form

```

let a = alloc() in [a]rlx := 0;
let b = alloc() in [b]rlx := 1;
([b]rlx := [a]rlx || [a]rlx := [b]rlx);
if [a]rlx < 20 then print [a]rlx

```

Figure 13. Program showing that range analysis is unsound under C11.

of thread-local reasoning, that of non-relational conjunctive invariants (i.e., invariants where each conjunct describes a property of only one variable). Intuitively, in the previous program, one would expect the invariant

$$[a]_{\text{rlx}} = 0 \wedge [b]_{\text{rlx}} = 0$$

to hold throughout the parallel composition since it holds initially and is preserved by every ‘reachable’ atomic statement of the program, arguing that the conditional stores are not reachable because the conditions are unsatisfiable according to the invariant. This kind of reasoning is performed by standard compiler optimizations such as “sparse conditional constant propagation” (Wegman and Zadeck 1991).

Note that with the W-RLX^* and R-RLX^* rules, we can easily prove that if we were to read $[a]$ at the end of the program, we would get 0. (To do so, pick $\mathcal{Q} := (\lambda x. x = 0)$ in the allocation rule for both locations.) This shows that these two rules are unsound under C11.

Observe that the same problematic execution remains consistent even if we strengthen either the relaxed reads to acquire/SC reads or (exclusively) the relaxed writes to release/SC writes. To make this execution inconsistent, we have to strengthen both the reads and the writes except at most one access. This means that even adding one of the W-RLX^* and R-RLX^* rules is unsound.

Global Range Analysis A concrete optimization that is unsound under C11 is global range analysis. Consider the program in Figure 13. An optimizing compiler may argue that the test $[a]_{\text{rlx}} < 20$ will always succeed because $[a]$ and $[b]$ store either 0 or 1, and therefore replace the conditional expression by the **then** branch. Somewhat surprisingly, under C11, this transformation introduces new behaviour and is therefore unsound. Because of the causal dependency cycle, the $[a]_{\text{rlx}}$ read can return any arbitrary value. Therefore, the transformed program can print any arbitrary value, while the original one could only print values less than 20.

A Crude Fix to the Model Since even this very basic reasoning is unsound for relaxed accesses, we decided to strengthen the C11 concurrency model with the following axiom stating that $\text{hb} \cup \text{rf}$ must be acyclic (i.e., its transitive closure must be irreflexive).

$$\text{acyclic}(\text{hb} \cup \{\{\text{rf}(a), a\} \mid a \in \mathcal{A}\}) \quad (\text{StrongAcyclicHB})$$

where $\text{acyclic}(R) \stackrel{\text{def}}{=} \nexists x \in \mathcal{A}. R^+(x, x)$.

```

let a = alloc() in [a]rlx := 0;
let b = alloc() in [b]rlx := 0;
( let x = [a]rlx in ) || ( let y = [b]rlx in )
[ b ]rlx := 1 || [ a ]rlx := 1

```

Figure 14. Program without a dependency cycle.

With this additional axiom, we can also show the soundness of the “starred” rules for relaxed memory accesses presented in the previous section. In contrast, the soundness of the other rules does not depend on this axiom.

Notice that when adding this strong acyclicity condition, we can drop the strictly weaker IrreflexiveHB axiom, as well as the $\neg\text{hb}(b, a)$ conjunct from the ConsistentRF axiom. We can further drop the slightly awkward ConsistentRFna axiom, and still have the soundness proof go through, because all the proof really needs to know is that the write precedes the read in some well-founded order. In the absence of causal cycles, this order need not be hb : we can instead take it to be $\text{hb} \cup \text{rf}$.

Simple though our proposed fix might seem, it is not perfect. Alas, the StrongAcyclicHB consistency axiom precludes the reordering of independent instructions, a transformation that compilers and processors with out-of-order execution units frequently perform. To illustrate the problem, consider the program in Figure 14, a slight variant of the program in Figure 11, where the writes to $[b]$ and $[a]$ are now independent of the earlier reads from $[a]$ and $[b]$ respectively. The problem is that when operating at the level of single executions, one cannot distinguish whether the $\text{hb} \cup \text{rf}$ cycle in the execution shown in Figure 12 constitutes a dependency cycle or not. If the execution comes from the program in Figure 11, the cycle should clearly be outlawed, but if it comes from the program of Figure 14, the cycle is harmless and should be allowed. Distinguishing these two cases is not easy and seems to require a radical change to the C11 model. Clearly, this lies beyond the scope of this paper.

7. Semantics and Soundness

In this section, we define the semantics of assertions and Hoare triples, and prove that our logic is sound with respect to the C11 memory model.

7.1 Semantics of Assertions

To define the meaning of separation logic assertions, we need an underlying separation algebra, i.e. a commutative partial monoid. To interpret the Acq and Rel assertions, our model of heaps will have to store assertions, which in turn represent sets of heaps. If we naively write down the domain equation, we will get an equation of the form,

$$\text{Heap}_{\text{spec}} \stackrel{?}{\cong} \text{Loc} \multimap (\dots + (\dots \times \mathbb{P}(\text{Heap}_{\text{spec}}))),$$

which does not have a solution in Set . Therefore, we either have to move to a more advanced category such as bounded

$$\begin{array}{l}
\mathcal{Q}_1 \oplus_{\text{acq}}^{b_1, b_2} \mathcal{Q}_2 \stackrel{\text{def}}{=} \begin{cases} \mathcal{Q}_1 & \text{if } b_1 \wedge b_2 \text{ and } \mathcal{Q}_1 = \mathcal{Q}_2 \\ \lambda v. \mathcal{Q}_1(v) * \mathcal{Q}_2(v) & \text{if } (\neg b_1 \vee \neg b_2) \text{ and } \mathbf{adef}(b_1, \mathcal{Q}_1, b_2, \mathcal{Q}_2) \\ \text{undef} & \text{if } \neg \mathbf{adef}(b_1, \mathcal{Q}_1, b_2, \mathcal{Q}_2) \end{cases} \\
h_1 \oplus' h_2 \stackrel{\text{def}}{=} \lambda \ell. \begin{cases} h_1(\ell) & \text{if } \ell \in \text{dom}(h_1) \setminus \text{dom}(h_2) \\ h_2(\ell) & \text{if } \ell \in \text{dom}(h_2) \setminus \text{dom}(h_1) \\ \text{NA}[v, k_1 + k_2] & \text{if } h_i(\ell) = \text{NA}[v, k_i] \text{ for } i = 1, 2 \text{ and } k_1 + k_2 \leq 1 \\ \text{Atom}[\lambda v. \mathcal{R}_1(v) \vee \mathcal{R}_2(v), \mathcal{Q}_1 \oplus_{\text{acq}}^{b_1, b_2} \mathcal{Q}_2, b_1 \vee b_2, b'_1 \vee b'_2] & \text{if } h_i(\ell) = \text{Atom}[\mathcal{R}_i, \mathcal{Q}_i, b_i, b'_i] \text{ for } i = 1, 2 \\ \text{undef} & \text{otherwise} \end{cases} \\
h_1 \oplus h_2 \stackrel{\text{def}}{=} \begin{cases} h_1 \oplus' h_2 & \text{if } \text{dom}(h_1 \oplus' h_2) = \text{dom}(h_1) \cup \text{dom}(h_2) \\ \text{undef} & \text{otherwise} \end{cases}
\end{array}$$

$$\begin{array}{l}
\mathbf{rval}(b, \mathcal{Q}) \stackrel{\text{def}}{=} \text{if } b \text{ then } \mathcal{Q} \text{ else } \lambda v. \text{emp} \\
\mathbf{adef}(b_1, \mathcal{Q}_1, b_2, \mathcal{Q}_2) \stackrel{\text{def}}{=} \\
\quad \forall v. \mathbf{rval}(b_2, \mathcal{Q}_1)(v) = \mathbf{rval}(b_1, \mathcal{Q}_2)(v) \\
\quad \vee \mathcal{Q}_1(v) = \mathcal{Q}_2(v) = \text{false}
\end{array}$$

Figure 15. Definition of heap composition, $h_1 \oplus h_2$.

ultrametric spaces (Birkedal et al. 2010), or change the equation to avoid the problematic recursion.

Here, for simplicity, we do the latter and cut the cycle by storing syntactic assertions, Assn , instead of semantic assertions, $\mathbb{P}(\text{Heap}_{\text{spec}})$, within heaps. Simply storing syntactic assertions is, however, insufficient because we want the heap model to form a separation algebra and to support the conversions rules REL-SPLIT and ACQ-SPLIT . To allow these conversions, we therefore have to store syntactic assertions up to associativity and commutativity of $*$ and \vee and their units. Furthermore, to support RMW-ACQ-SPLIT , we also need to equate $\text{false} * \text{false}$ and false . Formally, we define \sim to be the smallest equivalence relation on syntactic assertions, equating the following assertions:

- (S1) $\forall P, Q \in \text{Assn}. P * Q \sim Q * P,$
- (S2) $\forall P, Q, R \in \text{Assn}. P * (Q * R) \sim (P * Q) * R,$
- (S3) $\forall P \in \text{Assn}. P * \text{emp} \sim P,$
- (S4) $\text{false} * \text{false} \sim \text{false},$
- (S5) $\forall P, Q \in \text{Assn}. P \vee Q \sim Q \vee P,$
- (S6) $\forall P, Q, R \in \text{Assn}. P \vee (Q \vee R) \sim (P \vee Q) \vee R,$ and
- (S7) $\forall P \in \text{Assn}. (P \vee \text{false}) \sim (P \vee P) \sim P.$

where, for convenience, we have also included idempotence for disjunction. The model of heaps, $\text{Heap}_{\text{spec}}$, therefore is:

$$\begin{array}{l}
\text{Perm} \stackrel{\text{def}}{=} (0, 1] \quad \mathbb{B} \stackrel{\text{def}}{=} \{\text{true}, \text{false}\} \\
\mathcal{M} \stackrel{\text{def}}{=} \text{Val} \rightarrow \text{Assn} / \sim \\
\text{Heap}_{\text{spec}} \stackrel{\text{def}}{=} \text{Loc} \rightarrow \left(\begin{array}{l} \text{NA}[\mathbb{U} + (\text{Val} \times \text{Perm})] \\ + \text{Atom}[\mathcal{M} \times \mathcal{M} \times \mathbb{B} \times \mathbb{B}] \end{array} \right)
\end{array}$$

Each allocated location is either non-atomic or atomic. Non-atomic locations can either be uninitialized (represented by special symbol \mathbb{U}) or contain a value and a permission. Atomic locations contain two maps from values to syntactic assertions modulo \sim and two Boolean flags. The two maps represent the release and the acquire maps used to interpret the three assertion forms pertinent to RSL: $\text{Rel}(\ell, \mathcal{Q})$, $\text{Acq}(\ell, \mathcal{Q})$, and $\text{RMWAcq}(\ell, \mathcal{Q})$, with the first Boolean flag indicating whether the second map acts as a plain acquire

map or as an RMW-acquire map. The second Boolean flag records whether the location has been initialized or not.

Figure 15 defines the composition of two *logically disjoint* heaps, $h_1 \oplus h_2$. Note that two logically disjoint heaps can share some locations, provided that they store compatible information about them. For non-atomic locations, they should be initialized and have compatible permissions (i.e., whose sum does not exceed the full permission, 1). For atomic locations, the two heaps must contain compatible acquire maps, represented by the predicate $\mathbf{adef}(b_1, \mathcal{Q}_1, b_2, \mathcal{Q}_2)$. This predicate is somewhat complex because acquire maps represent plain acquire or RMW-acquire permissions depending on the relevant Boolean flag. The cases are:

- (Case $b_1 \wedge b_2$) we must have $\mathcal{Q}_1 = \mathcal{Q}_2$;
- (Case $b_1 \wedge \neg b_2$) we require that for all v , either $\mathcal{Q}_2(v) = \text{emp}$ or $\mathcal{Q}_1(v) = \mathcal{Q}_2(v) = \text{false}$;
- (Case $\neg b_1 \wedge b_2$) symmetrically to the previous case; and
- (Case $\neg b_1 \wedge \neg b_2$) no conditions.

Given these definitions, we can show that $(\text{Heap}_{\text{spec}}, \oplus, \emptyset)$ forms a separation algebra, which in turn means that it is a good model for separation logic assertions.

Lemma 1. $(\text{Heap}_{\text{spec}}, \oplus, \emptyset)$ forms a separation algebra. That is, \oplus is associative, commutative, and has \emptyset as its identity element.

In the proof of this lemma, property S4 is required to show associativity; replacing S4 with the more general property $\forall P \in \text{Assn}. P * \text{false} \sim \text{false}$ breaks associativity.

We remark that in contrast to most models for separation logic, our \oplus is not cancellative. For example, consider the heap $h_{\text{I}} = \{\ell \mapsto \text{Atom}[\text{False}, \text{Emp}, \text{false}, \text{true}]\}$. Clearly, $h_{\text{I}} \neq \emptyset$ and yet $h_{\text{I}} \oplus h_{\text{I}} = h_{\text{I}} = h_{\text{I}} \oplus \emptyset$. In practice, the lack of cancellativity does not affect reasoning about RSL assertions. It also does not mean that the heap model contains ‘junk’ information. Indeed, the heap h_{I} is used to model the assertion $\text{Init}(\ell)$, and we want $h_{\text{I}} \oplus h_{\text{I}} = h_{\text{I}}$ to validate INIT-SPLIT .

Definition 1 (Assertion Semantics).

Let $\llbracket - \rrbracket : \text{Assn} \rightarrow \mathbb{P}(\text{Heap}_{\text{spec}})$ be:

$$\begin{aligned}
\llbracket \text{false} \rrbracket &\stackrel{\text{def}}{=} \emptyset \\
\llbracket P \Rightarrow Q \rrbracket &\stackrel{\text{def}}{=} \{h \mid h \in \llbracket P \rrbracket \implies h \in \llbracket Q \rrbracket\} \\
\llbracket \forall x. P \rrbracket &\stackrel{\text{def}}{=} \{h \mid \forall v. h \in \llbracket P[v/x] \rrbracket\} \\
\llbracket \text{emp} \rrbracket &\stackrel{\text{def}}{=} \{\emptyset\} \\
\llbracket P * Q \rrbracket &\stackrel{\text{def}}{=} \{h_1 \oplus h_2 \mid h_1 \in \llbracket P \rrbracket \wedge h_2 \in \llbracket Q \rrbracket\} \\
\llbracket \text{Uninit}(\ell) \rrbracket &\stackrel{\text{def}}{=} \{\{\ell \mapsto \text{NA}[\mathbb{U}]\}\} \\
\llbracket \ell \xrightarrow{k} v \rrbracket &\stackrel{\text{def}}{=} \{\{\ell \mapsto \text{NA}[v, k]\}\} \\
\llbracket \text{Init}(\ell) \rrbracket &\stackrel{\text{def}}{=} \{\{\ell \mapsto \text{Atom}[\text{False}, \text{Emp}, \text{false}, \text{true}]\}\} \\
\llbracket \text{Rel}(\ell, \mathcal{Q}) \rrbracket &\stackrel{\text{def}}{=} \{\{\ell \mapsto \text{Atom}[\mathcal{Q}, \text{Emp}, \text{false}, \text{false}]\}\} \\
\llbracket \text{Acq}(\ell, \mathcal{Q}) \rrbracket &\stackrel{\text{def}}{=} \{\{\ell \mapsto \text{Atom}[\text{False}, \mathcal{Q}, \text{false}, \text{false}]\}\} \\
\llbracket \text{RMWAcq}(\ell, \mathcal{Q}) \rrbracket &\stackrel{\text{def}}{=} \{\{\ell \mapsto \text{Atom}[\text{False}, \mathcal{Q}, \text{true}, \text{false}]\}\}
\end{aligned}$$

where $\text{False} \stackrel{\text{def}}{=} \lambda v. \text{false}$ and $\text{Emp} \stackrel{\text{def}}{=} \lambda v. \text{emp}$.

Figure 16. Definition of the semantics of assertions.

Equipped with specification heaps, $\text{Heap}_{\text{spec}}$, we proceed to the semantics of assertions. These are given as a function $\llbracket - \rrbracket : \text{Assn} \rightarrow \mathbb{P}(\text{Heap}_{\text{spec}})$ in Figure 16.

A basic property of the assertion semantics, that justifies treating stored assertions up to \sim , is that \sim -related assertions have the same semantics.

Lemma 2. *If $P \sim Q$, then $\llbracket P \rrbracket = \llbracket Q \rrbracket$.*

Moreover, we can easily show that our model validates the logical entailments of Section 4.

Lemma 3. *The properties REL-SPLIT, ACQ-SPLIT, RMW-SPLIT, RMW-ACQ-SPLIT, and INIT-SPLIT hold universally.*

Finally, we say that an assertion is *precise* if and only if it uniquely determines a subheap where it holds. The definition is standard (O’Hearn 2007), but due of the lack of cancellativity of \oplus we require both the heaps satisfying the assertion to be equal ($h_1 = h'_1$) as well as their remainders ($h_2 = h'_2$). If \oplus were cancellative, then either of the equalities would suffice as it would imply the other.

Definition 2 (Precision). *An assertion is precise, denoted $\text{precise}(P)$, if and only if for all h_1, h'_1, h_2, h'_2 , if $h_1 \in \llbracket P \rrbracket$ and $h_2 \in \llbracket P \rrbracket$ and $h_1 \oplus h'_1 = h_2 \oplus h'_2 \neq \text{undef}$, then $h_1 = h'_1$ and $h_2 = h'_2$.*

7.2 Semantics of Hoare triples

We move on to the meaning of RSL triples, $\{P\} E \{y. Q\}$. To handle both models—the C11 standard one and the strengthened one of Section 6—we parametrize the definitions of the semantics of triples and all auxiliary definitions with respect to the model. For notational simplicity, however, we will present the definitions only for the strengthened model and we will note in text any differences for the standard C11 model.

Given an execution $\mathcal{X} = \langle \mathcal{A}, \text{lab}, \text{sb}, \text{rf}, \text{mo}, \text{sc} \rangle$, we define the helper functions: $\text{SBin}_{\mathcal{X}}(a)$, $\text{SBout}_{\mathcal{X}}(a)$, $\text{SWin}_{\mathcal{X}}(a)$,

and $\text{SWout}_{\mathcal{X}}(a)$, to get the set of sb/sw incoming/outgoing edges of an action $a \in \mathcal{A}$. Given also a set of actions, $V \subseteq \mathcal{A}$, we denote the set of its hb and rf predecessors as $\text{Pre}_{\mathcal{X}}(V)$.

$$\text{Pre}_{\mathcal{X}}(V) \stackrel{\text{def}}{=} \{a \mid \exists b \in V. \text{hb}(a, b) \vee a = \text{rf}(b)\}$$

This definition is very useful because we will generally be considering sets of actions V that are prefix-closed, namely $\text{Pre}_{\mathcal{X}}(V) \subseteq V$, and we will be growing such sets by adding one action at a time while maintaining prefix-closure. Doing so is always possible for consistent executions because of the StrongAcyclicHB axiom. In the standard C11 model, we have to resort to a stronger definition of $\text{Pre}_{\mathcal{X}}(V)$ that includes only the hb edges, not arbitrary rf edges as well.

$$\text{Pre}_{\mathcal{X}}^{\text{standard_C11}}(V) \stackrel{\text{def}}{=} \{a \mid \exists b \in V. \text{hb}(a, b)\}$$

To define the meaning of RSL triples, we will generally be annotating hb-edges with appropriate heaps. When doing so, however, it will be important to distinguish between happens-before edges that occur because of an sb-edge and those that occur because of an sw-edge. We therefore introduce the following definition that tags them accordingly.

Definition 3 (Tagged Happens Before). *Given an execution \mathcal{X} , let $\text{thb}_{\mathcal{X}}$ be a tagged union of $\text{sb}_{\mathcal{X}}$ and $\text{sw}_{\mathcal{X}}$, constructed as follows*

$$\begin{aligned}
\text{thb}_{\mathcal{X}} &\stackrel{\text{def}}{=} \{(\text{“sb”}, a, b) \mid (a, b) \in \text{sb}_{\mathcal{X}}\} \\
&\cup \{(\text{“sw”}, a, b) \mid (a, b) \in \text{sw}_{\mathcal{X}}\}
\end{aligned}$$

For a program expression, E , we denote $\mathcal{C}[\llbracket E \rrbracket]$ as the set of its *consistent contextual executions*. These executions are obtained by plugging in an execution of E in some arbitrary execution context, such that the whole execution is consistent, as follows.

$$\begin{aligned}
\mathcal{C}[\llbracket E \rrbracket] &\stackrel{\text{def}}{=} \{ \langle \text{res}, \mathcal{A}_{\text{ctx}}, \mathcal{A}_{\text{prg}}, \mathcal{X}, \text{fst}, \text{lst} \rangle \mid \exists \text{lab}_{\text{ctx}}, \text{lab}_{\text{prg}}. \\
&\quad \exists \text{sb}. \exists \text{sb}_{\text{prg}} = \text{sb} \cap (\mathcal{A}_{\text{prg}} \times \mathcal{A}_{\text{prg}}). \\
&\quad \mathcal{X} = \langle \mathcal{A}_{\text{ctx}} \uplus \mathcal{A}_{\text{prg}}, \text{lab}_{\text{ctx}} \cup \text{lab}_{\text{prg}}, \text{sb}, _, _, _ \rangle \\
&\quad \wedge \langle \text{res}, \mathcal{A}_{\text{prg}}, \text{lab}_{\text{prg}}, \text{sb}_{\text{prg}}, \text{fst}, \text{lst} \rangle \in \llbracket E \rrbracket \\
&\quad \wedge (\exists ! a. \text{sb}(a, \text{fst})) \wedge (\exists ! b. \text{sb}(\text{lst}, b)) \\
&\quad \wedge \text{Consistent}(\mathcal{X}) \}
\end{aligned}$$

In the definition of $\mathcal{C}[\llbracket E \rrbracket]$, we require that (1) the part of the execution corresponding to the expression matches its semantics, (2) *fst* has a unique sb-predecessor, (3) *lst* has a unique sb-successor, and (4) the entire execution is consistent. The requirements about the unique predecessor of *fst* and the unique successor of *lst* will be used for selecting unique edges responsible for carrying the expression’s precondition and postcondition.

To define the meaning of RSL triples, we will annotate the thb-edges of consistent contextual executions with heaps. We will call such functions, $\text{hmap} : \text{thb}_{\mathcal{X}} \rightarrow \text{Heap}_{\text{spec}}$,

Definition 4 (Local annotation validity). *Given an execution, $\mathcal{X} = \langle \mathcal{A}, \text{lab}, \text{sb}, \text{rf}, \text{mo}, \text{sc} \rangle$, a heap map, $hmap : \text{thb}_{\mathcal{X}} \rightarrow \text{Heap}_{\text{spec}}$, and a set of actions $V \subseteq \mathcal{A}$, the predicate $\text{Valid}(hmap, V)$ holds if and only if for all actions $a \in V$, there exist $\ell, v, \mathcal{Q}, \mathcal{Q}', \mathcal{Q}'', Z, h_1, h'_1, h_2, h_F, h_{\text{sink}}$ such that*

$$\begin{array}{l}
\left(\begin{array}{l}
\text{lab}(a) = \text{skip} \\
\wedge hmap(\text{SBin}_{\mathcal{X}}(a)) = hmap(\text{SBout}_{\mathcal{X}}(a)) \oplus h_{\text{sink}}
\end{array} \right) \\
\vee \left(\begin{array}{l}
\text{lab}(a) = W_Z(\ell, v) \wedge Z \in \{\text{rlx}, \text{rel}, \text{sc}\} \\
\wedge h_1 = \{\ell \mapsto \text{Atom}[\mathcal{Q}, \text{Emp}, \text{false}, _]\} \\
\wedge h'_1 = \{\ell \mapsto \text{Atom}[\mathcal{Q}, \text{Emp}, \text{false}, \text{true}]\} \\
\wedge h_2 = hmap(\text{SWout}_{\mathcal{X}}(a)) \oplus h_{\text{sink}} \wedge h_2 \in \llbracket \mathcal{Q}(v) \rrbracket \\
\wedge hmap(\text{SBin}_{\mathcal{X}}(a)) = h_1 \oplus h_2 \oplus h_F \\
\wedge hmap(\text{SBout}_{\mathcal{X}}(a)) = h'_1 \oplus h_F \\
\wedge (Z = \text{rlx} \implies \mathcal{Q}(v) = \text{emp})
\end{array} \right) \\
\vee \left(\begin{array}{l}
\text{lab}(a) = \text{RMW}_Z(\ell, v, v') \wedge Z \neq \text{na} \\
\wedge hmap(\text{SBin}_{\mathcal{X}}(a))(\ell) = \text{Atom}[_, \mathcal{Q}, \text{true}, \text{true}] \\
\wedge hmap(\text{SBin}_{\mathcal{X}}(a)) \oplus hmap(\text{SWin}_{\mathcal{X}}(a)) \\
= \{\ell \mapsto \text{Atom}[\mathcal{Q}', \text{Emp}, \text{false}, \text{false}]\} \oplus \\
hmap(\text{SBout}_{\mathcal{X}}(a)) \oplus hmap(\text{SWout}_{\mathcal{X}}(a)) \oplus h_{\text{sink}} \\
\wedge hmap(\text{SWin}_{\mathcal{X}}(a)) \in \llbracket \mathcal{Q}(v) \rrbracket \\
\wedge (hmap(\text{SWout}_{\mathcal{X}}(a)) \oplus h_{\text{sink}}) \in \llbracket \mathcal{Q}'(v') \rrbracket \\
\wedge (Z \in \{\text{rlx}, \text{rel}\} \implies \mathcal{Q}(v) = \text{emp}) \\
\wedge (Z \in \{\text{rlx}, \text{acq}\} \implies \mathcal{Q}'(v') = \text{emp})
\end{array} \right) \\
\vee \left(\begin{array}{l}
\text{lab}(a) = A(\ell) \\
\wedge hmap(\text{SBout}_{\mathcal{X}}(a)) = hmap(\text{SBin}_{\mathcal{X}}(a)) \oplus \\
\{\ell \mapsto \text{Atom}[\mathcal{Q}, \mathcal{Q}, _, \text{false}]\}
\end{array} \right) \\
\vee \left(\begin{array}{l}
\text{lab}(a) = A(\ell) \\
\wedge hmap(\text{SBout}_{\mathcal{X}}(a)) = hmap(\text{SBin}_{\mathcal{X}}(a)) \oplus \{\ell \mapsto \text{NA}[\mathbb{U}]\}
\end{array} \right) \\
\vee \left(\begin{array}{l}
\text{lab}(a) = W_{\text{na}}(\ell, v) \\
\wedge hmap(\text{SBin}_{\mathcal{X}}(a))(\ell) \in \{\text{NA}[\mathbb{U}], \text{NA}[_, 1]\} \\
\wedge hmap(\text{SBout}_{\mathcal{X}}(a)) = hmap(\text{SBin}_{\mathcal{X}}(a))[\ell \mapsto \text{NA}[v, 1]]
\end{array} \right) \\
\vee \left(\begin{array}{l}
\text{lab}(a) = R_{\text{na}}(\ell, v) \\
\wedge hmap(\text{SBin}_{\mathcal{X}}(a))(\ell) = \text{NA}[v, _] \\
\wedge hmap(\text{SBin}_{\mathcal{X}}(a)) = hmap(\text{SBout}_{\mathcal{X}}(a))
\end{array} \right) \\
\vee \left(\begin{array}{l}
\text{lab}(a) \in \{R_{\text{rlx}}(\ell, v), R_{\text{acq}}(\ell, v), R_{\text{sc}}(\ell, v)\} \\
\wedge hmap(\text{SWin}_{\mathcal{X}}(a)) \in \llbracket \mathcal{Q}(v) \rrbracket \wedge \text{precise}(\mathcal{Q}(v)) \\
\wedge hmap(\text{SBin}_{\mathcal{X}}(a)) = \{\ell \mapsto \text{Atom}[\text{False}, \mathcal{Q}, \text{false}, \text{true}]\} \oplus h_F \\
\wedge hmap(\text{SBout}_{\mathcal{X}}(a)) = hmap(\text{SWin}_{\mathcal{X}}(a)) \oplus h_F \oplus \\
\{\ell \mapsto \text{Atom}[\text{False}, \mathcal{Q}[v := \text{emp}], \text{false}, \text{true}]\}
\end{array} \right)
\end{array}$$

Definition 5 (Configuration safety). *Given sets of actions, \mathcal{A}_{ctx} and \mathcal{A}_{prg} , an execution $\mathcal{X} = \langle \mathcal{A}_{\text{prg}} \uplus \mathcal{A}_{\text{ctx}}, \text{lab}, \text{sb}, \text{rf}, \text{sc} \rangle$, a natural number, $n \in \mathbb{N}$, a set of actions, V , a heap map, $hmap : \text{Resp}_{\mathcal{X}}(V) \rightarrow \text{Heap}_{\text{spec}}$, a distinguished final action, $lst \in \mathcal{A}_{\text{prg}}$, and a set of heaps, Q , we define $\text{safe}_{\mathcal{X}}^n(V, hmap, \mathcal{A}_{\text{ctx}}, \mathcal{A}_{\text{prg}}, lst, Q)$ by structural recursion on n as follows:*

$\text{safe}_{\mathcal{X}}^0(V, hmap, \mathcal{A}_{\text{ctx}}, \mathcal{A}_{\text{prg}}, lst, Q)$ holds always.

$\text{safe}_{\mathcal{X}}^{n+1}(V, hmap, \mathcal{A}_{\text{ctx}}, \mathcal{A}_{\text{prg}}, lst, Q)$ holds if and only if the following conditions all hold:

- If $lst \in V$, then $hmap(\text{SBout}_{\mathcal{X}}(lst)) \in Q$; and
- For all $a \in \mathcal{A}_{\text{prg}} \setminus V$ such that $\text{Pre}_{\mathcal{X}}(\{a\}) \subseteq V$, there exists $hmap' : \text{Resp}_{\mathcal{X}}(\{a\}) \rightarrow \text{Heap}_{\text{spec}}$ such that $\text{Valid}(hmap \cup hmap', V \cup \{a\})$ and $\text{safe}_{\mathcal{X}}^n(V \cup \{a\}, hmap \cup hmap', \mathcal{A}_{\text{ctx}}, \mathcal{A}_{\text{prg}}, lst, Q)$; and
- For all $a \in \mathcal{A}_{\text{ctx}} \setminus V$ such that $\text{Pre}_{\mathcal{X}}(\{a\}) \subseteq V$, and all $hmap' : \text{Resp}_{\mathcal{X}}(\{a\}) \rightarrow \text{Heap}_{\text{spec}}$, if $\text{Valid}(hmap \cup hmap', V \cup \{a\})$, then $\text{safe}_{\mathcal{X}}^n(V \cup \{a\}, hmap \cup hmap', \mathcal{A}_{\text{ctx}}, \mathcal{A}_{\text{prg}}, lst, Q)$.

Figure 17. Definitions of annotation validity and configuration safety.

heap annotations or heap maps. For each thb -edge, it is important to decide who is responsible for choosing a heap to annotate that edge with: is it the program itself or is it its environment? Therefore, given a set of program actions $A \subseteq \mathcal{A}$, we define the set, $\text{Resp}_{\mathcal{X}}(A)$ of edges whose annotation is the responsibility of the program.

$$\text{Resp}_{\mathcal{X}}(A) \stackrel{\text{def}}{=} \bigcup_{a \in A} (\text{SBout}_{\mathcal{X}}(a) \cup \text{SWin}_{\mathcal{X}}(a))$$

This definition deserves some explanation.

First, as expected, the program is responsible for correctly annotating its outgoing sequenced-before edges. Conversely, it can assume that incoming sb -edges are correctly annotated. This part is consistent with the usual semantics of Hoare triples: the program may assume the precondition holds when starting its execution, and must establish the postcondition when returning.

What is perhaps a bit unusual is that the program is also responsible for the annotations on the *incoming* synchro-

nization edges, and not the outgoing ones. This is because when an acquire read synchronizes with a release write, it is the reader that ‘knows’ how much state ownership is to be transferred along the sw -edge. The writer simply knows how much total ownership is to be transferred away from itself, but not how this is to be distributed to the various readers that synchronize with the write.

In a slight abuse of notation, given a heap annotation, $hmap$, and a set of context edges, $S \subseteq \text{thb}_{\mathcal{X}}$, we will let $hmap(S) \stackrel{\text{def}}{=} \bigoplus_{x \in S \cap \text{dom}(hmap)} hmap(x)$.

Annotation Validity and Configuration Safety Figure 17 contains two important auxiliary definitions. First, we have annotation validity (Definition 4). A heap map, $hmap$ is *valid* up to a set of actions V , if and only if for every action $a \in V$, the annotation is locally valid around that action: basically the sum of the annotated heaps on the incoming edges should equal the sum of the annotated heaps on the outgoing edges, modulo the effect of action a .

Second, we have configuration safety (Definition 5), defined in the style of Vafeiadis (2011). Here, a configuration is a set of *visited* actions, $V \subseteq (\mathcal{A}_{\text{ctx}} \cup \mathcal{A}_{\text{prg}})$, and heap annotation, $hmap$, annotating precisely the thb-edges for which V is responsible. $\text{safe}_{\mathcal{X}}^n(V, hmap, \dots)$ asserts that such a configuration is safe for at least n further actions. Unless $n = 0$, a safe configuration must:

- Annotate the (unique) sb-outgoing edge from the command with a heap satisfying the postcondition in case the last action of the command is in V ;
- For any “ready-to-execute” action a of the command, it must be possible to extend the heap map so that it is safe also up to a for $n - 1$ actions; and
- For any “ready-to-execute” action of the context, any valid extension of the heap map should be safe for $n - 1$ actions.

The informal notion of action a being “ready-to-execute” is captured by the constraint that a has not yet been visited whereas all its predecessors have: $a \notin V \wedge \text{Pre}(\{a\}) \subseteq V$.

With these auxiliary definitions, we define the meaning of RSL triples as follows:

Definition 6 (Meaning of RSL triples).

*The Hoare triple, $\{P\} E \{y. Q\}$, holds if and only if for all $\langle res, \mathcal{A}_{\text{ctx}}, \mathcal{A}_{\text{prg}}, \mathcal{X}, fst, lst \rangle \in \mathcal{C}[E]$, for all $V \subseteq \mathcal{A}_{\text{ctx}}$ such that $\text{Pre}_{\mathcal{X}}(V) \subseteq V$, for all $hmap \in \text{Resp}_{\mathcal{X}}(V) \rightarrow \text{Heap}_{\text{spec}}$, for all $R \in \text{Assn}$, if $hmap(\text{SBin}_{\mathcal{X}}(fst)) \in \llbracket P * R \rrbracket$ and $\text{Valid}(hmap, V)$, then for all $n \in \mathbb{N}$, $\text{safe}_{\mathcal{X}}^n(V, hmap, \mathcal{A}_{\text{ctx}}, \mathcal{A}_{\text{prg}}, lst, Post)$, where $Post = \begin{cases} \llbracket Q[v/y] * R \rrbracket & \text{if } res = v \\ \text{Heap}_{\text{spec}} & \text{if } res = \perp. \end{cases}$*

The definition says that for any consistent contextual execution of E , all valid configurations annotating only the context edges and satisfying the precondition on the (unique) incoming sb-edge to the program, are safe for any number of steps. As is common in the definitions of the meaning of separation logic triples, the definition bakes in the frame rule—that is, it quantifies over all assertions R and star-conjoins R to the precondition and the postcondition.

7.3 Memory Safety and Race Freedom

The soundness proof of RSL consists of two parts. First, we have to show that every proof rule of §4 is a valid entailment according to the semantics of Hoare triples in Definition 6. Second, we have to show that RSL triples denote something useful for program executions, for example that they do not contain any data races nor any dangling reads.

We start with the second task as it is somewhat simpler. More specifically, we shall show that any consistent execution of a verified program under the true precondition is (a) memory safe, (b) has no uninitialized reads, (c) has no data races, and (d) if the program terminates, its postcondition is satisfiable. By memory safety, we mean that allocation of a

location must *happen before* any action reading or writing that location.

Definition 7. *An execution is memory safe if and only if $\forall a \in \mathcal{A}. \text{isaccess}_{\ell}(a) \implies \exists b. \text{hb}(b, a) \wedge \text{lab}(b) = A(\ell)$.*

Given a validly annotated execution by the heap map $hmap$, observe the following: (1) any action, a , accessing the location ℓ must have $\ell \in \text{dom}(hmap(\text{SBin}(a)))$; and (2) whenever a location is in the domain of the annotation of an edge leading to some action b , (i.e., when $\ell \in \text{dom}(hmap(_, _, b))$), then there must be an hb-earlier allocation action for that location. Putting these two together, we get memory safety for validly annotated executions.

Absence of reads from uninitialized locations follows by a similar argument. First, we say that a read action, a , reads from an uninitialized location if $\text{rf}(a) = \perp$, which from (ConsistentRFdom) means that there must be no previous write to that location. We can, however, observe that the annotation validity for read actions, a , requires that $hmap(\text{SBin}(a))(\ell) = \text{Atom}[_, _, _, \text{true}]$ (for atomic locations) or $hmap(\text{SBin}(a))(\ell) = \text{NA}[v, _]$ (for non-atomic locations). But, in order to get one of these heaps in a valid annotation, it must be the case that there was an hb-earlier write to the same location.

Proving race-freedom is slightly more involved. First, let us formalize exactly what race-freedom is. We say that two actions are *conflicting* if both access the same location, at least one of them is a write, and at least one of the accesses is non-atomic (i.e., atomic accesses do not conflict with one another). An execution is race-free if all conflicting actions are ordered by hb.

Definition 8. *Two actions $a \neq b$ are conflicting if there exists a location ℓ such that $\text{isaccess}_{\ell}(a)$ and $\text{isaccess}_{\ell}(b)$ and $\text{iswrite}(a) \vee \text{iswrite}(b)$, and $\text{mode}(a) = \text{na} \vee \text{mode}(b) = \text{na}$.*

Definition 9. *An execution is race-free if and only if for all conflicting actions $a, b \in \mathcal{A}$, we have $\text{hb}(a, b) \vee \text{hb}(b, a)$.*

To prove race-freedom, we need the notion of a set of transitions, \mathcal{T} , being pairwise independent. We say that \mathcal{T} is pairwise independent, if there exists no pair of transitions in \mathcal{T} such that one happens before the other.

Definition 10 (Independent Edges). *In an execution, \mathcal{X} , a set of transitions $\mathcal{T} \subseteq \text{thb}_{\mathcal{X}}$ is pairwise independent, denoted $\text{PairIndep}(\mathcal{T})$, if and only if for all $(_, a, a')$, $(_, b, b')$ in \mathcal{T} , we have $\neg \text{hb}_{\mathcal{X}}(a', b)$.*

The crux of the race freedom proof is the following independent heap compatibility lemma, which states that in every validly annotated execution, the heaps annotated at independent edges are \oplus -compatible.

Lemma 4 (Independent Heap Compatibility). *For every consistent execution, \mathcal{X} , heap map, $hmap : \text{Resp}_{\mathcal{X}}(\mathcal{A}_{\mathcal{X}}) \rightarrow \text{Heap}_{\text{spec}}$, and pairwise independent set of transitions, \mathcal{T} , if $\text{Valid}(hmap, \mathcal{A}_{\mathcal{X}})$ holds, then $\bigoplus_{x \in \mathcal{T}} hmap(x)$ is defined.*

To prove this lemma, we need the notion of the depth of a set of actions, which we take to be the number of its elements and its predecessors.

Definition 11 (Action Depth). *Given an execution, \mathcal{X} , the depth of a set of actions, $A \subseteq \mathcal{A}_{\mathcal{X}}$, which we denote as $\mathbf{D}_{\mathcal{X}}(A)$, is the number of actions in the set or that have happened before it, $|\bigcup_{n \geq 0} \underbrace{\text{Pre}_{\mathcal{X}}(\cdots \text{Pre}_{\mathcal{X}}(A) \cdots)}_n|$.*

The depth of actions satisfies this important property:

Lemma 5. *If $\text{hb}_{\mathcal{X}}(a, b)$, then $\mathbf{D}_{\mathcal{X}}(\{a\}) < \mathbf{D}_{\mathcal{X}}(\{b\})$.*

Lemma 4 is then proved by induction using the metric $\mathbf{D}_{\mathcal{X}}(\{a \mid (_, a, _) \in \mathcal{T}\})$, followed by case analysis on the action in \mathcal{T} with largest $\mathbf{D}_{\mathcal{X}}(\{-\})$ value.

Putting everything together, our main soundness theorem is stated in terms of complete consistent executions:

$$\begin{aligned} \mathcal{CC}[E] \stackrel{\text{def}}{=} \{ \langle \text{res}, \mathcal{X} \rangle \mid \exists a, b. \\ a \neq b \wedge \text{lab}_{\mathcal{X}}(a) = \text{lab}_{\mathcal{X}}(b) = \text{skip} \\ \wedge \langle \text{res}, \{a, b\}, _, \mathcal{X}, _, _ \rangle \in \mathcal{C}[E] \} \end{aligned}$$

where the program expression is put inside the trivial context providing it with an incoming sb-edge from a skip action and an outgoing sb-edge to a skip action. Here it is:

Theorem 1 (Adequacy). *Let $\{\text{true}\} E \{y. Q\}$. For every execution $\langle \text{res}, \mathcal{X} \rangle \in \mathcal{CC}[E]$, \mathcal{X} is memory safe, has no reads from uninitialized locations and no races. Moreover, if the execution is terminating, then Q holds of the result.*

7.4 Soundness of the Proof Rules

We move on to the proofs of soundness of the individual rules. For each rule, we have to prove that it is a valid entailment given the meaning of RSL triples (Definition 6). With the exception of R-ACQ and CAS*, these proofs are relatively straightforward because the conditions imposed by local validity are almost directly enforced by the proof rules.

The proofs of R-ACQ and CAS* are more complex because we also have to annotate the incoming sw-edges correctly and show that the annotation is valid not only for the program action under consideration, but also for the context actions at the other end—that is, for the write or RMW action with which the read or CAS synchronizes.

We start with the R-ACQ rule. Consider a consistent contextual execution where $\mathcal{A}_{\text{prg}} = \{a\}$ and $\text{lab}(a) = \text{R}_{\text{acq}}(\ell, v)$. We proceed with a case split. If $\mathcal{Q}(v) = \text{emp}$, we can simply annotate any incoming sw-edges with the empty heap and set $\text{hmap}'(\text{SBout}_{\mathcal{X}}(a)) = \text{hmap}(\text{SBin}_{\mathcal{X}}(a))$, which trivially preserves validity. When, however, $\mathcal{Q}(v) \neq \text{emp}$, the situation is much more difficult, because it is not immediately obvious that there is an incoming sw-edge that can be annotated in a way that satisfies the local validity conditions of both the acquire-read and the release-write (or RMW) at the other end. For this case, our proof works as follows.


First, as the precondition includes $\text{Init}(\ell)$, we know that there exists a write to ℓ that happens before the acquire read.

Lemma 6 (Init). *If $\text{Valid}(V, \text{hmap})$ and $\text{Pre}_{\mathcal{X}}(V) \subseteq V$ and $\text{hmap}(_, _, a)(\ell) = \text{Atom}[_, _, _, \text{true}]$, then there exists $c \in V$ such that $\text{lab}_{\mathcal{X}}(c) = \text{W}_{\ell}(_, _)$ and $(c, a) \in \text{hb}_{\mathcal{X}}$.*

Therefore, from the consistency axiom ConsistentRF, we get that $\exists w. \text{rf}(a) = w$. As $\text{Pre}_{\mathcal{X}}(\{a\}) \subseteq V$, we also know $w \in V$.

Next, we will show that w must be a plain atomic write that synchronizes with a . To see why this holds, observe that $\ell \in \text{dom}(\text{hmap}(\text{SBin}_{\mathcal{X}}(w)))$ holds as hmap is locally valid at $w \in V$. Now, informally, we can trace back through the $\text{thb}_{\mathcal{X}}$ edges to the point where for some node $c \in V$ such that $\text{hb}_{\mathcal{X}}(c, w)$, we have $\ell \notin \text{dom}(\text{hmap}(\text{SBin}_{\mathcal{X}}(c)))$ and yet $\ell \in \text{dom}(\text{hmap}(\text{SBout}_{\mathcal{X}}(c)))$. Since hmap is locally valid, the only way for this to happen is if $\text{lab}_{\mathcal{X}}(c) = \text{A}(\ell)$. Similarly, we can follow $\text{thb}_{\mathcal{X}}$ edges backwards from a and find a node $d \in V$ such that $\text{hb}_{\mathcal{X}}(d, a)$ and $\ell \notin \text{dom}(\text{hmap}(\text{SBin}_{\mathcal{X}}(d)))$ and $\ell \in \text{dom}(\text{hmap}(\text{SBout}_{\mathcal{X}}(d)))$. Again, since hmap is locally valid, $\text{lab}_{\mathcal{X}}(d) = \text{A}(\ell)$, and so from the consistency axiom ConsistentAlloc, we obtain that $c = d$. When tracing back from a , at each step we can show that there exist \mathcal{Q}' and b such that $\text{hmap}(t_{n+1})(\ell) = \text{Atom}[_, \mathcal{Q}', b, _]$ and either $b = \text{false} \wedge \mathcal{Q}'(v) \neq \text{emp}$ or $\mathcal{Q}'(v) = \text{false}$. So, in total, we get $\text{hmap}(\text{SBout}_{\mathcal{X}}(c))(\ell) = \text{Atom}[\mathcal{Q}', \mathcal{Q}', b, _]$ and either $b = \text{false} \wedge \mathcal{Q}'(v) \neq \text{emp}$ or $\mathcal{Q}'(v) = \text{false}$. Similarly, when tracing back from b , at each step we can show that whenever $\text{hmap}(t_{n+1})(\ell) = \text{Atom}[\mathcal{Q}', _, b, _]$, then there exist \mathcal{Q}'' and b' such that $\text{hmap}(t_n)(\ell) = \text{Atom}[\mathcal{Q}'', _, b', _]$ and $\mathcal{Q}''(v) \Rightarrow \mathcal{Q}(v)$ and $b' \Rightarrow b$. So, in total we get that there exist \mathcal{Q}'' and b' such that $\text{hmap}(\text{SBin}_{\mathcal{X}}(w))(\ell) = \text{Atom}[\mathcal{Q}'', _, b', _]$, and $\mathcal{Q}''(v) \Rightarrow \mathcal{Q}'(v)$ and $b' \Rightarrow b$. Since hmap is locally valid at w , $\exists h \in \llbracket \mathcal{Q}''(v) \rrbracket$; thus, $b' = b = \text{false}$ and $\mathcal{Q}''(v) \neq \text{emp}$, which means that w is a write that synchronizes with a .

We have the following picture: w synchronizes with a , but possibly also with some other reads $r_1, \dots, r_n \in V$ and perhaps even some reads not in V .



From the local validity of hmap at $w \in V$, we know that $(h_1 \oplus \cdots \oplus h_n \oplus h_{\text{sink}}) \in \llbracket \mathcal{Q}''(v) \rrbracket$, where each h_i is the heap annotated on the sw-edge from w to r_i . What remains to be shown is that we can split h_{sink} further; that is, we can find h', h'_{sink} such that $h_{\text{sink}} = h' \oplus h'_{\text{sink}}$ and $h' \in \llbracket \mathcal{Q}(v) \rrbracket$. Then, we annotate the (“sw”, w, a) edge with h' and $\text{SBout}_{\mathcal{X}}(a)$ with $h' \oplus \text{SBin}_{\mathcal{X}}(a)$, thereby ensuring local validity at both a and w . To find such a split, we rely on the following lemma.

Lemma 7 (Well-formedness). *Given a consistent execution \mathcal{X} , a prefix-closed set of actions, $V \subseteq \mathcal{A}_{\mathcal{X}}$ with $\text{Pre}_{\mathcal{X}}(V) \subseteq V$, a heap map, $\text{hmap} \in \text{Resp}_{\mathcal{X}}(V) \rightarrow \text{Heap}_{\text{spec}}$, that is locally valid with respect to V , $\text{Valid}(\text{hmap}, V)$, a pairwise*

independent set of transitions \mathcal{T} , such that $\{a \mid (_, a, _) \in \mathcal{T}\} \subseteq V$ and $\text{hmap}(\mathcal{T})(\ell) = \text{Atom}[_, \mathcal{Q}, _, _]$, an action, w , such that $\text{lab}(w) = \text{W}(\ell, v)$ and $\text{hmap}(\text{SBin}_{\mathcal{X}}(w))(\ell) = \text{Atom}[\mathcal{Q}'', _, _, _]$, and a partial map $\mathcal{R} : V \rightarrow \text{Assn}$, such that for all $a \in \text{dom}(\mathcal{R})$, a is a read action that synchronizes with w and acquires ownership of $\mathcal{R}(a)$, if moreover, $\{a \mid \exists(_, _, b) \in \mathcal{T} \wedge (a = b \vee \text{hb}(a, b))\} \cap \text{dom}(\mathcal{R}) = \emptyset$, then, $\mathcal{Q}''(v) \Rightarrow \mathcal{Q}(v) * \bigotimes_{r \in \text{dom}(\mathcal{R})} \mathcal{R}(r) * \text{true}$.

The proof of this lemma is rather technical and can be found in the Coq formalization. At a high level, however, it is similar to the proofs already described, using the depth metric to trace back the $\mathcal{T} \cup \{(\text{“sw”}, w, r) \mid r \in \text{dom}(\mathcal{R})\}$ edges until we reach c , the action that allocated ℓ .

Applying this lemma, we get that $(h_1 \oplus \dots \oplus h_n \oplus h_{\text{sink}}) \in \llbracket \mathcal{Q}(v) * \mathcal{R}(r_1) * \dots * \mathcal{R}(r_n) * \text{true} \rrbracket$, and since for all i , we also know that $\text{precise}(\mathcal{R}(r_i))$ and $h_i \in \llbracket \mathcal{R}(r_i) \rrbracket$, we obtain $h_{\text{sink}} \in \llbracket \mathcal{Q}(v) * \text{true} \rrbracket$, as required.

The proof of CAS is actually much simpler because there cannot be any resource-acquiring reads that synchronize with the write/RMW whence the CAS reads from. Details of this proof can be found in the Coq formalization.

7.5 The Coq Formalization

Our Coq development covers the entire soundness proof outlined in this section and follows the \LaTeX presentation very closely. To avoid excessive proof duplication, the definitions of configuration safety and triple validity are parametrized with respect to the memory model; that is, either the standard model or the one with the StrongAcyclicHB condition.

One notable difference is that in Coq we represent finite sets of actions, \mathcal{A} , as lists, and domain-restricted functions as functions over the full domain. For example, instead of $\text{lab} \in \mathcal{A} \rightarrow \text{Act}$, in Coq we have $\text{lab} : \text{AName} \rightarrow \text{Act}$, and add a consistency axiom stating that $\forall x \notin \mathcal{A}. \text{lab}(x) = \text{skip}$. Similarly, we define $\text{hmap} : \text{thb}(\text{AName} \times \text{AName}, \text{AName} \times \text{AName}) \rightarrow \text{Heap}_{\text{spec}}$, and in the definition of configuration safety, instead of saying that there exists a hmap' such that the configuration with $\text{hmap} \uplus \text{hmap}'$ is safe, we say that there exists hmap'' such that $\forall e \in \text{Resp}_{\mathcal{X}}(V). \text{hmap}''(e) = \text{hmap}(e)$ holds and the configuration with hmap'' is safe.

Another difference is that in Coq the treatment of assertions up to \sim is achieved by defining a syntactic assertion normalization function, norm , with the property that $P \sim Q \iff \text{norm}(P) = \text{norm}(Q)$. Then, we represent Assn/\sim as $\{P \in \text{Assn} \mid \text{norm}(P) = P\}$.

Finally, following Nanevski et al. (2010), we represent heaps as the option type $\text{Heap}_{\text{spec}} \cup \{\perp\}$, with \perp representing undefined heaps. This removes the ‘definedness’ side-conditions from the statements of commutativity and associativity of heap composition. In effect, we move the definedness checks to the semantics of assertions, where we ensure that $\perp \notin \llbracket P \rrbracket$ for any assertion P .

The formal development excluding standard libraries consists of about 3000 lines of definitions and statements

of lemmas and theorems, 5500 lines of proof, 500 lines of comments, and took the first author about two months to complete. It is worth pointing out that the formal proof revealed that a bug that we had missed in our earlier paper proofs: namely, the requirement that the ownership transfer governed by the R-ACQ rule to be precise. While we do not think that this side condition is strictly necessary for soundness in the absence of the conjunction rule, the current proof style fundamentally requires it.

8. Related Work and Conclusion

This paper introduced relaxed separation logic, a moderate extension of concurrent separation logic (O’Hearn 2007) with special primitives for handling C11’s acquire and release atomic accesses.

8.1 Related Work

About the C11 Model The C11 concurrency model is part of the C and C++ 2011 standards (ISO/IEC 9899:2011; ISO/IEC 14882:2011), and has been formalized by Batty et al. (2011). In a subsequent paper, Batty et al. (2012) simplified the C11 model in the absence of consume reads. It is this simplified model that we used in this paper.

In a recent paper, Batty et al. (2013) considered the notion of library atomicity in the context of the C11 memory model. While this work is largely orthogonal to our defining a program logic about C11, we expect that combining the two approaches will be fruitful as program logics are often the means for proving atomicity, at least in the SC setting.

Logics for Other Weak Memory Models We are aware of only three lines of work that define program logics over a relaxed memory model, none of which handles the C11 memory model.

- Ferreira et al. (2010) proved the soundness of concurrent separation logic (CSL) over a class of relaxed memory models, all satisfying the DRF-guarantee. In hindsight, their result is not surprising as the soundness of CSL over SC (sequential consistency) ensures that CSL-verified programs do not contain any data races, and hence whether the soundness proof is done over SC or over the relaxed memory model is irrelevant.
- Ridge (2010) developed a rely-guarantee proof system over x86-TSO and used it to verify a x86-TSO version of Simpson’s four slot algorithm, with all the results mechanized in the HOL theorem prover.
- Wehrman and Berdine (2011) proposed a variant of separation logic for x86-TSO featuring primitive assertions for modelling the state of the TSO buffers and both temporal and spatial separating conjunctions.

Besides obviously handling different memory models and being quite different program logics, there is a fundamental difference between the current work and these earlier pa-

pers on program logics for relaxed memory models. In this work, we define the meaning of Hoare triples directly over an axiomatic partial order semantics for concurrent programs, whereas the earlier works used an operational or an operationally flavoured trace semantics, very much like the traditional soundness proofs over SC. As a result, our soundness proof is completely different from the soundness proofs of the aforementioned papers.

8.2 Possible Future Research Directions

Being the first program logic for C11 concurrency, there are numerous opportunities for extending RSL, for example to deal with more advanced features of the C11 memory model, such as consume reads and memory fences. Similarly, one can also try to adapt to C11 setting more advanced program logics, such as RGSep (Vafeiadis and Parkinson 2007), concurrent abstract predicates (Dinsdale-Young et al. 2010), or CaReSL (Turon et al. 2013).

Initialization of Atomicity For simplicity, RSL tags locations as atomic or non-atomic and permits atomic accesses only on atomic locations and non-atomic accesses only on non-atomic locations. As a result of this choice, initialization writes to atomic accesses in our model also have to be atomic, whereas the C11 standard also allows non-atomic initialization writes to atomic location. To enable the verification of such programs, we should somehow allow the following ‘conversion’ rule

$$\frac{\mathcal{Q}(v) = \text{emp}}{\{\ell \mapsto v\} 0 \{ \text{Rel}(\ell, \mathcal{Q}) * \text{Acq}(\ell, \mathcal{Q}) * \text{Init}(\ell) \}}$$

Automation Another important research direction would be to develop techniques and tools for automating RSL proofs by adapting some of the work that has been done in automating standard separation logic (Distefano et al. 2006; Calcagno et al. 2009; Dudka et al. 2011).

Acknowledgments

We would like to thank Lars Birkedal, Arthur Charguéraud, Mike Dodds, Alexey Gotsman, Matthew Parkinson, Aaron Turon, John Wickerson, and the anonymous OOPSLA 2013 reviewers for their very useful comments that improved the content of this paper. The research was supported by the EC FP7 FET Young Explorers scheme via the project ADVENT.

References

M. Batty, S. Owens, S. Sarkar, P. Sewell, and T. Weber. Mathematizing C++ concurrency. In *POPL 2011*, pages 55–66. ACM, 2011.

M. Batty, K. Memarian, S. Owens, S. Sarkar, and P. Sewell. Clarifying and compiling C/C++ concurrency: From C++11 to POWER. In *POPL 2012*, pages 509–520. ACM, 2012.

M. Batty, M. Dodds, and A. Gotsman. Library abstraction for C/C++ concurrency. In *POPL 2013*, pages 235–248. ACM, 2013.

L. Birkedal, K. Støvring, and J. Thamsborg. The category-theoretic solution of recursive metric-space equations. *Theoretical Computer Science*, 411(47):4102–4122, 2010.

J. Boyland. Checking interference with fractional permissions. In *SAS 2003*, volume 2694 of *LNCS*, pages 55–72. Springer, 2003.

C. Calcagno, D. Distefano, and V. Vafeiadis. Bi-abductive resource invariant synthesis. In *APLAS*, volume 5904 of *LNCS*, pages 259–274. Springer, 2009.

T. Dinsdale-Young, M. Dodds, P. Gardner, M. Parkinson, and V. Vafeiadis. Concurrent abstract predicates. In *ECOOP 2010*, volume 6183 of *LNCS*, pages 504–528. Springer, 2010.

D. Distefano, P. W. O’Hearn, and H. Yang. A local shape analysis based on separation logic. In *TACAS*, volume 3920 of *LNCS*, pages 287–302. Springer, 2006.

K. Dudka, P. Peringer, and T. Vojnar. Predator: A practical tool for checking manipulation of dynamic data structures using separation logic. In *CAV*, volume 6806 of *LNCS*, pages 372–378. Springer, 2011.

R. Ferreira, X. Feng, and Z. Shao. Parameterized memory models and concurrent separation logic. In *ESOP 2010*, volume 6012 of *LNCS*, pages 267–286. Springer, 2010.

C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen. The essence of compiling with continuations. In *PLDI 1993*, pages 237–247. ACM, 1993.

ISO/IEC 14882:2011. Programming language C++, 2011.

ISO/IEC 9899:2011. Programming language C, 2011.

P. E. McKenney and B. Garst. N1525: Memory-order rationale, 2011. Available at <http://www.open-std.org/jtc1/sc22/wg14/www/docs/n1525.htm>.

A. Nanevski, V. Vafeiadis, and J. Berdine. Structuring the verification of heap-manipulating programs. In *POPL*, pages 261–274. ACM, 2010.

P. O’Hearn. Resources, concurrency, and local reasoning. *Theoretical Computer Science*, 375(1):271–307, 2007.

T. Ridge. A rely-guarantee proof system for x86-TSO. In *VSTTE 2010*, volume 6217 of *LNCS*, pages 55–70. Springer, 2010.

S. Sarkar, K. Memarian, S. Owens, M. Batty, P. Sewell, L. Maranget, J. Alglave, and D. Williams. Synchronising C/C++ and POWER. In *PLDI 2012*, pages 311–322. ACM, 2012.

A. Turon, D. Dreyer, and L. Birkedal. Unifying refinement and Hoare-style reasoning in a logic for higher-order concurrency. In *ICFP 2013*. ACM, 2013.

V. Vafeiadis. Concurrent separation logic and operational semantics. In *MFPS 2011*, volume 276 of *ENTCS*, pages 335–351. Elsevier, 2011.

V. Vafeiadis and M. Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR 2007*, volume 4703 of *LNCS*, pages 256–271. Springer, 2007.

M. N. Wegman and F. K. Zadeck. Constant propagation with conditional branches. *ACM Trans. Program. Lang. Syst.*, 13(2): 181–210, Apr. 1991.

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