



SecRSL: Security Separation Logic for C11 Release-Acquire Concurrency

PENGBO YAN, University of Melbourne, Australia

TOBY MURRAY, University of Melbourne, Australia

We present Security Relaxed Separation Logic (SecRSL), a separation logic for proving information-flow security of C11 programs in the Release-Acquire fragment with relaxed accesses. SecRSL is the first security logic that (1) supports weak-memory reasoning about programs in a high-level language; (2) inherits separation logic's virtues of compositional, local reasoning about (3) expressive security policies like value-dependent classification.

SecRSL is also, to our knowledge, the first security logic developed over an axiomatic memory model. Thus we also present the first definitions of information-flow security for an axiomatic weak memory model, against which we prove SecRSL sound. SecRSL ensures that programs satisfy a constant-time security guarantee, while being free of undefined behaviour.

We apply SecRSL to implement and verify the functional correctness and constant-time security of a range of concurrency primitives, including a spinlock module, a mixed-sensitivity mutex, and multiple synchronous channel implementations. Empirical performance evaluations of the latter demonstrate SecRSL's power to support the development of secure and performant concurrent C programs.

CCS Concepts: • **Theory of computation** → **Separation logic**; • **Security and privacy** → **Logic and verification**.

Additional Key Words and Phrases: Information-flow Security, Separation Logic, Weak Memory Consistency, Axiomatic Semantics

ACM Reference Format:

Pengbo Yan and Toby Murray. 2021. SecRSL: Security Separation Logic for C11 Release-Acquire Concurrency. *Proc. ACM Program. Lang.* 5, OOPSLA, Article 99 (October 2021), 26 pages. <https://doi.org/10.1145/3485476>

1 INTRODUCTION

Logics for proving that concurrent programs do not leak sensitive information have received much recent study [Ernst and Murray 2019; Frumin et al. 2021; Karbyshev et al. 2018; Murray et al. 2018, 2016; Schoepe et al. 2020]. A common thread of much recent work has been adapting ideas from concurrent separation logic [O'Hearn 2004] to reason about secure information flow. Indeed, recent security logics purposefully closely resemble traditional concurrent separation logics [Ernst and Murray 2019; Frumin et al. 2021]. Besides providing a familiar, compositional and elegant setting in which to carry out security proofs, such logics have also proved amenable to automated verification via symbolic execution [Ernst and Murray 2019].

Security separation logics have so far been confined to reasoning over sequentially-consistent memory models. There exist a handful of information-flow logics and type systems for weak memory consistency models [Mantel et al. 2014; Smith et al. 2019; Vaughan and Millstein 2012].

Authors' addresses: Pengbo Yan, University of Melbourne, Australia, pengbo.yan@unimelb.edu.au; Toby Murray, University of Melbourne, Australia, toby.murray@unimelb.edu.au.



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

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

2475-1421/2021/10-ART99

<https://doi.org/10.1145/3485476>

However they lack the local reasoning abilities, and consequent scalability, afforded by separation logics. They also lack separation logic's support for reasoning about resource ownership transfer with invariants, which is crucial for expressive reasoning. Additionally, these logics have targeted machine-level weak memory models, and so are not readily applicable for reasoning about programs in higher level languages.

In some sense, these limitations are not surprising. Logics for information-flow security are almost exclusively proved sound against *operational* semantic models. Yet formal weak memory models are often specified *axiomatically*, in which program behaviours are represented as mathematical objects constrained by the axioms of the weak memory model.

To our knowledge there do not even exist prior definitions of information-flow security for such models, without which one cannot even state the soundness theorems for any proposed logic. Even if one had such properties, another necessary but lacking ingredient for proving the soundness of such logics is suitable compositional, inductive properties that state the semantic meaning of the logic's judgements. The shape of such properties is well understood for traditional (non-security) separation logics [Doko and Vafeiadis 2016, 2017; Vafeiadis and Narayan 2013], many of which share very similar definitions. Yet little is known about what they should look like for security separation logics.

In this paper, we introduce *Security Relaxed Separation Logic* (SecRSL). SecRSL is the first concurrent separation logic for reasoning about secure information flow in a relaxed memory model. It targets the Release-Acquire fragment of C11 with relaxed accesses (Section 3), and thus supports reasoning about high-level programs for the first time. It also inherits the local reasoning, expressiveness, and compositionality of traditional concurrent separation logics.

SecRSL's design (Section 4) combines ideas from two prior logics: (1) Ernst and Murray's Security Concurrent Separation Logic (SecCSL) [Ernst and Murray 2019], an information-flow security analogue of traditional concurrent separation logic for sequential consistent concurrency; and (2) Vafeiadis and Narayan's Relaxed Separation Logic (RSL) [Vafeiadis and Narayan 2013]. Thus SecRSL inherits SecCSL's support for proving expressive security policies like those involving *value-dependent classification* [Murray et al. 2016], in which the sensitivity of one variable can change in response to changes in the value held by another. It also inherits RSL's ability to reason about ownership transfer with invariants via C11's Release-Acquire atomics.

To state the soundness of SecRSL (Section 5), we present what is to our knowledge the first definitions of information-flow security for an axiomatic weak memory model (Section 5.1). We believe it should be readily applicable to similar memory models [Doko and Vafeiadis 2016, 2017] with little modification, while providing guidance on how to structure such definitions for less similar axiomatic models. Our security definition protects against passive attackers who can observe the contents of **low** (public) memory locations, as well as attackers who can observe the program's memory access pattern (i.e. can see *which* locations are being accessed by the program but not the *values* being written to them). Thus SecRSL also protects against attackers who can mount cache side-channel attacks and provides a form of *constant-time security* [Barthe et al. 2019].

To prove SecRSL sound, we developed a novel, compositional definition of relational validity (Section 5.3) that encodes the meaning of SecRSL judgements. As with the top-level security properties, we believe this definition is of independent interest beyond the confines of our specific memory model, which we inherit from RSL [Vafeiadis and Narayan 2013].

We demonstrate SecRSL's power by using it to implement and verify the functional correctness and constant-time security of a range of concurrency primitives (Section 6). Specifically, we demonstrate SecRSL's support for traditional relaxed memory reasoning by showing how Vafeiadis and Narayan's RSL proof for a spinlock module can be replayed in SecRSL. Thus we prove it not only functionally correct but also to satisfy SecRSL's constant-time security guarantee. We

$$MP_1(x) \stackrel{\text{def}}{=} \begin{array}{l} \text{let } b = \text{alloc}_{\text{high}}() \text{ in} \\ \text{let } a = \text{alloc}_{\text{low}}() \text{ in} \\ \text{let } c = \text{alloc}_{\text{low}}() \text{ in} \\ \quad [a]_{\text{rlx}} := 0; \\ [b]_{\text{na}} := x; \quad \parallel \quad \text{repeat } [a]_{\text{acq}} \text{ end} \\ [a]_{\text{rel}} := 1; \quad \parallel \quad \text{let } y = [b]_{\text{na}} \text{ in} \\ \quad \quad \quad [c]_{\text{na}} := y \end{array}$$

Fig. 1. A message-passing program

then extend that example to implement a *mixed-sensitivity mutex* for protecting access to data of varying sensitivity. Finally, we implement and verify multiple implementations of a *synchronous channel* abstraction for transmitting data of varying sensitivity. We benchmark their performance against an (unverified) sequentially-consistent implementation. In doing so, we find that SecRSL enables significant performance gains of up to $\sim 90\%$ depending on platform. Thus demonstrating SecRSL’s power to support the development of secure and performant concurrent C programming abstractions.

All formal results in this paper have been mechanised in the Coq theorem prover. The theories are available as supplementary material [Yan 2021]. Various technical details have been omitted from this paper, which can be found in the appendices in the extended version [Yan and Murray 2021].

2 OVERVIEW

C11 Release-Acquire atomics. The program MP_1 in Fig. 1 is a slight adaptation of the classic *message-passing* program, and makes use of C11’s Release-Acquire atomics. This program has access to some local variable x . It creates two threads and uses the location b to send a message containing x from the left thread to the right one, where the location a is used to synchronise the two threads. The message-passing protocol initialises a to hold 0. Once b holds the message to be transferred, a is updated to hold 1 to signal to the right thread that the message (stored in b) is now ready to be read. The right thread busywaits until a is non-zero (the **repeat** command executes its body until it returns a non-zero value). The right thread then loads the message from b before writing it to the location c .

In the C11 Release-Acquire fragment we consider, loads and stores to memory locations are annotated with various modes. These affect the potential observable reorderings that are possible under the memory model. The left thread’s store of 1 to location a carries the *release* mode (**rel**), while the right thread’s load of a (inside the busyloop) carries the *acquire* mode (**acq**). Together these modes guarantee that if the acquire load reads the value 1, then the prior store of x to location b must have also completed. Hence, the subsequent load from b must read the value x and so x will be written to location c .

In SecRSL (as in RSL before it) two kinds of locations are distinguished: *atomic* locations and *non-atomic* locations. Modes like acquire and release can be used only with atomic locations; non-atomic locations must be accessed using the non-atomic mode (**na**). When reasoning about the program in Fig. 1 therefore, location a is treated as an atomic location while b and c are considered non-atomic locations.

The axiomatic C11 memory model defines a program execution as a set of actions related by several binary relations that impose ordering constraints on those events. Two possible executions

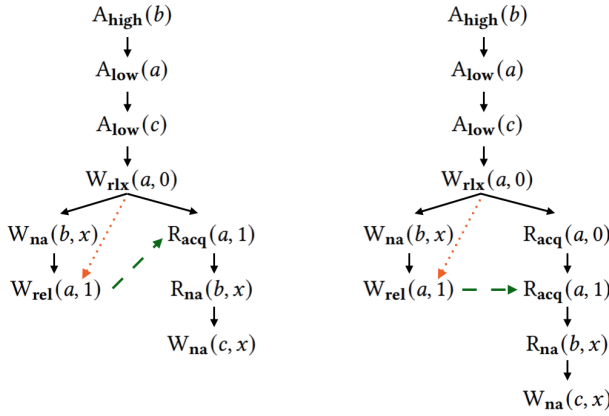


Fig. 2. Two executions of the $MP_1(x)$ program. Edges depict some key relations of the memory model (Section 3.2): *sb* (black, solid); *rf* (dashed, green); *mo* (dotted, orange). As with much prior work, we draw the arrows for the *rf* relation to depict data flow: for read and write actions r and w respectively, when $rf(r) = w$ we draw a (dashed, green) line from w to r .

of this program are depicted graphically in Fig. 2, where the nodes depict the actions and the edges depict various relations between them. The execution on the left in this figure arises when the loop executes only once; the one on the right executes the busyloop twice.

We formally define the Release-Acquire fragment of C11 and its memory model that SecRSL treats in Section 3, which we inherit from RSL [Vafeiadis and Narayan 2013].

Information-flow Security. What should information-flow security mean in this fragment of C11? As with SEC CSL [Ernst and Murray 2019], we assume the existence of a passive attacker who can observe certain parts of memory. Specifically, when they are allocated certain memory locations are deemed to be potentially attacker-observable. Such locations might hold data that will be publicly output or copied onto the network. We call these locations **low** and all other locations **high**.

Notice from the allocation statements in Fig. 1 that the atomic location a and the non-atomic location c are both **low**. Being conservative, atomic locations should be considered attacker-observable, because interactions with them (e.g. those with a in Fig. 1) affect thread scheduling and so are potentially observable to attackers who can measure time.

In this example, the location c is also considered attacker observable. As is standard in information-flow security, we assume the attacker knows the program’s source code. Hence, constants that appear in the program text are considered public. The program only writes public constants to a ; however it eventually writes x to location c . Hence, the program is information-flow secure only if x is public data.

We say that a program does not leak information (and so is information flow secure) if all executions of that program beginning from states that agree on the values of public (i.e. non-secret) data are indistinguishable to the attacker. Supposing the variable x holds public data initially, then the program MP_1 is secure if for all x, x' , if x and x' are equal, then $MP_1(x)$ is indistinguishable from $MP_1(x')$.

Due to the nondeterminism inherent in axiomatic semantic models, a single program can give rise to multiple executions: an implementation is free to choose any one of them while remaining compliant with the semantics. Indeed two possible behaviours of $MP_1(x)$ are depicted in Fig. 2. This nondeterminism creates potential challenges for secure information flow.

To understand why, suppose the initial state of the example was extended with an additional secret bit b and consider an implementation of the weak memory semantics that chooses to produce the left execution when $b = 1$ and to produce the right one otherwise. An attacker who can measure execution time would be able to distinguish these two executions since one performs an extra action.

Thus, as with prior work [Smith et al. 2019], we make the necessary assumption that the resolution of nondeterminism in the weak memory semantics does not itself reveal sensitive information. Since in our semantics nondeterminism arises from the selection of memory locations returned by the **alloc** expression, as well as from the concurrency semantics, we assume that neither the memory allocator nor the concurrency implementation (the compiler, operating system scheduler, hardware, etc.) reveal secrets to the attacker. Both assumptions are intuitively necessary in order to proceed with source level reasoning about information-flow security.

Our threat model also includes attackers who can observe not only the values written to **low** locations, but also the program's memory access pattern: such attackers can observe which locations the program is accessing even if they cannot see the values being written to them.

Thus as well as guaranteeing that indistinguishable values are written to **low** locations, SecRSL's top-level security property also requires that for two executions to be indistinguishable they must access the same locations in the same way at the same time. Additionally its rules prevent programs from branching on secrets. Thus it guarantees a form *constant-time security* [Barthe et al. 2019].

We formally define these guarantees via SecRSL's top-level security property in Section 5.1.

The Logic. Recall we said the example MP_1 is secure under the assumption that its argument is public if for all arguments x and x' that are equal, $MP_1(x)$ is indistinguishable from $MP_1(x')$. Thus SecRSL is a *relational* logic [Benton 2004] that in general relates the executions of two programs, and proves them indistinguishable.

SecRSL judgements have the general form

$$[HP], \{P\} E, E' \{(y, y')\}. Q, [(y, y'). HQ],$$

where P is the precondition, E and E' are the two programs we are proving indistinguishable, y and y' are the return values of E and E' , and Q is the postcondition. HP and HQ track the set of **high** locations before and after the execution.

Since SecRSL is relational, its separation logic assertions relate pairs of states [Ernst and Murray 2019]. The judgement above means that if E and E' are run from starting states related by P in which all locations in HP are unobservable to the attacker, then their executions will be indistinguishable to the attacker and, moreover, Q will relate the final states and HQ will track the final set of locations known to be unobservable to the attacker.

The security of MP_1 we can thus express in SecRSL as the judgement: $[\emptyset], \{\text{Low}(x, x')\} MP_1(x), MP_1(x') \{\text{true}\}, [\emptyset]$. SecRSL introduces the relational assertion $\text{Low}(e, e')$ to assert that expressions e and e' denote identical values. We note that HP and HQ *under-approximate* the set of **high** locations. This explains why the post-set in this judgement is empty.

We defer an explanation of the assertion forms of SecRSL and its rules to Section 4. However we note that it supports analogues of all of the reasoning principles of RSL. In particular, atomic locations like a in Fig. 1 carry (relational) invariants $Q(v)$ that can be used to reason about ownership transfer via Release-Acquire patterns. Being relational, these invariants can express security properties as well as functional ones.

Fig. 3 depicts a proof sketch for the security of MP_1 in SecRSL. We note that the invariant Q used for this proof encodes that when location a holds the value 1, then location b holds a public value. Hence this example demonstrates SecRSL's ability to support reasoning about value-dependent

$$\begin{array}{l}
Q(v) \stackrel{\text{def}}{=} (v = 0 \wedge \text{emp}) \vee (v = 1 \wedge ((b, b) \mapsto (x, x') \star \text{Low}(x, x'))); \quad Q' \stackrel{\text{def}}{=} Q[1 := \text{emp}] \\
\quad [\emptyset], \{\text{Low}(x, x')\} \\
\quad \text{let } b = \text{alloc}_{\text{high}}() \text{ in} \\
\quad [b], \{\text{Low}(x, x') \star \text{Uninit}(b, b)\} \\
\quad \text{let } a = \text{alloc}_{\text{low}}() \text{ in} \\
\quad [b], \{\text{Low}(x, x') \star \text{Uninit}(b, b) \star \text{Rel}(a, Q) \star \text{Acq}(a, Q)\} \\
\quad \text{let } c = \text{alloc}_{\text{low}}() \text{ in} \\
\quad [b], \{\text{Low}(x, x') \star \text{Uninit}(c, c) \star \text{Uninit}(b, b) \star \text{Rel}(a, Q) \star \text{Acq}(a, Q)\} \\
\quad [a]_{\text{rlx}} := 0; \\
\quad [b], \{\text{Low}(x, x') \star \text{Uninit}(c, c) \star \text{Uninit}(b, b) \star \text{Rel}(a, Q) \star \text{Acq}(a, Q) \star \text{Init}(a)\} \\
[b], \{\text{Low}(x, x') \star \text{Uninit}(b, b) \star \text{Rel}(a, Q)\} \quad [b], \{\text{Uninit}(c, c) \star \text{Acq}(a, Q) \star \text{Init}(a)\} \\
\quad [b]_{\text{na}} := x; \quad \text{repeat } [a]_{\text{acq}} \text{ end;} \\
[b], \{\text{Low}(x, x') \star (b, b) \mapsto (x, x') \star \text{Rel}(a, Q)\} \quad [b], \{\text{Low}(x, x') \star (b, b) \mapsto (x, x') \star \text{Uninit}(c, c) \\
\quad \star \text{Acq}(a, Q')\} \\
\quad [a]_{\text{rel}} := 1; \quad \text{let } y = [b]_{\text{na}} \text{ in} \\
[b], \{\text{Rel}(a, Q) \star \text{Init}(a)\} \quad [b], \{(b, b) \mapsto (y, y) \star \text{Uninit}(c, c) \star \text{Acq}(a, Q')\} \\
\quad [c]_{\text{na}} := y \\
\quad [b], \{(b, b) \mapsto (y, y) \star (c, c) \mapsto (y, y) \\
\quad \star \text{Acq}(a, Q')\} \\
\quad [\emptyset], \{\text{true}\}
\end{array}$$

Fig. 3. Secure message passing in SecRSL, proving $[\emptyset], \{\text{Low}(x, x')\} MP_1(x), MP_1(x') \{\text{true}\}, [\emptyset]$. To reduce clutter, we write proof sketches as if they operate over a single program. However, they should really be understood as operating over a pair of programs, where the second program is obtained by substituting each variable in the first program with its primed counterpart. This proof sketch uses notational shorthands defined on page 9.

classification policies, as well as the local, compositional reasoning with ownership transfer that makes concurrent separation logics so powerful.

3 LANGUAGE AND MEMORY MODEL

3.1 Programming Language

SecRSL is defined over the Release-Acquire fragment of C11 with relaxed accesses of RSL [Vafeiadis and Narayan 2013]. The only difference is that in SecRSL, each memory allocation expression $\text{alloc}_F()$ is annotated with a security level $F \in \{\text{high}, \text{low}\}$, which is used to define the top-level information-flow security property established by SecRSL. These annotations are *ghost information* and are ignored by the language semantics (see Section 3.2).

As with RSL [Vafeiadis and Narayan 2013], SecRSL programs are assumed to be in A-Normal form [Flanagan et al. 1993]. Thus an atomic expression, $e \in \text{AExp}$, is either a variable x or a value v (a memory location $\ell \in \text{Loc}$ or a number $n \in \mathbb{N}$). Program expressions, $E \in \text{Exp}$, include atomic expressions e , let-bindings, conditionals, loops, parallel composition, labelled memory allocations $\text{alloc}_F()$, loads $[e]_X$, stores $[e]_Y := e'$, and atomic compare-and-swap (CAS) instructions.

$$\begin{array}{ll}
v \in \text{Val} ::= \ell \mid n & \text{where } \ell \in \text{Loc}, n \in \mathbb{N} \\
e \in \text{AExp} ::= x \mid v & \text{where } x \in \text{Var}
\end{array}$$

$$\begin{aligned}
E \in \text{Exp} ::= & e \mid \text{let } x = E \text{ in } E' \mid \text{if } e \text{ then } E \text{ else } E' \\
& \mid \text{repeat } E \text{ end} \mid E_1 \parallel E_2 \mid \text{alloc}_F() \\
& \mid [e]_X \mid [e]_Y := e' \mid \text{CAS}_{Z,W}(e, e', e'') \\
\text{where } X \in & \{\text{sc, acq, rlx, na}\}, Y \in \{\text{sc, rel, rlx, na}\}, \\
Z \in & \{\text{sc, rel_acq, acq, rel, rlx}\}, \\
W \in & \{\text{sc, acq, rlx}\}, F \in \{\text{high, low}\}
\end{aligned}$$

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.

In the C11 Release-Acquire memory model, memory accesses are annotated by their mode, which affects the ordering guarantees they provide. Modes comprise: sequentially consistent (**sc**), acquire (**acq**), release (**rel**), combined release-acquire (**rel_acq**), relaxed (**rlx**), or non-atomic (**na**). Different kinds of memory accesses support different modes. For instance, reads cannot be releases, writes cannot be acquires, CASs cannot be non-atomic.

CAS is an atomic operation used heavily in lock-free concurrent algorithms. In high-level terms, it takes a location, ℓ , and two values, v' and v'' , as arguments. It atomically checks if the value in the memory location ℓ is v' . If it is then the CAS operation is said to succeed and it atomically stores v'' into memory location ℓ and returns v' (the old value). Otherwise the CAS is said to fail and it does not modify the memory and simply returns whatever non- v' value was stored in location ℓ . CAS expressions $\text{CAS}_{Z,W}(e, e', e'')$ are annotated with two access modes: one Z to be used for the successful case, and the other W for the unsuccessful case. The expressions e, e' and e'' respectively denote the values ℓ, v' and v'' mentioned above.

Finally, we write $[E]_{\text{na}}$ to abbreviate **let** $x = E$ **in** $[x]_{\text{na}}$; and we write $E_1; E_2$ for **let** $x = E_1$ **in** E_2 when x is not free in E_2 .

3.2 Memory Model

Like other relaxed memory models [Alglave et al. 2009; Batty et al. 2016, 2011; Mador-Haim et al. 2012; Manson et al. 2005; Sarkar et al. 2009] the semantics that our programming language inherits from RSL [Vafeiadis and Narayan 2013] is defined axiomatically, in which a program's semantics is defined as a set of its *executions*. Each execution is represented by a set of actions plus various binary relations on those actions. Actions in our semantics are identical to those in RSL, except that each allocate action $A_F(\ell)$ carries the security label F of the allocation expression $\text{alloc}_F()$ that generated it.

$$\begin{aligned}
\text{Act} ::= & \text{skip} \mid W_{(\text{sc}|\text{rel}|\text{rlx}|\text{na})}(\ell, v) \mid R_{(\text{sc}|\text{acq}|\text{rlx}|\text{na})}(\ell, v) \\
& \mid \text{RMW}_{(\text{sc}|\text{rel}|\text{rel_acq}|\text{acq}|\text{rlx})}(\ell, v, v') \mid A_{(\text{low}|\text{high})}(\ell)
\end{aligned}$$

Actions describe a program's interactions with memory. The skip action is a no-op, and represents local computation, as well as thread forks and joins. The other actions represent respectively: writes (generated by store expressions); reads (generated by load expressions); atomic read-modify-write actions (generated by CAS expressions); and allocations. Actions carry information about the memory location ℓ that was accessed and the values v, v' etc. read or written.

The binary relations over actions include [Vafeiadis and Narayan 2013] the *sequenced-before relation* sb , which relates actions according to the order that they appear in the program's textual control flow. We have $sb(a, b)$ if a immediately precedes b . The *reads-from relation* rf is a map from read actions r to write actions w : when $rf(r) = w$, it means that the value read by the action r was written to the location that is being read by the action w . The *memory-order* [Vafeiadis and Narayan 2013] relation mo (aka the *modification-order* [Batty 2014] relation) and the *sequential-consistency*

order relation sc are total orders: the former relates write actions to the same atomic location; the latter imposes a total order on all sc actions.

The arrows in Fig. 2 depict three of these relations: sb (solid, black); rf (dashed, green); and mo (dotted, orange).

In a program execution, each action is identified by a unique *name*, drawn from some countably infinite set $AName$. To give meaning to these opaque action names, the execution includes a *labelling* function lab that associates each with an action $a \in Act$. Then, formally, an execution is a tuple $\langle \mathcal{A}, lab, sb, rf, mo, sc \rangle$ where \mathcal{A} is a finite subset of action names from $AName$, lab is a function from $AName$ to Act , sb and mo are binary relations on elements of \mathcal{A} , and rf is a partial map between elements of \mathcal{A} .

Two additional relations are defined in terms of the others. The *synchronises-with* relation sw : intuitively $sw(w, r)$ holds when w is a release write and r is an acquire read that synchronises with w . We omit its full definition (see Vafeiadis and Narayan [2013, Figure 3]) for the sake of brevity. The *happens-before* relation hb formalises when one action must complete before another, and is the transitive closure of $sb \cup sw$.

The semantics of a program $E \in Exp$ is defined as usual for axiomatic weak memory models: the possible executions of E are calculated and then constrained by the axioms of the memory model to leave only those executions that are *consistent* with the model. We refer to Vafeiadis and Narayan [2013] for the axioms and write $Consistent(\langle \mathcal{A}, lab, sb, rf, mo, sc \rangle)$ when an execution $\langle \mathcal{A}, lab, sb, rf, mo, sc \rangle$ is consistent.

Executions are calculated by applying a semantic function $\llbracket \cdot \rrbracket$ that maps expressions $E \in Exp$ to a set $\llbracket E \rrbracket$ of their executions, each of which is a tuple $\langle res, \mathcal{A}, lab, sb, fst, lst \rangle$ where \mathcal{A} , lab and sb are as above, res is the expression's result, and fst and lst respectively denote the (action names of) the first and last actions of the expression in the sb order. We also refer readers to Vafeiadis and Narayan [2013, Section 3] for the details of the semantics, which are unchanged except that allocation expressions $alloc_F()$ produce allocation actions $A_F(\ell)$ that carry their security label F . Note that this allocation label is otherwise ignored by the semantics, and so does not influence program execution; it is merely propagated to the label of allocation actions to make it visible to SecRSL's top-level security property (Section 5.1).

Like RSL, some rules of SecRSL are valid only in a stronger memory model that strengthens the C11 axioms to exclude so-called “out of thin air” reads in programs with relaxed writes. The definition of this stronger memory model is identical to that of Vafeiadis and Narayan [2013, Section 6]. We refer to it as the *strengthened* memory model, while noting that this strengthening is quite standard and was used as the base for various separation logics that succeeded RSL [Doko and Vafeiadis 2016, 2017].

4 LOGIC

Like many other logics for proving secure information flow [Barthe et al. 2017; Benton 2004; Ernst and Murray 2019; Maillard et al. 2019], which necessarily relate pairs of program executions, SecRSL is a *relational* logic [Benton 2004]. In general, SecRSL reasons about a pair of A-Normal form programs E and E' to prove that they are indistinguishable to the attacker and thus do not leak sensitive information. (We define indistinguishability in Section 5.1.)

For example, consider some program $Prog_1(hi, lo)$ whose initial state comprises some secret hi , as well as some public data lo . Security requires that when run from two initial states that agree on lo but might differ on hi , the executions of $Prog_1$ are indistinguishable. Therefore, to prove this hypothetical program secure, we must prove that for all values lo, hi, lo', hi' , if lo and lo' are equal then the behaviours of $Prog_1(lo, hi)$ are indistinguishable to the attacker from the behaviours of $Prog_1(lo', hi')$.

SecRSL Judgement. Thus the SecRSL judgement:

$$[HP], \{P\} E, E' \{(y, y'). Q\}, [(y, y'). HQ]$$

where E and E' are expressions (programs). Intuitively, this judgement means that from the attacker's point of view, the executions of E are indistinguishable from those of E' . P (respectively Q) is a *relational* separation logic precondition [Ernst and Murray 2019] (respectively postcondition, which may refer to the values y and y' returned by the expressions E and E' respectively). We introduce SecRSL's assertion language shortly. HP (respectively HQ) under-approximates the set of memory locations known to be **high** (i.e. unobservable to the attacker) before (respectively after) E and E' execute. Like Q , HQ can also refer to the return values y and y' of E and E' respectively.¹ As with prior security logics (e.g. SecCSL [Ernst and Murray 2019]), the sets HP and HQ can depend on values read during the program and so support value-dependent location sensitivity [Murray et al. 2016] (i.e. locations whose classification depends on runtime values).

Notational Shorthands. For boolean condition B and SecRSL assertion P , we write $P \wedge B$ to abbreviate *if B then P else false*. When writing postconditions we will write $\{(y, y). Q\}$ instead of $\{(y, y'). Q \wedge y = y'\}$, or omit the “ (y, y') .” when the postcondition does not refer to the return values. To save space, we also avoid duplicating the “ (y, y') .” prefix in both the postcondition and the post-set of **high** locations, omitting it from the latter. We write $\{P\} E, E' \{(y, y'). Q\}$, when E and E' neither modify nor depend on the **high** locations set, which abbreviates $\forall H.P$. $[HP], \{P\} E, E' \{(y, y'). Q\}, [HP]$. We write $[HP], \{P\} E \{(y, y'). Q\}, [HQ]$, when talking about two identical programs E , to abbreviate $[HP], \{P\} E, E \{(y, y'). Q\}, [HQ]$. We also combine these abbreviations: $\{P\} E \{(y, y'). Q\}$ means $\forall H.P$. $[HP], \{P\} E, E \{(y, y'). Q\}, [HP]$. These conventions are used both for presenting the rules of the logic and to simplify proof sketches like that of Fig. 3.

SecRSL Assertions. SecRSL's relational assertions P, Q etc. are evaluated over pairs of memories, like SecCSL's assertions [Ernst and Murray 2019]. They include relational (i.e. 2-state) analogues of all the RSL assertions [Vafeiadis and Narayan 2013] and are defined below.

$$\begin{aligned} P, P_1, P_2 ::= & \text{false} \mid P_1 \Rightarrow P_2 \mid \forall(x, x'), P \mid (\ell, \ell') \mapsto (e, e') \\ & \mid \text{emp} \mid P_1 \star P_2 \mid \text{Init}(\ell) \mid \text{Uninit}(\ell, \ell') \mid \text{Low}(e, e') \\ & \mid \text{Rel}(\ell, Q) \mid \text{Acq}(\ell, Q) \mid \text{RMWAcq}(\ell, Q) \end{aligned}$$

In a judgement $[HP], \{P\} E, E' \{(y, y'). Q\}, [(y, y'). HQ]$ we call E the *left* program and E' the *right* program. SecRSL assertions are therefore evaluated in pairs of heaps (memories): the left (respectively right) memories are those encountered during the execution of E (respectively E'). Thus universal quantification quantifies over both a left and a right variable; points-to assertions $(\ell, \ell') \mapsto (e, e')$ state that in the left memory, the non-atomic location ℓ holds the value denoted by e (and likewise for the right memory with ℓ' and e'). Likewise $\text{Uninit}(\ell, \ell')$ means that the non-atomic location ℓ (ℓ') is uninitialised in the left (right) memory. $\text{Low}(e, e')$ means that e and e' denote equal values, and so are indistinguishable to the attacker and, hence, can be treated as public data that is safe to reveal to the attacker without violating security. It is analogous to SecCSL's *value-sensitivity* assertion [Ernst and Murray 2019].

As in RSL, the assertions $\text{Init}(\ell)$, $\text{Rel}(\ell, Q)$, $\text{Acq}(\ell, Q)$ and $\text{RMWAcq}(\ell, Q)$ refer to atomic locations ℓ . $\text{Init}(\ell)$ means that atomic location ℓ has been initialised; the others represent the permission to perform the corresponding action on that location. Q is an invariant, attached to atomic location ℓ , and is parameterised by the value v in memory at location ℓ . For instance, $\text{Acq}(\ell, Q)$ is the permission to perform an acquire load on location ℓ , obtaining the invariant $Q(v)$ where v is the

¹That SecRSL tracks which locations are known to be unobservable to the attacker, rather than tracking those that are known to be observable, might be surprising. We explain this design choice later in Section 5.4.

value read from the location; $\text{Rel}(\ell, Q)$ is the permission to perform a release write to location ℓ , provided that the invariant $Q(v)$ holds where v is the value being written; the $\text{RMWAcq}(\ell, Q)$ permission allows performing the CAS operation on location ℓ .

Unlike the aforementioned assertions for non-atomic locations, those for atomic locations do not need to mention pairs of locations (ℓ, ℓ') or pairs of values (v, v') . This is because SecRSL requires that each atomic location is **low** and that the access patterns of which locations are accessed are identical between the left and right executions (Section 2). Thus atomic locations will be proved to be identical in the left and right executions (both the location ℓ itself and the value v that it holds).

From these assertions the standard logical connectives true, \wedge , \vee , \neg , and \exists can be derived in the usual way. Moreover, SecRSL assertions enjoy the usual semantic equivalences of separation logic. In particular, we have that $\forall P, Q, R \in \text{Assn}$:

$$\begin{aligned}
& P \star Q \sim Q \star P \\
& P \star (Q \star R) \sim (P \star Q) \star R \\
& P \star \text{emp} \sim P \\
& \text{false} \star \text{false} \sim \text{false} \\
& P \vee Q \sim Q \vee P \\
& P \vee (Q \vee R) \sim (P \vee Q) \vee R \\
& (P \vee \text{false}) \sim (P \vee P) \\
& (P \vee P) \sim P
\end{aligned}$$

These equivalences are derived from the assertion semantics, which will be introduced in Section 5.2

$$\begin{array}{c}
\frac{}{\{\text{emp}\} \text{alloc}_{\text{low}}() \{(\ell, \ell).\text{Uninit}(\ell, \ell)\}} \text{A-L} \\
\frac{}{[\text{HF}], \{\text{emp}\} \text{alloc}_{\text{high}}() \{(\ell, \ell).\text{Uninit}(\ell, \ell)\}, [\{\ell\} \cup \text{HF}]} \text{A-H} \\
\frac{\text{if } \ell \notin \text{HF} \text{ then } G = \text{Low}(v, v') \text{ else } G = \text{emp} \\ (P = \text{Uninit}(\ell, \ell)) \vee (P = (\ell, \ell) \mapsto (_, _))}{[\text{HF}], \{P \star G\} [\ell]_{\text{na}} := v, [\ell]_{\text{na}} := v' \{(\ell, \ell) \mapsto (v, v')\}, [\text{HF}]} \text{NA-W} \\
\frac{\{(\ell, \ell) \mapsto (v, v')\} [\ell]_{\text{na}}, [\ell]_{\text{na}} \{(y, y').(\ell, \ell) \mapsto (v, v') \wedge y = v \wedge y' = v'\}}{\{P\} e, e' \{(y, y').P \wedge y = e \wedge y' = e'\}} \text{NA-R} \\
\frac{}{\{P\} e, e' \{(y, y').P \wedge y = e \wedge y' = e'\}} \text{VALUE}
\end{array}$$

Fig. 4. Non-Atomic Rules

Non-Atomic Rules. The rules for operations on non-atomic locations are shown in Fig. 4. Recall that our threat model assumes a memory allocator that does not leak sensitive information, i.e. the resolution of the nondeterministic choice about which memory location the $\text{alloc}()$ expression returns does not reveal secrets. Additionally, recall from Section 2 that SecRSL does not allow programs to branch on secrets. This ensures that all decisions about whether to allocate or not depend only on public data. Hence, when reasoning about memory allocation (rules A-L and A-H),

$$\begin{array}{c}
\frac{}{\{\text{emp}\} \text{alloc}_{\text{low}}() \{(\ell, \ell). \text{Rel}(\ell, Q) \star \text{Acq}(\ell, Q)\}} \text{A-R} \\
\frac{}{\{\text{emp}\} \text{alloc}_{\text{low}}() \{(\ell, \ell). \text{Rel}(\ell, Q) \star \text{RMWAcq}(\ell, Q)\}} \text{A-M} \\
\frac{}{\{\text{Acq}(\ell, Q) \star \text{Init}(\ell)\} [\ell]_{\text{rlx}} \{(v, v). \text{Acq}(\ell, Q)\}} \text{RLX-R} \\
\frac{}{\{\text{Rel}(\ell, Q) \star Q(v)\} [\ell]_{\text{rel}} := v \{\text{Rel}(\ell, Q) \star \text{Init}(\ell)\}} \text{REL-W} \\
\frac{\forall v. \text{precise}(Q(v))}{\{\text{Acq}(\ell, Q) \star \text{Init}(\ell)\} [\ell]_{\text{acq}} \{(v, v). Q(v) \star \text{Acq}(\ell, Q [v := \text{emp}])\}} \text{ACQ-R} \\
\frac{Q = \lambda x. (\text{if } x = v \text{ then emp else false})}{\{\text{Rel}(\ell, Q)\} [\ell]_{\text{rlx}} := v \{\text{Init}(\ell)\}} \text{RLX-W}^* \\
\frac{}{\{\text{Acq}(\ell, Q) \star \text{Init}(\ell)\} [\ell]_{\text{rlx}} \{(v, v). \text{Acq}(\ell, Q) \wedge Q(v) \neq \text{false}\}} \text{RLX-R}^* \\
\frac{
\begin{array}{l}
P \implies \text{Init}(\ell) \star \text{RMWAcq}(\ell, Q) \star \text{true} \\
P \star Q(v) \implies \text{Rel}(\ell, Q') \star Q'(v') \star R[v/y] \\
X \in \{\mathbf{rel}, \mathbf{rlx}\} \implies Q(v) = \text{emp} \\
X \in \{\mathbf{acq}, \mathbf{rlx}\} \implies Q'(v') = \text{emp} \\
\{P\} [\ell]_Y \{(y, y). y \neq v \Rightarrow R\}
\end{array}
}{\{P\} \text{CAS}_{X,Y}(\ell, v, v') \{(y, y). R\}} \text{CAS}^* \\
\frac{[HP], \{P\} E_t, E_t \{(y, y). Q\}, [HQ]}{[HP], \{P\} E_T, E_T \{(y, y). Q\}, [HQ]} \text{RELAX}, \text{ where } \mathbf{rlx} \sqsubseteq \mathbf{rel} \sqsubseteq \mathbf{sc} \text{ and } \mathbf{rlx} \sqsubseteq \mathbf{acq} \sqsubseteq \mathbf{sc}
\end{array}$$

Fig. 5. Atomic Rules

the postcondition can assert that the two locations ℓ and ℓ' returned in the left and right executions respectively are equal. Naturally, **high** allocations increase the set of known **high** locations (see rule A-H).

When writing to a non-atomic location ℓ , the rule NA-W ensures security by requiring that if the write might be visible to the attacker ($\ell \notin HF$) then the values v and v' being written must be indistinguishable: $\text{Low}(v, v')$.

In addition notice that the rules for reading and writing non-atomic locations require that the location ℓ being accessed is identical in both programs. This is required to provide SecRSL's constant-time guarantee of indistinguishability against attackers who can observe the memory access pattern of the program (Section 2).

The VALUE rule in Fig. 4 is analogous to its RSL counterpart.

Atomic Rules. The rules for operations on atomic memory locations are depicted in Fig. 5, and are each analogous to their RSL counterparts. Rules for allocating atomic locations (A-R and A-M), like their non-atomic counterparts, assert that the chosen location is identical (and thus public) in both executions, returning the appropriate permissions as in RSL. Notice that these rules are valid only for **low** allocations, since SecRSL requires that all atomic locations are **low** (Section 2). That is also the reason why each of these rules talks about two identical programs, since SecRSL

$$\begin{aligned}
& \text{emp} \iff \text{Low}(x, x) \\
& \text{Init}(\ell) \iff \text{Init}(\ell) \star \text{Init}(\ell) \\
& \text{Rel}(\ell, \lambda v. Q_1(v) \vee Q_2(v)) \iff \text{Rel}(\ell, Q_1) \star \text{Rel}(\ell, Q_2) \\
& \text{Acq}(\ell, \lambda v. Q_1(v) \star Q_2(v)) \iff \text{Acq}(\ell, Q_1) \star \text{Acq}(\ell, Q_2) \\
& \text{RMWAcq}(\ell, Q) \iff \text{RMWAcq}(\ell, Q) \star \text{RMWAcq}(\ell, Q) \\
& \frac{\forall v. (Q'(v) = \text{emp} \vee Q(v) = Q'(v) = \text{false})}{\text{RMWAcq}(\ell, Q) \iff \text{RMWAcq}(\ell, Q) \star \text{Acq}(\ell, Q')}
\end{aligned}$$

Fig. 6. Split Rules

$$\begin{aligned}
& \frac{(P \implies P') \wedge HP' \subseteq HP \quad \forall (y, y'). (Q' \implies Q) \wedge HQ \subseteq HQ' \quad [HP'], \{P'\} E, E' \{(y, y').Q'\}, [HQ']}{[HP], \{P\} E, E' \{(y, y').Q\}, [HQ]} \text{CONSEQ} \\
& \frac{[HP], \{P\} E, E' \{(y, y').Q\}, [HQ] \quad [HP], \{P'\} E, E' \{(y, y').Q'\}, [HQ]}{[HP], \{P \vee P'\} E, E' \{(y, y').Q \vee Q'\}, [HQ]} \text{DISJ} \\
& \frac{[H1], \{b \wedge b' \wedge P\} E_1, E'_1 \{(y, y').Q\}, [HQ] \quad [H2], \{\neg b \wedge \neg b' \wedge P\} E_2, E'_2 \{(y, y').Q\}, [HQ]}{E = \text{if } b \text{ then } E_1 \text{ else } E_2 \quad E' = \text{if } b' \text{ then } E'_1 \text{ else } E'_2} \text{IF} \\
& \frac{[H1], \{b \wedge b' \wedge P\} E_1, E'_1 \{(y, y').Q\}, [HQ] \quad [H2], \{\neg b \wedge \neg b' \wedge P\} E_2, E'_2 \{(y, y').Q\}, [HQ]}{[\text{if } b \text{ then } H1 \text{ else } H2], \{\text{Low}(b, b') \star P\} E, E' \{(y, y').Q\}, [HQ]} \text{IF} \\
& \frac{[HP], \{P\} E, E' \{(y, y').Q\}, [HQ]}{[HP], \{\exists x.P\} E, E' \{(y, y').\exists x.Q\}, [HQ]} \text{EX} \\
& \frac{[HP], \{P\} E, E' \{(y, y').Q\}, [HQ]}{[HP], \{P \star F\} E, E' \{(y, y').Q \star F\}, [HQ]} \text{FRAME} \\
& \frac{[HP], \{P_1\} E_1, E'_1 \{(y, y').Q_1\}, [HQ_1] \quad [HP], \{P_2\} E_2, E'_2 \{(y, y').Q_2\}, [HQ_2]}{[HP], \{P_1 \star P_2\} E_1 \parallel E_2, E'_1 \parallel E'_2 \{(y, y').Q_1 \star Q_2\}, [HQ_1 \cup HQ_2]} \text{PAR} \\
& \frac{\{P\} E, E' \{(y, y').Q \star \text{Low}(y, y')\} \quad Q[(0, 0)/(y, y')] \implies P}{\{P\} \text{ repeat } E \text{ end, repeat } E' \text{ end } \{(y, y').Q \wedge y \neq 0\}} \text{REPEAT} \\
& \frac{[HP_1], \{P_1\} E_1, E'_1 \{(x, x').Q_1\}, [HQ_1] \quad \forall (x, x'). [HQ_1], \{Q_1\} E_2, E'_2 \{(y, y').Q_2\}, [HQ_2]}{[HP_1], \{P_1\} \text{ let } x = E_1 \text{ in } E_2, \text{ let } x' = E'_1 \text{ in } E'_2 \{(y, y').Q_2\}, [HQ_2]} \text{LET}
\end{aligned}$$

Fig. 7. Structural Proof Rules

requires that all actions on **low** locations are identical between the two executions. As in RSL, performing a release write (rule REL-W) requires the permission $\text{Rel}(\ell, Q)$ as well as evidence that the invariant $Q(v)$ holds for the value v being written, which must be public (equal in both executions), as must the location ℓ being written to. Similarly, performing an acquire read (rule ACQ-R) requires the appropriate permission plus evidence that the location was initialised. Doing so obtains the invariant $Q(v)$ for the value v that was read, which is identical in both executions (because atomic locations must be **low**). As in RSL this rule requires that Q is precise, whose definition for SecRSL's relational assertions we postpone to Definition 5.6 in Section 5.2. The rules for relaxed reads and writes also follow their RSL counterparts. As there, the rules RLX-R*, RLX-W*, and CAS* (when $X \notin \{\mathbf{rel_acq}, \mathbf{sc}\}$) are only sound in the *strengthened* memory model (Section 3.2) of RSL and its successors [Doko and Vafeiadis 2016, 2017].

Split Rules. SecRSL assertions behave identically to their RSL counterparts. The assertions on atomic locations behave identically as in RSL (see Fig. 6): $\text{Init}(\ell)$ and $\text{RMWAcq}(\ell, Q)$ are duplicable, as is $\text{Rel}(\ell, Q)$ which can also be split along disjunctions in Q ; $\text{Acq}(\ell, Q)$ on the other hand can be split along separating conjunction in Q . $\text{RMWAcq}(\ell, Q)$ can also be split to produce a separate, but weak, $\text{Acq}(\ell, Q)$ permission.

Structural Rules. The structural rules for SecRSL are shown in Fig. 7. It supports all the standard rules from concurrent separation logic, like RSL before it. Notice that the consequence rule CONSEQ allows one to expand the pre-set HP of **high** locations, or reduce the post-set HQ . As in SEC CSL [Ernst and Murray 2019], the rules for conditionals and loops do not allow programs to branch on secrets. This is a necessary condition to enforce SecRSL's constant-time security (Section 2).

As with RSL, the rule for parallel composition PAR is not symmetric. This is because we inherit RSL's semantics for parallel composition in which, for $E_1 \parallel E_2$, the return value of E_2 is discarded. This is why the postcondition Q_2 in the second premise of the PAR rule does not refer to the return values of E_2 and E'_2 .

5 SOUNDNESS

Like all program logics, SecRSL's soundness is proved against the programming language semantics. Specifically, as with prior security separation logics like SEC CSL [Ernst and Murray 2019], we define a set of top-level security and safety properties over the programming language semantics that apply to entire program executions. We prove that if a program is verified in SecRSL then it enjoys these properties. We call this result SecRSL's *adequacy theorem* (Theorem 5.5), in Section 5.1.

In order to prove these properties, we developed a compositional and inductive definition of relational validity over the C11 weak memory model that encodes the semantic meaning of SecRSL judgements $[HP]$, $\{P\} E, E' \{(y, y'). Q\}$, $[HQ]$. We call this property *relational validity*. It holds when for every consistent execution (i.e. one that is consistent with the axioms of the memory model) of E , one can *inductively* construct a related consistent execution of E' that witnesses the security of both executions (by guaranteeing that they are indistinguishable) and their safety. The rules of the logic are then proved sound against this definition of relational validity.

Naturally, this relational validity rests on the semantics of SecRSL's assertions, which we present in Section 5.2. We present the definition of semantic validity for SecRSL's judgements in Section 5.3.

As mentioned, SecRSL's soundness proof was mechanised in Coq, and is available as supplementary material [Yan 2021].

5.1 Adequacy

SecRSL's adequacy theorem includes SecRSL's top-level security guarantees of *low safety* and *location safety*: the former guarantees that an attacker who can observe the contents of **low** memory locations throughout a program's execution cannot learn sensitive information. The latter strengthens this guarantee to also include attackers who can observe the memory-access pattern of the program, i.e. can observe which locations are being accessed but not the values being written to them. This latter property provides protection against attackers who e.g. can mount cache timing channel attacks, and so provides a form of *constant-time* security [Barthe et al. 2019].

To our knowledge, these are the first top-level confidentiality properties defined over an axiomatic weak memory semantics. We believe that our properties are of independent interest and could be adapted to other axiomatic weak memory models.

The top-level properties that SecRSL guarantees via its adequacy theorem also include that programs are memory safe, never read from uninitialised locations, and are free of data races. These properties SecRSL inherits from RSL [Vafeiadis and Narayan 2013].

Whole-program executions. SecRSL's top-level properties are defined over whole-program executions. For a program E , we denote this set of executions $CC[[E]]$ (the set of *complete consistent* executions [Vafeiadis and Narayan 2013]). As in RSL, whole-program executions of E are obtained from the consistent executions of E by adding two extra skip actions: one that precedes the execution of E and another that follows it.

Definition 5.1 (Whole-program executions). We denote by $CC[[E]]$ the set of whole-program executions of E :

$$CC[[E]] \stackrel{\text{def}}{=} \{ \langle res, \mathcal{A}, lab, sb, rf, mo, sc \rangle \mid \\ \exists a, b, \mathcal{A}_{prg}, lab_{prg}, sb_{prg}, fst, lst. a \neq b \wedge \mathcal{A} = \{a, b\} \uplus \mathcal{A}_{prg} \wedge \\ sb = sb_{prg} \cup (a, fst) \cup (lst, b) \wedge lab = \{a \rightarrow \text{skip}, b \rightarrow \text{skip}\} \cup lab_{prg} \wedge \\ \langle res, \mathcal{A}_{prg}, lab_{prg}, sb_{prg}, fst, lst \rangle \in [[E]] \wedge \text{Consistent}(\langle \mathcal{A}, lab, sb, rf, mo, sc \rangle) \}$$

Top-level Security Statement. A key challenge in phrasing the top-level adequacy statement over the C11 memory model, which would persist also for other axiomatic memory models, is dealing with the nondeterminism inherent in such models. Put simply, for two programs E and E' that (due to nondeterminism) each have a set of behaviours as their semantics, what does it mean for the behaviours of E to be indistinguishable from those of E' ?

Following prior work on security definitions for concurrent, nondeterministic programs [Murray and Lowe 2009] one way to define this is to assume that sources of nondeterminism will be resolved in the same way in the two programs, and so to compare only those pairs of executions that arise from resolving the nondeterminism identically in both programs. E and E' are indistinguishable when each such pair of executions is indistinguishable. If we take this approach for SecRSL, what assumptions would it hard-wire into the resulting adequacy theorem?

The sources of nondeterminism in the C11 memory model of Section 3.2 that we inherit from Vafeiadis and Narayan [2013] are the nondeterministic model of memory allocation and the weak memory concurrency semantics. The former abstracts away from how memory locations are allocated in the `alloc()` expression. The latter abstracts away from the concurrency inherent in the C11 language and, hence, the implementation choices of the compiler that generates code for the underlying instruction set architecture and its weak memory model, as well as the internal details of how threads are scheduled by the operating system and hardware.

Considering only pairs of executions in which nondeterminism is resolved identically therefore encodes into the security property the basic assumption of Section 2 that the memory allocator

never leaks sensitive information, and nor does the implementation choices of the weak memory model.

What does it mean for nondeterminism to have been resolved identically in two executions? To answer this question, let us rephrase what it means to prove indistinguishability of each pair of executions of E and E' in which nondeterminism has been resolved identically. This assumption can be stated equivalently by saying that if E' has the opportunity to make an identical choice to E , then it will do so (and vice-versa). Thus it suffices to show for each execution of E , there exists a corresponding execution of E' in which it is evident that the nondeterministic choices have been resolved identically, and then prove that these two executions are indistinguishable.

This is precisely what SecRSL's adequacy theorem does (as well as proving that both executions are safe).

To do so, we must define a sufficient condition that allows us to judge when the nondeterminism in two executions has been resolved identically. Recall that in our semantics (Section 3.2), each execution is a set of actions related by various binary relations. We can view each execution therefore as a directed graph (see e.g. Fig. 2), whose nodes are the actions and whose edges carry labels like sb , rf , etc.

We claim that nondeterminism has been resolved identically in two executions when (1) those executions are *isomorphic* (i.e. there is one-to-one mapping between the nodes of the two graphs and the two graphs have the same structure and edge relationships) and (2) if for every allocation action $A_F(\ell)$ in one graph, the isomorphic action of the other graph is identical (i.e. is also $A_F(\ell)$). Condition (1) implies identical resolution of nondeterminism resulting from the weak memory concurrency model, while (2) implies identical allocation decisions.

Thus SecRSL's adequacy theorem says that for every execution of E , there exists an execution of E' for which conditions (1) and (2) hold, and that the two executions are indistinguishable.

Since our definition of indistinguishability implies condition (2), SecRSL's adequacy theorem is stated more succinctly by saying that for each whole-program execution $\chi \in CC[[E]]$, there exists an isomorphic whole-program execution $\chi' \in CC[[E']]$ that is indistinguishable to χ .

Two executions are isomorphic when they share the same set of nodes and have identical edge relationships, including for the derived relations hb and sw . In the graph analogy, the set of nodes corresponds to the set of opaque action names \mathcal{A} . Hence we have:

Definition 5.2 (Isomorphic). Two whole-program executions $\chi = \langle res, \mathcal{A}, lab, sb, rf, mo, sc \rangle$ and $\chi' = \langle res', \mathcal{A}', lab', sb, rf, mo, sc \rangle$ are *isomorphic* if and only if for all $a, b \in \mathcal{A}$,
 $(hb_\chi(a, b) \iff hb_{\chi'}(a, b)) \wedge (sw_\chi(a, b) \iff sw_{\chi'}(a, b))$.

Execution indistinguishability is defined by the two aforementioned properties of location safety and low safety. The former implies condition (1) above (identical allocation decisions) and implies that the executions are indistinguishable to an attacker who can observe which locations are being accessed but not the contents of those locations. It does so by saying that the labels on isomorphic actions must either match or, if not, they must both be a non-atomic write (respectively read) to the same location.

Definition 5.3 (Location Safety). Two isomorphic whole-program executions $\chi = \langle res, \mathcal{A}, lab, sb, rf, mo, sc \rangle$ and $\chi' = \langle res', \mathcal{A}', lab', sb, rf, mo, sc \rangle$ satisfy *location safety* if and only if for all $a \in \mathcal{A}$. $lab(a) = lab'(a) \vee$
 $(\exists \ell v v'. lab(a) = R_{na}(\ell, v) \wedge lab'(a) = R_{na}(\ell, v')) \vee$
 $(\exists \ell v v'. lab(a) = W_{na}(\ell, v) \wedge lab'(a) = W_{na}(\ell, v'))$

The second indistinguishability property *low safety* defines indistinguishability against an attacker who can observe **low** memory locations. It says that any **low** location must be accessed

identically (including the operations performed on that location and the values written or read to it) between the two executions.

Definition 5.4 (Low Safety). Two isomorphic whole-program executions $\chi = \langle res, \mathcal{A}, lab, sb, rf, mo, sc \rangle$ and $\chi' = \langle res', \mathcal{A}, lab', sb, rf, mo, sc \rangle$ satisfy *low safety* if and only if for all locations ℓ for which $(\exists a \in \mathcal{A}. lab(a) = A_{\text{low}}(\ell) \vee lab'(a) = A_{\text{low}}(\ell))$ it is the case that for all $b \in \mathcal{A}$ such that $lab(b)$ or $lab'(b)$ is an action on ℓ , then $lab(b) = lab'(b)$.

With these definitions, SecRSL's adequacy theorem can be formally stated.

THEOREM 5.5 (ADEQUACY). *Let E and E' be two programs such that, for some Q and HQ , if a SecRSL judgement $\llbracket 0 \rrbracket, \{\text{true}\} E, E' \{(y, y').Q\}, [(y, y').HQ]$ holds, then, for every whole-program execution $\chi \in CC\llbracket E \rrbracket$ there exists some $\chi' \in CC\llbracket E' \rrbracket$ such that χ and χ' are isomorphic, and satisfy location safety and low safety. Moreover, each execution is memory safe, has no reads from uninitialised locations, and no data races.*

We omit the definitions of the safety properties (namely memory safety, plus absence of uninitialised reads and data races), since they are unchanged from RSL [Vafeiadis and Narayan 2013, Section 7.3].

5.2 Semantics of Assertions and Local Relational Validity

As in RSL and its descendants [Doko and Vafeiadis 2016, 2017], SecRSL assertions are given a semantics over heaps (memories) h . Because the axiomatic weak memory model defines executions only in terms of actions and relations between those actions, memories do not appear in the semantics. As with logics like RSL, we define what it means for a pair of executions χ, χ' to be valid with respect to a judgement $\llbracket HP \rrbracket, \{P\} E, E' \{(y, y).Q\}, \llbracket HQ \rrbracket$ by asserting the existence of memories that annotate the edges of the executions, and are consistent with the actions being performed, and satisfy the SecRSL assertions in the judgement.

We reuse the heap model of RSL [Vafeiadis and Narayan 2013, Section 7.1]. However, since SecRSL assertions are relational, each assertion is evaluated over a *pair* (h, h') of heaps [Ernst and Murray 2019; Yang 2007]. Thus the meaning of an assertion P we denote $\llbracket P \rrbracket$ and define $\llbracket - \rrbracket$ inductively as a function that maps an assertion to a set of pairs (h, h') for which the assertion holds.

The full definition appears in Yan and Murray [2021]. For instance, the meaning of the points-to assertion $(\ell, \ell') \mapsto (v, v')$ for non-atomic locations is analogous to its SecCSL counterpart and holds for a pair of heaps (h, h') when $\ell \mapsto v$ in h and $\ell' \mapsto v'$ in h' :

$$\llbracket (\ell, \ell') \mapsto (v, v') \rrbracket \stackrel{\text{def}}{=} \{ \{ \ell \mapsto \text{NA}[v] \}, \{ \ell' \mapsto \text{NA}[v'] \} \}$$

The assertion $\text{Low}(e, e')$ requires that e and e' are indistinguishable to the attacker and is equivalent to *if $e = e'$ then emp else false*:

$$\llbracket \text{Low}(e, e') \rrbracket \stackrel{\text{def}}{=} \text{if } e = e' \text{ then } \{(\emptyset, \emptyset)\} \text{ else } \emptyset$$

This assertion we include as SecRSL's analogue of SecCSL's *value sensitivity* assertion $\cdot :: \text{low}$.

From the relational assertion semantics, we lift the traditional notion of what it means for a separation logic assertion to be *precise* [O'Hearn 2004; Vafeiadis and Narayan 2013] to SecRSL's relational assertions.

Definition 5.6 (Precise). A SecRSL assertion P is *precise*, written $\text{precise}(P)$, if and only if for all h_1, h_2 such that $((h_1, _) \in \llbracket P \rrbracket \wedge (h_2, _) \in \llbracket P \rrbracket)$ or $((_, h_1) \in \llbracket P \rrbracket \wedge (_, h_2) \in \llbracket P \rrbracket)$, for all h_3, h_4 such that $h_1 \oplus h_3 = h_2 \oplus h_4 \neq \text{undef}$, we have $h_1 = h_2 \wedge h_3 = h_4$

In this definition \oplus refers to the heap addition operator of RSL's heap model [Vafeiadis and Narayan 2013, Figure 15].

Local Relational Validity. The concept of *local relational validity* is a key building block of our definition of semantic relational validity (i.e. the semantic meaning) of SecRSL judgements (which we present later in Section 5.3). It is the relational analogue of RSL's local validity [Vafeiadis and Narayan 2013, Definition 4]. It is defined for a pair of executions $\mathcal{X} = \langle res, \mathcal{A}, lab, sb, rf, mo, sc \rangle$ and $\mathcal{X}' = \langle res', \mathcal{A}, lab', sb, rf, mo, sc \rangle$ and a subset $V \subseteq \mathcal{A}$ of their actions whose validity is being asserted. Given heap annotation functions $hmap$ and $hmap'$ that annotate the sb and sw edges of \mathcal{X} and \mathcal{X}' respectively with heaps (to witness the memory of each program at that point during its execution), plus a **high** location set H , local relational validity asserts that the heap annotations are consistent with the actions $a \in V$ being performed at each point in those executions; plus imposing sufficient conditions to ensure that the executions are isomorphic and indistinguishable, and that each is safe. It also asserts that all **high** allocations must appear in H and all **low** allocations must not appear in H .

We relegate its definition to Yan and Murray [2021]. For executions $\mathcal{X} = \langle res, \mathcal{A}, lab, sb, rf, mo, sc \rangle$ and $\mathcal{X}' = \langle res', \mathcal{A}, lab', sb, rf, mo, sc \rangle$, actions $V \subseteq \mathcal{A}$, heap annotation functions $hmap$ and $hmap'$ and location set H , we write $\text{Valid}(\mathcal{X}, V, lab', hmap, hmap', H)$ when local relational validity holds for the subparts of the two executions identified by V .

Local relational validity considers just the subparts of the graphs identified by V in order to allow it to be inductively asserted over an ever-increasing part of the two executions, by the inductive definition of relational validity of SecRSL judgements. Finally, we now present that definition.

5.3 Relational Validity

We now define the semantic meaning of SecRSL's judgements $[HP]$, $\{P\} E, E' \{(y, y'). Q\}$, $[HQ]$. This definition is a novel, compositional and inductive relational validity property and, to our knowledge, the first such defined for an axiomatic weak memory semantics. As with the top-level security properties, we believe its design is of independent interest and should also be applicable to other axiomatic weak memory models. However, unlike the top-level definitions, it does not need to be trusted: it is an internal definition used to carry out the soundness proof of the logic; not a statement of any guarantees provided by the logic—instead those guarantees are provided by SecRSL's adequacy theorem (Theorem 5.5).

Recall that we refer to the semantic definition that encodes the meaning of SecRSL's judgements as *relational validity*. It is defined formally in Yan and Murray [2021]. Its formal definition uses various notations introduced in Vafeiadis and Narayan [2013]; however it is not necessary to parse the mathematics in order to understand it, which we carefully explain here.

In order to be compositional, relational validity considers executions of E and E' over all possible contexts. That is it considers (what we call) *contextual executions* in which the actions of E are embedded in those of some larger context. As is common in definitions of validity for separation logic judgements [Vafeiadis 2011], relational validity also quantifies over all possible frame assertions R , to be conjoined with the precondition P and postcondition Q , in order to derive the FRAME rule. In addition it quantifies over all possible supersets H of the initial **high**-location set HP .

Relational validity of the judgement $[HP]$, $\{P\} E, E' \{(y, y'). Q\}$, $[HQ]$ considers each consistent contextual execution of E . For each it considers all potential contexts of executions of E' in which the executions preceding E and E' satisfy local relational validity wrt H and the memories directly preceding E and E' satisfy $P \star R$. For each it then asserts *configuration safety* (defined in Yan and Murray [2021]). Configuration safety inductively asserts the existence of an execution of E' in this context whose final memories satisfy $Q \star R$. It requires that at each step of the induction

the execution of E' constructed so far satisfies local relational validity wrt the same part of the execution of E , under the assumption that the same is true for actions contributed by the contexts. In this way, configuration safety guarantees that the executions obtained when E and E' finish are isomorphic, indistinguishable, and that both are safe (i.e. free of undefined behaviour like reads from uninitialised memory or data races, etc.)

In order to define configuration safety, we had to develop an inductive characterisation that asserts *partial consistency* of the execution of E' constructed at each step of the induction. This property asserts a subset of the memory model's consistency axioms and can be applied inductively to an ever growing subpart of the execution of E' as it is being constructed. When combined together with local relational validity, partial consistency guarantees that for whole programs E' the entire constructed contextual execution will be consistent with the axioms of the memory model. Partial consistency is defined in [Yan and Murray \[2021\]](#).

5.4 Discussion

Soundness Proof. With the definition of relational validity, the soundness proof of SecRSL proceeds by proving the soundness of each of the rules against this definition. This proof follows a similar structure to that of RSL; however is considerably more complicated because of the additional need to construct the witness execution of E' .

The resulting Coq proof is $\sim 21,600$ lines, as compared to the original Coq proof for RSL which is $\sim 15,000$ lines. The two proofs share $\sim 6,000$ lines in common (basic libraries, C11 language definition and memory model) that, when excluded, make the SecRSL proof about $\sim 73\%$ larger.

*The **high** location set.* Readers familiar with logics for information-flow security might wonder why SecRSL tracks the locations known to be **high** (unobservable to the attacker). Most other logics instead track which locations are known to be observable to the attacker, whether through special assertions like SecCSL's *location sensitivity assertions* [[Ernst and Murray 2019](#)] or through a static labelling function or otherwise.

The reason is that the sets HP and HQ used for this purpose in SecRSL's judgements $[HP]$, $\{P\} E$, $E' \{(y, y). Q\}$, $[HQ]$ necessarily *under-approximate* the set of **high** locations. They must since the context in which E (or E') executes might have allocated additional memory locations unused by E . Were these sets instead used to track **low** locations, they would still need to under-approximate the true set of attacker-observable locations. However, under-approximating the set of attacker visible locations is not sound, as it would allow the logic to "forget" that a location was attacker-visible and so allow it to be written with sensitive data.

Points-To Assertion. Notice that all points-to assertions generated by SecRSL are of the form $(\ell, \ell) \mapsto (v, v')$ in which the *same* location ℓ is referenced, albeit with two possibly different values v and v' . This design choice is intentional and ensures SecRSL's constant-time guarantee (which requires that *which* memory locations are accessed by the program and the *order* in which they are accessed never depends on secrets). We purposefully chose to retain a pair of locations (ℓ, ℓ) to emphasise the relational nature of this assertion.

Indeed, while SecRSL's rules support compositional reasoning about partial programs, its adequacy theorem (Theorem 5.5) necessarily applies only to whole programs (a fundamental limitation it shares with prior separation logics defined for axiomatic memory models [[Doko and Vafeiadis 2016, 2017](#); [Vafeiadis and Narayan 2013](#)] that can give meaningful semantics only to whole programs). For this reason, modifying SecRSL to attempt to weaken its constant-time guarantee would not yield a more expressive logic.

Beyond SecRSL. We argue explicitly that the ideas underpinning SecRSL’s design and soundness proof should be readily applicable to other logics that extend the RSL memory model. Doko and Vafeiadis’ Fenced Separation Logic (FSL) [Doko and Vafeiadis 2016] is an obvious target. Indeed an information-flow security analogue of FSL would allow reasoning about a wider class of programs beyond the Release-Acquire fragment of C11 considered here.

FSL’s memory model is a small extension of the strengthened memory model of Section 3.2 to add support for fences. A security analogue of FSL would treat fence operations as potentially attacker observable, much like accesses to atomic locations in SecRSL (which, recall, must be **low**). With this insight, SecRSL’s adequacy statement would apply with almost no modification to the FSL semantics.

Just as RSL’s atomic location permissions (like $\text{Acq}(\ell, Q)$) were readily adapted to SecRSL’s relational setting while ensuring they behaved identically to their original counterparts, we conjecture the same should be true for the additional assertion modalities that FSL introduces for reasoning about fences.

The definitions of local relational validity (see Yan and Murray [2021]) could be applied to FSL by simply extending it to add a case for fence actions, while keeping its current structure. Given the similar structure of SecRSL’s soundness proof to that of RSL, and the close similarity of the soundness proofs for FSL and RSL, there is strong evidence to suggest that the soundness proof for a security analogue of FSL should follow a similar structure to that used in this paper. We leave its development for future work.

6 APPLYING THE LOGIC

We demonstrate SecRSL by verifying a number of case studies in Coq.

6.1 Verifying a Spinlock Module

For our first demonstration of SecRSL we show that, by virtue of its intentional similarity to RSL, SecRSL allows one to replay RSL proofs and, in doing so, obtain stronger guarantees than those provable in RSL.

Specifically we consider the spinlock case study of Vafeiadis and Narayan [2013], which we repeat below and modify only slightly to ensure that the location x that is allocated to create the lock is **low**.

$$\begin{aligned} \text{new_lock}() &\stackrel{\text{def}}{=} \text{let } x = \text{alloc}_{\text{low}}() \text{ in } [x]_{\text{rel}} := 1; x \\ \text{spin}(x) &\stackrel{\text{def}}{=} \text{repeat } [x]_{\text{rlx}} \text{ end} \\ \text{lock}(x) &\stackrel{\text{def}}{=} \text{repeat } \text{spin}(x); \text{CAS}_{\text{acq,rlx}}(x, 1, 0) \text{ end} \\ \text{unlock}() &\stackrel{\text{def}}{=} [x]_{\text{rel}} := 1 \end{aligned}$$

This module has a SecRSL specification that is almost identical to its RSL specification.

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

While syntactically identical, note that these specifications once expressed in SecRSL say that the lock module is not only correct, but also adheres to SecRSL’s constant-time security guarantee. Thus the lock module will not leak information.

<pre> new_MSM() $\stackrel{\text{def}}{=} \text{let } b = \text{alloc}_{\text{high}}() \text{ in} \text{let } a = \text{alloc}_{\text{low}}() \text{ in} [b]_{\text{na}} := 0; [a]_{\text{rel}} := 1;$</pre>	<pre> lock_MSM(a) $\stackrel{\text{def}}{=} \text{repeat} \text{let } x = (\text{repeat } [a]_{\text{rlx}} \text{ end}) \text{ in} \text{let } y = \text{CAS}_{\text{acq,rlx}}(a, x, 0) \text{ in} \text{if } (x == y) \text{ then } x \text{ else } 0 \text{end}$</pre>
<pre> unlock_MSM_L(a) $\stackrel{\text{def}}{=} [a]_{\text{rel}} := 1$</pre>	<pre> unlock_MSM_H(a) $\stackrel{\text{def}}{=} [a]_{\text{rel}} := 2$</pre>

Fig. 8. The mixed-sensitivity mutex implementation.

We also observe that these specifications correspond to the rules for the lock and unlock operations of SECCSL [Ernst and Murray 2019]. The proof of these specifications is sketched in Yan and Murray [2021].

It is no accident that this proof is almost syntactically identical to its RSL counterpart. This arises not only because the sketch follows the notational shorthands mentioned on page 9 and in the caption of Fig. 3, but—more importantly—because we carefully designed SecRSL to support all of RSL’s reasoning principles unchanged.

6.2 Mixed-Sensitivity Mutex

A common theme in much recent work on verified information flow security has been *mixed-sensitivity* (also known as *value-dependent classification*) data structures [Ernst and Murray 2019; Frumin et al. 2021; Lourenço and Caires 2015; Murray et al. 2018; Sison and Murray 2019]. These are ones that can hold data of varying sensitivity over time, where the data structure maintains information about the sensitivity of the data it currently contains.

We demonstrate SecRSL’s ability to reason about the implementations of such concurrent data structures for the first time, while taking account of weak memory effects. We extend the prior spinlock module so that it protects access to a mixed-sensitivity memory location. We refer to this resulting abstraction as a *mixed-sensitivity mutex*.

`new_MSM()` creates a mixed sensitivity mutex a . The `lock_MSM(a)` operation acquires access to the mixed-sensitivity location and returns a boolean indicating whether it currently holds sensitive (**high**) data or not. The module provides two operations, `unlock_MSM_H(a)` and `unlock_MSM_L(a)`, to relinquish access to the mixed-sensitivity location, depending on whether the data it now contains is **high** or **low** respectively.

Fig. 8 shows its implementation. The original spinlock writes 0 to location a to indicate that the lock is occupied (acquired) and 1 for unoccupied (free). The mixed-sensitivity mutex uses 0 to indicate that the mutex is occupied, while 1 and 2 indicate it is unoccupied and holding **low** (1) or **high** (2) data.

Defining the predicates $\text{ClassLow}(\ell)$ (respectively $\text{ClassHigh}(\ell)$) to denote when (non-atomic) location ℓ holds a possibly **high** (respectively definitely **low**) value, we verify the mixed-sensitivity mutex against the following specifications. Here the predicate $\text{LockMSM}(a)$ says that location a refers to a mixed-sensitivity mutex.

$$\begin{aligned}
& [], \{\text{emp}\} \text{new_MSM}(), \text{new_MSM}() \{\text{LockMSM}(a)\}, [b] \\
& \text{LockMSM}(a) \iff \text{LockMSM}(a) \star \text{LockMSM}(a) \\
& [b], \{\text{LockMSM}(a) \star \text{ClassLow}(b)\} \text{unlock_MSM_L}(a) \{\text{LockMSM}(a)\}, [b] \\
& [b], \{\text{LockMSM}(a) \star \text{ClassHigh}(b)\} \text{unlock_MSM_H}(a) \{\text{LockMSM}(a)\}, [b]
\end{aligned}$$

$$[b], \{\text{LockMSM}(a)\} \text{lock_MSM}(a) \{(y, y). \text{LockMSM}(a) \star (\text{if } y = 1 \text{ then ClassLow}(b) \text{ else ClassHigh}(b))\}, [b]$$

The first says that `new_MSM()` creates mixed-sensitivity mutexes; the second that `LockMSM(a)` is freely duplicable; the third requires that when unlocking with the `unlock_MSM_L(a)` operation, that a holds a **low** value (`ClassLow(a)`); the final specification says that after locking, the (necessarily **low**) return value y correctly indicates the sensitivity of the data held in location a .

The proof sketch for this example appears in [Yan and Murray \[2021\]](#).

6.3 Implementing Verified Synchronous Channels

Message-passing concurrency, in which concurrent threads transfer data over *channels* rather than (raw) shared memory, is a common programming abstraction and has been widely studied in the context of information-flow security [[Allen 1991](#); [Honda et al. 2000](#); [Karbyshv et al. 2018](#); [Murray and Lowe 2010](#); [Roscoe et al. 1994](#); [Roscoe 1995](#); [Terauchi 2008](#); [Zdancewic and Myers 2003](#)].

One of the most common message-passing abstractions is the *synchronous channel*, e.g. as widely studied in various process calculi like CSP, CCS, and the synchronous π -calculus, and implemented in various programming languages like Go and generalised by Ada’s rendezvous mechanism.

A synchronous channel allows data to be transmitted from a sending thread to a receiving thread. Both threads block until the other is ready, which means that this type of channel also forces the sender and receiver to synchronise each time that data is transmitted on the channel.

6.3.1 Release/Acquire Synchronous Channel. We implemented in C, and verified in SecRSL in Coq, a synchronous channel abstraction which is specifically designed to support transmitting data of varying sensitivity. The (inline) `new_channel()` operation creates a new channel; `send(v, ishigh)` synchronously sends the value v on the channel, where v ’s sensitivity is given by the boolean *ishigh*; `recv(d)` takes a pointer argument d and waits to receive the next value from the channel, which is written to location d and whose sensitivity is returned as the (boolean) return-value of `recv()`.

The code for the synchronous channel appears in [Fig. 9](#). The verified `new_channel()` operation yields distinct permissions (in the form of SecRSL predicates) to send and receive on the channel. The `new_channel()` operation creates three locations: an atomic location a and two non-atomic locations b and c . Location c is used only by the receiver. The sender begins with permission to perform a release write to location a , and ownership of non-atomic location b . The receiver has permission to perform an acquire read to a and owns the non-atomic location c .

A synchronous communication on the channel involves the sender writing to b the value v to be transmitted and then performing a release write to a to encode the sensitivity of v (similarly to the message-passing program in [Fig. 1](#)). In doing so, the sender transfers ownership of b to the receiver. The sender then busywaits for the receiver to return the ownership to the sender (so that it can be used to perform subsequent sends).

The receiver does so after waiting to receive on the channel, which involves polling a via acquire reads until its value changes. At this point the receiver has acquired ownership of b to perform a release write to a . The receiver reads b to learn the value v that was transmitted, and infers v ’s sensitivity from the value it read from a . Finally it releases the ownership it got from the sender by performing a release write to a .

Thus each send/receive pair involves a two-way transfer of location b , neatly illustrating SecRSL’s power for reasoning about concurrency abstractions involving ownership transfer.

To learn when a ’s value changes, the receiver keeps a local copy in location c of the most recent value that the receiver wrote to a . The sender encodes v ’s sensitivity (*ishigh*) by incrementing a either by 2 (if *ishigh* is true) or 1 (otherwise). The receiver releases the location it got from the sender by incrementing a by 3 from its original value (before it was modified by the sender). Thus

```

new_channel()  $\stackrel{\text{def}}{=} \begin{array}{l} \text{let } b = \text{alloc}_{\text{high}}() \text{ in} \\ \text{let } c = \text{alloc}_{\text{low}}() \text{ in} \\ \text{let } a = \text{alloc}_{\text{low}}() \text{ in} \\ [b]_{\text{na}} := 0; \\ [c]_{\text{na}} := 0; \\ [a]_{\text{rel}} := 0; \end{array}$ 
```

```

send( $v$ ,  $ishigh$ )  $\stackrel{\text{def}}{=} \begin{array}{l} \text{let } x = [b]_{\text{na}} \text{ in} \\ [b]_{\text{na}} := v; \\ [a]_{\text{rel}} := x + (ishigh ? 2 : 1); \\ \text{repeat} \\ \quad \text{let } z = [a]_{\text{acq}} \text{ in} \\ \quad \quad \text{if } z == x + 3 \text{ then } 1 \text{ else } 0 \\ \text{end} \\ [b]_{\text{na}} := x + 3; \end{array}$ 
```

```

recv( $d$ )  $\stackrel{\text{def}}{=} \begin{array}{l} \text{let } t = [c]_{\text{na}} \text{ in} \\ \text{let } lv = (\text{repeat} \\ \quad \text{let } z = [a]_{\text{acq}} \text{ in} \\ \quad \quad \text{if } (z == t + 1) || (z == t + 2) \text{ then } z \text{ else } 0 \\ \quad \text{end}) \text{ in} \\ \text{let } v = [b]_{\text{na}} \text{ in} \\ [d]_{\text{na}} := v; \\ [a]_{\text{rel}} := t + 3; \\ [c]_{\text{na}} := t + 3; \\ lv - t; \end{array}$ 
```

Fig. 9. The synchronous channel implementation. Here “||” denotes boolean “or” (disjunction).

the sender owns b whenever a 's value is divisible by 3; the receiver owns b otherwise, in which case its sensitivity is determined by the (nonzero) value of $(n \bmod 3)$.

The permission to send on the channel is encoded in the $\text{Sender}(n)$ predicate, where n denotes the value currently stored in location b , the sender's counter. Likewise, $\text{Recver}(n)$ is the permission to receive, where n records the current value stored in c , the receiver's counter. Reusing the predicates $\text{ClassLow}(\ell)$ and $\text{ClassHigh}(\ell)$ from Section 6.2 to denote when (non-atomic) location ℓ holds a possibly-**high** (respectively definitely **low**) value, we prove the following SecRSL specifications, for all n such that $(n \bmod 3) = 0$:

$$\begin{array}{c}
[], \{\text{emp}\} \\
\text{new_channel}(), \text{new_channel}() \\
\{ \text{Sender}(0) \star \text{Recver}(0) \}, [b] \\
\\
[b], \{ \text{Sender}(n) \star (\text{if } ishigh \text{ then } \text{emp} \text{ else } \text{Low}(v, v')) \} \\
\text{send}(v, ishigh), \text{send}(v', ishigh) \\
\{ \text{Sender}(n + 3) \}, [b] \\
\\
[b, d], \{ \text{Recver}(n) \star (\text{ClassHigh}(d)) \} \\
\text{recv}(d), \text{recv}(d) \\
\{(y, y') . \text{Recver}(n + 3) \star (\text{if } y = 1 \text{ then } \text{ClassLow}(d) \text{ else } \text{ClassHigh}(d)) \star \dots \\
\wedge (y = y') \wedge (y = 1 \vee y = 2)\}, [b, d] \\
\\
\text{ClassLow}(\ell) \stackrel{\text{def}}{=} \exists(v, v'), (\ell, \ell) \mapsto (v, v') \star \text{Low}(v, v') \\
\text{ClassHigh}(\ell) \stackrel{\text{def}}{=} \exists(v, v'), (\ell, \ell) \mapsto (v, v')
\end{array}$$

These say that `new_channel()` yields the permissions to send and receive, initialising the locations to 0. To call `send(v, ishigh)`, we must have the sender permission and *v*'s classification must match *ishigh*. After sending, the permission to subsequently send is returned to allow the sender to send repeatedly. The specification for `recv()` has a similar structure and additionally requires the location *d*, where the value received will be written, is valid. The “...” in the postcondition for `recv()` elides additional permissions that the logic tracks: namely the now-useless permission for `recv()` to read the alternative value (*n* + 1 or *n* + 2) from *a*, other than the value it *did* read and that will now never be written to *a* (see Yan and Murray [2021]). Of course one can always safely ignore the additional permission in proofs by employing the frame rule.

The proof sketch appears in Yan and Murray [2021].

6.3.2 Release/CAS Synchronous Channel. To showcase SecRSL's flexibility, we also implemented and verified an alternative synchronous channel implementation. Rather than using release/acquire synchronisation, this implementation instead uses release/CAS pairs. Doing so avoids the need for the counter *n* and leads to a simpler proof, albeit with a similar structure to the original. This implementation we also verified in Coq, against the following specifications, where Sender and Recver denote the permissions to send and receive respectively:

$$\begin{aligned}
 & [], \{ \text{emp} \} \\
 & \text{new_channel_CAS}(), \text{new_channel_CAS}() \\
 & \quad \{ \text{Sender} \star \text{Recver} \}, [b] \\
 \\
 & [b], \{ \text{Sender} \star (\text{if } \text{ishigh} \text{ then emp else Low}(v, v')) \} \\
 & \quad \text{send_CAS}(v, \text{ishigh}), \text{send_CAS}(v', \text{ishigh}) \\
 & \quad \{ \text{Sender} \}, [b] \\
 \\
 & [b, d], \{ \text{Recver} \star (\text{ClassHigh}(d)) \} \\
 & \quad \text{recv_CAS}(d), \text{recv_CAS}(d) \\
 & \{(y, y'). \text{Recver} \star (\text{if } y = 1 \text{ then ClassLow}(d) \text{ else ClassHigh}(d)) \\
 & \quad \wedge (y = y') \wedge (y = 1 \vee y = 2)\}, [b, d]
 \end{aligned}$$

6.3.3 Performance Comparison. We implemented both verified channel implementations in C, and performed a rudimentary comparison of their performance. As a baseline we implemented a sequentially-consistent version of the synchronous channel, shown in Yan and Murray [2021]. We constructed this implementation by taking the one from Fig. 9 and marking all shared variable accesses with the `sc` mode, in order to make its execution sequentially consistent (i.e. insulate it from weak memory effects). While this implementation is not verified, we posit that it might be verifiable using a suitable adaptation of the fine-grained security separation logic of Frumin et al. [2021].

The results are summarised in Table 1. We report maximum and minimum average bandwidths observed when transferring 1 GiB (1024 MiB) of random data over four trials, across various platforms. We also report the average number of rounds per second observed across all trials for each platform. The results include a version of the release/acquire channel implementation that transmits 8-byte values (long ints); all other implementations transmit 4-byte values (ints).

The release/acquire implementation consistently outperforms the others on both Intel platforms, with observed throughput improvements of up to 88%. Naturally doubling the size of the quantity transferred in each round tends to double the observed bandwidth. On ARMv8 AArch64, `acq` loads and `sc` loads both map to the LDAR Load-Acquire instruction; likewise `rel` and `sc` stores both map to the STLR Store-Release instruction [Ševčík and Sewell 2016]. We conjecture that this is why

Table 1. Performance comparison of the verified synchronous channel implementations. Platforms: *Intel MacBook*: MacBook 13-inch, 2020, 2 GHz Quad-Core Intel Core i5, MacOS. *Intel Server*: 8 CPU Xeon(R) Gold 6248 @ 2.50GHz, Linux x86_64 on VMWare hypervisor with full visualisation. *ARM Server*: 2 CPU Neoverse-N1, Linux ARM64 (aarch64), Amazon AWS EC2 t4g.micro instance.

Implementation	Observed Bandwidth MiB/sec (avg. rounds/sec)		
	Intel MacBook	Intel Server	ARM Server
Rel/Acq (4 bytes)	31–33 (8.3×10^6)	17–19 (4.7×10^6)	13–19 (4.1×10^6)
Rel/Acq (8 bytes)	60–65 (8.3×10^6)	33–36 (4.6×10^6)	23–26 (3.2×10^6)
Rel/CAS (4 bytes)	15–18 (4.2×10^6)	13–15 (3.7×10^6)	9–12 (2.7×10^6)
SC (4 bytes)	17–18 (4.8×10^6)	9.2–9.6 (2.5×10^6)	14–19 (4.2×10^6)

the release/acquire version was observed to perform no better on the ARM platform than the sequentially-consistent baseline. Indeed, the latter was observed to perform slightly better than the former; however, the results on this platform show considerable variability and this observed difference is well within the noise.

These empirical results demonstrate that, depending on the deployment platform, significant performance improvements can be obtained by utilising the C11 weak memory primitives, backed by the formal guarantees afforded by SecRSL. They clearly demonstrate the power of SecRSL over prior logics like SecCSL [Ernst and Murray 2019] which can reason only about data-race free programs, or even more recent logics [Frumin et al. 2021] that implicitly assume sequential consistency.

7 FURTHER RELATED WORK

Prior security logics and type systems for information-flow security on weak memory models include that of Vaughan and Millstein [2012] who developed a simple security type system for an operational semantics of the TSO memory model, and Mantel et al. [2014] who developed a transforming type system for ensuring security also on PSO and the IBM370. As security type systems, neither supported the precision afforded by a logic like SecRSL.

Smith et al. [2019] present a program logic for proving secure information flow of ARMv8. Unlike SecRSL which targets C11, they target a low level memory model. Like SecRSL, their logic supports reasoning about value-dependent classification. However, unlike SecRSL, theirs does not support local reasoning with ownership transfer and invariants. As Fig. 3 demonstrates, such support is vital for practical reasoning about expressive security policies.

Our security definitions over the C11 axiomatic memory model are an instance of noninterference [Goguen and Meseguer 1982] for a so-called “true concurrency” semantics. It would be interesting to compare how our definitions relate to those for other true concurrency models, e.g. those for Petri Nets [Baldan and Beggiano 2018; Baldan and Carraro 2014].

8 CONCLUSION

We presented SecRSL, a security separation logic for C11 Release-Acquire concurrency. SecRSL inherits RSL’s virtues of compositional, local reasoning about Release-Acquire atomics, plus SecCSL’s ability to reason about expressive security policies like value-dependent classification.

We also presented the first definition of information-flow security for an axiomatic weak memory model, against which we proved SecRSL sound. SecRSL ensures that programs satisfy a constant-time security guarantee, while being free of undefined behaviour.

We demonstrated SecRSL by using it to implement and verify the functional correctness and security of various concurrency primitives, including a spinlock module, a mixed-sensitivity mutex, and two synchronous channel implementations. Benchmarking the latter against an unverified sequentially-consistent implementation showed that SecRSL can enable significant performance gains.

Beyond these examples, we also believe (as Section 5.4 argues) these ideas are of interest—and can be readily applied—beyond the Release-Acquire fragment of C11.

9 DATA AVAILABILITY STATEMENT

The artifact [Yan 2021] accompanying this paper contains a Coq formalisation of SecRSL, its soundness proof, as well as example applications of the logic mentioned throughout this paper.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their insightful feedback on earlier drafts of this paper. This material is based upon work supported by the Commonwealth of Australia Defence Science and Technology Group, Next Generation Technologies Fund (NGTF).

REFERENCES

- Jade Alglave, Anthony Fox, Samin Ishtiaq, Magnus O Myreen, Susmit Sarkar, Peter Sewell, and Francesco Zappa Nardelli. 2009. The semantics of Power and ARM multiprocessor machine code. In *Proceedings of the 4th workshop on Declarative aspects of multicore programming*. 13–24.
- PG Allen. 1991. A comparison of non-interference and non-deducibility using CSP. In *IEEE Computer Security Foundations Workshop (CSFW)*. IEEE, 43–54.
- Paolo Baldan and Alessandro Beggiato. 2018. Multilevel transitive and intransitive non-interference, causally. *Theoretical Computer Science* 706 (2018), 54–82.
- Paolo Baldan and Alberto Carraro. 2014. Non-interference by unfolding. In *International Conference on Applications and Theory of Petri Nets and Concurrency*. Springer, 190–209.
- Gilles Barthe, Sandrine Blazy, Benjamin Grégoire, Rémi Hutin, Vincent Laporte, David Pichardie, and Alix Trieu. 2019. Formal verification of a constant-time preserving C compiler. *Proceedings of the ACM on Programming Languages* 4, POPL (2019), 1–30.
- Gilles Barthe, Thomas Espitau, Benjamin Grégoire, Justin Hsu, and Pierre-Yves Strub. 2017. Proving expected sensitivity of probabilistic programs. *Proceedings of the ACM on Programming Languages* 2, POPL (2017), 1–29.
- Mark Batty, Alastair F Donaldson, and John Wickerson. 2016. Overhauling SC atomics in C11 and OpenCL. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. 634–648.
- Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber. 2011. Mathematizing C++ concurrency. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. 55–66.
- Mark John Batty. 2014. *The C11 and C++11 Concurrency Model*. Ph.D. Dissertation. University of Cambridge.
- Nick Benton. 2004. Simple relational correctness proofs for static analyses and program transformations. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*. 14–25.
- Marko Doko and Viktor Vafeiadis. 2016. A program logic for C11 memory fences. In *International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI)*. Springer, 413–430.
- Marko Doko and Viktor Vafeiadis. 2017. Tackling real-life relaxed concurrency with FSL++. In *European Symposium on Programming (ESOP)*. Springer, 448–475.
- Gidon Ernst and Toby Murray. 2019. SecCSL: Security Concurrent Separation Logic. In *International Conference on Computer Aided Verification (CAV)*. 208–230.
- Cormac Flanagan, Amr Sabry, Bruce F Duba, and Matthias Felleisen. 1993. The essence of compiling with continuations. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. 237–247.
- Dan Frumin, Robbert Krebbers, and Lars Birkedal. 2021. Compositional Non-Interference for Fine-Grained Concurrent Programs. In *IEEE Symposium on Security & Privacy (S&P)*. To appear.
- Joseph Goguen and José Meseguer. 1982. Security Policies and Security Models. In *IEEE Symposium on Security & Privacy (S&P)*. IEEE Computer Society, Oakland, California, USA, 11–20.
- Kohei Honda, Vasco Vasconcelos, and Nobuko Yoshida. 2000. Secure information flow as typed process behaviour. In *European Symposium on Programming (ESOP)*. Springer, 180–199.

- Aleksandr Karbyshev, Kasper Svendsen, Aslan Askarov, and Lars Birkedal. 2018. Compositional Non-Interference for Concurrent Programs via Separation and Framing. In *International Conference on Principles of Security and Trust (POST)*.
- Luisa Lourenço and Luís Caires. 2015. Dependent Information Flow Types. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, Mumbai, India, 317–328.
- Sela Mador-Haim, Luc Maranget, Susmit Sarkar, Kayvan Memarian, Jade Alglave, Scott Owens, Rajeev Alur, Milo MK Martin, Peter Sewell, and Derek Williams. 2012. An axiomatic memory model for POWER multiprocessors. In *International Conference on Computer Aided Verification (CAV)*. 495–512.
- Kenji Maillard, Cătălin Hrițcu, Exequiel Rivas, and Antoine Van Muylder. 2019. The next 700 relational program logics. *Proceedings of the ACM on Programming Languages* 4, POPL (2019), 1–33.
- Jeremy Manson, William Pugh, and Sarita V Adve. 2005. The Java memory model. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, 378–391.
- Heiko Mantel, Matthias Perner, and Jens Sauer. 2014. Noninterference under weak memory models. In *IEEE Computer Security Foundations Symposium (CSF)*. IEEE, 80–94.
- Toby Murray and Gavin Lowe. 2009. On Refinement-Closed Security Properties and Nondeterministic Compositions. In *International Workshop on Automated Verification of Critical Systems (Electronic Notes in Theoretical Computer Science, Vol. 250)*. 49–68. <https://doi.org/10.1016/j.entcs.2009.08.017>
- Toby Murray and Gavin Lowe. 2010. Analysing the Information Flow Properties of Object-Capability Patterns. In *Formal Aspects of Security and Trust (Lecture Notes in Computer Science, Vol. 5983)*. Eindhoven, The Netherlands, 81–95. https://doi.org/10.1007/978-3-642-12459-4_7
- Toby Murray, Robert Sison, and Kai Engelhardt. 2018. COVERN: A Logic for Compositional Verification of Information Flow Control. In *IEEE European Symposium on Security and Privacy (EuroS&P)*. London, United Kingdom.
- Toby Murray, Robert Sison, Edward Pierzchalski, and Christine Rizkallah. 2016. Compositional Verification and Refinement of Concurrent Value-Dependent Noninterference. In *IEEE Computer Security Foundations Symposium (CSF)*. 417–431.
- Peter W O’Hearn. 2004. Resources, concurrency and local reasoning. In *International Conference on Concurrency Theory (CONCUR)*. Springer, 49–67.
- AW Roscoe, JCP Woodcock, and Lars Wulf. 1994. Non-interference through determinism. In *European Symposium on Research in Computer Security (ESORICS)*. Springer, 31–53.
- A William Roscoe. 1995. CSP and determinism in security modelling. In *IEEE Symposium on Security & Privacy (S&P)*. IEEE, 114–127.
- Susmit Sarkar, Peter Sewell, Francesco Zappa Nardelli, Scott Owens, Tom Ridge, Thomas Braibant, Magnus O Myreen, and Jade Alglave. 2009. The semantics of x86-CC multiprocessor machine code. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*.
- Daniel Schoepe, Toby Murray, and Andrei Sabelfeld. 2020. VERONICA: Expressive and Precise Concurrent Information Flow Security. In *IEEE Computer Security Foundations Symposium (CSF)*. IEEE, 79–94.
- Robert Sison and Toby Murray. 2019. Verifying That a Compiler Preserves Concurrent Value-Dependent Information-Flow Security. In *International Conference on Interactive Theorem Proving (ITP)*. 27:1–27:19.
- Graeme Smith, Nicholas Coughlin, and Toby Murray. 2019. Value-Dependent Information-Flow Security on Weak Memory Models. In *International Symposium on Formal Methods (FM)*. Springer, 539–555.
- Tachio Terauchi. 2008. A type system for observational determinism. In *IEEE Computer Security Foundations Symposium (CSF)*. IEEE, 287–300.
- Viktor Vafeiadis. 2011. Concurrent Separation Logic and Operational Semantics. In *Mathematical Foundations of Programming Semantics (MFPS)*. 335–351.
- Viktor Vafeiadis and Chinmay Narayan. 2013. Relaxed separation logic: A program logic for C11 concurrency. In *Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA)*. 867–884.
- Jeffrey A Vaughan and Todd Millstein. 2012. Secure information flow for concurrent programs under Total Store Order. In *IEEE Computer Security Foundations Symposium (CSF)*. IEEE, 19–29.
- Jaroslav Ševčík and Peter Sewell. 2016. C/C++11 mappings to processors. <https://www.cl.cam.ac.uk/~pes20/cpp/cpp0xmappings.html>. Accessed 2021-08-12.
- Pengbo Yan. 2021. *SecRSL: Security Separation Logic for C11 Release-Acquire Concurrency - Coq Formalisation*. <https://doi.org/10.5281/zenodo.5493554>
- Pengbo Yan and Toby Murray. 2021. SecRSL: Security Separation Logic for C11 Release-Acquire Concurrency (Extended version with technical appendices). arXiv:2109.03602 [cs.CR]
- Hongseok Yang. 2007. Relational separation logic. *Theoretical Computer Science* 375, 1-3 (2007), 308–334.
- Steve Zdancewic and Andrew C Myers. 2003. Observational determinism for concurrent program security. In *IEEE Computer Security Foundations Workshop (CSFW)*. IEEE, 29–43.