

SECCSL: Security Concurrent Separation Logic

Gidon Ernst and Toby Murray
{gidon.ernst,toby.murray}@unimelb.edu.au

University of Melbourne, Australia

Abstract. We present SECCSL, a concurrent separation logic for proving expressive, data-dependent information flow security properties of low-level programs. SECCSL is considerably more expressive, while being far simpler, than recent compositional information flow logics that cannot reason about pointers, arrays etc.

SECCSL inherits the structure of traditional concurrent separation logics by adopting a *relational* semantics for its assertions. Thus SECCSL reasoning can be automated via symbolic execution, in the style of contemporary verifiers like VeriFast, Viper, and Infer. We implement a prototype verification tool SECC that supports a subset of C. In doing so, we demonstrate for the first time the feasibility of automated reasoning about expressive information flow security policies over low-level code.

1 Introduction

Software verification successes abound, whether via interactive proof or via automatic program verifiers. While the former has yielded individual, deeply verified software artifacts [22,19,23] primarily by *researchers*, the latter appears to be having a growing impact on *industrial* software engineering [34,11,36].

At the same time, recent work has heralded major advancements in program logics for reasoning about secure *information flow* [32,21,31]—i.e. whether programs properly protect their secrets—yielding the first general program logics and proofs of information flow security for non-trivial concurrent programs [32]. Yet so far, such logics have remained confined to interactive proof assistants, making them practically inaccessible to industrial developers.

This is not especially surprising: the price that logics like [32] have tended to pay for their generality, especially with regards to support for expressive security policies, is complexity. Complexity naturally hinders automating the application of these logics. Worse, these logics reason only over very simple toy programming languages, which lack even pointers, arrays, and structures, for instance. Their complexity, we argue, makes scaling up these logics to real-world languages impractical. How, therefore, can we leverage the power of existing automatic deductive verification approaches for security proofs?

Doing so requires the development of a new logic for compositionally proving expressive security properties of concurrent programs that is radically simpler, more powerful, and straightforward to automate, than state-of-the-art logics. In this paper we present “Security Concurrent Separation Logic” (SECCSL),

SECCSL achieves this unprecedented combination of simplicity, power, and ease of automation by capturing core concepts of modern information flow logics, such as value-dependent variable sensitivity [47,25,29] and shared data invariants [32], in the familiar style of a concurrent separation logic (CSL) [38].

Prior work [13,18] has noted the promise of separation logic for reasoning about information flow yet, to date, that promise remains unrealised. Indeed, the only two prior encodings of information flow concepts into separation logics of which we are aware have overlooked crucial features like concurrency [13], or support for specifying strong security properties like (as we explain in Section 2) the ability to separately specify the sensitivity of *values* and memory *locations*. Some [18] lack soundness arguments altogether while others [13] fail to satisfy basic properties needed to feasibly automate them (see the discussion following Proposition 1 later). Designing a logic therefore with the right combination of features, with the right semantics, is non-trivial.

To manage this, SECCSL assertions have a *relational* interpretation [7,46] over a standard heap model. This allows them to canonically encode information flow concepts while maintaining the structure of traditional CSLs. We have mechanised the soundness of SECCSL in Isabelle/HOL [35]. To do so we devised a new, compositional information flow security property that, like SECCSL itself, is simpler and more powerful than its predecessors [32].

To demonstrate SECCSL’s capacity for automated application, we implemented the prototype tool SECC. SECC is the first automatic verifier for expressive information flow security for a subset of the C language. We target C because it dominates low-level security-critical code. SECC automates SECCSL reasoning via symbolic execution, in the style of contemporary verification tools like VeriFast [20], Viper [28], and Infer [10].

We present an overview of SECCSL’s main ideas in Section 2, and then the details in Section 3, before its soundness proof in Section 4. We discuss SECC in Section 5. Our Isabelle theories, the prototype tool SECC, and examples are available anonymously online [1]¹ and as attachment of the submission.

2 An Overview of SECCSL

2.1 Specifying Information Flow Control in SECCSL

Consider the program in Fig. 1. It maintains a global pointer `rec` to a shared record, protected by the lock `mutex`. The `is_classified` field of the record identifies the confidentiality of the record’s `data`: when `is_classified` is true, the value stored in the `data` field is confidential, and otherwise it is safe to release publicly. The left thread outputs the data in the record whenever it is public by writing to the (memory mapped) output device register pointer `OUTPUT_REG` (here also protected by `mutex`). The right thread updates the record, ensuring its contents is not confidential, here by clearing its `data`.

¹ <https://anonymous.4open.science/repository/42ff2a21-f974-476f-9e7d-980f9b63ff92>

```

/* globals shared between the two threads */
struct record { bool is_classified; int data; };
struct record * rec = /* ... initialisation omitted ... */;
volatile int * const OUTPUT_REG = /* memory-mapped IO device register */;

/* thread 1: output the record */
while(TRUE) {
    lock(mutex);
    if (!rec->is_classified)
        *OUTPUT_REG = rec->data;
    unlock(mutex);
}

/* thread 2: edit the record */
lock(mutex);
/* clear the record */
rec->is_classified = FALSE;
rec->data = 0;
unlock(mutex);

```

Fig. 1. Example of Concurrent Information Flow.

Suppose writing to the `OUTPUT_REG` register causes the value being written to be outputted to a publicly-visible location. Reasoning, then, that the example is secure requires capturing that (1) the `data` field of the record pointed to by `rec` is confidential precisely when the record's `is_classified` field says it is, and (2) `OUTPUT_REG` should never have confidential data written to it. Therefore the example only ever writes non-confidential data into `OUTPUT_REG`.

Condition (1) specifies the sensitivity of a data *value* in memory, whereas condition (2) specifies the sensitivity of the data that a memory *location* is permitted to hold. Prior security separation logics [13,18] reason only about value-sensitivity (condition (2)) but, as we explain below, both are needed. Like those prior logics, in SECCSL one specifies the sensitivity of the value denoted by an expression e via a security *label* ℓ : the assertion $e::\ell$ means that the sensitivity of the value denoted by expression e is at most ℓ . Security labels are drawn from a lattice with top element `high` (denoting the most confidential information), bottom element `low` (denoting public information), and ordered via \sqsubseteq : $\ell \sqsubseteq \ell'$ means that information labelled with ℓ' is at least as sensitive as that labelled by ℓ . Using this style of assertion, in conjunction with standard separation logic connectives (explained below), condition (1) can be specified as follows:

$$\exists c d. \text{rec} \mapsto (c, d) \wedge c::\text{low} \wedge d::(c ? \text{high} : \text{low}) \quad (1)$$

The \mapsto is the standard separation logic points-to connective: $e \mapsto e'$ means the memory location denoted by expression e holds the value denoted by e' . Thus (1) can be read as saying that the `rec` pointer points to a pair of values (c, d) . The first c (the value of the `is_classified` field) is public. The sensitivity of the second d (the value of the `data` field) is given by the value of the first c : it is `high` when c is true and is `low` otherwise. SECCSL makes reasoning in the presence of this kind of *value-dependent* sensitivity [47,25,29] considerably simpler than in prior logics for concurrent information flow [33,32], while supporting reasoning about pointers, structures and arrays at the same time.

To allow reasoning about location sensitivity, SECCSL augments the separation logic points-to predicate with a security label ℓ : $e \overset{\ell}{\mapsto} e'$ means that the value denoted by the expression e' is present in memory at the location denoted by e , and additionally that at all times the sensitivity of the value stored in that locations is never allowed to exceed ℓ . Thus in SECCSL, $e \mapsto e'$ abbreviates $e \overset{\text{high}}{\mapsto} e'$. In Fig. 1 the sensitivity of `OUTPUT_REG` can be specified as:

$$\exists v. \text{OUTPUT_REG} \overset{\text{low}}{\mapsto} v \quad (2)$$

2.2 Reasoning in SECCSL

SECCSL judgements have the form:

$$\ell_A \vdash \{P\} c \{Q\} \quad (3)$$

Here ℓ_A is the *attacker security level*, P is the program's precondition, c is the (concurrent) program command being executed, and Q is the program's postcondition. We defer a description of ℓ_A and the security property the judgement implies until Section 2.3.

As with traditional Hoare logics, intuitively the judgement means that if the program c begins in a state satisfying its precondition P then, when it terminates, the final state will satisfy its postcondition Q . As with traditional separation logics [41], the judgement also implies that the program's execution cannot abort (by e.g., dereferencing an invalid pointer).

As with traditional CSLs, SECCSL is geared towards reasoning over shared-memory programs that use lock-based synchronisation. Extending it to handle racy data accesses and relaxed memory models should be possible [44].

Each lock l has an associated invariant $\text{inv}(l)$, which is simply a predicate, like P or Q in (3), that describes the shared memory that the lock protects. In Fig. 1, where the lock `mutex` protects the shared pointer `rec` and `OUTPUT_REG`, the associated invariant $\text{inv}(\text{mutex})$ is simply the conjunction of (1) and (2).

$$(\exists c d. \text{rec} \mapsto (c, d) \wedge c :: \text{low} \wedge d :: (c ? \text{high} : \text{low})) \star (\exists v. \text{OUTPUT_REG} \overset{\text{low}}{\mapsto} v) \quad (4)$$

Here, the \star is *separating conjunction* from standard separation logic: $P \star Q$ asserts that the assertions P and Q both hold and, additionally, that the memory locations referenced by P and Q respectively do not overlap.

Thus SECCSL invariants, like SECCSL predicates, describe together both functional properties (e.g. `rec` is a valid pointer) and security concerns (e.g. the `OUTPUT_REG` location is publicly observable). Prior logics [32] for concurrent information flow have kept such concerns separate. SECCSL unifies them neatly together, increasing simplicity and expressiveness.

As with traditional CSLs, when acquiring a lock one gets to assume that the lock's invariant holds. Subsequently, when releasing the lock one must prove that the invariant has been re-established. For example, when reasoning about the code of the left-thread in Fig. 1, upon acquiring the `mutex`, SECCSL assumes (4)

holds, which allows proving that the loop body is secure. When reasoning about the right thread, one must prove that the invariant has been re-established when it releases the mutex. This is the reason e.g. that the right thread must clear the `data` field after setting `is_classified` to false.

Reasoning in SECCSL proceeds forward over the program text. When execution forks, as in Fig. 1, one reasons over each thread individually. For Fig. 1, SECCSL requires (see the IF rule of Fig. 3) proving that the guard of the if-condition is `low`, i.e. that the program is not branching on a secret (a potential timing channel—see Section 2.3 below). This follows straightforwardly from the invariant (4) (specifically the part “ $c :: \text{low}$ ”), which SECCSL assumes upon the left thread acquiring `mutex`. Secondly, when the write to `OUTPUT_REG` is performed, SECCSL requires (STORE rule of Fig. 3) that the expression that is being written to the location `OUTPUT_REG` has sensitivity `low`. This follows straightforwardly from the invariant and the guard of the if-statement. Finally, when the right thread releases `mutex`, it must have re-established the invariant (UNLOCK rule of Fig. 3), which here follows trivially.

2.3 Security Intuition and Informal Security Property

But what does security mean in SECCSL? Indeed, the SECCSL judgement (3) additionally implies that the program c does not leak any sensitive information during its execution to potential attackers.

The attacker security level ℓ_A in (3) represents an upper bound on the parts of the program’s memory that a potential, passive attacker is assumed to be able to observe before, during and after its execution. Intuitively this is all memory locations whose sensitivity is $\sqsubseteq \ell_A$. Which memory locations have sensitivity $\sqsubseteq \ell_A$ is defined by the precondition P and the lock invariants. Intuitively (see Definition 1 in Section 4), a memory location loc is visible to the ℓ_A attacker iff P or a lock invariant contains a judgement $e \xrightarrow{e_l} e'$ and in the program’s initial state e evaluates to loc and e_l evaluates to some $\ell \sqsubseteq \ell_A$.

Which data is sensitive and should not be leaked to the ℓ_A attacker is defined similarly by P and the lock invariants: an expression e is sensitive when P or a lock invariant contains a judgement $e :: e_l$ and in the program’s initial state e_l evaluates to some $\ell \not\sqsubseteq \ell_A$. Security, then, requires that in all intermediate states of the program’s execution no sensitive data can be inferred via the attacker-observable memory.

SECCSL proves a *compositional* security property that formalises this intuition. Since the property needs to be compositional with regards to concurrent execution, the resulting security property is *timing sensitive*, meaning that not only must the program never reveal sensitive data into attacker-observable memory locations but the times at which it updates these memory locations cannot depend on sensitive data. It is well-known that timing-insensitive security properties are not compositional under standard scheduling models [45,32]. For this reason SECCSL forbids programs from branching on sensitive values. This restriction could in principle be relaxed in the future via established techniques [27,26].

SECCSL’s top-level soundness theorem (Theorem 2) formalises the above intuitive definition of security in the style of a traditional *noninterference* [17] property that compares two program executions. To define security formally, SECCSL adopts a *relational* interpretation for the assertions P and Q (and the lock invariants) in which they are evaluated against pairs of execution states. Doing so allows SECCSL to retain the flavour and structure of a traditional concurrent separation logic without loss of power for reasoning about security.

In particular, under this relational semantics (see Fig. 2 described later in Section 3.1), when a pair of states satisfies an assertion P , it implies that the two states agree on the values of all non-sensitive expressions as defined by P . The formal noninterference security property is then roughly as follows (see Theorem 2 in Section 4): program c with precondition P is secure against the ℓ_A -attacker if, whenever executed twice from two initial states both satisfying P and the lock invariants (and so agreeing on the values of all data assumed to be initially observable to the ℓ_A attacker), in all intermediate pairs of states arrived at after running each execution for the same number of steps, the resulting states agree on all ℓ_A -visible memory locations as defined in precondition P and the initial lock invariants. This definition is timing sensitive because it compares executions that have executed the same number of steps.

3 The Logic SECCSL

3.1 Assertions

Pure expressions e that do not depend on the heap are composed of variables x , function applications, equations, and conditional expressions.

$$e ::= x \mid f(e_1, \dots, e_n) \mid e_1 = e_2 \mid \phi ? e_1 : e_2,$$

We assume a standard first-order many sorted typing discipline (not elaborated). We include pure formulas as boolean expressions $\phi: \mathbf{Bool}$, so that logical connectives are among the function symbols f , using standard notation.

We postulate that the logical signature contains a designated sort \mathbf{Label} , corresponding to the security label lattice, together with constants $\mathbf{low}, \mathbf{high}: \mathbf{Label}$ and binary predicate symbol $\sqsubseteq: \mathbf{Label} \times \mathbf{Label} \rightarrow \mathbf{Bool}$, whose interpretation $\sqsubseteq^{\mathcal{A}}$ satisfies the lattice axioms.

SECCSL’s assertions P, Q comprise pure boolean expressions ϕ , value sensitivity assertions, labelled points-to predicates (heap location sensitivity assertions), assertions guarded by (pure) conditionals, ordinary and separating conjunction, the empty heap, and existential quantification.

$$P ::= \phi \mid e :: e_l \mid e_p \xrightarrow{e_l} e_v \mid (\phi ? P : Q) \mid P \wedge Q \mid \mathbf{emp} \mid P \star Q \mid \exists x. P$$

Disjunction, negation, and implication are excluded because they cause issues for describing the set of ℓ -visible heap location to the ℓ -attacker, similarly to the problem of defining heap footprints for non-precise assertions [37,39,24]. Of course, these connectives can still occur inside pure expressions.

The standard expression semantics $\llbracket e \rrbracket_s$ evaluates e over a store s , which assigns values to variables x as $s(x)$. The interpretation $f^{\mathcal{A}}$ of a function symbol f is a function, given statically by a logical structure \mathcal{A} , which includes booleans $\mathbb{B} = \{true, false\}$ and assigns standard meaning to boolean connectives and if-then-else. We write $s \models \phi$ if $\llbracket \phi \rrbracket_s = true$.

The relational semantics of assertions, written $(s, h), (s', h') \models_{\ell} P$, is defined in Fig. 2 over two states (s, h) and (s', h') each consisting of a store and a heap. The semantics is defined against a the attacker security level ℓ (see Section 2.3).

Stores s and s' are related via $e :: e_l$. We require the expression e_l denoting the sensitivity to coincide on s and s' and whenever $\llbracket e_l \rrbracket_s \sqsubseteq^{\mathcal{A}} \ell$ holds, e must evaluate to the same value both states, (7). Heaps are related by $(s, h), (s', h') \models_{\ell} e_p \xrightarrow{e_l} e_v$, which similarly ensures that the two heap fragments are identical $h = h'$ when e_l says so, (8). Conditional assertions $\phi ? P : Q$ evaluate to P when ϕ holds for *both* stores, and to Q otherwise. The separating conjunction splits both heaps independently, (11). Similarly, the existential quantifier picks two values v and v' , (12). Whether parts of the split (resp. these two values) actually agree will depend on other assertions made.

Low Locations To capture strong security properties, we require a declarative specification of which heap locations are considered visible to the ℓ -attacker, when assertion P holds in some (initial) state (see Section 2.3). We denote this set $\text{low}_{\ell}(P, s)$, for initial store s . It is essentially a subset of P 's footprint.

Definition 1 (Low Locations).

$$\begin{aligned} \text{low}_{\ell}(\phi, s) &= \text{low}_{\ell}(e :: e_l, s) = \emptyset \\ \text{low}_{\ell}(e_p \xrightarrow{e_l} e_v, s) &= \begin{cases} \{\llbracket e_p \rrbracket_s\}, & \llbracket e_l \rrbracket_s \sqsubseteq^{\mathcal{A}} \ell \\ \emptyset, & \text{otherwise} \end{cases} \\ \text{low}_{\ell}(b ? P : Q, s) &= \begin{cases} \text{low}_{\ell}(P, s), & s \models b \\ \text{low}_{\ell}(Q, s), & \text{otherwise} \end{cases} \\ \text{low}_{\ell}(P \star Q, s) &= \text{low}_{\ell}(P \wedge Q, s) = \text{low}_{\ell}(P, s) \cup \text{low}_{\ell}(Q, s) \\ \text{low}_{\ell}(\exists x. P, s) &= \begin{cases} \text{low}_{\ell}(P, s), & \forall s v. \text{low}_{\ell}(P, s) = \text{low}_{\ell}(P, s(x \mapsto v)) \\ \emptyset, & \text{otherwise} \end{cases} \end{aligned}$$

Note that, by design, the definition does not give a useful result for existential like $\exists p v. p \xrightarrow{\text{low}} v$. This mirrors the usual difficulty of defining footprints for separation logic assertions with respect to non-precise assertions [37,39,24]. This restriction is not an issue in practice, as useful security properties must necessarily describe the attacker-visible portion of memory precisely.

3.2 Entailments

Although implication has been omitted from the assertion language, entailments $P \xRightarrow{\ell} Q$ between assertions still play a role in SECCSL via its Hoare logic

Using the abbreviation $s, h \models_{\ell} e_p \mapsto e_v \iff h = \{ \llbracket e_p \rrbracket_s \mapsto \llbracket e_v \rrbracket_s \}$

$$(s, h), (s', h') \models_{\ell} \mathbf{emp} \iff h = h' = \emptyset \quad (5)$$

$$(s, h), (s', h') \models_{\ell} \phi \iff s \models \phi \text{ and } s' \models \phi \quad (6)$$

$$(s, h), (s', h') \models_{\ell} e :: e_l \quad (7)$$

$$\iff \llbracket e_l \rrbracket_s = \llbracket e_l \rrbracket_{s'} \text{ and } \left(\llbracket e_l \rrbracket_s \sqsubseteq^A \ell \implies \llbracket e \rrbracket_s = \llbracket e \rrbracket_{s'} \right)$$

$$(s, h), (s', h') \models_{\ell} e_p \xrightarrow{e_l} e_v \quad (8)$$

$$\iff s, h \models_{\ell} e_p \mapsto e_v \text{ and } s', h' \models_{\ell} e_p \mapsto e_v \text{ and } (s, h), (s', h') \models_{\ell} e_p :: e_l \wedge e_v :: e_l$$

$$(s, h), (s', h') \models_{\ell} \phi ? P : Q \quad (9)$$

$$\iff \begin{cases} (s, h), (s', h') \models_{\ell} P, & \text{if } s \models \phi \text{ and } s' \models \phi \\ (s, h), (s', h') \models_{\ell} Q, & \text{otherwise} \end{cases}$$

$$(s, h), (s', h') \models_{\ell} P \wedge Q \quad (10)$$

$$\iff (s, h), (s', h') \models_{\ell} P \text{ and } (s, h), (s', h') \models_{\ell} Q$$

$$(s, h), (s', h') \models_{\ell} P \star Q \quad (11)$$

$$\iff \begin{aligned} &\text{there are disjoint sub-heaps } h_1, h_2 \text{ and } h'_1, h'_2 \\ &\text{with } h = h_1 \uplus h_2 \text{ and } h' = h'_1 \uplus h'_2 \\ &\text{such that } (s, h_1), (s', h'_1) \models_{\ell} P_1 \text{ and } (s, h_2), (s', h'_2) \models_{\ell} P_2 \end{aligned}$$

$$(s, h), (s', h') \models_{\ell} \exists x. P \quad (12)$$

$$\iff \begin{aligned} &\text{there are values } v, v' \\ &\text{such that } (s(x \mapsto v), h), (s'(x \mapsto v'), h') \models P \end{aligned}$$

Fig. 2. Relational semantics of assertions.

style consequence rule (CONSEQ in Fig. 3). We discuss entailment now as it sheds useful light on some consequences of SECCSL's relational semantics.

Definition 2 (Secure Entailment). $P \xRightarrow{\ell} Q$ holds iff

- $(s, h), (s', h') \models_{\ell} P$ implies $(s, h), (s', h') \models_{\ell} Q$ for all s, h and s', h'
- $\text{low}_{s_{\ell}}(P, s) \subseteq \text{low}_{s_{\ell}}(Q, s)$ for all s

The security level ℓ is used not just in the evaluation of the assertions but also to preserve the ℓ -attacker visible locations of P in Q . This reflects the intuition that P is stronger than Q , and so Q should make fewer assumptions than P about the attacker's observational powers.

Proposition 1.

$$e = e' \wedge e_l = e_l' \wedge e :: e_l \xRightarrow{\ell} e' :: e_l' \quad (13)$$

$$e :: e_l \wedge e_l \sqsubseteq e_l' \wedge e_l' :: \ell \xRightarrow{\ell} e :: e_l' \quad (14)$$

$$e_l :: \ell \xRightarrow{\ell} c :: e_l \quad \text{for a constant } c \quad (15)$$

$$e_1 :: e_l \wedge \dots \wedge e_n :: e_l \xRightarrow{\ell} f(e_1, \dots, e_n) :: e_l \quad \text{for } n > 0 \quad (16)$$

$$e_p \xrightarrow{e_l} e_v \wedge e_l \sqsubseteq \ell \xRightarrow{\ell} e_p \xrightarrow{e_l} e_v \wedge e_p :: e_l \wedge e_v :: e_l \quad (17)$$

$$P \xRightarrow{\ell} P' \text{ and } Q \xRightarrow{\ell} Q' \text{ implies } P \star Q \xRightarrow{\ell} P' \star Q' \quad (18)$$

Entailment (13) shows that sensitivity of values is compatible with equality. While one would expect (13) to hold naturally, this cannot be taken for granted and in fact the property fails in the prior security separation logic of [13], where labels are part of the semantics of expressions but are not compared by equality.

The second property (14) captures the intuition that less-sensitive data can always be used in contexts where more-sensitive data might be expected (but not vice-versa). Recall that e_l' here is an expression. The additional condition $e_l' :: \ell$ guarantees that this expression denotes a meaningful security level, i.e. evaluates identically in both states (cf. (7)). Note also the abuse of notation to let the semantic ℓ stand for some expression that denotes it.

Property (15) encodes that constants do not depend on any state; again the security level expression e_l must be meaningful, but trivially $c :: \ell$ when ℓ is constant, too.

Value sensitivity is congruent with function application (16). This is not surprising, as functions map arguments equal in both states to equal results. Yet, as with (13) above, this property may fail to hold in alternative logics that attach security levels to values. Note furthermore, that the reverse entailment is not true (a trivial counterexample is any constant function $\lambda x.c$).

Via (17), when $e_p \xrightarrow{e_l} e_v$ it follows that both the location e_p and the value e_v adhere to the level e_l , cf. (8). Note that the antecedent “ $e_p \xrightarrow{e_l} e_v$ ” is repeated in the consequent to ensure that the set of ℓ -attacker visible locations is preserved.

Finally, entailment is monotone with respect to separating conjunction (18), a key ingredient to decompose entailments spatially (cf. “subtraction rules” in [8, Section 4]). The remaining problem without spatial connectives can be encoded into first-order by introducing fresh variables for the second state, see e.g. [16]:

$$\phi \rightsquigarrow \phi \wedge \phi' \quad e :: e_l \rightsquigarrow e_l = e_l' \wedge (e_l \sqsubseteq \ell \rightarrow e = e') \quad \text{etc.} \quad (19)$$

3.3 Proof System

We consider a canonical concurrent programming language with a shared heap mediated by locks but not shared variables. Commands c comprise assignments to local variables, heap access (load and store), standard sequential programming constructs, as well as parallel composition and locking. For simplicity, we assume

$$\begin{array}{c}
\frac{x \notin \text{free}(e)}{\ell \vdash \{\mathbf{emp}\} x := e \{x = e\}} \text{ASSIGN} \\
\\
\frac{x \notin \text{free}(e_p, e_v, e_i)}{\ell \vdash \{e_p \xrightarrow{e_i} e_v\} x := [e_p] \{x = e_v \wedge e_p \xrightarrow{e_i} e_v\}} \text{LOAD} \\
\\
\frac{}{\ell \vdash \{e_v :: e_i \wedge e_p \xrightarrow{e_i} _ \} [e_p] := e_v \{e_p \xrightarrow{e_i} e_v\}} \text{STORE} \\
\\
\frac{}{\ell \vdash \{\mathbf{emp}\} \text{lock } l \{\text{inv}(l)\}} \text{LOCK} \qquad \frac{}{\ell \vdash \{\text{inv}(l)\} \text{unlock } l \{\mathbf{emp}\}} \text{UNLOCK} \\
\\
\frac{\ell \vdash \{P\} c_1 \{R\} \quad \ell \vdash \{R\} c_2 \{Q\}}{\ell \vdash \{P\} c_1; c_2 \{Q\}} \text{SEQ} \\
\\
\frac{\ell \vdash \{b \wedge P\} c_1 \{Q\} \quad \ell \vdash \{\neg b \wedge P\} c_2 \{Q\}}{\ell \vdash \{b :: \ell \wedge P\} \text{if } b \text{ then } c_1 \text{ else } c_2 \{Q\}} \text{IF} \\
\\
\frac{\ell \vdash \{b \wedge (b :: \ell) \wedge P\} c \{b :: \ell \wedge P\}}{\ell \vdash \{b :: \ell \wedge P\} \text{while } b \text{ do } c \{\neg b \wedge P\}} \text{WHILE} \\
\\
\frac{\ell \vdash \{P\} c \{Q\} \quad \text{modified}(c) \cap \text{free}(F) = \emptyset}{\ell \vdash \{P \star F\} c \{Q \star F\}} \text{FRAME} \\
\\
\frac{P \xrightarrow{\ell} P' \quad Q' \xrightarrow{\ell} Q \quad \ell \vdash \{P'\} c \{Q'\}}{\ell \vdash \{P\} c \{Q\}} \text{CONSEQ} \\
\\
\frac{\ell \vdash \{P_1\} c_1 \{P_1\} \quad \ell \vdash \{P_2\} c_2 \{P_2\} \quad \text{modified}(c_i) \cap \text{free}(c_j, P_j, Q_j) = \emptyset \text{ for } i \neq j}{\ell \vdash \{P_1 \star P_2\} c_1 \parallel c_2 \{Q_1 \star Q_2\}} \text{PAR}
\end{array}$$

Fig. 3. Proof Rules of SECCSL.

a static collection of valid lock identifiers l , each of which has an assertion as its associated invariant $\text{inv}(l)$, describing the protected portion of the heap.

$$\begin{aligned}
c ::= & x := e \mid x := [e_v] \mid [e_p] := e_v \mid \text{lock } l \mid \text{unlock } l \\
& c_1; c_2 \mid c_1 \parallel c_2 \mid \text{if } b \text{ then } c_1 \text{ else } c_2 \mid \text{while } b \text{ do } c
\end{aligned}$$

We postpone the (fairly standard) semantics to Section 4 where we present the soundness proof. We note that reads of volatile variables can be soundly reasoned about in this language using SECCSL by associating a fresh lock to each volatile variable and desugaring each read and write to a sequence of lock/access/unlock.

The SECCSL proof rules are shown in Fig. 3. They extend the standard rules of concurrent separation logic [38] (CSL) by additional side-conditions that amount to information flow checks $e :: e_i$ as part of the respective preconditions. The proof system of SECCSL is remarkably similar to the one of CSL.

Similarly to [43], without loss of generality we require that assignments (rule ASSIGN) are always to distinct variables, to avoid renaming in the assertions. Storing to a heap location through an e_l -sensitive location $e_p \xrightarrow{e_l} e_v$ (rule STORE) requires that the value e_v written to that location admits the corresponding security level e_l of the location e_p . Note that due to monotonicity (14) the security level does not have to match exactly. The rules for locking are standard [12]. The consequence rule (CONSEQ) uses entailment defined relative to ℓ (Definition 2). The parallel rule has the usual proviso that the variables modified in one thread cannot interfere with those relied on by the other and its pre-/postcondition.

To preclude information leaks through timing channels, the execution can branch on non-secret values only. This manifests in side conditions $b :: \ell$ for the respective branching condition b where, recall, ℓ is the attacker security level.

4 Security Definition & Soundness

The soundness theorem for SECCSL guarantees that if some triple $\ell \vdash \{P\} c \{Q\}$ is derived using the rules of Fig. 3, then: all executions of c started in a state satisfying precondition P are memory *safe*, partially *correct* with respect to postcondition Q , and moreover *secure* with respect to the sensitivity of values as denoted by P and Q and at all times respect the sensitivity of locations as denoted by P (see Section 2.3). Proof outlines are relegated to Appendix B. All results have been mechanised in Isabelle/HOL [35] and are available at [1].

The top-level security property of SECCSL is a noninterference condition [17]. Noninterference as a security property specifies, roughly, that for any pair of executions that start in states that agree on the values of all attacker-observable inputs, then, from the attacker’s point of view the resulting executions will be indistinguishable, i.e. all of the attacker visible observations will agree. In SECCSL, what is “attacker-observable” depends on the attacker level ℓ . The “inputs” are the expressions e , and the attacker-visible inputs are those expressions e whose sensitivity is given by $e :: \ell'$ judgements in the precondition P for which $\ell' \sqsubseteq \ell$. The attacker-visible observations are the contents of all memory locations in $\text{low}_\ell(P, s)$, for initial store s and precondition P .

ℓ Equivalence. Thus we define when two heaps are indistinguishable to the ℓ -attacker, in which case we say they are *ℓ -equivalent*.

Definition 3 (ℓ Equivalence). *Two heaps coincide on a set of locations A , written $h \stackrel{A}{\equiv} h'$, iff for all $a \in A$. $a \in \text{dom}(h) \cap \text{dom}(h')$ and $h(a) = h'(a)$. Two heaps h and h' are ℓ -equivalent wrt. store s and assertion P , if $h \stackrel{A}{\equiv} h'$ for $A = \text{low}_\ell(P, s)$.*

Then, the ℓ -validity of an assertion P in the relational semantics witnesses ℓ -equivalence between the corresponding heaps.

Lemma 1. *If $(s, h), (s', h') \models_\ell P$, then $h \stackrel{A}{\equiv} h'$ for $A = \text{low}_\ell(P, s)$.*

Furthermore, if $(s, h), (s', h') \models_\ell P$, then $\text{low}_\ell(P, s) = \text{low}_\ell(P, s')$ since the security levels in labeled points-to predicates must coincide on s and s' , cf. (8).

Semantics. Semantic configurations, denoted by k in the following, are one of three kinds: $(\mathbf{run} \ c, L, s, h)$ denotes a command c in a state s, h where L is a set of locks that are currently not held by any thread and can be acquired by c ; $(\mathbf{stop} \ L, s, h)$ similarly denotes a final state s, h with residual locks L , and \mathbf{abort} results from invalid heap access.

The single-step relation $(\mathbf{run} \ c, L, s, h) \xrightarrow{\sigma} k$ takes running configurations to successors k with respect to a schedule σ that resolves the non-determinism of parallel composition. The schedule σ is a list of *actions*: the action $\langle \tau \rangle$ represents the execution of atomic commands and the evaluation of conditionals; the actions $\langle 1 \rangle$ and $\langle 2 \rangle$ respectively denote the execution of the left- and right-hand sides of a parallel composition for a single step, and so define a deterministic scheduling discipline reminiscent of separation kernels [30]. For example, $(\mathbf{run} \ c_1 \ \parallel \ c_2, L, s, h) \xrightarrow{\langle 1 \rangle \cdot \sigma} (\mathbf{run} \ c'_1 \ \parallel \ c_2, L', s', h')$ if $(\mathbf{run} \ c_1, L, s, h) \xrightarrow{\sigma} (\mathbf{run} \ c'_1, L', s', h')$.

Configurations $(\mathbf{run} \ \mathbf{lock} \ l, L, s, h)$ can only be scheduled if $l \in L$ (symmetrically for \mathbf{unlock}) and otherwise block without a possible step. The rules are otherwise standard and can be found in Appendix A.

Executions $k_1 \xrightarrow{\sigma_1 \cdots \sigma_n}^* k_{n+1}$ chain several steps $k_i \xrightarrow{\sigma_i} k_{i+1}$ by accumulating the schedule. Note that since we are considering partial correctness only, the schedule is always finite and so are all executions.

Compositional Security. To prove that SECCSL establishes its top-level non-interference condition, we first define a compositional security condition that provides the central characterization of security for a command c with respect to precondition P and postcondition Q . That central, compositional property we denote $\mathbf{secure}_\ell^n(P, c, Q)$ and formalize below in Definition 4. It ensures that the first n steps (or fewer if the program terminates before that) are safe and preserve ℓ -equivalence of the heap locations specified initially in P , but in a way that is compositional across multiple execution steps, across multiple threads of execution and across different parts of the heap. It is somewhat akin, although more precise than, prior characterizations based on *strong low bisimulation* [42,15].

Disregarding the case when c terminates before the n -th step for a moment, for a pair of initial states (s_1, h_1) and (s'_1, h'_1) and initial set of locks L_1 , and a fixed schedule $\sigma = \sigma_1 \cdots \sigma_n$, $\mathbf{secure}_\ell^{n+1}(P_1, c_1, Q)$ requires that c performs a sequence of lockstep execution steps from each initial state

$$\begin{aligned} (\mathbf{run} \ c_i, L_i, s_i, h_i) &\xrightarrow{\sigma_i} (\mathbf{run} \ c_{i+1}, L_{i+1}, s_{i+1}, h_{i+1}) && \text{for } 1 \leq i \leq n && (20) \\ (\mathbf{run} \ c_i, L_i, s'_i, h'_i) &\xrightarrow{\sigma_i} (\mathbf{run} \ c_{i+1}, L_{i+1}, s'_{i+1}, h'_{i+1}) \end{aligned}$$

These executions must agree on the intermediate commands c_i and locks L_i and the i th pair of states must satisfy an intermediate assertion of the following form:

$$(s_i, h_i), (s'_i, h'_i) \models_\ell P_i \star F \star \mathbf{invs}(L_i) \quad \text{where } \mathbf{invs}(L_i) = \star_{l \in L_i} \mathbf{inv}(l_i) \quad (21)$$

Here P_i describes the part of the heap that command c_i is currently accessing. $\mathbf{invs}(L_i)$ is the set of lock invariants for the locks $l_i \in L_i$ not currently acquired.

Its presence ensures that whenever a lock is acquired that the associated invariant can be assumed to hold. Finally F is an arbitrary *frame*, an assertion that does not mention variables updated by c_i . Its inclusion allows the security property to compose with respect to different parts of the heap.

Moreover, each $P_{i+1} \star \text{invs}(L_{i+1})$ is required to preserve the sensitivity of all ℓ -visible heap locations of $P_i \star \text{invs}(L_i)$, i.e. so that $\text{low}_\ell(P_i \star \text{invs}(L_i), s_i) \subseteq \text{low}_\ell(P_{i+1} \star \text{invs}(L_{i+1}), s_{i+1})$. If some intermediate step $m \leq n$ terminates, then $P_{m+1} = Q$, ensuring the postcondition holds when the executions terminate. Lastly, neither execution is allowed to reach an **abort** configuration.

If the initial state satisfies $P_1 \star F \star \text{invs}(L_1)$ then (21) holds throughout the entire execution, and establishes the end-to-end property that any final state indeed satisfies the postcondition and that $\text{low}_\ell(P_1 \star \text{invs}(L_1), s_1) \subseteq \text{low}_\ell(P_i \star \text{invs}(L_i), s_i)$ with respect to the initially specified low locations.

The property $\text{secure}_\ell^n(P, c, Q)$ is defined recursively to match the steps of the lockstep execution of the program.

Definition 4 (Security).

- $\text{secure}_\ell^0(P_1, c_1, Q)$ holds always.
- $\text{secure}_\ell^{n+1}(P_1, c_1, Q)$ holds, iff for all pairs of states $(s_1, h_1), (s'_1, h'_1)$, frames F , and sets of locks L_1 , such that $(s_1, h_1), (s'_1, h'_1) \models_\ell P_1 \star F \star \text{invs}(L_1)$, and given two steps $(\text{run } c_1, L_1, s_1, h_1) \xrightarrow{\sigma} k$ and $(\text{run } c_1, L_1, s'_1, h'_1) \xrightarrow{\sigma} k'$ there exists an assertion P_2 and a pair of successor states with either of
 - $k = (\text{stop } L_2, s_2, h_2)$ and $k' = (\text{stop } L_2, s'_2, h'_2)$ and $P_2 = Q$
 - $k = (\text{run } c_2, L_2, s_2, h_2)$ and $k' = (\text{run } c_2, L_2, s'_2, h'_2)$ with $\text{secure}_\ell^n(P_2, c_2, Q)$ such that $(s_2, h_2), (s'_2, h'_2) \models_\ell P_2 \star F \star \text{invs}(L_2)$ and $\text{low}_\ell(P_1 \star \text{invs}(L_1), s_1) \subseteq \text{low}_\ell(P_2 \star \text{invs}(L_2), s_2)$ in both cases.

Two further technical side condition are imposed: c_1 doesn't modify variables occurring in $\text{invs}(L_1)$ and F (which guarantees that both remain intact), and the free variables in P_2 can only mention those already present in P_1, c_1 , or in any lock invariant (which guarantees that P_2 remains stable against concurrent assignments). Note that each step can pick a different frame F .

Lemma 2. $\ell \vdash \{P\} c \{Q\}$ implies $\text{secure}_\ell^n(P, c, Q)$ for every $n \geq 0$.

Safety, Correctness and Noninterference. Execution safety and correctness of with respect to pre- and postcondition follow straightforwardly from Lemma 2.

Corollary 1 (Safety). Given initial states $(s_1, h_1), (s'_1, h'_1) \models_\ell P \star \text{invs}(L_1)$ and two executions of a command c under the same schedule to resulting configurations k and k' respectively, then $\ell \vdash \{P\} c \{Q\}$ implies $k \neq \text{abort} \wedge k' \neq \text{abort}$.

Theorem 1 (Correctness). Given initial states $(s_1, h_1), (s'_1, h'_1) \models_\ell P \star \text{invs}(L_1)$ and two complete executions of a command c under the same schedule σ

$$\begin{aligned} & (\text{run } c, L_1, s_1, h_1) \xrightarrow{\sigma}^* (\text{stop } L_2, s_2, h_2) \\ & (\text{run } c_i, L_i, s'_i, h'_i) \xrightarrow{\sigma}^* (\text{stop } L_2, s'_2, h'_2) \end{aligned}$$

then $\ell \vdash \{P\} c \{Q\}$ implies $(s_2, h_2), (s'_2, h'_2) \models_\ell Q \star \text{invs}(L_2)$.

The top-level noninterference property also follows from Lemma 2. We state the noninterference property directly in the theorem in the interests of brevity.

Theorem 2 (Noninterference). *Given a command c , and initial states $(s_1, h_1), (s'_1, h'_1) \models_\ell P \star \text{invs}(L_1)$ then $\ell \vdash \{P\} c \{Q\}$ implies $h_i \stackrel{A}{\equiv} h'_i$, where $A = \text{low}_\ell(P, s_1)$, for all pairs of heaps h_i and h'_i arising from executing the same schedule from each initial state.*

5 SECC: Automating SECCSL

To demonstrate the ease by which SECCSL can be automated, we implemented SECC, available at [1]. SECC implements the logic from Section 3 for a subset of C, and is written in Scala². SECC is currently used to explore reasoning about example programs with interesting security policies. Thus its engineering has focused on features related to security reasoning (e.g. deciding when conditions $e :: e_i$ are entailed) rather than reasoning about complex data structures.

SECC automates SECCSL through symbolic execution [8], by computing strongest postconditions of commands with respect to a given precondition. Z3 [14] is used to discharge predicate logic verification conditions as well as those arising from security related side conditions, using (17) and the translation scheme (19) from Section 3.2. In addition to the logic from Section 3, SECC supports procedure modular verification with pre-/postconditions as usual.

While some issues of the C source language are not addressed (yet), such as integer overflow, those that impact directly on information flow security are taken into account. Specifically, the shortcut semantics of boolean operators `&&`, `||`, and ternary `_ ? _ : _` count as branching points and as such the left hand side resp. the test must not depend on sensitive data, similarly to the conditions of `if` statements and `while` loops.

A direct benefit of the integration of security levels into the assertion language is that it becomes possible to specify the sensitivity of data passed to library and operating system functions. For example, the execution time of `malloc(len)` would depend on the value of `len`, which can thus be required to satisfy `len :: low` by annotating its function header with an appropriate precondition, using SECC’s `requires` annotation. Likewise, SECC can reason about limited forms of declassification, in which external functions are trusted to safely release otherwise sensitive data, by giving them appropriate pre-/postconditions. For example, a password hashing library function prototype might be annotated with a postcondition asserting its result is `low`, via SECC’s `ensures` annotation.

SECC proves Fig. 1 secure, and correctly flags buggy variants as insecure, e.g., where the test in thread 1 is reversed, or when thread 2 does not clear the `data` field upon setting the `is_classified` to `FALSE`.

To compare SECC and SECCSL against the recent COVERN logic [32], we took a non-trivial example program that Murray et al. verified in COVERN, manually translated it to C, and verified it automatically using SECC.

² <https://www.scala-lang.org>

The original program³, written in COVERN’s tiny While language embedded in Isabelle/HOL, models the software functionality of a simplified implementation of the Cross Domain Desktop Compositor (CDDC) [6]. The CDDC is a device that facilitates interactions with multiple PCs, each of which runs applications at differing sensitivity, from a single keyboard, mouse and display. Its multi-threaded software handles routing of keyboard input to the appropriate PC and switching between the PCs via mouse gestures.

Verifying the C translation required adding SECCSL annotations for procedure pre-/postconditions and loop invariants. The C translation including those annotations is ~ 250 lines in length. The present, unoptimised, implementation of SECC verifies the resulting artifact in ~ 15 seconds. In contrast, the COVERN proof of this example requires ~ 600 lines of Isabelle/HOL definitions/specification, plus ~ 550 lines of Isabelle proof script.

6 Related Work

There has been much work targeting type systems and program logics for concurrent information flow. Karbyshev et al. [21] provide an excellent overview. Here we concentrate on work whose ideas are most closely related to SECCSL.

Security Separation Logics. Costanzo and Shao [13] propose a sequential separation logic for reasoning about information flow. Unlike SECCSL, theirs does not distinguish value and location sensitivity. Their separation logic assertions have a fairly standard (non-relational) semantics, at the price of having a *security-aware* language semantics that propagates security labels attached to values in the store and heap. As mentioned in Section 3.2, this has the unfortunate side-effect of breaking intuitive properties about sensitivity assertions. We conjecture that the absence of such properties would make their logic harder to automate than SECCSL, which SECC demonstrates is feasible. SECCSL avoids the aforementioned drawbacks by adopting a relational assertion semantics.

Eilers et al. [16] present a verification technique that can deal with a more general class of hyper-properties but does not support concurrency and without public heap locations. It is automated in the mature tool Viper. The approach is based on a product encoding of programs whereas our logic can be readily automated using standard techniques.

Gruetter and Murray [18] propose a security separation logic in Coq [9] for Verifiable C, the C subset of the Verified Software Toolchain [4,3]. However they provide no soundness proof for its rules and its feasibility to automate is unclear.

Separation-Like Approaches. Two recent compositional logics for concurrent information flow are the COVERN logic [32] and the type and effect system of Karbyshev et al. [21]. Both borrow ideas from separation logic. However, unlike SECCSL, neither is defined for languages with pointers, arrays etc.

Like SECCSL, COVERN proves a timing-sensitive security property. Location sensitivity is defined statically by value-dependent predicates, and value sensi-

³ https://bitbucket.org/covern/covern/src/master/examples/cddc/Example_CDDC_WhileLockLanguage.thy

tivity is tracked by a dependent security typing context Γ [33], relative to a Hoare logic predicate P over the entire shared memory. In COVERN locks carry non-relational invariants. In contrast, SECCSL unifies these elements together into separation logic assertions with a relational semantics. Doing so leads to a much simpler logic, amenable to automation, while supporting pointers, etc.

On the other hand, Karbyshev et al. [21] prove a timing-*insensitive* security property, but rely on primitives to interact with the scheduler to prevent leaks via scheduling decisions. Unlike SECCSL, which assumes a deterministic scheduling discipline, Karbyshev et al. support a wider class of scheduling policies. Their system also allows for racy data accesses between threads which SECCSL does not handle. Their system tracks resource ownership and transfer between threads at synchronisation points, similar to CSLs. Their resources include *labelled scheduler resources* that account for scheduler interaction, including when scheduling decisions become tainted by secret data—something that cannot occur in SECCSL’s deterministic scheduling model.

Prior logics for sequential languages, e.g. [2,5], have also adopted separation logic ideas to reason locally about memory, combining them with relational assertions similar to SECCSL’s $e :: e_l$ assertions. For instance, the agreement assertions $A(e)$ of [5] coincide with SECCSL’s $e :: \text{low}$. Unlike SECCSL, some of these logics support languages with explicit declassification actions [5].

CSL Soundness. The soundness proof for SECCSL follows the general structure of Vafeiadis’ [43] for CSL, which is also mechanised in Isabelle/HOL. There is, however, a technical difference: His analog of Definition 4, a recursive predicate called $\text{safe}_n(c, s, h, Q)$, refers to a semantic initial state s, h whereas we propagate a syntactic assertion (21) only. Our formulation has the benefit that some of the technical reasoning in the soundness proof is easier to automate. Its drawback is the need to impose technical side-conditions on the free variables of the frame F and the intermediate assertions P_i .

7 Conclusion

We presented SECCSL, a concurrent separation logic for proving expressive data-dependent information flow properties of programs. SECCSL is considerably simpler, yet handles features like pointers, arrays etc., which are out of scope for contemporary logics. It inherits the structure of traditional concurrent separation logics, and so like those logics can be automated via symbolic execution [20,28,10]. To demonstrate this, we implemented SECC, the first automatic verifier for expressive information flow security for a subset of the C.

Separation logic has proved to be a remarkably powerful vehicle for reasoning about programs, extending to weak memory concurrency [44], program synthesis [40] and many other domains. With SECCSL, we hope that in future the same possibilities might be opened to verified information flow security.

References

1. Isabelle files and tool demo for SecCSL (via Anonymous GitHub for Open-Science). <https://anonymous.4open.science/repository/42ff2a21-f974-476f-9e7d-980f9b63fff92> (2019)
2. Amtoft, T., Bandhakavi, S., Banerjee, A.: A logic for information flow in object-oriented programs. In: Proc. of Principles of Programming Languages (POPL). pp. 91–102. ACM (2006)
3. Appel, A.W., Dockins, R., Hobor, A., Beringer, L., Dodds, J., Stewart, G., Blazy, S., Leroy, X.: Program Logics for Certified Compilers. Cambridge University Press, New York, NY, USA (2014)
4. Appel, A.W., others: The Verified Software Toolchain. <https://github.com/PrincetonUniversity/VST> (2017)
5. Banerjee, A., Naumann, D.A., Rosenberg, S.: Expressive declassification policies and modular static enforcement. In: Proc. of Symposium on Security and Privacy (S&P). pp. 339–353. IEEE (2008)
6. Beaumont, M., McCarthy, J., Murray, T.: The cross domain desktop compositor: using hardware-based video compositing for a multi-level secure user interface. In: Annual Computer Security Applications Conference (ACSAC). pp. 533–545. ACM (2016)
7. Benton, N.: Simple relational correctness proofs for static analyses and program transformations. In: Proc. of Principles of Programming Languages (POPL). pp. 14–25. ACM (2004)
8. Berdine, J., Calcagno, C., O’hearn, P.W.: Symbolic execution with Separation Logic. In: Asian Symposium on Programming Languages and Systems (APLAS). pp. 52–68. Springer (2005)
9. Bertot, Y., Castéran, P.: Interactive Theorem Proving and Program Development. Coq’Art: The Calculus of Inductive Constructions. Texts in Theoretical Computer Science. An EATCS Series, Springer (2004)
10. Calcagno, C., Distefano, D.: Infer: An automatic program verifier for memory safety of c programs. In: NASA Formal Methods Symposium. pp. 459–465. Springer (2011)
11. Calcagno, C., Distefano, D., Dubreil, J., Gabi, D., Hooimeijer, P., Luca, M., O’Hearn, P., Papakonstantinou, I., Purbrick, J., Rodriguez, D.: Moving fast with software verification. In: NASA Formal Methods Symposium. pp. 3–11. Springer (2015)
12. Chlipala, A.: Formal Reasoning About Programs (2016)
13. Costanzo, D., Shao, Z.: A separation logic for enforcing declarative information flow control policies. In: Proc. of Principles of Security and Trust (POST). LNCS, vol. 8414, pp. 179–198. Springer (2014)
14. De Moura, L., Bjørner, N.: Z3: An efficient SMT solver. In: Proc. of Tools and Algorithms for the Construction and Analysis of Systems (TACAS). pp. 337–340. Springer (2008)
15. Del Tedesco, F., Sands, D., Russo, A.: Fault-resilient non-interference. In: Proc. of Computer Security Foundations Symposium (CSF). pp. 401–416. IEEE (2016)
16. Eilers, M., Müller, P., Hitz, S.: Modular product programs. In: Proc. of European Symposium on Programming (ESOP). pp. 502–529. Springer (2018)
17. Goguen, J., Meseguer, J.: Security policies and security models. In: Proc. of Symposium on Security and Privacy (S&P). pp. 11–20. Oakland, California, USA (Apr 1982)

18. Gruetter, S., Murray, T.: Short paper: Towards information flow reasoning about real-world C code. In: Proc. of Workshop on Programming Languages and Analysis for Security (PLAS). pp. 43–48. ACM (2017)
19. Gu, R., Shao, Z., Chen, H., Wu, X.N., Kim, J., Sjöberg, V., Costanzo, D.: CertiKOS: An extensible architecture for building certified concurrent OS kernels. In: Proc. of USENIX Symposium on Operating Systems Design and Implementation (OSDI) (Nov 2016)
20. Jacobs, B., Smans, J., Philippaerts, P., Vogels, F., Penninckx, W., Piessens, F.: VeriFast: A powerful, sound, predictable, fast verifier for C and Java. In: NASA Formal Methods Symposium. pp. 41–55. Springer (2011)
21. Karbyshev, A., Svendsen, K., Askarov, A., Birkedal, L.: Compositional non-interference for concurrent programs via separation and framing. In: Proc. of Principles of Security and Trust (POST) (2018)
22. Klein, G., Andronick, J., Elphinstone, K., Murray, T., Sewell, T., Kolanski, R., Heiser, G.: Comprehensive formal verification of an OS microkernel. ACM Transactions on Computer Systems **32**(1), 2:1–2:70 (Feb 2014)
23. Leroy, X.: Formal verification of a realistic compiler. Communications of the ACM **52**(7), 107–115 (2009)
24. Löding, C., Madhusudan, P., Murali, A., na, L.P.: A first order logic with frames, under submission. Available at: <http://madhu.cs.illinois.edu/F0FrameLogic.pdf>
25. Lourenço, L., Caires, L.: Dependent information flow types. In: Proc. of Principles of Programming Languages (POPL). pp. 317–328. Mumbai, India (Jan 2015)
26. Mantel, H., Sands, D.: Controlled declassification based on intransitive noninterference. In: Asian Symposium on Programming Languages and Systems (APLAS). pp. 129–145. Springer (2004)
27. Mantel, H., Sands, D., Sudbrock, H.: Assumptions and guarantees for compositional noninterference. In: Proc. of Computer Security Foundations Symposium (CSF). pp. 218–232. Cernay-la-Ville, France (Jun 2011)
28. Müller, P., Schwerhoff, M., Summers, A.J.: Viper: A verification infrastructure for permission-based reasoning. In: Proc. of Verification, Model Checking, and Abstract Interpretation (VMCAI). pp. 41–62. Springer (2016)
29. Murray, T.: Short paper: On high-assurance information-flow-secure programming languages. In: Proc. of Workshop on Programming Languages and Analysis for Security (PLAS). pp. 43–48 (2015)
30. Murray, T., Matichuk, D., Brassil, M., Gammie, P., Bourke, T., Seefried, S., Lewis, C., Gao, X., Klein, G.: seL4: from general purpose to a proof of information flow enforcement. In: Proc. of Symposium on Security and Privacy (S&P). pp. 415–429. San Francisco, CA (May 2013)
31. Murray, T., Sabelfeld, A., Bauer, L.: Special issue on verified information flow security. Journal of Computer Security **25**(4-5), 319–321 (2017)
32. Murray, T., Sison, R., Engelhardt, K.: COVERN: A logic for compositional verification of information flow control. In: Proc. of European Symposium on Security and Privacy (EuroS&P). London, United Kingdom (Apr 2018)
33. Murray, T., Sison, R., Pierzchalski, E., Rizkallah, C.: Compositional verification and refinement of concurrent value-dependent noninterference. In: Proc. of Computer Security Foundations Symposium (CSF). pp. 417–431 (Jun 2016)
34. Newcombe, C., Rath, T., Zhang, F., Munteanu, B., Brooker, M., Deardouff, M.: How amazon web services uses formal methods. Communications of the ACM **58**(4), 66–73 (2015)
35. Nipkow, T., Paulson, L., Wenzel, M.: Isabelle/HOL — A Proof Assistant for Higher-Order Logic, LNCS, vol. 2283 (2002)

36. O’Hearn, P.W.: Continuous reasoning: Scaling the impact of formal methods. In: Proc. of Logic in Computer Science (LICS). pp. 13–25. ACM (2018)
37. O’Hearn, P.W., Yang, H., Reynolds, J.C.: Separation and information hiding. ACM Transactions on Programming Languages and Systems (TOPLAS) **31**(3), 11 (2009)
38. O’hearn, P.W.: Resources, concurrency and local reasoning. In: International Conference on Concurrency Theory (CONCUR). pp. 49–67. Springer (2004)
39. Piskac, R., Wies, T., Zufferey, D.: Automating separation logic using SMT. In: Proc. of Computer Aided Verification (CAV). pp. 773–789. Springer (2013)
40. Polikarpova, N., Sergey, I.: Structuring the synthesis of heap-manipulating programs. Proceedings of the ACM on Programming Languages **3**(POPL), 72 (2019)
41. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: Proc. of Logic in Computer Science (LICS). pp. 55–74. IEEE (2002)
42. Sabelfeld, A., Sands, D.: Probabilistic noninterference for multi-threaded programs. In: Proc. of Computer Security Foundations Workshop (CSFW). pp. 200–214. IEEE (2000)
43. Vafeiadis, V.: Concurrent separation logic and operational semantics. In: Proc. of Mathematical Foundations of Programming Semantics (MFPS). pp. 335–351 (2011)
44. Vafeiadis, V., Narayan, C.: Relaxed separation logic: A program logic for C11 concurrency. In: Proc. of Object Oriented Programming Systems Languages & Applications (OOPSLA). pp. 867–884. ACM (2013)
45. Volpano, D., Smith, G.: Probabilistic noninterference in a concurrent language. Journal of Computer Security **7**(2,3), 231–253 (1999)
46. Yang, H.: Relational separation logic. Theoretical Computer Science **375**(1-3), 308–334 (2007)
47. Zheng, L., Myers, A.C.: Dynamic security labels and static information flow control. International Journal of Information Security **6**(2–3) (Mar 2007)

A Command Semantics

Symmetric parallel rules in which c_2 is scheduled under the action $\langle 2 \rangle$ omitted.

$$\begin{array}{c}
 \frac{s' = s(x \mapsto \llbracket e \rrbracket_s)}{(\mathbf{run} \ x := e, L, s, h) \xrightarrow{\langle \tau \rangle} (\mathbf{stop} \ L, s', h)} \quad \frac{\llbracket e \rrbracket_s \notin \text{dom}(h)}{(\mathbf{run} \ x := [e], L, s, h) \xrightarrow{\langle \tau \rangle} \mathbf{abort}} \\
 \frac{\llbracket e \rrbracket_s \in \text{dom}(h) \quad s' = s(x \mapsto h(\llbracket e \rrbracket_s))}{(\mathbf{run} \ x := [e], L, s, h) \xrightarrow{\langle \tau \rangle} (\mathbf{stop} \ L, s', h)} \quad \frac{\llbracket e_1 \rrbracket_s \notin \text{dom}(h)}{(\mathbf{run} \ [e_1] := e_2, L, s, h) \xrightarrow{\langle \tau \rangle} \mathbf{abort}} \\
 \frac{\llbracket e_1 \rrbracket_s \in \text{dom}(h) \quad h' = h(\llbracket e_1 \rrbracket_s \mapsto \llbracket e_2 \rrbracket_s)}{(\mathbf{run} \ [e_1] := e_2, L, s, h) \xrightarrow{\langle \tau \rangle} (\mathbf{stop} \ L, s, h')} \\
 \frac{l \in L \quad L' = L \setminus \{l\}}{(\mathbf{run} \ \mathbf{lock} \ l, L, s, h) \xrightarrow{\langle \tau \rangle} (\mathbf{stop} \ L', s, h)} \\
 \frac{l \notin L \quad L' = L \cup \{l\}}{(\mathbf{run} \ \mathbf{unlock} \ l, L, s, h) \xrightarrow{\langle \tau \rangle} (\mathbf{stop} \ L', s, h)} \quad \frac{(\mathbf{run} \ c_1, L, s, h) \xrightarrow{\sigma} \mathbf{abort}}{(\mathbf{run} \ c_1; c_2, L, s, h) \xrightarrow{\sigma} \mathbf{abort}} \\
 \frac{(\mathbf{run} \ c_1, L, s, h) \xrightarrow{\sigma} \mathbf{abort}}{(\mathbf{run} \ c_1 \parallel c_2, L, s, h) \xrightarrow{\langle 1 \rangle \cdot \sigma} \mathbf{abort}} \quad \frac{(\mathbf{run} \ c_1, L, s, h) \xrightarrow{\sigma} (\mathbf{stop} \ L', s', h')}{(\mathbf{run} \ c_1; c_2, L, s, h) \xrightarrow{\sigma} (\mathbf{run} \ c_2, L', s', h')} \\
 \frac{(\mathbf{run} \ c_1, L, s, h) \xrightarrow{\sigma} (\mathbf{run} \ c'_1, L', s', h')}{(\mathbf{run} \ c_1; c_2, L, s, h) \xrightarrow{\sigma} (\mathbf{run} \ c'_1; c_2, L', s', h')} \\
 \frac{(\mathbf{run} \ c_1, L, s, h) \xrightarrow{\sigma} (\mathbf{stop} \ L', s', h')}{(\mathbf{run} \ c_1 \parallel c_2, L, s, h) \xrightarrow{\langle 1 \rangle \cdot \sigma} (\mathbf{run} \ c_2, L', s', h')} \\
 \frac{(\mathbf{run} \ c_1, L, s, h) \xrightarrow{\sigma} (\mathbf{run} \ c'_1, L', s', h')}{(\mathbf{run} \ c_1 \parallel c_2, L, s, h) \xrightarrow{\langle 1 \rangle \cdot \sigma} (\mathbf{run} \ c'_1 \parallel c_2, L', s', h')} \\
 \frac{\text{if } s \models b \text{ then } c' = c_1 \text{ else } c' = c_2}{(\mathbf{run} \ \mathbf{if} \ b \ \mathbf{then} \ c_1 \ \mathbf{else} \ c_2, L, s, h) \xrightarrow{\langle \tau \rangle} (\mathbf{run} \ c', L, s, h)} \\
 \frac{s \not\models b}{(\mathbf{run} \ \mathbf{while} \ b \ \mathbf{do} \ c, L, s, h) \xrightarrow{\langle \tau \rangle} (\mathbf{stop} \ L, s, h)} \\
 \frac{s \models b}{(\mathbf{run} \ \underbrace{\mathbf{while} \ b \ \mathbf{do} \ c}_\omega, L, s, h) \xrightarrow{\langle \tau \rangle} (\mathbf{run} \ (c; \omega), L, s, h)}
 \end{array}$$

$$k \xrightarrow{\langle \rangle^*} k \quad \frac{k \xrightarrow{\sigma_1} k' \quad k' \xrightarrow{\sigma_2^*} k''}{k \xrightarrow{\sigma_1 \cdot \sigma_2^*} k''}$$

B Proofs

Proof of Lemma 1

If $(s, h), (s', h') \models_\ell P$, then $h \stackrel{A}{\equiv} h'$ for $A = \text{low}_{s_\ell}(P, s)$.

Proof. By induction on the structure of P , noting that $\text{low}_{s_\ell}(_, s)$ contains locations of the corresponding sub-heap only. \square

Proof of Lemma 2

$\ell \vdash \{P\} c \{Q\}$ implies $\text{secure}_\ell^n(P, c, Q)$ for every $n \geq 0$.

Proof (Outline). By induction on the derivation of the validity of the judgement. Noting that $n = 0$ is trivial, we may unfold the recursion of the security definition once to prove the base cases of assignment, load, store, and locking, which then follow from the respective side conditions of the proof rules.

For rules IF and WHILE, the side condition $b :: \ell$ guarantees that the test evaluates equivalently in the two states and thus execution proceeds with the same remainder program.

Except for IF, all remaining rules need a second induction on n to stepwise match security of the premise to security of the conclusion (e.g. over the steps of the first command in a sequential composition $c_1; c_2$).

The rule FRAME instantiates the frame F with the same assertion in each step, whereas PAR uses the frame F to preserve the current precondition P_2 of c_2 over steps of c_1 and vice-versa. \square

Proof of Corollary 1

Given a command c and initial states $(s_1, h_1), (s'_1, h'_1) \models_\ell P \star \text{invs}(L_1)$ and two executions under the same schedule to resulting configurations k and k' respectively, then $\ell \vdash \{P\} c \{Q\}$ implies $k \neq \mathbf{abort} \wedge k' \neq \mathbf{abort}$.

Proof. By induction on the number of steps n of the executions from $\text{secure}_\ell^n(P, c, Q)$ via Lemma 2. \square

Proof of Theorem 1

Given a command c and initial states $(s_1, h_1), (s'_1, h'_1) \models_\ell P \star \text{invs}(L_1)$ and two complete executions under the same schedule σ

$$\begin{aligned} & (\text{run } c, L_1, s_1, h_1) \xrightarrow{\sigma^*} (\text{stop } L_2, s_2, h_2) \\ & (\text{run } c_i, L_i, s'_i, h'_i) \xrightarrow{\sigma^*} (\text{stop } L_2, s'_2, h'_2) \end{aligned}$$

then $\ell \vdash \{P\} c \{Q\}$ implies $(s_2, h_2), (s'_2, h'_2) \models_\ell Q \star \text{invs}(L_2)$.

Proof. By induction on the number of steps n of the executions from $\text{secure}_\ell^n(P, c, Q)$ via Lemma 2. \square

Proof of Theorem 2

Given a command c , and initial states $(s_1, h_1), (s'_1, h'_1) \models_\ell P \star \text{invs}(L_1)$ then $\ell \vdash \{P\} c \{Q\}$ implies $h_i \stackrel{A}{\equiv} h'_i$, where $A = \text{low}_\ell(P, s_1)$, for all pairs of heaps h_i and h'_i arising from executing the same schedule from each initial state.

Proof. By induction on the number of steps i up to that state from $\text{secure}_\ell^i(P, c, Q)$ via Lemma 2 we have $\text{low}_\ell(P \star \text{invs}(L_1), s_1) \subseteq \text{low}_\ell(P_i \star \text{invs}(L_1), s_i)$ transitively over the prefix, where P_i and s_i are from the i -th state. The theorem then follows from Lemma 1 in Section 3.1. \square