

COMMCSL: Proving Information Flow Security for Concurrent Programs using Abstract Commutativity

MARCO EILERS, ETH Zurich, Switzerland

THIBAUT DARDINIER, ETH Zurich, Switzerland

PETER MÜLLER, ETH Zurich, Switzerland

Information flow security ensures that the secret data manipulated by a program does not influence its observable output. Proving information flow security is especially challenging for concurrent programs, where operations on secret data may influence the execution time of a thread and, thereby, the interleaving between threads. Such *internal timing channels* may affect the observable outcome of a program even if an attacker does not observe execution times. Existing verification techniques for information flow security in concurrent programs attempt to prove that secret data does not influence the relative timing of threads. However, these techniques are often restrictive (for instance because they disallow branching on secret data) and make strong assumptions about the execution platform (ignoring caching, processor instructions with data-dependent execution time, and other common features that affect execution time).

In this paper, we present a novel verification technique for secure information flow in concurrent programs that lifts these restrictions and does not make any assumptions about timing behavior. The key idea is to prove that all mutating operations performed on shared data commute, such that different thread interleavings do not influence its final value. Crucially, commutativity is required only for an *abstraction* of the shared data that contains the information that will be leaked to a public output. Abstract commutativity is satisfied by many more operations than standard commutativity, which makes our technique widely applicable.

We formalize our technique in COMMCSL, a relational concurrent separation logic with support for commutativity-based reasoning, and prove its soundness in Isabelle/HOL. We have implemented COMMCSL in HYPERVIPER, an automated verifier based on the Viper verification infrastructure, and demonstrate its ability to verify challenging examples.

1 INTRODUCTION

Reasoning about information flow is important to ensure the confidentiality and integrity of data. The main goal of information flow security is to ensure the absence of *value channels*, that is, unwanted information flows through the *result values* computed by a program. For sequential programs, preventing value channels is conceptually simple, and is enabled by various existing type systems, program logics, and static analyses.

Depending on the application scenario, it is sometimes desirable to also prevent side channels that let an attacker obtain information about secret values indirectly, by observing parameters such as execution time. Preventing timing channels requires ruling out *all possible* sources of secret-dependent timing differences, which may depend on complex (but standard) hardware features like caching [Stefan et al. 2013] and speculative execution [Kocher et al. 2019; Lipp et al. 2018].

For sequential programs, proving the absence of value channels does not involve reasoning about timing. Thus, for the many scenarios where attackers cannot observe execution time at all (e.g., because the time when data is made public is decoupled from its computation time, which is common for batch processing) or not with sufficient precision (e.g., because execution time is hidden by a laggy network), simple reasoning about values is sufficient even on standard hardware.

However, this is not the case for concurrent programs with shared memory, where timing differences may result in value channels, which can be observed even by attackers that cannot observe execution time. The program in Figure 1 illustrates this problem.

```

                                t1, t2 := 0
while (t1 < 100) do {           || while (t2 < h) do {
    t1 := t1 + 1                 ||    t2 := t2 + 1
}                                || }
s:=3                             || s:=4
                                || print(s)

```

Fig. 1. Example program: Timing channels become value channels.

Since no information about the secret value h is leaked to the output s through direct assignments or explicit control flow, each thread individually does not leak secret information. However, the execution time of the loop in the right thread depends on the value of h . This timing may affect the order in which the threads perform their assignment to s and, thus, which value is printed in the end. With a deterministic round-robin scheduler, the printed value leaks whether or not h is greater than 100. With a non-deterministic scheduler with a known probability distribution, the entire value of h is leaked probabilistically over multiple executions.

Information leaks that turn secret-dependent timing differences between threads into observable differences in the program state have long been recognized and are called *internal timing channels* [Volpano and Smith 1998]. Some existing solutions to this problem essentially eliminate shared memory between threads [Giffhorn and Snelting 2015; Vechev et al. 2010], which rules out many useful concurrent programming patterns. Others employ the techniques used to prevent standard timing channels. They forbid programs from including operations whose execution time depends on secret inputs [Murray et al. 2018; Sabelfeld and Sands 2000; Schoepe et al. 2020; Smith 2007]; in our example, they would reject the loop in the right thread. However, these techniques typically assume idealized hardware, where secret-depending branching is the only source of timing channels. In principle, they can be extended to standard hardware (e.g., one commonly prevents timing leaks through cache effects by forbidding all secret-dependent memory accesses [Jancar et al. 2022]). However, attempting to rule out *all* sources of timing differences requires information about compilers [Barthe et al. 2018] and the hardware on which code is to be executed [Andryscio et al. 2018; Cleemput et al. 2012], including hardware details manufacturers usually do not make public. As a result, it is virtually impossible to rule out timing channels with absolute certainty, and thus, existing techniques *cannot rule out value channels in concurrent programs on standard hardware with certainty* either, since they require eliminating timing channels.

This work. We present a novel technique to prove the absence of value channels in concurrent programs that does *not* require reasoning about timing and, thus, is sound even in the presence of complex modern hardware features. In the common case that attackers cannot observe execution time, our technique enables conceptually simple information flow security proofs. When the attacker *can* observe timing to a degree, our technique can be combined with orthogonal techniques to try to rule out timing channels. Crucially, our technique is *modular*, i.e., it does not require explicit reasoning about thread interleavings, which would make verification computationally infeasible.

Our technique is based on the observation that different, secret-dependent interleavings of shared data mutations lead to the same final result if the mutating operations on shared data performed by different threads *commute*. For instance, the program above is (rightfully) rejected by our technique because the two assignments to the shared variable s do not commute; however, a variation of this example where the concurrent assignments are modified so that the left thread atomically assigns $s + 3$ to s , and the right thread assigns $s + 4$, is allowed: This program always increases s by 7, since

the additions performed by the two threads commute and their (secret-dependent) order therefore does not influence the final result. Consequently, the adapted program contains no value channels.

Standard commutativity is a strong requirement that is satisfied only by a minority of operations of typical data structures. A key insight of our work is to restrict the commutativity requirement to those parts of the shared data that are required to be non-secret (or *public*). In our example, if s were secret (and consequently not printed at the end of the program) then the program would be information flow secure even though the assignments to s do not commute. To this end, we allow programmers to specify abstract views of shared data structures that contain the information that must remain public; all information not included in the abstraction is considered secret. It is then sufficient to prove *abstract commutativity*, that is, the order of two operations does not affect the abstract view of a data structure, but it may affect other aspects. For instance, two list-append operations commute under a set- or length-abstraction of the list, but do not commute on the concrete list (unless they append the same element).

Many programs expose public views on data structures that contain secret information. One common case is maintaining *internal data* that intentionally contains secret information (e.g., an employee database with individual salaries), and then exposing only a part of it (e.g., average salaries for reporting purposes); in this case, our technique allows specifying the projection to averages as the abstraction. Another common case is maintaining data that is *intended* to be public, but inadvertently tainted with secret information as a result of timing differences in a concurrent computation. For instance, a program may maintain a list of anonymized accounts, but the order in the list depends on the time it took to process the (secret) purchases of each account. In this case, we can use the multiset view of the list as the abstraction, and prove, for instance, that a sorted version of the list can be exposed without creating a value channel.

We formalize our technique as a relational concurrent separation logic with special constructs for reasoning about abstract commutativity of concurrent operations. This logic, COMMCSL, enables *modular* proofs of information flow security, does not require reasoning about the timing behavior of threads, and is amenable to automation via SMT solvers. Our logic supports an expressive assertion language that can express concepts like value-dependent secrecy [Murray et al. 2018]. In addition to proving existing applications secure, the ideas behind our technique can also be used as a guiding principle when building concurrent applications that deal with secret data.

Contributions and outline. We make the following contributions:

- We show how commutativity can be used to prove information flow security in shared-memory concurrent programs. Our verification technique allows threads to flexibly manipulate secret data and does not require reasoning (nor make assumptions) about the timing of executions, such that it is directly applicable to programs running on standard hardware.
- We introduce the generalized notion of abstract commutativity, which requires operations to commute only relative to a user-defined abstraction of the shared state. Abstract commutativity applies to a much wider set of common scenarios.
- We incorporate these techniques into COMMCSL, a novel concurrent separation logic that enables modular proofs of information flow security using abstract commutativity.
- We formalize our logic and prove its soundness in Isabelle/HOL [Nipkow et al. 2002].
- We implement COMMCSL in an automated, SMT-based verification tool called HYPERVIPER based on the Viper verification infrastructure [Müller et al. 2016].
- We evaluate our logic and tool by verifying the information flow security of several challenging, realistic examples.

This paper is organized as follows: Sec. 2 provides an informal overview of our technique. We formally define COMMCSL in Sec. 3, and summarize its soundness proof in Sec. 4. In Sec. 5, we

```

procedure targetSize(households) {
  // ensures: c is low
  n := |households|
  c := createCounter(0)
  (worker(households, 0, n/2, c)
   ||
   worker(households, n/2, n, c))
  return c
}

procedure worker(households, f, t, c) {
  for (i in f..t-1) {
    targets := countTargets(household[i])
    atomic:
      c.add(targets)
  }

  procedure countTargets(household: map[str, int]) {
    return household["nAdults"]
  }
}

```

Fig. 2. Example program: Multiple threads add values to a shared counter. households is an array containing customer data per household. The result of countTargets is low, but its execution time may be high, here due to potential hash collisions with other keys whose presence in the households map is secret.

describe the implementation of the logic in HYPERVIPER and use it to verify several challenging programming patterns. We discuss related work in Sec. 6 and conclude in Sec. 7.

2 OVERVIEW

In this section, we explain the central concepts behind COMMCSSL informally. We occasionally omit details for simplicity; the full logic with all features and checks will be presented in Sec. 3.

2.1 Problem Statement

Our goal in this paper is to prove that concurrent programs that have both high-sensitivity (high) and low-sensitivity (low) inputs do not leak information about the high inputs in their low output values¹. Formally, this can be expressed as *non-interference* [Volpano et al. 1996], a 2-safety *hyper-property* [Clarkson and Schneider 2010], i.e., a property of pairs of finite execution traces of the program ($s(x)$ denotes the value of variable x in program store s):

Definition 2.1. A program c with a set of input variables I and output variables O , of which some subsets $I_l \subseteq I$ and $O_l \subseteq O$ are low, satisfies non-interference iff for all s_1, s_2 and s'_1, s'_2 , if $\forall x \in I_l. s_1(x) = s_2(x)$ and $\langle c, s_1 \rangle \rightarrow^* \langle \text{skip}, s'_1 \rangle$ and $\langle c, s_2 \rangle \rightarrow^* \langle \text{skip}, s'_2 \rangle$, then $\forall x \in O_l. s'_1(x) = s'_2(x)$.

This definition captures value channels. As explained in the introduction, timing side channels can potentially be prevented using orthogonal techniques if necessary.

The example in Fig. 2 illustrates this property. The targetSize procedure determines the size of the target audience for a marketing campaign as the number of people in a given array of households that satisfy certain criteria. The procedure uses two worker threads that each iterate over half of the households, determine the number of target persons in each household, and add that number to a shared counter c . We assume that the number of households in the input as well as how many members of a household are in the target audience (i.e., the result value of countTargets) is low. However, the execution time of countTargets may depend on high data for a variety of reasons. For example, the implementation in Fig. 2 consists of a simple look-up in a hash map. Nevertheless its execution time may depend on hash collisions with other keys whose presence in the map is secret. Our goal is to prove that the output of targetSize is low.

2.2 Commutativity-Based Information Flow Reasoning

¹We limit our presentation to two security labels, high and low, instead of the general case of having arbitrary lattices of labels; however, techniques for verifying information flow security with two levels can be used to verify programs with arbitrary finite lattices by performing the verification multiple times, once for every element of the lattice.

```

procedure targets(households) returns (res) {
  // ensures: res is low
  n := |households|
  m := createMap()
  (worker(households, 0, n/2, m)
   ||
   worker(households, n/2, n, m))
  res := sort(toList(keys(m)))
}

procedure worker(households, f, t, m) {
  for (i in f..t-1) {
    adr, rsn := select(household[i])
    atomic:
      m.put(adr, rsn)
  }
}

```

Fig. 3. Example program: Multiple threads add values to a shared map.

The high data accessed in `countTargets` may affect the execution time of the procedure and, thereby, the thread schedule and the intermediate values of the shared counter `c`. Due to this internal timing channel, the value of `c` must be considered high *during* the execution of the worker threads. However, *after* the worker threads have terminated, the counter value can safely be considered to be low for two reasons. First, as per our assumption, each individual value added is low. Second, although high information may have affected the *order* in which different values are added to the counter, this order does not affect the final counter value because the updates *commute*. This example illustrates the central insight of this paper: *Internal timing channels do not affect the final outputs of a program if the modifications performed by different threads commute.*

Building on this insight, we present a verification technique for proving non-interference of concurrent programs. This technique enforces four central properties (in addition to standard checks that are sufficient to ensure non-interference for sequential programs):

- (1) Low initial value: For any shared data structure that is to be modified by multiple threads, the initial contents of the data structure is low.
- (2) Modifications are low: High data does not influence whether or not a thread performs a modification of the shared data at any given point in the program.
- (3) Modification arguments are low: Any data inserted into the shared data structure is low.
- (4) Commutativity: The modifications performed by different threads commute with each other.

These four properties are sufficient to ensure that, for any pair of executions of the verified program with identical low but potentially different high inputs, the *final* value of any shared data structure is identical after all threads have finished modifying it.

To understand why, consider two such executions: By property (1), the data structure will have the same contents in both executions when it is initially shared. By property (2), in both executions, the same number of modifications will be performed on the data structure (e.g., in our example, the same number of values will be added). By property (3), for each modification in the first execution, there is a matching modification in the second execution that uses the same arguments (e.g., if the value x has been added to the counter in the first execution at some point, then x is also added at some point in the second execution); the only difference is in the order of modifications. However, by property (4), reordering the modifications leaves the final value unchanged. Consequently, the final contents of the data structure are the result of performing the same commutative operations with the same arguments on the same initial value, and must therefore be identical.

These four properties form the basis of our verification technique. In the rest of this section, we will build on and expand on this basic idea: We will relax some of the properties to make our verification technique more complete and more widely applicable, and we will explain how we check each of the four properties on the program to be verified.

2.3 Abstract Commutativity

The methodology described so far is sound, but not sufficiently complete: First, it requires that *all* information stored in shared data structures must be low after all concurrent modifications, even information that is never leaked to a public output. Second, most mutating operations on common data structures do not commute and therefore do not satisfy our property (4) above.

As an example, consider a variation of the previous example in Fig. 3: Now, the shared data structure is a map, and each worker extracts a key-value pair per household, where the address (the key) is low, but the reason why an address was selected (the value) is high. Here, different invocations of put do not always commute: If two threads put the same key but different values, then the later put-operation will “win”, and its value will overwrite the previous value for said key. As a result, the final contents of the map allow an observer to conclude which put happened later and, thereby, draw conclusions about the secrets that cause the different thread interleavings.

However, the procedure `targets` does not actually return the entire contents of the map, but only (a sorted list representation of) its key set. Therefore, differences in values do not affect `targets`’ public output. Scenarios where programs expose some public views on the data they maintain are common in practice. To accommodate them, we allow programmers to define an *abstract view* of the shared data structure that is guaranteed to be low and, therefore, allowed to be leaked. By focusing on the relevant part of a shared data structure, we no longer have to demand commutativity of all concurrent modifications, but only *abstract commutativity*: commutativity modulo the abstract view. That is, we require that, if the *abstract* view of the shared data is low initially, then switching the order of any two modifications on the data does not affect the final *abstract* view of the data structure. This is fulfilled for the example in Fig. 3: Different put-operations do not commute w.r.t. the entire contents of the map, but they do commute w.r.t. the map’s key set.

Since we require only the abstract view of shared data to be low, we no longer have to demand that *all* arguments of map modifications have to be low: arguments that do not affect the abstract view of the data may contain high information. In the example, this allows us to insert high *values* into the map as long as all inserted *keys* are low, since values do not affect the key set. We use preconditions on mutating operations to specify which arguments must be low to ensure that the abstract view of the data remains low.

To summarize, we relax our four central properties as follows:

- (1) Low initial *abstract* value: The abstract view of the shared data is low when first shared.
- (2) Modifications are low: Unchanged².
- (3) Modification arguments fulfill sufficient precondition: (a) Arguments of each modification fulfill a given precondition that (b) ensures that the abstract view of the shared data remains low after the modification.
- (4) *Abstract* Commutativity: All modifications commute w.r.t. the abstract view of shared data.

By enabling these relaxations, abstract commutativity allows us to handle many more practical examples than standard commutativity. To demonstrate that the relaxed conditions are sufficient, we can make a similar argument as before to argue that the abstract view of each shared data structure will be low after all concurrent modifications on it have been performed.

2.4 Resource Specifications

Proving our four properties directly on the level of heap-manipulating programs is difficult: For example, showing that two modifications commute (even without an abstraction) on actual program states means proving the equivalence of two programs that perform the modifications in a different

²One could in principle restrict property (2) to modifications that affect the abstract view of the shared data structure. We omit this optimization since operations that do not affect the abstract value commute trivially and, therefore, do not complicate verification significantly.

$$\begin{array}{ll}
type_{MK}(v) \equiv K \multimap V & type_{MD}(v) \equiv K \multimap V \\
\alpha_{MK}(v) \equiv dom(v) & \alpha_{MD}(v) \equiv v \\
SharedActions_{MK} \equiv \{PUT\} & SharedActions_{MD} \equiv \emptyset \\
UniqueActions_{MK} \equiv \emptyset & UniqueActions_{MD} \equiv \{PUT_1, PUT_2\} \\
f_{PUT}(v, \langle key, val \rangle) \equiv m[k \mapsto v] & f_{PUT_i}(v, \langle key, val \rangle) \equiv m[k \mapsto v] \\
pre_{PUT}(\langle key, val \rangle) \equiv Low(key) & pre_{PUT_i}(\langle key, val \rangle) \equiv Low(key) \wedge Low(val) \\
& \wedge key \in range_i
\end{array}$$

Fig. 4. Left: Complete resource specification for the map example in Fig. 3, where K and V are the types of the keys and values in the map, respectively. Right: Alternative resource specification that allows two different threads to perform updates only in their own range of keys, which does not overlap with that of the other thread. We assume here that $range_1$ and $range_2$ do not overlap.

order. If these operations contain complex steps like pointer arithmetic, loops, or memory allocation (all of which might potentially be performed in real implementations of put methods on maps), it is non-trivial to even define equivalence of states (e.g., because of non-deterministic allocation), and even more difficult to perform such proofs [Koskinen and Bansal 2021].

We therefore do not check all of our properties *directly* on program states, but instead on pure mathematical values. We exploit the fact that, when working with separation logics, it is standard practice in the specification of data structures to use separation logic *predicates* that relate the contents of data structures to pure values [Parkinson and Bierman 2005]. For example, linked lists are typically specified using a predicate of the form $list(p, s)$, where p is the pointer to the start of the list and s is a mathematical sequence describing the contents of the list; $list(p, s)$ holds iff p points to a list whose contents are s . Methods that manipulate the data structures are then specified in terms of those pure values: for example, an append method of a list implementation would usually require the list predicate in its precondition for some abstract value s , and return it in its postcondition, with an abstract value that was extended by the appended value. We will assume that such specifications exist, and require additionally that the predicate uniquely determines the abstract value, which is typically the case for existing definitions. Thus, we exploit existing standard verification constructs to map our used data structures to pure values, enabling us to check properties (3) and (4) on the level of those pure values.

To check property (4), we must first determine which modifications are performed on the shared data. Scanning the entire program for such modifications would not be modular. Instead, we associate a set of legal operations with a data structure when we share it, and subsequently enforce that all modifications that are performed by any thread correspond to one of the legal operations.

We express these aspects of a shared data structure using a novel specification construct: a *resource specification* declares a *pure* data type (for the contents of a data structure in the standard separation logic style), an *abstract view* in the form of an abstraction function α that maps a value of the pure data type to another mathematical value that characterizes which aspects of the data structure ultimately have to be public, and a set of *actions* that may be performed on the shared data structure. Each action comes with a function that defines how the action modifies the (pure) value of the data structure and with a (relational) precondition that restricts the arguments of the action such that, after performing the action, the abstract view of the data structure remains low.

Fig. 4 (left) shows a resource specification for our map example: Its type is a partial mapping from keys to values, its only legal action is PUT, which updates the mapping, and the precondition of PUT requires the argument key (but not the value) to be low using an assertion which we will formally introduce later. We will discuss later what it means for an action to be shared or unique.

Crucially, such a resource specification is independent of any specific client program and of any specific implementation of a map data structure; it can be used in any program that uses a shared map whose key set is low in the end, and it can be combined with any separation logic predicate for arbitrary map implementations (e.g., tree-based, list-based, hash-based, etc.).

A resource specification is *valid* if all its actions commute modulo its abstraction function (property (4)), and every action’s precondition suffices to ensure the low-ness of the abstract view of the data (property (3b)); a formal definition of validity will follow in Sec. 4. The proof of validity has to be done only once per resource specification, and can be reused for different programs.

Crucially, our commutativity criterion (property (4)) enforces (abstract) *pairwise* commutativity of the actions in the resource specification, which avoids the (huge) effort of enumerating and comparing all possible interleavings of actions in a program.

2.5 Program Verification

Given a program and a resource specification that satisfies properties (3b) and (4), it remains to check properties (1), (2), and (3a) to prove that the program adheres to its resource specification.

Our logic for doing so, COMMCSL, is a concurrent separation logic (CSL) [O’Hearn 2004]. To enable reasoning about low and high values, we phrase COMMCSL as a relational logic that proves properties of two executions of the same program³. Like other concurrent separation logics, it requires showing that programs are data race free, by proving that each thread operates on its own partial heap separate from those of all other threads.

Invariants. As is standard in separation logics, heap locations that may be modified by different threads must be part of a shared resource, which is associated with an invariant that describes and constrains the resource’s partial heap. In addition, we use the invariant to also map the shared data to some pure value, as described before; that is, the invariant is used to connect a heap-based shared data structure in the verified program to a value of the type defined in the resource specification. We denote the invariant as $I(v)$, where v denotes the pure value the invariant maps its heap to. As explained before, for a linked list, $I(v)$ would typically be some predicate $list(p, v)$ for some pointer p and sequence v . When initially sharing the resource, the invariant $I(v)$ must be established for some v . We check property (1) by enforcing that at this point, $\alpha(v)$ is low.

Atomic modifications. As explained in Sec. 2.4, we check property (4) on the level of abstract actions rather than concrete implementations. For this to be sound, *all* modifications of a shared data structure must be reflected by one of its actions, which we verify as follows.

In COMMCSL (as in normal CSL), threads may modify the shared resource only in atomic-blocks; when entering such a block, they obtain the invariant, and on leaving the block, they have to re-establish the invariant and give it up again⁴. We use this mechanism to check that a program modifies a shared data structure only via the legal actions of the associated resource specification. Since a shared data structure may be modified only within an atomic-block, we can impose the following proof obligation: If, at the beginning of an atomic-block, $I(v)$ holds for some v , then at its end, $I(v')$ holds for some v' that is the result of applying one of the legal actions to the old value v , i.e., $v' = f_a(v, arg)$ for some action a and argument arg .

³While our technique could in principle also reason about more than two executions at a time, we focus only on pairs of executions, since that is sufficient for proving non-interference.

⁴*Obtaining* the invariant adds the partial heap described by the invariant to the current heap; *giving it up* amounts to checking that the current heap can be split into a partial heap that satisfies the invariant and a remainder, and removing the former. These logical steps are sometimes called *producing* and *consuming*, or *inhaling* and *exhaling* the invariant.

Guards. Let's turn to the remaining two properties, (2) and (3a). We could enforce property (2) by proving that there are no atomic-blocks under high guards (i.e., inside conditionals or loops whose conditions are high). However, this check would be overly conservative because it rules out implementations where the final value of a shared data structure is low even though intermediate states depend on a secret; our evaluation demonstrates that such examples occur in practice.

Therefore, we follow an alternative approach: we *record* which actions are performed on a shared resource, and then check property (2) retroactively on the recorded actions when the resource is unshared (and all concurrent modifications must therefore have finished). We will follow the same approach for property (3a) and, thus, also record the arguments of each performed action, such that we can check action preconditions later.

For each action, we introduce a separation logic resource, which we call a *guard*, to record how often and with which arguments this action has been performed. Like other separation logic resources, guards can be transferred between methods, split into fractional parts such that they can be shared between threads, and subsequently be recombined. Guards are parameterized with a multiset of arguments, which records the arguments of the actions performed on a shared resource so far. These parameters reflect *which* actions have been performed (as well as their arguments), but not their order, which is not known due to the influence of scheduling.

When sharing a resource, we obtain a guard for each action that is legal according to the resource specification. The parameters of these guards are initially empty because no actions have been performed yet. Performing an action then imposes a proof obligation that some fraction of the respective guard is held, and adds the argument to the guard's multiset.

Guards allow us to check properties (2) and (3a) at the time when a resource is unshared. Unsharing requires all guards to be held, so that all performed actions are known. Unsharing then consumes the guards, so that no threads have the ability to perform any actions after this point. We show property (2) simply by proving that the *cardinality* of the argument multiset (i.e., how often each action has been performed) is low.

To ensure property (3a), we need to prove that each execution of an action satisfies its precondition. Since these preconditions are in general relational (requiring, for instance, an argument to be low), this proof must match the execution of an action in one run of the program with an execution in the other run. Proving preconditions retroactively when we unshare a resource gives the proof more freedom which executions of an action to match, which makes the proof technique more complete. Retroactive checking of preconditions is possible because they are expressed only over argument values and therefore state-independent.

Concretely, when unsharing, we require showing for each action a that there is a bijection between the elements of the multiset of arguments s in one program execution and its elements in the other execution, such that each pair of elements together fulfills the relational precondition of the action. We denote this fact via the assertion $PRE_a(s)$. For the PUT action of the map example, this means checking that there is a bijection that maps every key-value pair in the multiset in the first execution to a pair with the same key (but potentially a different value) in the second execution, since the precondition of PUT requires the keys to be low (but not the values).

2.6 Verification of the Map Example

We now show the entire proof for the map example from Fig. 3 in Fig. 5 (where we assume that our resource specification from Fig. 4 (left) is valid). The code indentation in targets indicates where the resource is shared. When sharing the resource, we must establish and give up the invariant $I(v)$, which we instantiate to $Map(m, v)$ (a separation logic predicate for a map, which relates the map's contents to the pure partial mapping v), and show that $\alpha(v)$ is low (property (1)), which is the case because the map is empty in both executions. We then share the map, meaning that we

```

1  procedure targets(households) {
2    // ensures  $\exists v'. \text{List}(\text{res}, v') * \text{Low}(v')$ 
3    n := |households|
4    m := createMap()
5    {Map(m, empty)}
6     $\models \{ \text{Map}(m, \text{empty}) * \text{Low}(\text{dom}(\text{empty})) \}$ 
7    // share
8    {guardPUT( $\emptyset^\#, 1$ )}
9     $\models \{ \text{guard}_{\text{PUT}}(\emptyset^\#, \frac{1}{2}) * \text{guard}_{\text{PUT}}(\emptyset^\#, \frac{1}{2}) \}$ 
10   worker(households, 0, n/2, m) ||
11   worker(households, n/2, n, m)
12   { $\exists s_1, s_2. \text{guard}_{\text{PUT}}(s_1, \frac{1}{2}) * \text{PRE}_{\text{PUT}}(s_1) * \text{guard}_{\text{PUT}}(s_2, \frac{1}{2}) * \text{PRE}_{\text{PUT}}(s_2)$ }
13    $\models \{ \exists s. \text{guard}_{\text{PUT}}(s, 1) * \text{PRE}_{\text{PUT}}(s) \}$ 
14   // unshare
15   { $\exists v. \text{Map}(m, v) * \text{Low}(\text{dom}(v))$ }
16   return sort(toList(keys(m)))
17   { $\exists v, v'. \text{Map}(m, v) * \text{Low}(\text{dom}(v)) * \text{List}(\text{res}, v') * \text{Low}(v')$ }
18 }

```

```

1  procedure worker(households, f, t, m) {
2    // requires guardPUT( $\emptyset^\#, \frac{1}{2}$ )
3    // ensures  $\exists s'. \text{guard}_{\text{PUT}}(s', \frac{1}{2}) * \text{PRE}_{\text{PUT}}(s')$ 
4    {guardPUT( $\emptyset^\#, \frac{1}{2}$ )}
5     $\models \{ \text{guard}_{\text{PUT}}(\emptyset^\#, \frac{1}{2}) * \text{PRE}_{\text{PUT}}(\emptyset^\#) \}$ 
6    for (i in f..t-1) {
7      { $\exists s'. \text{guard}_{\text{PUT}}(s', \frac{1}{2}) * \text{PRE}_{\text{PUT}}(s')$ }
8      adr, rsu := select(household[i])
9      { $\exists s'. \text{guard}_{\text{PUT}}(s', \frac{1}{2}) * \text{PRE}_{\text{PUT}}(s') * \text{Low}(\text{adr})$ }
10     atomic:
11       { $\exists v. \text{Map}(m, v)$ }
12       m.put(adr, rsu)
13       { $\exists v. \text{Map}(m, v[\text{adr} \mapsto \text{rsu}])$ }
14        $\models \{ \text{Map}(m, f_{\text{PUT}}(v, \langle \text{adr}, \text{rsu} \rangle)) \}$ 
15       { $\exists s'. \text{guard}_{\text{PUT}}(s' \cup^\# \{ \langle \text{adr}, \text{rsu} \rangle \}^\#, \frac{1}{2}) * \text{PRE}_{\text{PUT}}(s') * \text{pre}_{\text{PUT}}(\langle \text{adr}, \text{rsu} \rangle)$ }
16        $\models \{ \exists s'. \text{guard}_{\text{PUT}}(s' \cup^\# \{ \langle \text{adr}, \text{rsu} \rangle \}^\#, \frac{1}{2}) * \text{PRE}_{\text{PUT}}(s' \cup^\# \{ \langle \text{adr}, \text{rsu} \rangle \}^\#) \}$ 
17     }
18   }
19 }
20 }

```

Fig. 5. Proof outline for the example from Fig. 3, verified against the resource specification from Fig. 4 (left). $I(v)$ is defined to be $\text{Map}(m, v)$, where Map is assumed to be a pre-existing separation logic predicate that relates the contents of a map m to a mathematical value v . Superscript # denotes multiset operations. $\text{guard}_{\text{PUT}}(s, f)$ denotes an f -fraction of the guard for action PUT with argument multiset s .

obtain the guard for the PUT action (the only action in our resource specification), with an empty argument multiset. We split this guard into two parts and give each part to a worker.

Inside worker, we prove the loop invariant $\text{guard}_{\text{PUT}}(s', \frac{1}{2}) * \text{PRE}_{\text{PUT}}(s')$ for some s' (the argument multisets in proofs are typically existentially quantified, s.t. it is not necessary to track actual values in specifications). That is, we collect all arguments with which we performed the PUT action in s' and show that each such argument satisfies PUT's precondition. As explained above, we are *required* to prove $\text{PRE}_{\text{PUT}}(s')$ only later when we unshare the resource, but in this example, it is convenient to maintain this property throughout. Initially, $\text{PRE}_{\text{PUT}}(s')$ holds trivially since s' is empty.

When the worker enters an atomic-block, it obtains the invariant $I(v)$ for some v (that is, $\text{Map}(m, v)$), and has to show at the end of the block that $I(v')$ holds, where $v' = f_{\text{PUT}}(v, \text{arg})$ and $\text{arg} = \langle \text{adr}, \text{rsu} \rangle$. We assume that the specification of the put-method is sufficiently strong to prove this. The proof rule for atomic-blocks adds arg to the argument multiset of the guard. Since we can prove at this point that $\text{pre}_{\text{PUT}}(\text{arg})$ holds (i.e., adr is low), and $\text{PRE}_{\text{PUT}}(s')$ held for the old argument multiset s' , we can show that the new argument multiset $s'' = s' \cup^\# \{ \text{arg} \}^\#$ also fulfills $\text{PRE}_{\text{PUT}}(s'')$, and thus we maintain our loop invariant (which is also the postcondition of worker).

Back in targets, we unshare the resource, which consumes the guard. To do so, we must recombine the fractional guards into a single one, whose argument multiset s is the union of the argument multisets of the individual guards, s_1 and s_2 . Now we *must* prove that $\text{PRE}_{\text{PUT}}(s)$ holds, which follows from the postcondition of the workers, once the guards are combined. We may now assume that the abstract view $\alpha(v)$ of the new value v of the shared map is low, i.e., we know that the map's key set is low and can prove the postcondition of targets (assuming an appropriate specification for the procedures that extract the key set from the map and turn it into a sorted list).

2.7 Unique Actions

The technique we presented so far requires that *all* modifications of the shared resource commute (modulo abstraction). However, in asymmetric thread collaborations, there are often actions that

$$c ::= x := e \mid x := [e] \mid [e] := e \mid x := \text{alloc}(e) \mid \text{skip} \mid c; c \mid \text{if } (b) \text{ then } \{c\} \text{ else } \{c\} \\ \mid \text{while } (b) \text{ do } \{c\} \mid c \parallel c \mid \text{atomic } c$$

Fig. 6. Programming language. e ranges over integer-typed expressions, b over boolean-typed ones.

are performed by only one thread (e.g., in a producer-consumer scenario, a single producer thread might add values to a queue, whereas multiple consumer threads read and remove values from the queue); we call such actions *unique* actions (as opposed to *shared* actions that are performed by multiple threads). Since the order of applications of a unique action does not depend on scheduling, it is not necessary for unique actions to commute with themselves.

As an example, consider a variation of the map example where different workers work on different ranges of keys, and where both the keys and values PUT into the map are low. In such a program, it is never the case that a PUT of one thread overwrites a value that was PUT by a different thread, and therefore differences in timing do not lead to different final map contents.

We can support such examples by defining an alternative resource specification that declares n different PUT actions, where the precondition of each action PUT_i requires its argument key to be in a range separate from that of all other PUT actions (and both key and value to be low), as shown in Fig. 4 (right) for $n = 2$. We declare each of these actions to be unique, which means that they may be performed only by a single thread. For resource specification validity, unique actions are required to (abstractly) commute with all actions *except* themselves. That is, for example, PUT_1 has to commute with PUT_2 , but *not* with itself. We can now define α to be the identity function (i.e., use no abstraction at all), and prove programs secure even if they leak the entire map⁵.

During program verification, we do not allow splitting the guards for unique actions, which ensures that only a single thread may perform them. In addition, guards for unique actions record the *sequence* of previous arguments, since their overall order is now no longer dependent on scheduling, and $\text{PRE}_a(s)$ for a unique action a and argument sequence s requires that the length of s is low, and for each index i , the values of $s[i]$ in both executions fulfill the action's precondition.

3 LOGIC

In this section, we formalize our technique in COMMCSL, a relational concurrent separation logic with support for commutativity-based reasoning.

3.1 Language

We formalize our logic for an imperative, concurrent language with a mutable heap, whose commands are defined in Fig. 6. We assume as given some expression language for integer-typed expressions e and boolean expressions b , with the usual operations. Commands like assignments, conditionals, loops, parallel composition and sequential composition are standard. Programs interact with the heap via read and store commands $x := [e]$ and $[e] := e$, as well as $x := \text{alloc}(e)$, which allocates a single heap location, initializes its value to e , and assigns the resulting pointer to x ; an extension to allocating multiple heap locations at once is straightforward. Like other CSLs, we formalize our logic for a single shared resource; threads can use the command `atomic` c to atomically execute the command c while having access to the shared resource. However, our approach is not limited to having a single shared resource, and multiple resources are supported in our implementation analogous to other CSLs [Brookes 2004; O'Hearn 2004].

⁵As an aside, this example also shows that the same data structure can have different resource specifications that have different demands and give different guarantees for the sensitivity of the final result.

Program states have the form $\langle s, h \rangle$, where the store s is a map from names to (integer) values and the heap h is a partial map from locations (natural numbers) to integer values. Expression evaluation is deterministic and total, meaning in particular that references to uninitialized variables evaluate to some default value; we denote the value of expression e in store s by $\llbracket e \rrbracket(s)$. Program configurations have the form $c, \langle s, h \rangle$ or `abort`; the latter represents a failed computation (which will never be reachable for verified programs). We use a small-step semantics with transitions of the form $c, \langle s, h \rangle \rightarrow c', \langle s', h' \rangle$ (for non-aborting steps). Our semantics is identical to that of other concurrent separation logics [Vafeiadis 2011]; its rules are shown in App. A.1.

3.2 Resource Specifications

To simplify the formalization, we assume that our resource specification allows arbitrarily many unique actions but only *one* shared action, which is not restrictive, since one can merge multiple shared actions into one s.t. the argument selects which action to perform. Given this assumption, a resource specification for a resource value of type T has the form $\langle \alpha, f_a, F_{au} \rangle$, where α is the abstraction function of type $T \rightarrow T_\alpha$ (for some mathematical type T_α), f_a is the shared action, whose precondition we denote as pre_{a_s} , and F_{au} is a family of unique actions, indexed by some finite type I ; we will generally write f_{a_i} as a synonym for $F_{au}(i)$ for all $i \in I$, and denote the precondition of unique action f_{a_i} by pre_{a_i} . For each action a , f_a is a function of type $T \rightarrow T_{arg_a} \rightarrow T$, where T_{arg_a} is the argument type of the function. Both preconditions are *relational* preconditions on the argument of the action, i.e., they can require that (aspects of) the arguments are low. Thus, they are boolean-typed functions that take two arguments of type T_{arg_a} , where the two arguments denote the values of the argument in the two executions. Note that the preconditions *cannot* constrain the resource value itself; we explain this limitation and show how to work around it in App. D.

A resource specification is *valid* iff (A) every action's relational precondition is sufficient to preserve the low-ness of the abstract view of the resource value (property (3b) from Sec. 2), and (B) all *relevant* pairs of actions commute w.r.t. the abstract view (property (4) from Sec. 2). Relevant pairs of execution are the shared action paired with all actions including itself, and every unique action paired with all unique actions *except* itself; this reflects the fact that unique actions do not have to commute with themselves. Formally, we define validity as follows:

Definition 3.1. A resource specification $\langle \alpha, f_a, f_{au} \rangle$ is valid iff

- (A) For all actions a , values v, v' , and arguments arg, arg' , if $\alpha(v) = \alpha(v')$ and $pre_a(arg, arg')$, then $\alpha(f_a(v, arg)) = \alpha(f_a(v', arg'))$.
- (B) For all pairs of actions a, a' in $\{(a_s, a_s)\} \cup \{(a_s, a_i) \mid i \in I\} \cup \{(a_i, a_j) \mid i, j \in I \wedge i \neq j\}$ and all arguments arg, arg' , if $\alpha(v) = \alpha(v')$ then $\alpha(f_{a'}(f_a(v, arg), arg')) = \alpha(f_a(f_{a'}(v, arg'), arg))$.

3.3 Extended Heaps

While our semantics works on ordinary program heaps h , we define our assertions and our logic on *extended heaps*, an enriched notion of heaps that represent both fractional permissions [Boyland 2003] and guards: An extended heap gh is a triple $\langle ph, gs, Gu \rangle$. ph is a standard *permission heap* [Bornat et al. 2005; Vafeiadis 2011] that can express *partial ownership* of a heap location, that is, a partial map from locations to pairs $\langle r, v \rangle$ of positive rational numbers of at most 1 and values v . For example, a permission heap that maps location l to $\langle v, \frac{1}{2} \rangle$ denotes a half permission to the heap location l , where value v is stored. Partial ownership allows multiple threads to concurrently *read* the same heap location, since reading a location requires only some positive permission amount, whereas modifying its value requires a permission of 1.

The shared action guard state gs and the family of unique action guard states Gu , where $Gu(i) = gu_i$, are specific to our technique. A guard for an action represents the *right* to perform that action, and it tracks the arguments with which said action has already been performed.

Every *unique* guard state gu_i is either \perp or a *sequence* of argument values of the unique action a_i , and represents (when a resource is shared) the entire sequence of arguments with which the unique action has been performed. Since only one thread is allowed to execute a unique action, the order of executions is known and, thus, we can track the arguments in a sequence (as opposed to a multiset for shared actions). Moreover, unique guard states cannot be split or combined; that is, the thread that performs the unique action will have the entire sequence of argument values in its unique guard state, and all other threads will have a guard state of \perp . Thus, when adding two unique guard states, if one has a non- \perp value, then the other must be \perp , otherwise addition is undefined (see App. B.1 for the formal definition of guard heap addition).

Addition of guard heap families is defined pointwise. We write \perp for guard heap families that are \perp for all i , and $[i \mapsto gu]$ for a guard heap family whose value is gu for i , and \perp for all other indices.

The *shared* guard state gs is either \perp or a pair $\langle r, args_s \rangle$, where r is a positive rational number of at most 1 and $args_s$ is a *multiset* of argument values. The shared action may be performed by multiple threads if each of those threads has a positive fractional shared guard state. Thus, the shared guard state represents the (potentially partial) knowledge of the multiset of arguments with which the shared action has been executed so far. If r is 1, then $args_s$ denotes *all* arguments with which the action has been performed so far. On the other hand, two threads might have guard states with $r = \frac{1}{2}$ each and respective arguments $args_{s1}$ and $args_{s2}$, which each contain all arguments with which *that thread* has performed the shared action; then, in total, the action has been performed with the union of the multisets $args_{s1} \cup^\# args_{s2}$. In general, when adding the arguments of two non- \perp shared guard states, we take the union of the argument multisets, or of one state if the other is \perp (see App. B.1). The sum of two permission heaps $ph \oplus ph'$ is standard and shown in App. B.1; intuitively, permission amounts are added (to a value of at most one) and values are unchanged.

The sum of two extended heaps $\langle ph, gs, gu \rangle \oplus \langle ph', gs', gu' \rangle$ is defined as $\langle ph \oplus ph', gs \oplus gs', gu \oplus gu' \rangle$ iff the sums of all its components are defined. Finally, an extended heap can be *normalized* to a normal heap (which we denote by $norm(gh)$), by taking the permission heap and removing the permission amounts. That is, the normalized form of $\langle ph, gs, gu \rangle$ has the domain $dom(ph)$, and for each location $l \in dom(ph)$, the normalized heap has the value v s.t. $ph(l) = \langle _, v \rangle$.

3.4 Assertions

Our assertion language is defined as follows:

$$P, Q ::= emp \mid b \mid e \mapsto^r e \mid P * Q \mid P \wedge Q \mid \exists x. P \mid sguard(r, e) \mid uguard_i(e) \mid b \Rightarrow P \mid Low(e)$$

where P and Q range over assertions, r ranges over positive rationals up to 1, b ranges over boolean expressions, and e ranges over expressions (including sequence- and multiset-typed expression, to describe the arguments of the two kinds of guard-assertions). Assertions are *relational*, i.e., their validity (Fig. 7) is defined over pairs of states, allowing them to express that expressions are low, that is, equal in both states.

All basic separation logic assertions have their usual meaning, but applied to both states. For instance, standard points-to-assertions $e_1 \mapsto^r e_2$ represent a permission amount of r to location e_1 with a value of e_2 . Separating conjunctions $P * Q$ hold if the extended heaps of *both* states can be separated into partial extended heaps that fulfill the conjuncts. Existentials are interpreted such that there can be *different* values for the quantified variable in both states, that is, $\exists x. e \mapsto^1 x$ expresses that e points to potentially different values in the two states (i.e., it might be high).

$(s_1, gh_1), (s_2, gh_2) \models emp$	\iff	$dom(ph_1) = \emptyset \wedge dom(ph_2) = \emptyset$
$(s_1, gh_1), (s_2, gh_2) \models b$	\iff	$\llbracket b \rrbracket_{s_1} = 1 \wedge \llbracket b \rrbracket_{s_2} = 1$
$(s_1, gh_1), (s_2, gh_2) \models e_1 \mapsto^r e_2$	\iff	$gh_1 = \{ \llbracket e_1 \rrbracket_{s_1} \mapsto \langle r, \llbracket e_2 \rrbracket_{s_1} \rangle \} \wedge gh_2 = \{ \llbracket e_1 \rrbracket_{s_2} \mapsto \langle r, \llbracket e_2 \rrbracket_{s_2} \rangle \}$
$(s_1, gh_1), (s_2, gh_2) \models P * Q$	\iff	$\exists gh'_1, gh''_1, gh'_2, gh''_2. gh_1 = gh'_1 \oplus gh''_1 \wedge gh_2 = gh'_2 \oplus gh''_2 \wedge$ $(s_1, gh'_1), (s_2, gh'_2) \models P \wedge (s_1, gh''_1), (s_2, gh''_2) \models Q$
$(s_1, gh_1), (s_2, gh_2) \models P \wedge Q$	\iff	$(s_1, gh_1), (s_2, gh_2) \models P \wedge (s_1, gh_1), (s_2, gh_2) \models Q$
$(s_1, gh_1), (s_2, gh_2) \models \exists x. P$	\iff	$\exists v_1, v_2. (s_1[x \mapsto v_1], gh_1), (s_2[x \mapsto v_2], gh_2) \models P$
$(s_1, gh_1), (s_2, gh_2) \models sguard(r, e)$	\iff	$gs_1 = \langle r, \llbracket e \rrbracket_{s_1} \rangle \wedge gs_2 = \langle r, \llbracket e \rrbracket_{s_2} \rangle \wedge$ $Gu_1 = \perp \wedge Gu_2 = \perp \wedge dom(ph_1) = \emptyset \wedge dom(ph_2) = \emptyset$
$(s_1, gh_1), (s_2, gh_2) \models uguard_i(e)$	\iff	$Gu_1 = [i \mapsto \llbracket e \rrbracket_{s_1}] \wedge Gu_2 = [i \mapsto \llbracket e \rrbracket_{s_2}] \wedge$ $gs_1 = \perp \wedge gs_2 = \perp \wedge dom(ph_1) = \emptyset \wedge dom(ph_2) = \emptyset$
$(s_1, gh_1), (s_2, gh_2) \models b \Rightarrow P$	\iff	$\llbracket b \rrbracket_{s_1} = \llbracket b \rrbracket_{s_2} \wedge \llbracket b \rrbracket_{s_1} \Rightarrow (s_1, gh_1), (s_2, gh_2) \models P$
$(s_1, gh_1), (s_2, gh_2) \models Low(e)$	\iff	$\llbracket e \rrbracket_{s_1} = \llbracket e \rrbracket_{s_2}$

Fig. 7. Assertion validity, where $gh_i = \langle ph_i, gs_i, Gu_i \rangle$.

The assertion $sguard(r, e_a)$ represents a fractional amount r of the guard for the shared action with the multiset of arguments e_a ; it holds in a state with store s iff the shared guard state is $\langle r, \llbracket e_a \rrbracket_s \rangle$, the permission heap is empty, and the unique guard state is \perp . Thus, a state fulfilling $sguard(r, e_a)$ can be added to any other state with a compatible shared guard state. Similarly, $uguard_i(e_a)$ represents the guard for the unique action with index i and holds iff the sequence of arguments is $\llbracket e_a \rrbracket_s$, the shared guard state is \perp , and the permission heap is empty.

Finally, the relational assertion $Low(e)$ states that expression e is low. When used in an implication $b \Rightarrow Low(e)$, it can also express *value-dependent sensitivity* [Murray et al. 2018]. For example, a data structure might contain pairs of booleans and other values, where the boolean expresses the sensitivity of the other value; then, at runtime, clients could retrieve values from the data structure and e.g. output them to different channels based on their sensitivity.

We define $PRE(e)$ in terms of other assertion constructs⁶. $PRE_s(e)$, for the shared action, expresses that there is a bijection between the elements of the multiset-typed expression e in one state and the elements of e in the other state, s.t. each element in one state and its corresponding element in the other state together fulfill the relational precondition pre_{a_s} of the shared action. Formally, we define this notion recursively, by stating that there is an element x in e in one execution, and an element x' in e in the second execution, s.t. they together fulfill pre_{a_s} , and after removing both values from their respective multisets, $PRE(e)$ holds again for the resulting smaller multisets:

Definition 3.2. If the argument type of the shared action is T_{arg_s} , and e is a multiset of T_{arg_s} , then

$$PRE_s(e) = \begin{cases} true & \text{if } e = \emptyset^\# \\ \exists x \in e. pre_{a_s}(x) \wedge PRE_s(e \setminus^\# \{x\}) & \text{otherwise} \end{cases} \quad (1)$$

Since pre_{a_s} is relational, $PRE_s(e)$ is also relational. Moreover, since the existential quantifier may be interpreted differently in both executions, $PRE_s(e)$ expresses the existence of a bijection between the argument multiset, as discussed in Sec. 2.5.

For each unique action, we require that for the sequence of arguments e , the elements at each index fulfill the relational precondition of the unique action:

$$PRE_i(e) = Low(|e|) \wedge \bigwedge_{j \in [0, |e|)} pre_{a_i}(e[j]) \quad (2)$$

We call an assertion *unary* if it does not restrict pairs of state relative to one another. That is, P is unary if for all s_1, gh_1, s_2, gh_2 , if $(s_1, gh_1), (s_1, gh_1) \models P$ and $(s_2, gh_2), (s_2, gh_2) \models P$ then

⁶In our Isabelle formalization, $PRE(e)$ is actually defined as another assertion in the assertion language.

also $(s_1, gh_1), (s_2, gh_2) \models P$. Note that all assertions that do not syntactically contain any $\text{Low}(e)$ -assertions (or $\text{PRE}(e)$ -assertions that are defined in terms of $\text{Low}(e)$) are unary.

Finally, $\text{noguard}(P)$ states that the assertion P holds only for states whose guard states are all \perp .

3.5 Resource Contexts

A *resource context* combines a resource specification with an assertion that defines the shared heap data structure and maps it to a single mathematical value. That is, a resource context Γ has the form $\langle \alpha, f_{as}, F_{au}, I(x) \rangle$, where the *invariant* $I(x)$ is a parameterized unary assertion defining a valid state of the shared data structure with the value x . The value of x must be uniquely defined by the invariant, i.e., we require that if $(s_1, gh_1), (s_2, gh_2) \models I(v)$ and $(s_1, gh_1), (s_2, gh_2) \models I(v')$ then $v = v'$.

Given a resource context Γ , we can now formally connect the value of the guard states to the permission heap, using the notion of *consistency*. Consistency expresses that the value of the shared data structure on the heap is a possible result of applying the recorded actions with the recorded arguments in some order. We say that a state with shared guard state $\langle r, args_s \rangle$ and unique guard states $args_i$ for every unique action a_i is consistent with a context Γ from an initial value v_0 iff (1) its shared guard state is complete (i.e., $r = 1$), (2) $I(v)$ holds in the state, for some v , and (3) if FU_i is the sequence of applications of f_{a_i} with the arguments in $args_i$ (i.e., the first application in FU_i applies f_{a_i} with argument $args_i[0]$, etc.) for all i , and similarly, FS is the sequence of applications of f_{as} with the arguments in *some* permutation of $args_s$, then there is some interleaving of FS and all FU_i s.t. applying this interleaving to the initial value v_0 results in a final value of v .

3.6 Proof Rules

Our proof rules define a judgment of the form $\Gamma_{\perp} \vdash \{P\}c\{Q\}$, where Γ_{\perp} is either \perp or some resource context Γ . Intuitively, it expresses that a command c , when executed from two states that fulfill P and either no initial shared resource (if Γ_{\perp} is \perp) or a shared resource fulfilling Γ , will not abort, and if they terminate, the resulting states will fulfill Q . The formal definition will follow in Sec. 4.

Figure 8 shows the proof rules of COMMCSL. All rules for basic commands, e.g., those for assignments, heap accesses, as well as sequential and parallel composition, and the rules for framing, existentials, and consequence, are standard. The rules for loops and conditionals are similar to other relational logics, as they have to take into account the possibility of two executions taking different branches (discussed below).

The share rule captures both the share and unshare operation. It can be used only when the current resource context is \perp , i.e., when there is no shared data structure. Once the share rule is used, threads can access a shared data structure, and must use guards to justify their actions on it and record the arguments of the actions they perform. The share rule enforces that the new context's resource specification is valid, the resource invariant initially holds, and the value x of the resource is low modulo abstraction (property (1) from Sec. 2.3). The postcondition of the conclusion expresses that the invariant will hold again, with a new value x' , whose abstract view is also low. In its premise, it gives access to the guards for the shared and the unique actions, initially with an empty argument multiset resp. argument sequences. It requires that after executing c , the guards are present again; in particular, the guard for the shared action must be present in its entirety, and therefore record all arguments of the shared action in its multiset. It also requires that PRE holds for both the sequences of arguments of the unique actions and the multiset of arguments of the shared action, meaning that we can now, retroactively, show that for each application of an action, the relational precondition was fulfilled (property (3a)).

There are two (very similar) atomic-rules, one for the shared action and one for the unique actions. Both require that the resource context is not \perp , i.e., there is currently a shared resource. Both also require that the guard for the respective action is initially present, which ensures that no

thread can modify the shared data structure without holding the appropriate guard. For the shared action, any positive fraction of the guard is sufficient; for the unique actions, the whole unsplittable guard is required. In the postcondition of the conclusion, the guard records the new argument with which the action has been performed in its multiset or sequence, respectively. The rule's premise requires showing that, assuming the invariant holds initially for some value x_o , it holds again after the atomic modification, with a new value that is the result of applying the respective action with the respective argument. We impose side conditions that require the assertions P and Q to not contain any guards (that is, any guards held before the atomic block must be framed away using the frame rule), and that the invariant $I(x)$ is precise [O'Hearn et al. 2004], i.e., fulfilled by at most one sub-heap of any given heap. This condition is not limiting in practice, see App. B.2.

The remaining main difference between our logic and standard CSL is in the rules for conditionals and loops. There are two rules for conditionals: The first can be used if the condition can be shown to be low, i.e., both executions will use the same branch. If this is not the case, the second rule must be used, which limits the postcondition Q to be unary. This restriction prevents indirect flows; for example, for the command `if (h) then { $x:=1$ } else { $x:=0$ }`, where h is high, this rule does not allow proving the postcondition $\text{Low}(x)$, which holds for each branch individually, but does not hold if one execution executes the then-branch and the other execution chooses the else-branch.

The same approach is used for loops: If the loop condition can be shown to be low initially and to remain low after each loop iteration, one can use a relational (non-unary) loop invariant. Alternatively, if the loop condition may be high, we restrict the invariant to be unary⁷.

4 SOUNDNESS

We have formalized COMMCSL and proved it sound in Isabelle/HOL ; here, we give a high-level overview of our proof.

We build on Vafeiadis's soundness proof of CSL [Vafeiadis 2011] and extend it to enable relational reasoning. Its basic idea is to define a predicate $\text{safe}_n(P, c, Q)$, which (ignoring frames and the resource invariant for now) intuitively expresses that c , when executed for n steps from any state satisfying P , will not abort and, if it terminates after those n steps, will end up satisfying Q . Hoare triples are then defined to hold if safe_n holds for all n . The same idea is used in the soundness proof for SecCSL [Ernst and Murray 2019], a concurrent separation logic for non-interference. As explained before, the central difference between our logic and SecCSL is that SecCSL enforces that two executions with identical low inputs have the same control flow, s.t. the scheduler makes the same decisions in both executions, and therefore (in an idealized scenario) no internal timing channels exist. Since we allow high data to influence control flow and therefore timing and scheduling, we have to prove that, given two executions with identical low inputs with arbitrary, potentially *different* schedules (and different control flow per thread), the program's public output will be low.

As a result, the inductive argument used in existing proofs does not work for our setting. We solve this problem by formally separating the proof of safety from the proof of relational properties. We define a predicate $\text{safe}_n(\Gamma_\perp, c, \sigma, \Sigma)$, where Γ_\perp is, as before, either \perp or a resource context Γ , σ ranges over program states with extended heaps of the form $\langle s, gh \rangle$, and Σ is a set of such program states. Intuitively, safe_n expresses (1) that executing c from σ , extended with a heap satisfying the resource invariant, for n steps will not abort, (2) that if it terminates, will result in *one* of the states in Σ (since, due to non-deterministic thread scheduling, there may be more than one final state), and (3) that if there is a shared resource, the state remains consistent w.r.t. some initial value v_0 , i.e.,

⁷It is possible to generalize these rules to a more complete rule that allows using a relational invariant for some *prefix* of the loop iterations [Sousa and Dillig 2016], but we do not do this here because loop handling is orthogonal to our main contribution.

$$\begin{array}{c}
\frac{\Gamma_{\perp} = \Gamma \Rightarrow x \notin \text{fv}(\Gamma)}{\Gamma_{\perp} \vdash \{P[e/x]\}x := e\{P\}} \text{ (ASSIGN)} \quad \frac{x \notin \text{fv}(e) \quad \Gamma_{\perp} = \Gamma \Rightarrow x \notin \text{fv}(\Gamma)}{\Gamma_{\perp} \vdash \{\text{emp}\}x := \text{alloc}(e)\{x \mapsto^1 e\}} \text{ (NEW)} \\
\\
\frac{x \notin \text{fv}(e_1, e_2) \quad \Gamma_{\perp} = \Gamma \Rightarrow x \notin \text{fv}(\Gamma)}{\Gamma_{\perp} \vdash \{e_1 \mapsto^r e_2\}x := [e]\{e_1 \mapsto^r e_2 * x = e_2\}} \text{ (READ)} \quad \frac{}{\Gamma_{\perp} \vdash \{e_1 \mapsto^1 _ \}[e_1] := e_2\{e_1 \mapsto^1 e_2\}} \text{ (WRITE)} \\
\\
\frac{\Gamma_{\perp} \vdash \{P \wedge b\}c_1\{Q\} \quad \Gamma_{\perp} \vdash \{P \wedge \neg b\}c_2\{Q\}}{\Gamma_{\perp} \vdash \{P \wedge \text{Low}(b)\} \text{if } (b) \text{ then } \{c_1\} \text{ else } \{c_2\}\{Q\}} \text{ (IF1)} \\
\\
\frac{\Gamma_{\perp} \vdash \{P \wedge b\}c_1\{Q\} \quad \Gamma_{\perp} \vdash \{P \wedge \neg b\}c_2\{Q\} \quad \text{unary } Q}{\Gamma_{\perp} \vdash \{P\} \text{if } (b) \text{ then } \{c_1\} \text{ else } \{c_2\}\{Q\}} \text{ (IF2)} \\
\\
\frac{\Gamma_{\perp} \vdash \{P \wedge b\}c_1\{P \wedge \text{Low}(b)\}}{\Gamma_{\perp} \vdash \{P \wedge \text{Low}(b)\} \text{while } (b) \text{ do } \{c_1\}\{P \wedge \neg b\}} \text{ (WHILE1)} \quad \frac{\Gamma_{\perp} \vdash \{P \wedge b\}c_1\{P\} \quad \text{unary } P}{\Gamma_{\perp} \vdash \{P\} \text{while } (b) \text{ do } \{c_1\}\{P \wedge \neg b\}} \text{ (WHILE2)} \\
\\
\frac{\Gamma_{\perp} \vdash \{P\}c_1\{R\} \quad \Gamma_{\perp} \vdash \{R\}c_2\{Q\}}{\Gamma_{\perp} \vdash \{P\}c_1; c_2\{Q\}} \text{ (SEQ)} \quad \frac{}{\Gamma_{\perp} \vdash \{P\} \text{skip}\{P\}} \text{ (SKIP)} \\
\\
\frac{\Gamma_{\perp} \vdash \{P_1\}c_1\{Q_1\} \quad \Gamma_{\perp} \vdash \{P_2\}c_2\{Q_2\} \quad \text{fv}(P_1, c_1, Q_1) \cap \text{mod}(c_2) = \emptyset \quad \text{fv}(P_2, c_2, Q_2) \cap \text{mod}(c_1) = \emptyset \quad \Gamma_{\perp} = \Gamma \Rightarrow \text{fv}(\Gamma) \cap \text{mod}(c_1, c_2) = \emptyset \quad P_1 \text{ is precise or } P_2 \text{ is precise}}{\Gamma_{\perp} \vdash \{P_1 * P_2\}c_1 || c_2\{Q_1 * Q_2\}} \text{ (PAR)} \\
\\
\frac{P \Rightarrow P' \quad \Gamma_{\perp} \vdash \{P'\}c\{Q'\} \quad Q' \Rightarrow Q}{\Gamma_{\perp} \vdash \{P\}c\{Q\}} \text{ (CONS)} \quad \frac{\text{fv}(R) \cap \text{mod}(c) = \emptyset \quad \Gamma_{\perp} \vdash \{P\}c\{Q\} \quad P \text{ is precise or } R \text{ is precise}}{\Gamma_{\perp} \vdash \{P * R\}c\{Q * R\}} \text{ (FRAME)} \\
\\
\frac{x \notin \text{fv}(c) \quad \Gamma_{\perp} = \Gamma \Rightarrow x \notin \text{fv}(\Gamma) \quad \Gamma_{\perp} \vdash \{P\}c\{Q\}}{\Gamma_{\perp} \vdash \{\exists x. P\}c\{\exists x. Q\}} \text{ (EXISTS)} \\
\\
\frac{\Gamma = \langle \alpha, f_{as}, f_{au}, I(x) \rangle \quad \Gamma \text{ is valid} \quad I(x) \text{ is unary and precise} \quad \Gamma \vdash \{P * \text{sguard}(1, \emptyset^\#) * \text{UniqueEmpty}\}c\{Q * \text{sguard}(1, x_s) * \text{PRE}_s(x_s) * \text{UniquePre}\}}{\perp \vdash \{I(x) * \text{Low}(\alpha(x)) * P\}c\{\exists x'. I(x') * \text{Low}(\alpha(x')) * Q\}} \text{ (SHARE)} \\
\\
\frac{\Gamma = \langle \alpha, f_{as}, f_{au}, I(x) \rangle \quad I(x) \text{ is unary and precise} \quad x_v \notin \text{fv}(P) \cup \text{fv}(Q) \cup \text{fv}(I(x)) \quad x_s, x_a, x_v \notin \text{mod}(c) \quad \text{noguard}(P) \quad \text{noguard}(Q) \quad \perp \vdash \{P * I(x_v)\}c\{Q * I(f_{as}(x_v, x_a))\}}{\Gamma \vdash \{P * \text{sguard}(r, x_s)\} \text{atomic } c\{Q * \text{sguard}(r, x_s \cup^\# \{x_a\})\}} \text{ (ATOMICSHR)} \\
\\
\frac{\Gamma = \langle \alpha, f_{as}, f_{au}, I(x) \rangle \quad I(x_v) \text{ is unary and precise} \quad x_v \notin \text{fv}(P) \cup \text{fv}(Q) \cup \text{fv}(I(x)) \quad x_s, x_a, x_v \notin \text{mod}(c) \quad \text{noguard}(P) \quad \text{noguard}(Q) \quad \perp \vdash \{P * I(x_v)\}c\{Q * I(f_{ai}(x_v, x_a))\}}{\Gamma \vdash \{P * \text{uguard}_i(x_s)\} \text{atomic } c\{Q * \text{uguard}_i(x_s ++ [x_a])\}} \text{ (ATOMICUNQ)}
\end{array}$$

Fig. 8. Proof rules of COMMCSL. The SHARE rule and the two ATOMIC rules implement all checks specific to our technique. The rules for conditionals and loops have been adjusted for information flow reasoning. All other rules are as in standard CSL. Recall that the family of unique actions f_{au} is indexed by a (finite) type I , and let us denote the elements of I by i_0, \dots, i_n . We write *UniqueEmpty* as an abbreviation for $\text{uguard}_{i_0}(\emptyset) * \dots * \text{uguard}_{i_n}(\emptyset)$, and *UniquePre* as an abbreviation for $\exists x_{i_0}, \dots, x_{i_n}. \text{uguard}_{i_0}(x_{i_0}) * \text{PRE}_u(x_{i_0}) * \dots * \text{uguard}_{i_0}(x_{i_n}) * \text{PRE}_u(x_{i_n})$. We use $\text{fv}(P)$ to denote the free variables in the assertion P , and $\text{mod}(c)$ for the set of variables modified by command c .

that the state of the resource heap continues to be a possible result of applying the actions with the arguments specified in the guards to the initial value v_0 . We provide the formal definition of safety in App. C. Crucially, this predicate makes a statement only about an individual execution.

We can now define the validity of a Hoare triple $\Gamma_\perp \models \{P\}c\{Q\}$ by stating that any two executions from a pair of states satisfying P are *safe* for an arbitrary number of steps, and *any* pair of final states from those two executions will fulfill Q . More precisely, the Hoare triple holds if there is some function Ψ that maps any initial state σ to a set of all possible final states Σ , and for any σ_1, σ_2 fulfilling P , all pairs of states in $\Psi(\sigma_1) \times \Psi(\sigma_2)$ fulfill Q :

Definition 4.1. $\Gamma_\perp \models \{P\}c\{Q\}$ holds iff there is some Ψ s.t. for all n, σ , if $\sigma, \sigma \models P$, we have $\text{safe}(n, \Gamma_\perp, c, \sigma, \Psi(\sigma))$, and for all $\sigma_1, \sigma_2, \sigma'_1, \sigma'_2$, if $\sigma_1, \sigma_2 \models P$ and $\sigma'_1 \in \Psi(\sigma_1)$ and $\sigma'_2 \in \Psi(\sigma_2)$, then $\sigma'_1, \sigma'_2 \models Q$.

THEOREM 4.2. (Soundness) *If $\Gamma_\perp \vdash \{P\}c\{Q\}$ then $\Gamma_\perp \models \{P\}c\{Q\}$.*

As a corollary, we prove that this theorem actually implies that Hoare triples make true statements about pairs of program executions from an initial state without a shared resource, using \widehat{gh} to range over extended heaps whose guard states are \perp and whose permission heaps have full permission of every location in their domain:

COROLLARY 4.3. *If $\perp \vdash \{P\}c\{Q\}$ and $\langle s_1, \widehat{gh}_1 \rangle, \langle s_2, \widehat{gh}_2 \rangle \models P$, then c does not abort from any $\langle s_i, \text{norm}(\widehat{gh}_i) \rangle$, and if $c, \langle s_i, \text{norm}(\widehat{gh}_i) \rangle \rightarrow^* \text{skip}, \langle s'_i, h_i \rangle$ for some s'_i, h_i and $h_i = \text{norm}(\widehat{gh}_i)$ for some \widehat{gh}_i , then $\langle s'_1, \widehat{gh}_1 \rangle, \langle s'_2, \widehat{gh}_2 \rangle \models Q$.*

5 IMPLEMENTATION AND EVALUATION

We have implemented our technique in **HYPERVIPER**, an automated prototype verifier based on the Viper verification infrastructure [Müller et al. 2016]. **HYPERVIPER** supports a richer language than the one we use in this paper; in particular, instead of parallel composition commands, it allows dynamic thread creation using `fork` and `join` commands. Additionally, **HYPERVIPER** supports *multiple* resources in a single program (which can each be associated with different resource specifications) and allows threads to explicitly acquire and release different resources.

HYPERVIPER's level of automation is similar to existing automated verifiers: Users must provide method pre- and postcondition as well as loop invariants, in an assertion language that is similar to the one shown in this paper, and additionally have to declare resource specifications. Given these specifications and some annotations indicating key proof steps (e.g., which action is performed by an atomic block, and how to split and merge shared action guards; see App. E for an example), our tool automatically either verifies the program or indicates potential errors.

Internally, **HYPERVIPER** encodes the validity constraints for all resource specifications as well as all other proof obligations imposed by our logic into the Viper intermediate language. To encode relational proof obligations, it uses a modular product program construction [Eilers et al. 2018], and combines it with existing encodings for concurrent programs [Leino and Müller 2009] in a sound way [Eilers et al. 2021]. Subsequently, it automatically verifies the generated program using one of Viper's backend verifiers and, ultimately, the Z3 SMT solver [de Moura and Bjørner 2008].

To demonstrate the practical usefulness of our verification technique, we have applied **HYPERVIPER** to a number of example programs that represent a variety of applications that concurrently manipulate shared data structures with secret data. The first 15 examples model specific applications, which have in common that they spawn a number of worker threads to modify some shared data structure and subsequently output the abstract view of the shared data structure. All examples explicitly model data structures containing different kinds of secret information (e.g., employee

Example	Data structure	Abstraction	LOC	Ann.	T
Count-Vaccinated	Counter, increment	None	44	46	10.15
Figure 2	Integer, add	None	129	95	10.90
Count-Sick-Days	Integer, add	None	52	45	13.67
Figure 1	Integer, arbitrary	Constant	29	20	1.52
Mean-Salary	List, append	Mean	80	84	14.10
Email-Metadata	List, append	Multiset	82	75	16.70
Patient-Statistic	List, append	Length	73	70	4.92
Debt-Sum	List, append	Sum	76	81	14.45
Sick-Employee-Names	Treeset, add	None	105	113	28.43
Website-Visitor-IPs	Listset, add	None	74	69	6.20
Figure 3	HashMap, put	Key set	129	96	10.37
Sales-By-Region	HashMap, disjoint put	None	129	104	12.37
Salary-Histogram	HashMap, increment value	None	135	109	13.78
Count-Purchases	HashMap, add value	None	137	109	11.73
Most-Valuable-Purchase	HashMap, conditional put	None	140	118	17.87
1-Producer-1-Consumer	Queue	Consumed sequence	82	88	3.23
Pipeline	Two queues	Consumed sequences	122	100	3.66
2-Producers-2-Consumers	Queue	Produced multiset	130	134	8.45

Table 1. Evaluated examples. We show the used data structure and the operation(s) we allow on it; for the last three examples, these are producing and consuming data. LOC are lines of code, not including specifications, Ann. are only specifications and proof annotations. T is the verification time in seconds, averaged over 5 runs.

names and salaries, patient health data, or user activity on a website). Depending on the example, the data contains some low parts (e.g., whether or not a patient has been vaccinated), or some *aspect* of the secret data is low (e.g., the range of an employee’s salary or the number of—otherwise secret—purchases they have made).

The last three examples model variations of general parallel programming patterns; here, we leave the data and computation abstract and show instead that the pattern is safe for *all* possible implementations where the individual roles do not directly leak secret data themselves. Table 1 shows, for each example, the total lines of code (not including whitespace), the number of lines used for annotations (i.e., specifications and proof annotations), and the verification time in seconds (measured on an 8-core AMD Ryzen 6850U with 32GB of RAM running Ubuntu 22.04 on a warmed-up JVM, averaged over five runs), as well as the used data structure and abstraction. The examples range from 29 to 140 lines of code, with a similar number of lines of specifications and proof annotations. Typical verification times are between 3 and 30 seconds. Thus, we conclude that commutativity-based reasoning can be automated efficiently. In the remainder of this section, we will discuss how different aspects of our technique allow us to verify the diverse set of examples.

Precise action definitions. For several examples, we can prove that the actions of all threads commute by carefully defining the legal actions. For example, while puts on maps with the same keys do not commute in general, in the example Salary-Histogram, each put *increments* the value for the key it is modifying (i.e., it increments the number of employees that fall within the specific salary range represented by the key). As a result, the updates commute. Similarly, the example Most-Valuable-Purchase iterates through purchase records and finds the highest price each user has ever paid: Here, each thread updates the map from user IDs to prices *only if* its purchase value is greater than the one currently in the map, which again results in commuting map updates.

Abstraction. In half of our examples, we use an abstraction to relax the commutativity requirement. For most examples, this abstraction states which part of the computed internal (and partially secret) data we want to make public. For example, we compute a list of employee names and salaries but leak only the mean salary, and we collect a list of individual debts of a person but leak only the overall sum of their debt, not individual amounts or creditors. In one example (Email-Metadata), we abstract a list to a multiset to prove that we may leak the list after sorting it, which eliminates the secret-dependent item order. For the producer-consumer and pipeline examples, our abstraction does not return a view on the actual shared data, but on *ghost data* added for verification purposes: It does not state that the current data in the queue is low, which is not of interest, since the final state of the queue will always be empty, but that the sequence of *consumed* items must be low (or, if there are multiple producers and consumers, the multiset view of that sequence). We also show that we can support the original version of Figure 1 where each thread sets the shared variable to a different value, but only if said value is not leaked.

Resource specifications. The examples Sick-Employee-Names and Website-Visitor-IPs both add low values to sets (names of sick employees and IP addresses of website visitors, respectively), but use different set implementations. Since resource specifications abstract over concrete implementations of data structures, we can reuse the same resource spec for both examples.

Unique actions. In three examples, we use unique actions to exploit that some actions are performed only by a single thread, and their order is therefore not affected by thread interleavings. In the example Sales-By-Region, different threads process data from different regions, and add it to a shared map. Since the keys are region-specific, writes from different threads never conflict, and thus all updates commute. For the single producer-consumer and the pipeline examples, we exploit that there is only a single producer and a single consumer thread per shared queue to prove that the sequence of consumed values is low; if there were multiple such threads (as is the case in the multiple producer-consumer example), the order of consumed items could be affected by secret data (and thus, only its multiset view is low).

Retroactive checking of action arguments. For the multiple producer-consumer example, the number of items consumed by each individual consumer depends on scheduling and therefore potentially high data. However, the *total* number of consumed items is low. Thus, checking that the number of performed consume actions is low *after* joining all threads, when the overall number of consume actions is known, allows us to prove this example secure. In the pipeline example, a thread produces low data into a first queue, a middle thread consumes this data, transforms it, and produces the result into a second queue, from where a third thread consumes it. While all threads are executing, the middle thread does not know that the data it reads from the first queue is low; we learn that only once the first queue is unshared. Thus, we would not be able to prove that the middle thread fulfills the second queue’s precondition (that produced data is low) while all threads are running, but we *can* prove this precondition retroactively after unsharing the first queue.

High branches. Circa half of our examples have secret-dependent execution time due to branches on high data, and would thus be rejected by all existing techniques, even if the attacker cannot observe timing.

6 RELATED WORK

Researchers have developed a plethora of type systems [Smith 2007], static analyses [Giffhorn and Snelting 2015], program transformations [Eilers et al. 2021], and program logics [Ernst and Murray 2019; Murray et al. 2018] to verify information flow security of concurrent programs, as well as multiple definitions of information flow security in this setting. Bisimulation-based

properties [Focardi and Gorrieri 1995] and observational determinism [Zdancewic and Myers 2003] are properties of (sets of) traces, which assume that attackers can observe either low program variables or low events *during* the execution of the program, unlike our setting, where we assume that the attacker can observe *only* the public output of the program.

In our setting, the standard property is non-interference (which is what our logic guarantees), or weaker versions of it like *possibilistic* [Smith and Volpano 1998] and *probabilistic* [Sabelfeld and Sands 2000] non-interference, which guarantee, respectively, that secret data does not influence either the set of possible final outputs or their probabilities. Since possibilistic non-interference is too weak in practice, most existing techniques target either traditional or probabilistic non-interference. As discussed previously, they achieve this goal by entirely preventing secret-dependent execution time differences, under the idealized assumption that only high branches can lead to such differences [Eilers et al. 2021; Ernst and Murray 2019; Murray et al. 2018; Schoepe et al. 2020; Smith 2006, 2007]. These techniques can be applied to more realistic settings by also preventing other sources of timing differences, but preventing them entirely on standard hardware is very complex and requires strong assumptions about the used compiler and hardware. Our technique has the unique advantage of not requiring any reasoning about execution time in order to prove information flow security for output values, which makes it applicable independently of the used hardware. Our technique can also be used to complement existing approaches in order to verify programs where secret data *does* lead to timing differences (and therefore cannot be verified with existing techniques), but these timing differences do not influence the final output.

Most of the aforementioned techniques either do not support a mutable heap (e.g. [Murray et al. 2018; Smith 2007] or work for programs using locks to protect shared memory like our logic (e.g. [Eilers et al. 2021; Ernst and Murray 2019]), but some recent logics target more complex settings, like fine-grained concurrency [Frumin et al. 2021] or relaxed memory models [Yan and Murray 2021]. We believe that our approach also extends to fine-grained concurrency, since the general idea of making sure that concurrent changes commute also applies in this setting.

Instead of statically proving that a program does not leak high data, some techniques aim to prevent unwanted information leakage for arbitrary programs at runtime, including leakage as a result of internal timing channels by modifying the language runtime [Vassena et al. 2019], using information flow aware concurrency primitives [Stefan et al. 2012], or by transforming the original program [Russo et al. 2006]. Such techniques can allow running programs that would be rejected by static techniques, but at the cost of requiring specific runtimes or affecting execution performance.

Commutativity-based reasoning has been used for other applications, particularly in the context of parallelization. Bansal et al. [2018] infer commutativity conditions on data types, and Brutschy et al. [2017] use commutativity to check if concurrent database accesses of an application are *serializable* [Bernstein et al. 1987]. In both cases, it is useful to distinguish between the concrete representation of data in memory (and the respective individual reads and writes) and the higher-level data structure it represents (and its higher-level operations, like appending to a list), which is a kind of abstraction. In contrast, we use abstraction to exclude *secret* parts of the high-level state of a data structure from our commutativity requirement. Additionally, we show only that the final state of shared data is identical (modulo abstraction) after all modifications, whereas in other contexts, it is also important that queries in intermediate states return the same results.

7 CONCLUSION AND FUTURE WORK

We have presented COMMCSL, a concurrent separation logic for verifying non-interference using abstract commutativity. We proved COMMCSL sound in Isabelle/HOL, and our evaluation shows that COMMCSL can be automated efficiently and is able to verify common examples. For future work, we suggest applying reasoning based on (abstract) commutativity to other settings like

fine-grained concurrency, and to explore when it is sound to assume that shared data is low *while* threads are still performing concurrent updates.

REFERENCES

- Marc Andryscio, Andres Nötzli, Fraser Brown, Ranjit Jhala, and Deian Stefan. 2018. Towards Verified, Constant-time Floating Point Operations. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, Toronto, ON, Canada, October 15-19, 2018*, David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang (Eds.). ACM, 1369–1382. <https://doi.org/10.1145/3243734.3243766>
- Kshitij Bansal, Eric Koskinen, and Omer Tripp. 2018. Automatic Generation of Precise and Useful Commutativity Conditions. In *Tools and Algorithms for the Construction and Analysis of Systems - 24th International Conference, TACAS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 10805)*, Dirk Beyer and Marieke Huisman (Eds.). Springer, 115–132. https://doi.org/10.1007/978-3-319-89960-2_7
- Gilles Barthe, Benjamin Grégoire, and Vincent Laporte. 2018. Secure Compilation of Side-Channel Countermeasures: The Case of Cryptographic "Constant-Time". In *31st IEEE Computer Security Foundations Symposium, CSF 2018, Oxford, United Kingdom, July 9-12, 2018*. IEEE Computer Society, 328–343. <https://doi.org/10.1109/CSF.2018.00031>
- Philip A. Bernstein, Vassos Hadzilacos, and Nathan Goodman. 1987. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley. <http://research.microsoft.com/en-us/people/philbe/ccontrol.aspx>
- Richard Bornat, Cristiano Calcagno, Peter W. O'Hearn, and Matthew J. Parkinson. 2005. Permission accounting in separation logic. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, Long Beach, California, USA, January 12-14, 2005*, Jens Palsberg and Martin Abadi (Eds.). ACM, 259–270. <https://doi.org/10.1145/1040305.1040327>
- John Boyland. 2003. Checking Interference with Fractional Permissions. In *Static Analysis, 10th International Symposium, SAS 2003, San Diego, CA, USA, June 11-13, 2003, Proceedings (Lecture Notes in Computer Science, Vol. 2694)*, Radhia Cousot (Ed.). Springer, 55–72. https://doi.org/10.1007/3-540-44898-5_4
- Stephen D. Brookes. 2004. A Semantics for Concurrent Separation Logic. In *CONCUR 2004 - Concurrency Theory, 15th International Conference, London, UK, August 31 - September 3, 2004, Proceedings (Lecture Notes in Computer Science, Vol. 3170)*, Philippa Gardner and Nobuko Yoshida (Eds.). Springer, 16–34. https://doi.org/10.1007/978-3-540-28644-8_2
- Lucas Brutschy, Dimitar K. Dimitrov, Peter Müller, and Martin T. Vechev. 2017. Serializability for eventual consistency: criterion, analysis, and applications. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM, 458–472. <https://doi.org/10.1145/3009837.3009895>
- Michael R. Clarkson and Fred B. Schneider. 2010. Hyperproperties. *Journal of Computer Security* 18, 6 (2010), 1157–1210.
- Jeroen V. Cleemput, Bart Coppens, and Bjorn De Sutter. 2012. Compiler mitigations for time attacks on modern x86 processors. *TACO* 8, 4 (2012), 23:1–23:20. <https://doi.org/10.1145/2086696.2086702>
- Leonardo Mendonça de Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings (Lecture Notes in Computer Science, Vol. 4963)*, C. R. Ramakrishnan and Jakob Rehof (Eds.). Springer, 337–340. https://doi.org/10.1007/978-3-540-78800-3_24
- Marco Eilers, Severin Meier, and Peter Müller. 2021. Product Programs in the Wild: Retrofitting Program Verifiers to Check Information Flow Security. In *Computer Aided Verification - 33rd International Conference, CAV 2021, Virtual Event, July 20-23, 2021, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 12759)*, Alexandra Silva and K. Rustan M. Leino (Eds.). Springer, 718–741. https://doi.org/10.1007/978-3-030-81685-8_34
- Marco Eilers, Peter Müller, and Samuel Hitz. 2018. Modular Product Programs. In *Programming Languages and Systems - 27th European Symposium on Programming, ESOP 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings (Lecture Notes in Computer Science, Vol. 10801)*, Amal Ahmed (Ed.). Springer, 502–529. https://doi.org/10.1007/978-3-319-89884-1_18
- Gidon Ernst and Toby Murray. 2019. SecCSL: Security Concurrent Separation Logic. In *Computer Aided Verification - 31st International Conference, CAV 2019, New York City, NY, USA, July 15-18, 2019, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 11562)*, Isil Dillig and Serdar Tasiran (Eds.). Springer, 208–230. https://doi.org/10.1007/978-3-030-25543-5_13
- Riccardo Focardi and Roberto Gorrieri. 1995. A Taxonomy of Security Properties for Process Algebras. *J. Comput. Secur.* 3, 1 (1995), 5–34. <https://doi.org/10.3233/JCS-1994/1995-3103>
- Dan Frumin, Robbert Krebbers, and Lars Birkedal. 2021. Compositional Non-Interference for Fine-Grained Concurrent Programs. In *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*. IEEE, 1416–1433. <https://doi.org/10.1109/SP40001.2021.00003>

- Dennis Giffhorn and Gregor Snelting. 2015. A new algorithm for low-deterministic security. *Int. J. Inf. Sec.* 14, 3 (2015), 263–287. <https://doi.org/10.1007/s10207-014-0257-6>
- Jan Jancar, Marcel Fourné, Daniel De Almeida Braga, Mohamed Sabt, Peter Schwabe, Gilles Barthe, Pierre-Alain Fouque, and Yasemin Acar. 2022. “They’re not that hard to mitigate”: What Cryptographic Library Developers Think About Timing Attacks. In *2022 IEEE Symposium on Security and Privacy, S&P 2022*.
- Paul Kocher, Jann Horn, Anders Fogh, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom. 2019. Spectre Attacks: Exploiting Speculative Execution. In *2019 IEEE Symposium on Security and Privacy, SP 2019, San Francisco, CA, USA, May 19-23, 2019*. IEEE, 1–19. <https://doi.org/10.1109/SP.2019.00002>
- Eric Koskinen and Kshitij Bansal. 2021. Decomposing Data Structure Commutativity Proofs with $\$m!\$n\$$ -Differencing. In *Verification, Model Checking, and Abstract Interpretation - 22nd International Conference, VMCAI 2021, Copenhagen, Denmark, January 17-19, 2021, Proceedings (Lecture Notes in Computer Science, Vol. 12597)*, Fritz Henglein, Sharon Shoham, and Yakir Vizel (Eds.). Springer, 81–103. https://doi.org/10.1007/978-3-030-67067-2_5
- K. Rustan M. Leino and Peter Müller. 2009. A Basis for Verifying Multi-threaded Programs. In *Programming Languages and Systems, 18th European Symposium on Programming, ESOP 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings (Lecture Notes in Computer Science, Vol. 5502)*, Giuseppe Castagna (Ed.). Springer, 378–393. https://doi.org/10.1007/978-3-642-00590-9_27
- Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Anders Fogh, Jann Horn, Stefan Mangard, Paul Kocher, Daniel Genkin, Yuval Yarom, and Mike Hamburg. 2018. Meltdown: Reading Kernel Memory from User Space. In *27th USENIX Security Symposium, USENIX Security 2018, Baltimore, MD, USA, August 15-17, 2018*, William Enck and Adrienne Porter Felt (Eds.). USENIX Association, 973–990. <https://www.usenix.org/conference/usenixsecurity18/presentation/lipp>
- Peter Müller, Malte Schwerhoff, and Alexander J. Summers. 2016. Viper: A Verification Infrastructure for Permission-Based Reasoning. In *Verification, Model Checking, and Abstract Interpretation - 17th International Conference, VMCAI 2016, St. Petersburg, FL, USA, January 17-19, 2016. Proceedings (Lecture Notes in Computer Science, Vol. 9583)*, Barbara Jobstmann and K. Rustan M. Leino (Eds.). Springer, 41–62. https://doi.org/10.1007/978-3-662-49122-5_2
- Toby C. Murray, Robert Sison, and Kai Engelhardt. 2018. COVERN: A Logic for Compositional Verification of Information Flow Control. In *2018 IEEE European Symposium on Security and Privacy, EuroS&P 2018, London, United Kingdom, April 24-26, 2018*. IEEE, 16–30. <https://doi.org/10.1109/EuroSP.2018.00010>
- Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. 2002. Isabelle/HOL - A Proof Assistant for Higher-Order Logic. *Lecture Notes in Computer Science*, Vol. 2283. Springer. <https://doi.org/10.1007/3-540-45949-9>
- Peter W. O’Hearn. 2004. Resources, Concurrency and Local Reasoning. In *CONCUR 2004 - Concurrency Theory, 15th International Conference, London, UK, August 31 - September 3, 2004. Proceedings (Lecture Notes in Computer Science, Vol. 3170)*, Philippa Gardner and Nobuko Yoshida (Eds.). Springer, 49–67. https://doi.org/10.1007/978-3-540-28644-8_4
- Peter W. O’Hearn, Hongseok Yang, and John C. Reynolds. 2004. Separation and Information Hiding. In *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (Venice, Italy) (POPL ’04)*. Association for Computing Machinery, New York, NY, USA, 268–280. <https://doi.org/10.1145/964001.964024>
- Matthew J. Parkinson and Gavin M. Bierman. 2005. Separation logic and abstraction. In *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, Long Beach, California, USA, January 12-14, 2005*, Jens Palsberg and Martin Abadi (Eds.). ACM, 247–258. <https://doi.org/10.1145/1040305.1040326>
- Alejandro Russo, John Hughes, David A. Naumann, and Andrei Sabelfeld. 2006. Closing Internal Timing Channels by Transformation. In *Advances in Computer Science - ASIAN 2006. Secure Software and Related Issues, 11th Asian Computing Science Conference, Tokyo, Japan, December 6-8, 2006, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 4435)*, Mitsu Okada and Ichiro Satoh (Eds.). Springer, 120–135. https://doi.org/10.1007/978-3-540-77505-8_10
- Andrei Sabelfeld and David Sands. 2000. Probabilistic Noninterference for Multi-Threaded Programs. In *Proceedings of the 13th IEEE Computer Security Foundations Workshop, CSFW ’00, Cambridge, England, UK, July 3-5, 2000*. IEEE Computer Society, 200–214. <https://doi.org/10.1109/CSFW.2000.856937>
- Daniel Schoepe, Toby Murray, and Andrei Sabelfeld. 2020. VERONICA: Expressive and Precise Concurrent Information Flow Security (Extended Version with Technical Appendices). CoRR abs/2001.11142 (2020). arXiv:2001.11142 <https://arxiv.org/abs/2001.11142>
- Geoffrey Smith. 2006. Improved typings for probabilistic noninterference in a multi-threaded language. *J. Comput. Secur.* 14, 6 (2006), 591–623. <http://content.iospress.com/articles/journal-of-computer-security/jcs273>
- Geoffrey Smith. 2007. Principles of Secure Information Flow Analysis. In *Malware Detection*, Mihai Christodorescu, Somesh Jha, Douglas Maughan, Dawn Song, and Cliff Wang (Eds.). Advances in Information Security, Vol. 27. Springer, 291–307. https://doi.org/10.1007/978-0-387-44599-1_13
- Geoffrey Smith and Dennis M. Volpano. 1998. Secure Information Flow in a Multi-Threaded Imperative Language. In *POPL ’98, Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Diego, CA, USA,*

- January 19-21, 1998, David B. MacQueen and Luca Cardelli (Eds.). ACM, 355–364. <https://doi.org/10.1145/268946.268975>
- Marcelo Sousa and Isil Dillig. 2016. Cartesian hoare logic for verifying k-safety properties. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2016, Santa Barbara, CA, USA, June 13-17, 2016*, Chandra Krintz and Emery D. Berger (Eds.). ACM, 57–69. <https://doi.org/10.1145/2908080.2908092>
- Deian Stefan, Pablo Buiras, Edward Z. Yang, Amit Levy, David Terei, Alejandro Russo, and David Mazières. 2013. Eliminating Cache-Based Timing Attacks with Instruction-Based Scheduling. In *Computer Security - ESORICS 2013 - 18th European Symposium on Research in Computer Security, Egham, UK, September 9-13, 2013. Proceedings (Lecture Notes in Computer Science, Vol. 8134)*, Jason Crampton, Sushil Jajodia, and Keith Mayes (Eds.). Springer, 718–735. https://doi.org/10.1007/978-3-642-40203-6_40
- Deian Stefan, Alejandro Russo, Pablo Buiras, Amit Levy, John C. Mitchell, and David Mazières. 2012. Addressing covert termination and timing channels in concurrent information flow systems. In *ACM SIGPLAN International Conference on Functional Programming, ICFP'12, Copenhagen, Denmark, September 9-15, 2012*, Peter Thiemann and Robby Bruce Findler (Eds.). ACM, 201–214. <https://doi.org/10.1145/2364527.2364557>
- Viktor Vafeiadis. 2011. Concurrent Separation Logic and Operational Semantics. In *Twenty-seventh Conference on the Mathematical Foundations of Programming Semantics, MFPS 2011, Pittsburgh, PA, USA, May 25-28, 2011 (Electronic Notes in Theoretical Computer Science, Vol. 276)*, Michael W. Mislove and Joël Ouaknine (Eds.). Elsevier, 335–351. <https://doi.org/10.1016/j.entcs.2011.09.029>
- Marco Vassena, Gary Soeller, Peter Amidon, Matthew Chan, John Renner, and Deian Stefan. 2019. Foundations for Parallel Information Flow Control Runtime Systems. In *Principles of Security and Trust - 8th International Conference, POST 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings (Lecture Notes in Computer Science, Vol. 11426)*, Flemming Nielson and David Sands (Eds.). Springer, 1–28. https://doi.org/10.1007/978-3-030-17138-4_1
- Martin T. Vechev, Eran Yahav, Raghavan Raman, and Vivek Sarkar. 2010. Automatic Verification of Determinism for Structured Parallel Programs. In *Static Analysis - 17th International Symposium, SAS 2010, Perpignan, France, September 14-16, 2010. Proceedings (Lecture Notes in Computer Science, Vol. 6337)*, Radhia Cousot and Matthieu Martel (Eds.). Springer, 455–471. https://doi.org/10.1007/978-3-642-15769-1_28
- Dennis M. Volpano, Cynthia E. Irvine, and Geoffrey Smith. 1996. A Sound Type System for Secure Flow Analysis. *J. Comput. Secur.* 4, 2/3 (1996), 167–188. <https://doi.org/10.3233/JCS-1996-42-304>
- Dennis M. Volpano and Geoffrey Smith. 1998. Probabilistic Noninterference in a Concurrent Language. In *Proceedings of the 11th IEEE Computer Security Foundations Workshop, Rockport, Massachusetts, USA, June 9-11, 1998*. IEEE Computer Society, 34–43. <https://doi.org/10.1109/CSFW.1998.683153>
- Pengbo Yan and Toby Murray. 2021. SecRSL: security separation logic for C11 release-acquire concurrency. *Proc. ACM Program. Lang.* 5, OOPSLA (2021), 1–26. <https://doi.org/10.1145/3485476>
- Steve Zdancewic and Andrew C. Myers. 2003. Observational Determinism for Concurrent Program Security. In *16th IEEE Computer Security Foundations Workshop (CSFW-16 2003), 30 June - 2 July 2003, Pacific Grove, CA, USA*. IEEE Computer Society, 29. <https://doi.org/10.1109/CSFW.2003.1212703>

A APPENDIX

A.1 Operational semantics

Fig. 9 shows the operational small-step semantics of our language, which are standard and taken from Vafeiadis [2011].

B FULL LOGIC

In this section, we provide additional definitions and explanations not shown in the main part of the paper.

B.1 Additional Definitions

The sum of two guard heaps for a unique action a_i is defined as follows:

$$gu_i \oplus gu'_i = \begin{cases} gu_i & \text{if } gu'_i = \perp \\ gu'_i & \text{if } gu_i = \perp \\ \text{undefined} & \text{otherwise} \end{cases} \quad (3)$$

$$\begin{array}{c}
\frac{h(\llbracket e \rrbracket(s)) = v}{\langle x := [e], (s, h) \rangle \rightarrow \langle \text{skip}, (s[x \mapsto v], h) \rangle} \text{(READ)} \quad \frac{\llbracket e \rrbracket(s) \notin \text{dom}(h)}{\langle x := [e], (s, h) \rangle \rightarrow \text{abort}} \text{(READA)} \\
\\
\frac{\llbracket e_1 \rrbracket(s) \in \text{dom}(h)}{\langle [e_1] := e_2, (s, h) \rangle \rightarrow \langle \text{skip}, (s, h[\llbracket e_1 \rrbracket(s) \mapsto h[\llbracket e_2 \rrbracket(s)]] \rangle} \text{(WRITE)} \quad \frac{\llbracket e_1 \rrbracket(s) \notin \text{dom}(h)}{\langle [e_1] := e_2, (s, h) \rangle \rightarrow \text{abort}} \text{(WRITEA)} \\
\\
\frac{l \notin \text{dom}(h)}{\langle x := \text{alloc}(e), (s, h) \rangle \rightarrow \langle \text{skip}, (s[x \mapsto l], h[l \mapsto \llbracket e \rrbracket(s)]) \rangle} \text{(ALLOC)} \\
\\
\frac{}{\langle \text{skip}; c_2, (s, h) \rangle \rightarrow \langle c_2, (s, h) \rangle} \text{(SEQ1)} \quad \frac{\langle c_1, (s, h) \rangle \rightarrow \langle c'_1, (s', h') \rangle}{\langle c_1; c_2, (s, h) \rangle \rightarrow \langle c'_1; c_2, (s', h') \rangle} \text{(SEQ2)} \\
\\
\frac{}{\langle \text{abort}; c_2, (s, h) \rangle \rightarrow \text{abort}} \text{(SEQA)} \quad \frac{}{\langle x := e, (s, h) \rangle \rightarrow \langle \text{skip}, (s[x \mapsto \llbracket e \rrbracket(s)], h) \rangle} \text{(ASSIGN)} \\
\\
\frac{\llbracket b \rrbracket(s)}{\langle \text{if } (b) \text{ then } \{c_1\} \text{ else } \{c_2\}, (s, h) \rangle \rightarrow \langle c_1, (s, h) \rangle} \text{(IF1)} \quad \frac{\neg \llbracket b \rrbracket(s)}{\langle \text{if } (b) \text{ then } \{c_1\} \text{ else } \{c_2\}, (s, h) \rangle \rightarrow \langle c_2, (s, h) \rangle} \text{(IF2)} \\
\\
\frac{}{\langle \text{while } (e) \text{ do } \{c\}, (s, h) \rangle \rightarrow \langle \text{if } (b) \text{ then } \{c; \text{while } (e) \text{ do } \{c\} \} \text{ else } \{\text{skip}\}, (s, h) \rangle} \text{(LOOP)} \\
\\
\frac{\langle c_1, (s, h) \rangle \rightarrow \langle c'_1, (s', h') \rangle}{\langle c_1 || c_2, (s, h) \rangle \rightarrow \langle c'_1 || c_2, (s', h') \rangle} \text{(PAR1)} \quad \frac{\langle c_2, (s, h) \rangle \rightarrow \langle c'_2, (s', h') \rangle}{\langle c_1 || c_2, (s, h) \rangle \rightarrow \langle c_1 || c'_2, (s', h') \rangle} \text{(PAR2)} \\
\\
\frac{}{\langle \text{skip} || \text{skip}, (s, h) \rangle \rightarrow \langle \text{skip}, (s, h) \rangle} \text{(PAR3)} \\
\\
\frac{\langle c_1, (s, h) \rangle \rightarrow \text{abort}}{\langle c_1 || c_2, (s, h) \rangle \rightarrow \text{abort}} \text{(PARA1)} \quad \frac{\langle c_2, (s, h) \rangle \rightarrow \text{abort}}{\langle c_1 || c_2, (s, h) \rangle \rightarrow \text{abort}} \text{(PARA2)} \\
\\
\frac{\langle c, (s, h) \rangle \rightarrow^* \langle \text{skip}, (s', h') \rangle}{\langle \text{atomic } c, (s, h) \rangle \rightarrow \langle \text{skip}, (s', h') \rangle} \text{(ATOM)} \quad \frac{\langle c, (s, h) \rangle \rightarrow^* \text{abort}}{\langle \text{atomic } c, (s, h) \rangle \rightarrow \text{abort}} \text{(ATOMA)}
\end{array}$$

Fig. 9. Operational semantics.

That is, at least one guard heap must have value \perp , and the sum simply contains the non- \perp value (if any exists).

Addition of shared guard heaps computes the multiset union of the respective argument multisets of the two summands:

$$gs \oplus gs' = \begin{cases} \langle r + r', \text{args} \cup^\# \text{args}' \rangle & \text{if } gs = \langle r, \text{args} \rangle \wedge gs' = \langle r', \text{args}' \rangle \wedge r + r' \leq 1 \\ gs & \text{if } gs' = \perp \\ gs' & \text{if } gs = \perp \\ \text{undefined} & \text{otherwise} \end{cases} \quad (4)$$

Finally, as stated before, permission heap addition is standard:

$$(ph \oplus ph')(l) = \begin{cases} ph(l) \oplus ph'(l) & \text{if } l \in \text{dom}(ph) \cap \text{dom}(ph') \\ ph(l) & \text{if } l \in \text{dom}(ph) \wedge l \notin \text{dom}(ph') \\ ph'(l) & \text{if } l \in \text{dom}(ph') \wedge l \notin \text{dom}(ph) \end{cases} \quad (5)$$

where

$$\langle r, v \rangle \oplus \langle r', v' \rangle = \begin{cases} \langle r + r', v \rangle & \text{if } r + r' \leq 1 \wedge v = v' \\ \text{undefined} & \text{otherwise} \end{cases} \quad (6)$$

B.2 Precision Side Conditions

Rules FRAME, PAR, SHARE, ATOMICSHR, and ATOMICUNQ contain precision requirements that do not exist in their standard CSL equivalents. While we believe these restrictions can be lifted to something weaker, *some* restrictions on the assertions are necessary in our setting; existing logics do not require them because they do not reason about pairs of executions with potentially different schedules.

The actual condition we require at different points in the proof (e.g., in the frame rule for $P_1 = P$ and $P_2 = R$) is the following: If $gh, gh' \models P_1 * P_2$ and $gh = gh_1 \oplus gh_2$, where $gh_1, gh'_1 \models P_1$ and $gh_2, gh'_2 \models P_2$, and similarly $gh' = gh'_1 \oplus gh'_2$, where $gh'_1, gh'_1 \models P_1$ and $gh'_2, gh'_2 \models P_2$ then we must have $gh_1, gh'_1 \models P_1$ and $gh_2, gh'_2 \models P_2$.

While this condition does not hold in general, it holds if P_1 or P_2 is precise, which is why we introduce the precision requirements. These requirements could be omitted throughout the logic by instead restricting the assertion language (in particular, allowing only specific forms of disjunctions and existential quantification that always fulfill the condition); this is what we do in our implementation, meaning that no such side condition has to be checked there. In our experience, restricting the assertion language does not limit the expressiveness of assertions in a practically meaningful way.

B.3 Noguard Side Conditions

The noguard restrictions in the atomic rules are required as a result of our definition of safety (Definition C.1), which is essentially trivially true for states with an empty resource context but non-empty guard states, since generally guards are meaningful only if there is a shared resource. However, in the atomic-rules, if P were to contain a guard assertion, then its premise would be a Hoare triple with an empty resource context and a guard in the precondition. Thus, an induction proof would fail at this point, since the premise does not guarantee anything about the execution of c .

To prevent this problem without having to complicate the definition of safety only for this one case, we enforce that P and Q contain no guards here. In practice, this side condition is not restrictive; for any atomic statement, all guards held by the current thread that are not used for the action performed by the atomic statement can be framed around it, and it is never necessary to have another guard in P to verify the body of the atomic statement.

C DETAILS ON SOUNDNESS

Here, we provide our formal definition of the safety predicate safe_n . We use \widehat{gh} to range over extended heaps whose guard states are \perp and whose permission heaps have full permission of every location in their domain, and \bar{gh} to range over extended heaps whose permission heaps have full permission of every location in their domain (but who may have arbitrary guard states):

Definition C.1. $\text{safe}_0(\Gamma_\perp, c, \langle s, gh \rangle, \Sigma)$ holds always.

$\text{safe}_{n+1}(\perp, c, \langle s, gh \rangle, \Sigma)$ holds iff

- (1) If $c = \text{skip}$ then $\langle s, gh \rangle \in \Sigma$.
- (2) c will never abort from any state $\langle s, \text{norm}(\widehat{gh}) \rangle$ s.t. $\widehat{gh} = gh \oplus gh_f$ for some gh_f .

- (3) For all $\widehat{gh}, gh_f, c', s', h'$, where $\widehat{gh} = gh \oplus gh_f$, if $c, \langle s, \text{norm}(\widehat{gh}) \rangle \rightarrow c', \langle s', h' \rangle$, then there is some gh' and some \widehat{gh}' s.t. $gh' = \text{norm}(\widehat{gh}')$ and $\widehat{gh}' = gh' \oplus gh_f$, and $\text{safe}_n(\perp, c', \langle s', gh' \rangle, \Sigma)$. $\text{safe}_{n+1}(\Gamma, c, \langle s, gh \rangle, \Sigma)$ holds iff
- (1) If $c = \text{skip}$ then $\langle s, gh \rangle \in \Sigma$.
 - (2) c will never abort from any state $\langle s, \text{norm}(\widehat{gh}) \rangle$ s.t. $\widehat{gh} = gh \oplus gh_f \oplus gh_f$ for some gh_f, gh_f s.t. gh_f satisfies $I(v)$ for some v and \widehat{gh} is consistent with Γ from some v_0 .
 - (3) For all $\widehat{gh}, gh_f, gh_f, c', s', h', v_0$, where $\widehat{gh} = gh \oplus gh_f \oplus gh_f$ and gh_f satisfies $I(v)$ for some v and \widehat{gh} is consistent with Γ from v_0 , if $c, \langle s, \text{norm}(\widehat{gh}) \rangle \rightarrow c', \langle s', h' \rangle$, then there is some $gh', gh_f', \widehat{gh}'$ s.t. gh_f' fulfills $I(v')$ for some v' , \widehat{gh} is consistent with Γ from v_0 , $gh' = \text{norm}(\widehat{gh}')$ and $\widehat{gh}' = gh' \oplus gh_f' \oplus gh_f$, and $\text{safe}_n(\Gamma, c', \langle s', gh' \rangle, \Sigma)$.

That is, the two non-trivial cases where $n \neq 0$ both contain three almost identical clauses: The first requires that if the state is final, it is contained in Σ ; the second requires that the program will not abort from the current state extended with an arbitrary frame gh_f and, if there is a shared resource, a resource heap gh_f that satisfies the invariant. The extension by an arbitrary heap is standard in CSL and necessary to prove the frame rule; additionally, here, it is used to “complete” the permission heap to contain full permissions for every location in its domain. Additionally, recall that, since the operational semantics is defined on ordinary heaps without permissions or guard states, we have to normalize the extended heap first before talking about executions from this state. The third clause requires that performing a single step from the given state (again, extended with an arbitrary frame gh_f and, if there is a shared resource, a resource heap gh_f , and normalized to get a normal heap) preserves the frame, and that recursively, the next n steps will also be safe from the new state. Additionally, in the second case (where there is a shared resource), it requires that the state remains consistent w.r.t. some initial value v_0 , i.e., that the state of the resource heap continues to be a possible result of applying the actions with the arguments specified in the guards to the initial value v_0 .

To prove soundness of the rule SHARE, we prove a lemma that expresses that, for a valid resource specification, any possible order of actions in two executions that fulfills the precondition predicates PRE_s and PRE_i for all i will lead to the same final abstract value:

LEMMA C.2. *If Γ is valid, $\alpha(v_0) = \alpha(v'_0)$, v is consistent with v_0 through some sequence of shared action applications with argument multiset args_s and unique action applications with argument sequences args_i for all i , and similarly v' is consistent with Γ from v'_0 through some sequence of shared action applications with argument multiset args'_s and unique action applications with argument sequences args'_i s.t. $\text{PRE}_s(\text{args}_s, \text{args}'_s)$ and $\text{PRE}_i(\text{args}_i, \text{args}'_i)$ for all i , then $\alpha(v) = \alpha(v')$.*

D PARTIAL ACTIONS

Our actions have to be total w.r.t. the value of the shared resource. In this section, we will explain the reason for this requirement, and how restrictive it is in practice. Consider an example consisting of two *producer* threads which produce some data and put it in a shared queue, and a single *consumer* thread that reads the data from the queue and subsequently processes it in some way. For illustration purposes, we limit the size of the buffer to one.

Both producing and consuming an item are now possible only in certain states of the shared resource: The consumer can consume an item only when the queue is non-empty, and the producers can produce an item only when it is empty (since we limited the size of the buffer to 1). Thus, in such an implementation, threads would have to block before acquiring the shared resource until




Fig. 10. Two executions of a producer-consumer example (left and right). When using pairwise swaps to permute one into the other, some resulting intermediate states (marked in red) are impossible (single-element buffers with more than one or less than zero elements).

they can execute their action; this could be supported e.g. using a command `atomiccwhere` that waits until e is true before executing c atomically. As a result, the natural definitions of the produce and consume actions are *partial* functions, which is not allowed by our framework.

The reason for this is that, if one were to allow partial actions and adapt the abstract commutativity requirement s.t. it requires actions to commute only when all involved function applications are defined, our technique would become unsound. as illustrated by the example in Figure 10, which shows two possible traces of this program (left and right, ignore the rest of the figure for now), where both producer threads produce only a single item (thread τ_1 produces the number 1, and thread τ_3 produces the number 3, and the consumer thread τ_2 consumes both), starting from an empty buffer. In the execution on the left, τ_1 produces its item first, whereas in the execution on the right, τ_3 is faster. The depicted states have the form l, s , where l is the buffer (we write $[]$ when the buffer is empty and $[v]$ for a buffer containing value v), and s is the sequence of produced values. Note that the multiset of the produced items is identical in both executions, whereas the sequence of produced items is not.

In our scenario, the two executions, left and right, are the *only* possible executions of our program, all other schedules are impossible: Since the buffer is initially empty, the first action *has to* be a produce-action (the consumer would block until the queue is non-empty), and after a produce-action has happened, the buffer is full, meaning that the next action can only be a consume-action. However, as a result, if produce and consume are partial functions, it is impossible to use pairwise swaps to permute one execution to the other using pairwise swaps: In fact, no swaps are possible at all.

However, the soundness of our technique depends on the ability to permute one sequence of actions into the other via pairwise swaps, since our commutativity criterion only requires pairwise commutativity for all actions. The example illustrates this: A weakened abstract commutativity criterion where any two actions f_1 and f_2 have to commute only when both orders of applications $f_1(f_2(x))$ and $f_2(f_1(x))$ are defined is unsound here, since it is trivially fulfilled for *any* abstraction,

$$\begin{aligned}
\text{type}_{\text{PC}}(v) &\equiv \langle \text{Either}[\text{Nat}, \text{Seq}[\text{Int}]], \text{Seq}[\text{Int}] \rangle \\
\alpha_{\text{PC}}(\langle v_1, v_2 \rangle) &\equiv \text{ms}(v_2) \\
\text{SharedActions}_{\text{PC}} &\equiv \{\text{PROD}\} \\
\text{UniqueActions}_{\text{PC}} &\equiv \{\text{CONS}\} \\
f_{\text{PROD}}(v, a) &\equiv \begin{cases} \langle \text{Right}(xs ++ [a]), s ++ [a] \rangle & \text{if } v = \langle \text{Right}(xs), s \rangle \\ \langle \text{Right}([], s ++ [a]) \rangle & \text{if } v = \langle \text{Left}(-1), s \rangle \\ \langle \text{Left}(-n), s ++ [a] \rangle & \text{if } v = \langle \text{Left}(-(n+1)), s \rangle \end{cases} \\
\text{pre}_{\text{PROD}}(a) &\equiv \text{Low}(a) \\
f_{\text{CONS}}(v, a) &\equiv \begin{cases} \langle \text{Right}(xs), s \rangle & \text{if } v = \langle \text{Right}(x :: xs), s \rangle \\ \langle \text{Left}(-1), s \rangle & \text{if } v = \langle \text{Right}([], s) \rangle \\ \langle \text{Left}(-(n+1)), s \rangle & \text{if } v = \langle \text{Left}(-n), s \rangle \end{cases} \\
\text{pre}_{\text{CONS}}(a) &\equiv \text{true}
\end{aligned}$$

Fig. 11. Producer-consumer resource specification. The second part of the tuple represents the sequence of produced items, whereas the first part represents the current contents of the buffer/queue. We mark real queue contents as *right* and queues with negative contents as *left*. The abstraction function returns the multiset view of the sequence of produced items.

since it is never the case that both orders are defined: In particular, $\text{cons}(\text{cons}(x))$ and $\text{prod}(\text{prod}(x))$ are never defined, and for any x , only one of $\text{prod}(\text{cons}(x))$ and $\text{cons}(\text{prod}(x))$ is defined. As a result, no check would prevent us from, for example, declaring the entire sequence of produced items (as opposed to its multiset abstraction) to be low, which would be unsound.

Thus, we require actions to be total functions on the resource value, but users can still support examples like the shown one by *artificially* making actions total, possibly by modifying the type of the shared state in such a way that it can represent orders of actions that cannot happen in reality. In Figure 10 (and in the producer-consumer example in our evaluation), we did this by 1) allowing the queue to be longer than one element, and 2) allowing it to contain less than zero elements (in which case we only track the (negative) length of the sequence). That is, we define consuming in such a way that consuming from an empty queue leaves its state as -1 (and consuming from a queue whose size is already $-n$ results in a state $-(n+1)$). Conversely, producing when the queue is full adds an item to the queue, producing an item when a queue has state -1 results in an empty queue, and producing when the state is $-(n+1)$ results in state $-n$. We show the resulting resource specification, with its ghost state, in Figure 11.

We then prove that these definitions of the produce- and consume-actions commute modulo our abstraction, which they do when using the multiset-abstraction (but not when using the set of produced items directly). Now, conceptually, it is possible to permute the execution on the left into the one on the right using swaps of neighboring actions, as shown in Figure 10 in the middle (we show two possible sequences of swaps that have the same result). Both of these swaps lead to intermediate states that cannot occur in any real run (highlighted in red), but the existence of these sequences, along with our commutativity criterion, guarantees that each such swap will leave (the abstract view of) the final state unchanged; therefore, we can conclude that both real executions, left and right, will lead to end results that are equal modulo abstraction.

Crucially, *any* way of making the action functions total is valid, as long as the functions result in a resource specification that fulfills our validity criteria. In practice, we have found that there are a few typical patterns that allow one to accomodate most relevant cases: For example, to model reading from an empty data structure, it is generally sufficient to store its (negative) length instead of its contents, as in our example.

E HYPERVIPER EXAMPLE

Here, we show a slightly simplified version of the relevant parts of an example similar to Figure 3 encoded into HYPERVIPER. First, users have to declare resource specifications that define allowed actions as well as the abstract view of the resource:

```
// Resource specification
specType MapLock {
  type Map[Int, Int]
  invariant(l, v) = [l.lockMap |-> ?mp && isMap(mp) && v == mapValue(mp)]
  alpha(v): Set[Int] = keys(v)

  actions = [(Put, Pair[Int, Int], duplicable)]

  action Put(v, arg)
    requires low(fst(arg))
    { (put(v, fst(arg), snd(arg))) }

  noLabels = N()
}
```

Then, the program itself is verified using said resource specifications. We first show the main method which spawns a variable number of worker threads (using fork and join statements). Guard assertions `sguard` correspond directly to the ones shown in this paper, but HYPERVIPER uses a function `sguardArgs` to refer to the argument of the current guard. Note that guards are explicitly associated with resources, since HYPERVIPER supports multiple shared resource in one program, and that shared action guards have to be explicitly split (in the first loop) and merged (in the second loop) using ghost statements. Sometimes, it is necessary to assert lemmas about e.g. multisets in order to prove that specific splits are allowed, since the verifier cannot prove these completely without user guidance.

```
// general definitions
function N(): Int
  ensures result > 1

function nTargets(): Int

// class TargetInfo
field addr: Int
field reason: Int

define targetInfo(r) (
  [r.addr |-> _] && [r.reason |-> ?r && low(r)]
)

field lockMap : Ref

method main(inputs: Seq[Seq[Ref]])
  requires lowEvent && low(|inputs|) && |inputs| == N()
  requires forall ip: Int :: ip >= 0 && ip < N() ==> low(|inputs[ip]|)
  requires forall ip: Int, jp: Int :: ip >= 0 && ip < N() && jp >= 0 && jp < |inputs[ip]|
    ==> targetInfo(inputs[ip][jp])
{
  var m : Ref
```

```

m := createMap()
var l : Lock
l := new(lockMap)
l.lockMap := m
share[MapLock](l, empty())
var i : Int := 0
var threads : Array
threads := emptyArray(|inputs|)
while (i < |inputs|)
  invariant i >= 0 && i <= |inputs| && low(i)
  invariant forall j : Int :: j >= 0 && j < |inputs| && j >= i
    ==> [loc(threads, j).arr_loc |-> _]
  invariant forall j : Int :: j >= 0 && j < i
    ==> joinable(loc(threads, j), inputs[j], l, j)
  invariant sguard[MapLock, Put](l, intervalSet(i, N())) &&
    sguardArgs[MapLock, Put](l, intervalSet(i, N())) == Multiset[Pair[Int, Int]]()
  invariant forall ip: Int, jp: Int :: ip >= i && ip < N() && jp >= 0 && jp < |inputs[ip]|
    ==> targetInfo(inputs[ip][jp])
{
  var t: Thread
  assert Set(i) union intervalSet(i+1, N()) == intervalSet(i, N())
  split[MapLock, Put](l, Set(i), intervalSet(i+1, N()),
    Multiset[Pair[Int, Int]](), Multiset[Pair[Int, Int]]())
  t := fork worker(inputs[i], l, i)
  loc(threads, i).arr_loc := t
  fold joinableArrayEntry(loc(threads, i), inputs[i], l, i)
  i := i + 1
}

i := 0
while (i < |inputs|)
  invariant i >= 0 && i <= |inputs| && low(i)
  invariant forall j : Int :: j >= 0 && j < |inputs| && j >= i
    ==> joinable(loc(threads, j), inputs[j], l, j)
  invariant forall j : Int :: j >= 0 && j < i ==> [loc(threads, j).arr_loc |-> _]
  invariant sguard[MapLock, Put](l, intervalSet(0, i)) &&
    allPre[MapLock, Put](sguardArgs[MapLock, Put](l, intervalSet(0, i)))
{
  unfold joinableArrayEntry(loc(threads, i), inputs[i], l, i)
  join[worker](loc(threads, i).arr_loc)
  assert Set(i) union intervalSet(0, i) == intervalSet(0, i + 1)
  merge[MapLock, Put](l, Set(i), intervalSet(0, i))
  i := i + 1
}
unshare[MapLock](l)

var res: Seq[Int]
res := toSeq(keys(mapValue(l.lockMap)))
print(res)
}

```

Finally, for the worker method, we prove that the invariant that all arguments of the Put action fulfill its precondition is maintained. The with-command, which acquires a shared resource and then atomically executes its body, has to be annotated with the action it is supposed to perform.

```

method worker(inputs: Seq[Ref], l: Lock, lbl: Int)
  requires lowEvent && low(|inputs|) && sguard[MapLock, Put](l, Set(lbl)) &&
    sguardArgs[MapLock, Put](l, Set(lbl)) == Multiset[Pair[Int, Int]]()
  requires forall jp: Int :: jp >= 0 && jp < |inputs| ==> targetInfo(inputs[jp])
  ensures sguard[MapLock, Put](l, Set(lbl)) &&
    allPre[MapLock, Put](sguardArgs[MapLock, Put](l, Set(lbl)))
{

```

```

var i : Int := 0
while (i < |inputs|)
  invariant i >= 0 && i <= |inputs| && low(i) && sguard[MapLock, Put](l, Set(lbl)) &&
    allPre[MapLock, Put](sguardArgs[MapLock, Put](l, Set(lbl)))
  invariant forall jp: Int :: jp >= 0 && jp < |inputs| ==> targetInfo(inputs[jp])
{
  var k: Int
  var v: Int
  k := inputs[i].addr
  v := inputs[i].reason
  with[MapLock] l performing Put(pair(k, v)) at lbl {
    var tmp : Map[Int, Int]
    tmp := mapValue(l.lockMap)
    mapPut(l.lockMap, k, v)
  }
  i := i + 1
}

```