

Concurrent Abstract Predicates

– long version –

(Draft, submitted for publication)

Thomas Dinsdale-Young¹, Mike Dodds², Philippa Gardner¹,
Matthew Parkinson², and Viktor Vafeiadis²

¹ Imperial College, London

² University of Cambridge

Abstract. Abstraction is key to understanding and reasoning about large computer systems. Abstraction is easy to achieve if the relevant data structures are disjoint, but rather difficult when they are partially shared, as is the case for concurrent modules. We present a program logic for reasoning abstractly about data structures that gives us a fiction of disjointness and permits compositional reasoning. The internal details of a module are completely hidden from the client by *concurrent abstract predicates*. We reason about a module’s implementation using separation logic with permissions, and provide abstract specifications for use by client programs using concurrent abstract predicates. We illustrate our abstract reasoning by building two implementations of a lock module on top of hardware instructions, and two implementations of a concurrent set module on top of the lock module.

1 Introduction

When designing large physical systems, we use both abstraction and locality to hide irrelevant aspects of the system. For example, when building a house in London, we do not consider the gravitational forces exerted by the brick molecules, nor what the weather is like in Paris. Similarly, we use abstraction and locality when designing and reasoning about large computer systems. Locality allows us to consider small parts of a system in isolation, while abstraction gives us a structured view of the system, because components can be used purely in terms of their abstract properties.

Locality can be often used directly to achieve a measure of abstraction. Using separation logic [18] we can prove that a module runs within a particular structure, disjointly from other modules. If the structure is manipulated only by module functions, it can be represented just in terms of its abstract properties, by an abstract predicate [23]. For example, we can implement a set by a linked list. This set can be represented by an abstract predicate asserting that, for example, “the set is $\{5, 6\}$ ”. A client can then reason about the set without reasoning about its internal implementation: given “the set is $\{5, 6\}$ ”, the client can infer that, after deleting 6, that “the set is $\{5\}$ ”.

Often, however, fine-grained reasoning at the abstract level cannot be captured by traditional abstract predicates. It can be captured by the abstract reasoning introduced by context logic [4, 5], which reasons at the same level of abstraction as the module. With our set-module example, the client can reason about a set in terms of its individual elements, manipulating the element 6 separately from the rest of the set. This fine-grained abstraction is not supported using traditional abstract predicates, because locality in the abstraction need not correspond to locality in the implementation [8]. If the set module is implemented as a list, then to manipulate individual elements of the set, the implementation traverses the global list structure. Therefore, individual elements are not represented disjointly in the implementation, and fine-grained reasoning is not possible using simple abstract predicates.

For concurrency, fine-grained reasoning is essential for compositional reasoning about modules shared between threads. For example, we may wish to manipulate individual elements of the set in separate threads, whereas, in the implementation, threads may access a shared list structure simultaneously. This sharing breaks the abstract view of disjointness between the elements. Even the humble lock breaks disjointness, since acquiring the lock necessarily involves a race to decide who acquires the lock. Existing systems for concurrent reasoning either assume locks as primitive or do not provide abstraction for locks.

We present a program logic which allows abstraction in the presence of sharing, introducing *concurrent abstract predicates*. These predicates present a fiction of disjointness [8]; that is, they can be used *as if* each abstract predicate represents disjoint resource, whereas in fact resources are shared between predicates. For example, given a set implemented as a linked list we can write abstract predicates asserting “the set contains 5, which I control” and “the set contains 6, which I control”. Both of these predicates assert properties about the same shared structure, and both can be used at the same time by separate concurrent threads: for example, elements can be deleted concurrently from a set.

Concurrent abstract predicates capture information about the permitted changes to the shared structure. In the case of the set predicates, each predicate gives the thread full control over a particular element of the set. Only the thread owning the predicate can remove this element. We implement this control using resource permissions [9], with the property that the permissions must ensure that a predicate is *self-stable*: that is, immune from interference from the surrounding environment. Predicates are thus able to specify independent properties about the data, even though the data is shared.

With our program logic, a module implementation can be verified using a low-level specification. This low-level specification can be abstracted to give a high-level specification expressed using concurrent abstract predicates. Clients of the module can then be verified purely in terms of this abstract specification, without reference to the module’s implementation. We demonstrate this methodology by providing two implementations of a lock library satisfying the same abstract lock specification, and then using this specification to build two implementations of a concurrent set satisfying the same abstract set specification. At each level, we

reason entirely abstractly, avoiding reasoning about the implementation of the preceding level. We therefore demonstrate that concurrent abstract predicates provide the necessary abstraction for compositional reasoning about concurrent systems.

In this paper, we use a very simple low-level language with pointers and functions for exposition. The work naturally extends to object-oriented languages using Parkinson and Bierman’s logic for Java [23]. The object-oriented features are orthogonal to concurrency abstraction we present in this paper.

Paper Structure. In §2 we provide an informal development of our approach using the simple example of an abstract lock specification. We show how our approach can be used to validate two implementations of a lock. In §3 we give an implementation of a set, which makes use of a lock conforming to our abstract specification. We validate the set using our abstract lock specification. In §4 we give a syntax for logical assertions and a proof system for judgements about programs. We define soundness for our logic and sketch our soundness argument. Finally, in §5 we draw conclusions and relate our approach to other work.

In this long version we include appendices give full technical details for our approach. Appendix A defines the operational semantics of our programming language. This semantics omits logical annotations, and deals purely with concrete state. Appendix B gives a full soundness argument for the proof system.

2 Informal Development

We develop our core idea: to define abstract specifications for concurrent modules and prove that concrete module implementations satisfy these specifications. We motivate our work using a lock module, one of the simplest examples of concurrent resource sharing. We define an abstract specification for locks, and give two lock implementations satisfying the specification.

2.1 Lock Specification

A typical lock module has the functions $\text{lock}(x)$ and $\text{unlock}(x)$, for acquiring and releasing a lock respectively. It also has a mechanism for constructing locks, such as $\text{makelock}(n)$, which allocates a lock followed by a contiguous block of memory of size n . We specify these functions using the following local Hoare triples:

$$\begin{array}{lll} \{\text{isLock}(x)\} & \text{lock}(x) & \{\text{isLock}(x) * \text{Locked}(x)\} \\ \{\text{Locked}(x)\} & \text{unlock}(x) & \{\text{emp}\} \\ \{\text{emp}\} & \text{makelock}(n) & \left\{ \begin{array}{l} \exists x. \text{ret} = x \wedge \text{isLock}(x) * \text{Locked}(x) \\ * (x+1) \mapsto _ * \dots * (x+n) \mapsto _ \end{array} \right\} \end{array}$$

This abstract specification, which is independent of the underlying implementation, is presented by the module to the client.³ The assertions `isLock(x)` and `Locked(x)` are abstract predicates. `isLock(x)` asserts that the lock can be acquired by the thread, while `Locked(x)` asserts that the thread holds the lock. Below we give a concrete interpretation of these predicates for a simple compare-and-swap lock. The logical connective $*$ is the separating conjunction from separation logic: an assertion $p * q$ asserts that the state can be split disjointly into two parts, one satisfying p and the other satisfying q .

The abstract predicates satisfy the following axioms, which the module also presents to the client:

$$\begin{aligned} \text{isLock}(x) &\iff \text{isLock}(x) * \text{isLock}(x) \\ \text{Locked}(x) * \text{Locked}(x) &\iff \text{false} \end{aligned}$$

The first axiom allows the client to share freely the knowledge that x is a lock that can be locked.⁴ The second axiom implies that a lock can only be locked once. With the separation logic interpretation of triples (given in §4.5), the client can infer that, if `lock(x)` is called twice in succession, then the program will not terminate as the post-condition is not satisfiable.

2.2 Example: A Compare-and-Swap Lock

Consider a simple compare-and-swap lock implementation:

<pre>lock(x) { local b; do ⟨b := !CAS(&x, 0, 1)⟩ while(b) }</pre>	<pre>unlock(x) { ⟨[x] := 0⟩ }</pre>	<pre>makelock(n) { local x := alloc(n+1); [x] := 1; return x; }</pre>
---	---------------------------------------	---

Interpretation of Abstract Predicates. We relate the lock implementation to our lock specification by giving a concrete interpretation of the abstract predicates. The predicates are not just interpreted as assertions about the internal state of the module, but also as assertions about the internal *interference* of the module: that is, how concurrent threads can modify shared parts of the internal state.

To describe this internal interface, we extend separation logic with two assertions, the shared region assertion $\boxed{P}_{I(\vec{x})}^r$ and the permission assertion $[A]_\pi^r$. The shared region assertion $\boxed{P}_{I(\vec{x})}^r$ specifies that there is a shared region of memory, identified by label r , and that the entire shared region satisfies P . The shared

³ This specification resembles those used in the work of Gotsman *et al.* [13] and Hobor *et al.* [17] on dynamically-allocated locks. However, our system is much more general, as it can reason about abstract specifications other than locks.

⁴ We do not record the splittings of `isLock(x)`, although we could use permissions [3, 2] to explicitly track this information.

state is indivisible so that all threads maintain a consistent view of it. This is expressed by the logical equivalence $\boxed{P}_{I(\vec{x})}^r * \boxed{Q}_{I(\vec{x})}^r \Leftrightarrow \boxed{P \wedge Q}_{I(\vec{x})}^r$ for shared regions. The possible actions on the state are declared by the environment $I(\vec{x})$.

The permission assertion $[A]_\pi^r$ specifies that the thread has permission π to perform action A over region r , provided the action is declared in the environment. Following Boyland [3], the permission π can be: the fractional permission, $\pi \in (0, 1)$, denoting that both the thread and the environment can do the action; or the full permission, 1, denoting that the thread can do the action but the environment cannot.⁵ We now have the machinery to interpret our lock predicates concretely:

$$\begin{aligned} \text{isLock}(x) &\equiv \exists r, \pi. [\text{LOCK}]_\pi^r * \boxed{(x \mapsto 0 * [\text{UNLOCK}]_1^r) \vee x \mapsto 1}_{I(r,x)}^r \\ \text{Locked}(x) &\equiv \exists r. [\text{UNLOCK}]_1^r * \boxed{x \mapsto 1}_{I(r,x)}^r \end{aligned}$$

The abstract predicate $\text{isLock}(x)$ is interpreted by the concrete, implementation-specific assertion on the right-hand side. This specifies that the local state contains the permission $[\text{LOCK}]_\pi^r$, meaning that the thread can acquire the lock. It also asserts that the shared region satisfies the module's invariant: either the lock is unlocked ($x \mapsto 0$) and the region holds the full permission $[\text{UNLOCK}]_1^r$ to unlock the lock; or the lock is locked ($x \mapsto 1$) and the unlocking permission is gone (the thread that acquired the lock will have it).

Meanwhile, the abstract predicate $\text{Locked}(x)$ is interpreted as the permission assertion $[\text{UNLOCK}]_1^r$ in the local state, giving the current thread full permission to unlock the lock in region r , and the shared region assertion, stating that the lock is locked ($x \mapsto 1$).

The actions permitted on the lock's shared region are declared in $I(r, x)$. Actions describe how either the current thread or the environment may change the shared state. They have the form $A: P \rightsquigarrow Q$, where assertion P describes the part of the shared state required to do the action and Q describes the part of the state after the action. The actions for the lock module are

$$I(r, x) \stackrel{\text{def}}{=} \left(\begin{array}{l} \text{LOCK: } x \mapsto 0 * [\text{UNLOCK}]_1^r \rightsquigarrow x \mapsto 1, \\ \text{UNLOCK: } x \mapsto 1 \rightsquigarrow x \mapsto 0 * [\text{UNLOCK}]_1^r \end{array} \right)$$

The LOCK action requires that the shared region contains the unlocked lock ($x \mapsto 0$), and full permission $[\text{UNLOCK}]_1^r$ to unlock the lock. The result of the action is to lock the lock ($x \mapsto 1$) and to move the full unlock permission to the thread's local state ($[\text{UNLOCK}]_1^r$ has gone from the shared state). The movement of $[\text{UNLOCK}]_1^r$ into local state allows the locking thread to release the lock afterwards. Note that local state is not explicitly represented in the action; since interference only happens on shared state, actions do not need to be prescriptive about local state.

⁵ The semantics also contains a zero permission, 0, denoting that the thread may not do the action but the environment may. However, this is not a useful assertion to make, so we forbid it.

The UNLOCK action requires that the shared region r contains the locked lock ($x \mapsto 1$). The result of the action is to unlock the lock ($x \mapsto 0$) and move the $[\text{UNLOCK}]_1^r$ permission into the shared state. The thread must have $[\text{UNLOCK}]_1^r$ in its local state in order to move it to the shared state as a result of the action.

Notice that UNLOCK is self-referential. The action moves exclusive permission on itself out of local state. Consequently, a thread can only apply UNLOCK once (intuitively, a thread can only release a lock once without locking it again). In §4.2, we discuss how our semantics supports such self-referential actions.

The abstract predicates must be *self-stable* with respect to the actions: that is, for any action permitted by the module (actions in $I(r, x)$), the predicate will remain true. Self-stability ensures that a client can use these predicates without having to consider the module's internal interference. For example, assume that the predicate $\text{Locked}(x)$ is true. There are two actions the environment can perform that can potentially affect the location x :

- LOCK, but this action does not apply, as x has value 1 in the shared state of $\text{Locked}(x)$; and
- UNLOCK, but this action also does not apply, as full permission on it is in the local state of $\text{Locked}(x)$.

The module implementer must show that the concrete interpretation of the predicates satisfies the axioms presented to the client. In our example, the first axiom, that only a single $\text{Locked}(x)$ can exist, follows from the presence in the local state of full permission on UNLOCK. The second, that $\text{isLock}(x)$ can be split, follows from the fact that non-exclusive permissions can be arbitrarily subdivided and that $*$ behaves additively on shared region assertions.

Verifying the Lock Implementation. Given the definitions above, the lock implementation can be verified against its specification; see Fig. 1 and Fig. 2.

For the unlock case, the atomic update $\langle [x] := 0 \rangle$ is allowed, because it can be viewed as performing the UNLOCK action, full permission for which is in the local state. The third assertion specifies that the permission $[\text{UNLOCK}]_1^r$ has moved from the local state to the shared region r as stipulated by the unlock action. This assertion is not, however, stable under interference from the environment since another thread could acquire the lock. It does imply the fourth assertion, which is stable under such interference. The semantics of assertions allows us to forget about the shared region, resulting in the postcondition, **emp**.

For the lock case, the key proof step is the atomic compare-and-swap command in the loop. If successful, this command updates the location $[x]$ in the shared state region from 0 to 1. This update is allowed because of the permission $[\text{LOCK}]_\pi^r$ in the local state and the action in $I(r, x)$. The postcondition of the CAS specifies that either location x has value 1 and the unlock permission has moved into the local state as stipulated by the LOCK action, or nothing has happened and the precondition is still satisfied. This postcondition is stable and so the Hoare triple is valid.

For the makelock case, the key proof step is the creation of a fresh shared region and its associated permissions. Our proof system includes a *repartitioning*

$$\begin{array}{c|c}
\begin{array}{l}
\{\text{isLock}(x)\} \\
\text{lock}(x) \{ \\
\left\{ \begin{array}{l} \exists r, \pi. \\ (x \mapsto 0 * [\text{UNLOCK}]_1^r \vee x \mapsto 1)_{I(r,x)}^r \\ * [\text{LOCK}]_\pi^r \end{array} \right\} \\
\text{local } b; \\
\text{do} \\
\left\{ \begin{array}{l} \exists r, \pi. \\ (x \mapsto 0 * [\text{UNLOCK}]_1^r \vee x \mapsto 1)_{I(r,x)}^r \\ * [\text{LOCK}]_{\pi \in (0,1]}^r \end{array} \right\} \\
\langle b := !\text{CAS}(\&x, 0, 1); \\
\left\{ \begin{array}{l} \exists r, \pi. \left(\boxed{x \mapsto 1}_{I(r,x)}^r * [\text{LOCK}]_\pi^r * [\text{UNLOCK}]_1^r * b = \text{false} \right) \vee \\ \left(\left((x \mapsto 0 * [\text{UNLOCK}]_1^r \vee x \mapsto 1)_{I(r,x)}^r * [\text{LOCK}]_\pi^r * b = \text{true} \right) \right) \end{array} \right\} \\
\text{while}(b) \\
\left\{ \exists r. \boxed{x \mapsto 1}_{I(r,x)}^r * [\text{LOCK}]_\pi^r * [\text{UNLOCK}]_1^r * b = \text{false} \right\} \\
\} \\
\{\text{isLock}(x) * \text{Locked}(x)\}
\end{array}
&
\begin{array}{l}
\{\text{Locked}(x)\} \\
\text{unlock}(x) \{ \\
\left\{ \exists r. [\text{UNLOCK}]_1^r * \boxed{x \mapsto 1}_{I(r,x)}^r \right\} \\
\langle [x] := 0; \\
\left\{ \exists r. \left(x \mapsto 0 * [\text{UNLOCK}]_1^r \right)_{I(r,x)}^r \right\} \\
\left\{ \begin{array}{l} \exists r. \\ (x \mapsto 0 * [\text{UNLOCK}]_1^r \\ \vee x \mapsto 1)_{I(r,x)}^r \end{array} \right\} \\
\} \\
\{\text{emp}\}
\end{array}
\end{array}$$

Fig. 1. Verifying the compare-and-swap lock implementation: lock and unlock.

implication \implies , which enables us to repartition the state between regions and to create regions. In particular, we have the implication

$$P \implies \exists r. \boxed{P}_{I(\vec{x})}^r * \text{all}(I(\vec{x}))$$

which creates the fresh shared region r and full permission for all of the actions defined in $I(\vec{x})$ (denoted by $\text{all}(I(\vec{x}))$). In our example, we have

$$x \mapsto 1 \implies \exists r. \boxed{x \mapsto 1}_{I(r,x)}^r * [\text{LOCK}]_1^r * [\text{UNLOCK}]_1^r$$

The final postcondition results from the definitions of $\text{isLock}(x)$ and $\text{Locked}(x)$, and recalling that $\boxed{x \mapsto 1}_{I(r,x)}^r \Leftrightarrow \boxed{x \mapsto 1}_{I(r,x)}^r * \boxed{x \mapsto 1}_{I(r,x)}^r$.

2.3 The Proof System

We give an informal description of the proof system, with the formal details given in §4. Judgements in our proof system have the form $\Delta; \Gamma \vdash \{P\}C\{Q\}$, where Δ contains predicate definitions and axioms, while Γ presents abstract specifications of the functions used by C . The local Hoare triple $\{P\}C\{Q\}$ has the fault-avoiding partial-correctness interpretation advocated by separation logic: if the program C is run from a state satisfying P then it will not fault, but will either terminate in a state satisfying Q or not terminate at all.

```

{emp}
makelock(n) {
  local x := alloc(n + 1);
  {x ↦ _ * (x + 1) ↦ _ * ... * (x + n) ↦ _}
  [x] := 1;
  {x ↦ 1 * (x + 1) ↦ _ * ... * (x + n) ↦ _}
  {∃r. x ↦ 1I(r,x)r * [LOCK]Ir * [UNLOCK]Ir * (x + 1) ↦ _ * ... * (x + n) ↦ _}
  return x;
}
{∃x. ret = x ∧ isLock(x) * Locked(x) * (x + 1) ↦ _ * ... * (x + n) ↦ _}

```

Fig. 2. Verifying the compare-and-swap lock implementation: makelock.

The proof rule for atomic commands is

$$\frac{\vdash_{\text{SL}} \{p\} C \{q\} \quad \Delta \vdash P \Longrightarrow^{\{p\}\{q\}} Q \quad \Delta \vdash \text{stable}(P, Q)}{\Delta; \Gamma \vdash \{P\} \langle C \rangle \{Q\}} \text{ (ATOMIC)}$$

The bodies of atomic commands do not contain other atomic commands, nor do they contain parallel composition. They can thus be specified using separation logic. The first premise, $\vdash_{\text{SL}} \{p\} C \{q\}$, is therefore a valid triple in sequential separation logic, where p, q denote separation logic assertions that do not specify predicates, shared regions or interference.

The second premise, $\Delta \vdash P \Longrightarrow^{\{p\}\{q\}} Q$, is more complex. It says that the interference allowed by P enables the state to be repartitioned to Q , given the change to memory given by $\{p\}\{q\}$. In our example, when the CAS performs the update the change is $\{x \mapsto 0\}\{x \mapsto 1\}$. We also require that P and Q are *stable*, so that they cannot be falsified by concurrently executing threads. Precondition and postcondition stability is a general requirement that our proof rules have, which for presentation purposes we keep implicit in the rest of this paper.

The implication $P \Longrightarrow Q$ used earlier for constructing a new region is a shorthand for $P \Longrightarrow^{\{\text{emp}\}\{\text{emp}\}} Q$, i.e. a repartitioning where no concrete state changes. We use this implication in the rule of consequence, which we use to move permissions between regions.

$$\frac{\Delta \vdash P \Longrightarrow P' \quad \Delta; \Gamma \vdash \{P'\} C \{Q'\} \quad \Delta \vdash Q' \Longrightarrow Q}{\Delta; \Gamma \vdash \{P\} C \{Q\}} \text{ (CONSEQ)}$$

We now introduce a rule that allows us to combine a verified module with a verified client to obtain a complete verified system. The idea is that clients of the module are verified with respect to the specification of the module, without reference to the internal interference and the concrete predicate definitions.

Our proof system for programs includes abstract specifications for functions. In previous work on verifying fine-grained programs [25], interference had to be specified explicitly for each function. Here we can prove a specification for a

module, and then represent the specification abstractly, without mentioning the interference internal to the module.

As we have seen, our predicates can describe the internal interference of a module. Given this, we can define high-level specifications for a module where abstract predicates correspond to invariant assertions about the state of the module (that is, they are ‘self-stable’). As these abstract assertions are invariant, we can hide the predicate definitions and treat the specifications as abstract.

The following proof rule expresses the combination of a module with a client, hiding the module’s internal predicate definitions. (In §4.5, we show that this rule is a consequence of more fundamental rules in our proof system).

$$\frac{\Delta \vdash \{P_1\}C_1\{Q_1\} \quad \dots \quad \Delta \vdash \{P_n\}C_n\{Q_n\} \quad \Delta \vdash \Delta' \quad \Delta'; \{P_1\}f_1\{Q_1\}, \dots, \{P_n\}f_n\{Q_n\} \vdash \{P\}C\{Q\}}{\vdash \{P\} \text{ let } f_1 = C_1 \dots f_n = C_n \text{ in } C \{Q\}}$$

This rule defines a module consisting of functions $f_1 \dots f_n$ and uses it to verify a client specification $\{P\}C\{Q\}$.⁶ The rule should be read as follows:

If

- the implementation C_i of f_i satisfies the specification $\{P_i\}C_i\{Q_i\}$ under predicate assumptions Δ , for each i ;
- the axioms exposed to the client in Δ' are satisfied by the predicate assumptions Δ ; and
- the specifications $\{P_1\}f_1\{Q_1\}, \dots, \{P_n\}f_n\{Q_n\}$ and just the predicate assumptions Δ' can be used to prove the client $\{P\}C\{Q\}$;

then the composed system satisfies $\{P\} \text{ let } f_1 = C_1 \dots f_n = C_n \text{ in } C \{Q\}$.

Using this rule, we can define an abstract module specification and use this specification to verify a client program. Any implementation satisfying the specification can be used in the same place. We are only required to show that the module implementation satisfies the specification.

2.4 Example: A Ticketed Lock

We have shown that our lock specification is satisfied by a simple compare-and-swap lock. We now consider a more complex algorithm: a lock that issues tickets to clients contending for the lock. This algorithm is the standard lock algorithm used in current versions of Linux, and provides fairness guarantees for threads contending for the lock. Despite the fact that the ticketed lock is quite different from the compare-and-swap lock, we will show this module also implements our abstract lock specification.

The lock operations are defined as follows:

⁶ We could equally have applied this to a class using [23] rather than a set of functions. However, to keep the presentation simple we just use functions.

```

lock(x) {
  ⟨int i := INCR(x.next);⟩
  while(i ≠ x.owner) { }
}

unlock(x) {
  ⟨x.owner++;⟩
}

makelock(n) {
  local x := alloc(n+2);
  (x+1).owner := 0;
  (x+1).next := 1;
  return (x+1);
}

```

Here field names are encoded as offsets ($\text{.next} = 0$, $\text{.owner} = -1$).

The implementation assumes an atomic operation INCR that reads a location and increments the stored value. To acquire the lock, a client atomically increments $x.\text{next}$ and reads it into a variable i . The value of i becomes the client's ticket. The client waits for $x.\text{owner}$ to equal its ticket value i . Once this is the case, the client holds the lock. The lock is released by incrementing $x.\text{owner}$.

The algorithm is correct because (1) each ticket is held by at most one client and (2) only the thread holding the lock can increment $x.\text{owner}$.

Interpretation of Abstract Predicates. The actions for the ticketed lock are:

$$T(t, x) \stackrel{\text{def}}{=} \left(\begin{array}{l} \text{TAKE: } \exists k. ([\text{NEXT}(k)]_1^t * x.\text{next} \mapsto k \rightsquigarrow x.\text{next} \mapsto (k+1)) \\ \text{NEXT}(k): x.\text{owner} \mapsto k \rightsquigarrow x.\text{owner} \mapsto (k+1) * [\text{NEXT}(k)]_1^t \end{array} \right)$$

Intuitively, TAKE corresponds to taking a ticket value from $x.\text{next}$, and $\text{NEXT}(k)$ corresponds to releasing the lock when $x.\text{owner} = k$. The shared state contains permissions on $\text{NEXT}(k)$ for all the values of k not currently used by active threads. Note the $\exists k$ is required to connect the old and new values of the next field in the TAKE action.

The concrete interpretation of the predicates is as follows:

$$\begin{aligned} \text{isLock}(x) &\equiv \exists t. \left[\begin{array}{l} \exists k, k'. x.\text{owner} \mapsto k * x.\text{next} \mapsto k' \\ * k \leq k' * \textcircled{*} k'' > k'. [\text{NEXT}(k'')]_1^t \end{array} \right]_{T(t, x)}^t * [\text{TAKE}]_\pi^t \\ \text{Locked}(x) &\equiv \exists t, k. \left[x.\text{owner} \mapsto k * \text{true} \right]_{T(t, x)}^t * [\text{NEXT}(k)]_1^t \end{aligned}$$

($\textcircled{*}$ is the lifting of $*$ to sets; it is the separating (multiplicative) version of \forall .)

$\text{isLock}(x)$ requires values $x.\text{next}$ and $x.\text{owner}$ to be in the shared state, and that a permission on $\text{NEXT}(k)$ is in the shared state for each value greater than the current ticket $x.\text{next}$. It also requires a permission on TAKE to be in local state. $\text{Locked}(x)$ requires just that there is an exclusive permission on $\text{NEXT}(k)$ in local state for the current value, k , of $x.\text{owner}$.

Self-stability of $\text{Locked}(x)$ is ensured by the fact that the predicate holds full permission on the action $\text{NEXT}(k)$, and the action TAKE cannot affect the $x.\text{owner}$ field. Self-stability for $\text{isLock}(x)$ is ensured by the fact that TAKE preserves membership of the invariant for region t .

The predicate axioms follow immediately from the predicate definitions, as for the compare-and-swap lock.

$$\begin{array}{c|c}
\begin{array}{l}
\{\text{isLock}(x)\} \\
\text{lock}(x) \{ \\
\left\{ \begin{array}{l} \exists t, k, k'. [\text{TAKE}]_{\pi}^t \wedge k \leq k' \\ * \left[\begin{array}{l} x.\text{owner} \mapsto k * x.\text{next} \mapsto k' * \\ \bigotimes k'' > k'. [\text{NEXT}(k'')]_1^t * \text{true} \end{array} \right]_{T(t,x)}^t \\ \end{array} \right\} \\
\langle \text{int } i := \text{INCR}(x.\text{next}); \rangle \\
\left\{ \begin{array}{l} \exists t, k, k'. [\text{TAKE}]_{\pi}^t * [\text{NEXT}(i)]_1^t \wedge k \leq i \leq k' \\ * \left[\begin{array}{l} x.\text{owner} \mapsto k * x.\text{next} \mapsto k' * \\ \bigotimes k'' > k'. [\text{NEXT}(k'')]_1^t * \text{true} \end{array} \right]_{T(t,x)}^t \\ \end{array} \right\} \\
\text{while}(i \neq x.\text{owner}); \\
\left\{ \begin{array}{l} \exists t, k'. [\text{TAKE}]_{\pi}^t * [\text{NEXT}(i)]_1^t \wedge i \leq k' \\ * \left[\begin{array}{l} x.\text{owner} \mapsto i * x.\text{next} \mapsto k' * \\ \bigotimes k'' > k'. [\text{NEXT}(k'')]_1^t * \text{true} \end{array} \right]_{T(t,x)}^t \\ \end{array} \right\} \\
\} \\
\{\text{isLock}(x) * \text{Locked}(x)\}
\end{array}
&
\begin{array}{l}
\{\text{Locked}(x)\} \\
\text{unlock}(x) \{ \\
\left\{ \begin{array}{l} \exists t, k. \\ \left[\begin{array}{l} x.\text{owner} \mapsto k * \text{true} \\ * [\text{NEXT}(k)]_1^t \end{array} \right]_{T(t,x)}^t \\ \end{array} \right\} \\
\langle x.\text{owner}++ \rangle; \\
\left\{ \begin{array}{l} \exists t, k. \\ \left[\begin{array}{l} x.\text{owner} \mapsto (k+1) * \\ [\text{NEXT}(k)]_1^t * \text{true} \end{array} \right]_{T(t,x)}^t \\ \end{array} \right\} \\
\} \\
\{\text{emp}\}
\end{array}
\end{array}$$

Fig. 3. Proofs for the ticketed lock module operations: lock and unlock.

Verifying the Lock Implementation. Given the definitions above, the ticketed lock implementation can be verified against the lock specification, as shown in Fig. 3. The proofs follow the intuitive structure sketched above for the actions. That is, $\text{lock}(x)$ pulls a ticket and a permission out of the shared state, and $\text{unlock}(x)$ returns it to the shared state. (We omit the proof of makelock , which is similar to the previous example.)

3 Composing Abstract Specifications

In the previous section we showed that our system can be used to present abstract specifications for concurrent modules. In this section, we show how these concurrent specifications can be used to verify client programs, which may themselves be modules satisfying abstract specifications. We illustrate this by defining a concurrent set module specification and giving two implementations. The set implementations depend for their correctness on the lock specification presented in the previous section.

3.1 A Set Module Specification

A typical set module has three functions: $\text{contains}(h, v)$, $\text{add}(h, v)$ and $\text{remove}(h, v)$. These functions have the following abstract specifications:

$$\begin{aligned} \{\text{in}(h, v)\} \quad & \text{contains}(h, v) \quad \{\text{in}(h, v) * \text{ret} = \text{true}\} \\ \{\text{out}(h, v)\} \quad & \text{contains}(h, v) \quad \{\text{out}(h, v) * \text{ret} = \text{false}\} \\ \{\text{own}(h, v)\} \quad & \text{add}(h, v) \quad \{\text{in}(h, v)\} \\ \{\text{own}(h, v)\} \quad & \text{remove}(h, v) \quad \{\text{out}(h, v)\} \end{aligned}$$

Here $\text{in}(h, v)$ is an abstract predicate stating that the set at h contains v . Correspondingly, $\text{out}(h, v)$ says that the set does not contain v . We define $\text{own}(h, v)$ as the disjunction of these two predicates.

These assertions do not only capture knowledge about the set, but also exclusive permission to alter the set by changing whether v belongs to it. Consequently $\text{out}(h, v)$ is not simply the negation of $\text{in}(h, v)$. The exclusivity of permissions is captured by the module's axiom:

$$\text{own}(h, v) * \text{own}(h, v) \implies \text{false}$$

We can reason disjointly about set predicates, even though they may be implemented by a single shared structure. For example, consider the command

$$\text{remove}(h, v_1) \parallel \text{remove}(h, v_2)$$

$$\begin{array}{c} \{\text{own}(h, v_1) * \text{own}(h, v_2)\} \\ \{\text{own}(h, v_1)\} \parallel \{\text{own}(h, v_2)\} \\ \text{remove}(h, v_1) \parallel \text{remove}(h, v_2) \\ \{\text{out}(h, v_1)\} \parallel \{\text{out}(h, v_2)\} \\ \{\text{out}(h, v_1) * \text{out}(h, v_2)\} \end{array}$$

This command should succeed if it has the permissions to change the values v_1 and v_2 (where $v_1 \neq v_2$), and it should yield a set without v_1 and v_2 . This intuition is captured by the proof outline shown in Fig. 4.

Fig. 4. Proof outline for set module client.

3.2 Example: The Coarse-grained Set

Consider a coarse-grained set implementation, based on the lock module and the sequential set operations $\text{scontains}(y, v)$, $\text{sadd}(y, v)$ and $\text{sremove}(y, v)$.

$\text{contains}(h, v) \{$ $\quad \text{lock}(h.\text{lock});$ $\quad \text{scontains}(h.\text{set}, v);$ $\quad \text{unlock}(h.\text{lock});$ $\}$	$\text{add}(h, v) \{$ $\quad \text{lock}(h.\text{lock});$ $\quad \text{sadd}(h.\text{set}, v);$ $\quad \text{unlock}(h.\text{lock});$ $\}$	$\text{remove}(h, v) \{$ $\quad \text{lock}(h.\text{lock});$ $\quad \text{sremove}(h.\text{set}, v);$ $\quad \text{unlock}(h.\text{lock});$ $\}$
--	--	--

Interpretation of Abstract Predicates. We assume a sequential set predicate $\text{Set}(y, xs)$ that asserts that the sequential set at location y contains values xs . The predicate Set cannot be split, and so must be held by one thread at once. This enforces sequential behaviour. The sequential set operations have the following specifications with respect to Set :

$$\begin{aligned} \{\text{Set}(y, vs)\} \quad \text{scontains}(y, v) \quad & \{\text{Set}(y, vs) * \text{ret} = (v \in vs)\} \\ \{\text{Set}(y, vs)\} \quad \text{sadd}(y, v) \quad & \{\text{Set}(y, \{v\} \cup vs)\} \\ \{\text{Set}(y, vs)\} \quad \text{sremove}(y, v) \quad & \{\text{Set}(y, vs \setminus \{v\})\} \end{aligned}$$

In the set implementation, the predicate Set is held in the shared state when the lock is not locked. Then when the lock is acquired by a thread, the predicate is pulled into the thread's local state so that it can be modified according to the sequential set specification. When the lock is released, the predicate is returned to the shared state. The actions for the set module are

$$C(s, h) \stackrel{\text{def}}{=} \left(\begin{array}{l} \text{CHANGE}(v): \left(\begin{array}{l} \exists vs, ws. \text{Set}(h.\text{set}, vs) \\ * [\text{SGAP}(ws)]_1^s \wedge \\ vs \setminus \{v\} = ws \setminus \{v\} \end{array} \right) \rightsquigarrow \text{Locked}(h.\text{lock}) \\ \text{SGAP}(ws): \text{Locked}(h.\text{lock}) \rightsquigarrow \text{Set}(h.\text{set}, ws) * [\text{SGAP}(ws)]_1^s \end{array} \right)$$

The $\text{SGAP}(vs)$ action allows the thread to return the set containing vs to the shared state. The $\text{CHANGE}(v)$ action allows a thread to acquire the set from the shared state. To do so, the thread must currently hold the lock. It gives up the permission to release the lock in exchange for the set. The thread also acquires the permission $[\text{SGAP}(ws)]_1^s$, which allows it to re-acquire the lock permission by relinquishing the set, having only changed whether or not v is in the set.

To define concrete interpretations for the predicates, we first define $P_\in(h, v)$ and $P_\notin(h, v)$:

$$\begin{aligned} \text{allgaps}(s) &\equiv \textcircled{*} ws. [\text{SGAP}(ws)]_1^s \\ P_\triangleleft(h, v, s) &\equiv \exists vs. v \triangleleft vs \wedge \left(\begin{array}{l} (\text{allgaps}(s) * \text{Set}(h.\text{set}, vs)) \\ \vee \text{Locked}(h.\text{lock}) * ([\text{SGAP}(vs)]_1^s - \textcircled{*} \text{allgaps}(s)) \end{array} \right) \\ &\quad \text{where } \triangleleft = \in \text{ or } \triangleleft = \notin \end{aligned}$$

The predicate allgaps defines the set of all SGAP permissions. $P_\in(h, v, s)$ asserts that the shared state s contains either the set with contents vs , where $v \in vs$, and all possible SGAP permissions; or it contains the Locked predicate and is missing one of the SGAP permissions. The missing SGAP permission records the contents of the set when it is eventually released. $P_\notin(h, v, s)$ defines the case where $v \notin vs$.

The concrete definitions of $\text{in}(h, v)$ and $\text{out}(h, v)$ are as follows:

$$\begin{aligned} \text{in}(h, v) &\equiv \exists s. \text{isLock}(h.\text{lock}) * [\text{CHANGE}(v)]_1^s * \boxed{P_\in(h, v, s)}_{C(s, h)}^s \\ \text{out}(h, v) &\equiv \exists s. \text{isLock}(h.\text{lock}) * [\text{CHANGE}(v)]_1^s * \boxed{P_\notin(h, v, s)}_{C(s, h)}^s \end{aligned}$$

The $\text{in}(h, v)$ predicate gives a thread the permissions needed to acquire the lock, $\text{isLock}(h.\text{lock})$, and to change whether v is in the set, $[\text{SCHANGE}(v)]_1^s$. The shared state is described by the predicate $P_\in(h, v, s)$. The $\text{out}(h, v)$ predicate is defined analogously to $\text{in}(h, v)$, but with a negation on the set contains.

$\text{in}(h, v)$ and $\text{out}(h, v)$ are self-stable. For $\text{in}(h, v)$, the only actions available to another thread are $\text{SCHANGE}(w)$, where $w \neq v$, and $\text{SGAP}(vs)$, where $v \in vs$. The assertion $P_\in(h, v, s)$ is invariant under both of these changes: $\text{SCHANGE}(w)$ requires the disjunct $\text{allgaps} * \text{Set}(h.\text{set}, vs)$ to hold and leaves the disjunct $\text{Locked}(h.\text{lock}) * ([\text{SGAP}(vs)]_1 - \otimes \text{allgaps}(s))$ holding; $\text{SGAP}(vs)$ does the reverse. Similar arguments hold for $\text{out}(h, v)$.

The predicate axiom holds as a consequence of the fact that exclusive permissions cannot be combined.


Verifying the Set Implementation. Given the definitions above, we can verify the implementations of the set module. Fig. 5 shows a proof of $\text{add}(h, v)$ when the value is not in the set. The case where the value is in the set, and the proofs of remove and contains follow a similar structure.

The most interesting steps of this proof are those before and after the operation $\text{sadd}(h.\text{lock})$, when the permissions $[\text{SCHANGE}(v)]_1^s$ and $[\text{SGAP}(vs)]_1^s$ are used to repartition between shared and local state. These steps are purely logical repartitioning of the state.

3.3 The Fine-grained Set

Our previous implementation of a concurrent set used a single global lock. We now consider a set implementation that uses a sorted list with one lock per node in the list. Our algorithm (adapted from [15, §9.5]) is given in Fig. 6.

The three module functions use the function $\text{locate}(h, x)$ that traverses the list from the head h up to the position for a node holding value x , whether or not such a node is present. It begins by locking the initial node of the list. It then moves down the list by hand-over-hand locking. The algorithm first locks the node following its currently held node, and then releases the previously-held lock. The following diagram illustrates this pattern of locking:



No thread can access a node locked by another thread, or traverse past a locked node. Consequently, a thread cannot overtake any other threads accessing the list. Nodes can be added and removed from locked segments of the list. If a thread locks a node, then a new node can be inserted directly after it, as long as it preserves the sorted nature of the list. Also, if a thread has locked two nodes in sequence, then the second can be removed.

Since nodes of the list are deleted when an element is removed from the set, the lock module must provide a mechanism for disposing of locked blocks, $\text{disposelock}(x, n)$. For brevity, we omitted details of this in our exposition. In order for disposal to be sound, no other thread must be able to access the lock,

```

{out(h, v)}
add(h, v)
{
   $\exists s. \text{isLock}(h.\text{lock}) * [\text{SCHANGE}(v)]_1^s * \boxed{P_{\notin}(h, v, s)}_{C(s, h)}^s$ 
  lock(h.lock)
  {
     $\exists s. \text{isLock}(h.\text{lock}) * \text{Locked}(h.\text{lock}) * [\text{SCHANGE}(v)]_1^s * \boxed{P_{\notin}(h, v, s)}_{C(s, h)}^s$ 
    // use SCHANGE permission to extract Set predicate and SGAP permission
    {
       $\exists s. \text{isLock}(h.\text{lock}) * [\text{SGAP}(vs \cup \{v\})]_1^s * [\text{SCHANGE}(v)]_1^s * \text{Set}(h.\text{set}, vs)$ 
      *  $\boxed{\text{Locked}(h.\text{lock}) * ([\text{SGAP}(vs \cup \{v\})]_1^s - \otimes \text{allgaps}(s))}_{C(s, h)}^s$ 
    }
    sadd(h.set, v)
    {
       $\exists s. \text{isLock}(h.\text{lock}) * [\text{SGAP}(vs \cup \{v\})]_1^s * [\text{SCHANGE}(v)]_1^s * \text{Set}(h.\text{set}, vs \cup \{v\})$ 
      *  $\boxed{\text{Locked}(h.\text{lock}) * ([\text{SGAP}(vs \cup \{v\})]_1^s - \otimes \text{allgaps}(s))}_{C(s, h)}^s$ 
    }
    // use SGAP permission to put back Set and SGAP permission
    {
       $\exists s. \text{isLock}(h.\text{lock}) * \text{Locked}(h.\text{lock}) * [\text{SCHANGE}(v)]_1^s * \boxed{P_{\in}(h, v, s)}_{C(s, h)}^s$ 
      unlock(h.lock)
      {
         $\exists s. \text{isLock}(h.\text{lock}) * [\text{SCHANGE}(v)]_1^s * \boxed{P_{\in}(h, v, s)}_{C(s, h)}^s$ 
      }
    }
  }
}
{in(h, v)}

```

Fig. 5. Proof of the add(h, v) operation for the coarse-grained set module.

```

remove(h, x) {
  local p, c, z;
  (p, c) := locate(h, x);
  if (c.value == x) {
    lock(c);
    z := c.next;
    p.next := z;
    disposelock(c, 2);
  }
  unlock(p);
}

contains(h, x) {
  local p, c, b;
  (p, c) := locate(h, x);
  b := (c.value == x);
  unlock(p);
  return b;
}

add(h, x) {
  local p, c, z;
  (p, c) := locate(h, x);
  if (c.value != x) {
    z := makelock(2);
    unlock(z);
    z.value := x;
    z.next := c;
    p.next := z;
  }
  unlock(p);
}

locate(h, x) {
  local p, c;
  p := h;
  lock(p);
  c := p.next;
  while (c.value < x) {
    lock(c);
    unlock(p);
    p := c;
    c := p.next;
  }
  return(p, c);
}

```

Fig. 6. Lock-coupling list algorithm.

and so we require the `isLock` predicate to record permissions. Creating a lock gives full permission, locking requires only fractional permission and disposing a lock requires full permission in this scheme. Lock disposal is specified as

$$\left\{ \begin{array}{l} \text{isLock}(x, 1) * \text{Locked}(x) \\ * (x + 1) \mapsto _ * \dots * (x + n) \mapsto _ \end{array} \right\} \quad \text{disposelock}(x, n) \quad \{\text{emp}\}$$

and for both lock implementations we have given, the function simply deallocates the memory block.

Interpretation of Abstract Predicates. We first define some basic predicates for describing the structure of a list. List nodes consist of a lock, a value field and a link field. Because owning a node gives the permission to lock its successor, our node predicate includes the `isLock` predicate of the *next* node.

$$\begin{aligned} \text{link}(a, b) &\equiv (a.\text{next} \mapsto b) * \text{isLock}(b, 1) & \text{val}(a, v) &\equiv (a.\text{value} \mapsto v) \\ \text{node}(a, b, v) &\equiv \text{val}(a, v) * \text{link}(a, b) \end{aligned}$$

The actions for the fine-grained set module are defined by the interference environment $F(s, h)$, defined by the following assertion:

$$\text{LGAP}(n, t, x): \left\{ \begin{array}{l} \text{Locked}(n) \rightsquigarrow \text{link}(n, t) * [\text{LGAP}(n, t, x)]_1^s \\ \forall v_1, v_2. \left(\begin{array}{l} \text{Locked}(n) \\ * \text{val}(n, v_1) \\ * \text{val}(t, v_2) \end{array} \right) \rightsquigarrow \left(\begin{array}{l} [\text{LGAP}(n, t, x)]_1^s * \text{node}(y, t, x) \\ * \text{link}(n, y) * \text{val}(n, v_1) \\ * \text{val}(t, v_2) * \wedge v_1 < x < v_2 \end{array} \right) \\ \left(\begin{array}{l} \text{Locked}(n) * \\ \text{Locked}(t) * \text{val}(t, x) \end{array} \right) \rightsquigarrow \left(\begin{array}{l} [\text{LGAP}(n, t, x)]_1^s * \\ [\text{LGAP}(t, y, x)]_1^s * \text{link}(n, y) \end{array} \right) \end{array} \right.$$

$$\text{LCHANGE}(x): \quad \forall n, t. [\text{LGAP}(n, t, x)]_1^s * \text{link}(n, t) \rightsquigarrow \text{Locked}(n)$$

Intuitively, the $\text{LGAP}(n, t, x)$ actions allow a link to be replaced between n and t , such that it preserves the list's contents up to inclusion of x . Locked (and so, hidden) nodes of the list correspond to permissions on the action LGAP .

The first LGAP action simply puts the link back. The second updates the link to point to a new node y whose tail is t and value is x . We require that inserting this node maintains the ascending order of the list. This action is the basis of adding x to the set. The final action allows the link to be updated to skip an element, if that element is x . This is the basis of removing x from the set.

The $\text{LCHANGE}(x)$ action allows a thread holding the lock to a node n to access it, by obtaining $\text{link}(n, t)$, and to modify it up to adding or removing a node with value x , by obtaining the permission $[\text{LGAP}(n, t, x)]_1^s$.

The predicates $L_{\in}(h, x, s)$ and $L_{\notin}(h, x, s)$ correspond to lists containing and not containing x . Their definitions are given in Fig. 7, along with auxiliary

$$\begin{aligned}
\text{lsg}(x, y, S, []) &\equiv x = y \wedge S = \emptyset \\
\text{lsg}(x, y, S \uplus \{(x, z)\}, v :: vs) &\equiv \text{Locked}(x) * \text{val}(x, v) * \text{lsg}(z, y, S, vs) \\
\text{lsg}(x, y, S, v :: vs) &\equiv x \neq y \wedge \text{node}(x, v, z) * \text{lsg}(z, y, S, vs) \\
\text{slsg}(x, y, S, vs) &\equiv \text{lsg}(x, y, S, vs) \wedge \text{sorted}(vs) \wedge \exists vs'. vs = -\infty :: vs' @ [\infty] \\
\text{gaps}(S, s) &\equiv \bigotimes (x, y) \notin S. \bigotimes z. [\text{LGAP}(x, y, z)]_1^s * \\
&\quad \bigotimes (x, y). \in S. \exists w. \bigotimes z \neq w. [\text{LGAP}(x, y, z)]_1^s \wedge \\
&\quad \forall x, y, w, z. (x, y) \in S \wedge (w, z) \in S \Rightarrow (x = w \Leftrightarrow y = z) \\
\text{mygaps}(z, s) &\equiv \bigotimes x, y. [\text{LGAP}(x, y, z)]_1^s * \text{true} \\
L_{\triangleleft}(h, x, s) &\equiv \exists xs, S. x \triangleleft xs \wedge \text{slsg}(h, \text{nil}, S, xs) * (\text{gaps}(S, s) \wedge \text{mygaps}(x, s)) \\
&\quad \text{where } \triangleleft = \in \text{ or } \triangleleft \neq
\end{aligned}$$

Fig. 7. Predicates for lock-coupling list.

predicates. The `slsg` ('sorted list with gaps') predicate represents a list with locked segments. The `gaps` predicate tracks the unused GAP permissions in the shared state. In both cases, a carrier set S records the locked sections in the list.

The concrete interpretation of the predicates for this implementation of the set module are defined as follows:

$$\begin{aligned}
\text{in}(h, x) &\equiv \exists s. \exists \pi > 0. \text{isLock}(h, \pi) * [\text{LCHANGE}(x)]_1^s * \boxed{L_{\in}(h, x, s)}_{F(s, h)}^s \\
\text{out}(h, x) &\equiv \exists s. \exists \pi > 0. \text{isLock}(h, \pi) * [\text{LCHANGE}(x)]_1^s * \boxed{L_{\notin}(h, x, s)}_{F(s, h)}^s
\end{aligned}$$

The predicate axiom holds as a consequence of the fact that exclusive permissions cannot be combined.

Verifying the set implementation Fig. 8 gives a sketch-proof for `locate(x)`. The main loop searches through the list checking if the current element, c , is less than the element being searched for x . At the start of the loop, the thread has p locked, signified by having the $\text{LGAP}(p, c, x)$ permission in its local state. The first step of the loop locks the new current element c . This is allowed as the $\text{link}(p, c)$ predicate contains a `isLock` permissions for c . After, locking the thread additionally has a `Locked(c)` predicate in its local state.

The proof then uses the action in the local state $\text{LCHANGE}(x)$, to swap the local `Locked` predicate, for the LGAP action associate with c and the link between c and the next node. In some sense, we are publishing that we have locked this node. This thread now has two nodes locked in the list and can `UNLOCK` the first. This unlocking first involves putting the LGAP permission into the shared state, which returns the `Locked` predicate for that node to the threads local state. This allows the previous node to be safely unlocked.

As with the coarse-grained set, locking and unlocking are two-step processes. First the lock is acquired in local state, then a permission is used to change

```

locate(h, x) {
  {in(h, x)}
  p := h;
  lock(p);
  c := p.next;
  while (c.value < x) {
    {
      
$$\left\{ \begin{array}{l} \exists s. \left[ \begin{array}{l} (\exists v. \text{val}(c, v) * v < x) \wedge \exists xs, S. x \in xs \wedge \{(p, c)\} \subseteq S \wedge \text{smsg}(h, \text{nil}, S, xs) \\ * (\text{gaps}(S, s) \wedge ([\text{LGAP}(p, c, x)]_1^s - \textcircled{*} \text{mygaps}(x, s))) \end{array} \right]_{F(s, h)}^s \\ * \text{link}(p, c) * [\text{LGAP}(p, c, x)]_1^s * \exists \pi > 0. \text{isLock}(h, \pi) * [\text{LCHANGE}(x)]_1^s \end{array} \right\}$$

      lock(c);
      {
        
$$\left\{ \begin{array}{l} \exists s. \left[ \begin{array}{l} (\exists v. \text{val}(c, v) * v < x) \wedge \exists xs, S. x \in xs \wedge \{(p, c)\} \subseteq S \wedge \text{smsg}(h, \text{nil}, S, xs) \\ * (\text{gaps}(S, s) \wedge ([\text{LGAP}(p, c, x)]_1^s - \textcircled{*} \text{mygaps}(x, s))) \end{array} \right]_{F(s, h)}^s \\ * \text{link}(p, c) * [\text{LGAP}(p, c, x)]_1^s * \exists \pi > 0. \text{isLock}(h, \pi) * [\text{LCHANGE}(x)]_1^s * \text{Locked}(c) \end{array} \right\}$$

        // By LCHANGE using Locked(c)
        {
          
$$\left\{ \begin{array}{l} \exists s, z. \left[ \begin{array}{l} (\exists v. \text{val}(c, v) * v < x) \\ \wedge \exists xs, S. x \in xs \wedge \{(p, c), (c, z)\} \subseteq S \wedge \text{smsg}(h, \text{nil}, S, xs) \\ * \left( \text{gaps}(S, s) \wedge \left( \left( [\text{LGAP}(p, c, x)]_1^s \right) - \textcircled{*} \text{mygaps}(x, s) \right) \right) \end{array} \right]_{F(s, h)}^s \\ * \text{link}(p, c) * [\text{LGAP}(p, c, x)]_1^s * [\text{LGAP}(c, z, x)]_1^s * \exists \pi > 0. \text{isLock}(h, \pi) \\ * [\text{LCHANGE}(x)]_1^s * \text{link}(c, z) \end{array} \right\}$$

          // By LGAP(p, c, x)
          {
            
$$\left\{ \begin{array}{l} \exists s, z. \left[ \begin{array}{l} (\exists v. \text{val}(c, v) * v < x) \wedge \exists xs, S. x \in xs \wedge \{(c, z)\} \subseteq S \wedge \text{smsg}(h, \text{nil}, S, xs) \\ * (\text{gaps}(S, s) \wedge ([\text{LGAP}(c, z, x)]_1^s - \textcircled{*} \text{mygaps}(x, s))) \end{array} \right]_{F(s, h)}^s \\ * \text{Locked}(p) * [\text{LGAP}(c, z, x)]_1^s * \exists \pi > 0. \text{isLock}(h, \pi) * [\text{LCHANGE}(x)]_1^s * \text{link}(c, z) \end{array} \right\}$$

            unlock(p);
            p := c;
            c := p.next;
            {
              
$$\left\{ \begin{array}{l} \exists s. \left[ \begin{array}{l} (\exists v. \text{val}(p, v) * v < x) \wedge \exists xs, S. x \in xs \wedge \{(p, c)\} \subseteq S \wedge \text{smsg}(h, \text{nil}, S, xs) \\ * (\text{gaps}(S, s) \wedge ([\text{LGAP}(p, c, x)]_1^s - \textcircled{*} \text{mygaps}(x, s))) \end{array} \right]_{F(s, h)}^s \\ * [\text{LGAP}(p, c, x)]_1^s * \exists \pi > 0. \text{isLock}(h, \pi) * [\text{LCHANGE}(x)]_1^s * \text{link}(p, c) \end{array} \right\}$$

            }
          }
        {
          
$$\left\{ \begin{array}{l} \exists s. \left[ \begin{array}{l} (\exists v. \text{val}(p, v) * v < x) \wedge (\exists v'. \text{val}(c, v') * v' \geq x) \wedge \exists xs, S. x \in xs \wedge \{(p, c)\} \subseteq S \\ \wedge \text{smsg}(h, \text{nil}, S, xs) * (\text{gaps}(S, s) \wedge ([\text{LGAP}(p, c, x)]_1^s - \textcircled{*} \text{mygaps}(x, s))) \end{array} \right]_{F(s, h)}^s \\ * [\text{LGAP}(p, c, x)]_1^s * \exists \pi > 0. \text{isLock}(h, \pi) * [\text{LCHANGE}(x)]_1^s * \text{link}(p, c) \end{array} \right\}$$

        }
        return(p, c);
      }
    }
  }
}

```

Fig. 8. Proof for locate(x).

$\{ \text{in}(h, x) \}$
 $(p, c) := \text{locate}(x); \quad // \text{ Use spec. from Fig. 8}$

$$\left\{ \begin{array}{l} \exists s. \left[\begin{array}{l} (\exists v. \text{val}(p, v) * v < x) \wedge (\exists v'. \text{val}(c, v') * v' \geq x) \wedge \exists xs, S. x \in xs \wedge \{(p, c)\} \subseteq S \\ \wedge \text{smsg}(h, \text{nil}, S, xs) * (\text{gaps}(S, s) \wedge ([\text{LGAP}(p, c, x)]_1^s - \otimes \text{mygaps}(x, s))) \end{array} \right]_{F(s, h)}^s \\ * [\text{LGAP}(p, c, x)]_1^s * \exists \pi > 0. \text{isLock}(h, \pi) * [\text{LCHANGE}(x)]_1^s * \text{link}(p, c) \end{array} \right\}$$

if $(c.\text{value} == x)$ **{**
 $\text{lock}(c);$

$$\left\{ \begin{array}{l} \exists s. \left[\begin{array}{l} (\text{val}(c, x) * \text{true}) \wedge \exists xs, S. x \in xs \wedge \{(p, c)\} \subseteq S \wedge \text{smsg}(h, \text{nil}, S, xs) \\ * (\text{gaps}(S, s) \wedge ([\text{LGAP}(p, c, x)]_1^s - \otimes \text{mygaps}(x, s))) \end{array} \right]_{F(s, h)}^s \\ * [\text{LGAP}(p, c, x)]_1^s * \exists \pi > 0. \text{isLock}(h, \pi) * [\text{LCHANGE}(x)]_1^s * \text{link}(p, c) * \text{Locked}(c) \end{array} \right\}$$

 $// \text{ By LCHANGE using Locked}(c)$

$$\left\{ \begin{array}{l} \exists s, z. \left[\begin{array}{l} (\text{val}(c, x) * \text{true}) \wedge \exists xs, S. x \in xs \wedge \{(p, c), (c, z)\} \subseteq S \wedge \text{smsg}(h, \text{nil}, S, xs) \\ * \left(\text{gaps}(S, s) \wedge \left(\left([\text{LGAP}(p, c, x)]_1^s \right) * [\text{LGAP}(c, z, x)]_1^s \right) - \otimes \text{mygaps}(x, s) \right) \end{array} \right]_{F(s, h)}^s \\ * [\text{LGAP}(p, c, x)]_1^s * [\text{LGAP}(c, z, x)]_1^s * \text{link}(p, c) * \text{link}(c, z) \\ * \exists \pi > 0. \text{isLock}(h, \pi) * [\text{LCHANGE}(x)]_1^s \end{array} \right\}$$

 $z := c.\text{next};$
 $p.\text{next} := z;$

$$\left\{ \begin{array}{l} \exists s, z. \left[\begin{array}{l} (\text{val}(c, x) * \text{true}) \wedge \exists xs, S. x \in xs \wedge \{(p, c), (c, z)\} \subseteq S \wedge \text{smsg}(h, \text{nil}, S, xs) \\ * \left(\text{gaps}(S, s) \wedge \left(\left([\text{LGAP}(p, c, x)]_1^s \right) * [\text{LGAP}(c, z, x)]_1^s \right) - \otimes \text{mygaps}(x, s) \right) \end{array} \right]_{F(s, h)}^s \\ * [\text{LGAP}(p, c, x)]_1^s * [\text{LGAP}(c, z, x)]_1^s * \text{link}(p, z) * \text{isLock}(c, 1) * (c.\text{next} \mapsto z) \\ * \exists \pi > 0. \text{isLock}(h, \pi) * [\text{LCHANGE}(x)]_1^s \end{array} \right\}$$

 $// \text{ By LGAP}(p, c, x)$

$$\left\{ \begin{array}{l} \exists s, z. \left[\begin{array}{l} \exists xs, S. x \notin xs \wedge \text{smsg}(h, \text{nil}, S, xs) * (\text{gaps}(S, s) \wedge \text{mygaps}(x, s)) \end{array} \right]_{F(s, h)}^s \\ * \text{Locked}(p) * \text{Locked}(c) * \text{val}(c, x) * \text{isLock}(c, 1) * (c.\text{next} \mapsto z) \\ * \exists \pi > 0. \text{isLock}(h, \pi) * [\text{LCHANGE}(x)]_1^s \end{array} \right\}$$

 $\text{disposelock}(c, 2);$
 $\}$

$$\left\{ \begin{array}{l} \exists s, z. \left[\begin{array}{l} \exists xs, S. x \notin xs \wedge \text{smsg}(h, \text{nil}, S, xs) * (\text{gaps}(S, s) \wedge \text{mygaps}(x, s)) \end{array} \right]_{F(s, h)}^s \\ * \text{Locked}(p) * \exists \pi > 0. \text{isLock}(h, \pi) * [\text{LCHANGE}(x)]_1^s \end{array} \right\}$$

 $\text{unlock}(p);$
 $\{ \text{out}(h, x) \}$

Fig. 9. Proof-sketch for $\text{remove}(x)$.

the shared state to reflect this locked status. Similarly, for the unlock, first the permission is used to extract the knowledge that the node is locked, and then it is locally unlocked.

4 Semantics and Soundness

4.1 Assertion Syntax

Recall from §2.3 that our proof judgments have the form $\Delta, \Gamma \vdash \{P\} C \{Q\}$. Here, P and Q are *assertions* in the set **Assn**. We also define a set of *basic assertions*, **BAssn**, which omit permissions, regions and predicates. Regions in assertions are annotated by *interference assertions* in the set **IAssn**. Δ is an *axiom definition* in the set **Axioms**. Finally, Γ is a *function specification* in the set **Triples**. The syntax is defined as follows.

$$\begin{aligned}
(\text{Assn}) \quad P, Q &::= \text{emp} \mid E_1 \mapsto E_2 \mid P * Q \mid P \multimap Q \mid \text{false} \mid P \Rightarrow Q \mid \exists x. P \mid \\
&\quad [\gamma(E_1, \dots, E_n)]_\pi^R \mid \boxed{P}_I^R \mid \alpha(E_1, \dots, E_n) \mid \otimes x. P \\
(\text{BAssn}) \quad p, q &::= \text{emp} \mid E_1 \mapsto E_2 \mid p * q \mid p \multimap q \mid \text{false} \mid p \Rightarrow q \mid \exists x. p \\
(\text{IAssn}) \quad I &::= \gamma(\vec{x}) : \exists \vec{y}. (P \rightsquigarrow Q) \mid I_1, I_2 \\
(\text{Axioms}) \quad \Delta &::= \emptyset \mid \forall \vec{x}. P \implies Q \mid \forall \vec{x}. \alpha(\vec{x}) \equiv P \mid \Delta_1, \Delta_2 \\
(\text{Triples}) \quad \Gamma &::= \emptyset \mid \Gamma, \{P\} f \{Q\}
\end{aligned}$$

In the above definitions, γ ranges over the set of action names, **AName**; α ranges over the set of abstract predicate names, **PName**; x and y range over the set of logical variables, **Var**; and f ranges over the set of function names, **FName**. We assume an appropriate syntax for expressions, $E, R, \pi \in \text{Expr}$, including basic arithmetic operators.

4.2 Assertion Model

Let $(\text{Heap}, \uplus, \emptyset)$ be any separation algebra [6] representing *machine states* (or *heaps*). Typically, we take **Heap** to be the set of finite partial functions from locations to values, and \uplus to be the union of partial functions with disjoint domains: the standard separation logic model. We let h denote a heap.

Our assertions include permissions which specify the possible interactions with shared regions. Hence, we define **LState**, the set of *logical states*, which pair heaps with permission assignments (elements of **Perm**, defined below).

$$l \in \text{LState} \stackrel{\text{def}}{=} \text{Heap} \times \text{Perm}$$

Assertions make an explicit (logical) division between shared state, which can be accessed by all threads, and thread-local state, which is private to a thread and cannot be subject to interference. Shared state is divided into *regions*. Intuitively,

each region can be seen as the internal state of a single shared structure, i.e. a single lock, set, etc.

Each region is identified by a *region name*, r , from the set \mathbf{RName} . A region is also associated with a syntactic interference assertion, from the set \mathbf{IAssn} , that determines how threads may modify the shared region. A *shared state* in \mathbf{SState} is defined to be a finite partial function mapping region names to logical states and interference assertions:

$$s \in \mathbf{SState} \stackrel{\text{def}}{=} \mathbf{RName} \xrightarrow{\text{fin}} (\mathbf{LState} \times \mathbf{IAssn})$$

A *world* in \mathbf{World} is a pair of a local (logical) state and a shared state, subject to a well-formedness condition. Informally, the well-formedness condition ensures that all parts of the state are disjoint and that each permission corresponds with an appropriate region; we defer the formal definition.

$$w \in \mathbf{World} \stackrel{\text{def}}{=} \{(l, s) \in \mathbf{LState} \times \mathbf{SState} \mid \text{wf}((l, s))\}$$

Given a logical state l , we write l_H and l_P to stand for the heap and permission assignment components respectively. We also write w_L and w_S to stand for the local and shared components of the world w respectively.

Recall from §2.2 that actions can be self-referential. For example, the action **UNLOCK** moves the permission $[\mathbf{UNLOCK}]_1^r$ from shared to local state. Our semantics breaks the loop by distinguishing between the syntactic representation of an action and its semantics. Actions are represented syntactically by *tokens*, consisting of the region name, the action name and a sequence of value parameters (e.g. $[\mathbf{CHANGE}(v)]_1^s$ corresponds to the token (s, \mathbf{CHANGE}, v) with permission 1):

$$t, (r, \gamma, \vec{v}) \in \mathbf{Token} \stackrel{\text{def}}{=} \mathbf{RName} \times \mathbf{AName} \times \mathbf{Val}^*$$

The semantics of a token is defined by an *interference environment* (see §4.4).

Permission assignments in \mathbf{Perm} associate each token with a *permission value* from the interval $[0, 1]$ determining whether the associated action can occur.

$$pr \in \mathbf{Perm} \stackrel{\text{def}}{=} \mathbf{Token} \rightarrow [0, 1]$$

Intuitively, 1 represents full, exclusive permission, 0 represents no permission, and the intermediate values represent partial, non-exclusive permission.⁷

The composition operator \oplus on $[0, 1]$ is defined as addition, with the proviso that the operator is undefined if the two permissions add up to more than 1. Composition on \mathbf{Perm} is the obvious lifting: $pr \oplus pr' \stackrel{\text{def}}{=} \lambda t. pr(t) \oplus pr'(t)$. Composition on logical states is defined by lifting composition for heaps and permission

⁷ This is the fractional permission model of Boyland [3]. As in Dodds *et al.* [9], we can extend our permission structure with a fourth possibility called a *deny*: a non-exclusive permission *prohibiting* both the environment and thread from performing the action. This requires only minor changes to the rely and guarantee functions presented later in this section – other definitions stay the same.

assignments: $l \oplus l' \stackrel{\text{def}}{=} (l_H \uplus l'_H, l_P \oplus l'_P)$. Composition on worlds is defined by composing local states and requiring that shared states be identical:

$$w \oplus w' \stackrel{\text{def}}{=} \begin{cases} (w_L \oplus w'_L, w_S) & \text{if } w_S = w'_S \\ \perp & \text{otherwise.} \end{cases}$$

We write $\mathbf{0}_{\text{Perm}}$ for the empty permission (which assigns zero permission to every token, i.e. $\lambda t.0$), and $[t \mapsto \pi]$ for the permission mapping the token t to π and all other tokens to 0.

The operator $\lfloor w \rfloor$ collapses a world w to a logical state. Here \oplus is the natural lifting of \oplus to sets. The operator $\llbracket w \rrbracket$ further collapses w to a heap:

$$\lfloor w \rfloor \stackrel{\text{def}}{=} w_L \oplus \left(\bigoplus_{r \in \text{dom}(w_S)} w_S(r) \right) \quad \llbracket w \rrbracket \stackrel{\text{def}}{=} (\lfloor w \rfloor)_H$$

The *action domain* of an interference assertion, $\text{adom}(I)$ is the set of action names and parameters that the specification allows.

$$\begin{aligned} \text{adom}(\gamma(x_1, \dots, x_n) : \exists \vec{y}. (P \rightsquigarrow Q)) &\stackrel{\text{def}}{=} \{(\gamma, (v_1, \dots, v_n)) \mid v_i \in \text{Val}\} \\ \text{adom}((I_1, I_2)) &\stackrel{\text{def}}{=} \text{adom}(I_1) \cup \text{adom}(I_2) \end{aligned}$$

We are finally in a position to define well-formedness for a world, $\text{wf}(w)$.

$$\begin{aligned} \text{wf}(w) &\stackrel{\text{def}}{\iff} \exists l \in \text{LState}. l = \lfloor w \rfloor \wedge \\ &\quad \forall r, \gamma, \vec{v}. l_P(r, \gamma, \vec{v}) > 0 \implies \exists l', I. w_S(r) = (l', I) \wedge (\gamma, \vec{v}) \in \text{adom}(I) \end{aligned}$$

4.3 Assertion Semantics

Fig. 10 presents the semantics of assertions, $\llbracket P \rrbracket_{\delta, i}$. We first define a weaker semantics $\langle P \rangle_{\delta, i}$ that does not enforce well-formedness, then define $\llbracket P \rrbracket_{\delta, i}$ by restricting it to the set of well-formed worlds. The semantics of assertions depends on the semantics of expressions, $\llbracket - \rrbracket_- : \text{Expr} \times \text{Interp} \rightarrow \text{Val}$. We have not formally defined this, so we assume an appropriate semantics. The semantics of IAssns can also depend on the semantics of free variables. We define $\langle - \rangle_- : \text{IAssns} \times \text{Interp} \rightarrow \text{IAssns}$ to replace the free variables with their values.

The semantics is parameterized by a *predicate environment*, δ , mapping abstract predicates to their semantic definitions, and an *interpretation*, i , mapping logical variables to values:

$$\begin{aligned} \delta \in \text{PEnv} &\stackrel{\text{def}}{=} \text{PName} \times \text{Val}^* \rightarrow \mathcal{P}(\text{World}) \\ i \in \text{Interp} &\stackrel{\text{def}}{=} \text{Var} \rightarrow \text{Val} \end{aligned}$$

We assume that $\text{RName} \cup (0, 1] \subset \text{Val}$, so that variables may range over region names and fractions.

The cell assertion \mapsto , the separating star $*$ and the existential separating implication $-\otimes$ are standard. The quantifier \bigotimes is infinitary version of $*$; that is,

$$\begin{aligned}
\langle\!\langle - \rangle\!\rangle_{-,-} &: \text{Assn} \times \text{PEnv} \times \text{Interp} \rightarrow \mathcal{P}(\text{LState} \times \text{SState}) \\
\langle\!\langle E_1 \mapsto E_2 \rangle\!\rangle_{\delta,i} &\stackrel{\text{def}}{=} \{(l, s) \mid \text{dom}(l_H) = \{\llbracket E_1 \rrbracket_i\} \wedge l_H(\llbracket E_1 \rrbracket_i) = \llbracket E_2 \rrbracket_i\} \\
\langle\!\langle \text{emp} \rangle\!\rangle_{\delta,i} &\stackrel{\text{def}}{=} \{(\emptyset, \mathbf{0}_{\text{Perm}}), s \mid s \in \text{SState}\} \\
\langle\!\langle P_1 * P_2 \rangle\!\rangle_{\delta,i} &\stackrel{\text{def}}{=} \{w_1 \oplus w_2 \mid w_1 \in \langle\!\langle P_1 \rangle\!\rangle_{\delta,i} \wedge w_2 \in \langle\!\langle P_2 \rangle\!\rangle_{\delta,i}\} \\
\langle\!\langle P_1 -\circ P_2 \rangle\!\rangle_{\delta,i} &\stackrel{\text{def}}{=} \{w \mid \exists w_1, w_2. w_2 = w \oplus w_1 \wedge w_1 \in \langle\!\langle P_1 \rangle\!\rangle_{\delta,i} \wedge w_2 \in \langle\!\langle P_2 \rangle\!\rangle_{\delta,i}\} \\
\langle\!\langle \bigotimes x. P \rangle\!\rangle_{\delta,i} &\stackrel{\text{def}}{=} \bigcup_W \left\{ \bigoplus_v W(v) \mid \forall v. W(v) \in \langle\!\langle P \rangle\!\rangle_{\delta,i[x \mapsto v]} \right\} \\
\langle\!\langle [\gamma(E_1, \dots, E_n)]_\pi^R \rangle\!\rangle_{\delta,i} &\stackrel{\text{def}}{=} \{(\emptyset, [(\gamma, \llbracket R \rrbracket_i, \llbracket E_1 \rrbracket_i, \dots, \llbracket E_n \rrbracket_i) \mapsto \llbracket \pi \rrbracket_i]), s \mid s \in \text{SState}\} \\
\langle\!\langle \boxed{P}_I^R \rangle\!\rangle_{\delta,i} &\stackrel{\text{def}}{=} \{(\emptyset, \mathbf{0}_{\text{Perm}}), s \mid \exists l. (l, s) \in \langle\!\langle P \rangle\!\rangle_{\delta,i} \wedge s(\llbracket R \rrbracket_i) = (l, \langle\!\langle I \rangle\!\rangle_i)\} \\
\langle\!\langle \alpha(E_1, \dots, E_n) \rangle\!\rangle_{\delta,i} &\stackrel{\text{def}}{=} \delta(\alpha, \llbracket E_1 \rrbracket_i, \dots, \llbracket E_n \rrbracket_i) \\
\llbracket - \rrbracket_{-,-} &: \text{Assn} \times \text{PEnv} \times \text{Interp} \rightarrow \mathcal{P}(\text{World}) \\
\llbracket P \rrbracket_{\delta,i} &\stackrel{\text{def}}{=} \{(l, s) \in \langle\!\langle P \rangle\!\rangle_{\delta,i} \mid \text{wf}((l, s))\}
\end{aligned}$$

Fig. 10. Semantics of assertions. The cases for conjunction, implication, existential, etc. are standard, simply distributing over the local and shared state.

the multiplicative analogue of \forall . The empty assertion **emp** asserts that the local state and permission are empty, but that the shared state can contain anything.

Abstract predicates, $\alpha(E_1, \dots, E_n)$, are used to encapsulate concrete properties. For example, in the lock specification (§2.1), we used **Locked**(x) to assert that x is held by the current thread. The meaning of an abstract predicate is simply looked up the in the predicate environment, δ .

The permission assertion $[\gamma(E_1, \dots)]_\pi^R$ states that the token $(\llbracket R \rrbracket_i, \gamma, \llbracket E_1 \rrbracket_i \dots)$ is associated with permission value $\llbracket \pi \rrbracket_i$.

A shared-state assertion \boxed{P}_I^R asserts that P holds for region $\llbracket R \rrbracket_i$ in the shared state, and that the region's interference is given by the interference assertion, $\langle\!\langle I \rangle\!\rangle_i$. For example, in the compare-and-swap lock implementation (§2.2), P asserts that the shared state for a lock is either locked or unlocked, and I defines the meaning of actions **LOCK** and **UNLOCK**. We use $\langle\!\langle I \rangle\!\rangle_i$ as we need to bind the location x and region r to the correct values.

Separating conjunction behaves as conventional (non-separating) conjunction between shared-state assertions over the same region; that is: $\boxed{P}_I^r * \boxed{Q}_I^r \iff \boxed{P \wedge Q}_I^r$. We permit nesting of shared-state assertions. However, nested assertions can always be rewritten in equivalent unnested form:

$$\boxed{\boxed{P}_I^r * Q}_{I'}^{r'} \iff \boxed{P}_I^r * \boxed{Q}_{I'}^{r'}$$

In this paper, we only use nesting to ensure that shared and unshared elements can be referred to by a single abstract predicate. In future, nesting may be useful for defining mutually recursive modules.

4.4 Environment Semantics

An interference assertion defines the actions that are permitted over a region. For example, in the compare-and-swap lock implementation, (§2.2) the assertion $I(r, x)$ defines the action LOCK as $x \mapsto 0 * [\text{UNLOCK}]_1^r \rightsquigarrow x \mapsto 1$.

Semantically, an interference assertion is defined as a map from tokens to sets of shared-state pairs (what we call an *interference environment*):

$$\llbracket - \rrbracket_- : \text{IAssn} \times \text{PEnv} \rightarrow \text{Token} \rightarrow \mathcal{P}(\text{SState} \times \text{SState})$$

A primitive interference assertion defines an interference environment that maps the token (r, γ, \vec{v}) to an action relation corresponding to transitions from states satisfying $\llbracket P \rrbracket_I^r$ to $\llbracket Q \rrbracket_I^r$. The relation does not involve local state, and only the region r of the shared state changes. The action LOCK defines a relation from shared states where the lock region is unlocked, to ones where it is locked. Composition of interference assertions is defined by union of relations.

$$\begin{aligned} \llbracket \gamma(\vec{x}) : \exists \vec{y}. (P \rightsquigarrow Q) \rrbracket_\delta(r, \gamma', \vec{v}) &\stackrel{\text{def}}{=} \left\{ (s, s') \left| \begin{array}{l} \gamma' = \gamma \wedge (\forall r' \neq r. s(r') = s'(r')) \wedge \\ \exists l, l', l_0, I. s(r) = (l \oplus l_0, I) \wedge \\ s'(r) = (l' \oplus l_0, I) \wedge \\ \exists \vec{v}'. (l, s) \in \llbracket P \rrbracket_{\delta, [\vec{x} \mapsto \vec{v}, \vec{y} \mapsto \vec{v}']} \wedge \\ (l', s') \in \llbracket Q \rrbracket_{\delta, [\vec{x} \mapsto \vec{v}, \vec{y} \mapsto \vec{v}']} \end{array} \right. \right\} \\ \llbracket I_1, I_2 \rrbracket_\delta(r, \gamma, \vec{v}) &\stackrel{\text{def}}{=} \llbracket I_1 \rrbracket_\delta(r, \gamma, \vec{v}) \cup \llbracket I_2 \rrbracket_\delta(r, \gamma, \vec{v}) \end{aligned}$$

Given a region name r and an interference assertion I , $\text{all}(I, r)$ is the logical state assigning full permission to all tokens with region r whose semantics is defined by I .

$$\text{all}(I, r) \stackrel{\text{def}}{=} \left(\emptyset, \bigoplus_{(\gamma, \vec{v}) \in \text{adom}(I)} [(r, \gamma, \vec{v}) \mapsto 1] \right)$$

The guarantee G_δ describes what updates to the world the thread is allowed to perform. The thread can update its local state as it pleases, but changes to the shared state must be according to some action for which the thread has sufficient permission ($(w_L)_P(r, \gamma, \vec{v}) > 0$). For example, in the compare-and-swap lock proof (§2.2) the thread must hold permission 1 on UNLOCK before unlocking the shared lock. Without this restriction other threads could potentially unlock the lock. It is also very important that updates preserve the total amount permission in the world ($\lfloor w \rfloor_P = \lfloor w' \rfloor_P$), so that the thread cannot acquire permission to do an action out of thin air.

Moreover, the thread can create a new region by giving away some of its local state and gaining full permission on the newly created region. This is described by G^c . Conversely, it can destroy any region that it fully owns and grab ownership of the state it protects (cf. $(G^c)^{-1}$).

$$\begin{aligned} G^c &\stackrel{\text{def}}{=} \left\{ (w, w') \left| \begin{array}{l} \exists r, I, \ell_1, \ell_2. \ r \notin \text{dom}(w_S) \wedge w'_S = w_S[r \mapsto (\ell_1, I)] \wedge \\ w_L = \ell_1 \oplus \ell_2 \wedge w'_L = \ell_2 \oplus \text{all}(I, r) \end{array} \right. \right\} \\ G_\delta &\stackrel{\text{def}}{=} \left\{ (w, w') \left| \begin{array}{l} \exists r, \gamma, \vec{v}. (w_S, w'_S) \in \llbracket (w_S(r))_2 \rrbracket_\delta(r, \gamma, \vec{v}) \wedge \\ (w_L)_P(r, \gamma, \vec{v}) > 0 \wedge \lfloor w \rfloor_P = \lfloor w' \rfloor_P \end{array} \right. \right\} \cup G^c \cup (G^c)^{-1} \end{aligned}$$

We now define formally the notion of *update implication*, $P \Longrightarrow_{\delta}^{\{p\}\{q\}} Q$. This states that a rewrite from p to q is permitted in a larger world satisfying P , and that the resulting world satisfies Q . As the definition of update implication erases the divisions between regions, it permits repartitioning between regions. Repartitioning is sound provided it is permitted by the guarantee, $(G_{\delta})^*$.

Definition 1 (Update Implication). $P \Longrightarrow_{\delta}^{\{p\}\{q\}} Q$ holds iff, for every world w_1 in $\llbracket P \rrbracket_{\delta,i}$, there exists a heap h_1 in $\llbracket p \rrbracket_i$ and a residual heap h' such that

- $h_1 \oplus h' = \llbracket w_1 \rrbracket$; and
- for every heap h_2 in $\llbracket q \rrbracket_i$, there exists a world w_2 in $\llbracket Q \rrbracket_{\delta,i}$ such that
 - $h_2 \oplus h' = \llbracket w_2 \rrbracket$; and
 - the update is allowed by the guarantee, i.e. $(w_1, w_2) \in (G_{\delta})^*$.

Note that if $p = q = \text{emp}$, then the update implication preserves the concrete state, and only allows the world to be repartitioned. Hence we call this special case *repartitioning implication* and write $P \Longrightarrow_{\delta} Q$ as a shorthand for $P \Longrightarrow_{\delta}^{\{\text{emp}\}\{\text{emp}\}} Q$. Recall from the proof of the compare-and-swap lock implementation (§2.2) that repartitioning implication was used to create a new shared regions when making a lock.

The rely R_{δ} describes the possible world updates that the environment can do. Intuitively, it models interference from other threads. At any point, it can update the shared state by performing one of the actions in any one of the shared regions r , provided that the environment potentially has permission to perform that action. For this to be possible, the world must contain less than the total permission ($\lfloor w \rfloor_P(r, \gamma, \vec{v}) < 1$). This models the fact that some other thread's local state could contain permission $\pi > 0$ on the action.

In addition, the environment can create a new region (cf. R^c) or can destroy an existing region (cf. $(R^c)^{-1}$) provided that no permission for that region exists elsewhere in the world.

$$\begin{aligned}
 R^c &\stackrel{\text{def}}{=} \left\{ (w, w') \mid \begin{array}{l} \exists r, \ell, I. r \notin \text{dom}(w_S) \wedge w'_L = w_L \wedge w'_S = w_S[r \mapsto (\ell, I)] \wedge \\ \lfloor w' \rfloor \text{ defined} \wedge (\forall \gamma, \vec{v}. \lfloor w' \rfloor_P(r, \gamma, \vec{v}) = 0) \end{array} \right\} \\
 R_{\delta} &\stackrel{\text{def}}{=} \left\{ (w, w') \mid \begin{array}{l} \exists r, \gamma, \vec{v}. (w_S, w'_S) \in \llbracket (w_S(r))_2 \rrbracket_{\delta}(r, \gamma, \vec{v}) \wedge \\ w_L = w'_L \wedge \lfloor w \rfloor_P(r, \gamma, \vec{v}) < 1 \end{array} \right\} \cup R^c \cup (R^c)^{-1}
 \end{aligned}$$

These definitions allow us to define stability of assertions. We say that an assertion is *stable* if and only if it cannot be falsified by the interference from other threads that it permits.

Definition 2 (Stability). $\text{stable}_{\delta}(P)$ holds iff for all w, w' , and i , if $w \in \llbracket P \rrbracket_{\delta,i}$ and $(w, w') \in R_{\delta}$, then $w' \in \llbracket P \rrbracket_{\delta,i}$.

Similarly, we say that a predicate environment is stable if and only if all the predicates it defines are stable.

Definition 3 (Predicate Environment Stability). $\text{pstable}(\delta)$ holds iff for all $X \in \text{rng}(\delta)$, for all w and w' , if $w \in X$ and $(w, w') \in R_{\delta}$, then $w' \in X$.

$$\begin{array}{c}
\frac{\vdash_{\text{SL}} \{p\} C \{q\} \quad \Delta \vdash P \Rightarrow^{\{p\}\{q\}} Q}{\Delta; \Gamma \vdash \{P\} \langle C \rangle \{Q\}} \text{ (ATOMIC)} \qquad \frac{\vdash_{\text{SL}} \{p\} C \{q\}}{\Delta; \Gamma \vdash \{p\} C \{q\}} \text{ (PRIM)} \\
\\
\frac{\Delta; \Gamma \vdash \{P_1\} C_1 \{Q_1\} \quad \Delta; \Gamma \vdash \{P_2\} C_2 \{Q_2\}}{\Delta; \Gamma \vdash \{P_1 * P_2\} C_1 \parallel C_2 \{Q_1 * Q_2\}} \text{ (PAR)} \qquad \frac{\{P\} f \{Q\} \in \Gamma}{\Delta; \Gamma \vdash \{P\} f \{Q\}} \text{ (CALL)} \\
\\
\frac{\Delta; \Gamma \vdash \{P\} C \{Q\} \quad \Delta \vdash \text{stable}(R)}{\Delta; \Gamma \vdash \{P * R\} C \{Q * R\}} \text{ (FRAME)} \qquad \frac{\Delta \vdash P \Rightarrow P' \quad \Delta \vdash Q' \Rightarrow Q}{\Delta; \Gamma \vdash \{P\} C \{Q\}} \text{ (CONSEQ)} \\
\\
\frac{\Delta \vdash \Delta' \quad \Delta'; \Gamma \vdash \{P\} C \{Q\}}{\Delta; \Gamma \vdash \{P\} C \{Q\}} \text{ (PRED-I)} \qquad \frac{\Delta \vdash \text{stable}(R) \quad \alpha \notin \Delta, \Gamma, P, Q \quad \Delta, (\forall \vec{x}. \alpha(\vec{x}) \equiv R); \Gamma \vdash \{P\} C \{Q\}}{\Delta; \Gamma \vdash \{P\} C \{Q\}} \text{ (PRED-E)} \\
\\
\frac{\Delta; \Gamma \vdash \{P_1\} C_1 \{Q_1\} \quad \dots \quad \Delta; \Gamma \vdash \{P_n\} C_n \{Q_n\} \quad \Delta; \{P_1\} f_1 \{Q_1\}, \dots, \{P_n\} f_n \{Q_n\}, \Gamma \vdash \{P\} C \{Q\}}{\Delta; \Gamma \vdash \{P\} \text{let } f_1 = C_1 \dots f_n = C_n \text{ in } C \{Q\}} \text{ (LET)}
\end{array}$$

Fig. 11. Selected proof rules.

A syntactic predicate environment, Δ , is defined in the semantics as a set of stable predicate environments:

$$\begin{aligned}
\llbracket \emptyset \rrbracket &\stackrel{\text{def}}{=} \{\delta \mid \text{pstable}(\delta)\} & \llbracket \Delta_1, \Delta_2 \rrbracket &\stackrel{\text{def}}{=} \llbracket \Delta_1 \rrbracket \cap \llbracket \Delta_2 \rrbracket \\
\llbracket \forall \vec{x}. \alpha(\vec{x}) \equiv P \rrbracket &\stackrel{\text{def}}{=} \{\delta \mid \text{pstable}(\delta) \wedge \forall \vec{v}. \delta(\alpha, \vec{v}) = \llbracket P \rrbracket_{\delta, [\vec{x} \mapsto \vec{v}]} \} \\
\llbracket \forall \vec{x}. P \Rightarrow Q \rrbracket &\stackrel{\text{def}}{=} \{\delta \mid \text{pstable}(\delta) \wedge \forall \vec{v}. \llbracket P \rrbracket_{\delta, [\vec{x} \mapsto \vec{v}]} \subseteq \llbracket Q \rrbracket_{\delta, [\vec{x} \mapsto \vec{v}]} \}
\end{aligned}$$

4.5 Programming Language and Proof System

We define a proof system for deriving local Hoare triples for a simple concurrent imperative programming language of commands:

$$\begin{aligned}
\text{(Cmd)} \quad C ::= & \text{skip} \mid c \mid f \mid \langle C \rangle \mid C_1; C_2 \mid C_1 + C_2 \mid C^* \mid C_1 \parallel C_2 \mid \\
& \text{let } f_1 = C_1 \dots f_n = C_n \text{ in } C
\end{aligned}$$

We require that atomic statements $\langle C \rangle$ are not nested and that function names $f_1 \dots f_n$ for a let are pairwise distinct. Here c stands for basic commands, modelled semantically as subsets of $\mathcal{P}(\text{Heap} \times \text{Heap})$ satisfying the locality conditions of Calcagno *et al.* [6].

Judgements about such programs have the form $\Delta; \Gamma \vdash \{P\} C \{Q\}$. This judgement asserts that, beginning in a state satisfying P , interpreted under predicate definitions satisfying Δ , the program C using procedures specified by Γ will not fault and, if it terminates, the final state will satisfy Q .

A selection of the proof rules for our Hoare-style program logic are shown in Fig. 11. These rules are modified from RGSep [25] and deny-guarantee [9]. All

of the rules in our program logic carry an implied assumption that the pre- and post-conditions of their judgements are stable.

The judgement $\vdash_{\text{SL}} \{p\} C \{q\}$ appearing in the ATOMIC and PRIM rules is a judgment in standard sequential separation logic. The other minor judgments in the rules are defined semantically to quantify over all $\delta \in \llbracket \Delta \rrbracket$:

- $\Delta \vdash P \Longrightarrow^{\{p\}\{q\}} Q$ means $\forall \delta \in \llbracket \Delta \rrbracket . P \Longrightarrow_{\delta}^{\{p\}\{q\}} Q$ (and similarly without a superscript);
- $\Delta \vdash \text{stable}(P)$ means $\forall \delta \in \llbracket \Delta \rrbracket . \text{stable}_{\delta}(P)$; and
- $\Delta \vdash \Delta'$ means $\forall \delta \in \llbracket \Delta \rrbracket . \delta \in \llbracket \Delta' \rrbracket$.

To reason about predicate assumptions, we introduce two rules, PRED-I and PRED-E. The PRED-I rule allows the assumptions about the predicate definitions to be weakened. If a triple is provable with assumptions Δ' , then it must be provable under stronger assumptions Δ . The PRED-E rule allows the introduction of predicate definitions. For this to be sound, the predicate name α must not be used anywhere in the existing definitions and assertions. We require that recursively defined predicates are satisfiable; otherwise the premise of a proof rule could be vacuously true. We ensure this holds by checking that only positive occurrences of the predicate appear in its definition.

The ATOMIC and CONSEQ rule were discussed in depth in §2.3. That section also discussed a rule for modules, which can be derived as follows:

$$\frac{\Delta \vdash \{P_1\} C_1 \{Q_1\} \quad \dots \quad \frac{\Delta \vdash \Delta' \quad \Delta'; \{P_1\} f_1 \{Q_1\}, \dots \vdash \{P\} C \{Q\}}{\Delta; \{P_1\} f_1 \{Q_1\}, \dots \vdash \{P\} C \{Q\}} \text{PRED-I}}{\frac{\Delta \vdash \{P\} \text{let } f_1 = C_1 \dots f_n = C_n \text{ in } C \{Q\}}{\vdash \{P\} \text{let } f_1 = C_1 \dots f_n = C_n \text{ in } C \{Q\}} \text{LET}} \text{PRED-E}$$

The PAR rule is the key rule for disjoint concurrency. In this rule we exploit our fiction of disjointness to prove safety for concurrent programs. Our set-up allows us to define a simple parallel rule while capturing fine-grained interactions.

4.6 Judgement Semantics and Soundness

We define the meaning of judgements in our proof system with respect to a transition relation $C, h \xrightarrow{\eta} C', h'$ defining the operational semantics of our language. The transition is parameterised with a *function environment*, η , mapping function names to their definitions. We also define a faulting relation $C, h \xrightarrow{\eta} \text{fault}$.

$$\begin{aligned} \eta &\in \text{FEnv} \stackrel{\text{def}}{=} \text{FName} \rightarrow \text{Cmd} \\ \xrightarrow{\quad} &\in \text{OpTrans} \stackrel{\text{def}}{=} \mathcal{P}(\text{FEnv} \times \text{Cmd} \times \text{Heap} \times \text{Cmd} \times \text{Heap}) \\ \xrightarrow{\quad} \text{fault} &\in \text{OpFault} \stackrel{\text{def}}{=} \mathcal{P}(\text{FEnv} \times \text{Cmd} \times \text{Heap}) \end{aligned}$$

To define the meaning of judgments, we first define the notion of a logical configuration $(C, w, \eta, \delta, i, Q)$ being safe for at least n steps:

Definition 4 (Configuration safety). $C, w, \eta, \delta, i, Q \text{ safe}_0$ always holds; and $C, w, \eta, \delta, i, Q \text{ safe}_{n+1}$ iff the following four conditions hold:

1. For all w' , if $(w, w') \in R_\delta$ then $C, w', \eta, \delta, i, Q \text{ safe}_n$;
2. $\neg((C, \llbracket w \rrbracket) \xrightarrow{\eta} \text{fault})$;
3. For all C' and h' , if $(C, \llbracket w \rrbracket) \xrightarrow{\eta} (C', h')$, then there exists w' such that $h' = \llbracket w' \rrbracket$, $(w, w') \in (G_\delta)^*$, and $C', w', \eta, \delta, i, Q \text{ safe}_n$; and
4. If $C = \mathbf{skip}$, then there exists w' such that $\llbracket w \rrbracket = \llbracket w' \rrbracket$, $(w, w') \in (G_\delta)^*$, and $w' \in \llbracket Q \rrbracket_{\delta, i}$.

This definition says that a configuration is safe provided that: (1) changing the w in a way that respects the rely is still safe; (2) the program cannot fault; (3) if the program can make a step, then there should be an equivalent step in the logical world that is allowed by the guarantee; and (4) if the configuration has terminated, then the post-condition should hold. The use of $(G_\delta)^*$ in the third and fourth conjuncts allows the world to be repartitioned.

We define the meaning of judgements in terms of configuration safety:

Definition 5 (Judgement Semantics). $\Delta; \Gamma \models \{P\} C \{Q\}$ holds iff

$$\forall i, n. \forall \delta \in \llbracket \Delta \rrbracket. \forall \eta \in \llbracket \Gamma \rrbracket_{n, \delta, i}. \forall w \in \llbracket P \rrbracket_{\delta, i}. C, w, \eta, \delta, i, Q \text{ safe}_{n+1},$$

where $\llbracket \Gamma \rrbracket_{n, \delta, i} \stackrel{\text{def}}{=} \{\eta \mid \forall \{P\} f \{Q\} \in \Gamma. \forall w \in \llbracket P \rrbracket_{\delta, i}. \eta(f), w, \eta, \delta, i, Q \text{ safe}_n\}$.

Theorem 1 (Soundness). If $\Delta; \Gamma \vdash \{P\} C \{Q\}$, then $\Delta; \Gamma \models \{P\} C \{Q\}$.

We prove this by structural induction on the proof rules. See Appendix B for full details. The most interesting case is the PAR rule. The simplicity of this rule embodies the compositionality of our logic. To prove it sound, the following lemma is required.

Lemma 1 (Abstract state locality).

If $(C, \llbracket w_1 \oplus w_2 \rrbracket) \xrightarrow{\eta} (C', h)$ and $C, w_1, \eta, \delta, i, Q \text{ safe}_{n+1}$, then $\exists w'_1, w'_2$ such that $(C, \llbracket w_1 \rrbracket) \xrightarrow{\eta} (C', \llbracket w'_1 \rrbracket)$ and $h = \llbracket w'_1 \oplus w'_2 \rrbracket$ and $(w_1, w'_1) \in G_\delta$ and $(w_2, w'_2) \in R_\delta$.

Proof. To prove this, we require that basic commands obey a concrete locality assumption. We must also prove that rely and guarantee obey similar locality lemmas. The lemma then follows from the definition of configuration safety. The full proof is given in Appendix B, Lemma 9. \square

The lemma shows that a program execution only affects the parts of the world required for it to run safely. The soundness of PAR follows because the lemma shows that threads are safely contained within their abstract footprints.

5 Conclusions and Related Work

We have presented a logic that provides abstraction in a concurrent setting. It brings together three streams of research: (1) abstract predicates [23], using predicates to abstract the internal details of a module or class; (2) deny-guarantee [9], reasoning about concurrent programs using rely-guarantee and separation logic; and (3) abstract local reasoning [8], formalising the fiction of disjointness obtained when reasoning at the abstract level.

Our work on concurrent abstract predicates has been strongly influenced by O’Hearn’s concurrent separation logic (CSL) [21]. CSL takes statically allocated locks as a primitive. With CSL, we can reason about programs with these locks as if they are disjoint from each other, even though they interfere on a shared state. CSL therefore provides a key example of the fiction of disjointness. The CSL approach has been extended to deal with dynamically-allocated locks [13, 17] and re-entrant locks [14], by providing new proofs rules and assertions. This work uses invariants to abstract the state over which interference acts. Parkinson *et al.* [22] have shown how the CSL approach can be used to reason about concurrent algorithms that do not use locks. However, representing the complex interactions of concurrent algorithms in a single invariant can require auxiliary state making proofs less tractable.

Jacobs and Piessens [19] have developed an approach to reasoning abstractly based on CSL. They use a logic similar to Gotsman [13] to reason about dynamically allocated locks with associated invariants. Their logic uses auxiliary state to express the temporal nature of interference. To deal with auxiliary state in a modular way they add a special implication that allows the auxiliary state to be changed in any way that satisfies the invariant. This implication is remarkably similar to our repartitioning implication \Rightarrow . Their implication can be used in the specification of a module, allowing the auxiliary state of the client to be updated during the module’s execution. We believe that this technique could be soundly applied with our repartitioning implication, which would simplify the use of the lock specification in the two set algorithms.

An alternative approach to using invariants is to abstract interference over the shared state by relations modelling the interaction of different threads: the rely-guarantee method [20]. There have been two recent logics that combine RG with separation logic: RGSep [25] and SAGL [11]. Both logics allow more elegant proofs of concurrent algorithms than the invariant-based proofs, but they have the serious drawback that interference on the shared state cannot be abstracted. Feng’s Local Rely-Guarantee [10] improves the locality of RGSep and SAGL, by scoping interference with a precise footprint and having special rules for framing and hiding interference, but this approach still has no support for abstraction. Meanwhile, Dinsdale-Young, Gardner and Wheelhouse used RGSep to analyse a concurrent B-tree algorithm [7], and conjectured ideas about abstraction which led to this current collaboration.

We have built on deny-guarantee reasoning [9], a reformulation of rely-guarantee introduced to reason about dynamically scoped concurrency. This paper has combined a simplified form of deny-guarantee with the earlier work on RGSep.

Deny-guarantee reasoning is related to the ‘state guarantees’ of Bierhoff et al. [1] in linear logic, which are also splittable permissions describing how a state can be updated (e.g., read-only, read-write,...).

The concept of the fiction of disjointness presented here arose from Gardner *et al.*’s work on abstract local reasoning about program modules using context logic [4, 5, 12]. Recently, Dinsdale-Young, Gardner and Wheelhouse have shown how to justify such abstract local reasoning about modules by relating it to separation logic reasoning about their non-local implementations [8]. Reasoning abstractly about modules thus provides a fiction of separation, which is not necessarily mirrored in their implementations.

We should point out that proofs in our proof system seem to be slightly more complex in practice than RGSep and SAGL, as can be seen by comparing the lock-coupling list proof with the RGSep one [25]. This might just be the penalty that we pay for having greater modularity although, as we acquire more experience with doing proofs using concurrent abstract predicates, we expect to be able to reduce this complexity.

An alternative approach to abstracting concurrent algorithms is to use linearisability [16]. Linearisability provides a fiction of atomicity allowing “reason[ing] about properties of concurrent objects given just their (sequential) specifications” [16]. With linearisability, we are forced to step outside the program logic at module boundaries and fall back on operational reasoning. In contrast, with concurrent abstract predicates we are able to write modular proofs within the one logical formalism.

Acknowledgements Thanks to Richard Bornat, Alexey Gotsman, Suresh Jaganathan, Mohammad Raza, John Wickerson and Mark Wheelhouse for useful discussions and feedback. This work was supported by an EPSRC doctoral training account (Dinsdale-Young), EPSRC grant EP/F019394/1 (Dodds and Parkinson), an RAEng/EPSRC research fellowship (Parkinson) and a Microsoft Research Cambridge/RAEng senior research fellowship (Gardner).

References

1. K. Bierhoff and J. Aldrich. Modular typestate checking of aliased objects. In *OOPSLA*, pages 301–320, 2007.
2. R. Bornat, C. Calcagno, P. O’Hearn, and M. Parkinson. Permission accounting in separation logic. In *POPL*, pages 259–270, 2005.
3. J. Boyland. Checking interference with fractional permissions. In *SAS*, 2003.
4. C. Calcagno, P. Gardner, and U. Zarfaty. Context logic and tree update. In *POPL*, 2005.
5. C. Calcagno, P. Gardner, and U. Zarfaty. Local reasoning about data update. *Festschrift Computation, Meaning and Logic: Articles dedicated to Gordon Plotkin*, 172, 2007.
6. C. Calcagno, P. W. O’Hearn, and H. Yang. Local action and abstract separation logic. In *Symp. on Logic in Comp. Sci. (LICS’07)*, pages 366–378, 2007.

7. T. Dinsdale-Young, P. Gardner, and M. Wheelhouse. Local reasoning about a concurrent B*-list algorithm. Talk and technical report, see <http://www.doc.ic.ac.uk/~td202/>, 2009.
8. T. Dinsdale-Young, P. Gardner, and M. Wheelhouse. Abstract local reasoning. Imperial technical report, 2010.
9. M. Dodds, X. Feng, M. Parkinson, and V. Vafeiadis. Deny-guarantee reasoning. In *ESOP*, 2009.
10. X. Feng. Local rely-guarantee reasoning. In *POPL*, 2009.
11. X. Feng, R. Ferreira, and Z. Shao. On the relationship between concurrent separation logic and assume-guarantee reasoning. In *ESOP*, 2007.
12. P. Gardner, G. Smith, M. Wheelhouse, and U. Zarfaty. Local reasoning about dom. In *PODS*, 2008.
13. A. Gotsman, J. Berdine, B. Cook, N. Rinetzky, and M. Sagiv. Local reasoning for storable locks and threads. In *APLAS*, 2007.
14. C. Haack, M. Huisman, and C. Hurlin. Reasoning about Java’s Reentrant Locks. In *APLAS*, pages 171–187, 2008.
15. M. Herlihy and N. Shavit. *The Art of Multiprocessor Programming*. Elsevier, 2008.
16. M. P. Herlihy and J. M. Wing. Linearizability: a correctness condition for concurrent objects. *TOPLAS*, 12(3):463–492, 1990.
17. A. Hobor, A. W. Appel, and F. Z. Nardelli. Oracle semantics for concurrent separation logic. In *ESOP*, 2008.
18. S. S. Ishtiaq and P. W. O’Hearn. BI as an assertion language for mutable data structures. In *POPL*, pages 14–26, Jan. 2001.
19. B. Jacobs and F. Piessens. Modular full functional specification and verification of lock-free data structures. Technical Report CW 551, Katholieke Universiteit Leuven, Department of Computer Science, June 2009.
20. C. B. Jones. Annotated bibliography on rely/guarantee conditions. <http://homepages.cs.ncl.ac.uk/cliff.jones/ftp-stuff/rg-hist.pdf>, 2007.
21. P. W. O’Hearn. Resources, concurrency and local reasoning. *TCS*, 2007.
22. M. Parkinson, R. Bornat, and P. O’Hearn. Modular verification of a non-blocking stack. In *POPL*, pages 297–302, Jan. 2007.
23. M. J. Parkinson and G. M. Bierman. Separation logic and abstraction. In *POPL*, pages 247–258, 2005.
24. V. Vafeiadis. *Modular Fine-Grained Concurrency Verification*. PhD thesis, University of Cambridge, July 2007.
25. V. Vafeiadis and M. Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR*, pages 256–271, 2007.

A Language Semantics

The operational semantics of our programming language is given in Fig. 12. This semantics is largely taken from [24]. However, this prior semantics required two operational semantics – a logical semantics with proof annotations to separate and a machine semantics erasing such annotations. Using Vafeiadis’s new approach to soundness of RGSep, Definition 4, we can avoid this requirement for two semantics.

$$\begin{array}{c}
\frac{(C, h) \xrightarrow{\eta[f_1 \mapsto C_1 \dots f_n \mapsto C_n]} (C', h')}{(\text{let } f_1 = C_1 \dots f_n = C_n \text{ in } C, h) \xrightarrow{\eta} (\text{let } f_1 = C_1 \dots f_n = C_n \text{ in } C', h')} \\
\\
\frac{}{(\text{let } \dots \text{ in } \mathbf{skip}, h) \xrightarrow{\eta} (\mathbf{skip}, h)} \quad \frac{f \in \text{dom}(\eta)}{(f, h) \xrightarrow{\eta} (\eta(f), h)} \quad \frac{c(h, h')}{(c, h) \xrightarrow{\eta} (\mathbf{skip}, h')} \\
\\
\frac{(C, h) \xrightarrow{\eta} (C_1, h')}{(C; C', h) \xrightarrow{\eta} (C_1; C', h')} \quad \frac{}{(\mathbf{skip}; C, h) \xrightarrow{\eta} (C, h)} \quad \frac{}{(C^*, h) \xrightarrow{\eta} (\mathbf{skip} + (C; C^*), h)} \\
\\
\frac{}{(C_1 + C_2, h) \xrightarrow{\eta} (C_1, h)} \quad \frac{}{(C_1 + C_2, h) \xrightarrow{\eta} (C_2, h)} \quad \frac{(C, h) \xrightarrow{\eta^*} (\mathbf{skip}, h')}{(\langle C \rangle, h) \xrightarrow{\eta} (\mathbf{skip}, h')} \\
\\
\frac{(C_1, h) \xrightarrow{\eta} (C'_1, h')}{(C_1 \| C_2, h) \xrightarrow{\eta} (C'_1 \| C_2, h')} \quad \frac{(C_2, h) \xrightarrow{\eta} (C'_2, h')}{(C_1 \| C_2, h) \xrightarrow{\eta} (C_1 \| C'_2, h')} \quad \frac{}{(\mathbf{skip} \| \mathbf{skip}, h) \xrightarrow{\eta} (\mathbf{skip}, h)} \\
\\
\frac{(C_1, h) \xrightarrow{\eta} \text{fault}}{(C_1; C_2, h) \xrightarrow{\eta} \text{fault}} \quad \frac{(C_1, h) \xrightarrow{\eta} \text{fault}}{(C_1 \| C_2, h) \xrightarrow{\eta} \text{fault}} \quad \frac{(C_2, h) \xrightarrow{\eta} \text{fault}}{(C_1 \| C_2, h) \xrightarrow{\eta} \text{fault}} \\
\\
\frac{f \notin \text{dom}(\eta)}{(f, h) \xrightarrow{\eta} \text{fault}} \quad \frac{(\neg \exists h'. c(h, h'))}{(c, h) \xrightarrow{\eta} \text{fault}} \quad \frac{(C, h) \xrightarrow{\eta^*} \text{fault}}{(\langle C \rangle, h) \xrightarrow{\eta} \text{fault}}
\end{array}$$

Fig. 12. Operational semantics.

We assume a simple Dijkstra-style language with constructs for sequential composition, nondeterministic choice and looping. Programs in the given in paper can be encoded into this language for proof purposes. Primitive updates in this language are modelled by a function $c: \mathbf{States} \times \mathbf{States}$.

Theorem 2 (Well-definedness). *For any program C , state h and environment η , there exists a transition $(C, h) \xrightarrow{\eta} (C', h')$ or $(C, h) \xrightarrow{\eta} \text{fault}$.*

Proof. Holds trivially by induction for most cases. The only interesting case is the rule dealing with atomic commands. For this, we observe that a derivation sequence beginning with C cannot result in the program $\langle C \rangle$, so the derivation must be well-founded. \square

B Soundness of proof system

Lemma 2 (Skip safety). *If $\text{stable}_\delta(Q)$ and $w \models_\delta Q$ then $\mathbf{skip}, w, \eta, \mathbb{I}, \delta, Q$ safe $_n$.*

Proof. Prove by induction on n . The 0 case holds by definition. Consider now the inductive case. By the definition of stability, any w' such that $(w, w') \in R_\delta$ must

satisfy $w' \models Q$. Consequently the first clause holds by the inductive assumption. Clauses 2 and 3 hold trivially, as no reductions from **skip** exist in the semantics. Clause 4 holds trivially by picking w' as w . \square

Assumption 1 (Primitive locality) *We assume that elements of the primitive update function c : $\text{States} \times \text{States}$ are local. That is:*

1. *If no h_2 exists such that $(h_1 \oplus h, h_2) \in c$, then no h'_2 exists such that $(h_1, h'_2) \in c$, and*
2. *if $(h_1 \oplus h, h_2) \in c$ and there exists an h_3 such that $(h_1, h_3) \in c$ then there exists an h_4 such that $(h_1, h_4) \in c$ and $h_2 = h \oplus h_4$.*

Lemma 3 (Concrete locality). *We need both single-step and multi-step versions of locality.*

- *If $(C, h \oplus h') \xrightarrow{\eta} (C_2, h_2)$ and $\neg((C, h) \xrightarrow{\eta} \text{fault})$, then $\exists h'_2$ such that $(C, h) \xrightarrow{\eta} (C_2, h'_2)$ and $h'_2 \oplus h' = h_2$.*
- *If $(C, h \oplus h') \xrightarrow{\eta^*} (C_2, h_2)$ and $\neg((C, h) \xrightarrow{\eta^*} \text{fault})$, then $\exists h'_2$ such that $(C, h) \xrightarrow{\eta^*} (C_2, h'_2)$ and $h'_2 \oplus h' = h_2$.*

Proof. Proved by induction over the structure of a derivation or derivation sequence.

Consider the first case, consisting of a single derivation $(C, h \oplus h') \xrightarrow{\eta} (C_2, h_2)$. For most of the rules, locality holds simply by the first clause of the inductive assumption. For the two rules dealing with primitive updates, locality is ensured by Assumption 1. For the rule dealing with atomic statements, locality is ensured by the second clause of the inductive assumption.

Now consider the second case, where we have a multi-step derivation $(C, h \oplus h') \xrightarrow{\eta^*} (C_2, h_2)$. Any such derivation consists of k single-step derivations. We can thus prove the result by appeal to the two inductive assumptions. \square

Lemma 4 (Fault locality). *If $(C, h \oplus h') \xrightarrow{\eta} \text{fault}$, then $(C, h) \xrightarrow{\eta} \text{fault}$.*

Proof. As in the previous lemma, the proof needs both single-step and multi-step versions of locality. The result is proved by induction over the structure of a faulting derivation or derivation sequence.

- *If $(C, h \oplus h') \xrightarrow{\eta} \text{fault}$, then $(C, h) \xrightarrow{\eta} \text{fault}$.*
- *If $(C, h \oplus h') \xrightarrow{\eta^*} \text{fault}$, then $(C, h) \xrightarrow{\eta^*} \text{fault}$.*

The first clause holds when we have a single-step faulting derivation, $(C, h \oplus h') \xrightarrow{\eta} \text{fault}$. For most of the rules, locality holds immediately by the first clause of the inductive assumption. For the rule dealing with primitive update, locality is ensured by Assumption 1, as in the previous lemma. For the rule dealing with atomic statements, locality holds by the second clause of the inductive assumption.

To prove the second clause, we observe that any faulting derivation sequence $(C, h \oplus h') \xrightarrow{\eta^*} \text{fault}$ must consist of a length- k derivation $(C, h \oplus h') \xrightarrow{\eta^k} (C', h'')$

and a single derivation $(C', h'') \xrightarrow{\eta} \text{fault}$. We can thus show the result by applying Lemma 3 and the first inductive assumption. \square

Lemma 5 (Guarantee locality).

If $(w_1, w'_1) \in G_\delta$ and $w_1 \oplus w_2$ is defined, then there exists w'_2 such that $(w_1 \oplus w_2, w'_1 \oplus w'_2) \in G_\delta$ and $w'_1 \oplus w'_2$ is defined.

Proof. By the definition of state composition, $(w_1)_S = (w_1 \oplus w_2)_S$ and $(w'_1)_S = (w'_1 \oplus w'_2)_S$. Also, by permission composition, if $((w_1)_L)_P(r, \gamma, \vec{v}) \in (0, 1]$ then $((w_1 \oplus w_2)_L)_P(r, \gamma, \vec{v}) \in (0, 1]$. For region creation, it follows that the new region cannot be in the domain of $(w_1)_S$. For the region removal, $((w_1)_L)_P$ must contain all the permissions on the region being removed, hence so will $((w_1 \oplus w_2)_L)_P$. The rest follows immediately by the definition of guarantee. \square

Lemma 6 (Guarantee locality II). *If $(w_1, w'_1) \in G_\delta$ and $(w_2, w'_2) \in R_\delta$ and $w_1 \oplus w_2$ defined and $w'_1 \oplus w'_2$ defined, then $(w_1 \oplus w_2, w'_1 \oplus w'_2) \in G_\delta$.*

Proof. Similar to previous. \square

Lemma 7 (Rely locality).

If $w = w_1 \oplus w_2$ and $(w, w') \in R_\delta$ then $\exists w'_1, w'_2$ such that $(w_1, w'_1) \in R_\delta$ and $(w_2, w'_2) \in R_\delta$ and $w' = w'_1 \oplus w'_2$.

Proof. By the definition of state composition, $(w_1)_S = (w_2)_S = (w_1 \oplus w_2)_S$, and $(w'_1)_S = (w'_2)_S = (w'_1 \oplus w'_2)_S$. Also, by permission composition, if $((w_1 \oplus w_2)_L)_P(r, \gamma, \vec{v}) \in [0, 1]$ then $((w_1)_L)_P(r, \gamma, \vec{v}) \in [0, 1]$ and $((w_2)_L)_P(r, \gamma, \vec{v}) \in [0, 1]$. If w allows a region to be created or removed, then so will both w_1 and w_2 . The rest follows immediately by the definition of rely. \square

Lemma 8 (Containment of rely in guarantee).

If $w_1 \oplus w_2$ is defined and $(w_1, w'_1) \in G_\delta$ then there exists a w'_2 such that $(w_2, w'_2) \in R_\delta$ and $w'_1 \oplus w'_2$ is defined.

Proof. By the definition of state composition $(w_1)_S = (w_1 \oplus w_2)_S$. By the definition of permission composition, if $((w_1)_L)_P(r, \gamma, \vec{v}) \in (0, 1]$ and $w_1 \oplus w_2$ is defined, then $((w_2)_L)_P(r, \gamma, \vec{v}) \in [0, 1]$. We define w'_2 as $w'_2 = ((w_2)_L, (w'_1)_S)$. For region creation this is trivial. For region removal, we know $((w_1)_L)_P$ must contain all the permissions for the removed region, therefore $((w_2)_L)_P$ cannot contain any permissions on the region, hence it can be removed. The rest holds immediately by the definition of rely. \square

Lemma 9 (Abstract state locality).

If $(C, \llbracket w_1 \oplus w_2 \rrbracket) \xrightarrow{\eta} (C', h)$ and $C, w_1, \eta, \delta, i, Q$ safe _{$n+1$} , then $\exists w'_1, w'_2$ such that $(C, \llbracket w_1 \rrbracket) \xrightarrow{\eta} (C', \llbracket w'_1 \rrbracket)$ and $h = \llbracket w'_1 \oplus w'_2 \rrbracket$ and $(w_1, w'_1) \in G_\delta$ and $(w_2, w'_2) \in R_\delta$.

Proof. We first construct a remainder state h' from the local portion of w_2 , so $h' = \lfloor (w_2)_L \rfloor$. By the definition of composition $\llbracket w_1 \oplus w_2 \rrbracket = \llbracket w_1 \rrbracket \oplus h'$.

By appeal to Lemma 3 (Concrete single-step locality) there exists a h'' such that $C, \llbracket w_1 \rrbracket \xrightarrow{\eta} C', h''$ and $h' \oplus h'' = h$. By the definition of safety there must exist a w'_1 such that $\llbracket w_1 \rrbracket = h''$ and $(w_1, w'_1) \in G_\delta$.

We now define w'_2 as $((w_2)_L, (w'_1)_S)$ – that is, the local state from w_2 and shared state from w'_1 . It therefore holds trivially that $\llbracket w'_1 \rrbracket \oplus h' = \llbracket w'_1 \oplus w'_2 \rrbracket$, proving the main requirement. It remains to prove that $(w_2, w'_2) \in R_\delta$. This follows immediately by appeal to Lemma 8. \square

Lemma 10 (Parallel safety decomposition).

*If $w = w_1 \oplus w_2$, and $C_1, w_1, \eta, \delta, i, Q_1 \text{ safe}_n$ and $C_2, w_2, \eta, \delta, i, Q_2 \text{ safe}_n$, and $\text{stable}_\delta(Q_1)$ and $\text{stable}_\delta(Q_2)$ then $C_1 \parallel C_2, w, \eta, \delta, i, Q_1 * Q_2 \text{ safe}_n$.*

Proof. By induction. The zero case holds trivially. In the inductive case we must prove $C_1 \parallel C_2, w, \eta, \delta, i, Q_1 * Q_2 \text{ safe}_{n+1}$. We break down the definition of safety as follows (the fourth clause is satisfied trivially as the program is not **skip**):

1. $(w, w') \in R_\delta \implies C_1 \parallel C_2, w', \eta, \delta, i, Q_1 * Q_2 \text{ safe}_n$
2. $\neg((C_1 \parallel C_2, \llbracket w \rrbracket) \xrightarrow{\eta} \text{fault})$
3. $(C_1 \parallel C_2, \llbracket w \rrbracket) \xrightarrow{\eta} (C', h') \implies \exists w'. \llbracket w' \rrbracket = h' \wedge (w, w') \in (G_\delta)^* \wedge C', w', \eta, \delta, i, Q_1 * Q_2 \text{ safe}_n$

For the first clause, we assume that $(w, w') \in R_\delta$. By Lemma 7 (Rely locality) $\exists w'_1, w'_2$ such that $(w_1, w'_1) \in R_\delta$ and $(w_2, w'_2) \in R_\delta$ and $w' = w'_1 \oplus w'_2$. Hence, with the assumptions we get $C_2, w'_2, \eta, \delta, i, Q_2 \text{ safe}_n$ and $C_1, w'_1, \eta, \delta, i, Q_1 \text{ safe}_n$. Hence, by the inductive assumption this proves the clause.

For the second clause, we observe that if the thread faults in a parallel composition, then either C_1 or C_2 must have faulted. By Lemma 4 (Concrete fault locality) the threads must also fault in $\llbracket w_1 \rrbracket$ or $\llbracket w_2 \rrbracket$ respectively. But by the safety assumptions for C_1 and C_2 this cannot occur.

For the third clause, there are two cases: either $C_1 \parallel C_2 = \text{skip} \parallel \text{skip}$, or one of C_1 or C_2 is reducible. We first consider the **skip** case. We pick the value of w' as w . By the semantics, $C' = \text{skip}$ and $h' = \llbracket w \rrbracket$. The first and second requirements of this clause hold trivially, and the third holds by Lemma 2 (Skip safety) and the assumptions of safety.

Now consider the case where C_1 or C_2 is reducible. We only consider C_1 ; the C_2 case is identical. It must hold that $(C_1, \llbracket w \rrbracket) \xrightarrow{\eta} (C'_1, h')$ and $C' = C'_1 \parallel C_2$. By appeal to Lemma 9 (Abstract locality) we have that $(C_1, \llbracket w_1 \rrbracket) \xrightarrow{\eta} (C'_1, \llbracket w'_1 \rrbracket)$ and $(w_1, w'_1) \in G_\delta$ and $(w_2, w'_2) \in R_\delta$ and $\llbracket w'_1 \oplus w'_2 \rrbracket = h'$. Consequently by the safety assumptions we have $C'_1, w'_1, \eta, \delta, i, Q_1 \text{ safe}_n$. and $C_2, w'_2, \eta, \delta, i, Q_2 \text{ safe}_n$. Hence by inductive hypothesis, we have $C'_1 \parallel C_2, w'_1 \oplus w'_2, \eta, \delta, i, Q_1 * Q_2 \text{ safe}_n$. By Lemma 6 (Guarantee locality II), we get $(w_1 \oplus w_2, w'_1 \oplus w'_2) \in G_\delta$ \square

Lemma 11. *If $C, w_1, \eta, \delta, i, Q \text{ safe}_n$ and $w = w_1 \oplus w_2$ and $w_2, i \models F$ and $\text{stable}_\delta(F)$, then $C, w, \eta, \delta, i, Q * F \text{ safe}_n$.*

Proof. By induction on n . The zero case is trivial. In the inductive case we must prove $C, w, \eta, \delta, i, Q * F \text{ safe}_{n+1}$, which we break down as follows

1. $(w, w') \in R_\delta \implies C, w', \eta, \delta, i, Q * F \text{ safe}_n$
2. $\neg((C, \llbracket w \rrbracket) \xrightarrow{\eta} \text{fault})$
3. $(C, \llbracket w \rrbracket) \xrightarrow{\eta} (C', h') \implies \exists w'. \llbracket w' \rrbracket = h' \wedge (w, w') \in (G_\delta)^* \wedge C', w', \eta, \delta, i, Q * F \text{ safe}_n$
4. $C = \mathbf{skip} \implies \exists w'. (w, w') \in (G_\delta)^* \wedge \llbracket w \rrbracket = \llbracket w' \rrbracket \wedge w', i \models_\delta Q * F$

For the first clause, we can assume $(w, w') \in R_\delta$ which with Lemma 7 (Rely locality) gives $(w_1, w'_1) \in R_\delta$ and $(w_2, w'_2) \in R_\delta$ and $w' = w'_1 \oplus w'_2$. As F is stable, we know $w'_2, i \models F$. By the inductive assumption we thus have $C, w'_1, \eta, \delta, i, Q * F \text{ safe}_n$, which completes the case.

The second clause holds by the contrapositive of Lemma 4 (Concrete fault locality).

For the third clause, we assume a transition $C, \llbracket w \rrbracket \xrightarrow{\eta} C', h'$. By appeal to Lemma 9 (Abstract state locality) there exists a transition $C, \llbracket w_1 \rrbracket \xrightarrow{\eta} C', h''$ and there exist worlds w'_1, w'_2 such that $\llbracket w'_1 \oplus w'_2 \rrbracket = h'$. This satisfies the first requirement. It also holds that $(w_1, w'_1) \in G_\delta$, so the second requirement follows by Lemma 5 (Guarantee locality). The third requirement follows immediately from the inductive assumption.

For the fourth clause assume $C = \mathbf{skip}$. By assumption we know $(w_1, w'_1) \in (G_\delta)^*$ and $\llbracket w_1 \rrbracket = \llbracket w'_1 \rrbracket$ and $w'_1, i \models_\delta Q$. By Lemma 5 (Guarantee locality) we know $(w_1 \oplus w_2, w'_1 \oplus w'_2) \in (G_\delta)^*$ and $(w_2, w'_2) \in (R_\delta)^*$. As F is stable, we know $w'_2, i \models F$, and hence $w'_1 \oplus w'_2, i \models Q * F$ as required. \square

Lemma 12 (Frame safety). *If $\delta, \eta \models_n \{P\} C \{Q\}$ and $\text{stable}_\delta(F)$, then $\delta, \eta \models_n \{P * F\} C \{Q * F\}$.*

Proof. We pick a w, i such that $w, i \models P * F$. Note that by the definition of $*$, $\exists w_1, w_2$ such that $w_1, i \models P$, $w_2, i \models F$, and $w = w_1 \oplus w_2$. Thus the result follows directly from Lemma 11. \square

Lemma 13.

If $\delta, \eta \models_n \{P'\} C \{Q\}$, and if $(w, w') \in (R_\delta)^$ then there exists w'' such that $(w', w'') \in G_\delta$ and $\llbracket w' \rrbracket = \llbracket w'' \rrbracket$ and $w'' \in \llbracket P' \rrbracket_{i, \delta}$, then $C, w, \eta, \delta, i, Q \text{ safe}_n$.*

Proof. Induction on n . The zero case is trivial. For the inductive case we must prove $C, w, \eta, \delta, i, Q \text{ safe}_{n+1}$, which we break down to

1. $(w, w') \in R_\delta \implies C, w', \eta, \delta, i, Q \text{ safe}_n$
2. $\neg((C, \llbracket w \rrbracket) \xrightarrow{\eta} \text{fault})$
3. $(C, \llbracket w \rrbracket) \xrightarrow{\eta} (C', h') \implies \exists w'. \llbracket w' \rrbracket = h' \wedge (w, w') \in (G_\delta)^* \wedge C', w', \eta, \delta, i, Q \text{ safe}_n$
4. $C = \mathbf{skip} \implies \exists w'. (w, w') \in (G_\delta)^* \wedge \llbracket w \rrbracket = \llbracket w' \rrbracket \wedge w', i \models_\delta Q$

First clause follows directly from induction. Second and third clauses hold trivially as $\llbracket _ \rrbracket$ is preserved by guarantee steps. Fourth clause is trivial. \square

Lemma 14 (Pre-partitioning safety). *If $P \Longrightarrow_{\delta} P'$ and $\delta, \eta \models_n \{P'\}C\{Q\}$ and $\text{stable}_{\delta}(P)$, then $\delta, \eta \models_n \{P\}C\{Q\}$.*

Proof. Follows directly from Lemma 13, and the definitions of \Longrightarrow and **stable**. \square

Lemma 15. *If $C, w, \eta, \delta, i, Q'$ safe_n and $Q' \Longrightarrow_{\delta} Q$, then C, w, η, δ, i, Q safe_n .*

Proof. By induction on n . First three clauses of the definition follow trivially. For the fourth, we can assume

$$\exists w'. (w, w') \in (G_{\delta})^* \wedge \llbracket w \rrbracket = \llbracket w' \rrbracket \wedge w', i \models_{\delta} Q'$$

We must prove

$$\exists w'. (w, w') \in (G_{\delta})^* \wedge \llbracket w \rrbracket = \llbracket w' \rrbracket \wedge w', i \models_{\delta} Q$$

This follows by definition of \Longrightarrow and transitivity. \square

Lemma 16 (Post-partitioning safety). *If $Q' \Longrightarrow_{\delta} Q$ and $\delta, \eta \models_n \{P\}C\{Q'\}$, then $\delta, \eta \models_n \{P\}C\{Q\}$.*

Proof. Follows directly from previous lemma. \square

Lemma 17 (Atomic safety). *If $\vdash_{\text{SL}} \{p\}C\{q\}$ and $P \Longrightarrow_{\delta}^{\{p\}\{q\}} Q$, and $\text{stable}_{\delta}(P)$ and $\text{stable}_{\delta}(Q)$, then $\delta; \eta \models_n \{P\} \langle C \rangle \{Q\}$*

Proof. By induction on n . The zero case is trivial. For the inductive case we assume a w such that $w, i \models_{\delta} P$. We now need to prove that $\langle C \rangle, w, \eta, \delta, i, Q$ safe_{n+1} , which we break down as follows (the fourth clause is trivial as the command is not **skip**)

1. $(w, w') \in R_{\delta} \Longrightarrow \langle C \rangle, w', \eta, \delta, i, Q$ safe_n
2. $\neg(\langle C \rangle, \llbracket w \rrbracket \xrightarrow{\eta} \text{fault})$
3. $\langle C \rangle, \llbracket w \rrbracket \xrightarrow{\eta} C', h' \Longrightarrow \exists w'. \llbracket w' \rrbracket = h' \wedge (w, w') \in (G_{\delta})^* \wedge C', w', \eta, \mathbb{I}, \delta, Q$ safe_n

The first clause holds by the stability of P and the inductive assumption.

By the semantics of $\vdash_{\text{SL}} \{p\}C\{q\}$, for any h such that $h \models p$, it holds that $\neg((C, h) \xrightarrow{\eta}^* \text{fault})$. By the definition of update implication, there must exist states h', h'' such that $\llbracket w \rrbracket = h' \oplus h''$ and $h' \models p$. The second clause therefore follows by the contrapositive of Lemma 4 (Concrete fault locality).

For clause 3, we assume that there exists a transition $\langle C \rangle, \llbracket w \rrbracket \xrightarrow{\eta} C', h_3$. By reduction rules, $C' = \text{skip}$. By Lemma 4 (Concrete fault locality) there exists a transition $\langle C \rangle, h' \xrightarrow{\eta} \text{skip}, h_4$ such that $h_4 \oplus h'' = h_3$. By the semantics of $\vdash_{\text{SL}} \{p\}C\{q\}$, it must hold that $h_4 \models q$. By the definition of update implication, there exists a w_2 such that $w_2 \models Q$, $(w, w_2) \in G_{\delta}$, and $h_4 \oplus h_3 = \llbracket w_2 \rrbracket$. Consequently $h_3 = \llbracket w_2 \rrbracket$, proving the first two requirements. As Q is stable, the third requirement is ensured by Lemma 2 (Skip safety). \square

Lemma 18 (Predicate elimination). *If $\Delta, \forall \bar{x}. \alpha(\bar{x}) \equiv R; \Gamma \vdash \{P\}C\{Q\}$ and $\alpha \notin \Gamma, \Delta, P, Q$, then $\Delta; \Gamma \vdash \{P\}C\{Q\}$.*

Proof. We first claim that:

$$\delta' \in \llbracket \Delta \rrbracket \iff \exists \delta. \delta \in \llbracket \Delta, \forall \bar{x}. \alpha(\bar{x}) \equiv R \rrbracket \wedge \forall \vec{v}. \delta'[(\alpha, \vec{v}) \mapsto \perp] = \delta[(\alpha, \vec{v}) \mapsto \perp]$$

when $\Delta \vdash \text{stable}(R)$. This can be easily proved by appeal to the predicate definition semantics. Note that stability is required otherwise the set of δ s is empty.

We then make use of this result to prove that for any n, \mathbb{I} , and $\delta' \in \llbracket \Delta \rrbracket$, there exists a $\delta. \delta \in \llbracket \Delta \wedge \forall \bar{x}. \alpha(\bar{x}) \Leftrightarrow R \rrbracket$ such that

$$\llbracket \Gamma \rrbracket_{n, \delta'} = \llbracket \Gamma \rrbracket_{n, \delta}$$

and for any assertion P , it holds that

$$w \models_{\delta'} P \iff w \models_{\delta} P$$

Both results can be proved by simple appeal to the semantics. These results are sufficient to prove our main result, as other values are unaffected by the selection of δ . \square

Theorem 3 (Soundness). *If $\Delta; \Gamma \vdash \{P\}C\{Q\}$, then $\Delta; \Gamma \models \{P\}C\{Q\}$.*

Proof. Proved by induction over proof rules. We consider in detail the SKIP, PAR, FRAME, ATOMIC, CONSEQ and PRED-E rules. Other rules follow trivially by the inductive assumption.

$$\frac{}{\Delta; \Gamma \models \{P\} \text{ skip } \{P\}} \text{ (SKIP)}$$

Holds by Lemma 2 (Skip safety).

$$\frac{\Delta; \Gamma \models \{P_1\} C_1 \{Q_1\} \quad \Delta; \Gamma \models \{P_2\} C_2 \{Q_2\}}{\Delta; \Gamma \models \{P_1 * P_2\} C_1 \parallel C_2 \{Q_1 * Q_2\}} \text{ (PAR)}$$

Holds by Lemma 10 (Parallel safety decomposition).

$$\frac{\Delta; \Gamma \models \{P\} C \{Q\} \quad \text{stable}(F)}{\Delta; \Gamma \models \{P * F\} C \{Q * F\}} \text{ (FRAME)}$$

Holds by Lemma 12 (Frame safety).

$$\frac{\vdash_{\text{SL}} \{p\} C \{q\} \quad \Delta \vdash P \Rightarrow_{\delta}^{\{p\}\{q\}} Q}{\Delta; \Gamma \models \{P\} \langle C \rangle \{Q\}} \text{ (ATOMIC)}$$

Holds by Lemma 17 (Atomic safety).

$$\frac{\Delta \vdash P \equiv P' \quad \Delta; \Gamma \models \{P'\} C \{Q'\} \quad \Delta \vdash Q' \equiv Q}{\Delta; \Gamma \models \{P\} C \{Q\}} \text{ (CONSEQ)}$$

For simplicity we treat the P and Q implications separately - no loss of generality results. The P case follows from Lemma 14 (Pre-partitioning safety), while the Q case follows from Lemma 16 (Post-partitioning safety).

$$\frac{\text{stable}(R) \quad \alpha \notin \Gamma, \Delta, P, Q \quad \Delta, \forall \bar{x}. \alpha(\bar{x}) \equiv R; \Gamma \models \{P\} C \{Q\}}{\Delta; \Gamma \models \{P\} C \{Q\}} \text{ (PRED-E)}$$

Holds by appeal to Lemma 18 (Predicate elimination).

□