

Concurrent Abstract Predicates

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

¹ Imperial College, London

{td202,pg}@doc.ic.ac.uk

² University of Cambridge

{md466,mjp41,vv216}@cl.cam.ac.uk

Abstract. Abstraction is key to understanding and reasoning about large computer systems. Abstraction is simple to achieve if the relevant data structures are disjoint, but rather difficult when they are partially shared, as is often the case for concurrent modules. We present a program logic for reasoning abstractly about data structures that provides 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 physical systems, we use abstraction and locality to hide irrelevant details. For example, when building a house in London, we do not consider the gravitational forces on individual 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. Abstraction gives us a structured view of the system, because components can be represented by their abstract properties.

With locality, we can directly provide some degree of abstraction. Using separation logic [16], we can prove that a module operates only on a particular data structure, not accessed by other modules. If the structure is manipulated only by module functions, it can be represented just in terms of its abstract properties, using abstract predicates [21]. For example, we can give a specification for a set using an abstract predicate to assert that “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, after deleting 6, that “the set is $\{5\}$ ”.

This combination of abstract predicates with the low-level locality from separation logic supports coarse-grained reasoning about modules, where each data structure is represented by a single abstract predicate. However, we often need to reason about modules in a fine-grained manner, where many abstract predicates refer to properties of the same data structure. For example, a fine-grained

specification for the set module would allow element 6 to be manipulated separately from the rest of the set. Fine-grained reasoning of this sort is advocated by context logic [3].

Fine-grained abstractions often cannot be achieved using traditional abstract predicates. This is because separation in the abstraction (union, in the case of the set module) need not correspond to separation in the implementation [7]. If the set module is implemented using an array, and each element of the set is represented with a disjoint element of the array, then the high-level and low-level separation correspond. However, if the set module is implemented as a singly-linked list, then the implementation must traverse the global list to manipulate individual set elements. Individual set elements are not represented disjointly in the implementation, and fine-grained reasoning is not possible with traditional abstract predicates combined with separation logic.

In this paper, we present a program logic that allows fine-grained abstraction in the presence of sharing, by introducing *concurrent abstract predicates*. These predicates present a fiction of disjointness [7]; 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 by 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 predicates assert properties about the same shared structure, and both can be used at the same time by separate 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 [8], 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 are shared.

With our program logic, a module implementation can be verified against a high-level specification expressed using concurrent abstract predicates. Clients of the module can then be verified purely in terms of this high-level specification, without reference to the module’s implementation. We demonstrate this by presenting two implementations of a lock module satisfying the same abstract lock specification, and 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. Hence, concurrent abstract predicates provide the necessary abstraction for compositional reasoning about concurrent systems.

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 implementations satisfying the specification.

2.1 Lock Specification

A typical lock module has the functions `lock(x)` and `unlock(x)`. It also has a mechanism for constructing locks, such as `makeLock(n)`, which allocates a lock and a contiguous block of memory of size `n`. We specify these functions as:

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

This abstract specification, which is presented by the module to the client, is independent of the underlying implementation.¹ The assertions `isLock(x)` and `Locked(x)` are abstract predicates. `isLock(x)` asserts that the lock at `x` can be acquired by the thread with this assertion, while `Locked(x)` asserts that the thread holds the lock. We use the separating conjunction, `*`, from separation logic: `p * q` asserts that the state can be split disjointly into two parts, one satisfying `p` and the other satisfying `q`. Later, we give a concrete interpretation of these predicates for a simple compare-and-swap lock.

The module presents the following abstract predicate axioms to the client:

$$\text{isLock}(x) \iff \text{isLock}(x) * \text{isLock}(x) \quad (1)$$

$$\text{Locked}(x) * \text{Locked}(x) \iff \text{false} \quad (2)$$

The first axiom allows the client to share freely the knowledge that `x` is a lock.² 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.

```

lock(x) {
  local b;
  do ⟨b := ¬CAS(&x, 0, 1)⟩
  while(b)
}

unlock(x) {
  ⟨[x] := 0⟩
}

makeLock(n) {
  local x := alloc(n+1);
  [x] := 1;
  return x;
}

```

(Here angle brackets denote atomic statements.)

¹ This specification resembles those used in the work of Gotsman *et al.* [11] and Hobor *et al.* [15] on dynamically-allocated locks.

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

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 module 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 state is indivisible so that all threads maintain a consistent view of it. This is expressed by the logical equivalence $\boxed{P}_{I(\vec{x})}^r * \boxed{Q}_{I(\vec{x})}^r \Leftrightarrow \boxed{P \wedge Q}_{I(\vec{x})}^r$. 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 [2], 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, $\pi = 1$, denoting that the thread can do the action but the environment cannot.³ We now have the machinery to interpret our lock predicates concretely:

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

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 has 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 modify 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

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

³ The state model also contains a zero permission, 0, denoting that the thread may not do the action but the environment may.

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.

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 must 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 implementer of the module must show that the concrete interpretation of the predicates satisfies the axioms presented to the client. In our example, axiom 1, that only a single $\text{Locked}(x)$ can exist, follows from the presence in the local state of full permission on UNLOCK. Axiom 2, 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 post-condition, `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 referred to by x in the shared state region from 0 to 1. This update is allowed because of

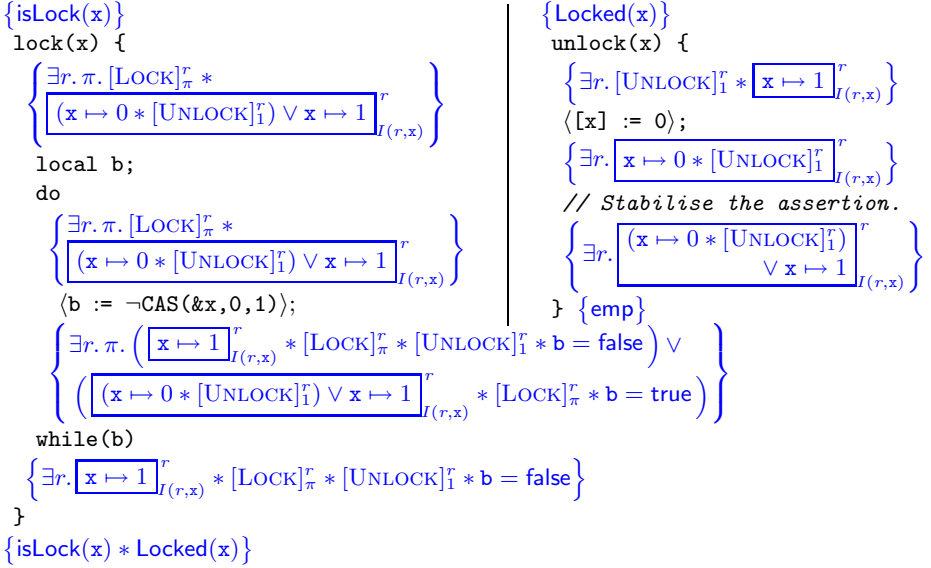


Fig. 1. Verifying the compare-and-swap lock implementation: `lock` and `unlock`

the permission $[\text{LOCK}]_{\pi}^r$ in the local state and the action in $I(r, x)$. The post-condition 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 pre-condition is still satisfied. This post-condition 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 *repartitioning operator*, denoted by \Longrightarrow , which enables us to repartition the state between regions and to create regions. In particular, we have that:

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

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

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

The final post-condition results from the definitions of `isLock(x)` and `Locked(x)`.

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, and 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

```

{emp}
makelock(n) {
  local x := alloc(n + 1);
  {x ↦ _ * (x + 1) ↦ _ * ... * (x + n) ↦ _}
  [x] := 1;
  {x ↦ 1 * (x + 1) ↦ _ * ... * (x + n) ↦ _}
  // Create shared lock region.
  {∃r. x ↦ 1I(r,x)r * [LOCK]1r * [UNLOCK]1r * (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`

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.

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 triple in sequential separation logic, where p, q denote separation logic assertions that do not specify predicates, shared regions or interference.

The second premise, $\Delta \vdash P \Longrightarrow^{\{p\}\{q\}} Q$, says that the interference allowed by P enables the state to be repartitioned to Q , given the change to memory specified 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. Pre-condition and post-condition stability is a general requirement that our proof rules have, which for presentation purposes we keep implicit in the rest of the paper.

The repartitioning arrow $P \Longrightarrow Q$ 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 repartitioning in the rule of consequence to move resources between regions. The operator \Longrightarrow includes conventional implication, so this rule of consequence subsumes the traditional one.

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

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 concurrent algorithms [23], 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.

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

This rule defines a module consisting of functions $f_1 \dots f_n$ and uses it to verify a client specification $\{P\}C\{Q\}$.

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 now consider another, more complex lock algorithm: a lock that issues tickets to clients contending for the lock. This algorithm is 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:

```
lock(x) {                               unlock(x) {                             makelock(n) {
  <int i := INCR(x.next);>                <x.owner++;>                               local x := alloc(n+2);
  while(<i ≠ x.owner>) skip;              }                                           (x+1).owner := 0;
}                                                                                     (x+1).next := 1;
                                                                                     return (x+1);
                                           }
}
```

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

The implementation assumes an atomic operation INCR that increments a stored value and returns the original value. To acquire the lock, a client atomically increments `x.next` and reads it to a variable `i`. The value of `i` becomes the client’s ticket. The client waits for `x.owner` to equal its ticket value `i`. Once this is the case, the client holds the lock. The lock is released by incrementing `x.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.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.next`, and NEXT(k) corresponds to releasing the lock when `x.owner` = k . The shared state contains permissions on 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:

$$\begin{aligned} \text{isLock}(x) &\equiv \exists t. \exists \pi. \boxed{\exists k, k'. x.\text{owner} \mapsto k * x.\text{next} \mapsto k' * \\ &\quad k \leq k' * \otimes k'' \geq k'. [\text{NEXT}(k'')]_1^t * \text{true}}_{T(t,x)}^t * [\text{TAKE}]_\pi^t \\ \text{Locked}(x) &\equiv \exists t, k. \boxed{x.\text{owner} \mapsto k * \text{true}}_{T(t,x)}^t * [\text{NEXT}(k)]_1^t \end{aligned}$$

(\otimes is the lifting of $*$ to sets; it is the multiplicative analogue of \forall .)

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

Self-stability of `Locked`(x) is ensured by the fact that the predicate holds full permission on the action NEXT(k), and the action TAKE cannot affect the `x.owner` field. Self-stability for `isLock`(x) is ensured by the fact that its definition is invariant under the action TAKE.

The axioms follow trivially from the predicate definitions, as in the CAS lock.

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, `lock`(x) pulls a ticket and a permission out of the shared state, and `unlock`(x) returns it to the shared state. (We omit the proof of `makeLock`, which is similar to the previous example.)

$\text{out}(h, v)$ 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.

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

For example, the above 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.

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

Fig. 4. Proof outline for the set module client

3.2 Example: The Coarse-Grained Set

Consider a coarse-grained set implementation, based on the lock module described in §2.1 and the sequential set operations $\text{scontains}(h, v)$, $\text{sadd}(h, v)$ and $\text{sremove}(h, v)$.

<pre>contains(h, v) { lock(h.lock); r := scontains(h.set, v); unlock(h.lock); return r; }</pre>	<pre>add(h, v) { lock(h.lock); sadd(h.set, v); unlock(h.lock); }</pre>	<pre>remove(h, v) { lock(h.lock); sremove(h.set, v); unlock(h.lock); }</pre>
---	--	--

Interpretation of Abstract Predicates. We assume a coarse-grained 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 :

$$\begin{array}{l} \{\text{Set}(h, vs)\} \text{ scontains}(h, v) \{\text{Set}(h, vs) * \text{ret} = (v \in vs)\} \\ \{\text{Set}(h, vs)\} \text{ sadd}(h, v) \{\text{Set}(h, \{v\} \cup vs)\} \\ \{\text{Set}(h, vs)\} \text{ sremove}(h, v) \{\text{Set}(h, vs \setminus \{v\})\} \end{array}$$

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

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

The $\text{SGAP}(ws)$ action allows the thread to return the set containing ws 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.

We first define the auxiliary predicates $\text{allgaps}(s)$, $P_{\in}(h, v, s)$ and $P_{\notin}(h, v, s)$:

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

allgaps defines the set of all SGAP permissions. $P_{\in}(h, v, s)$ is used to assert that the shared state 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 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, $[\text{CHANGE}(v)]_1^s$. 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 \notin in place of \in .

$\text{in}(h, v)$ and $\text{out}(h, v)$ are self-stable. For $\text{in}(h, v)$, the only actions available to another thread are $\text{CHANGE}(w)$, 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: $\text{CHANGE}(w)$ requires the disjunct $\text{allgaps} * \text{Set}(h.\text{set}, vs)$ to hold and leaves the disjunct $\text{Locked}(h.\text{lock}) * ([\text{SGAP}(vs)]_1^s - \textcircled{*} \text{allgaps}(s))$ 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 $\text{sadd}(h.\text{set}, v)$, when the permissions $[\text{CHANGE}(v)]_1^s$ and $[\text{SGAP}(vs)]_1^s$ are used to repartition between shared and local state. These steps are purely logical repartitioning of the state.

```

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


```

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

3.3 Example: 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 [13, §9.5]) is given in Fig. 6. We omit `contains` for space reasons.

The three module functions use the function `locate(h, v)` that traverses the sorted list from the head `h` up to the position for a node holding value `v`, 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.

Proof sketch We can verify this algorithm with our logic. The details are given in the technical report [5]. Here, we just present the intuition behind the proof.

```

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

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

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

Fig. 6. Lock-coupling list algorithm

As with the course-grained example, we have actions `LCHANGE` and `LGAP`, parameterised by value v . An `LCHANGE` permission allows a thread to take locked parts of the list out of the shared state into its local state, acquiring `LGAP` permissions and giving up the appropriate `Locked` predicates. These `LGAP` permissions allow the thread to return the parts of the list it acquired, having possibly inserted or removed a node with value v , and to regain the `Locked` predicates. We can then give the definition for $\text{in}(h, v)$, or for $\text{out}(h, v)$, as shared regions where the list starting at h contains, or does not contain, the value v , respectively. Both predicates include the full permission for `LCHANGE` on v . The list definition must track gaps for the currently locked segments. These gaps correspond to missing `LGAP` permissions.

4 Semantics and Soundness

We present the model for interpreting our assertions and program judgements, and sketch a proof of soundness of our logic. Details of the proof are given in [5].

4.1 Assertion Syntax

Recall from §2.3 that our proof judgements 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:

$$\begin{aligned}
 (\text{Assn}) \quad P, Q &::= \text{emp} \mid E_1 \mapsto E_2 \mid P * Q \mid P -\otimes Q \mid \text{false} \mid P \Rightarrow Q \mid \exists x. P \mid \\
 &\quad [\gamma(E_1, \dots, E_n)]_{\pi}^r \mid \boxed{P}_I^r \mid \alpha(E_1, \dots, E_n) \mid \otimes x. P \\
 (\text{BAssn}) \quad p, q &::= \text{emp} \mid E_1 \mapsto E_2 \mid p * q \mid p -\otimes q \mid \text{false} \mid p \Rightarrow q \mid \exists x. p \mid \otimes x. P \\
 (\text{IAssn}) \quad I &::= \gamma(\vec{x}) : \exists \vec{y}. (P \rightsquigarrow Q) \mid I_1, I_2
 \end{aligned}$$

(Axioms) $\Delta ::= \emptyset \mid \forall \vec{x}. P \implies Q \mid \forall \vec{x}. \alpha(\vec{x}) \equiv P \mid \Delta_1, \Delta_2$

(Triples) $\Gamma ::= \emptyset \mid \Gamma, \{P\}f\{Q\}$

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

4.2 Assertion Model

Let $(\mathbf{Heap}, \uplus, \emptyset)$ be any separation algebra [4] representing *machine states* (or *heaps*). Typically, we take \mathbf{Heap} to be the standard heap model: the set of finite partial functions from locations to values, where \uplus is the union of partial functions with disjoint domains. We let h denote a heap.

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

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

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 to an appropriate region; we defer the formal definition of well-formedness:

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

Given a logical state l , we write $l_{\mathbf{H}}$ and $l_{\mathbf{P}}$ to stand for the heap and permission assignment components respectively. We also write $w_{\mathbf{L}}$ and $w_{\mathbf{S}}$ 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 \mathbf{UNLOCK} moves the permission $[\mathbf{UNLOCK}]_1^r$ from local to shared 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.* the permission $[\text{SCHANGE}(v)]_1^s$ pertains to the token $(s, \text{SCHANGE}, v)$):

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

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

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

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

Intuitively, 1 represents full, exclusive permission, 0 represents no permission, and the intermediate values represent partial, non-exclusive permission.⁴

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

$$w \oplus w' \stackrel{\text{def}}{=} \begin{cases} (w_{\text{L}} \oplus w'_{\text{L}}, w_{\text{S}}) & \text{if } w_{\text{S}} = w'_{\text{S}} \\ \perp & \text{otherwise.} \end{cases}$$

We write $\mathbf{0}_{\text{Perm}}$ for the empty permission assignment (which assigns zero permission value 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.

We define the operator $\llbracket (l, s) \rrbracket$ which collapses a pair of a logical state l and shared state s to a single logical state, and the operator $\llbracket (l, s) \rrbracket_{\text{H}}$ which gives the heap component of $\llbracket (l, s) \rrbracket$; we use \bigoplus , the natural lifting of \oplus to sets:

$$\llbracket (l, s) \rrbracket \stackrel{\text{def}}{=} l \oplus \left(\bigoplus_{r \in \text{dom}(s)} s(r) \right) \quad \llbracket (l, s) \rrbracket_{\text{H}} \stackrel{\text{def}}{=} (\llbracket (l, s) \rrbracket)_{\text{H}}$$

The *action domain* of an interference assertion, $\text{adom}(I)$, is the set of action names and their appropriate parameters:

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

We are finally in a position to define well-formedness of worlds, $\text{wf}(l, s)$:

$$\begin{aligned} \text{wf}(l, s) &\stackrel{\text{def}}{\iff} \llbracket (l, s) \rrbracket \text{ is defined } \wedge \\ &\forall r, \gamma, \vec{v}. \llbracket (l, s) \rrbracket_{\text{P}}(r, \gamma, \vec{v}) > 0 \implies \exists l', I. s(r) = (l', I) \wedge (\gamma, \vec{v}) \in \text{adom}(I) \end{aligned}$$

⁴ This is the fractional permission model of Boyland [2]. With minimal changes we could add a *deny* permission prohibiting both the environment and thread from performing the action (see Dodds *et al.* [8]). We can achieve much the same effect in the Boyland-style system by placing a full permission in the shared state.

$$\begin{aligned}
(-)_{-,-} & : \text{Assn} \times \text{PEnv} \times \text{Interp} \rightarrow \mathcal{P}(\text{LState} \times \text{SState}) \\
\langle E_1 \mapsto E_2 \rangle_{\delta,i} & \stackrel{\text{def}}{=} \left\{ (l, s) \mid \begin{array}{l} \text{dom}(l_H) = \{ \llbracket E_1 \rrbracket_i \} \wedge l_H(\llbracket E_1 \rrbracket_i) = \llbracket E_2 \rrbracket_i \\ \wedge l_P = \mathbf{0}_{\text{Perm}} \wedge s \in \text{SState} \end{array} \right\} \\
\langle \text{emp} \rangle_{\delta,i} & \stackrel{\text{def}}{=} \{ (\emptyset, \mathbf{0}_{\text{Perm}}), s \mid s \in \text{SState} \} \\
\langle P_1 * P_2 \rangle_{\delta,i} & \stackrel{\text{def}}{=} \{ w_1 \oplus w_2 \mid w_1 \in \langle P_1 \rangle_{\delta,i} \wedge w_2 \in \langle P_2 \rangle_{\delta,i} \} \\
\langle P_1 -\otimes P_2 \rangle_{\delta,i} & \stackrel{\text{def}}{=} \{ w \mid \exists w_1, w_2. w_2 = w \oplus w_1 \wedge w_1 \in \langle P_1 \rangle_{\delta,i} \wedge w_2 \in \langle P_2 \rangle_{\delta,i} \} \\
\langle \bigotimes x. P \rangle_{\delta,i} & \stackrel{\text{def}}{=} \bigcup_W \left\{ \bigoplus_v W(v) \mid \forall v. W(v) \in \langle P \rangle_{\delta,i[x \mapsto v]} \right\} \\
\langle [\gamma(E_1, \dots, E_n)]_{\pi}^r \rangle_{\delta,i} & \stackrel{\text{def}}{=} \left\{ ((\emptyset, [\gamma, \llbracket r \rrbracket_i, \llbracket E_1 \rrbracket_i, \dots, \llbracket E_n \rrbracket_i] \mapsto \llbracket \pi \rrbracket_i), s) \mid \begin{array}{l} s \in \text{SState} \wedge \\ \llbracket \pi \rrbracket_i \in (0, 1] \end{array} \right\} \\
\langle \llbracket P \rrbracket_I^r \rangle_{\delta,i} & \stackrel{\text{def}}{=} \{ ((\emptyset, \mathbf{0}_{\text{Perm}}), s) \mid \exists l. (l, s) \in \langle P \rangle_{\delta,i} \wedge s(\llbracket r \rrbracket_i) = (l, \llbracket I \rrbracket_i) \} \\
\langle \alpha(E_1, \dots, E_n) \rangle_{\delta,i} & \stackrel{\text{def}}{=} \delta(\alpha, \llbracket E_1 \rrbracket_i, \dots, \llbracket E_n \rrbracket_i) \\
\\
[-]_{-,-} & : \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 P \rangle_{\delta,i} \mid \text{wf}((l, s)) \}
\end{aligned}$$

Fig. 7. Semantics of assertions. The cases for conjunction, implication, existential, etc. are standard, simply distributing over the local and shared state.

4.3 Assertion Semantics

Fig. 7 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, and just assume an appropriate semantics. The semantics of IAssns can also depend on the semantics of free variables. We define $\llbracket - \rrbracket_- : \text{IAssns} \times \text{Interp} \rightarrow \text{IAssns}$ to replace the free variables with their values.

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

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

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 the iterated version of $*$; that is, the finite, multiplicative analogue of \forall . The empty assertion emp asserts that the local state and permission assignment 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 $\text{Locked}(x)$ to assert that x is held by the current thread. The meaning of an abstract predicate is simply given by 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, $\llbracket I \rrbracket_i$. 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 $\llbracket I \rrbracket_i$ 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 the 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 defines 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 \boxed{P}_I^r to \boxed{Q}_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.

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

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 defined in 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 which updates, from world w to w' , the thread is allowed to perform. A thread can update its local state as it pleases, but any change to a shared region, r , must correspond to an action, $\gamma(\vec{v})$, for which the thread has sufficient permission, $(w_L)_P(r, \gamma, \vec{v}) > 0$. For example, in the CAS lock proof (§2.2), the thread must hold permission 1 on UNLOCK before unlocking. Without this restriction, other threads could potentially unlock the lock.

It is important that each update preserves the total amount of permission in the world, that is, $\lfloor w \rfloor_P = \lfloor w' \rfloor_P$, so that threads cannot acquire permissions out of thin air. This does *not* hold for heaps, as we permit memory allocation.

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. $(G^c)^{-1}$).

$$G^c \stackrel{\text{def}}{=} \left\{ (w, w') \mid \begin{array}{l} \exists r, I, \ell_1, \ell_2. \ r \notin \text{dom}(w_S) \wedge w'_S = w_S[r \mapsto (\ell_1, I)] \wedge \\ w_L = \ell_1 \oplus \ell_2 \wedge w'_L = \ell_2 \oplus \text{all}(I, r) \end{array} \right\}$$

$$G_\delta \stackrel{\text{def}}{=} \left\{ (w, w') \mid \begin{array}{l} ((\exists r, \gamma, \vec{v}. (w_S, w'_S) \in \llbracket (w_S(r))_2 \rrbracket_\delta(r, \gamma, \vec{v}) \wedge \\ (w_L)_P(r, \gamma, \vec{v}) > 0) \vee w_S = w'_S) \wedge \lfloor w \rfloor_P = \lfloor w' \rfloor_P \end{array} \right\} \cup G^c \cup (G^c)^{-1}$$

Some permitted updates do not modify the heap, but simply repartition it between shared regions. This is captured by $\overline{G}_\delta \stackrel{\text{def}}{=} G_\delta \cap \{(w, w') \mid \llbracket w \rrbracket = \llbracket w' \rrbracket\}$. In practice, we allow an unlimited number of repartitionings in a single step, only one of which actually modifies the heap. This is captured by \widehat{G}_δ , defined as:

$$\widehat{G}_\delta \stackrel{\text{def}}{=} (\overline{G}_\delta)^*; G_\delta; (\overline{G}_\delta)^*$$

We now define the notion of *repartitioning* with respect to an update from p to q , written $P \Longrightarrow_\delta^{\{p\}\{q\}} Q$. This asserts any world w_1 satisfying P can be collapsed to a heap $\llbracket w_1 \rrbracket$ that has a subheap h_1 satisfying the separation logic assertion p ; furthermore, when this subheap is replaced with any subheap h_2 that satisfies q , the resulting heap can be reconstructed into a world w_2 that satisfies Q and for which the transition from w_1 to w_2 is permitted by the guarantee, \widehat{G}_δ . The guarantee limits the repartitioning that takes place between the regions.

Definition 1 (Repartitioning). $P \Longrightarrow_\delta^{\{p\}\{q\}} Q$ holds iff, for every variable interpretation i and 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. $(w_1, w_2) \in \widehat{G}_\delta$.

Note that, if $p = q = \text{emp}$, then the repartitioning preserves the concrete state, and only allows the world to be repartitioned. We 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 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 ($\llbracket w \rrbracket_{\text{P}}(r, \gamma, \vec{v}) < 1$). 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.

$$R^c \stackrel{\text{def}}{=} \left\{ (w, w') \mid \begin{array}{l} \exists r, \ell, I. r \notin \text{dom}(w_{\text{S}}) \wedge w'_{\text{L}} = w_{\text{L}} \wedge w'_{\text{S}} = w_{\text{S}}[r \mapsto (\ell, I)] \wedge \\ \llbracket w' \rrbracket \text{ defined} \wedge (\forall \gamma, \vec{v}. \llbracket w' \rrbracket_{\text{P}}(r, \gamma, \vec{v}) = 0) \end{array} \right\}$$

$$R_{\delta} \stackrel{\text{def}}{=} \left\{ (w, w') \mid \begin{array}{l} \exists r, \gamma, \vec{v}. (w_{\text{S}}, w'_{\text{S}}) \in \llbracket (w_{\text{S}}(r))_2 \rrbracket_{\delta}(r, \gamma, \vec{v}) \wedge \\ w_{\text{L}} = w'_{\text{L}} \wedge \llbracket w \rrbracket_{\text{P}}(r, \gamma, \vec{v}) < 1 \end{array} \right\} \cup R^c \cup (R^c)^{-1}$$

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). *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). *pstable(δ) holds iff for all $X \in \text{ran}(\delta)$, for all w and w' , if $w \in X$ and $(w, w') \in R_{\delta}$, then $w' \in X$.*

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:

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

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

Fig. 8. Selected proof rules

We require that atomic statements $\langle C \rangle$ are not nested and that function names $f_1 \dots f_n$ for a let are distinct. Here c stands for basic commands, modelled semantically as subsets of $\mathcal{P}(\text{Heap} \times \text{Heap})$ satisfying the locality conditions of [4].

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. 8. These rules are modified from RGSep [23] and deny-guarantee [8]. 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 ATOMIC and PRIM is a judgement in standard sequential separation logic. The other minor judgements 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 $\llbracket \Delta \rrbracket \subseteq \llbracket \Delta' \rrbracket$.

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 predicate definitions are satisfiable; otherwise the premise of a proof rule could be vacuously true. We ensure this by requiring that all occurrences of the predicate in its definition are positive.

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. The ATOMIC and CONSEQ rule were discussed in §2.3. That section also discussed a rule for modules, which can be derived using PRED-I, LET and PRED-E.

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 judgements, 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. $\forall w', \text{ if } (w, w') \in (R_\delta)^* \text{ then } C, w', \eta, \delta, i, Q \text{ safe}_n$;
2. $\neg((C, \llbracket w \rrbracket) \xrightarrow{\eta} \text{fault})$;
3. $\forall C', h', \text{ if } (C, \llbracket w \rrbracket) \xrightarrow{\eta} (C', h'), \text{ then there } \exists w' \text{ such that } (w, w') \in \widehat{G}_\delta, h' = \llbracket w' \rrbracket \text{ and } C', w', \eta, \delta, i, Q \text{ safe}_n$; and
4. $\text{ if } C = \text{skip}, \text{ then } \exists w' \text{ such that } \llbracket w \rrbracket = \llbracket w' \rrbracket, (w, w') \in \widehat{G}_\delta, \text{ and } w' \in \llbracket Q \rrbracket_{\delta, i}$.

This definition says that a configuration is safe provided that: (1) changing the world 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 \widehat{G}_δ in the third and fourth conjuncts allows the world to be repartitioned.

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 $\Delta; \Gamma \models \{P\} C \{Q\}$.*

Proof is by structural induction. Most interesting is PAR, which embodies the compositionality of our logic. The proof requires the following lemma.

Lemma 1 (Abstract state locality)

If $(C, \llbracket w_1 \oplus w_2 \rrbracket) \xrightarrow{\eta} (C', h)$ and $C, w_1, \eta, \delta, i, Q$ safe_n , then $\exists w'_1, w'_2$ such that $(C, \llbracket w_1 \rrbracket) \xrightarrow{\eta} (C', \llbracket w'_1 \rrbracket)$, $h = \llbracket w'_1 \oplus w'_2 \rrbracket$, $(w_1, w'_1) \in \widehat{G}_\delta$, and $(w_2, w'_2) \in (R_\delta)^*$.

Proof. We require that basic commands obey a concrete locality assumption. We must prove that the rely and guarantee obey similar locality lemmas. The lemma then follows from the definition of configuration safety. \square

This lemma shows that program only affects the resources required for it to run safely: that is, programs are safely contained within their abstract footprints. The soundness of PAR follows immediately.

5 Conclusions and Related Work

Our program logic allows fine-grained abstraction in a concurrent setting. It brings together three streams of research: (1) abstract predicates [21] for abstracting the internal details of a module or class; (2) deny-guarantee [8] for reasoning about concurrent programs; and (3) context logic [3, 7] for fine-grained reasoning at the module level.

Our work on concurrent abstract predicates has been strongly influenced by O’Hearn’s concurrent separation logic (CSL) [19]. CSL takes statically allocated locks as a primitive. With CSL, we can reason about programs with 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 with new proof rules and assertions to deal with dynamically-allocated locks [11, 15] and re-entrant locks [12]. Parkinson *et al.* [20] have shown how the CSL approach can be used to reason about concurrent algorithms that do not use locks. However, representing the concurrent interactions in an invariant can require intractable auxiliary state.

Jacobs and Piessens [17] have developed an approach to reasoning abstractly that is based on CSL for dynamically allocated locks [11]. Their logic uses auxiliary state to express the temporal nature of interference. To deal modularly with auxiliary state they add a special implication that allows the auxiliary state to be changed in any way that satisfies the invariant. This implication is similar to our repartitioning operator \Longrightarrow . Unlike our operator, theirs can be used in a module specification, allowing a client’s auxiliary state to be updated during the module’s execution. Our operator could be extended with this technique, which may simplify the use of the lock specification in the set algorithms.

An alternative to using invariants is to abstract interference over the shared state by relations modelling the interaction of different threads: the rely-guarantee method [18]. There have been two recent logics that combine RG with separation logic: RGSep [23] and SAGL [10]. Both allow more elegant proofs of concurrent algorithms than the invariant-based approaches, but they have the drawback that interference on the shared state cannot be abstracted. Feng’s Local Rely-Guarantee [9] improves the locality of RGSep and SAGL by scoping interference with a precise footprint, but this approach still has no support

for abstraction. Many of our ideas on abstraction originated in Dinsdale-Young, Gardner and Wheelhouse’s work on using RGSep to analyse a concurrent B-tree algorithm [6, 22].

We have combined RGSep with resource permissions, as first introduced for deny-guarantee reasoning [8]. Deny-guarantee is a reformulation of rely-guarantee allowing reasoning about dynamically scoped concurrency. 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.

We have also used ideas from context logic [3], a high-level logic for fine-grained local reasoning about program modules. Recent work in context logic has shown how implementations of modules can be verified, by relating local reasoning about module specifications with the low-level reasoning about implementations [7]. As presented here, these implementations break away from the fiction of disjointness presented by the module specifications.

Proofs in our proof system are slightly more complex in practice than RGSep and SAGL, as can be seen by comparing the lock-coupling list proof with the RGSep one [23]. This may 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 [14]. Linearisability provides a fiction of atomicity allowing “reason[ing] about properties of concurrent objects given just their (sequential) specifications” [14]. 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 a single logical formalism.

Acknowledgements. Thanks to Richard Bornat, Alexey Gotsman, Peter O’Hearn, Suresh Jagannathan, Mohammad Raza, Noam Rinetzky, Mark Wheelhouse, John Wickerson and the anonymous referees for useful feedback. This work was supported by an EPSRC DTA (Dinsdale-Young), an EPSRC Programme Grant EP/H008373/1 (Dinsdale-Young and Gardner), an EPSRC grant EP/F019394/1 (Dodds and Parkinson), an RAEng/EPSC research fellowship (Parkinson) and a Microsoft Research Cambridge/RAEng senior research fellowship (Gardner).

References

- [1] Bierhoff, K., Aldrich, J.: Modular typestate checking of aliased objects. In: OOPSLA, pp. 301–320 (2007)
- [2] Boyland, J.: Checking interference with fractional permissions. In: Cousot, R. (ed.) SAS 2003. LNCS, vol. 2694. Springer, Heidelberg (2003)
- [3] Calcagno, C., Gardner, P., Zarfaty, U.: Local reasoning about data update. *Festschrift Computation, Meaning and Logic: Articles dedicated to Gordon Plotkin*, 172 (2007)
- [4] Calcagno, C., O’Hearn, P.W., Yang, H.: Local action and abstract separation logic. In: Symp. on Logic in Comp. Sci (LICS 2007), pp. 366–378 (2007)

- [5] Dinsdale-Young, T., Dodds, M., Gardner, P., Parkinson, M., Vafeiadis, V.: Concurrent abstract predicates. Technical Report 777, University of Cambridge Computer Laboratory (2010)
- [6] Dinsdale-Young, T., Gardner, P., Wheelhouse, M.: Local reasoning about a concurrent B*-list algorithm. Talk and unpublished report (2009), <http://www.doc.ic.ac.uk/~pg/>
- [7] Dinsdale-Young, T., Gardner, P., Wheelhouse, M.: Locality refinement. Technical Report DTR10-8, Imperial College London (2010)
- [8] Dodds, M., Feng, X., Parkinson, M., Vafeiadis, V.: Deny-guarantee reasoning. In: Castagna, G. (ed.) ESOP 2009. LNCS, vol. 5502, pp. 363–377. Springer, Heidelberg (2009)
- [9] Feng, X.: Local rely-guarantee reasoning. In: POPL (2009)
- [10] Feng, X., Ferreira, R., Shao, Z.: On the relationship between concurrent separation logic and assume-guarantee reasoning. In: De Nicola, R. (ed.) ESOP 2007. LNCS, vol. 4421, pp. 173–188. Springer, Heidelberg (2007)
- [11] Gotsman, A., Berdine, J., Cook, B., Rinetzky, N., Sagiv, M.: Local reasoning for storable locks and threads. In: Shao, Z. (ed.) APLAS 2007. LNCS, vol. 4807, pp. 19–37. Springer, Heidelberg (2007)
- [12] Haack, C., Huisman, M., Hurlin, C.: Reasoning about Java’s Reentrant Locks. In: Ramalingam, G. (ed.) APLAS 2008. LNCS, vol. 5356, pp. 171–187. Springer, Heidelberg (2008)
- [13] Herlihy, M., Shavit, N.: The Art of Multiprocessor Programming. Elsevier, Amsterdam (2008)
- [14] Herlihy, M.P., Wing, J.M.: Linearizability: a correctness condition for concurrent objects. TOPLAS 12(3), 463–492 (1990)
- [15] Hobor, A., Appel, A.W., Nardelli, F.Z.: Oracle semantics for concurrent separation logic. In: Drossopoulou, S. (ed.) ESOP 2008. LNCS, vol. 4960, pp. 353–367. Springer, Heidelberg (2008)
- [16] Ishtiaq, S.S., O’Hearn, P.W.: BI as an assertion language for mutable data structures. In: POPL, pp. 14–26 (January 2001)
- [17] Jacobs, B., Piessens, F.: Modular full functional specification and verification of lock-free data structures. Technical Report CW 551, Katholieke Universiteit Leuven, Department of Computer Science (June 2009)
- [18] Jones, C.B.: Annotated bibliography on rely/guarantee conditions (2007), <http://homepages.cs.ncl.ac.uk/cliff.jones/ftp-stuff/rg-hist.pdf>
- [19] O’Hearn, P.W.: Resources, concurrency and local reasoning. TCS (2007)
- [20] Parkinson, M., Bornat, R., O’Hearn, P.: Modular verification of a non-blocking stack. In: POPL, pp. 297–302 (January 2007)
- [21] Parkinson, M.J., Bierman, G.M.: Separation logic and abstraction. In: POPL, pp. 247–258 (2005)
- [22] Pinto, P.: Reasoning about B^{Link} trees. Advanced masters ISO project, Imperial College London (2010), Supervised by Dinsdale-Young, Gardner and Wheelhouse
- [23] Vafeiadis, V., Parkinson, M.: A marriage of rely/Guarantee and separation logic. In: Caires, L., Vasconcelos, V.T. (eds.) CONCUR 2007. LNCS, vol. 4703, pp. 256–271. Springer, Heidelberg (2007)