

Modular Reasoning about Separation of Concurrent Data Structures

Kasper Svendsen¹, Lars Birkedal¹, and Matthew Parkinson²

¹ IT University of Copenhagen
{kasv,birkedal}@itu.dk

² Microsoft Research Cambridge
mattpark@microsoft.com

Abstract. In a concurrent setting, the usage protocol of standard separation logic specifications are not refinable by clients, because standard specifications abstract all information about potential interleavings. This breaks modularity, as libraries cannot be verified in isolation, since the appropriate specification depends on how clients intend to use the library.

In this paper we propose a new logic and a new style of specification for thread-safe concurrent data structures. Our specifications allow clients to refine usage protocols and associate ownership of additional resources with instances of these data structures.

1 Introduction

Why? One of the challenges of specifying the abstract behavior of a library is that the appropriate specification depends on the context in which the library is going to be used. Consider a simple bag library with operations to push and pop elements from the bag. In a sequential setting the standard separation logic specification is:

$$\begin{aligned} & \{\mathbf{bag}_e(x, X)\} \text{ x.Push}(y) \{\mathbf{bag}_e(x, X \cup \{y\})\} \\ & \{\mathbf{bag}_e(x, X)\} \text{ x.Pop}() \{\text{ret. } (X = \emptyset \wedge \text{ret} = \text{null} \wedge \mathbf{bag}_e(x, X)) \vee \\ & \quad (\exists Y. X = Y \cup \{\text{ret}\} \wedge \mathbf{bag}_e(x, Y))\} \\ & \mathbf{bag}_e(x, X) * \mathbf{bag}_e(x, Y) \Rightarrow \perp \end{aligned}$$

Here \mathbf{bag}_e is an abstract predicate, i.e., implicitly existentially quantified, so that clients cannot depend on its definition [2], x is a reference to a bag object, and X and Y range over multisets of elements. The implication in the third line expresses that the \mathbf{bag}_e predicate cannot be duplicated. Hence this specification enforces that clients follow a strict usage protocol, with a single exclusive owner of the bag object. On the other hand, this specification allows the owner of the bag to track the exact contents of the bag. In other words, $\mathbf{bag}_e(x, X)$ asserts full ownership of the bag and that the bag contains exactly the objects in the multiset X .

Now consider a client of the bag library and suppose this client wants to implement a bag of independent tasks scheduled for execution. This client might

not care about the exact contents of the bag, only that each task in the bag owns the resources necessary to perform its task. In addition, this client might wish to share the bag to allow multiple users to schedule tasks for execution. Thus this client might prefer the following specification for shared bags:

$$\begin{aligned} & \{ \text{bag}_s(x, P) * P(y) \} \text{x.Push}(y) \{ \text{bag}_s(x, P) \} \\ & \{ \text{bag}_s(x, P) \} \text{x.Pop}() \{ \text{ret. } \text{bag}_s(x, P) * (\text{ret} = \text{null} \vee P(\text{ret})) \} \\ & \text{bag}_s(x, P) \Rightarrow \text{bag}_s(x, P) * \text{bag}_s(x, P) \end{aligned}$$

This specification allows more sharing, but it does not track the exact contents of the bag. Instead, it allows clients to associate additional resources with each element of the bag using the P predicate, and to freely share the bag as expressed by the implication in the third line. Clients thus transfer $P(y)$ to the bag when pushing y , and receive $P(\text{ret})$ from the bag, when pop returns a non-null element.

In a sequential first-order setting without reentrancy, the standard separation logic specification suffices. Using techniques from fictional separation logic [11], clients can refine the standard specification to allow the additional sharing of the shared bag specification. However, in a concurrent setting, it is easy to come up with a non-thread-safe implementation (without synchronization), that satisfies the standard specification (as it enforces a single exclusive owner), but not the shared bag specification. Hence, in a higher-order concurrent setting with reentrancy, this type of refinement is unsound!

What? The key challenge is to provide a logic that enables clients to refine the specifications to their requirements in a concurrent setting. In this paper we propose such a logic, called Higher-Order Concurrent Abstract Predicates (HOCAP), and a new style of specification for thread-safe concurrent data structures.¹ This style of specification allows clients to refine the usage protocol and associate ownership of additional resources with instances of the data structure, in a concurrent higher-order setting.

How? Observe first that while it is not sound to refine specifications to allow *more sharing* in a concurrent setting, it is sound to refine specifications to permit *less sharing*. Thus we will start with a weak specification that allows unrestricted sharing of instances of the data structure, and then let clients refine this specification as needed.

To reason about sharing we partition the state into *regions*, with *protocols* governing how the state in each region is allowed to evolve, following earlier work on concurrent abstract predicates [5]. Our new program logic, HOCAP, also uses *phantom fields* – a logical construct akin to auxiliary variables, that only occur in the logic.

To support abstract refinement of library specifications, we propose to verify the implementation using a region to share the concrete state of the implementation, with a fixed protocol that *relates* the concrete state of the implementation

¹ We consider a concurrent data structure thread-safe if each of its methods has one or more synchronization points, where the abstract effects of the method appear to take affect. See Related Work for a discussion of the relation to linearizability.

with an abstract description of the state of the data structure. To refine this specification, clients define a region of their own, with a protocol on the *abstract state* of the data structure. For soundness, these two regions must evolve in lock-step and *synchronize* when the abstract state changes (in synchronization points). We do so by giving each region a half permission to a shared phantom field; synchronization can then be enforced since updating a phantom field requires full permission. Half permissions have previously been used to synchronize local and shared state [14]; here we are using it to synchronize two shared regions.

For the bag example, we introduce a phantom field `cont` that contains the abstract state of the bag: a multiset of references to the elements in the bag. The bag constructor also returns a half permission to the phantom field `cont`:

$$\frac{}{\{\text{emp}\}\text{new Bag}()\{\text{ret. bag}(\text{ret}) * \text{ret}_{\text{cont}} \xrightarrow{1/2} \emptyset\}}$$

Here $\text{ret}_{\text{cont}} \xrightarrow{1/2} \emptyset$ asserts partial ownership of the phantom `cont` field. Since the client obtains half the `cont` permission upon calling the constructor, the library cannot update the `cont` field on its own.

The protocol governing the bag `x` thus relates the concrete state of the bag with its abstract state (the value of the `cont` field):

$$(\exists X. x_{\text{cont}} \xrightarrow{1/2} X * \text{list}(x, X)) \quad \rightsquigarrow \quad (\exists X. x_{\text{cont}} \xrightarrow{1/2} X * \text{list}(x, X))$$

This protocol permits any atomic update to the region containing the internal state of bag `x` from a state satisfying the left side of \rightsquigarrow to a state satisfying the right side.

To allow the library to update `cont` in synchronization points, we therefore transfer the library's half-permission to the client and require the client to update the phantom field with the abstract effects of the method, and then transfer a half-permission back to the library. When the client updates the phantom field, the client is forced to prove that the abstract effects of the method is permitted by whatever protocols the client may have imposed on the abstract state.

We express the update to the phantom `cont` field using a *view-shift* [4]. Conceptually, a view-shift corresponds to a step in the instrumented semantics that does not change the concrete machine state. View-shifts, written $P \sqsubseteq Q$, thus generalize assertion implication by allowing updates to phantom fields (given sufficient permission) and ownership transfer between the local state and shared regions.

The bag push method thus requires the client to provide a view-shift, to update the abstract state from X to $X \cup \{y\}$ in the synchronization point:

$$\frac{\forall X. x_{\text{cont}} \xrightarrow{1/2} X * P \sqsubseteq x_{\text{cont}} \xrightarrow{1/2} X \cup \{y\} * Q}{\{\text{bag}(x) * P\}x.\text{Push}(y)\{\text{bag}(x) * Q\}}$$

Here, P and Q are universally quantified and thus picked by the client. Hence, the client can use P and Q to perform further updates of the instrumented state in the synchronization point and relate the new abstract state with its local state. We thus refer to P and Q as synchronization pre- and postconditions.

Likewise, the bag pop operation requires two view-shifts; one, in case the bag is empty in the synchronization point, and another, in case the bag is non-empty in the synchronization point:

$$\frac{\begin{array}{l} x_{\text{cont}} \xrightarrow{1/2} \emptyset * P \sqsubseteq x_{\text{cont}} \xrightarrow{1/2} \emptyset * Q(\text{null}) \\ \forall X. \forall y. x_{\text{cont}} \xrightarrow{1/2} X \cup \{y\} * P \sqsubseteq x_{\text{cont}} \xrightarrow{1/2} X * Q(y) \end{array}}{\{\text{bag}(x) * P\}x.\text{Pop}()\{\text{bag}(x) * Q(\text{ret})\}}$$

Finally, the **bag** predicate is freely duplicable:

$$\text{bag}(x) \Rightarrow \text{bag}(x) * \text{bag}(x)$$

Note that since P and Q are universally quantified — our logic is *higher order* — the client could potentially pick instantiations referring to the library’s region, thus introducing self-referential region assertions. We can illustrate this problem by instantiating P with an assertion that itself refers to the bag in the specification of **Push**. Since $\text{bag}(x)$ asserts that there exists a shared region that owns half the x_{cont} field, it follows that $\text{bag}(x) * x_{\text{cont}} \mapsto _ \Rightarrow \text{false}$. Hence, by instantiating P with $\text{bag}(x) * x_{\text{cont}} \xrightarrow{1/2} _$, we can derive the postcondition **false** from the specification of **Push**.

To prevent this, we introduce a notion of *region type* and a notion of *support*, as an over-approximation of the types of regions a given assertion refers to. Our formal bag specification (presented in Section 3) thus imposes support restrictions on P and Q to ensure the client does not introduce self-referential region assertions.

Another key challenge we address is higher-order protocols. Higher-order protocols are crucial to allow clients to associate ownership of additional resources with shared data structures. For example, to derive the shared bag specification from the generic specification, we use a second region with a protocol that requires clients to transfer ownership of $P(x)$, when pushing x into the bag:

$$(\exists X. x_{\text{cont}} \xrightarrow{1/2} X * \otimes_{y \in X} P(y)) \rightsquigarrow (\exists X. x_{\text{cont}} \xrightarrow{1/2} X * \otimes_{y \in X} P(y))$$

Again, P is a predicate variable and could be instantiated to refer to the state and protocol of this and other regions – making the above protocol a higher-order protocol. We also use region types to break a circularity introduced by higher-order protocols. In particular, instead of assigning protocols to individual regions, we assign parameterized protocols to region types. This allows us to reason about higher-order protocols that refer to the region types – and thus, implicitly, the protocol – of other regions. We show that this well-behaved subset of higher-order protocols, called *state-independent* protocols, suffices for sophisticated libraries, such as the Joins library [16].

To summarize, our new logic and specification methodology allows clients to refine the usage protocol of the bag. It also allows clients to transfer ownership of resources to the bag, by transferring them to a client region synchronized with the abstract state of the bag.

More details and examples can be found in the extended version of this article, which is available at <http://www.itu.dk/people/kasv/hocap-ext.pdf>.

Related Work. Jacobs and Piessens introduced the idea of parameterizing the specification of concurrent methods with ghost code, to be executed in synchronization points [10]. Here we build on their idea, using a much stronger logic based on CAP [5], to address the main problem with their approach.

Instead of regions with protocols, Jacobs and Piessens use ghost objects – data structures built from ghost variables – with handles that represent partial information about the data structure and permissions to modify it. While these handles provide support for reasoning about the state of shared ghost objects, they lack the ability to associate ownership of additional state with ghost objects. Instead, Jacobs and Piessens use the lock invariant of the lock protecting the concurrent data structure to associate ownership of additional state.

However, this approach is problematic without proper storable locks. In particular, Jacobs and Piessens logic and model of storable locks only supports lock labels parameterized over simple types (i.e., not assertions). This forces the *client* to create the synchronization primitive, so that the *client* can pick a lock invariant containing both the state of the concurrent data structure and any additional resources the client may wish to associated with the data structure. This breaks abstraction, by exposing internal implementation details to the client (the synchronization primitive used) and it requires the client to reprove the shared bag specification every time it is needed. Hence, Jacobs and Piessens cannot derive the shared bag specification. We solve this problem using higher-order protocols.

CAP was designed to verify concurrent data structures [5]. However, the original specifications and proofs are non-modular in the sense that implementations have been verified against unrefinable specifications with fixed usage protocols.

Recently, Dodds et. al. introduced a higher-order variant of CAP to give a generic specification for a library for deterministic parallelism [6]. While their proofs make explicit use of nested region assertions and higher-order protocols, the authors failed to recognize the semantic difficulties these features introduce. Consequently, their reasoning is unsound. In particular, their higher-order representation predicates are not stable.

Another approach for achieving modular reasoning is to prove concurrent implementations to be contextual refinements of coarse-grained counterparts – thus taking the coarse-grained counterparts as specifications. Previous efforts for proving such contextual refinements have mostly focused on indirect proofs through a linearizability property on traces of concurrent libraries [9,7]. So far, this approach lacks support for transfer of ownership of resources between client and library. More recently, there has been work on proving such contextual refinements directly, using logical relations [20]. Unless combined with a program logic, both of these approaches restrict all reasoning to statements about contextual refinement or contextual equivalence. As our approach demonstrates, if a Hoare-style specification is what we are ultimately interested in, then contextual refinement is unnecessary; what we really want is a generic specification that is refinable by clients.

Conceptually, linearizability aims to provide a fiction of atomicity to clients of concurrent libraries. Our approach does not. Instead, we aim to allow clients to reason about changes of the abstract state in synchronization points *inside* concurrent libraries. To illustrate the distinction, consider an extension of the bag library with a `Push2(x, y)` method that takes two elements and pushes them one at a time (i.e., with the implementation `Push(x); Push(y)`). This method is not linearizable, as it has two synchronization points. However, it still has a natural specification expressed in terms of two view-shifts, one for each synchronization point:

$$\frac{\forall X. x_{\text{cont}} \xrightarrow{1/2} X * P \sqsubseteq x_{\text{cont}} \xrightarrow{1/2} X \cup \{y\} * Q}{\forall X. x_{\text{cont}} \xrightarrow{1/2} X * Q \sqsubseteq x_{\text{cont}} \xrightarrow{1/2} X \cup \{z\} * R} \{ \text{bag}(x) * P \} x. \text{Push2}(y, z) \{ \text{bag}(x) * R \}$$

From this specification, a client can derive a natural shared bag specification:

$$\{ \text{bag}_s(x, P) * P(y) * P(z) \} x. \text{Push2}(y, z) \{ \text{bag}_s(x, P) \}$$

Contributions. We propose a new style of specification for thread-safe concurrent data structures. Using protocol synchronization, this style of specification allows clients to refine the usage protocol of concurrent data structures. Moreover, using nested region assertions and state-independent higher-order protocols, our specification style allows clients to associate additional resources with the data structure.

Technically, we realize the ideas by developing HOCAP, a higher-order separation logic for a subset of C^\sharp featuring named delegates and fork concurrency. The logic allows two or more protocols to be synchronized and evolve in lock-step. In addition, we support nested region assertions, state-independent higher-order protocols, and guarded recursive assertions. We present a step-indexed model of the logic and use it to prove the logic sound. We emphasize that unlike earlier versions of CAP, our logic includes sufficient proof rules for carrying out all proofs (including stability proofs) of examples *in the logic*, i.e., without passing to the semantics.

Lastly, in the extended version we demonstrate the power and utility of the logic by verifying a library for executing tasks in parallel, based on Doug Lea’s Fork/Join framework [12]. We have also used the logic to specify and verify the Joins library [16] and clients thereof, which will be described in a separate paper.

2 The Logic

Our logic is a general program logic for a subset of C^\sharp , featuring delegates referring to named methods² and an atomic compare-and-swap statement. New threads are

² Anonymous delegates in C^\sharp may capture the l -values of free variables and hence the semantics and logic for anonymous methods is non-trivial, see our earlier paper [18]. Those semantic issues are orthogonal to what we discuss in the present paper and hence we omit anonymous delegates here.

allocated via a fork statement that forks a delegate. Each thread has a private stack, but all threads share a common heap. We use an interleaving semantics.

The specification logic is an intuitionistic higher-order logic over a simply typed term language, and the assertion logic an intuitionistic higher-order separation logic over the same simply typed term language. Types are closed under the usual type constructors, \rightarrow , \times , and $+$. Basic types include the type of assertions, Prop , the type of specifications, Spec , the type of C^\sharp values, Val , and the type of fractional permissions, Perm .

2.1 Concurrent Abstract Predicates

Recall that the basic idea behind CAP is to provide an abstraction of possible interference from concurrently executing threads, by partitioning the state into regions, with protocols governing how the state in each region is allowed to evolve. Requiring all assertions to be *stable* – i.e., closed under protocols – and proving all specifications with respect to arbitrary stable frames, then achieves thread-local reasoning about shared mutable state.

Following earlier work on CAP [5], we use a shared region assertion, written $\boxed{\text{P}}^{r,t,a}$, which asserts that r is a region and that the resources in region r satisfy the assertion P . Unlike earlier versions, the region assertion is also annotated with a region type t and a protocol argument a , since we assign parameterized protocols to region types instead of regions, as mentioned above. Region assertions are freely duplicable and thus satisfy,

$$\boxed{\text{P}}^{r,t,a} \Leftrightarrow \boxed{\text{P}}^{r,t,a} * \boxed{\text{P}}^{r,t,a} \quad (1)$$

Protocols consist of *named actions* and updates to a shared region require *ownership* of a named action justifying the update. Protocols are specified using protocol assertions, written $\text{protocol}(t, l)$. Here t is a region type and l is a parametric protocol. We use the following notation for a parametric protocol l with parameter a and named actions $\alpha_1, \dots, \alpha_n$:

$$l(a) = (\alpha_1 : (\Delta_1). P_1 \rightsquigarrow Q_1; \dots; \alpha_n : (\Delta_n). P_n \rightsquigarrow Q_n)$$

Here Δ_i is a context of logical variables relating the action precondition P_i with the action postcondition Q_i . The action α_i thus allows updates from states satisfying P_i to states satisfying Q_i . We use $l(a)[\alpha_i]$ to refer to the definition of the α_i action in protocol l applied to argument a . Hence, $l(a)[\alpha_i] = (\Delta_i). P_i \rightsquigarrow Q_i$. We use $\boxed{\text{P}}^{r,t,a}$ as shorthand for $\boxed{\text{P}}^{r,t,a} * \text{protocol}(t, l)$.

We can distinguish different client roles in protocols through ownership of named actions. An action assertion $[\alpha]_\pi^r$ asserts fractional ownership of the named action α on region r with fraction π . Fractions are used to allow multiple clients to use the same action. We can split or reassemble action assertions using the following property,

$$[\alpha]_{p+q}^r \Leftrightarrow [\alpha]_p^r * [\alpha]_q^r \quad (2)$$

where $p, q, p + q$ are terms of type Perm – permissions in $(0, 1]$.

An assertion \mathbf{p} is *stable* if it is closed under interference from the environment. In the absence of self-referential region assertions and higher-order protocols, the region assertion, $\boxed{\mathbf{P}}^{r,t,a}$ is stable if \mathbf{P} is closed under all $l(a)$ actions:³

$$\forall \tilde{y}. \text{valid}(\mathbf{P} \wedge P_i(\tilde{y}) \Rightarrow \perp) \vee \text{valid}(Q_i(\tilde{y}) \Rightarrow \mathbf{P})$$

for all i , where $l(a)[\alpha_i] = (\tilde{x})$. $P_i(\tilde{x}) \rightsquigarrow Q_i(\tilde{x})$.

Example. To illustrate reasoning about sharing, consider a counter with read and increment methods. Since the count can only be increased, this counter satisfies the specification of a monotonic counter [15]:

$$\begin{array}{l} \{\text{counter}(x, n)\} \ x.\text{Increment}() \ \{\text{counter}(x, n + 1)\} \\ \{\text{counter}(x, n)\} \ \ x.\text{Read}() \ \ \{\text{ret. counter}(x, \text{ret}) * n \leq \text{ret}\} \\ \text{counter}(x, n) \Rightarrow \text{counter}(x, n) * \text{counter}(x, n) \end{array}$$

Here $\text{counter}(x, n)$ asserts that n is a lower-bound on the current count. Hence we expect that this predicate can be freely duplicated, as expressed by the third line above.

To verify a counter implementation against this specification, we place the current count in a shared region, with a protocol that allows the current count to be increased. Assertions about lower bounds are thus invariant under the protocol. If the counter implementation maintains the current count in field `count`, then we can specify the counter protocol as follows:

$$\text{counter}(x, n) \stackrel{\text{def}}{=} \exists r, \pi. [\text{INCR}]_{\pi}^r * \boxed{\exists m. n \leq m * x.\text{count} \mapsto m}^{r, \text{Counter}, x}$$

where l is a parametric protocol with parameter x and a single action `INCR`, that allows the count field of x to be increased:

$$l(x) = (\text{INCR} : (m, k : \mathbf{N}). x.\text{count} \mapsto m * m \leq k \rightsquigarrow x.\text{count} \mapsto k)$$

Here we have used a fixed region type `Counter` for the counter region r . Since fractional permissions can always be split (2), and region assertions always duplicated (1), it follows that $\text{counter}(x, n) \Rightarrow \text{counter}(x, n) * \text{counter}(x, n)$, as required by the specification. Since the shared region assertion in $\text{counter}(x, n)$ contains no self-referential region assertions or higher-order protocols, to prove it stable, it suffices to show that,

$$\begin{array}{l} \forall m, k. \text{valid}((\exists m : \mathbf{N}. n \leq m * x.\text{count} \mapsto m) \wedge (x.\text{count} \mapsto m * m \leq k) \Rightarrow \perp) \vee \\ \text{valid}(x.\text{count} \mapsto k \Rightarrow (\exists m : \mathbf{N}. n \leq m * x.\text{count} \mapsto m)) \end{array}$$

This follows easily by case analysis on $n \leq k$. Lastly, to verify the implementation of `Increment` and `Read`, we have to prove they satisfy the protocol, namely that they do not decrease the current count. This is easy.

³ This is a formula in the specification logic; \mathbf{P} and \mathbf{Q} are assertions and for an assertion \mathbf{P} , $\text{valid}(\mathbf{P})$ is the specification that expresses that \mathbf{P} is valid in the assertion logic.

2.2 Higher-Order Concurrent Abstract Predicates

As the above example illustrates, we can use CAP to reason about a shared counter by imposing a protocol on the shared count field. Since this is a protocol on a primitive resource (the `count` field), first-order CAP suffices. To reason about examples, such as the shared bag, which associates ownership of general resources – through the `P` predicate – with a shared bag, we need Higher-Order CAP. In particular, to define the `bags` predicate requires region and protocol assertions containing the predicate variable `P`.

To support modular reasoning about region and protocol assertions containing predicate and assertion variables, ideally, we want to treat predicate and assertion variables as black boxes. For instance, consider the assertion,

$$Q \stackrel{\text{def}}{=} \boxed{P}^{r,t,-} * \text{protocol}(t, l) \quad (3)$$

where `l` is the parametric protocol $l(-) = (\tau : P \rightsquigarrow P)$ expressed in terms of the assertion variable `P`. Treating `P` as a black box, `Q` is clearly stable if `P` is stable, as `Q` asserts that `P` holds of the resources in region `r`, which is clearly closed under the protocol `l`. However, in general `P` could itself be instantiated with region and protocol assertions, introducing the possibility of *self-referential region assertions* and turning `l` into a *higher-order protocol*. This makes reasoning significantly more challenging. In particular, some self-referential region assertions do not admit modular stability proofs: it is possible to instantiate `P` with stable assertions for which `Q` is not stable. Furthermore, higher-order protocols introduce a circularity in the definition of the model.

Self-referential Region Assertions. To see how self-referential region assertions can break the modularity of stability proofs, consider assertion `P` below:

$$P \stackrel{\text{def}}{=} x \mapsto 0 * \boxed{y \mapsto 0}^{r,t',-} * \text{protocol}(t', J),$$

where `J` is the protocol with a single α action that allows the `y` variable to be changed from 0 to 1, provided region `r` owns variable `x` and `x` is zero:

$$J(-) = \left(\alpha : \boxed{x \mapsto 0}^{r,t,-} * y \mapsto 0 \rightsquigarrow \boxed{x \mapsto 0}^{r,t,-} * y \mapsto 1 \right)$$

Then `P` is stable, because `P` asserts full ownership of the `x` variable, ensuring that the environment cannot perform the α action, as `x` cannot also be owned by region `r`. However, the region assertion `Q` defined above is not stable when instantiated with this `P`, as $\boxed{P}^{r,t,-}$ asserts that region `r` *does* own `x`, thus allowing the environment to perform the α action. As this example illustrates, some self-referential region assertions thus do not admit modular stability proofs. A similar problem occurs when reasoning about atomic updates to shared regions.

Support. To ensure modular reasoning about stability and atomic updates to shared regions, we require clients to explicitly prove that their instantiations of predicate variables do not introduce self-referential region assertions. To facilitate these proofs, we introduce a notion of support, which gives an over-approximation of the types of regions a given assertion refers to.

An assertion P is supported by a set of region types A , if P is invariant under arbitrary changes to the state and protocol of any region of a region type not in A . To support modular reasoning about hierarchies of concurrent libraries, instead of reasoning directly in terms of sets of regions, we introduce a partial order on region types and reason in terms of upwards-closed sets of region types. More formally, we introduce a new type, $RType$, of region types with a partial order $\leq : RType \times RType \rightarrow Spec$, with a bottom element $\perp : RType$ and finite meets. We say that an assertion P is dependent on region type t if it is supported by the set of region types greater than or equal to t . We introduce two new specification assertions, $dep, indep : RType \times Prop \rightarrow Spec$ for asserting that an assertion is dependent and independent of a given region type, respectively. The inference rules for dep and $indep$ are fairly natural. For instance, if P is dependent on region type t_1 , then $\boxed{P}^{r,t_2,a}$ is dependent on the greatest lower bound, of t_1 and t_2 .

Whenever we reason about region assertions, $\boxed{P}^{r,t,a}$ we thus require that P is independent of the region type t . This excludes self-referential region assertions through protocols (such as in (3)), and through nested region assertions (such as $\boxed{\boxed{P}^{r,t,a}}^{r,t,a}$).

Stability. General higher-order protocols would introduce a circularity in the definition of the model. We break this circularity by exploiting the indirection of region types – i.e., that we assign protocols to region types instead of individual regions. This allows us to support protocols with assertions about the region types of regions, but without assertions about the protocols assigned to those region types. Technically, we enforce this restriction by ignoring protocol assertions in action pre- and postconditions when interpreting protocols. The parameterized higher-order protocol l ,

$$l(x) = (x \mapsto 0 * \text{protocol}(t, J) \rightsquigarrow x \mapsto 1 * \text{protocol}(t, J))$$

is thus interpreted as $l(x) = (x \mapsto 0 \rightsquigarrow x \mapsto 1)$. The interpretation simply ignores the $\text{protocol}(t, J)$ assertion (See definition of act in the technical report [19]).

In the absence of self-referential region assertions, a region assertion $\boxed{P}^{r,t,a}$ is stable under the α action, if P is closed under the action pre- and postcondition of the α action of $l(a)$ and l is a first-order protocol. If l is a higher-order protocol, then the assertion $\boxed{P}^{r,t,a}$ is stable under the α action, if P is closed under the

action pre- and postcondition of the α action of $I(\mathbf{a})$ and P is also *protocol-pure*. We thus have the following proof rule for stability:

$$\frac{\begin{array}{c} I(\mathbf{a})[\alpha] = (\tilde{x}).I_p(\tilde{x}) \rightsquigarrow I_q(\tilde{x}) \quad \forall \tilde{x}. \text{valid}(P \wedge I_p(\tilde{x}) \Rightarrow \perp) \vee \text{valid}(I_q(\tilde{x}) \Rightarrow P) \\ \text{indep}_t(P) \quad \text{indep}_t(Q) \quad \text{stable}(P * Q) \quad \text{pure}_{\text{protocol}}(P) \quad \text{pure}_{\text{state}}(Q) \end{array}}{\text{stable}_{\alpha}^r \left(\boxed{P}^{r,t,a} * Q \right)} \text{SA}$$

Here $\text{pure}_{\text{protocol}}$ and $\text{pure}_{\text{state}}$ are propositions in the specification logic; $\text{pure}_{\text{protocol}}(P)$ expresses that P is invariant under any changes to protocols and $\text{pure}_{\text{state}}(P)$ expresses that P is invariant under any change to the local or shared state. The SA proof rule thus allows us to prove stability of region assertions, by first “pulling out” any protocol assertions, Q , from the region assertion. We say that an assertion is *expressible using state-independent protocols* if the protocol assertions can be “pulled out” in this sense. Formally,

$$\text{sip} \stackrel{\text{def}}{=} \lambda P : \text{Prop}. \exists Q, R : \text{Prop}. \text{valid}(P \Leftrightarrow Q * R) \wedge \text{pure}_{\text{protocol}}(Q) \wedge \text{pure}_{\text{state}}(R)$$

In particular, if $P \Leftrightarrow Q * R$ and $\text{pure}_{\text{state}}(R)$, then $\boxed{P}^{r,t,a} \Leftrightarrow \boxed{Q}^{r,t,a} * R$. Thus, if $\text{sip}(P)$, then $\boxed{P}^{r,t,a}$ can be rewritten to a form that satisfies the $\text{pure}_{\text{protocol}}$ premise of the SA rule. Expressibility using state-independent protocols is closed under conjunction and separating conjunction, but in general not under disjunction or existential quantification. To achieve closure under existential quantification, $\exists x : X. P(x)$, we have to impose a stronger restriction on the predicate family P . Namely, P has to be uniformly expressible using state-independent protocols:

$$\begin{aligned} \text{usip}_X \stackrel{\text{def}}{=} \lambda P : X \rightarrow \text{Prop}. \exists R : \text{Prop}. \exists Q : X \rightarrow \text{Prop}. \text{pure}_{\text{state}}(R) \wedge \\ \forall x \in X. (P(x) \Leftrightarrow Q(x) * R) \wedge \text{pure}_{\text{protocol}}(Q(x)) \end{aligned}$$

Then we have that $\text{usip}_X(P) \Rightarrow \text{sip}(\exists x \in X. P(x))$.

2.3 View-Shifts

Phantom State. Proofs in Hoare logic often employ auxiliary variables [13], as an abstraction of the history of execution and state. To support this style of reasoning, without changing the formal operational semantics, we instrument our abstract semantics with phantom fields.

We thus extend our logic with a phantom points-to assertion, written $x_f \overset{p}{\mapsto} v$, which asserts partial ownership, with fraction p , of the phantom field f on object x , and that the current value of the phantom field is v .

Phantom fields live in the instrumented state and are thus updated through view-shifts. Updating a phantom field requires full ownership of the field ($x_f \overset{1}{\mapsto} v_1 \sqsubseteq_{\perp} x_f \overset{1}{\mapsto} v_2$).⁴ A fractional phantom field permission can be split and re-assembled arbitrarily. As a partial fraction only confers read-only ownership, two

⁴ The view-shift is annotated with the \perp region type; we explain the reason for such annotations on view-shifts in the following.

partial fractional assertions must agree on the current value of a given phantom field ($x_f \stackrel{P_1}{\mapsto} v_1 * x_f \stackrel{P_2}{\mapsto} v_2 \Rightarrow v_1 = v_2$). To create a phantom field f we require that the field does not already exist, so that we can take full ownership of the field. We thus require all phantom fields of an object o to be created simultaneously when o is first constructed (in the proof rule for constructors, see the technical report [19]).

Simultaneous Updates. To support synchronization of two regions by splitting ownership of a common phantom field, we need to update the value of the phantom field in both regions *simultaneously*. Previous versions of CAP have only supported sequences of *independent* updates to *single* regions. To support synchronization of protocols we thus extend CAP with support for simultaneous updates of *multiple* regions.

We have chosen a semantics that requires that updates of regions have the same action granularity (you cannot have one simultaneous update of two regions, where the update of one region is justified by one action, and the update of the other region is justified by two actions). This is a choice; it simplifies stability proofs, but it means that we must explicitly track the regions that may have been updated by a view-shift. We thus index the view-shift relation with a region type t . The indexed view-shift relation, \sqsubseteq_t , thus describes a *single* update that, in addition to updating the local state, may update *multiple* shared regions with region types not greater than or equal to t , where each update must be justified by a *single* action. The indexed view-shift relation is thus *not* transitive.

Figure 1 contains a selection of proof rules for view-shifts. The two main rules, VSNOPEN and VSOPEN, are used to open a region, to allow access to the resources in that shared region. Both rules allow us to open a region and perform a nested view-shift on the contents of that region. This is how we reason about simultaneous updates to multiple regions in the logic. Rule VSNOPEN allows the nested view-shift to modify further regions, while VSOPEN does not (note the use of region type \perp on the nested view shift in VSOPEN). Both rules require a proof the update is possible –

$$P_1 * P_2 \sqsubseteq_{t_1 \sqcap t_2} Q_1 * Q_2 \quad \text{and} \quad P_1 * P_2 \sqsubseteq_{\perp} Q_1 * Q_2,$$

respectively – and a proof that the update is allowed by the protocol, denoted

$$\boxed{P_1}^{r, t_1, a} * P_2 \rightsquigarrow^{r, t_2} \boxed{Q_1}^{r, t_1, a} * Q_2$$

and explained below.

Since actions owned by shared regions cannot be used to perform updates to shared regions, the VSNOPEN rule further requires that P_1 does not assert ownership of any local action permissions ($\text{pure}_{\text{perm}}(P_1)$). This ensures that no local action permissions from P_1 were used to justify any actions performed in the nested view-shift. Since VSOPEN does not allow the nested view-shift to update any regions, this restriction is unnecessary for the VSOPEN rule.

$$\begin{array}{c}
 \frac{\text{pure}_{\text{perm}}(P_1) \quad \text{indep}_{t_1 \sqcap t_2}(P_1, P_2, Q_1, Q_2) \quad t_2 \not\leq t_1}{\boxed{P_1}^{r, t_1, a} * P_2 \rightsquigarrow^{r, t_2} \boxed{Q_1}^{r, t_1, a} * Q_2 \quad P_1 * P_2 \sqsubseteq_{t_1 \sqcap t_2} Q_1 * Q_2} \text{VSNOOPEN} \\
 \boxed{P_1}^{r, t_1, a} * P_2 \sqsubseteq_{t_2} \boxed{Q_1}^{r, t_1, a} * Q_2 \\
 \\
 \frac{\text{indep}_{t_1 \sqcap t_2}(P_1, P_2, Q_1, Q_2) \quad t_2 \not\leq t_1}{\boxed{P_1}^{r, t_1, a} * P_2 \rightsquigarrow^{r, t_2} \boxed{Q_1}^{r, t_1, a} * Q_2 \quad P_1 * P_2 \sqsubseteq_{\perp} Q_1 * Q_2} \text{VSOOPEN} \\
 \boxed{P_1}^{r, t_1, a} * P_2 \sqsubseteq_{t_2} \boxed{Q_1}^{r, t_1, a} * Q_2 \\
 \\
 \frac{P \sqsubseteq_t Q \quad \text{stable}(R)}{P * R \sqsubseteq_t Q * R} \text{VSFRAME} \qquad \frac{P \sqsubseteq_{t_1} Q \quad t_1 \leq t_2}{P \sqsubseteq_{t_2} Q} \text{VSWEAKEN}
 \end{array}$$

Fig. 1. Selected view-shift proof rules

Update Allowed. The update allowed relation, $P \rightsquigarrow^{r,t} Q$, asserts that the update described by P and Q to region r is justified by an action owned by P .

Thus the basic proof rule for the update allowed relation is:

$$\frac{\text{indep}_{t_2}(P(\tilde{v}), Q(\tilde{v})) \quad t_2 \not\leq t_1 \quad l(a)[\alpha] = (\tilde{x}). P(\tilde{x}) \rightsquigarrow Q(\tilde{x})}{\boxed{P(\tilde{v})}^{r, t_1, a} * [\alpha]_{\pi}^r \rightsquigarrow^{r, t_2} \boxed{Q(\tilde{v})}^{r, t_1, a} * [\alpha]_{\pi}^r} \text{UA ACT}$$

Since the update allowed relation simply asserts that any update described by P and Q is allowed, it satisfies a slightly non-standard rule of consequence, that allows strengthening of both the pre- and postcondition. From this non-standard rule-of-consequence, it follows that the update allowed relation satisfies a frame rule that allows arbitrary changes to the context:

$$\frac{P \Rightarrow P' \quad P' \rightsquigarrow^{r,t} Q' \quad Q \Rightarrow Q'}{P \rightsquigarrow^{r,t} Q} \text{UA CONSEQ} \qquad \frac{P \rightsquigarrow^{s,t} Q}{P * R_1 \rightsquigarrow^{s,t} Q * R_2} \text{UAF}$$

3 Concurrent Bag

We now return to the concurrent bag from the introduction, and show how to formalize the informal specification from the introduction. Next, we show how to derive the two bag specifications from the introduction, using protocol synchronization, nested region assertions, and higher-order protocols.

Specification. In the introduction we proposed a refineable bag specification with phantom variables to force protocol synchronization and with view-shifts to synchronize client and library in synchronization points. In the formal specification we restrict the synchronization pre- and postconditions, P and Q , using

region types, to ensure that the client's instantiation does not introduce self-referential region assertions. Upon creation of new bag instances, the client picks a region type t for that bag instance and the client is then required to prove that all its synchronization pre- and postconditions are independent of region type t . The formal refinable bag specification is:

$$\begin{array}{c}
\overline{\{\text{emp}\}\text{new Bag}()\{\text{ret. bag}(t, \text{ret}) * \text{ret}_{\text{cont}} \xrightarrow{1/2} \emptyset\}} \\
\text{stable}(P) \quad \text{stable}(Q) \quad \text{indep}_t(P) \quad \text{indep}_t(Q) \\
\forall x. x_{\text{cont}} \xrightarrow{1/2} \emptyset * P(x) \sqsubseteq_t x_{\text{cont}} \xrightarrow{1/2} \emptyset * Q(x, \text{null}) \\
\forall X. \forall x, y. x_{\text{cont}} \xrightarrow{1/2} X \cup \{y\} * P(x) \sqsubseteq_t x_{\text{cont}} \xrightarrow{1/2} X * Q(x, y) \\
\hline
\{\text{bag}(t, x) * P(x)\}x.\text{Pop}()\{\text{ret. bag}(t, x) * Q(x, \text{ret})\} \\
\text{stable}(P) \quad \text{stable}(Q) \quad \text{indep}_t(P) \quad \text{indep}_t(Q) \\
\forall X. \forall x, y. x_{\text{cont}} \xrightarrow{1/2} X * P(x, y) \sqsubseteq_t x_{\text{cont}} \xrightarrow{1/2} X \cup \{y\} * Q(x, y) \\
\hline
\{\text{bag}(t, x) * P(x, y)\}x.\text{Push}(y)\{\text{bag}(t, x) * Q(x, y)\} \\
\hline
\text{bag}(t, x) \Leftrightarrow \text{bag}(t, x) * \text{bag}(t, x) \qquad \text{dep}_t(\text{bag}(t, x))
\end{array}$$

The indep_t assumptions on the synchronization pre- and postconditions ensure that P and Q do not introduce self-referential region assertions. Furthermore, the index on the view-shifts, \sqsubseteq_t , ensures that the granularity of actions match between the library and any client protocols.

Exclusive Owner. We now show how to derive the standard specification with a single exclusive owner. This specification is very simple to derive; we simply let the exclusive owner of the bag keep the $1/2$ permission of the phantom field containing the abstract state of the bag: $\text{bag}_e(t, x, X) \stackrel{\text{def}}{=} \text{bag}(t, x) * x_{\text{cont}} \xrightarrow{1/2} X$.

Shared Bag. The derivation of the shared bag specification is more interesting, as it uses both protocol synchronization and higher-order protocols. We begin by formalizing the shared bag specification in our logic:

$$\begin{array}{c}
\frac{\text{dep}_r(P)}{\text{dep}_{r\text{it}}(\text{bag}_s(t, x, P))} \qquad \frac{\text{stable}(P) \quad \text{indep}_t(P) \quad \text{usip}_{\text{Val}}(P)}{\{\text{emp}\}\text{new Bag}()\{\text{ret. bag}_s(t, \text{ret}, P)\}} \\
\hline
\{\text{bag}_s(t, x, P) * P(y)\}x.\text{Push}(y)\{\text{bag}_s(t, x, P)\} \\
\hline
\{\text{bag}_s(t, x, P)\}x.\text{Pop}()\{\text{ret. bag}_s(t, x, P) * (\text{ret} = \text{null} \vee P(\text{ret}))\} \\
\hline
\text{bag}_s(t, x, P) \Leftrightarrow \text{bag}_s(t, x, P) * \text{bag}_s(t, x, P)
\end{array}$$

This corresponds to the specification from the introduction, except with restrictions on predicate P to ensure it is expressible using state-independent protocols and does not introduce self-referential protocol or region assertions.

With these restrictions on P we can now derive the shared bag specification from our generic specification. The idea is to introduce a new region containing the state associated with each element currently in the bag:

$$\begin{aligned} \mathbf{bag}_s(t, x, P) &\stackrel{\text{def}}{=} \exists r : \text{RId}. \exists \pi : \text{Perm}. \exists t_1, t_2 : \text{RType}. \\ &\quad t \leq t_1 \wedge t \leq t_2 \wedge t_1 \not\leq t_2 \wedge t_2 \not\leq t_1 \wedge \text{indep}_t(P) \wedge \text{usip}(P) \wedge \\ &\quad \mathbf{bag}(t_1, x) * \boxed{\mathbf{q}(x, P)}_{(P)}^{r, t_2, x} * [\text{UPD}]_{\pi}^r \\ \mathbf{q}(x, P) &\stackrel{\text{def}}{=} \exists X : \mathcal{P}_m(\text{Val}). x_{\text{cont}} \xrightarrow{1/2} X * \otimes_{y \in X} P(y) \\ \mathbf{l}(P)(x) &\stackrel{\text{def}}{=} (\text{UPD} : \mathbf{q}(x, P) \rightsquigarrow \mathbf{q}(x, P)) \end{aligned}$$

The parametric protocol $\mathbf{l}(P)$ allows the bag to be changed arbitrarily, provided the region still contains the state associated with each element currently in the bag. From the assumption that each $P(x)$ is stable and that $\text{usip}_{\text{Val}}(P)$ it follows that $\mathbf{q}(x, P)$ is stable and $\text{sip}(\mathbf{q}(x, P))$. Hence, there exists $R, S : \text{Prop}$ such that $\mathbf{q}(x, P) \Leftrightarrow R * S$, $\text{pure}_{\text{protocol}}(S)$ and $\text{pure}_{\text{state}}(R)$. Thus, $\mathbf{bag}_s(t, x, P)$ is equivalent to the following assertion:

$$\exists r, \pi, t_1, t_2. t \leq t_1 \wedge t \leq t_2 \wedge t_1 \not\leq t_2 \wedge t_2 \not\leq t_1 \wedge \mathbf{bag}(t_1, x) * \boxed{\mathbf{S}}_{(P)}^{r, t_2, x} * R * [\text{UPD}]_{\pi}^r$$

Hence, to prove $\mathbf{bag}_s(t, x, P)$ stable, it suffices to prove stability of $\boxed{\mathbf{S}}_{(P)}^{r, t_2, x} * R$. Applying rule SA, it thus suffices to prove,

$$\text{valid}(\mathbf{q}(x, P) \wedge S \Rightarrow \perp) \vee \text{valid}(\mathbf{q}(x, P) \Rightarrow S)$$

and the right disjunct follows easily from the assumption that $\mathbf{q}(x, P) \Leftrightarrow R * S$.

To derive the shared bag specification for `push`, we thus have to transfer the resources associated with the element being pushed, $P(y)$, to the client region containing the element resources. We thus instantiate P and Q in the generic bag specification with $P(y) * \boxed{\mathbf{q}(x, P)}_{(P)}^{r, t_2, x} * [\text{UPD}]_{\pi}^r$ and $\boxed{\mathbf{q}(x, P)}_{(P)}^{r, t_2, x} * [\text{UPD}]_{\pi}^r$, respectively.

We thus have to provide a view-shift to synchronize the abstract state of the library protocol with our client protocol r :

$$\begin{aligned} \forall X : \mathcal{P}_m(\text{Val}). x_{\text{cont}} \xrightarrow{1/2} X * P(y) * \boxed{\mathbf{q}(x, P)}_{(P)}^{r, t_2, x} * [\text{UPD}]_{\pi}^r &\sqsubseteq_{t_1} \\ x_{\text{cont}} \xrightarrow{1/2} (X \cup \{y\}) * \boxed{\mathbf{q}(x, P)}_{(P)}^{r, t_2, x} * [\text{UPD}]_{\pi}^r & \end{aligned}$$

Since $x_{\text{cont}} \xrightarrow{1/2} X * P(y) * [\text{UPD}]_{\pi}^r$ and $\mathbf{q}(x, P)$ are all independent of region type t , by rule VSOPEN it suffices to prove that the change to region r is allowed and possible. The update is easily shown to be allowed by the UPD action, using the UAACT rule and update action frame rule (UAF). To show the possibility of the view shift it suffices to prove that:

$$\begin{aligned} x_{\text{cont}} \xrightarrow{1/2} X * P(y) * \exists Z : \mathcal{P}_m(\text{Val}). x_{\text{cont}} \xrightarrow{1/2} Z * \otimes_{z \in Z} P(z) * [\text{UPD}]_{\pi}^r &\sqsubseteq_{\perp} \\ x_{\text{cont}} \xrightarrow{1/2} (X \cup \{y\}) * \exists Z : \mathcal{P}_m(\text{Val}). x_{\text{cont}} \xrightarrow{1/2} Z * \otimes_{z \in Z} P(z) * [\text{UPD}]_{\pi}^r & \end{aligned}$$

which follows easily, as $x_{\text{cont}} \xrightarrow{1/2} X * x_{\text{cont}} \xrightarrow{1/2} Z \Rightarrow X = Z$.

Note that to provide a view-shift to synchronize the abstract state of the library protocol with the client protocol, we were essentially forced to update the phantom field `cont` in the client region, which in turn forced us to transfer ownership of $P(y)$ to the client region.

4 Semantics

In this section we sketch the model and the interpretation of our logic. Due to lack of space, we focus on parts presented in Section 2. The full model, interpretation and accompanying soundness proof can be found in the technical report [19].

The presentation of the model is strongly inspired by the Views framework presentation [4]. The model is an instance of the Views framework extended with step-indexing to model guarded recursion, and thread local state to model dynamic allocation of threads.

The basic structure of the model is defined below. Assertions are modeled as step-indexed predicates on instrumented states (\mathcal{M}). Instrumented states consist of three components, a local state, a shared state and an action model. The local state specifies the current local resources. The shared state is further partitioned into regions and each region consists of a local state, a region type and a protocol parameter. The action model maps region types to parameterized protocols, which are functions from a tuple containing a protocol argument, a region identifier and an action identifier to an action. Lastly, actions are modeled as certain step-indexed relations on shared states. In particular, actions are *not* relations on shared states *and* action models, and thus do not support general higher-order protocols. Actions do however support state-independent protocols, through the region type indirection.

$$\begin{aligned}
\text{LState} &\stackrel{\text{def}}{=} \text{Heap} \times \text{PHeap} \times \text{Cap} & \text{SState} &\stackrel{\text{def}}{=} \text{RId} \rightarrow (\text{LState} \times \text{RType} \times \text{Val}) \\
\mathcal{M} &\stackrel{\text{def}}{=} \text{LState} \times \text{SState} \times \text{AMod} & \text{AMod} &\stackrel{\text{def}}{=} \text{RType} \rightarrow ((\text{Val} \times \text{RId} \times \text{AId}) \rightarrow \text{Act}) \\
\text{Cap} &\stackrel{\text{def}}{=} \{f \in \text{RId} \times \text{AId} \rightarrow [0, 1] \mid \exists R \subseteq_{\text{fin}} \text{RId}. \forall r \in \text{RId} \setminus R. \forall \alpha \in \text{AId}. f(r, \alpha) = 0\} \\
\text{Act} &\stackrel{\text{def}}{=} \{R \in \mathcal{P}(\mathbb{N} \times \text{SState} \times \text{SState}) \mid \\
&\quad \forall (i, s_1, s_2) \in R. \forall j \leq i. \forall r \in \text{RId} \setminus \text{dom}(s_2). \forall n \in \text{RType}. \forall l, l' \in \text{LState}. \\
&\quad \quad s_1 \leq s_2 \wedge (j, s_1, s_2) \in R \wedge (j, s_1, s_2[r \mapsto (l', n)]) \in R \wedge \\
&\quad \quad (j, s_1[r \mapsto (l, n)], s_2[r \mapsto (l', n)]) \in R\} \\
\text{Prop} &\stackrel{\text{def}}{=} \{U \in \mathcal{P}(\mathbb{N} \times \mathcal{M}) \mid \forall (i, m_1) \in U. \forall j \leq i. \forall m_2 \in \mathcal{M}. \\
&\quad \quad (m_1 =_j m_2 \vee m_1 \leq m_2) \Rightarrow (j, m_2) \in U\} \\
\text{Spec} &\stackrel{\text{def}}{=} \{U \in \mathcal{P}(\mathbb{N}) \mid \forall i \in U. \forall j \leq i. j \in U\}
\end{aligned}$$

The semantics of both the assertion logic and specification logic is step-indexed. The specification logic is step-indexed to allow reasoning about mutual recursion. The assertion logic is step-indexed to support nested triples (which embed specifications in the assertion logic) [17] and guarded recursive predicates [1,3]. Specifications are thus modeled as downwards closed subsets of numbers, and assertions

are modeled as step-indexed predicates on instrumented states, that are downwards closed in the step-index and upwards closed in \mathcal{M} . The upwards closure in \mathcal{M} ensures that assertions are closed under allocation of new regions and protocols (the ordering \leq on \mathcal{M} is defined as expected). To define guarded recursive functions and predicates, the types of our logic are modeled as sets with a step-indexed equivalence relation, $=_i$, and terms and predicates are modeled as non-expansive functions. However, as this part of the model is mostly orthogonal to CAP, we will elide the details, which can be found in the technical report [19].

Comparison with Previous Models of CAP. The original model of (first-order) CAP [5] employed a syntactic treatment of actions to break a circularity in the definition of worlds. Our model follows the previous model of higher-order CAP (without higher-order protocols) [6] in treating actions semantically. However, to support higher-order protocols we introduce a new indirection, in the form of region types. Actions are thus relations on shared states, which include the region types of allocated regions. Actions can thus implicitly refer to the protocol on regions through the region type indirection. While previous work has only considered CAP for a *first-order programming language*, our HOCAP is for a *higher-order programming language*. We thus step-index both the specification and assertion logic, instead of just the specification logic.

Model Operations. Separating conjunction is interpreted as the lifting of the partial commutative $\bullet_{\mathcal{M}}$ function to \mathbf{Prop} (point-wise in the step-index). The $\bullet_{\mathcal{M}}$ function expresses how to compose two instrumented states. Two instrumented states are combinable if they agree on the shared state and action model, by combining their local states, using $\bullet_{\mathbf{LState}}$. Local states are combined using the standard combination function, $\bullet_{\mathbb{W}}$, on disjoint partial functions, on the heap and phantom heap component, and by point-wise summing up the action permissions.

While assertions are modeled as step-indexed predicates on instrumented states, which include phantom fields, protocols, and regions, the operational semantics operates on concrete states, which are simply heaps. The main soundness theorem (Theorem 1) expresses that any step in the concrete semantics has a corresponding step in the instrumented semantics. This is expressed in terms of an erasure function, $[-] \in \mathcal{M} \rightarrow \mathbf{Heap}$, that erases the instrumentation from an instrumented state. The erasure of an instrumented state is simply the combination of the local state and all shared regions.

$$\begin{aligned} \llbracket (l, s) \rrbracket &\stackrel{\text{def}}{=} l \bullet_{\mathbf{LState}} \prod_{r \in \text{dom}(s)} s(r).l \\ \llbracket (l, s, \varsigma) \rrbracket &\stackrel{\text{def}}{=} \begin{cases} h & \text{if } (h, ph, c) = \llbracket (l, s) \rrbracket \text{ and } \pi_1(\text{dom}(ph)) \subseteq \text{objs}(h) \\ \text{undef} & \text{otherwise} \end{cases} \end{aligned}$$

Interference. The interference relation $R_i^A \subseteq \mathcal{M} \times \mathcal{M}$ describes possible interference from the environment. It is defined as the reflexive, transitive closure of the single-action interference relation, \hat{R}_i^A (defined below), that describes possible environment interference using at most one action on each region. Defining

R_i^A as the reflexive, transitive closure of \hat{R}_i^A forces a common action granularity on updates to multiple regions with protocols referring to each other. In addition to the step-index $i \in \mathbb{N}$, the single-action interference relation is also indexed by a set $A \in \mathcal{P}(\text{RType})$ of region types of those regions that are allowed to change and that actions justifying those changes are allowed to depend on.

$$\begin{aligned} (l_1, s_1, \varsigma_1) \hat{R}_i^A (l_2, s_2, \varsigma_2) \quad \text{iff} \quad & l_1 = l_2 \wedge s_1 \leq s_2 \wedge \varsigma_1 \leq \varsigma_2 \wedge \lceil (l_1, s_1) \rceil \text{ defined} \wedge \\ & (\forall r \in \text{dom}(s_1). s_1(r) = s_2(r) \vee (\exists \alpha. s_1(r).t \in A \wedge \\ & (\lceil (l_1, s_1) \rceil.c)(r, \alpha) < 1 \wedge (i, s_1|_A, s_2|_A) \in \varsigma_1(s_1(r).t)(s_1(r).a, r, \alpha))) \\ s|_A \stackrel{\text{def}}{=} \lambda r \in \text{RId}. \quad & \begin{cases} s(r) & \text{if } r \in \text{dom}(s) \text{ and } s(r).t \in A \\ \text{undef} & \text{otherwise} \end{cases} \end{aligned}$$

In particular, the \hat{R}_i^A relation expresses that the environment is not allowed to change the local state ($l_1 = l_2$), but it is allowed to allocate new regions and protocols ($s_1 \leq s_2$ and $\varsigma_1 \leq \varsigma_2$). Furthermore, the environment is allowed to update the resources of any region r with a region type in A ($s_1(r).t \in A$), provided the update is justified by an action α that is partially owned by the environment ($\lceil (l_1, s_1) \rceil(r, \alpha) < 1$).

An assertion is stable if it is closed under interference to all region types:

$$\text{stable}(p) \stackrel{\text{def}}{=} \{i \in \mathbb{N} \mid \forall j \leq i. \forall (m_1, m_2) \in R_j^{\text{RType}}. (j, m_1) \in p \Rightarrow (j, m_2) \in p\}$$

Previous models of CAP have only permitted multiple independent updates, whereas our model supports multiple dependent updates. Previous models thus lack the A -index that we use to enforce a common action granularity on updates to multiple dependent regions.

View-Shifts. View-shifts describe a step in the instrumented semantics that correspond to a no-op in the concrete semantics. To perform a view-shift from p to q we thus have to prove that for every concrete state c in the erasure of some instrumented state $m \in p$ there exists an instrumented state $m' \in q$ such that c is in the erasure of m' .

$$\begin{aligned} p \sqsubseteq_t q \stackrel{\text{def}}{=} \{i \in \mathbb{N} \mid \forall m \in \mathcal{M}. \forall j \in \mathbb{N}. 0 \leq j \leq i \Rightarrow \\ \lfloor p * \{(j, m)\} \rfloor_j \subseteq \lfloor q * \{(j, m') \mid m \hat{R}_j^{\{t' \mid t \preceq t'\}} m'\} \rfloor_j \} \end{aligned}$$

To allow framing on view-shifts (rule VSFRAME in Section 2.3) we bake in framing under certain stable frames. The frames in question depend on the region index $t \in \text{RType}$. In particular, \sqsubseteq_t permits a single simultaneous update of multiple regions with region types not greater than or equal to t , each justified by a single action. Hence, we require that \sqsubseteq_t is closed under arbitrary frames that are stable under a single simultaneous update of multiple regions with region types not greater than or equal to t , each justified by a single action, i.e., $\hat{R}^{\{t' \mid t \preceq t'\}}$.

Support. In Section 2.2 we introduced specification logic assertions `indep` and `dep`, to internalize a notion of region type support in the logic, to allow explicit

proofs of the absence of self-referential region assertions. Their meaning is defined in terms of the following supp assertion, which asserts that p is supported by the set of region types $A \in \mathcal{P}(\text{RType})$. Formally, $\text{supp}_A(p)$ asserts that p is closed under arbitrary shared states that agree on all regions of type A ($s|_A = s'|_A$) and arbitrary action models that are A equivalent ($\varsigma \equiv_A \varsigma'$).

$$\text{supp}_A(p) \stackrel{\text{def}}{=} \{i \in \mathbf{N} \mid \forall j \leq i. \forall (j, (l, s, \varsigma)) \in p. \forall s'. \forall \varsigma'. \\ s|_A = s'|_A \wedge \varsigma \equiv_A \varsigma' \Rightarrow (j, (l, s', \varsigma')) \in p\}$$

Intuitively, two action models are considered A -equivalent if they agree on the regions of types in A (but they are allowed to differ on regions of types not in A). An assertion p is then dependent on region type $t \in \text{RType}$ if p is supported by the set of region types greater than or equal to t , and independent if it is supported by the set of region types not greater than or equal to t :

$$\text{dep}_t(p) \stackrel{\text{def}}{=} \text{supp}_{\{t' \mid t \leq t'\}}(p) \qquad \text{indep}_t(p) \stackrel{\text{def}}{=} \text{supp}_{\{t' \mid t \not\leq t'\}}(p)$$

Purity. To reason about state-independent protocols and nested view-shifts we have introduced several types of purity; namely, state, protocol and permission purity. Since our assertion logic is intuitionistic, we interpret purity as closure under arbitrary changes to the state, protocols, and permissions, respectively. For instance, $\text{pure}_{\text{prot}}(p) \stackrel{\text{def}}{=} \{i \in \mathbf{N} \mid \forall j \leq i. \forall (j, (l, s, \varsigma)) \in p. \forall \varsigma'. (j, (l, s, \varsigma')) \in p\}$.

Soundness. The main soundness theorem expresses that for any derivable Hoare triple, $\{p\}\bar{c}\{q\}$, if \bar{c} is executed with a local stack s as thread t , with a global heap h that is in the erasure of some instrumented state in $p(s)$, then, if t (and any threads t may have forked) terminates, then the terminal heap h' is in the erasure of some instrumented state in $q(s')$, where s' is the terminal stack of t .

Theorem 1. *If $\Gamma \vdash (\Delta).\{P\}\bar{c}\{Q\}$ then for all $\vartheta \in \llbracket \Gamma \rrbracket$, thread identifiers $t \in \text{TIId}$, stacks $s \in \llbracket \Delta \rrbracket$, and heaps $h \in \llbracket \llbracket \Gamma; \Delta \vdash P : \text{Prop} \rrbracket(\vartheta, s) \rrbracket$, if*

$$(h, \{(t, s, \bar{c})\}) \rightarrow (h', \{(t, s', \text{skip})\}) \uplus T'$$

and T' is irreducible then $h' \in \llbracket \llbracket \Gamma; \Delta \vdash Q : \text{Prop} \rrbracket(\vartheta, s') \rrbracket$.

5 Conclusion and Future Work

We have proposed a new style of specification for thread-safe data structures that allows the client to refine the specification with a usage protocol, in a concurrent setting. We have shown how to apply it to the bag and concurrent runner example. To realize this style of specification we have presented a new higher-order separation logic with Concurrent Abstract Predicates, that supports state-independent higher-order protocols and synchronization of multiple regions. We have also used the logic to specify and verify Joins, a sophisticated library implemented using higher-order code and shared mutable state.

We have demonstrated that our logic and style of specification scales to implementations of fine-grained concurrent data structures without helping [8]. Future work includes investigating concurrent data structures that use helping.

References

1. Appel, A., Melliès, P.-A., Richards, C., Vouillon, J.: A very modal model of a modern, major, general type system. In: Proc. of POPL (2007)
2. Biering, B., Birkedal, L., Torp-Smith, N.: BI-Hyperdoctrines, Higher-order Separation Logic, and Abstraction. ACM TOPLAS (2007)
3. Birkedal, L., Møgelberg, R., Schwinghammer, J., Støvring, K.: First Steps in Synthetic Guarded Domain Theory: Step-Indexing in the Topos of Trees. In: Proc. of LICS (2011)
4. Dinsdale-Young, T., Birkedal, L., Gardner, P., Parkinson, M., Yang, H.: Views: Compositional Reasoning for Concurrent Programs. In: Proceedings of POPL (2013)
5. Dinsdale-Young, T., Dodds, M., Gardner, P., Parkinson, M.J., Vafeiadis, V.: Concurrent Abstract Predicates. In: D’Hondt, T. (ed.) ECOOP 2010. LNCS, vol. 6183, pp. 504–528. Springer, Heidelberg (2010)
6. Dodds, M., Jagannathan, S., Parkinson, M.J.: Modular reasoning for deterministic parallelism. In: Proceedings of POPL, pp. 259–270 (2011)
7. Filipović, I., O’Hearn, P., Rinetzky, N., Yang, H.: Abstraction for Concurrent Objects. In: Castagna, G. (ed.) ESOP 2009. LNCS, vol. 5502, pp. 252–266. Springer, Heidelberg (2009)
8. Herlihy, M., Shavit, N.: The Art of Multiprocessor Programming. Morgan Kaufmann (2008)
9. Herlihy, M.P., Wing, J.M.: Linearizability: a correctness condition for concurrent objects. ACM TOPLAS 12, 463–492 (1990)
10. Jacobs, B., Piessens, F.: Expressive modular fine-grained concurrency specification. In: Proceedings of POPL, pp. 271–282 (2011)
11. Jensen, J.B., Birkedal, L.: Fictional Separation Logic. In: Seidl, H. (ed.) ESOP 2012. LNCS, vol. 7211, pp. 377–396. Springer, Heidelberg (2012)
12. Lea, D.: A java fork/join framework. In: Proceedings of the ACM 2000 Conference on Java Grande, JAVA 2000, pp. 36–43. ACM (2000)
13. Owicki, S.S.: Axiomatic Proof Techniques for Parallel Programs. PhD thesis, Cornell (1975)
14. Parkinson, M., Bornat, R., O’Hearn, P.: Modular verification of a non-blocking stack. SIGPLAN Not. 42(1) (2007)
15. Pilkiewicz, A., Pottier, F.: The essence of monotonic state. In: Proceedings of TLDI, pp. 73–86 (2011)
16. Russo, C.V.: The Joins Concurrency Library. In: Hanus, M. (ed.) PADL 2007. LNCS, vol. 4354, pp. 260–274. Springer, Heidelberg (2007)
17. Schwinghammer, J., Birkedal, L., Reus, B., Yang, H.: Nested Hoare Triples and Frame Rules for Higher-Order Store. LMCS 7(3:21) (2011)
18. Svendsen, K., Birkedal, L., Parkinson, M.: Verifying Generics and Delegates. In: D’Hondt, T. (ed.) ECOOP 2010. LNCS, vol. 6183, pp. 175–199. Springer, Heidelberg (2010)
19. Svendsen, K., Birkedal, L., Parkinson, M.: Higher-order Concurrent Abstract Predicates. Technical report, IT University of Copenhagen (2012), <http://www.itu.dk/people/kasv/hocap-tr.pdf>
20. Turon, A., Thamsborg, J., Ahmed, A., Birkedal, L., Dreyer, D.: Logical Relations for Fine-Grained Concurrency. In: Proceedings of POPL (2013)