Verifying Custom Synchronisation Constructs
Using Higher-Order Separation Logic

MIKE DODDS, University of York, UK
SURESH JAGANNATHAN, Purdue University, Indiana
MATTHEW J. PARKINSON, Microsoft Research, UK
KASPER SVENDSEN, Aarhus University, Denmark
LARS BIRKEDAL, Aarhus University, Denmark

Synchronisation constructs lie at the heart of any reliable concurrent program. Many such constructs are standard – e.g., locks, queues, stacks, and hash-tables. However, many concurrent applications require custom synchronisation constructs with special-purpose behaviour. These constructs present a significant challenge for verification. Like standard constructs, they rely on subtle racy behaviour, but unlike standard constructs, they may not have well-understood abstract interfaces. As they are custom-built, such constructs are also far more likely to be unreliable.

This paper examines the formal specification and verification of custom synchronisation constructs. Our target is a library of channels used in automated parallelization to enforce sequential behaviour between program statements. Our high-level specification captures the conditions necessary for correct execution; these conditions reflect program dependencies necessary to ensure sequential behaviour. We connect the high-level specification with the low-level library implementation, to prove that a client's requirements are satisfied. Significantly, we can reason about program and library correctness without breaking abstraction boundaries.

To achieve this, we use a program logic called iCAP (impredicative Concurrent Abstract Predicates) based on separation logic. iCAP supports both high-level abstraction and low-level reasoning about races. We use this to show that our high-level channel specification abstracts three different, increasingly complex low-level implementations of the library. iCAP's support for higher-order reasoning lets us prove that sequential dependencies are respected, while iCAP's next-generation semantic model lets us avoid ugly problems with cyclic dependencies.

Categories and Subject Descriptors: D.2.4 [**Software Engineering**]: Program Verification—*Correctness proofs*; D.3.3 [**Programming Languages**]: Language Constructs and Features—*Concurrent programming structures*

General Terms: Languages, Theory, Verification

Additional Key Words and Phrases: Separation Logic, Concurrent Abstract Predicates, Concurrency

1. INTRODUCTION

Concurrent programming is challenging because it requires programmers to parcel work into useful units, and weave suitable concurrency control to coordinate access to shared data. Coordination is generally performed by *synchronisation constructs*. In order that programmers can build and reason about concurrent programs, it is essential that these synchronisation constructs hide implementation details behind specifications, allowing clients to reason about correctness in terms of abstract, rather than concrete, behaviour.

For standard synchronisation constructs – e.g. locks, queues, stacks – abstract specifications are well-studied. However, many concurrent applications depend on non-standard, custom synchronisation constructs. These may have poorly-defined abstract interfaces, while at the same time depending on complex racy implementation behaviour. Verifying these constructs requires a technique that can build up strong abstractions, reason about the logical distribution of data between threads, and at the same time deal with the intricacies of low-level concurrency. This is our objective in this paper.

Our target is to verify one such custom concurrency construct: barriers used for automated parallelisation. In *deterministic* parallelisation, code regions in a sequential program are executed concurrently. While the parallelized program is internally nondeterministic, control constructs are used to ensure that it exhibits the same deterministic observable behaviour as its sequential counterpart. Automatic parallelisation of this kind has been well-studied for loop-intensive numerical computations. However, it is also possible to extract parallelism from irregularly structured sequential programs [Bocchino et al. 2009; Rinard and Lam 1992; Welc et al. 2005].

One way to implement deterministic parallelism is through compiler-injected barriers [Navabi et al. 2008]. We can think of these barriers as enforcing the original sequential program dependencies on shared resources. A resource could be any program variable, data structure, memory region, lock, etc. for which resource ownership guarantees are essential in order to enforce deterministic semantics. While the intuition behind using such barriers is quite simple, there are many possible implementations, and verifying that an implementation enforces the correct behaviour is challenging for several reasons:

— *Custom data-structures.* To enable the maximum level of parallelism, barriers are implemented using custom data-structures that collect and summarise signals.
— *Non-local signalling.* The patterns of signalling in a barrier implementation are highly non-local. To access a resource, a barrier must wait until all logically preceding threads have indicated that it is safe to do so. However, threads are locally unaware of this context.
— *Out-of-order signalling.* The parallelisation process will strive to identify the earliest point in a thread's execution path from where a resource is no longer required. In some cases, this means threads can release resources without ever acquiring them, so that subsequent signalling of this resource by its predecessor can bypass it altogether.
— *Shared read access.* Barriers may treat reads and writes differently to ensure preservation of sequential behaviour. Although many reads can be performed concurrently, they must be sequentialized with respect to writes. Moreover, reads must be sequentialized with respect to other reads, if there is an intervening write.
— *Higher-order specifications.* Abstractly, channels can be used to control access to any kind of resource for which ownership is important. Thus, the natural specification is higher-order: the resource is a parameter to the specification. Channels may control access to other channels, or even later stages of the same channel.

In this paper, we show how to tackle these verification challenges. We use *impredicative concurrent abstract predicates* (iCAP), a recent program logic that enables abstract, higher-order reasoning about concurrent libraries [Svendsen and Birkedal 2014a]. This allows us to reason about both high-level properties and low-level implementation details. iCAP sup-

ports fine-grained reasoning about concurrent behaviour, meaning that each thread can be permitted exactly the behaviour it needs. Furthermore, reasoning in iCAP is *local*, meaning even shared state can be encapsulated and abstracted.

The result of our work is a verified high-level specification for barriers, independent of their low-level implementation. Using iCAP, we have proved that three very different low-level implementations satisfy the same high-level specification. In the presence of runtime thread creation and dynamic (heap-allocated) data, our specification must be both generic and dynamic, since it must be possible to construct barriers at runtime that control access to arbitrary resources. To allow this, we use iCAP's higher-order quantification mechanism to encode complex patterns of resource redistribution. It is worth emphasising that the barriers we look at were not designed with verification in mind; we have developed the specification to suit the application, not vice versa.

In this paper we focus on just the verification problem for barriers, but in a companion paper, we define a parallelising program analysis which injects appropriate barriers [Botinčan et al. 2013]. Our work here contributes to the eventual goal of a fully specified and verified system for deterministic parallelism. More generally, access to concurrent data is often controlled by custom synchronisation constructs, and our work in this paper demonstrates how to reason soundly about such bespoke concurrency constructs.

Contribution

This paper substantially revises and expands our conference paper [Dodds et al. 2011]. Our main contributions relative to [Dodds et al. 2011] are as follows:

— A revised higher-order abstract specification for the custom synchronisation barriers used in deterministic parallelism. Our new specification is cleaner and more general.
— New proofs of this specification for simplified, out-of-order and summarising barrier implementations, written using the iCAP proof system [Svendsen and Birkedal 2014a]. The first two implementations were proved correct in [Dodds et al. 2011], while the proof of the summarising version is entirely novel.
— A self-contained presentation of the formal iCAP proof system and a tutorial on how to use iCAP to verify concurrent software. The formal iCAP proof system is not presented nor explained in the iCAP conference paper [Svendsen and Birkedal 2014a].
— An encoding in iCAP of constructs we call *saved propositions*. These serve some of the functions of auxiliary variables capable of storing propositions, and allow us to reason about resource transfer and splitting without altering the proof system.
— A new application of Wickerson et al.'s *explicit stabilization* [Wickerson et al. 2010] to reason about the stability of complex separation logic assertions.
— A verified application of our barriers: a tree-based key-value store. This demonstrates a non-trivial dynamic pattern of signalling, as signals are activated at branches only if they are visited by a thread. It also demonstrates the need for higher-order reasoning: in the logical definition of a tree, the induction passes through a higher-order parameter.

This paper also corrects a subtle logical problem in [Dodds et al. 2011], discovered by Svendsen a year after publication. This arose as a result of a circularity in the model of higher-order propositions, which rendered several important reasoning steps unsound. The details are discussed in §8, but we emphasise that this problem could not have been solved by the higher-order separation logics available in 2011. The development of iCAP was in part motivated by resolving this kind of problem; in this paper, we show that iCAP can be used to verify tricky practical algorithms.

Paper Structure

§2 discusses related work. §3 introduces the behaviour of barriers informally, and defines our abstract specification. It also discusses an example application, a tree-based key-value store.

§4 gives a very simple barrier implementation, and shows how it can be verified with respect to the core of the specification. This section also serves as a tutorial introduction to iCAP, the logic we use for verification. §5 discusses how the specification can be extended to cover the splitting of resources offered by a channel. §6 gives a more complicated implementation where channels are arranged into chains, and verifies the full abstract specification. §7 gives an optimised implementation where signals between channels are summarised, and verifies it. §8 explores the problems with our conference paper [Dodds et al. 2011], and how we have addressed them. Some of the subsidiary lemmas are proved in full in the appendices.

2. RELATED WORK

iCAP is a new logic for verifying complicated concurrent algorithms [Svendsen and Birkedal 2014a; 2014b]. Although we have focussed in this paper on barriers used for deterministic parallelism [Welc et al. 2005; Berger et al. 2010; Bocchino et al. 2009; Navabi et al. 2008], our intention is to illustrate how iCAP can be used to specify and verify novel concurrency constructs in general.

Prior to 2011, most work on concurrent separation logic considered concurrency constructs as primitive in the logic. This begins with O'Hearn's work on concurrent separation logic [O'Hearn 2007], which takes statically allocated locks as a primitive. CSL has been extended to deal with dynamically-allocated locks [Gotsman et al. 2007; Hobor et al. 2008; Jacobs and Piessens 2009] and re-entrant locks [Haack et al. 2008]. Others have extended separation logic or similar logics with primitive channels [Hoare and O'Hearn 2008; Bell et al. 2009; Villard et al. 2010; Leino et al. 2010], and event driven programs [Krishnaswami et al. 2010]. There are important disadvantages to handling each distinct concurrency construct with a new custom logic:

— Developing a custom logic might be acceptable for standard synchronisation constructs such as locks, but it is infeasible for every domain-specific construct.
— Embedding each construct as primitive in the logic provides no means for verifying implementations of the construct.
— Each custom logic handles one fixed kind of construct, with no means of verifying programs that use multiple concurrency constructs.

iCAP solves all three problems. New synchronisation constructs can be introduced as libraries and given abstract specifications that abstract over the internal data representation and state through abstract predicates. Implementations can be verified against these abstract specifications by giving these predicates concrete definitions (our paper does precisely this for barriers). As new constructs can be freely introduced as libraries, clients are free to combine multiple concurrency constructs as needed. Furthermore, using iCAP's higher-order quantification, specifications can abstract over arbitrary predicates, including those defined by other concurrency constructs. This allows us to support separate reasoning about each construct, while still allowing them to interact cleanly. For instance, abstract lock predicates defined by a lock library can freely be transferred through our channels. (See [Dinsdale-Young et al. 2010] for an example of such a lock specification).

iCAP descends from our earlier Concurrent Abstract Predicates (CAP) logic [Dinsdale-Young et al. 2010]. CAP combined the explicit treatment of concurrent interference from rely-guarantee [Jones 1983; Feng et al. 2007; Vafeiadis 2007] and abstraction through abstract predicates [Parkinson and Bierman 2005], with a rich system of protocols based on capabilities [Dodds et al. 2009]. iCAP extends CAP with higher-order propositions and an improved system of concurrent protocols [Svendsen and Birkedal 2014a]. iCAP's step-indexed semantics is supported by an underlying theory called the *topos of trees* [Birkedal et al. 2012].

Recent years have seen a great deal of work on concurrent logics, many of which take inspiration from CAP. Complex concurrency constructs have been verified before in CAP-like

logics, e.g. concurrent B-trees in [da Rocha Pinto et al. 2011]. The proof in [da Rocha Pinto et al. 2011] is mostly concerned with complex manipulations of the B-tree structure. In comparison, our barrier implementations are relatively simple, and a large proportion of our proof concerns changes in *ownership* to support our higher-order specification. The verification of the Joins library in [Svendsen et al. 2013] has similarities to our work. Both papers deal with barriers using higher-order separation logic. However, the implementations and specifications are substantially different – for example our implementation permits chains of channels, and our specification deals with resource-splitting. The authors of [Svendsen et al. 2013] are also co-authors on this paper, and iCAP was largely developed as a improvement on the HOCAP logic introduced there.

Two significant alternative logics to iCAP are CaReSL [Turon et al. 2013] and TaDA [da Rocha Pinto et al. 2014]. Like iCAP, both extend CAP with richer protocols. Unlike iCAP, both are primarily aimed at proving atomicity / linearizability, and confine themselves to second-order logic only. This makes them less suitable for our purposes. It is plausible that many of the proofs in this paper could be recast into these logics. However, we would have to constrain the higher-order parameters from our specification with some kind of explicit stratification. We would expect proofs to be significantly more complex as a result of the bookkeeping needed to track this stratification.

Another logic aimed at fine-grained concurrent data-structures is FSCL [**?**]. This is defined through a shallow embedding into Coq's Calculus of Inductive Constructions, and thus supports definition of higher-order specifications. However, FCSLs reasoning principles for higher-order specifications are weaker than those of iCAP. In particular, FCSL lacks support for impredicative protocols and assertions in the heap, both of which result in circularities that would manifest as universe inconsistency errors in Coq. Thus, while FCSL can define our proposed barrier specification, we believe that FCSL would be unable to verify an implementation against it. Impredicative protocols and assertions in the heap are fundamental to the generic higher-order specifications that we require. For example, in §3.6, we apply our barrier specification in a key-value store: here the inductive definition of a tree passes through a higher-order parameter, representing the fact that subtrees must be acquired through channel communication.

3. A SPECIFICATION FOR DETERMINISTIC PARALLELISM

In this section, we describe the intuitive behaviour of a library of barriers for enforcing deterministic parallelism that forms our case study. Based on this, we define a high-level specification for barriers – the full abstract specification is given in §3.4. These barriers are based on the ones used for deterministic parallelism in [Navabi et al. 2008]. In [Botinčan et al. 2013] we use our abstract specification in a proof-based parallelizing analysis that is guaranteed to preserve sequential behaviour.

We assume that code sections believed to be amenable for parallelization have been identified, and the program split accordingly into threads. We assume a total logical ordering on threads, such that executing the threads serially in the logical order gives the same result as the original (unparallelized) program.

Barriers are associated with resources (e.g., program variables, data structures, etc.) that are to be shared between concurrently-executing program segments. There are two sorts of barriers. A *signal* barrier notifies logically later threads that the current thread will no longer use the resource. A *wait* barrier blocks until all logically prior threads have signalled that they will no longer use the resource (i.e., have issued signals).

We assume barriers are injected by an analysis which ensures that all salient data dependencies in the sequential program are respected. For example, suppose we run two instances of the function `f` in sequence (here `sleep(rand())` waits for an unknown period of time).

```
    void f(int *x, int *y, int v) {            *x = 0;
      if(*x < 10) {                            *y = 0;
        *y = *y + v;
        *x = *x + v;                           f(x,y,5);
        sleep(rand());                         f(x,y,11);
      } else {
        sleep(rand());
      }
    }
```

When this program terminates, location `x` and `y` will both hold 16.

The second call to `f` will wait for the first call to finish its arbitrarily long `sleep`, even though the first call will do nothing more once it wakes. An analysis could parallelize this function by passing control between the two earlier. The parallelized functions `f1` and `f2` are given below. We run both concurrently, but require that `f1` passes control of `x` and `y` to `f2` *before* sleeping, allowing `f2` to continue executing.

FUNCTION DEFINITIONS: PROGRAM BODY:

```
 f1(x,y,v,i) {              f2(x,y,v,i) {         *x = 0; *y = 0;
   if(*x < 10) {              wait(i);            chan *i = newchan();
     *y = *y + v;             if(*x < 10) {
     *x = *x + v;               *y = *y + v;      f1(x,y,5,i) || f2(x,y,11,i);
     signal(i);                 *x = *x + v;
     sleep(rand());             sleep(rand());
   } else {                   } else {
     signal(i);                 sleep(rand());
     sleep(rand());          }
   }                       }
 }
```

The barriers in `f1` and `f2` ensure that the two threads wait exactly until the resources they require can be safely modified, without violating sequential program dependencies. The correct ordering is enforced by barriers that communicate through a channel; in the example, `newchan` creates the channel `i`. Assuming the barriers are correctly implemented, the resulting behaviour is equivalent to the original sequential program, with `x` and `y` both holding 16.

3.1. Verifying a Client Program

How can we verify that our parallelized program based on `f1` and `f2` satisfies the same specification as the original sequential program? Typically (e.g. in [Navabi et al. 2008]) one would incorporate signalling machinery as part of a parallelization program analysis. Clients would then reason about program behaviour using the operational semantics of the barrier implementation. Validating the correctness of parallelization with respect to the sequential program semantics would therefore require a detailed knowledge of the barrier implementation. Any changes to the implementation could entail reproving the correctness of the parallelization analysis.

In contrast, we reason about program behaviour in terms of abstract specifications for `signal`, `wait` and `newchan`. Such an approach has the advantages that (1) implementors can modify their underlying implementation and be sure that relevant program properties are preserved by the implementation, and (2) client proofs (in this case, proofs involving compiler correctness) can be completed without knowledge of the underlying implementation.

We will reason about `f1` and `f2` using separation logic, which lets us precisely control the allocation of resources to threads over time. Assertions in separation logic denote resources: heap cells and data-structures, but also abstract resources like channel ends. For example, we write the following assertion to denote that `x` points to value v and `y` to value v': $x \mapsto v * y \mapsto v'$ The *separating conjunction* $*$ asserts that `x` and `y` are distinct. As well as capturing information about the current state of resources, assertions in separation logic also capture *ownership*. Thus the assertion $x \mapsto v * y \mapsto v'$ in an invariant for a thread implicitly states that the thread has exclusive access to `x` and `y`.

To reason about the parallel composition of threads, we can use the PAR rule of concurrent separation logic [O'Hearn 2007]:

$$\frac{\{P_1\}\, C_1\, \{Q_1\} \qquad \{P_2\}\, C_2\, \{Q_2\}}{\{P_1 * P_2\}\, C_1 \,\|\, C_2\, \{Q_1 * Q_2\}} \ \text{PAR}$$

To verify `f1` and `f2`, we must encode the fact that `f1` gives up access to `x` and `y` by calling `signal(i)`, while `f2` retrieves access to them by calling `wait(i)`. We encode these two facts with two predicates, `recv` and `send`, corresponding to the *promised resource*, the resource that can be acquired from logically earlier threads, and the *required resource*, the resource that must be supplied to logically later threads. We read these as follows:

$\mathsf{recv}(\mathtt{i}, P)$ –   By calling `wait` on `i`, the thread will acquire a resource satisfying the assertion P.

$\mathsf{send}(\mathtt{i}, P)$ –   By calling `signal` on `i` when holding a resource satisfying P, the thread will lose the resource P.

These predicates are *abstract*; each instantiation of the library will define them differently. The client only depends on an abstract specification that captures their intuitive meaning:

$$\{\mathsf{emp}\} \quad \mathtt{i\ =\ newchan()} \quad \{\mathsf{send}(\mathtt{i}, P) * \mathsf{recv}(\mathtt{i}, P)\}$$
$$\{\mathsf{send}(\mathtt{i}, P) * \mathsf{stable}(P) * P\} \quad \mathtt{signal(i)} \quad \{\mathsf{emp}\}$$
$$\{\mathsf{recv}(\mathtt{i}, P) * \mathsf{stable}(P)\} \quad \mathtt{wait(i)} \quad \{P\}$$

The assigned variable `i` stands for `newchan`'s return value – i.e. the address of the new channel. We also use this notation in the specification of `extend`, below.

This specification of `newchan` is implicitly universally quantified for all assertions P, meaning that we can construct a channel for *any* iCAP assertion. The same is true for other operation specifications: unless otherwise stated, predicates are universally quantified and thus can be chosen freely.

In general, the universally quantified P can be instantiated with assertions about shared resources. In this case we need to establish that these assertions are *stable*, i.e., invariant under changes potentially performed by other threads. This is expressed by the stability assertion, $\mathsf{stable}(P)$, in the preconditions of `signal` and `wait`. Stability will be explained in §4.1. For the moment, note that if P is a thread-local assertion, such as $x \mapsto v$, then P is trivially stable, as these assertions assert exclusive ownership of the underlying resource.

Note that we do not require stability in the specification of `newchan` – i.e. the proposition P in $\{\mathsf{send}(\mathtt{i}, P) * \mathsf{recv}(\mathtt{i}, P)\}$ need not be invariant. This gives us more freedom, because P can be be dynamically modified before a corresponding resource is supplied – see the sections below on splitting and renunciation. We only require that when a resource is eventually supplied or recieved, it is stable.

New `recv` and `send` predicates can be constructed at run-time using `newchan`, meaning we can construct an arbitrarily large number of channels for use in the program. Given these two predicates, we can give the following specifications for `f1` and `f2`. (Here we specialise

```
{x ↦ 0 * y ↦ 0 * send (i, x ↦ 5 * y ↦ 5)}            {recv (i, x ↦ 5 * y ↦ 5)}
f1(x,y,5,i) {                                         f2(x,y,11,i)
 if(*x < 10) {                                         wait(i); // Channel spec.
   *y = *y + 5; *x = *x + 5;                           {x ↦ 5 * y ↦ 5}
   {x ↦ 5 * y ↦ 5 * send (i, x ↦ 5 * y ↦ 5)}          if(*x < 10) {
   signal(i); // Channel spec.                           *y = *y + 11; *x = *x + 11;
   {emp}                                               {x ↦ 16 * y ↦ 16}
   sleep(rand());                                        sleep(rand());
 } else ... // Contradicts x<10.                       } else ... // Contradicts x<10.
}                                                     }
{emp}                                                 {x ↦ 16 * y ↦ 16}
```

Fig. 1. Proofs for f1 (left) and f2 (right).

to the particular parameter values of 5 / 11; it would be easy to generalise).

$$\{x \mapsto 0 * y \mapsto 0 * \mathsf{send}(i,\ x \mapsto 5 * y \mapsto 5)\} \quad \mathtt{f1(x,y,5,i)} \quad \{\mathsf{emp}\}$$
$$\{\mathsf{recv}(i,\ x \mapsto 5 * y \mapsto 5)\} \quad \mathtt{f2(x,y,11,i)} \quad \{x \mapsto 16 * y \mapsto 16\}$$

The send predicate in the specification for f1 says that the thread must supply the resources x and y such that they both contain the value 5. The specification for f2 says that the thread can receive x and y containing the value 5. Fig. 1 gives sketch-proofs for these two specifications.

Given this specification, the proof for the main program goes through as follows:

```
{x ↦ _ * y ↦ _}
 *x = 0; *y = 0; chan *i = newchan();
{x ↦ 0 * y ↦ 0 * send(i, x ↦ 5 * y ↦ 5) * recv(i, x ↦ 5 * y ↦ 5)}
{x ↦ 0 * y ↦ 0 * send(i, x ↦ 5 * y ↦ 5)} ‖ {recv(i, x ↦ 5 * y ↦ 5)}  ⎫
         f1(x,y,5,i)                      ‖       f2(x,y,11,i)         ⎬  PAR rule
            {emp}                         ‖    {x ↦ 16 * y ↦ 16}       ⎭  application.
{x ↦ 16 * y ↦ 16}
```

This proof establishes that the parallelized version of the program satisfies the same specification as the sequential original.

3.2. Splitting Waiters

It is often useful for several threads to receive resources via the same channel. This kind of sharing is safe as long as the promised resources are split disjointly. It would be unsafe for two threads to both gain access to x at the same time, but it is safe for one thread to access x while another accesses y. Consider the following three threads:

```
*x = *y + 1; ‖ wait(i); ‖ wait(i);
signal(i)    ‖ z = *x   ‖ *y = 4
```

The first thread signals on i to indicate that it has finished with both x and y. The other two threads both wait on this signal, and each use a different aspect of the promised resource.

To support splitting, we add a property to the specification allowing threads to divide promised resources:

$$\{\mathsf{recv}(a, P * Q) * \mathsf{stable}(P * Q)\} \quad \langle\texttt{skip}\rangle \quad \{\mathsf{recv}(a, P) * \mathsf{recv}(a, Q)\}$$

This axiom states that when a thread has been promised a resource that consists of two parts, access can be split between two threads, potentially before the resource is available. This is achieved by splitting a single promise for a resource consisting of two disjoint parts into two promises, one for each part.

Note that the splitting property is *not* a logical entailment – applying it requires an operational step, `skip`. This is because the property manipulates a shared higher-order resource: $\mathsf{recv}(a, P)$. To avoid problems with self-reference, iCAP requires that such manipulations correspond with operational steps – this anomaly is discussed when we introduce iCAP in §4.1. Thus we have to assume that every application of the splitting specification is associated with a `skip`. We discuss whether this assumption is justified in §4.1.

3.3. Chains and renunciation

To allow many threads to access related resources in sequence, we can construct a *chain* of channels. A `wait` barrier called on a channel waits for `signal` barriers on *all* preceding channels. We use the ordering in a chain to model the logical ordering between a sequence of parallelized threads. A chain initially consists of a singleton channel constructed using `newchan`. We introduce an operation `extend` which takes as its argument an existing channel, and creates a new channel immediately preceding it in the chain.

Connecting channels into chains creates a new opportunity for parallelism: the ability to *renounce* access to a resource without acquiring it first. In the simple specification given above, a thread can only call `signal` if it has acquired the required resource from its predecessors. However, this is often unnecessary. For example, consider the following three threads:

```
                ||  if (b != 0) {   ||
 *x = *x + a;   ||    wait(iEnd);   ||  wait(jEnd);
 signal(iStart) ||    *x = *x + b;  ||  r = *x
                ||  }               ||
                ||  signal(jStart); ||
```

Here we give the start and end points of channels distinct names: `iStart` and `iEnd` are the start and end points of channel `i` (resp. `j`). The channels `i` and `j` are arranged in a chain. The second thread waits on `iEnd` only if it needs to access `x`. Otherwise it signals immediately, even if the first thread has not signalled. Without renunciation the thread would have to insert a wait confirming that the first thread had signalled.

Chains. To support chains in our specification, we introduce an order predicate '$\_ \prec \_$' which represents the order between links in the chain. $x \prec y$ asserts that the channel x is earlier in the chain than channel y. We use two axioms about the ordering of channels:

$$x \prec y \implies x \prec y * x \prec y \qquad \textit{(duplication)}$$
$$x \prec y * y \prec z \implies x \prec z \qquad \textit{(transitivity)}$$

The precondition for the three-thread example above would therefore be as follows. The predicate `iEnd` $\prec$ `jStart` represents the relationship between the two channels `i`/`j`:

$$\left\{ \begin{array}{l} \mathsf{send}(\texttt{iStart}, x \mapsto \_) * \mathsf{recv}(\texttt{iEnd}, x \mapsto \_) * \\ \mathsf{send}(\texttt{jStart}, x \mapsto \_) * \mathsf{recv}(\texttt{jEnd}, x \mapsto \_) * \texttt{iEnd} \prec \texttt{jStart} \end{array} \right\}$$

The abstract specification of `extend` takes a `send` predicate and a set of order predicates about earlier channels E, and a set of order predicates about later channels L. The function returns a pair of channels (a,b), that are later than all the channels before x and before all the channels after x, and a is before b in the chain. Note this means a single channel's start and end-point may have different identifiers. It also creates `recv`, `send`, and order predicates representing the new channel.

$$\left\{\begin{array}{l}\mathsf{send}(\mathtt{x},P)\, * \\ \text{\Large$\circledast$}_{e\in E}\, e \prec \mathtt{x} * \text{\Large$\circledast$}_{l\in L}\, \mathtt{x} \prec l\end{array}\right\} \quad (\mathtt{a,b})\mathtt{=extend(x)} \quad \left\{\begin{array}{l}\mathsf{send}(\mathtt{a},Q) * \mathsf{recv}(\mathtt{a},Q) * \mathsf{send}(\mathtt{b},P) \\ * \,\mathtt{a} \prec \mathtt{b} * \text{\Large$\circledast$}_{e\in E}\, e \prec \mathtt{a} * \text{\Large$\circledast$}_{l\in L}\, \mathtt{b} \prec l\end{array}\right\}$$

Using this, we can construct the precondition for the three-thread example as follows:

$$\big\{\mathsf{emp}\big\}$$
```
jEnd = newchan();
```
$$\big\{\mathsf{send}(\mathtt{jEnd}, x \mapsto \_) * \mathsf{recv}(\mathtt{jEnd}, x \mapsto \_)\big\}$$
```
(iEnd, jStart) = extend(jEnd);
iStart = iEnd;
```
$$\left\{\begin{array}{l}\mathsf{send}(\mathtt{iStart}, x \mapsto \_) * \mathsf{recv}(\mathtt{iEnd}, x \mapsto \_) * \\ \mathsf{send}(\mathtt{jStart}, x \mapsto \_) * \mathsf{recv}(\mathtt{jEnd}, x \mapsto \_) * \mathtt{iEnd} \prec \mathtt{jStart}\end{array}\right\}$$

Renunciation. To support renunciation, we add an axiom allowing threads to satisfy required resources using earlier promised resources:

$$\{\mathsf{recv}(x,P) * \mathsf{send}(y, P * Q) * x \prec y\} \quad \langle\mathtt{skip}\rangle \quad \{\mathsf{send}(y,Q)\}$$

By giving up the ability to acquire the P resource on the x channel, we can forward the P resource to partially discharge our send obligation on a subsequent channel y. If the initial send obligation on y requires us to supply a resource with two disjoint parts P and Q, after renouncing P to y, the obligation on y reduces to Q.

In the three-thread example, we can use this specification to justify signalling on `jStart` without waiting on `iEnd`:

$$\big\{\mathsf{recv}(\mathtt{iEnd}, x \mapsto \_) * \mathsf{send}(\mathtt{jStart}, x \mapsto \_) * \mathtt{iEnd} \prec \mathtt{jStart}\big\}$$
```
// Renunciation axiom
```
$$\big\{\mathsf{send}(\mathtt{jStart}, \mathsf{emp})\big\}$$
```
signal(jStart); // signal specification
```
$$\big\{\mathsf{emp}\big\}$$

3.4. Full abstract specification

Figure 2 shows our client-facing abstract specification for deterministic parallelism. It introduces the extra predicates and axioms to support chains, renunciation and splitting.

Note that in §5 we define a more general specification which is more convenient when verifying channel implementations. Specifically, it uses *explicit stabilization* rather than `stable` assertions (see §4.1). The client-facing specification given in Figure 2 is less general, but also less complex, and sufficient to verify all our examples.

3.5. Adding Forward Extension

The specification in Fig. 2 allows chain extension 'backwards', by inserting a channel immediately before an existing `send`. However, in order to impose a sequential order, it is often useful to allows extension in the other direction, at the end of the chain. Fortunately, it is

SPECIFICATIONS:

$$\{\mathsf{emp}\} \quad \mathtt{i = newchan()} \quad \{\mathsf{recv}(\mathsf{i}, P) * \mathsf{send}(\mathsf{i}, P)\}$$

$$\{\mathsf{send}(\mathsf{i}, P) * \mathsf{stable}(P) * P\} \quad \mathtt{signal(i)} \quad \{\mathsf{emp}\}$$

$$\{\mathsf{recv}(\mathsf{i}, P) * \mathsf{stable}(\mathsf{P})\} \quad \mathtt{wait(i)} \quad \{P\}$$

$$\left\{\begin{matrix}\mathsf{send}(\mathtt{x}, P) * \\ \circledast_{e \in E}\, e \prec \mathtt{x} * \circledast_{l \in L}\, \mathtt{x} \prec l\end{matrix}\right\} \quad \mathtt{(b,a)=extend(x)} \quad \left\{\begin{matrix}\mathsf{send}(\mathtt{b}, Q) * \mathsf{recv}(\mathtt{b}, Q) * \mathsf{send}(\mathtt{a}, P) \\ * \mathtt{b} \prec \mathtt{a} * \circledast_{e \in E}\, e \prec \mathtt{b} * \circledast_{l \in L}\, \mathtt{a} \prec l\end{matrix}\right\}$$

AXIOMS:

$$x \prec y \quad \Longrightarrow \quad x \prec y * x \prec y$$

$$x \prec y * y \prec z \quad \Longrightarrow \quad x \prec z$$

$$\{\mathsf{recv}(x, P) * \mathsf{send}(y, P * Q) * x \prec y\} \quad \langle\mathtt{skip}\rangle \quad \{\mathsf{send}(y, Q)\}$$

$$\{\mathsf{recv}(a, P * Q) * \mathsf{stable}(P * Q)\} \quad \langle\mathtt{skip}\rangle \quad \{\mathsf{recv}(a, P) * \mathsf{recv}(a, Q)\}$$

Fig. 2. Abstract specification for deterministic parallelism.

simple to implement a wrapper library with this behaviour on top of our abstract specification. We use an object called seqChan to represent the end of the chain, with specification as follows:

$$\{\mathsf{seqChan}(\mathtt{x}, P)\} \quad \mathtt{cw,cs = extendSC(x)} \quad \{\mathsf{recv}(\mathtt{cw}, P) * \mathsf{send}(\mathtt{cs}, Q) * \mathtt{cw} \prec \mathtt{cs} * \mathsf{seqChan}(\mathtt{x}, Q)\}$$

$$\{P\} \quad \mathtt{sc = newSeqChan()} \quad \{\mathsf{seqChan}(\mathtt{sc}, P)\}$$

To implement a seqChan object, we store a pair of channels. The first, cWait, is the recv point for the resource P. The second, cEnd is a dummy send channel used to allow extension. In order to extend, the wrapper library calls extend(cEnd) – this yields in a new channel that can replace cWait. The seqChan object is thus defined as follows:

$$\mathsf{seqChan}(\mathtt{x}, P) \triangleq \exists cw, ce.\, \mathtt{x.cWait} \mapsto cw * \mathtt{x.cEnd} \mapsto ce * \mathsf{recv}(cw, P) * \mathsf{send}(ce, \_) * cw \prec ce$$

The details of the implementation and verification of the wrapper library are given in Fig. 3. Note that all the reasoning in this proof uses the abstract specification in Fig. 2 – we can use this specification to build and verify richer families of synchronisation constructs.

3.6. Example: Concurrent Key-value Store

We will now use our channels to implement a simple tree-based key-value store. We will focus on two methods: an add/update method which inserts a key-value pair, and a reverse lookup method which returns the key associated with a particular value, if it exists. These methods are called sequentially, but our implementation permit concurrent updates and lookups internally within the data-structure. The methods have the following specification:

$$\{\mathsf{Tree}(\mathtt{x}, m)\} \quad \mathtt{add(x, k, v);} \quad \{\mathsf{Tree}(\mathtt{x}, m[\mathtt{k} \mapsto \mathtt{v}])\}$$

$$\{\mathsf{Tree}(\mathtt{x}, m)\} \quad \mathtt{k = revlookup(x, v);} \quad \{\mathsf{Tree}(\mathtt{x}, m) * (m(\mathtt{k}) = \mathtt{v} \lor (\mathtt{k} = -1 \land \mathtt{v} \notin img(m)))\}$$

The second parameter of the Tree predicate represents the contents of the key-value store, i.e. it is a partial function $m \colon \mathsf{Key} \rightharpoonup \mathsf{Val}$.

This specification requires that methods are called sequentially, and indeed makes it appear that the each method runs entirely sequentially. For example, the postcondition of add(x,k,v) is a Tree updated so that the store maps key k to value v. However, internally the add operation forks a new thread to apply the update, then immediately returns. In other words, internally add and revlookup operations run concurrently, but our channels enforce the illusion of sequential access.

```
 1  struct seqChan {
 2    chan * cWait;
 3    chan * cEnd;
 4  }
 5
 6  seqChan * newSeqChan() {
 7    chan c1,c2;
```
$\{P\}$
```
 9    seqChan * sc = new(seqChan);
```
$\{\mathsf{sc.cWait} \mapsto \_ * \mathsf{sc.cEnd} \mapsto \_ * P\}$
```
11    c2 = newchan();
```
$\{\mathsf{sc.cWait} \mapsto \_ * \mathsf{sc.cEnd} \mapsto \_ * P * \mathsf{recv}(\mathsf{c2},Q) * \mathsf{send}(\mathsf{c2},Q)\}$
```
13    c1,c2 = extend(c2);
```
$\{\mathsf{sc.cWait} \mapsto \_ * \mathsf{sc.cEnd} \mapsto \_ * P * \mathsf{recv}(\_,Q) * \mathsf{send}(\mathsf{c1},P) * \mathsf{recv}(\mathsf{c1},P) * \mathsf{send}(\mathsf{c2},Q) * \mathsf{c1} \prec \mathsf{c2}\}$
```
15    signal(c1);
```
$\{\mathsf{sc.cWait} \mapsto \_ * \mathsf{sc.cEnd} \mapsto \_ * \mathsf{recv}(\_,Q) * \mathsf{recv}(\mathsf{c1},P) * \mathsf{send}(\mathsf{c2},Q) * \mathsf{c1} \prec \mathsf{c2}\}$
```
17    sc->cWait = c1;
18    sc->cEnd = c2;
```
$\{\mathsf{sc.cWait} \mapsto \mathsf{c1} * \mathsf{sc.cEnd} \mapsto \mathsf{c2} * \mathsf{recv}(\_,Q) * \mathsf{recv}(\mathsf{c1},P) * \mathsf{send}(\mathsf{c2},Q) * \mathsf{c1} \prec \mathsf{c2}\}$
```
20    // Drop useless recv predicate for end of the channel
```
$\{\mathsf{seqChan}(\mathsf{sc},P)\}$
```
22    return sc;
23  }
24
25  (chan *,chan *) extendSC(seqChan* sc) {
26    chan ce,cs,cw;
```
$\{\mathsf{seqChan}(\mathsf{sc},P)\}$

$\{\exists cw, ce.\, \mathsf{sc.cWait} \mapsto cw * \mathsf{sc.cEnd} \mapsto ce * \mathsf{recv}(cw,P) * \mathsf{send}(ce,\_) * cw \prec ce\}$
```
29    cw = sc->cWait;
```
$\{\exists ce.\, \mathsf{sc.cWait} \mapsto \mathsf{cw} * \mathsf{sc.cEnd} \mapsto ce * \mathsf{recv}(\mathsf{cw},P) * \mathsf{send}(ce,\_) * \mathsf{cw} \prec ce\}$
```
31    cs,ce = extend(sc->cEnd);
```
$\left\{ \begin{array}{l} \mathsf{sc.cWait} \mapsto \mathsf{cw} * \mathsf{sc.cEnd} \mapsto \_ * \mathsf{recv}(\mathsf{cw},P) * \mathsf{send}(\mathsf{cs},Q) * \mathsf{recv}(\mathsf{cs},Q) \\ * \mathsf{send}(\mathsf{ce},\_) * \mathsf{cw} \prec \mathsf{cs} * \mathsf{cs} \prec \mathsf{ce} \end{array} \right\}$
```
33    sc->cWait = cs;
34    sc->cEnd = ce;
```
$\left\{ \begin{array}{l} \mathsf{sc.cWait} \mapsto \mathsf{cs} * \mathsf{sc.cEnd} \mapsto \mathsf{ce} * \mathsf{recv}(\mathsf{cw},P) * \mathsf{send}(\mathsf{cs},Q) * \mathsf{recv}(\mathsf{cs},Q) \\ * \mathsf{send}(\mathsf{ce},\_) * \mathsf{cw} \prec \mathsf{cs} * \mathsf{cs} \prec \mathsf{ce} \end{array} \right\}$

$\{\mathsf{recv}(\mathsf{cw},P) * \mathsf{send}(\mathsf{cs},Q) * \mathsf{cw} \prec \mathsf{cs} * \mathsf{seqChan}(\mathsf{sc},Q)\}$
```
37    return (cw,cs);
38  }
```

Fig. 3. Source code and proof outlines for seq channels

Algorithm. Keys and values are stored in a standard binary tree data-structure, with left and right subtrees partitioned according to the value stored at a given node. Each subtree stores a seqChan in field `seqc`, which is used to sequentialise concurrent access to it.

```
struct Node {                    Tree * newTree() {
  int key;                         Tree * t = new(Tree);
  int val;                         t->node = null;
  Tree * left;                     t->seqc = newSeqChan();
  Tree * right;                   }
};
                                 Node * mkNode(int k, int v) {
struct Tree {                      Node * n = new(Node);
  Node * node;                     n->val = v;
  seqChan * seqc;                  n->key = k;
}                                  n->left = newTree();
                                   n->right = newTree();
                                 }
```

The two main methods are defined as follows. Both call `extendSC` to register the current operation in the chain, then call a concurrent helper method which uses the channels to synchronise. In the case of `addNode`, the helper method is forked into another thread. These helper methods, `addNode` and `revlookupNode`, are defined in Fig. 4.

```
add(Tree* t, int k, int v) {
  cw,cs = extendSC(t->seqc);
  fork(addNode(&(t->node),k,v,cw,cs));
}

int revlookup(Tree* t, int v) {
  cw,cs = extendSC(t->seqc);
  k = revlookupNode(&(t->node), v, cw, cs);
}
```

Channel operations are highlighted red in Fig. 4. Both methods can be understood sequentially, if these operations are ignored[1]:

— `addNode()` searches from the root for the position of the chosen key – the order on nodes allows it to locate a unique position. On finding the appropriate subtree, it either creates a new node if one is absent, or updates the value if a node exists.
— `revLookupNode()` searches the entire tree for a key associated with the value. It checks the current node, and then calls itself recursively on the left and right subtrees. Finding an appropriate value causes the function to return and the search to end.

The channel operations enforce sequential access for both operations when searching down the tree. Before accessing a subtree, an operation must register its presence using `extendSC`, and then wait for all preceding threads using `wait`.

For example, `addNode` first waits for any preceding operations accessing this subtree (line 4). After taking a step left or right down the tree, the method registers its presence in the new subtree by extending the chain (line 19). It then signals it has left the parent node, allowing other operations to access it (line 20).

---

[1]We might imagine the program being parallelised by injecting these constructs, either by an expert or a sufficiently clever automated analysis.

```
1  addNode(Tree * curr, int k, int v, chan *cw, chan *cs) {
2    chan c1,c2;
3    while(true) {
4      wait(cw);
5      if(curr->node == null) {
6        Node * n = mkNode(k,v);
7        curr->node = n;
8        signal(cs);
9        return;
10     } else if (curr->node->key < k) {
11       curr = curr->node->left;
12     } else if (curr->node->key > k) {
13       curr = curr->node->right;
14     } else {
15       curr->node->val = v;
16       signal(cs);
17       return;
18     }
19     c1,c2 = extendSC(curr->seqc);
20     signal(cs);
21     cw = c1;
22     cs = c2;
23   }
24 }
25
26
27 int revlookupNode(Tree * curr, int v, chan cw, chan cs) {
28     wait(cw);
29     if(curr->node == null) {
30       signal(cs);
31       return -1;
32     } else {
33         if (curr->node->val == v) {
34           signal(cs);
35           return k;
36         }
37         cl = curr->left;
38         cr = curr->right;
39         c1,c2 = extendSC(cl->seqc);
40         c3,c4 = extendSC(cr->seqc);
41         signal(cs);
42         r = revlookupNode(cl, v, c1, c2)
43         if(r == -1) {
44             return revlookupNode(cr, v, c3, c4)
45         } else {
46             signal(c4); //Renunciation
47             return r;
48         }
49     }
50 }
```

Fig. 4.   Concurrent helper methods for the key-value store.

`revlookupNode` follows a similar pattern, but an operation registers its presence in both left and right subtrees, because it may need to search both (lines 22–23). If the value is located in the left subtree, then the right subtree will never be searched, and can be signalled without waiting (line 42) – in other words, this is an example of renunciation.

Note that the need to wait for a return value sequentialises calls to `revlookup`. However, this restriction could be lifted with a little more work: `revlookup` could return a channel identifier through which its result value could be passed. A client program could then initiate multiple calls to `revlookup` in sequence, and wait for all the results at once. Internally, these calls would run in parallel in precisely the same way as the current implementation. This would break the illusion of sequential execution, as clients would need to manage return-value channels. However, this specification could be verified with largely identical internal reasoning to the current version.

Verifying the key-value store. The predicate Tree representing the store data-structure is defined as follows:

$$\mathsf{Tree}(x, m) \triangleq \exists c.\, \mathsf{seqChan}(c, \exists r.\, x.\mathtt{node} \mapsto r * \mathsf{Node}(r, m)) * x.\mathtt{seqc} \mapsto c$$
$$\mathsf{Node}(x, m) \triangleq \exists l, r, k.\, x.\mathtt{key} \mapsto k * x.\mathtt{val} \mapsto v * x.\mathtt{left} \mapsto l * x.\mathtt{right} \mapsto r *$$
$$\mathsf{Tree}(l, m_l) * \mathsf{Tree}(r, m_r) * m_l; [k : v]; m_r = m * keys(m_l) < k < keys(m_r)$$
$$\mathsf{Node}(\mathtt{null}, m) \triangleq m = \emptyset$$

The Tree predicate consists of a pointer to a seqChan predicate, which promises the rest of the tree. Thus, by using `extendSC` and `wait`, a thread can gain access to the remainder of the tree. The Node predicate splits the map m between its key/value, and the left and right subtrees. Each subtree is represented recursively by a Tree predicate.

The key step in the proof lifts the specification of `extendSC` to the Tree predicate. The correctness of this step follows trivially from the definition of Tree.

$$\left\{\mathsf{Tree}(\mathtt{t}, m)\right\}$$
```
  cw,cs = extendSC(t->seqc);
```
$$\left\{\begin{aligned} &\mathsf{Tree}(\mathtt{t}, m') * \mathsf{recv}\,(\mathtt{cw}, \exists x.\, \mathtt{curr.node} \mapsto x * \mathsf{Node}(x, m)) \\ &* \mathsf{send}\,(\mathtt{cs}, \exists x.\, \mathtt{curr.node} \mapsto x * \mathsf{Node}(x, m')) * \mathtt{cw} \prec \mathtt{cs} \end{aligned}\right\}$$

In the precondition, the current thread is promised a tree containing map m. This proof step promises that m will then be updated to some other map m'. This m' can be chosen as needed to represent the correct update – the choice is entirely arbitrary. However, before the channel can be signalled, the channel specification requires that the map m' is established, which enforces the obligation to make a reasonable choice. This step also gives a handle to the old promise, `cw`, and a handle used to establish the new promise, `cs`.

This proof step is used in the `add` and `addNode` functions. For example, in `add`, the promised map is updated to $m[k : v]$ as follows. This guarantees a correctly-updated data-structure to the next thread accessing the store, but defers the actual data-structure modification.

$$\left\{\mathsf{Tree}(\mathtt{t}, m)\right\}$$
```
add(Tree* t, int k, int v) {
  cw,cs = extendSC(t->seqc);
```
$$\left\{\begin{aligned} &\mathsf{Tree}(\mathtt{t}, m[\mathtt{k} : \mathtt{v}]) * \mathsf{recv}\,(\mathtt{cw}, \exists x.\, \mathtt{t.node} \mapsto x * \mathsf{Node}(x, m)) \\ &* \mathsf{send}\,(\mathtt{cs}, \exists x.\, \mathtt{t.node} \mapsto x * \mathsf{Node}(x, m[\mathtt{k} : \mathtt{v}])) \end{aligned}\right\}$$
```
  fork(addNode(t, k, v, cw, cs));
}
```

$$\left\{ \mathsf{Tree}(\mathtt{t}, m[\mathtt{k}:\mathtt{v}]) \right\}$$

The proof outline for `addNode` can be found in Fig. 5. Each iteration of the main loop updates one subtree. The loop invariant (line 6) consists of a `recv` predicate for the previous version of the subtree, and a `send` predicate for the subtree updated with the key/value pair. Calling `wait` grants access to the previous version of the subtree. The algorithm then examines the node contents and branches on the results.

If the current node is `null`, then the map m must be empty. Therefore, `send` can be satisfied by just creating a single node containing the key/value pair (line 14). The case where the node already exists is similar.

In the recursive case, the algorithm first promises to update the tree by calling `extendSC` on the appropriate subtree (line 19). It then uses `signal` to indicate it has finished with the current node, and then recurses to the appropriate subtree.

More specifically, suppose the algorithm takes the left branch. In order to signal, it must first satisfy the `send` on `cs`, which requires a `Node` predicate with the key/value pair inserted. It is easy to prove that the map can be partitioned into subtrees less-than and greater-than the current key – these two maps are written m_l and m_r (line 22). As the required key must be added to the left subtree, the right subtree need not be updated. Calling `extendSC` promises to updates the left subtree with the key/value pair, which in turn satisfies the `Node` predicate required by `send`. This means the algorithm can call signal (line 23).

The proof structure for `revlookup` is similar to the one for `add`: by extending the chain with `extendSC`, the program gains the right to access the tree structure, while still permitting other threads to work on the tree at the same time.

$$\left\{ \mathsf{Tree}(\mathtt{t}, m) \right\}$$
```
int revlookup(Tree* t, int v) {
  cw,cs = extendSC(t->seqc);
```
$$\left\{ \begin{array}{l} \mathsf{Tree}(\mathtt{x}, m) * \mathsf{recv}\left(\mathtt{cw}, \exists x.\, \mathtt{t.node} \mapsto x * \mathsf{Node}(x, m)\right) \\ * \mathsf{send}\left(\mathtt{cs}, \exists x.\, \mathtt{t.node} \mapsto x * \mathsf{Node}(x, m)\right) \end{array} \right\}$$
```
  k = revlookupNode(&(t->node), v, cw, cs);
}
```
$$\left\{ \mathsf{Tree}(\mathtt{x}, m) * (m(\mathtt{k}) = \mathtt{v} \lor (\mathtt{k} = -1 \land \mathtt{v} \notin img(m))) \right\}$$

The proof outline for the `revlookupNode` helper function can be found in Figures 6 and 7. The broad structure of the reasoning is similar to addNode, albeit expressed inductively rather than iteratively. Each step down the tree retrieves a subtree by calling `wait`, and then releases the portion of the tree it has already searched using `signal`. Because subtrees are associated with channels, the algorithm only accesses the portion of the tree it needs, leaving the remainder of the tree available to other threads.

One significant step is the renunciation that takes place on line 40. Here the algorithm takes requirement to supply a subtree to `c4` and satisfies it using the `recv` predicate for `c3`. In other words, the resource is released without ever being acquired by the thread.

Summary. Thus, using our abstract specification for channels, we have verified the behaviour of the parallelised key-value store. Crucially, even though this program features many threads running at once, with complex communication between threads, each individual thread is able to reason locally, without dealing with other threads or the implementation of the barriers. Furthermore, we are able to verify a function specification which hides the existence of internal concurrency. In [Botinčan et al. 2013] we use the same abstract specification as the basis for a general parallelisation analysis.

```
1 {recv (cw, ∃x. curr.node ↦ x * Node(x, m)) * send (cs, ∃x. curr.node ↦ x * Node(x, m[k : v]))}
2 addNode(Tree * curr, int k, int v, chan *cw, chan *cs) {
3   Tree * t;
4   chan c1,c2;
5   while(true) {
6 {recv (cw, ∃x. curr.node ↦ x * Node(x, m)) * send (cs, ∃x. curr.node ↦ x * Node(x, m[k : v]))}
7     wait(cw);
8 {∃x. curr.node ↦ x * Node(x, m) * send (cs, ∃x. curr.node ↦ x * Node(x, m[k : v]))}
9     if(curr->node == null) {
10    {curr.node ↦ null * m = ∅ * send (cs, ∃x. curr.node ↦ x * Node(x, m[k : v]))}
11        Node * n = mkNode(k,v);
12        curr->node = n;
13    {∃x. curr.node ↦ x * Node(x, [k : v]) * send (cs, ∃x. curr.node ↦ x * Node(x, [k : v]))}
14        signal(cs);
15        return;
16    } else if (curr->node->key < k) {
17        curr = curr->node->left;
```

$$\left\{ \begin{array}{l} \exists c, x.\, c.\text{node} \mapsto x * \exists r, k', v'.\, x.\text{key} \mapsto k' * x.\text{val} \mapsto v' * x.\text{left} \mapsto \text{curr} * x.\text{right} \mapsto r \\ * \, \text{Tree}(\text{curr}, m_l) * \text{Tree}(r, m_r) * m_l; [k' : v']; m_r = m * keys(m_l) < k' < keys(m_r) \\ * \, \text{k} < k' * \text{send} (\text{cs}, \exists x.\, c.\text{node} \mapsto x * \text{Node}(x, m[\text{k} : \text{v}])) \end{array} \right\}$$

```
19        c1,c2 = extendSC(curr->seqc);
```

$$\left\{ \begin{array}{l} \exists c, x.\, c.\text{node} \mapsto x * \exists r, k', v'.\, x.\text{key} \mapsto k' * x.\text{val} \mapsto v' * x.\text{left} \mapsto \text{curr} * x.\text{right} \mapsto r \\ * \, \text{Tree}(\text{curr}, m_l[\text{k} : \text{v}]) * \text{Tree}(r, m_r) * m_l[\text{k} : \text{v}]; [k' : v']; m_r = m * keys(m_l) < k' < keys(m_r) \\ * \, \text{send} (\text{cs}, \exists x.\, c.\text{node} \mapsto x * \text{Node}(x, m[\text{k} : \text{v}])) \\ * \, \text{recv} (\text{c1}, \exists x.\, \text{curr.node} \mapsto x * \text{Node}(x, m_l)) * \text{send} (\text{c2}, \exists x.\, \text{curr.node} \mapsto x * \text{Node}(x, m_l[\text{k} : \text{v}])) \end{array} \right\}$$

```
21        // definition of Node predicate
```

$$\left\{ \begin{array}{l} \exists c, x.\, c.\text{node} \mapsto x * \text{Node}(x, m[\text{k} : \text{v}]) * \text{send} (\text{cs}, \exists x.\, c.\text{node} \mapsto x * \text{Node}(x, m[\text{k} : \text{v}])) \\ * \, \text{recv} (\text{c1}, \exists x.\, \text{curr.node} \mapsto x * \text{Node}(x, m_l)) * \text{send} (\text{c2}, \exists x.\, \text{curr.node} \mapsto x * \text{Node}(x, m_l[\text{k} : \text{v}])) \end{array} \right\}$$

```
23        signal(cs);
24 {recv (c1, ∃x. curr.node ↦ x * Node(x, m_l)) * send (c2, ∃x. curr.node ↦ x * Node(x, m_l[k : v]))}
25        cw = c1;
26        cs = c2;
27 {recv (cw, ∃x. curr.node ↦ x * Node(x, m_l)) * send (cs, ∃x. curr.node ↦ x * Node(x, m_l[k : v]))}
28    } else if (curr->node->key > k) {
29        // Similar to left side
30        ...
31    } else {
32        curr->node->val = v;
33        signal(cs);
34        return;
35    }
36   }
37 }
```

Fig. 5.   Proof outline for `addNode`.

```
1   {recv (cw, ∃x. curr.node ↦ x * Node(x, m)) * send (cs, ∃x. curr.node ↦ x * Node(x, m))}
2   int revlookupNode(Tree * curr, int v, chan cw, chan cs) {
3       wait(cw);
4   {∃x. curr.node ↦ x * Node(x, m) * send (cs, ∃x. curr.node ↦ x * Node(x, m))}
5       if(curr->node == null) {
6     {∃x. curr.node ↦ x * Node(x, m) * send (cs, ∃x. curr.node ↦ x * Node(x, m)) ∧ m = ∅}
7         signal(cs);
8     {v ∉ img(m)}
9         return -1;
10      } else {
11  {∃x. curr.node ↦ x * Node(x, m) * send (cs, ∃x. curr.node ↦ x * Node(x, m))}
12          if (curr->node->val == v) {
13            k = curr->node->key;
14        {∃x. curr.node ↦ x * Node(x, m) * send (cs, ∃x. curr.node ↦ x * Node(x, m)) ∧ m = [k : v]}
15              signal(cs);
16        {m(k) = v}
17              return k;
18          }
19  {∃x. curr.node ↦ x * Node(x, m) * send (cs, ∃x. curr.node ↦ x * Node(x, m))}
20          cl = curr->left;
21          cr = curr->right;
22          c1,c2 = extendSC(cl->seqc);
23          c3,c4 = extendSC(cr->seqc);
```

$$
24 \quad \left\{ \begin{array}{l} \exists x.\,\mathsf{curr.node} \mapsto x * \mathsf{Node}(x, m) * \mathsf{send}\,(\mathsf{cs}, \exists x.\,\mathsf{curr.node} \mapsto x * \mathsf{Node}(x, m)) \\ * \,\mathsf{recv}\,(\mathsf{c1}, \exists x.\,\mathsf{cl.node} \mapsto x * \mathsf{Node}(x, m_l)) * \mathsf{send}\,(\mathsf{c2}, \exists x.\,\mathsf{cl.node} \mapsto x * \mathsf{Node}(x, m_l)) \\ * \,\mathsf{recv}\,(\mathsf{c3}, \exists x.\,\mathsf{cr.node} \mapsto x * \mathsf{Node}(x, m_r)) * \mathsf{send}\,(\mathsf{c4}, \exists x.\,\mathsf{cr.node} \mapsto x * \mathsf{Node}(x, m_r)) \\ * \,\exists k', v'.\, m = m_l; [k' : v']; m_r * v \neq v' \end{array} \right\}
$$

```
25          signal(cs);
```

$$
26 \quad \left\{ \begin{array}{l} \mathsf{recv}\,(\mathsf{c1}, \exists x.\,\mathsf{cl.node} \mapsto x * \mathsf{Node}(x, m_l)) * \mathsf{send}\,(\mathsf{c2}, \exists x.\,\mathsf{cl.node} \mapsto x * \mathsf{Node}(x, m_l)) \\ * \,\mathsf{recv}\,(\mathsf{c3}, \exists x.\,\mathsf{cr.node} \mapsto x * \mathsf{Node}(x, m_r)) * \mathsf{send}\,(\mathsf{c4}, \exists x.\,\mathsf{cr.node} \mapsto x * \mathsf{Node}(x, m_r)) \\ * \,\exists k', v'.\, m = m_l; [k' : v']; m_r \end{array} \right\}
$$

```
27          r = revlookupNode(cl, v, c1, c2)
```

$$
28 \quad \left\{ \begin{array}{l} \mathsf{recv}\,(\mathsf{c3}, \exists x.\,\mathsf{cr.node} \mapsto x * \mathsf{Node}(x, m_r)) * \mathsf{send}\,(\mathsf{c4}, \exists x.\,\mathsf{cr.node} \mapsto x * \mathsf{Node}(x, m_r)) \\ * \,\exists k', v'.\, m = m_l; [k' : v']; m_r * v \neq v' * (m_l(\mathsf{r}) = \mathsf{v} \lor (\mathsf{v} \notin img(m_l) \land r = -1)) \end{array} \right\}
$$

```
29      }
30  }
```

Fig. 6.   Proof outline for `revlookupNode` (completed in Fig. 7)

```
31  ⎧ recv (c3, ∃x. cr.node ↦ x * Node(x, mr)) * send (c4, ∃x. cr.node ↦ x * Node(x, mr)) ⎫
    ⎩ * ∃k', v'. m = ml; [k' : v']; mr * v ≠ v' * (ml(r) = v ∨ (v ∉ img(ml) ∧ r = −1))      ⎭
32          if(r == -1) {
33  ⎧ recv (c3, ∃x. cr.node ↦ x * Node(x, mr)) * send (c4, ∃x. cr.node ↦ x * Node(x, mr)) ⎫
    ⎩ * ∃k', v'. m = ml; [k' : v']; mr * v ≠ v' * (ml(−1) = v ∨ (v ∉ img(ml)))             ⎭
34              r = revlookupNode(cr, v, c3, c4)
35  ⎧ ∃k', v'. m = ml; [k' : v']; mr * v ≠ v' * (ml(−1) = v ∨ (v ∉ img(ml)))  ⎫
    ⎩ * (mr(r) = v ∨ (v ∉ img(mr) ∧ r = −1))                                   ⎭
36  { (m(−1) = v ∨ (v ∉ img(m) ∧ r = −1)) ∨ (m(r) = v) }
37              return r;
38          } else {
39  ⎧ recv (c3, ∃x. cr.node ↦ x * Node(x, mr)) * send (c4, ∃x. cr.node ↦ x * Node(x, mr)) ⎫
    ⎩ * ∃k', v'. m = ml; [k' : v']; mr * v ≠ v' * ml(r) = v                                ⎭
40              <skip> // Renunciation
41  { send (c4, true) * ∃k', v'. m = ml; [k' : v']; mr * v ≠ v' * ml(r) = v }
42              signal(c4);
43  { ∃k', v'. m = ml; [k' : v']; mr * v ≠ v' * ml(r) = v }
44  { m(r) = v }
45              return r;
46          }
47      }
48  }
```

Fig. 7. Proof outline for `revlookupNode` (continued from Fig. 6)

4. PROOF STRATEGY

This section provides an intuitive introduction to our general proof approach, and to iCAP, the reasoning system our proofs are based on. We introduce iCAP using a tutorial-style presentation by verifying a simple channel implementation against a simplified barrier specification consisting of strengthened versions of the first three axioms of our abstract specification:

$$\{\text{emp}\}\ \text{i = newchan()}\ \{\text{recv}(\text{i}, P) * \text{send}(\text{i}, P)\}$$

$$\{\text{send}(\text{i}, P) * \text{stable}(P) * P\}\quad \text{signal(i)}\quad \{\text{emp}\}$$

$$\{\text{recv}(\text{i}, P) * \text{stable}(P)\}\quad \text{wait(i)}\quad \{P\}$$

By avoiding splitting, chain extension, and renunciation, we can illustrate the basic features of iCAP in a straightforward manner.

Figure 8 gives a simple barrier implementation. Each channel has a `flag` field representing the current state of the channel. Each send / recv pair is associated with one such structure in the heap. The `signal` simply sets the flag, while the `wait` loops until the flag is set. In §5 and §6 we reintroduce the necessary extra reasoning to verify our full abstract specification.

4.1. iCAP tutorial

Impredicative Concurrent Abstract Predicates (iCAP) is a separation logic variant intended for verifying concurrent higher-order programs [Svendsen and Birkedal 2014a; 2014b]. In this

```
struct chan {                        signal(chan *x) {
  int flag;                            x->flag := 1;
}                                    }

chan *newchan() {                    wait(chan *x) {
  chan *x := new(chan);                while(x->flag == 0)
  x->flag := 0;                          skip;
  return x;                          }
}
```

Fig. 8.   Implementation of the barrier library.

section we introduce the iCAP concepts and proof rules needed to understand the proofs in this paper. The full iCAP proof system is included in Appendix C.

We do not present iCAP's step-indexed semantics in this paper. Instead we reason entirely using iCAP's proof rules, and treat the semantics as a black box. Soundness of these proof rules is proved in the iCAP article [Svendsen and Birkedal 2014a] with some auxiliary lemmas proved in the associated Technical Report [Svendsen and Birkedal 2014b].

In this section, we employ a tutorial-style presentation, using the verification of the simplified barrier to motivate and illustrate concepts and proof rules as they are introduced. When presenting proof rules we will elide typing contexts and logical contexts.

Regions. To handle concurrency, iCAP extends separation logic with *regions* containing resources shared between threads. Conceptually, the state in iCAP consists of a local component and a finite number of shared regions, each governed by a *protocol*, as illustrated below.
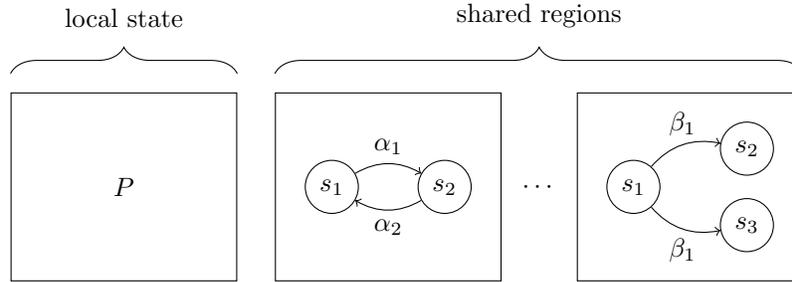
In the state illustrated above, we own the local resource P exclusively and can thus access it non-atomically. Resources owned by regions are shared between every thread and must therefore be accessed atomically and updated according to the protocol of the given region.

Protocols require that threads sharing a region behave as expected: for example, only releasing a lock once it has been acquired. They are expressed as a *labelled transition system* and an *interpretation function*. The labelled transition system defines the abstract states of the given region and describes how the abstract state may evolve, while the interpretation function maps each abstract state to a corresponding concrete heap resource. An assertion about a region has the following form:

$$\mathsf{region}(R, T, I, r)$$

In this assertion, R is the set of abstract states which the region could currently occupy. These possible states are taken from a larger set, fixed when the region is created. I maps from abstract states to invariants, also written in iCAP's assertion language. Intuitively,

$I(x)$ describes the resources contained in the region when it is in abstract state $x \in R$. To allow multiple distinct regions, r is a unique identifier for this region.

The remaining field, T, is a transition relation over abstract states, with transitions labelled with *actions*. T and I express the protocol that all threads must adhere to when accessing the region. Threads are only allowed to move a region from one abstract state to another if there exists a path in the labelled transition system T labelled with actions *permitted* to the thread. Permitted actions are tracked using *tokens*. These are linear objects created at the same time as a region. By issuing threads different tokens, we grant them different abilities over the shared region. For instance, the following assertion asserts ownership of a set and change token: $[\mathsf{set}]_i^{r_1} * [\mathsf{change}]_j^{r_2}$. Here r_1 and r_2 are the identifiers for the associated regions, while i and j are fractional parameters tracking how ownership of each token is shared.

A token with full permission ($i = 1$) asserts exclusive ownership of the action and thus ensures that no other thread can use the given action to change the abstract state of the shared region. Partial permission ($i < 1$) allows the owner to use the action, but does not exclude other currently executing threads from also using the action. Tokens can be split if the permission value is preserved: $[\alpha]_{i+j}^r \Leftrightarrow [\alpha]_i^r * [\alpha]_j^r$. In contrast, region assertions can be duplicated arbitrarily:

$$\mathsf{region}(R, T, I, r) \implies \mathsf{region}(R, T, I, r) * \mathsf{region}(R, T, I, r)$$

This allows arbitrarily many threads to access a shared region, but the ability to modify the region is restricted to threads holding the appropriate tokens.

Defining send *and* recv. To verify the implementation shown in Figure 8 with respect to the abstract specification we must first give concrete definitions to the abstract predicates send and recv.

The idea is to introduce a shared region for each channel, governing the internal state of that channel (the flag field) and ownership of the promised resource P. The shared region will have three possible abstract states: Low, High and Done. In the Low state the flag is low and the promised resource has yet to be provided. In the High state the flag is high and the promised resource has been sent, but not yet received and is thus conceptually owned by the channel. Lastly, in the Done state the flag is high and the promised resource has been sent and received. Each abstract state is associated with an invariant by the interpretation function I_{b}, defined as follows:

$$
\begin{aligned}
I_{\mathsf{b}}(x, P)(\mathsf{Low}) &\triangleq x.\mathrm{flag} \mapsto 0 \\
I_{\mathsf{b}}(x, P)(\mathsf{High}) &\triangleq x.\mathrm{flag} \mapsto 1 * \mathsf{stable}(P) * P \\
I_{\mathsf{b}}(x, P)(\mathsf{Done}) &\triangleq x.\mathrm{flag} \mapsto 1
\end{aligned}
$$

Here x is the location of the channel, while P is the resource controlled by the channel. Note that to move from the Low to the High abstract state the client must transfer ownership of P to the shared region, in addition to setting the flag. Likewise, to move from the High to the Done abstract state the client may transfer P out of the shared region and into its local state.

We want to ensure that only the sender (i.e., the owner of the send resource) is allowed to transition the abstract state from Low to High and only the receiver from High to Done. This will force the sender to transfer P to the shared region upon setting the flag and allow the receiver to take ownership of P once the flag has been set. In iCAP, we express this by labelling the Low to High transition with a set action and giving the sender exclusive ownership of the set action. The transition relation T_{b} thus has the following form. Each transition corresponds to an operation that can be performed on the channel.
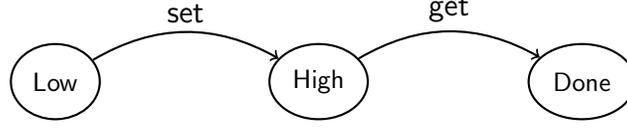
Using these definitions, we can give the predicates send and recv an interpretation:

$$\mathsf{send}(x, P) \triangleq \exists r.\ \mathsf{region}(\{\mathsf{Low}\}, T_\mathsf{b}, I_\mathsf{b}(x, P), r) * [\mathsf{set}]_1^r$$

$$\mathsf{recv}(x, P) \triangleq \exists r.\ \mathsf{region}(\{\mathsf{Low}, \mathsf{High}\}, T_\mathsf{b}, I_\mathsf{b}(x, P), r) * [\mathsf{get}]_1^r$$

The send predicate asserts that the abstract state is Low, since the promised resource has not been sent yet. On the other hand, recv asserts that the abstract state is either Low or High, but not Done, as the receiver does not know whether the promised resource has been sent yet, only that it has not been received yet. $[\mathsf{set}]_1^r$ and $[\mathsf{get}]_1^r$ are tokens allowing the thread to take particular transitions in T_b. The send predicate allows the sender to set the flag and supply the promised resource, while recv allows the receiver to retrieve the promised resource.

Stability. Region assertions allow us to describe our knowledge about the current abstract state of a region. However, since regions are shared, concurrently executing threads may update the abstract state of regions, invalidating our region assertions in the process. An assertion is said to be *stable* if it is invariant under possible interference from the environment. Since the concrete send resource defined above asserts exclusive ownership of the set action, the environment cannot use this transition to update the abstract state. Hence, if we own the send resource, the environment cannot invalidate our knowledge that the current abstract state is Low; thus, the send resource is stable.

More generally, a region assertion, $\mathsf{region}(R, T, I, r)$, is stable if R is closed under all transitions in T labelled with actions potentially owned by the environment. This is captured by the following iCAP proof rule.

$$(\forall \alpha \notin A.\ \forall x \in X.\ T(\alpha)(x) \subseteq X) \implies \mathsf{stable}(\mathsf{region}(X, T, I, r) * \circledast_{\alpha \in A} [\alpha]_1^r)$$

Since the assertion asserts exclusive ownership of all the action in A ($\circledast_{\alpha \in A} [\alpha]_1^r$), the environment cannot own any actions in A. Hence, the region assertion is stable if X is closed under any transitions α not in A, as expressed by the assumption.

Stability is closed under several of the usual connectives of higher-order separation logic ($\bot, \top, \vee, \wedge, \forall, \exists, =_\tau, \mathsf{emp}, *$), but generally not under implication and separating implication. To avoid reasoning about stability of separating implication, we use Wickerson et al's explicit stabilisation operators, $\lfloor - \rfloor$ and $\lceil - \rceil$, instead of requiring P to be stable [Wickerson et al. 2010]. We explicitly stabilise using $\lfloor P \rfloor$ and $\lceil P \rceil$ which are stable by construction for any assertion P. $\lfloor P \rfloor$ stands for the weakest assertion stronger than P that is stable and $\lceil P \rceil$ for the strongest assertion weaker than P that is stable. Thus, if P is already stable, explicitly stabilising P does nothing:

$$\mathsf{stable}(P) \implies (\lfloor P \rfloor \iff P \iff \lceil P \rceil)$$

However, in general, only the right implications hold: $\lfloor P \rfloor \Rightarrow P \Rightarrow \lceil P \rceil$. Explicit stabilisation operators are semi-distributive over separating conjunction:

$$\lfloor P \rfloor * \lfloor Q \rfloor \implies \lfloor P * Q \rfloor \qquad\qquad \lceil P * Q \rceil \implies \lceil P \rceil * \lceil Q \rceil$$

As a result, they are easy to move around in proofs.

The concrete recv resource defined above is thus also easily shown to be stable, as $\{\mathsf{Low}, \mathsf{High}\}$ is closed under the set transition, which is the only transition potentially owned by the environment.

When reasoning about non-atomic statements, iCAP requires that the pre- and postcondition is stable, to account for possible interference from the environment. We use standard Hoare triples, written $\{P\}\ C\ \{Q\}$, when reasoning about non-atomic statements and angled triples, written $\langle P\rangle\ C\ \langle Q\rangle$, when reasoning about atomic statements. For every standard Hoare triple, $\{P\}\ C\ \{Q\}$, we implicitly have to prove stability of P and Q. For the majority of the proof outlines in this article, these stability proofs are trivial and will be omitted. The following structural rule allows us to switch from atomic to non-atomic triples:

$$\mathsf{atomic}(C) \wedge \mathsf{stable}(P) \wedge \mathsf{stable}(Q) \wedge \langle P\rangle\ C\ \langle Q\rangle \implies \{P\}\ C\ \{Q\}$$

The predicate atomic holds for any command which is assumed to be atomic by iCAP: these are CAS, field read, field assignment and stack assignment.

Higher-order shared resources. The idea of shared resources that must satisfy an invariant or evolve following a protocol could equally be expressed in prior logics such as CAP or RGSep [Vafeiadis and Parkinson 2007; Dinsdale-Young et al. 2010]. The distinction with iCAP is that it is based on a higher-order separation logic and supports shared higher-order resources – i.e., shared regions containing shared resources. For instance, the send resource defined above is parametric in the promised resource P, which the client is free to instantiate with a shared resource – for example, another synchronisation construct.

It is well-known that reasoning about shared higher-order resources is difficult (for example, see the problems with our previous paper, discussed in §8). Intuitively, this is because the semantics of protocols is defined in terms of the semantics of assertions, but assertions are defined in terms of protocols. To avoid this problematic circularity, iCAP stratifies the construction of the semantic domain of protocols using step-indexing. As a result of this stratification, assertions and specifications are only given meaning with respect to a step-index.

In the case of self-referential protocols, step-indexing breaks the circularities by interpreting a region assertion at step-index $n + 1$ in terms of the semantics of the region interpretation at step-index n. For instance, a region assertion $\mathsf{region}(\{x\}, T, I, r)$ holds at step-index $n + 1$, if the shared resources owned by region r satisfies $I(x)$ at step-index n. To capture this syntactically, iCAP introduces a 'later' modality, written $\triangleright$. The assertion $\triangleright P$ holds at step-index $n + 1$, if P holds at step-index n. The region assertion $\mathsf{region}(\{x\}, T, I, r)$ thus expresses that the shared region r currently owns the resources described by $\triangleright I(x)$.

The semantics of Hoare triples connects the step-indexes of the pre- and postcondition with execution steps. Disregarding environment interference, a specification $\{P\}C\{Q\}$ is valid, if for any number of steps n and any initial state in P at step-index n, if C terminates in $k \leq n$ steps, then the terminal states satisfies Q at step-index $n - k$. The effects of this stratification and connection of step-indices with operational steps, is that reasoning involving shared resources require a corresponding step in the program execution. Intuitively, the logic must 'wait a step' before examining resources transferred out of shared regions, to prevent cycles in the reasoning. Formally, this is expressed by the frame rule for atomic commands (AFrame) given below, which allows the removal of a later from the frame. In particular, if $\triangleright P$ holds for the frame at step-index n then P holds for the frame at step-index $n - 1$. Thus, since the atomic command requires exactly one step of execution, P holds in the post-condition, as the post-condition is evaluated at step-index $n - 1$.

$$\mathsf{stable}(R) \wedge \langle P\rangle\ C\ \langle Q\rangle \implies \langle P * \triangleright R\rangle\ C\ \langle Q * R\rangle \qquad \text{(AFrame)}$$

One intuition for the step-indexed model is that an assertion P holds at step-index n if the assertion that it makes about the state cannot be invalidated with n or less steps of execution. Consequently, if an assertion P holds at step-index $n + 1$ it also holds at step-index n, since any assertion that can be invalidated in n or less steps, can also be invalidated in $n + 1$ or less steps. This property is expressed by the (SMono) rule given below, which

expresses that if P holds now, it also holds with one less step of execution left ($\triangleright P$).

$$P \implies \triangleright P \qquad\qquad (\text{SMono})$$

Using (SMono) and the rule of consequence, one can derive the standard CAP frame rule. The assertion $\triangleright P$ is stable if P is stable.

The effect of introducing $\triangleright$ is that accesses to shared resources in regions coincide with operational steps in the program. This breaks the circularity and gives iCAP a well-defined semantics.

However, it also means that splitting and renunciation must be associated with an explicit `skip` instruction to justify the transfer of shared resources in and out of regions. While these `skip` instructions are crucial to the well-definedness of protocols, they can typically be eliminated once we consider a whole program. In particular, for pre- and postconditions expressible in first-order separation logic, iCAP is adequate with respect to first-order separation logic [Svendsen and Birkedal 2014b, Theorem 1] and in first-order separation logic `skip` instructions can freely be eliminated. Hence, if P and Q are expressible in first-order separation logic and $\vdash_{iCAP} \{P\}\ C\ \{Q\}$, then $\vdash_{SL} \{P\}\ \tilde{C}\ \{Q\}$, where $\tilde{C}$ is C stripped of `skip` instructions.

The later operator distributes over conjunction, disjunction, separating conjunction, stabilization brackets. It semi-distributes over implication and separating implication, i.e:

$$\triangleright(P \Rightarrow Q) \implies (\triangleright P \Rightarrow \triangleright Q) \qquad\qquad \triangleright(P \mathbin{-\!\!*} Q) \implies (\triangleright P \mathbin{-\!\!*} \triangleright Q)$$

Later also distributes over existential and universal quantification over non-empty types. See Appendix C.2.2 for proof rules. In proof outlines we sometimes apply these properties silently.

Reasoning about shared regions. All statements that access resources owned by shared regions must be atomic and obey the protocols of the regions involved. This is achieved using structural rules that allow shared resources to be treated as local resources for the duration of an atomic statement. We refer to these rules as "region opening" rules and as "entering" and "exiting" a shared region, when applying these rules in proof outlines.

Conceptually, these rules require us to prove 1) that we own sufficient action permissions to justify any potential updates of the abstract state and 2) that we transfer the appropriate resources back to the shared region after the atomic statement. To illustrate, consider verifying the signal method of the simplified barrier implementation:

$$\{\mathsf{send}(\mathtt{i}, P) * \mathsf{stable}(P) * P\}\ \texttt{i->flag := 1}\ \{\mathsf{emp}\}$$

Since raising the flag is an atomic statement and both the pre- and postcondition is stable, we can switch to atomic triples. This frees us from reasoning about stability and possible interference during the execution of the atomic statement. After unfolding the send resource we are thus left with the following proof obligation:

$$\langle \mathsf{region}(\{\mathsf{Low}\}, T_\mathsf{b}, I_\mathsf{b}(\mathtt{x}, P), r) * [\mathsf{set}]_1^r * \mathsf{stable}(P) * P \rangle\ \texttt{x->flag := 1}\ \langle \mathsf{emp} \rangle$$

The precondition asserts that the flag field is owned by the shared region r. To update the field we are thus forced to open the shared region governing the given channel and thus to respect its protocol. Intuitively, upon raising the flag the abstract state changes from Low to High. We can thus strengthen the proof obligation:

$$\langle \mathsf{region}(\{\mathsf{Low}\}, T_\mathsf{b}, I_\mathsf{b}(\mathtt{x}, P), r) * [\mathsf{set}]_1^r * \mathsf{stable}(P) * P \rangle$$

$$\texttt{x->flag := 1}$$

$$\langle \mathsf{region}(\{\mathsf{High}\}, T_\mathsf{b}, I_\mathsf{b}(\mathtt{x}, P), r) \rangle$$

Note that this postcondition is not stable under `get` transitions. However, since we are reasoning about an atomic statement, the postcondition is not required to be stable.

To discharge the above proof obligation we first have to prove that we are allowed to update the abstract state from Low to High. Since the precondition asserts exclusive ownership of the set transition, this reduces to proving that there exists a set-labelled path from Low to High in T_b, which is true by definition. Secondly, we must prove that `x->flag := 1` does indeed transform the resources associated with abstract state Low to those associated with abstract state High, according to $I_b(x, P)$. We are thus left with the following proof obligation:

$$\langle \rhd I_b(x, P)(\mathsf{Low}) * [\mathsf{set}]_1^r * \mathsf{stable}(P) * P \rangle \ \mathtt{x\text{-}>flag \ := \ 1} \ \langle \rhd I_b(x, P)(\mathsf{High}) \rangle$$

Given local ownership of the shared resources for the abstract state Low we must transfer back resources corresponding to abstract state High after the execution of the atomic statement. Since protocols are implicitly interpreted one step later, the shared resources are only available one step later in the precondition and only have to be provided one step later in the postcondition. After unfolding $I_b(x, P)$ and applying SMono we are left with the following proof obligation:

$$\langle (\rhd \mathsf{x.flag} \mapsto 0) * [\mathsf{set}]_1^r * \mathsf{stable}(P) * P \rangle \ \mathtt{x\text{-}>flag \ := \ 1} \ \langle \mathsf{x.flag} \mapsto 1 * \mathsf{stable}(P) * P \rangle$$

Conceptually, this is provable because step-indexing only affects assertions that can generate problematic circularities in the domain. Primitive points-to assertions such as $x.f \mapsto v$ only affect addresses and values, and are thus independent of the step-indexing. This is captured by the structural (LPoints) rule given below, which allows us to remove a $\rhd$ from the pre-condition of a points-to assertion.

$$\langle x.f \mapsto y \rangle \ C \ \langle Q \rangle \implies \langle \rhd x.f \mapsto y \rangle \ C \ \langle Q \rangle \qquad \text{(LPoints)}$$

The general iCAP proof rule for accessing shared regions is given below:

$$\frac{\forall x \in X. \ (x, f(x)) \in \overline{T(A)} \qquad \forall x \in X. \ \langle P * \circledast_{\alpha \in A} [\alpha]_{g(\alpha)}^r * \rhd I(x) \rangle \ C \ \langle Q(x) * \rhd I(f(x)) \rangle^{\mathcal{E}}}{\langle P * \circledast_{\alpha \in A} [\alpha]_{g(\alpha)}^r * \mathsf{region}(X, T, I, r) \rangle \ C \ \langle \exists x. \ Q(x) * \mathsf{region}(\{f(x)\}, T, I, r) \rangle^{\mathcal{E} \uplus \{r\}}} \ \text{(AOpen)}$$

It generalizes the above example by considering a set of possible initial abstract states X and allowing the client to transfer ownership of resources in and out of the shared region using P and Q. To apply the rule we must define a function f that for every possible initial abstract state $x \in X$ defines the desired terminal abstract state $f(x)$. The first premise asserts that for every possible initial abstract state x there exists a path from x to $f(x)$ labelled with actions from the set A (we use $\overline{R}$ as notation for the reflexive, transitive closure of the relation R). Since the precondition asserts non-exclusive ownership of every action in A, this ensures that we are allowed to update the abstract state from x to $f(x)$ for every possible initial state $x \in X$. The g function, which records the fractional ownership of each action $\alpha \in A$, ensures that we own the same action fractions in the assumption and conclusion of the rule. We implicitly require that $g(\alpha) \neq 0$ for all $\alpha \in A$. The second premise ensures that C does indeed transform the resources associated with the abstract state x to $f(x)$ for every possible initial abstract state $x \in X$.

Since opening a shared region grants local ownership of the region's current resources, in general it would be unsound to have two nested openings of the same region in the proof tree, as this would duplicate the region's resources. To avoid this, we annotate the postcondition of atomic triples with a *region mask*, $\mathcal{E}$, of regions that may be opened. To open the region r, the above proof rule requires that r is in the region mask and removes it from the region mask in the premise, to ensure r is not opened again.

Verifying `signal` *and* `wait`. Figures 10 and 11 sketch proofs for `signal()` and `wait()` respectively. The proof of `wait` (Fig. 11) uses similar reasoning to `signal`. The main difference is that the region has two initial abstract states, Low and High, and that the thread

holds the token get, allowing it to transition from High to Done. We deal with the Low and High cases separately – see bottom left and right of Fig. 11. In the Low case, the resource has not been sent yet and we close the region in the Low state again. In the High case, the resource has been sent and we use the thread's get token to take ownership of P and close the region in the Done state.

iCAP is an intuitionistic separation logic, meaning it admits weakening, and in particular that $P*Q \implies Q$. We often use this property to dispose of unwanted predicates – intuitively, we can forget they exist. For example, we delete a redundant set token on line 11 of Figure 10.

View-shifts. The region opening rule given above allows a region to be opened for the duration of an atomic statement. It is also possible to open a shared region and close it immediately before the next statement is executed. This is for instance useful to abstractly describe transfer of resources in and out of shared regions. iCAP expresses such updates of the abstract state that do not affect the concrete state using the *view-shift* operator, $\sqsubseteq$.

For instance, the following view-shift expresses that if the channel is in the abstract state High and we own the get action, then we can change the abstract state to Done and take ownership of $\triangleright P$.

$$\mathsf{region}(\{\mathsf{High}\}, T_{\mathsf{b}}, I_{\mathsf{b}}(x, P), r) * [\mathsf{get}]_1^r \sqsubseteq \mathsf{region}(\{\mathsf{Done}\}, T_{\mathsf{b}}, I_{\mathsf{b}}(x, P), r) * [\mathsf{get}]_1^r * \triangleright P$$

View-shifts generalize standard implication and satisfy a generalized rule of consequence:

$$P_1 \sqsubseteq P_2 \wedge \{P_2\}\ C\ \{Q_2\} \wedge Q_2 \sqsubseteq Q_1 \implies \{P_1\}\ C\ \{Q_1\} \qquad \text{(ACONSQ)}$$

We can thus use view-shifts to factor manipulations of shared resources that do not affect the concrete state. The view-shift region opening rule is very similar to the region opening rule for atomic statements and also requires proving that any updates of the abstract state are permitted and that resources are transformed correctly:

$$\frac{\forall x \in X.\ (x, f(x)) \in \overline{T(A)} \qquad \forall x \in X.\ P * \circledast_{\alpha \in A}[\alpha]_{g(\alpha)}^r * \triangleright I(x)\rangle \sqsubseteq^{\mathcal{E}} Q * \triangleright I(f(x))}{P * \circledast_{\alpha \in A}[\alpha]_{g(\alpha)}^r * \mathsf{region}(X, T, I, r)\rangle \sqsubseteq^{\mathcal{E} \uplus \{r\}} Q * \mathsf{region}(\{f(x)\}, T, I, r)} \text{ (VOPEN)}$$

Just as for atomic commands, we need to ensure that already opened regions cannot be opened again. We thus index view-shifts with a region mask annotation ($\mathcal{E}$) that describes the set of regions that may be opened.

Allocation of shared regions is also described using view shifts. To allocate a shared region in initial abstract state x with interpretation map I we must transfer the resource $I(x)$ to the shared region. This is expressed by the following view-shift axiom:

$$(\forall x.\ \mathsf{stable}(I(x))) \implies I(x) \sqsubseteq \exists r.\ \mathsf{region}(\{x\}, T, I, r) * \circledast_{\alpha \in A}[\alpha]_1^r \qquad \text{(VALLOC)}$$

The resources owned by a shared region must always be stable. This is enforced upon allocation of the region, as expressed by the above axiom. Upon allocating a shared region, we can take exclusive ownership of any set of actions A on r.

While iCAP supports recursively-defined higher-order shared resources, soundness of iCAP depends crucially on the fact that the transition systems associated with each region are *not* recursively-defined. This is enforced when allocating a new region. Formally, the VALLOC rule includes a side-condition on the labelled transition system, T, which is given as relation on abstract states indexed by an action identifier, to enforce this. This side-condition is trivially satisfied by every labelled transition system expressible in first-order logic with equality. Since this is the case for all the transition systems employed in this article, we will not go into details about this technical side-condition. We refer the reader to the iCAP article for details [Svendsen and Birkedal 2014a].

Verifying `newchan`. Figure 9 shows a sketch proof of `newchan`. First the concrete channel data structure is allocated. This allows ownership of the `flag` field to be transferred to a newly allocated channel region, using the VALLOC and ACONSQ rule.

```
1  {emp}
2  newchan() {
3    chan *x := new(chan);
4    x->flag := 0;
5    {x.flag ↦ 0}
6    // introduce a later using SMono
7    ⟨▷I_b(x, P)(Low)⟩
8    // allocate new channel region in the Low state
9    ⟨∃r. region({Low}, T_b, I_b(x, P), r) * [set]_1^r * [get]_1^r⟩
10   // duplicate and weaken region assertion
11   ⟨∃r. region({Low}, T_b, I_b(x, P), r) * [set]_1^r * region({Low, High}, T_b, I_b(x, P), r) * [get]_1^r⟩
12   // definition of recv and send predicates
13   {recv(x, P) * send(x, P)}
14   return x;
15 }
16 {recv(ret, P) * send(ret, P)}
```

Fig. 9.   Proof of `newchan()` using simple predicate definitions.

```
1  {send(x, P) * stable(P) * P}
2  signal(chan *x) {
3    // definition of send.
4    ⟨region({Low}, T_b, I_b(x, P), r) * [set]_1^r * stable(P) * P⟩
5    // enter the region.
6    ⟨▷(x.flag ↦ 0) * [set]_1^r * stable(P) * P⟩
7    // drop ▷ using LPoints
8    ⟨x.flag ↦ 0 * [set]_1^r * stable(P) * P⟩
9    x->flag := 1;
10   ⟨x.flag ↦ 1 * [set]_1^r * stable(P) * P⟩
11   // Close region, perform set transition. Delete set token using weakening.
12   ⟨region({High}, T_b, I_b(x, P), r)⟩
13   // stabilise the region's abstract state.
14   {region({High, Done}, T_b, I_b(x, P), r)}
15 }
16 {emp}
```

Fig. 10.   Proof of `signal()` using simple predicate definitions.

PROOF BODY:

1 $\left\{\mathsf{recv}(\mathbf{x}, P) * \mathsf{stable}(P)\right\}$

2 `wait(chan *x) {`

3 `    int b;`

4 `    do {`

5 $\left\{\mathsf{region}(\{\mathsf{Low}, \mathsf{High}\}, T_\mathsf{b}, I_\mathsf{b}(\mathbf{x}, P), r) * [\mathsf{get}]_1^r * \mathsf{stable}(P)\right\}$

6 `        // case-split on` Low / High.

7 $\left\langle(\mathsf{region}(\{\mathsf{Low}\}, T_\mathsf{b}, I_\mathsf{b}(\mathbf{x}, P), r) \vee \mathsf{region}(\{\mathsf{High}\}, T_\mathsf{b}, I_\mathsf{b}(\mathbf{x}, P), r)) * [\mathsf{get}]_1^r * \mathsf{stable}(P)\right\rangle$

8 `        b = x->flag; //` Low & High *cases given below.*

9 $\left\langle\begin{array}{l}(\mathsf{region}(\{\mathsf{Low}, \mathsf{High}\}, T_\mathsf{b}, I_\mathsf{b}(\mathbf{x}, P), r) * [\mathsf{get}]_1^r \wedge \mathsf{b} = 0) \ \vee \\ (\mathsf{region}(\{\mathsf{Done}\}, T_\mathsf{b}, I_\mathsf{b}(\mathbf{x}, P), r) * [\mathsf{get}]_1^r * \mathsf{stable}(P) * P \wedge \mathsf{b} = 1)\end{array}\right\rangle$

10 `    } while (b == 0)`

11 `  } // Delete region assertions using weakening.`

12 $\left\{P\right\}$

LOW CASE:

1 $\left\langle\mathsf{region}(\{\mathsf{Low}\}, T_\mathsf{b}, I_\mathsf{b}(\mathbf{x}, P), r) * [\mathsf{get}]_1^r\right\rangle$

2 *// open region*

3 $\left\langle\triangleright(\mathbf{x}.\mathsf{flag} \mapsto 0) * [\mathsf{get}]_1^r\right\rangle$

4 *// drop $\triangleright$ and read value.*

5 `b = x->flag;`

6 $\left\langle\mathbf{x}.\mathsf{flag} \mapsto 0 * [\mathsf{get}]_1^r \wedge \mathsf{b} = 0\right\rangle$

7 *// Close region with* get *transition.*

8 *// Stabilise region state.*

9 $\left\langle\begin{array}{c}\mathsf{region}(\{\mathsf{Low}, \mathsf{High}\}, T_\mathsf{b}, I_\mathsf{b}(\mathbf{x}, P), r) \\ * [\mathsf{get}]_1^r \wedge \mathsf{b} = 0\end{array}\right\rangle$

HIGH CASE:

1 $\left\langle\mathsf{region}(\{\mathsf{High}\}, T_\mathsf{b}, I_\mathsf{b}(\mathbf{x}, P), r) * [\mathsf{get}]_1^r\right\rangle$

2 *// open region*

3 $\left\langle\triangleright(\mathbf{x}.\mathsf{flag} \mapsto 1 * \mathsf{stable}(P) * P) * [\mathsf{get}]_1^r\right\rangle$

4 *// drop $\triangleright$, and read value.*

5 `b = x->flag;`

6 $\left\langle\mathbf{x}.\mathsf{flag} \mapsto 1 * \mathsf{stable}(P) * P * [\mathsf{get}]_1^r \wedge \mathsf{b} = 1\right\rangle$

7 *// close region with* get *transition.*

8 $\left\langle\begin{array}{c}\mathsf{region}(\{\mathsf{Done}\}, T_\mathsf{b}, I_\mathsf{b}(\mathbf{x}, P), r) * \\ \mathsf{stable}(P) * P * [\mathsf{get}]_1^r \wedge \mathsf{b} = 1\end{array}\right\rangle$

Fig. 11. Proof of `wait()` using simple predicate definitions.

5. SPLITTING CHANNELS

In this section, we extend our proof of the simple implementation to cover *splitting* – intuitively, promised resources can be divided between threads before they are sent. As the simple implementation sequentialises signalling, we leave extension and renunciation to §6.

Strengthened abstract specification. To manage promised split resources inside the proof, we use separating implication, $-\!*$. However, the stability assertion stable is difficult to reason about because it is does not distribute with respect to $*$. Explicit stabilization operators, $\lceil - \rceil$ and $\lfloor - \rfloor$ are easier to reason about, and we therefore define a channel specification written in terms of explicit stabilization, from which our client-facing abstract specification can be derived. This strengthened specification is given in Figure 12.

To derive the weaker specification given in §3.4 from this stronger one, we define $\mathsf{send}_w(x, P)$ in the weaker specification as $\exists P'. \ \mathsf{send}_s(x, P') \wedge \mathsf{valid}(P \Rightarrow P')$, using the send_s predicate from the stronger specification. The valid predicate allows us to utilise facts about iCAP axioms in our library definitions and specifications. For further discussion, see Appendix C.2.1. Predicates other than send are lifted to the new specification without modification.

Function specifications then follow directly from $\mathsf{stable}(P) \implies (\lfloor P \rfloor \Leftrightarrow P \Leftrightarrow \lceil P \rceil)$.

SPECIFICATIONS:

$$\{\mathsf{emp}\} \quad \mathtt{i\ =\ newchan()} \quad \{\mathsf{recv}(\mathtt{i},P) * \mathsf{send}(\mathtt{i},P)\}$$

$$\{\mathsf{send}(\mathtt{i},P) * \lfloor P \rfloor\} \quad \mathtt{signal(i)} \quad \{\mathsf{emp}\}$$

$$\{\mathsf{recv}(\mathtt{i},P)\} \quad \mathtt{wait(i)} \quad \{\lceil P \rceil\}$$

$$\left\{ \begin{matrix} \mathsf{send}(\mathtt{x},P) * \\ \circledast_{e \in E}\, e \prec \mathtt{x} * \circledast_{l \in L}\, \mathtt{x} \prec l \end{matrix} \right\} \quad \mathtt{(b,a)=extend(x)} \quad \left\{ \begin{matrix} \mathsf{send}(\mathtt{b},Q) * \mathsf{recv}(\mathtt{b},Q) * \mathsf{send}(\mathtt{a},P) \\ * \mathtt{b} \prec \mathtt{a} * \circledast_{e \in E}\, e \prec \mathtt{b} * \circledast_{l \in L}\, \mathtt{a} \prec l \end{matrix} \right\}$$

AXIOMS:

$$x \prec y \quad \implies \quad x \prec y * x \prec y$$

$$x \prec y * y \prec z \quad \implies \quad x \prec z$$

$$\{\mathsf{recv}(x,P) * \mathsf{send}(y,Q) * x \prec y\} \quad \langle\mathtt{skip}\rangle \quad \{\mathsf{send}(y, P \mathbin{-\!\!*} Q)\}$$

$$\{\mathsf{recv}(a,P) * \lfloor \lceil P \rceil \mathbin{-\!\!*} (P_1 * P_2) \rfloor\} \quad \langle\mathtt{skip}\rangle \quad \{\mathsf{recv}(a,P_1) * \mathsf{recv}(a,P_2)\}$$

Fig. 12.   Full specification with explicit stabilisation.

To derive the weakened renunciation axiom, we reason as follows:

$\{\mathsf{send}_w(x, P * Q) * \mathsf{recv}(y,P) * x \prec y\}$

$\iff \{\exists P'.\, \mathsf{send}_s(x, P') * \mathsf{valid}((P * Q) \Rightarrow P') * \mathsf{recv}(y,P) * x \prec y\}$ (Definition)

$\iff \{\exists P'.\, \mathsf{send}_s(x, P') * \mathsf{valid}(Q \Rightarrow (P \mathbin{-\!\!*} P')) * \mathsf{recv}(y,P) * x \prec y\}$ (Adjoint)

    $\langle\mathtt{skip}\rangle$ (Renunciation spec)

    $\{\exists P'.\, \mathsf{send}_s(x, P \mathbin{-\!\!*} P') * \mathsf{valid}(Q \Rightarrow (P \mathbin{-\!\!*} P'))\}$ (Definition)

$\iff \{\mathsf{send}_w(Q)\}$

In our strengthened specification, splitting is expressed by the following axiom:

$$\{\mathsf{recv}(a,P) * \lfloor \lceil P \rceil \mathbin{-\!\!*} (P_1 * P_2) \rfloor\} \quad \langle\mathtt{skip}\rangle \quad \{\mathsf{recv}(a,P_1) * \mathsf{recv}(a,P_2)\}$$

To derive the weaker splitting axiom note that if $\mathsf{stable}(P * Q)$ then:

$$\top \Rightarrow \lfloor \top \rfloor \Rightarrow \lfloor P * Q \mathbin{-\!\!*} P * Q \rfloor \Rightarrow \lfloor \lceil P * Q \rceil \mathbin{-\!\!*} P * Q \rfloor$$

In the remainder of the paper, we verify our implementations against the more general specification given in Figure 12.

Abstract state. To begin the proof, we first define the set of abstract states. In the absence of splitting, the entire promised resource is transferred to a single recipient. For each channel, the resource can therefore be either unsent, sent but not received, or received. In Section 4 we thus introduced three abstract states (Low, High and Done) to represent these three situations.

Splitting means that promised resources can be logically divided between recipients. The abstract state must therefore also track how the promised resource has been split and (if it has been sent) which recipients have taken ownership. Intuitively, we parameterise the abstract state with sets of propositions describing the splitting. We would then have abstract states $\mathsf{Low}(\mathcal{I})$ and $\mathsf{High}(\mathcal{I})$, with $\mathcal{I} \in \mathcal{P}_{\mathit{fin}}(\mathsf{Prop})$. Each $P \in \mathcal{I}$ represents a promise to a recipient. In the abstract state $\mathsf{Low}(\mathcal{I})$, the whole resource, which has yet to be sent, has been split into a set of promised resources $\mathcal{I}$. In the abstract state $\mathsf{High}(\mathcal{I})$, the entire promised resource has been sent, and the portions still in $\mathcal{I}$ have yet to be received.

In the presence of splitting, $\mathsf{recv}(a,P)$ only confers the right to receive and take ownership of the portion of the resource represented by P. We capture this by indexing the split and

receive transition with a proposition describing the associated resource. We thus have two transitions: send for sending the entire promised resource and change_P for splitting the resource P or taking ownership of P.

$$\mathsf{send} : \qquad \mathsf{Low}(\mathcal{I}) \rightsquigarrow \mathsf{High}(\mathcal{I})$$

$$\mathsf{change}_P : \ \mathsf{Low}(\mathcal{I} \uplus \{P\}) \rightsquigarrow \mathsf{Low}(\mathcal{I} \uplus \{P_1, P_2\})$$
$$\mathsf{change}_P : \ \mathsf{High}(\mathcal{I} \uplus \{P\}) \rightsquigarrow \mathsf{High}(\mathcal{I} \uplus \{P_1, P_2\})$$
$$\mathsf{change}_P : \ \mathsf{High}(\mathcal{I} \uplus \{P\}) \rightsquigarrow \mathsf{High}(\mathcal{I})$$

Note that when splitting P using the change_P transition, the transition system does not enforce that $P \twoheadrightarrow P_1 * P_2$. Rather, this will be enforced by the interpretation function for the abstract states.

The transition system described in this section unfortunately cannot be expressed directly in iCAP. This is because iCAP's abstract states and transitions cannot be directly indexed by propositions. It is unclear how this restriction could be lifted in iCAP's step-indexing framework.

Instead, iCAP supports *saved propositions*, an encoding which allows propositions to be associated with identifiers and stored. To formalize the above transition system, we index states and transitions with these identifiers. This skirts the restrictions imposed by step-indexing, and allows the reasoning we want. In the following section we introduce the necessary logical machinery.

5.1. Saved Propositions

A saved proposition, written $r \stackrel{\pi}{\Longmapsto} P$, associates an identifier r with a proposition P. By introducing the indirection from identifiers to propositions we lose some properties. Most importantly, we cannot easily unify saved propositions: given $r \stackrel{\pi_1}{\Longmapsto} P$ and $r \stackrel{\pi_2}{\Longmapsto} Q$, in general it does not hold that $P = Q$. However, saved propositions still satisfy enough properties that we can verify the splitting axiom.

In addition to the identifier r and proposition P, we also have a fractional parameter $\pi \in (0, 1]$ which records how the saved proposition has been shared between threads. In other words, it serves the same role as fractional permissions for heap cells in standard separation logic [Bornat et al. 2005].

We require that saved propositions satisfy the following three properties:

$$\mathsf{emp} \implies \exists r.\ r \stackrel{1}{\Longmapsto} P \tag{1}$$

$$r \stackrel{\pi_1}{\Longmapsto} P * r \stackrel{\pi_2}{\Longmapsto} P \iff \begin{cases} r \stackrel{\pi_1 + \pi_2}{\Longmapsto} P & \textit{if } (\pi_1 + \pi_2) \leq 1 \\ \mathsf{false} & \textit{otherwise} \end{cases} \tag{2}$$

$$r \stackrel{\pi_1}{\Longmapsto} P * r \stackrel{\pi_2}{\Longmapsto} Q \implies r \stackrel{\pi_1}{\Longmapsto} P * r \stackrel{\pi_2}{\Longmapsto} Q * (\triangleright P \Rightarrow \triangleright Q) \tag{3}$$

Property (1) allows us to create a saved proposition for an arbitrary proposition P. *Property* (2) says that saved propositions are linear, meaning we can split and join them without worrying about unwanted duplication. Observe that the fractions π_1 and π_2 are used to track splitting. *Property* (3) says that holding two saved propositions on the same region allow us to convert from one to another. This is a corrected version of the unification property discussed above. The iCAP later operator '$\triangleright$' is needed because we use shared regions internally in the definition of saved propositions.

The following additional property is derivable from property (3) – proved in Lemma A.1.

$$\left(\begin{array}{c} r \overset{\pi_1}{\longmapsto} P * r \overset{\pi_2}{\longmapsto} Q * \\ (X \mathrel{-\!\!*} \rhd(Q * Y)) * (P \mathrel{-\!\!*} Z) \end{array} \right) \quad \Longrightarrow \quad r \overset{\pi_1}{\longmapsto} P * r \overset{\pi_2}{\longmapsto} Q * (X \mathrel{-\!\!*} \rhd(Z * Y)) \tag{4}$$

This says that we can apply property (3) inside separating implications. This is useful when modifying a resource embedded into a larger assertion.

We can encode and verify saved propositions as predicates in iCAP – they do not require any extension of the logic. Our encoding is given in Appendix A. In our encoding, region identifiers are used as identifiers for saved propositions – this is because internally saved propositions are encoded by regions.

5.2. Predicate Definitions for send and recv

Once again, we begin by defining the structure of a region. Abstract states are now terms of the form $\mathsf{Low}(\mathcal{I})$ and $\mathsf{High}(\mathcal{I})$, where $\mathcal{I}$ is a finite set of identifiers in $\mathcal{P}_{\mathit{fin}}(\mathsf{RId})$. We use LoHi to stand for either Low or High. The $\mathcal{I}$ parameter represents the set of outstanding obligations, i.e. the resources that other threads expect to be supplied. As described above, we use saved propositions to give an interpretation to these sets of region identifiers. If we have the abstract state $\mathsf{Low}(\mathcal{I})$ or $\mathsf{High}(\mathcal{I})$, then each $i \in \mathcal{I}$ corresponds to a resource promised to some thread. To find out *what* resource P is expected, we examine the associated saved proposition $i \overset{\pi}{\longmapsto} P$.

Actions in the transition relation T_{m} are of the form send and $\mathsf{change}(i)$, where i is the identifier for a saved proposition. The first action, send, sets the flag, and simply moves from Low to High. The second, $\mathsf{change}(i)$, is both the action of taking the resource associated to region i when the flag is high, and splitting the resource associated with identifier i to the resource required by i_1 and i_2. These new identifiers i_1/i_2 can be chosen arbitrarily: however, the invariant mapping I_{m} ensures they are properly associated with saved propositions.

$$T_{\mathsf{m}}(\mathsf{send}) \triangleq \{(\mathsf{Low}(\mathcal{I}), \mathsf{High}(\mathcal{I}))\}$$

$$T_{\mathsf{m}}(\mathsf{change}(i)) \triangleq \{(\mathsf{High}(\mathcal{I} \uplus \{i\}), \mathsf{High}(\mathcal{I}))\} \cup \{(\mathsf{Low}(\mathcal{I} \uplus \{i\}), \mathsf{Low}(\mathcal{I} \uplus \{i_1, i_2\}))\}$$
$$\cup \{(\mathsf{High}(\mathcal{I} \uplus \{i\}), \mathsf{High}(\mathcal{I} \uplus \{i_1, i_2\}))\}$$

The invariants associated with Low and High are defined as follows:

$$I_{\mathsf{m}}(x, r, P)(\mathsf{Low}(\mathcal{I})) \triangleq x.\mathsf{flag} \mapsto 0 * \mathsf{waiting}(P, \mathcal{I}) * \mathsf{changes}_r(\mathcal{I})$$

$$I_{\mathsf{m}}(x, r, P)(\mathsf{High}(\mathcal{I})) \triangleq x.\mathsf{flag} \mapsto 1 * \mathsf{ress}(\mathcal{I}) * \mathsf{changes}_r(\mathcal{I})$$

$$\textit{where}$$

$$\mathsf{waiting}(P, \mathcal{I}) \triangleq \exists Q \colon \mathcal{I} \to \mathsf{Prop}. \lfloor (\rhd P) \mathrel{-\!\!*} \rhd \underset{i \in \mathcal{I}}{\circledast} . Q(i) \rfloor * \underset{i \in \mathcal{I}}{\circledast} . i \overset{1/2}{\longmapsto} Q(i)$$

$$\mathsf{ress}(\mathcal{I}) \triangleq \underset{i \in \mathcal{I}}{\circledast} \exists R. (i \overset{1/2}{\longmapsto} R) * \lceil \rhd R \rceil$$

$$\mathsf{changes}_r(\mathcal{I}) \triangleq \underset{i \notin \mathcal{I}}{\circledast} . [\mathsf{change}(i)]_1^r$$

The definitions here use three auxiliary predicates: $\mathsf{waiting}$, standing for resources that have been promised but not supplied; ress, standing for resources once they have been supplied; and $\mathsf{changes}$, standing for change tokens for unused identifiers. The set $\mathsf{changes}$ can be seen as a 'library' of tokens which are not currently in use, and are currently not used by any thread. Having this library of tokens allows new saved propositions to be added when splitting.

```
1  {emp}
2  chan *newchan() {
3    chan *x := new(chan);
4    x->flag := 0;
5  {x.flag ↦ 0}
6    // Create floor assertion
7  {x.flag ↦ 0 * ⌊(▷P) -* (▷P)⌋}
8    skip; // Create saved proposition
9  {∃r₁. r₁ ↦¹ P * x.flag ↦ 0 * ⌊(▷P) -* (▷P)⌋}
10   // Create channel region
11 {∃r₁, r₂. r₁ ↦^{1/2} P * creg(x, r₂, P, {Low({r₁})}) * [send]₁^{r₂} * [change(r₁)]₁^{r₂}}
12   // Satisfy the send predicate
13 {∃r₁, r₂. send(x, P) * r₁ ↦^{1/2} P * creg(x, r₂, P, {Low({r₁})}) * [change(r₁)]₁^{r₂}}
14 } // Satisfy the recv predicate
15 {send(x, P) * recv(x, P)}
```

Fig. 13.   Proof of `newchan()` w.r.t. the full specification.

The representation of Low consists of the flag, change tokens, and the $\mathsf{waiting}$ predicate. $\mathsf{waiting}(P, \mathcal{I})$ requires the existence of a mapping Q from region identifiers, to propositions representing obligations to other threads. The obligations for different threads are tied together using fractional saved propositions $i \overset{1/2}{\Longmapsto} Q(i)$. The assertion $\lfloor (\triangleright P) \twoheadrightarrow \triangleright \bigotimes_{i \in \mathcal{I}}. Q(i) \rfloor$ means that supplying the resource P will satisfy each obligation $Q(i)$.

The representation of the High state consists of the flag and change tokens, and the ress predicate. $\mathsf{ress}(\mathcal{I})$ pairs together fractional saved propositions, $i \overset{1/2}{\Longmapsto} R$ with resources $\lceil \triangleright R \rceil$. The other half of each saved proposition is held by the thread that has been promised the resource through the recv predicate (see below). This ensures that all threads that have been promised resources can claim them.

We use a shorthand for the region assertion in our definitions and proofs:

$$\mathsf{creg}(x, r, P, S) \triangleq \mathsf{region}(S, T_{\mathsf{m}}, I_{\mathsf{m}}(x, r, P), r)$$

The definition of the send predicate is now straightforward. It asserts that the region is in a Low state, and holds the unique permission to perform the send action.

$$\mathsf{send}(x, P) \triangleq \exists r. \mathsf{creg}(x, r, P, \{\mathsf{Low}(\mathcal{I}) \mid \mathsf{true}\}) * [\mathsf{send}]_1^r$$

The definition of $\mathsf{recv}(x, Q)$ predicate is more complex. It includes $r' \overset{1/2}{\Longmapsto} Q$, half the permission on the saved proposition Q. It also asserts that r' is one of identifiers recorded in the region. This ensures that the resource retrieved from the shared region is the correct one, i.e. the one that was promised (see the next section for the reasoning steps involved).

$$\mathsf{recv}(x, Q) \triangleq \exists R, r, r'. \mathsf{creg}(x, r, R, \{\mathsf{LoHi}(\mathcal{I}) \mid r' \in \mathcal{I}\}) * r' \overset{1/2}{\Longmapsto} Q * [\mathsf{change}(r')]_1^r$$

5.3. Proofs of `newchan()`, `signal()`, `wait()`, and the Splitting Axiom

Proving `newchan`. The proof of `newchan` (Fig. 13) allocates a region containing the concrete channel state. Most steps in the proof are straightforward. The challenging ones are

1 $\left\{ \mathsf{send}(\mathbf{x}, P) * \lfloor P \rfloor \right\}$

2 `signal(chan *x) {`

3 $\left\{ \mathsf{creg}(\mathbf{x}, r, P, \{\mathsf{Low}(\mathcal{I}') \mid \mathsf{true}\}) * [\mathsf{send}]_1^r * \lfloor P \rfloor \right\}$

4   `// open the region.` $\mathsf{Low}(\mathcal{I})$ `is an arbitrary member of` $\{\mathsf{Low}(\mathcal{I}') \mid \mathsf{true}\}$`.`

5      $\left\langle \triangleright I_{\mathsf{m}}(\mathbf{x}, r, P)(\mathsf{Low}(\mathcal{I})) * [\mathsf{send}]_1^r * \lfloor P \rfloor \right\rangle$

6   `// invariant definition.`

7      $\left\langle [\mathsf{send}]_1^r * \lfloor P \rfloor * \triangleright (\mathbf{x}.\mathsf{flag} \mapsto 0 * \mathsf{waiting}(P, \mathcal{I}) * \mathsf{changes}_r(\mathcal{I})) \right\rangle$

8   `// pull out points-to using` $LBin$ `and` $LPoints$`.`

9      $\left\langle [\mathsf{send}]_1^r * \lfloor P \rfloor * \mathbf{x}.\mathsf{flag} \mapsto 0 * \triangleright (\mathsf{waiting}(P, \mathcal{I}) * \mathsf{changes}_r(\mathcal{I})) \right\rangle$

10   `x->flag := 1; // rewrite flag, drop` $\triangleright$ `using` $AFrame$

11     $\left\langle [\mathsf{send}]_1^r * \lfloor P \rfloor * \mathbf{x}.\mathsf{flag} \mapsto 1 * \mathsf{waiting}(P, \mathcal{I}) * \mathsf{changes}_r(\mathcal{I}) \right\rangle$

12   `// Lemma` 5.1`, add a` $\triangleright$ `using` $SMono$`.`

13     $\left\langle [\mathsf{send}]_1^r * \triangleright (\mathbf{x}.\mathsf{flag} \mapsto 1 * \mathsf{ress}(\mathcal{I}) * \mathsf{changes}_r(\mathcal{I})) \right\rangle$

14   `// Invariant definition.`

15     $\left\langle [\mathsf{send}]_1^r * \triangleright I_{\mathsf{m}}(\mathbf{x}, r, P)(\mathsf{High}(\mathcal{I})) \right\rangle$

16   `// close the region using` `send` `transition.`

17 $\left\{ [\mathsf{send}]_1^r * \mathsf{creg}(\mathbf{x}, r, P, \{\mathsf{High}(\mathcal{I}') \mid \mathsf{true}\}) \right\}$

18 `} // delete redundant assertions using weakening.`

19 $\left\{ \mathsf{emp} \right\}$

Fig. 14.   Proof of `signal()` w.r.t. the full specification.

line 6 – creating the stabilised assertion – and line 10 – the view-shift which creates the region itself. In line 6, we need the following implication:

$$\mathsf{emp} \implies \lfloor (\triangleright P) \mathbin{-\!\!*} (\triangleright P) \rfloor$$

To show this holds, we observe that for any X, $\mathsf{emp} \implies X \mathbin{-\!\!*} X$, and that emp is always stable. The implication then follows by monotonicity of explicit stabilisation brackets, $(A \Rightarrow B) \implies (\lfloor A \rfloor \Rightarrow \lfloor B \rfloor)$.

Line 10 requires us to prove the following view-shift:

$$r_1 \overset{1/2}{\Longmapsto} P * \mathbf{x} \mapsto 0 * \lfloor (\triangleright P) \mathbin{-\!\!*} (\triangleright P) \rfloor$$
$$\sqsubseteq \quad \exists r_2.\, \mathsf{creg}(\mathbf{x}, r_2, P, \{\mathsf{Low}(\{r_1\}, \emptyset)\}) * [\mathsf{send}]_1^{r_2} * [\mathsf{change}(r_1)]_1^{r_2}$$

To prove this, we appeal to iCAP's VAlloc rule, which controls construction of new regions (see §4.1 for its definition). Intuitively, this requires that the resources available satisfy the initial abstract state, defined by I_{m}. Note that all the change tokens apart from $\mathsf{change}(r_1)$ are stored in the token library predicate $\mathsf{changes}$, held inside the new region.

Proving `signal`. The proof of `signal` (Fig. 14) works by opening the channel region (line 4), merging in the supplied resource $\lfloor P \rfloor$ to give the promised resources (line 12), and closing the region again (line 16). When we close the region we also need to confirm that the transition from $\mathsf{Low}(\mathcal{I})$ to $\mathsf{High}(\mathcal{I})$ is allowed, but this is simple: it's the only transition associated to send by T_{m}. The trickiest step is the merging of the resource into the region (line 12), embodied by the following lemma.

LEMMA 5.1. $\lfloor P \rfloor * \mathsf{waiting}(P, \mathcal{I}) \sqsubseteq \mathsf{ress}(\mathcal{I})$

```
1  {recv(x, P)}
2  wait(chan *x) {
3    {∃R, r, r'. r' ⟼^{1/2} P * [change(r')]₁^r * creg(x, r, R, {LoHi(I') | r' ∈ I'})}
4      // Open the region. Only consider the High case.
5      ⟨r' ∈ I ∧ r' ⟼^{1/2} P * [change(r')]₁^r * ▷Iₘ(x, r, R)(High(I))⟩
6        // Apply invariant definition
7        // Pull out points-to using LBin and LPoints
8      ⟨r' ∈ I ∧ r' ⟼^{1/2} P * [change(r')]₁^r * x.flag ↦ 1 * ▷(ress(I) * changes_r(I))⟩
9        assume(x->flag == 1) // drop ▷ using AFrame, push in [change(r')] perm.
10     ⟨r' ∈ I ∧ r' ⟼^{1/2} P * x.flag ↦ 1 * ress(I) * changes_r(I \ {r'})⟩
11       // Lemma 5.2, add ▷ using SMono
12     ⟨▷⌈P⌉ * ▷(x.flag ↦ 1 * ress(I \ {r'}) * changes_r(I \ {r'}))⟩
13       // Invariant definition
14     ⟨▷⌈P⌉ * ▷Iₘ(x, r, R)(High(I \ {r'}))⟩
15     // close the region using change(r') transition.
16   {▷⌈P⌉ * ∃R, r. creg(x, r, R, {High(I) | true})}
17     // Delete redundant assertions using weakening.
18   {▷⌈P⌉}
19  }
```

Fig. 15. Proof of `wait()` w.r.t. the full specification.

PROOF. ⌊P⌋ * waiting(P, I)

$\qquad \sqsubseteq$ ⌊▷P⌋ * ∃Q: I → Prop. ⌊▷P -* ▷ ⊛_{i∈I}. Q(i)⌋ * (⊛_{i∈I}. i ⟼^{1/2} Q(i))

$\qquad \sqsubseteq$ ∃Q: I → Prop. ⌊▷ ⊛_{i∈I}. Q(i)⌋ * (⊛_{i∈I}. i ⟼^{1/2} Q(i))

$\qquad \sqsubseteq$ ∃Q: I → Prop. (⊛_{i∈I}. ⌈▷Q(i)⌉) * (⊛_{i∈I}. i ⟼^{1/2} Q(i))

$\qquad \sqsubseteq$ (⊛_{i∈I}. ∃R. i ⟼^{1/2} R * ⌈▷R⌉)

$\qquad \sqsubseteq$ ress(I)

All of the steps in this proof are in fact standard implications (this is also true of all the other view-shift lemmas unless stated explicitly). To prove the second step, we appeal to the fact that ⌊−⌋ is semi-distributitive over separating conjunction, ⌊A⌋ * ⌊B⌋ $\implies$ ⌊A * B⌋, and *modus ponens* for separating implication, A * (A -* B) $\implies$ B. The third step follows from the fact that ⌈−⌉ is weaker than ⌊−⌋, and ⌈−⌉ is semi-distributitive over separating conjunction, ⌈A * B⌉ $\implies$ ⌈A⌉ * ⌈B⌉. □

Proving `wait`. In the proof of `wait` (Fig. 15) we open the shared region (line 4), extract the required resource (line 11) and close the region again (line 15). For simplicity, we assume that the abstract state is High; if not, the algorithm spins doing nothing until it is the case. Each promised resource is associated with a region identifier i in the set I; removing the resource is modelled abstractly by removing i. This abstract transition is allowed by the [change] permission. The key step in the proof is extracting the resource (line 11), embodied by the following lemma.

$$1 \quad \left\{ \mathsf{recv}(\mathrm{x}, P) * \lfloor \lceil P \rceil \twoheadrightarrow P_1 * P_2 \rfloor \right\}$$

$$2 \quad \left\{ \exists R, r, r'. \, r' \stackrel{1/2}{\longmapsto} P * [\mathsf{change}(r')]_1^r * \mathsf{creg}(\mathrm{x}, r, R, \{\mathsf{LoHi}(\mathcal{I}') \mid r' \in \mathcal{I}'\}) * \lfloor \lceil P \rceil \twoheadrightarrow P_1 * P_2 \rfloor \right\}$$

3   // *Open the region using an arbitrary state containing* r'

$$4 \quad \left\langle \exists R, r, r'. \, r' \stackrel{1/2}{\longmapsto} P * [\mathsf{change}(r')]_1^r * \triangleright I_\mathsf{m}(\mathrm{x}, r, R)(\mathsf{LoHi}(\mathcal{I} \uplus \{r'\})) * \lfloor \lceil P \rceil \twoheadrightarrow P_1 * P_2 \rfloor \right\rangle$$

5   `skip`   // *Use AFrame to remove* $\triangleright$

$$6 \quad \left\langle \exists R, r, r'. \, r' \stackrel{1/2}{\longmapsto} P * [\mathsf{change}(r')]_1^r * I_\mathsf{m}(\mathrm{x}, r, R)(\mathsf{LoHi}(\mathcal{I} \uplus \{r'\})) * \lfloor \lceil P \rceil \twoheadrightarrow P_1 * P_2 \rfloor \right\rangle$$

7   // *Lemma* 5.5

$$8 \quad \left\langle \begin{array}{l} \exists R, r, r_1, r_2. \, r_1 \stackrel{1/2}{\longmapsto} P_1 * r_1 \stackrel{1/2}{\longmapsto} P_2 * [\mathsf{change}(r_1)]_1^r * [\mathsf{change}(r_2)]_1^r \\ * \, I_\mathsf{m}(\mathrm{x}, r, R)(\mathsf{LoHi}(\mathcal{I} \uplus \{r_1, r_2\})) \end{array} \right\rangle$$

9   // *Close the region using* $\mathsf{change}(r')$ *transition.*

$$10 \quad \left\{ \begin{array}{l} \exists R, r, r_1, r_2. \, r_1 \stackrel{1/2}{\longmapsto} P_1 * r_1 \stackrel{1/2}{\longmapsto} P_2 * [\mathsf{change}(r_1)]_1^r * [\mathsf{change}(r_2)]_1^r \\ * \, \mathsf{creg}(\mathrm{x}, r, R, \{\mathsf{LoHi}(\mathcal{I}') \mid r_1, r_2 \in \mathcal{I}'\}) \end{array} \right\}$$

11 // *Definition of predicates.*

$$12 \quad \left\{ \mathsf{recv}(\mathrm{x}, P_1) * \mathsf{recv}(\mathrm{x}, P_2) \right\}$$

Fig. 16. Proof outline for splitting axiom.

LEMMA 5.2. $r \stackrel{1/2}{\longmapsto} P * \mathsf{ress}(\mathcal{I}) \wedge r \in \mathcal{I} \; \sqsubseteq \; \mathsf{ress}(\mathcal{I} \setminus \{r\}) * \triangleright \lceil P \rceil$

PROOF. $r \stackrel{1/2}{\longmapsto} P * \mathsf{ress}(\mathcal{I} \uplus \{r\})$

$\sqsubseteq \; r \stackrel{1/2}{\longmapsto} P * \mathsf{ress}(\mathcal{I}) * \exists R. \, r \stackrel{1/2}{\longmapsto} R * \lceil \triangleright R \rceil$ (Definition of ress)

$\sqsubseteq \; r \stackrel{1/2}{\longmapsto} P * \mathsf{ress}(\mathcal{I}) * \exists R. \, r \stackrel{1/2}{\longmapsto} R * \lceil \triangleright P \rceil$ (Property 3, mono of $\lceil - \rceil$)

$\sqsubseteq \; \mathsf{ress}(\mathcal{I}) * \triangleright \lceil P \rceil$ ($\lceil \triangleright P \rceil \implies \triangleright \lceil P \rceil$)

                                                                  □

Proving the Splitting Axiom. In our specification, splitting must always be associated with a `skip` step. It should now be clear why we need this: a `skip` step allows us to enter the shared region and get rid of $\triangleright$. We present the proof outline in Figure 16. The core of the proof is two lemmas which express splitting in the Low and High cases.

LEMMA 5.3 (LOW SPLITTING).

$$r \stackrel{1/2}{\longmapsto} P * \lfloor \lceil P \rceil \twoheadrightarrow P_1 * P_2 \rfloor * \mathsf{waiting}(R, \mathcal{I} \uplus r)$$
$$\sqsubseteq \; \exists r_1, r_2. \, r_1 \stackrel{1/2}{\longmapsto} P_1 * r_2 \stackrel{1/2}{\longmapsto} P_2 * \mathsf{waiting}(R, \mathcal{I} \uplus \{r_1, r_2\})$$

Proof.

$$r \overset{1/2}{\Longmapsto} P * \lfloor \lceil P \rceil \rightarrow\!\!* P_1 * P_2 \rfloor * \mathsf{waiting}(R, \mathcal{I} \uplus r)$$

(Predicate definitions, extract saved propositions)

$$\sqsubseteq \exists Q'.\ r \overset{1/2}{\Longmapsto} P * r \overset{1/2}{\Longmapsto} Q' * \lfloor \lceil P \rceil \rightarrow\!\!* P_1 * P_2 \rfloor *$$
$$\exists Q : \mathcal{I} \rightarrow \mathsf{Prop}. \lfloor (\rhd R) \rightarrow\!\!* \rhd Q' * \rhd \bigcircledast_{i \in \mathcal{I}}. Q(i) \rfloor * (\bigcircledast_{i \in \mathcal{I}}. i \overset{1/2}{\Longmapsto} Q(i))$$

(Unify P/Q', apply $P \implies \lceil P \rceil$)

$$\sqsubseteq \lfloor \lceil \rhd P \rceil \rightarrow\!\!* (\rhd P_1) * (\rhd P_2) \rfloor *$$
$$\exists Q : \mathcal{I} \rightarrow \mathsf{Prop}. \lfloor (\rhd R) \rightarrow\!\!* \lceil \rhd P \rceil * \rhd \bigcircledast_{i \in \mathcal{I}}. Q(i) \rfloor * (\bigcircledast_{i \in \mathcal{I}}. i \overset{1/2}{\Longmapsto} Q(i))$$

(Modus ponens, create saved props r_1/r_2)

$$\sqsubseteq \exists r_1, r_2.\ r_1 \overset{1}{\Longmapsto} P_1 * r_2 \overset{1}{\Longmapsto} P_2 \land r_1, r_2 \notin \mathcal{I} *$$
$$\exists Q : \mathcal{I} \rightarrow \mathsf{Prop}. \lfloor (\rhd R) \rightarrow\!\!* (\rhd P_1) * (\rhd P_2) * \rhd \bigcircledast_{i \in \mathcal{I}}. Q(i) \rfloor * (\bigcircledast_{i \in \mathcal{I}}. i \overset{1/2}{\Longmapsto} Q(i))$$

(Push saved props into $\mathcal{I}$)

$$\sqsubseteq \exists r_1, r_2.\ r_1 \overset{1/2}{\Longmapsto} P_1 * r_2 \overset{1/2}{\Longmapsto} P_2 *$$
$$\exists Q : \mathcal{I} \uplus \{r_1, r_2\} \rightarrow \mathsf{Prop}. \lfloor (\rhd R) \rightarrow\!\!* \rhd \bigcircledast_{i \in \mathcal{I} \uplus \{r_1, r_2\}}. Q(i) \rfloor * (\bigcircledast_{i \in \mathcal{I} \uplus \{r_1, r_2\}}. i \overset{1/2}{\Longmapsto} Q(i))$$

(Predicate definitions)

$$\sqsubseteq \exists r_1, r_2.\ r_1 \overset{1/2}{\Longmapsto} P_1 * r_2 \overset{1/2}{\Longmapsto} P_2 * \mathsf{waiting}(R, \mathcal{I} \uplus \{r_1, r_2\})$$

□

Lemma 5.4 (High splitting).

$$r \overset{1/2}{\Longmapsto} P * \lfloor \lceil P \rceil \rightarrow\!\!* P_1 * P_2 \rfloor * \mathsf{ress}(\mathcal{I} \uplus r)$$
$$\sqsubseteq \exists r_1, r_2.\ r_1 \overset{1/2}{\Longmapsto} P_1 * r_2 \overset{1/2}{\Longmapsto} P_2 * \mathsf{ress}(\mathcal{I} \uplus \{r_1, r_2\})$$

Proof.

$$r \overset{1/2}{\Longmapsto} P * \lfloor \lceil P \rceil \rightarrow\!\!* P_1 * P_2 \rfloor * \mathsf{ress}(\mathcal{I} \uplus r)$$

$\sqsubseteq (\lceil P \rceil \rightarrow\!\!* P_1 * P_2) * \mathsf{ress}(\mathcal{I}) * \rhd \lceil P \rceil$ (Lemma 5.2)

$\sqsubseteq ((\rhd \lceil P \rceil) \rightarrow\!\!* (\rhd P_1) * (\rhd P_2)) * \mathsf{ress}(\mathcal{I}) * \rhd \lceil P \rceil$ (SMono, dist $\rhd$ over $\rightarrow\!\!*$)

$\sqsubseteq \mathsf{ress}(\mathcal{I}) * (\rhd P_1) * (\rhd P_2)$ (Modus ponens)

$\sqsubseteq \mathsf{ress}(\mathcal{I}) * (\rhd P_1) * (\rhd P_2) * \exists r_1, r_2.\ r_1 \overset{1}{\Longmapsto} P_1 * r_2 \overset{1}{\Longmapsto} P_2$ (Property 1)

$\sqsubseteq \mathsf{ress}(\mathcal{I} \uplus \{r_1, r_2\}) * \exists r_1, r_2.\ r_1 \overset{1/2}{\Longmapsto} P_1 * r_2 \overset{1/2}{\Longmapsto} P_2$ (Sub-lemma)

The last step in building the new ress predicate consists of two applications of the following sub-lemma:

$$\mathsf{ress}(\mathcal{I}) * \rhd P * r \overset{1}{\Longmapsto} P$$
$$\sqsubseteq (\bigcircledast_{i \in \mathcal{I}} \exists Q. i \overset{1/2}{\Longmapsto} Q * \lceil \rhd Q \rceil) * \lceil \rhd P \rceil * r \overset{1}{\Longmapsto} P$$
$$\sqsubseteq (\bigcircledast_{i \in \mathcal{I} \uplus \{r\}} \exists Q. i \overset{1/2}{\Longmapsto} Q * \lceil \rhd Q \rceil) * r \overset{1/2}{\Longmapsto} P$$
$$\sqsubseteq \mathsf{ress}(\mathcal{I} \uplus \{r\}) * r \overset{1/2}{\Longmapsto} P$$

□

These two lemmas are combined as follows.

LEMMA 5.5. $\quad r' \overset{1/2}{\Longmapsto} P * I_{\mathsf{m}}(x, r, R)(\mathsf{LoHi}(\mathcal{I} \uplus \{r'\})) * [\mathsf{change}(r')]_r^1 * \lfloor \lceil P \rceil \twoheadrightarrow P_1 * P_2 \rfloor$

$$\sqsubseteq$$

$$\exists r_1, r_2. \, [\mathsf{change}(r_1)]_r^1 * r_1 \overset{1/2}{\Longmapsto} P_1 * [\mathsf{change}(r_2)]_r^1 * r_2 \overset{1/2}{\Longmapsto} P_2 *$$
$$\triangleright I_{\mathsf{m}}(x, r, R)(\mathsf{LoHi}(\mathcal{I} \uplus \{r_1, r_2\}))$$

PROOF. We case-split on whether LoHi is Low or High. The two proofs are given by Lemma 5.3 and 5.4 and some rearrangement of the change permissions. □

6. CHAINS AND RENUNCIATION

The simple barrier implementation verified in §4 and §5 does not consider an order of channels. In this section, we verify an implementation that supports chains of channels and early renunciation. Recall that renunciation is expressed by the following axiom:

$$\{\mathsf{recv}(x, P) * \mathsf{send}(y, Q) * x \prec y\} \quad \langle \mathtt{skip} \rangle \quad \{\mathsf{send}(y, P \twoheadrightarrow Q)\}$$

In the chain of channels, $x \prec y$ states that x is earlier than y. This axiom states that the required resource Q can be partially or totally satisfied using the earlier promised resource P (if $P \Leftrightarrow Q$, then $(P \twoheadrightarrow Q) \Leftrightarrow \mathsf{emp}$, i.e. no more resources need to be supplied by the client in order for the signal to be set).

In our new implementation channels are arranged into a chain represented by a linked list. Calls to `signal` do not block, and can complete in any order consistent with the specification. However, renunciation means later channels may depend on resources promised earlier in the chain. To ensure renounced resources are available, `wait` checks all predecessors in the chain. The implementation is defined as follows:

```
struct chan {                      chan *newchan() {
  int flag;                          chan *x = new(chan);
  chan *prev;                        x->flag = 0;
}                                    x->prev = NULL;
                                     return x;
signal(chan *x) {                  }
  x->flag = 1;
}                                  extend(chan *x) {
                                     chan *z = new(chan);
wait(chan *x) {                      z->flag = 0;
  chan *c = x;                       z->prev = x->prev;
  while(c != NULL) {                 x->prev = z;
    while(c->flag == 0) skip;        return (z,x);
    c = c->prev;                   }
  }
}
```

Calling `signal` sets the current channel flag to 1, then exits immediately. When `wait` is called, it blocks until every bit earlier in the chain is set. To do this, it follows `prev` fields, waiting for each `flag` field before accessing the preceding location. To add extra nodes to the chain, `extend` allocates a new channel and then inserts it immediately before the channel passed as an argument.

6.1. Abstract State

Our fundamental approach remains the same as for the previous proof. That is, shared channels are represented by abstract states, e.g. $\mathsf{High}(\mathcal{I})$. Resource obligations are represented by sets of identifiers which are tied to saved propositions, e.g. the members of $\mathcal{I}$. Modifications

to the shared channel by the thread and environment are represented by transitions over these abstract states. Predicates send, recv are defined as constraints on the abstract state.

The difference with the new implementation is that operations access multiple channels along the chain. As a result, the abstract state cannot be a single channel: instead it is an ordered sequence of channel nodes from the set CNode:

$$[\text{node}(x),\ \text{node}(y),\ \text{node}(z),\ldots]$$

Here x, y, z are addresses, and channel nodes are ordered $z \prec y \prec x$. Note that the list is reversed with respect to chain order: nodes closer to the tail precede than those closer to the head. (We do this because pointers in the underlying list go in this direction).

Each CNode plays a similar role to an individual channel in the previous section. Therefore, each has a state High / Low, and a set $\mathcal{I}$ representing splittings of the promised resource. In addition, to handle renunciation each node records a set $\mathcal{W}$ of identifiers for resources promised to it through renunciation. Formally, channel nodes have the following structure:

$$
\begin{aligned}
\text{CNode} \quad &\triangleq \quad \langle \\
&loc \in \text{Addr}, &&\text{(physical address)} \\
&res \in \text{RId}, &&\text{(region ID for sent resource)} \\
&\mathcal{I} \in \mathcal{P}_{fin}(\text{RId}), &&\text{(region IDs for promised resources)} \\
&flg \in \{\text{High}, \text{Low}\}, &&\text{(flag status)} \\
&\mathcal{W} \in \mathcal{P}_{fin}(\text{RId}), &&\text{(region IDs for earlier renounced resources)} \\
&\rangle
\end{aligned}
$$

Each CNode represents one channel in the chain, so the abstract state of the barrier is an *abstract chain* consisting of a finite sequence in CNode^+.

We assume that the CNode locations in any abstract chain are pairwise-distinct. Thus, where convenient we sometimes treat an abstract chain as a function from locations to tuples: i.e. $rs(x)$ gives some tuple $(r, \mathcal{I}, f, \mathcal{W})$. Given a CNode s, we sometimes write $s.flg$, $s.\mathcal{I}$ etc. to identify the appropriate components of the tuple. We use 0 and 1 to represent the Low and High flag state, respectively.

Most of the abstract transitions are just operations on individual channels in the chain: these are liftings of the single-channel operations defined the previous section. For example, the operation set just involves rewriting the flag field for some channel from Low to High, leaving the state otherwise unchanged.

Renunciation is the most interesting case as it involves *two* channels in the chain. The abstract specification for renunciation uses a resource from a recv predicate to satisfy a send predicate – the former must be earlier in the chain than the latter. Correspondingly, the renun abstract transition copies an identifier from the earlier channel's promise set, to the later channel's renunciation set. For example:

$$
\begin{aligned}
[\ldots \langle x, i, \mathcal{I}, 0, \mathcal{W} \rangle \ldots \langle y, i', \mathcal{I}' \uplus \{r\}, 0, \mathcal{W}' \rangle \ldots] \quad &\rightsquigarrow \\
[\ldots \langle x, i'', \mathcal{I}, 0, \mathcal{W} \uplus \{r\} \rangle \ldots &\langle y, i', \mathcal{I}' \cup \{r\}, 0, \mathcal{W}' \rangle \ldots]
\end{aligned}
$$

In the proof, there will exist saved propositions $r \overset{1}{\Mapsto} P$ and $i \overset{1}{\Mapsto} Q$ – the former records the promised resource from y, while the latter records the resource required to signal the channel x. Intuitively, after this transition, the resource P can no longer be claimed: it will be used to satisfy channel x. This corresponds to the send predicate disappearing in the renunciation axiom. The identifier i also changes to i'' – the associated saved proposition will now be $i'' \overset{1}{\Mapsto} (P \mathrel{-\!*} Q)$, refecting the fact that P no longer needs to be supplied.

This transition is purely abstract in the same way that splitting is: nothing has changed in the concrete representation. All that has changed is the way that the threads agree to use resources.

Chain extension also involves multiple channels. The abstract specification takes as its precondition a send and two sets E and L representing nodes earlier and later in the chain:

$$\left\{ \begin{matrix} \mathsf{send}(\mathsf{x}, P) \, * \\ \circledast_{e \in E}\, e \prec \mathsf{x} * \circledast_{l \in L}\, \mathsf{x} \prec l \end{matrix} \right\} \quad \mathtt{(b,a)=extend(x)} \quad \left\{ \begin{matrix} \mathsf{send}(\mathsf{b}, Q) * \mathsf{recv}(\mathsf{b}, Q) * \mathsf{send}(\mathsf{a}, P) \\ *\, \mathsf{b} \prec \mathsf{a} * \circledast_{e \in E}\, e \prec \mathsf{b} * \circledast_{l \in L}\, \mathsf{a} \prec l \end{matrix} \right\}$$

In the abstract state, this corresponds to the following transition:

$$[\dots \mathit{nodes\ in}\ L \dots\ \langle \mathsf{x}, i, \mathcal{I}, 0, \mathcal{W} \rangle \dots \mathit{nodes\ in}\ E \dots] \quad \rightsquigarrow$$
$$[\dots \mathit{nodes\ in}\ L \dots\ \langle \mathsf{a}, i, \mathcal{I}, 0, \mathcal{W} \rangle, \langle \mathsf{b}, i', \mathcal{I}', 0, \emptyset \rangle \dots \mathit{nodes\ in}\ E \dots]$$

The new node is inserted immediately preceding the parameter x, with the remaining structure of the abstract state remaining unchanged.

6.2. Definitions and predicates

Abstract state predicates. The introduction of renunciation makes it important that resources used later in the chain are available as promised earlier in the chain. More concretely, given an abstract chain $x \cdot xs$, every identifier in the set $x.\mathcal{W}$ must be available from some node in xs (i.e. in some set $\mathcal{I}$):

$$\mathsf{available}([\,]) \triangleq \emptyset, \qquad \mathsf{available}(s \cdot xs) \triangleq (\mathsf{available}(xs) \setminus s.\mathcal{W}) \uplus s.\mathcal{I}$$

$$\mathsf{wf}([\,]) \triangleq \mathsf{true}, \qquad \mathsf{wf}(s \cdot xs) \triangleq \mathsf{wf}(xs) \wedge s.\mathcal{W} \subseteq \mathsf{available}(xs) \wedge s.\mathcal{I} \cap s.\mathcal{W} = \emptyset \wedge$$
$$\forall s' \in xs.\, s.\mathcal{I} \cap s'.\mathcal{I} = \emptyset \wedge s.\mathcal{W} \cap s'.\mathcal{W} = \emptyset$$

The predicate available constructs the set of identifiers that have been promised earlier in the chain, and that have not been taken by some other earlier channel. Well-formedness, wf, then requires that the set of identifiers available from earlier in the chain includes those required by the current channel.

We also define two predicates over abstract chains, ctrue and cconf. The first asserts that all the flags in the abstract chain have been set, while the second furthermore asserts that all sets of waited-for renounced resources are empty. When the latter holds, any resources promised by this node must be available for retrieval.

$$\mathsf{ctrue}(rs) \triangleq \forall e \in rs.\, e.\mathit{flg} = 1$$
$$\mathsf{cconf}(rs) \triangleq \forall e \in rs.\, e.\mathit{flg} = 1 \wedge e.\mathcal{W} = \emptyset$$

Finally, we use $rs_1 \overset{pr}{\rightarrow}{}^* rs_2$ to denote that the abstract chain rs_2 can be derived from rs_1 by cancelling out renounced resources with the corresponding promises. This is used in the proof to show that renounced resources can eventually be satisfied with real resources, once all the flags in the chain have been set. $\overset{pr}{\rightarrow}{}^*$ is the transitive-reflexive closure of a relation which cancels a single promise to a later node using a renounced resource from an earlier node. For example:

$$[\dots \langle x, i, \mathcal{I}, 1, \mathcal{W} \uplus \{r\} \rangle, \dots \langle y, i', \mathcal{I}' \uplus \{r\}, 1, \mathcal{W}' \rangle \dots] \quad \overset{pr}{\longrightarrow}$$
$$[\dots \langle x, i, \mathcal{I}, 1, \mathcal{W} \rangle, \dots \langle y, i', \mathcal{I}', 1, \mathcal{W}' \rangle \dots]$$

$$\begin{aligned}
\mathsf{renun}_c(x, rs, rs') &\triangleq rs(x).\mathit{flg} = 0 \wedge \exists r', w.\, rs' = (rs \blacktriangleleft_x \blacktriangleleft_{res} r') \blacktriangleleft_x \blacktriangleleft_{\mathcal{W}} (\bullet \uplus w) \\
\mathsf{set}_c(x, rs, rs') &\triangleq rs(x).\mathit{flg} = 0 \wedge rs' = rs \blacktriangleleft_x \blacktriangleleft_{\mathit{flg}} (1) \\
\mathsf{ext}_c(x, rs, rs') &\triangleq rs = (rs_1 \cdot a \cdot rs_2) \wedge rs' = (rs_1 \cdot a \cdot b \cdot rs_2) \wedge \\
&\qquad a.\mathit{loc} = x \wedge a.\mathit{flg} = 0 \wedge b.\mathit{flg} = 0 \wedge b.\mathcal{W} = \emptyset \\
\mathsf{split}_c(x, r, rs, rs') &\triangleq r \in rs(x).\mathcal{I} \wedge \exists r_2, r_3.\, rs' = rs \blacktriangleleft_x \blacktriangleleft_{\mathcal{I}} ((\bullet \setminus \{r\}) \uplus \{r_2, r_3\}) \\
\mathsf{sat}_c(rs, rs') &\triangleq rs = (rs_1 \cdot rs_2) \wedge rs' = (rs_1 \cdot rs_2') \wedge \mathsf{cconf}(rs_2') \wedge rs_2 \xrightarrow{pr}{}_* rs_2' \\
\mathsf{get}_c(x, r, rs, rs') &\triangleq rs = (rs_1 \cdot rs(x) \cdot rs_2) \wedge \mathsf{cconf}(rs(x) \cdot rs_2) \wedge rs' = rs \blacktriangleleft_x \blacktriangleleft_{\mathcal{I}} (\bullet \setminus \{r\})
\end{aligned}$$

$$\begin{aligned}
T_{\mathsf{c}}(\mathsf{send}(x)) &\triangleq \{(a, b) \mid \mathsf{wf}(b) \wedge (\mathsf{renun}_c(x, a, b) \vee \mathsf{set}_c(x, a, b) \vee \mathsf{ext}_c(x, a, b))\} \\
T_{\mathsf{c}}(\mathsf{change}(x, r)) &\triangleq \{(a, b) \mid \mathsf{wf}(b) \wedge (\mathsf{split}_c(x, r, a, b) \vee \mathsf{sat}_c(a, b) \vee \mathsf{get}_c(x, r, a, b))\}
\end{aligned}$$

Fig. 17. Definition of T_{c}, the transition relation for the chained channel implementation.

We define the relation formally as follows:

$$\begin{aligned}
rs \xrightarrow{pr} rs' \quad\triangleq\quad &\exists x, y, r. \\
&r \in rs(x).\mathcal{W} \wedge r \in rs(y).\mathcal{I} \wedge (x, y) \in \mathsf{ord}(rs) \wedge \\
&rs' = (rs \blacktriangleleft_x \blacktriangleleft_{\mathcal{W}} (\bullet \setminus \{r\})) \blacktriangleleft_y \blacktriangleleft_{\mathcal{I}} (\bullet \setminus \{r\})
\end{aligned}$$

We write $(x, y) \in \mathsf{seq}(rs)$ to say that the two addresses x and y appear adjacent in the sequence rs, and $(x, y) \in \mathsf{ord}(rs)$ to say just that they are ordered in rs.

Lenses. The notation $\blacktriangleleft$ is a *lens* allowing a single field of a chain to be updated without modifying the remainder of the chain. We use lenses to make our definitions more compact. Lenses are a notation borrowed from functional programming which we use to update one field of an object, while preserving the remainder of it. By object, we mean either a tuple or a map – we treat tuples as maps from field-names to values. Recall that we can also treat abstract chains as maps from locations to tuples as convenient.

We define the lens notation as follows. In the following, let x be the tuple / map we wish to update. Let i/j be values in the domain (i.e field names for a tuple). Let f be an expression with $\bullet$ standing for the value to be updated. Then the lens notation is defined as follows:

$$x \blacktriangleleft_i f \triangleq x[i \mapsto (f[x(i)/\bullet])] \qquad\qquad x \blacktriangleleft_i \blacktriangleleft_j f \triangleq x[i \mapsto (x(i) \blacktriangleleft_j f)]$$

On the left, the value associated with index / field name f in object x is updated to $f[x(i)/\bullet]$. On the right, we show we can stack lenses, allowing us to update fields deeper inside the object – here we update field j of field i of object x.

To give an example, in the preceding section, we had the following lens expression:

$$(rs \blacktriangleleft_x \blacktriangleleft_{\mathcal{W}} (\bullet \setminus \{r\})) \blacktriangleleft_y \blacktriangleleft_{\mathcal{I}} (\bullet \setminus \{r\})$$

Note that here we are using $\mathcal{W}$ / $\mathcal{I}$ as identifiers for particular tuple components. This expression denotes the chain rs, with identifier r removed from both the set of renounced resources at CNode x and the set of promised resource at CNode y.

Predicate definitions. As usual, we begin by defining the structure of the shared region. Abstract states have the form $\mathsf{Chain}(rs)$, where rs is an abstract chain. Actions have the form $\mathsf{send}(x)$ and $\mathsf{change}(x, r)$, where x is an address, and r a region identifier. The transition

$$\text{chainds}(x \cdot y \cdot rs) \triangleq x.loc \mapsto \{\text{prev} = y.loc; \text{flag} = x.flg\} * \text{chainds}(y \cdot rs)$$

$$\text{chainds}(x \cdot null) \triangleq x.loc \mapsto \{\text{prev} = \mathsf{NULL}; \text{flag} = x.flg\}$$

$$\text{resource}(\mathcal{I}, \mathcal{W}) \triangleq \exists Q \colon \mathcal{I} \to \mathsf{Prop}, \ R \colon \mathcal{W} \to \mathsf{Prop}.$$
$$\circledast_{i \in \mathcal{I}}. \ i \xmapsto{1/2} Q(i) * \circledast_{w \in \mathcal{W}}. \ w \xmapsto{1/2} R(w)$$
$$* \lfloor (\triangleright \circledast_{w \in \mathcal{W}}. \ R(w)) \twoheadrightarrow \triangleright \circledast_{i \in \mathcal{I}}. \lceil Q(i) \rceil \rfloor$$

$$\text{chainres}(x \cdot rs) \triangleq \text{resource}(x.\mathcal{I}, x.\mathcal{W} \uplus \{x.res \mid \neg x.flg\}) * \text{chainres}(rs)$$

$$\text{chainres}(null) \triangleq \mathsf{emp}$$

$$\text{uS}(rs) \triangleq \{x \mid (x, \_, \_, 0, \_) \in rs\}$$

$$\text{uC}(rs) \triangleq \{(x, i) \mid (x, \_, \mathcal{I}, \_, \_) \in rs \land i \in \mathcal{I} \land \neg\exists(y, \_, \_, \_, \mathcal{W}) \in rs. \ i \in \mathcal{W}\}$$

$$\text{unused}(r, rs) \triangleq (\circledast \, x \notin \text{uS}(rs). [\text{send}(x)]_1^r) * (\circledast(x, r') \notin \text{uC}(rs). [\text{change}(x, r')]_1^r)$$

Fig. 18. Predicates used in defining the state of a region.

relation T_{c} is defined in Fig. 17. We assume that physical addresses are used uniquely, so where convenient we use chains as finite functions of type

$$\mathsf{Addr} \xrightarrow{fin} (\mathsf{RId} \times \mathcal{P}(\mathsf{RId}) \times \{\mathsf{High}, \mathsf{Low}\} \times \mathcal{P}(\mathsf{RId}))$$

The transition relation defines six kinds of transitions in Fig. 17. For send we have renunciation, which adds an element to $\mathcal{W}$; setting the flag; and extending the chain, which creates a new CNode b. For change we have splitting; satisfying the renounced resource set, which sets $\mathcal{W}$ to $\emptyset$ and pulls the resources out of earlier chain elements; and pulling out a resource.

To translate from an abstract chain to a concrete invariant, we define three predicates: chainds, chainres, and unused (defined in Fig. 18). The predicate chainds represents the linked list underpinning the implementation. Each link in the chain has the appropriate prev and flag values set determined by the corresponding CNode in the chain.

The predicate chainres represents the resources that are communicated through the chain. The key predicate is resource, which ties together a set of promised resources $\mathcal{I}$ and a set of resources waited for $\mathcal{W}$. Note that the set $\mathcal{W}$ of resources waited for includes both renounced resources, *and* the resource supplied by the preceding channel – these are unioned by chainres. The core of the resource predicate is the following assertion:

$$\lfloor (\triangleright \circledast_{w \in \mathcal{W}}. \ R(w)) \twoheadrightarrow \triangleright \circledast_{i \in \mathcal{I}}. \lceil Q(i) \rceil \rfloor$$

Here Q and R map identifiers to propositions. Leaving aside explicit stabilization and $\triangleright$, this assertion has a straightforward intuition: supplying all the resources waited for (those with identifiers in $\mathcal{W}$) results in the resources promised (those with identifiers in $\mathcal{I}$).

When there are no resources waited for, i.e. $\mathcal{W} = \emptyset$, the resource predicate can be simplified to just the promised resources:

LEMMA 6.1. $\text{resource}(\mathcal{I}, \emptyset) \ \sqsubseteq \ \circledast_{i \in \mathcal{I}}. \exists Q \colon \mathsf{Prop}. \ i \xmapsto{1/2} Q(i) * \lceil \triangleright Q(i) \rceil$

The unused predicate stands for the set of unused permissions (similar to changes in the previous proof). We define this using $\text{uS}(rs)$, the set of used send permissions, and $\text{uC}(rs)$, the set of used change permissions.

The representation function for the region, I_{c}, is defined as follows:

$$I_{\mathsf{c}}(r)(\mathsf{Chain}(rs)) \triangleq \text{chainds}(rs) * \text{chainres}(rs) * \text{unused}(r, rs)$$

As before, we use a shorthand for the region assertion in our definitions and proofs:

$$\mathsf{oreg}(r, S) \triangleq \mathsf{region}(S, T_\mathsf{c}, I_\mathsf{c}(r), r)$$

We can now define the send, recv, and ordering predicates.

$$\mathsf{send}(x, P) \triangleq \exists r_1, r_2.\, \mathsf{oreg}\left(r_1, \{\mathsf{Chain}(rs) \mid \mathsf{wf}(rs) \wedge rs(x) = (r_2, \_, 0, \_)\}\right)$$
$$* \; r_2 \overset{1/2}{\longmapsto} P * [\mathsf{send}(x)]_1^{r_1}$$

$$\mathsf{recv}(x, P) \triangleq \exists r_1, r_2.\, \mathsf{oreg}(r_1, \{\mathsf{Chain}(rs) \mid rs(x) = (\_, \mathcal{I}, \_, \_) \wedge r_2 \in \mathcal{I} \wedge \mathsf{wf}(rs)\})$$
$$* \; r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(x, r_2)]_1^{r_1}$$

$$x \prec y \triangleq \exists r.\, \mathsf{oreg}(r, \{\mathsf{Chain}(rs) \mid (y, x) \in \mathsf{ord}(rs)\})$$

The structure of the send and recv predicates is similar to the un-chained proof. Both send and recv contain saved propositions for their proposition parameter P. The predicate definitions ensure the identifier for this saved proposition is embeddeded into the abstract chain correctly. Meanwhile, the $\prec$ predicate is a straightforward lifting of the ord predicate on abstract chains.

6.3. Proving `signal`, `wait` and `extend`

The majority of the proof concerns manipulations of resource obligations, rather than reads and writes to the underlying data-structure. To help with proof clarity, as far as possible we factor reads and writes into small, separate specifications.

Proving `signal`. The sketch-proof is shown in Fig. 19 – it is similar in structure to the one in §5.3. The main additional challenge is to show that resources are supplied to the appropriate point in the chain. To do this, we use the following lemma, which says that supplying the resource $\lfloor P \rfloor$ and an associated saved proposition removes the need to supply r. This is then sufficient to allow the flag to be set.

LEMMA 6.2. $\quad r \overset{1/2}{\longmapsto} P * \lfloor P \rfloor * \mathsf{resource}(\mathcal{I}, \mathcal{W} \uplus \{r\}) \;\sqsubseteq\; \mathsf{resource}(\mathcal{I}, \mathcal{W})$

PROOF. Given in Appendix B. □

In the proof of `signal`, this lemma is used to show that the appropriate resource has been supplied (Fig. 19, line 12). By factoring logical resource transfer away from the physical signalling, we simplify the proof structure considerably.

The rest of the proof consists of manipulating predicates. We pull out the node associated with `x` and set the flag (lines 1–8). Once the resource has been supplied on line 12, the remainder of the proof closes the region again.

Proving `wait`. The sketch-proof is given in Fig. 20. The three most important steps are checking that all preceding flags in the chain are set (lines 5–12), checking that renounced resources have been supplied (line 15), and retrieving the resource from the chain (line 20). The last two of these require helper lemmas, given below.

Resources that are renounced earlier in the chain can be used to satisfy required resources later in the chain. These resources are represented by the set $\mathcal{W}$ in the abstract state of a `cnode`. Renounced resources need not be supplied when `signal` is called, but they must be available before `wait` returns. To ensure this, the implementation of `wait` checks all the preceding flags in the chain. Once all preceding flags are set, all the resources should be available. However, proving this is subtle, because renounced resources may themselves be satisfied by resources renounced earlier in the chain.

$_1$ $\Big\{ \mathsf{send}(\mathrm{x}, P) * \lfloor P \rfloor \Big\}$

$_2$ *// Definition of* send.

$_3$ $\left\{ \begin{array}{l} \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{send}(\mathrm{x})]_1^{r_1} * \lfloor P \rfloor * \\ \mathsf{oreg}\,(r_1, \{\mathsf{Chain}(rs) \mid \mathsf{wf}(rs) \land rs(\mathrm{x}) = (r_2, \mathcal{I}, 0, \mathcal{W})\}) \end{array} \right\}$

$_4$ *// Enter the region.*

$_5$ $\left\langle \begin{array}{l} \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{send}(\mathrm{x})]_1^{r_1} * \lfloor P \rfloor * \\ \rhd \Big(\mathsf{chainds}(rs) * \mathsf{chainres}(rs) * \mathsf{unused}(r_1, rs) \land \mathsf{wf}(rs) \land rs(\mathrm{x}) = (r_2, \mathcal{I}, 0, \mathcal{W}) \Big) \end{array} \right\rangle$

$_6$ *// Split up using chain axioms.*

$_7$ $\left\langle \begin{array}{l} \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{send}(\mathrm{x})]_1^{r_1} * \lfloor P \rfloor * \\ \rhd \left(\begin{array}{l} \exists rs_1, rs_2, a.\, \mathsf{chainds}(rs_1) * \mathrm{x} \mapsto \{\mathrm{prev} = \mathsf{hd}(rs_2).loc; \mathrm{flag} = 0\} * \mathsf{chainds}(rs_2) * \\ \mathsf{chainres}(rs) * \mathsf{unused}(r_1, rs) \land rs = rs_1 \cdot a \cdot rs_2 \land \mathsf{wf}(rs) \land a = (\mathrm{x}, r_2, \mathcal{I}, 0, \mathcal{W}) \end{array} \right) \end{array} \right\rangle$

$_8$ `x->flag = 1;` *// Set the flag, drop the* $\rhd$ *using Aframe.*

$_9$ $\left\langle \begin{array}{l} \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{send}(\mathrm{x})]_1^{r_1} * \lfloor P \rfloor * \\ \exists rs_1, rs_2, a.\, \mathsf{chainds}(rs_1) * \mathrm{x} \mapsto \{\mathrm{prev} = \mathsf{hd}(rs_2).loc; \mathrm{flag} = 1\} * \mathsf{chainds}(rs_2) * \\ \mathsf{chainres}(rs) * \mathsf{unused}(r_1, rs) \land rs = rs_1 \cdot a \cdot rs_2 \land \mathsf{wf}(rs) \land a = (\mathrm{x}, r_2, \mathcal{I}, 0, \mathcal{W}) \end{array} \right\rangle$

$_{10}$ *// Pull out* resource *predicate from* chainres *predicate.*

$_{11}$ $\left\langle \begin{array}{l} \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{send}(\mathrm{x})]_1^{r_1} * \lfloor P \rfloor * \\ \exists rs_1, rs_2, a.\, \mathsf{chainds}(rs_1) * \mathrm{x} \mapsto \{\mathrm{prev} = \mathsf{hd}(rs_2).loc; \mathrm{flag} = 1\} * \mathsf{chainds}(rs_2) \\ \qquad\qquad * \mathsf{chainres}(rs_1) * \mathsf{resource}(\mathcal{I}, \mathcal{W} \uplus \{r_2\}) * \mathsf{chainres}(rs_2) \\ \qquad * \mathsf{unused}(r_1, rs) \land rs = rs_1 \cdot a \cdot rs_2 \land \mathsf{wf}(rs) \land a = (\mathrm{x}, r_2, \mathcal{I}, 0, \mathcal{W}) \end{array} \right\rangle$

$_{12}$ *// Apply Lemma* 6.2 *-- use* $\lfloor P \rfloor$ *to remove* r_2 *from the second parameter of* resource.

$_{13}$ $\left\langle \begin{array}{l} \exists r_1, r_2.\, [\mathsf{send}(\mathrm{x})]_1^{r_1} * \\ \exists rs_1, rs_2, a.\, \mathsf{chainds}(rs_1) * \mathrm{x} \mapsto \{\mathrm{prev} = \mathsf{hd}(rs_2).loc; \mathrm{flag} = 1\} * \mathsf{chainds}(rs_2) \\ \qquad\qquad * \mathsf{chainres}(rs_1) * \mathsf{resource}(\mathcal{I}, \mathcal{W}) * \mathsf{chainres}(rs_2) \\ \qquad * \mathsf{unused}(r_1, rs) \land rs = rs_1 \cdot a \cdot rs_2 \land \mathsf{wf}(rs) \land a = (\mathrm{x}, r_2, \mathcal{I}, 0, \mathcal{W}) \end{array} \right\rangle$

$_{14}$ *// Collapse the chain, add* $\rhd$.

$_{15}$ $\left\langle \begin{array}{l} \exists r_1, r_2. \\ \rhd (\mathsf{chainds}(rs') * \mathsf{chainres}(rs') * \mathsf{unused}(r_1, rs') \land rs' = rs[\mathrm{x} \mapsto (r_2, \mathcal{I}, 1, \mathcal{W})]) \end{array} \right\rangle$

$_{16}$ *// Close the region using transition* set.

$_{17}$ *// Well-formed structure of chain hasn't changed.*

$_{18}$ $\Big\{ \exists r_1, r_2.\, \mathsf{oreg}(r_1, \{\mathsf{Chain}(rs) \mid \mathsf{wf}(rs) \land rs(\mathrm{x}) = (r_2, \mathcal{I}, 1, \mathcal{W})\}) \Big\}$

$_{19}$ *// Delete all assertions using weakening.*

$_{20}$ $\Big\{ \mathsf{emp} \Big\}$

Fig. 19. Sketch-proof for `signal` with out-of-order signalling.

1. $\left\{ \mathsf{recv}(\mathsf{x}, P) \right\}$
2. `wait(x){`
3.   `chan *c = x;`
4.   *// Definition of* `recv`*.*
5. $\left\{ \begin{array}{c} \exists r_1, r_2.\, \mathsf{oreg}\,(r_1, \{\mathsf{Chain}(rs) \mid \exists \mathcal{I}'.\, rs(\mathsf{x}) = (\_, \mathcal{I}', \_, \_) \wedge r_2 \in \mathcal{I}' \wedge \mathsf{wf}(rs)\}) \\ * \, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(\mathsf{x}, r_2)]_1^{r_1} \wedge \mathsf{c} = \mathsf{x} \end{array} \right\}$
6.   `while(c != NULL) {`
7. $\left\{ \begin{array}{l} \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(\mathsf{x}, r_2)]_1^{r_1} * \\ \mathsf{oreg}\left(r_1, \left\{ \mathsf{Chain}(rs) \; \middle| \; \begin{array}{c} \exists \mathcal{I}'.\, \mathsf{wf}(rs) \wedge \exists i.\, rs[i] = (\mathsf{x}, \_, \mathcal{I}' \uplus r_2, \_, \_) \wedge \\ \exists m < \mathrm{len}(rs).\, rs[m] = (\mathsf{c}, \_, \_, \_, \_) \wedge \\ \forall j.\, i \le j < m \implies rs[j] = (\_, \_, \_, 1, \_) \end{array} \right\} \right) \end{array} \right\}$
8.   `while(!c->flag) skip;`
9.   `c = c->prev;`
10.   `}`
11. *// Flags up to* `x` *are now set. Stable because flags cannot get unset.*
12. $\left\{ \begin{array}{l} \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(\mathsf{x}, r_2)]_1^{r_1} * \\ \mathsf{oreg}\left(r_1, \left\{ \mathsf{Chain}(rs) \; \middle| \; \begin{array}{c} \exists \mathcal{I}'.\, \mathsf{wf}(rs) \wedge \exists i.\, rs[i] = (\mathsf{x}, \_, \mathcal{I}' \uplus r_2, \_, \_) \wedge \\ \forall j.\, i \le j < \mathrm{len}(rs) \implies rs[j] = (\_, \_, \_, 1, \_) \end{array} \right\} \right) \end{array} \right\}$
13. $\langle \mathsf{skip} \rangle;$ *// Enter the region and drop* $\rhd$ *using AFrame.*
14. $\left\langle \begin{array}{l} \exists rs_1, rs_2, a.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(\mathsf{x}, r_2)]_1^{r_1} * \mathsf{chainds}(rs) * \mathsf{chainres}(rs) * \mathsf{unused}(r_1, rs) \wedge \\ \mathsf{wf}(rs) \wedge rs = rs_1 \cdot a \cdot rs_2 \wedge a = (\mathsf{x}, r, \mathcal{I}' \uplus r_2, \_, \mathcal{W}) \wedge \mathsf{ctrue}(a \cdot rs_2) \end{array} \right\rangle$
15. *// Apply Lemma* 6.3 *to convert from* ctrue *to* cconf*, denoting that the*
16. *// resource* P *is now available to be claimed.*
17. $\left\langle \begin{array}{l} \exists rs_1, rs_2, rs_2', a.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(\mathsf{x}, r_2)]_1^{r_1} * \\ \mathsf{chainds}(rs) * \mathsf{chainres}(rs_1 \cdot rs_2') * \mathsf{unused}(r_1, rs) \wedge \mathsf{wf}(rs) \wedge rs = rs_1 \cdot a \cdot rs_2 \wedge \\ a = (\mathsf{x}, r, \mathcal{I}' \uplus r_2, \_, \mathcal{W}) \wedge \mathsf{cconf}(rs_2') \wedge a \cdot rs_2 \overset{pr}{\rightarrow}^* rs_2' \end{array} \right\rangle$
18. *// Pull out the* $\mathsf{resource}$ *predicate for chain node* `x`*.*
19. $\left\langle \begin{array}{l} \exists rs_1, rs_2, rs_2', a.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(\mathsf{x}, r_2)]_1^{r_1} * \\ \mathsf{chainds}(rs) * \mathsf{chainres}(rs_1) * \mathsf{resource}(\mathcal{I}' \uplus r_2, \emptyset) * \mathsf{chainres}(rs_2') * \mathsf{unused}(r_1, rs) \wedge \\ \mathsf{wf}(rs) \wedge rs = rs_1 \cdot a \cdot rs_2 \wedge a = (\mathsf{x}, r, \mathcal{I}' \uplus r_2, \_, \mathcal{W}) \wedge \mathsf{cconf}(rs_2') \wedge a \cdot rs_2 \overset{pr}{\rightarrow}^* rs_2' \end{array} \right\rangle$
20. *// Apply Lemma* 6.4 *to retrieve the resource* $\lceil \rhd P \rceil$*.*
21. $\left\langle \begin{array}{l} \lceil \rhd P \rceil * \exists rs_1, rs_2, rs_2', a.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(\mathsf{x}, r_2)]_1^{r_1} * \\ \mathsf{chainds}(rs) * \mathsf{chainres}(rs_1) * \mathsf{resource}(\mathcal{I}', \emptyset) * \mathsf{chainres}(rs_2') * \mathsf{unused}(r_1, rs) \wedge \\ \mathsf{wf}(rs) \wedge rs = rs_1 \cdot a \cdot rs_2 \wedge a = (\mathsf{x}, r, \mathcal{I}' \uplus r_2, \_, \mathcal{W}) \wedge \mathsf{cconf}(rs_2') \wedge a \cdot rs_2 \overset{pr}{\rightarrow}^* rs_2' \end{array} \right\rangle$
22. *// Close the region, use transitions* sat *and* get*.*
23. $\left\{ \begin{array}{l} \exists r_1, r_2.\, \lceil \rhd P \rceil * \\ \mathsf{oreg}\left(r_1, \left\{ \mathsf{Chain}(rs') \; \middle| \; \begin{array}{c} \mathsf{wf}(rs') \wedge rs = rs_1 \cdot a \cdot rs_2 \wedge a \cdot rs_2 \overset{pr}{\rightarrow}^* rs_2' \wedge \\ \mathsf{cconf}(rs_2') \wedge rs' = (rs_1 \cdot rs_2') \blacktriangleleft_x \blacktriangleleft_{\mathcal{I}} (\bullet \setminus \{r_2\}) \end{array} \right\} \right) \end{array} \right\}$
24. `}`*// Use* $\lceil \rhd P \rceil \Rightarrow \rhd \lceil P \rceil$*, and delete redundant assertions using weakening.*
25. $\left\{ \rhd \lceil P \rceil \right\}$

Fig. 20. Sketch-proof for `wait` with out-of-order signalling.

To establish the required resources are available, we use the following lemma. This says that a chainres predicate for a chain where all the flags are set can be transformed into one where pending resources have been resolved (asserted by ctrue and cconf respectively).

LEMMA 6.3. $\mathsf{chainres}(rs) \wedge \mathsf{wf}(rs) \wedge \mathsf{ctrue}(rs)$

$\qquad \sqsubseteq\ \exists rs'.\, \mathsf{chainres}(rs') \wedge \mathsf{cconf}(rs') \wedge rs \overset{pr}{\rightarrow}^* rs' \wedge \mathsf{wf}(rs')$

PROOF. Given in Appendix B. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ □

We apply this lemma on line 15 of the sketch-proof.

Once we've established that the resources are available, we use the following lemma to extract the appropriate resource from the resource predicate:

LEMMA 6.4. $\mathsf{resource}(\mathcal{I} \uplus r_2, \emptyset) * r_2 \overset{1/2}{\Longmapsto} P \ \sqsubseteq\ \mathsf{resource}(\mathcal{I}, \emptyset) * \lceil \triangleright P \rceil$

PROOF. Given in Appendix B $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad$ □

This lemma says that, given resource and an identifier r_2 in $\mathcal{I}$ such that all required resources are available, the resource $\triangleright P$ associated with r_2 can be retrieved. We apply this lemma on line 20 of the sketch-proof.

Proving extend. The sketch-proof is given in Fig. 21. The key steps in the proof are creating a new node to add to the chain (lines 3–9), stitching the new node into the chain itself (line 12), then satisfying the required invariants for the region (lines 18–22).

It is important that new saved propositions are *fresh* – that is, their identifiers have not been used elsewhere in the chain. We use the following lemma to show new identifiers are fresh:

LEMMA 6.5.

$\{\mathsf{oreg}(r, \mathcal{T}) * r' \overset{1}{\Longmapsto} P\} \quad \langle\mathtt{skip}\rangle \quad \{\mathsf{oreg}(r, \mathcal{T} \cap \{\mathsf{Chain}(rs) \mid r' \notin rs\}) * r' \overset{1}{\Longmapsto} P\}$

PROOF. Each identifier r'' used in rs is associated with a fractional saved proposition $r'' \overset{1/2}{\Longmapsto} P$. We case-split on the finite set of possible equalities and appeal to the linearity of saved propositions (Property 2). The skip is required by iCAP because we access the internal state of the region r. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ □

The following lemma uses this freshness property, along with the freshness of allocated locations to show that we can we can retrieve the required permissions from unused (the 'library' of unused permissions). We use this lemma on line 16 of the sketch-proof.

LEMMA 6.6.

$rs = rs_1 \cdot (x, r_2, \mathcal{I}, 0, \mathcal{W}) \cdot rs_2 \wedge \mathsf{unused}(r_1, rs) \wedge z, r', r'' \notin rs$

$\quad \sqsubseteq\quad \mathsf{unused}(r_1, rs_1 \cdot (x, r_2, \mathcal{I}, 0, \mathcal{W}) \cdot (z, r', \{r''\}, 0, \emptyset) \cdot rs_2) * [\mathsf{send}(z)]_1^{r_1} * [\mathsf{change}(z, r'')]_1^{r_1}$

PROOF. Appeal to the definition of unused. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ □

On line 18, we close the region. The resulting chain is well-formed because the new region has no elements in its renunciation set $\mathcal{W}$, and the rest of the chain is preserved. The chain is stable because we hold the send permission on x and z, meaning these channels cannot be extended or renounced.

Proving newchan. Omitted: this proof is similar to extend. However it is simpler: we only need to construct a new point in the chain, and not update the existing chain to take account of it.

1 $\left\{ \mathsf{send}(\mathbf{x}, P) * \circledast_{e \in E}\, e \prec \mathbf{x} * \circledast_{l \in L}\, \mathbf{x} \prec l \right\}$

2 `extend(x){` *// Frame off order predicates*

3   `chan *z = new(chan); z->flag = 0;`

4   `z->prev = x->prev;` *// Stable as thread holds exclusive $[\mathsf{send}(x)]$ permission.*

5 $\left\{ \begin{array}{l} \exists r_1, r_2, x'.\ r_2 \overset{1/2}{\Longmapsto} P * [\mathsf{send}(\mathbf{x})]_1^{r_1} * \mathbf{z} \mapsto \{\mathrm{prev} = x'; \mathrm{flag} = 0\} * \\ \mathsf{oreg}\left(r_1, \{\mathsf{Chain}(rs) \mid \mathsf{wf}(rs) \wedge rs(\mathbf{x}) = (r_2, \_, 0, \_) \wedge (\mathbf{x}, x') \in \mathsf{seq}(rs) \wedge \mathbf{z} \notin rs\}\right) \end{array} \right\}$

6   `skip; skip;` *// Make saved props, use Lemma 6.5 to show they are fresh.*

7 $\left\{ \begin{array}{l} \exists r_1, r_2, x', r', r''.\ r_2 \overset{1/2}{\Longmapsto} P * [\mathsf{send}(\mathbf{x})]_1^{r_1} * \mathbf{z} \mapsto \{\mathrm{prev} = x'; \mathrm{flag} = 0\} * r' \overset{1}{\Longmapsto} Q * r'' \overset{1}{\Longmapsto} Q * \\ \mathsf{oreg}\left(r_1, \{\mathsf{Chain}(rs) \mid \mathsf{wf}(rs) \wedge rs(\mathbf{x}) = (r_2, \_, 0, \_) \wedge (\mathbf{x}, x') \in \mathsf{seq}(rs) \wedge \mathbf{z}, r', r'' \notin rs\}\right) \end{array} \right\}$

8   *// Create a node in the chain.*

9 $\left\{ \begin{array}{l} \exists r_1, r_2, x', r', r''.\ r_2 \overset{1/2}{\Longmapsto} P * [\mathsf{send}(\mathbf{x})]_1^{r_1} * \\ \mathsf{oreg}\left(r_1, \{\mathsf{Chain}(rs) \mid \mathsf{wf}(rs) \wedge rs(\mathbf{x}) = (r_2, \_, 0, \_) \wedge (\mathbf{x}, x') \in \mathsf{seq}(rs) \wedge \mathbf{z}, r', r'' \notin rs\}\right) * \\ \qquad r' \overset{1/2}{\Longmapsto} Q * r'' \overset{1/2}{\Longmapsto} Q * \mathbf{z} \mapsto \{\mathrm{prev} = x'; \mathrm{flag} = 0\} * \mathsf{resource}(\{r''\}, \{r'\}) \end{array} \right\}$

10   *// Enter the region, split the chain up.*

11 $\left\langle \begin{array}{l} \exists r_1, r_2, x', r', r'', h, rs_1, rs_2.\ r_2 \overset{1/2}{\Longmapsto} P * [\mathsf{send}(\mathbf{x})]_1^{r_1} * \\ \triangleright \left(\begin{array}{l} \mathsf{chainds}(rs_1) * \mathbf{x} \mapsto \{\mathrm{prev} = x'; \mathrm{flag} = 0\} * \mathsf{chainds}(rs_2) * \mathsf{chainres}(rs) * \\ \mathsf{unused}(r_1, rs) \wedge rs = rs_1 \cdot e \cdot rs_2 \wedge \mathsf{wf}(rs) \wedge e = (\mathbf{x}, r_2, \mathcal{I}, 0, \mathcal{W}) \wedge \mathbf{z}, r', r'' \notin rs \end{array} \right) * \\ \qquad r' \overset{1/2}{\Longmapsto} Q * r'' \overset{1/2}{\Longmapsto} Q * \mathbf{z} \mapsto \{\mathrm{prev} = \mathsf{hd}(rs_2).loc; \mathrm{flag} = 0\} * \mathsf{resource}(\{r''\}, \{r'\}) \end{array} \right\rangle$

12   `x->prev = z;` *// Stitch the new cnode into the chain using LPoints / AFrame.*

13 $\left\langle \begin{array}{l} \exists r_1, r_2, x', r', r'', h, rs_1, rs_2.\ r_2 \overset{1/2}{\Longmapsto} P * [\mathsf{send}(\mathbf{x})]_1^{r_1} * \\ \mathsf{chainds}(rs_1) * \mathbf{x} \mapsto \{\mathrm{prev} = \mathbf{z}; \mathrm{flag} = 0\} * \mathsf{chainds}(rs_2) * \mathsf{chainres}(rs) * \\ \mathsf{unused}(r_1, rs) \wedge rs = rs_1 \cdot e \cdot rs_2 \wedge \mathsf{wf}(rs) \wedge e = (\mathbf{x}, r_2, \mathcal{I}, 0, \mathcal{W}) \wedge \mathbf{z}, r', r'' \notin rs * \\ r' \overset{1/2}{\Longmapsto} Q * r'' \overset{1/2}{\Longmapsto} Q * \mathbf{z} \mapsto \{\mathrm{prev} = \mathsf{hd}(rs_2).loc; \mathrm{flag} = 0\} * \mathsf{resource}(\{r''\}, \{r'\}) \end{array} \right\rangle$

14   *// Fold the* $\mathsf{chainds}$ *predicate back up.*

15 $\left\langle \begin{array}{l} \exists r_1, r_2, x', r', r'', h, rs_1, rs_2.\ r_2 \overset{1/2}{\Longmapsto} P * [\mathsf{send}(\mathbf{x})]_1^{r_1} * r' \overset{1/2}{\Longmapsto} Q * r'' \overset{1/2}{\Longmapsto} Q \\ \mathsf{chainds}(rs_1 \cdot e \cdot e' \cdot rs_2) * \mathsf{chainres}(rs_1 \cdot e \cdot e' \cdot rs_2) * \mathsf{unused}(r_1, rs) \wedge \\ rs = rs_1 \cdot e \cdot rs_2 \wedge \mathsf{wf}(rs) \wedge e = (\mathbf{x}, r_2, \mathcal{I}, 0, \mathcal{W}) \wedge e' = (\mathbf{z}, r', \{r''\}, 0, \emptyset) \wedge \mathbf{z}, r', r'' \notin rs \end{array} \right\rangle$

16   *// Get* $[\mathsf{send}(z)] * [\mathsf{change}(z, r'')]$ *from* $\mathsf{unused}(r_1, rs)$ *using Lemma 6.6.*

17 $\left\langle \begin{array}{l} \exists r_1, r_2, r', r'', h, rs_1, rs_2. \\ \mathsf{chainds}(rs_1 \cdot e \cdot e' \cdot rs_2) * \mathsf{chainres}(rs_1 \cdot e \cdot e' \cdot rs_2) * \mathsf{unused}(r_1, rs_1 \cdot e \cdot e' \cdot rs_2) \\ \qquad \wedge rs = rs_1 \cdot e \cdot rs_2 \wedge \mathsf{wf}(rs) \wedge e = (\mathbf{x}, r_2, \mathcal{I}, 0, \mathcal{W}) \wedge e' = (\mathbf{z}, r', \{r''\}, 0, \emptyset) \\ \qquad \wedge r_2 \overset{1/2}{\Longmapsto} P * [\mathsf{send}(\mathbf{x})]_1^{r_1} * r' \overset{1/2}{\Longmapsto} Q * r'' \overset{1/2}{\Longmapsto} Q * [\mathsf{send}(\mathbf{z})]_1^{r_1} * [\mathsf{change}(\mathbf{z}, r'')]_1^{r_1} \end{array} \right\rangle$

18   *// Close the region using the* ext *transition.*

19 $\left\{ \begin{array}{l} \exists r_1, r_2, r', r'', h. \\ \mathsf{oreg}\left(r_1, \{\mathsf{Chain}(rs') \mid rs' = rs_1 \cdot (\mathbf{x}, r_2, \_, 0, \_) \cdot (\mathbf{z}, r', \mathcal{I}', 0, \_) \cdot rs_2 \wedge r'' \in \mathcal{I}' \wedge \mathsf{wf}(rs')\}\right) \\ \qquad \wedge r_2 \overset{1/2}{\Longmapsto} P * [\mathsf{send}(\mathbf{x})]_1^{r_1} * r' \overset{1/2}{\Longmapsto} Q * r'' \overset{1/2}{\Longmapsto} Q * [\mathsf{send}(\mathbf{z})]_1^{r_1} * [\mathsf{change}(\mathbf{z}, r'')]_1^{r_1} \end{array} \right\}$

20   `return (z,x);`

21 `}` *// Frame on order predicates. Fold invariant into predicates.*

22 $\left\{ \mathsf{send}(\mathbf{z}, Q) * \mathsf{recv}(\mathbf{z}, Q) * \mathsf{send}(\mathbf{x}, P) * \mathbf{z} \prec \mathbf{x} * \circledast_{e \in E}\, e \prec \mathbf{z} * \circledast_{l \in L}\, \mathbf{x} \prec l \right\}$

Fig. 21. Sketch-proof of `extend` with out-of-order signalling.

6.4. Proving Renunciation and Splitting Axioms

Renunciation. The axiom is defined as follows:

$$\{\mathsf{recv}(x, P) * \mathsf{send}(y, Q) * x \prec y\} \quad \langle\mathtt{skip}\rangle \quad \{\mathsf{send}(y, P \twoheadrightarrow Q)\}$$

The sketch-proof is given in Fig. 22. Internally each predicate contains a view on the same shared region, and the first step of the proof consists of conjoining these three views to give a single stable view on the shared structure (line 3). The remaining steps are supplying the renounced resource to the shared region (line 13), and closing the region to give a new send predicate (line 15).

In order to conjoin the regions arising from the send, recv and the order predicates, they need to operate over the same region. Although the predicates do not expose region names, we know from order predicates that all of the regions share common elements in their chain addresses. We therefore use an extra lemma to show that pairs of such regions must be the same:

LEMMA 6.7.

$$\{\mathsf{oreg}(r, \{\mathsf{Chain}(rs) \mid x \in rs\}) * \mathsf{oreg}(r', \{\mathsf{Chain}(rs') \mid x \in rs'\})\} \quad \langle\mathtt{skip}\rangle \quad \{r = r'\}$$

PROOF. Given in Appendix B. □

We use this lemma on line 3, Fig. 22. The conjoined region that arises from this lemma (line 6) is stable because elements cannot be reordered with respect to each other once they are in the chain, and because exclusive [chain] and [send] permissions are held for x and y respectively.

When we push the renounced resource into the $\mathsf{resource}$ predicate (line 13) we use the following lemma to show that the renunciation set $\mathcal{W}$ is updated appropriately:

LEMMA 6.8.　$\mathsf{resource}(\mathcal{I}, \mathcal{W} \uplus \{r\}) * r \overset{1/2}{\Longmapsto} Q * r' \overset{1/2}{\Longmapsto} P$

$$\sqsubseteq \quad \exists r''. \, \mathsf{resource}(\mathcal{I}, \mathcal{W} \uplus \{r', r''\}) * r'' \overset{1/2}{\Longmapsto} (P \twoheadrightarrow Q)$$

PROOF. Given in Appendix B. □

Note that the identifier r used for sending resources is replaced with a fresh identifier r'' because the associated invariant is changed from Q to $P \twoheadrightarrow Q$. Internally this corresponds to deleting one saved proposition through weakening, and then creating another.

On line 15 we close the region. The resulting chain is well-formed because the identifier we selected was previously unused for renunciation – we get this from the definition of unused. Furthermore, the remainder of the chain stays the same. The resulting chain is trivially stable because the exclusive [send] permission is held.

Splitting. The axiom is defined as follows:

$$\{\mathsf{recv}(a, P) * \lfloor \lceil P \rceil \twoheadrightarrow (P_1 * P_2) \rfloor\} \quad \langle\mathtt{skip}\rangle \quad \{\mathsf{recv}(a, P_1) * \mathsf{recv}(a, P_2)\}$$

A sketch-proof is given in Fig. 23. The key step is splitting one element of the promised resource set $\mathcal{I}$ for a node (line 6). To do this, we use the following lemma, which states that the saved proposition r_2 can be exchanged for new saved propositions r_3 and r_4 if the appropriate resource $\lfloor \lceil P \rceil \twoheadrightarrow (P_1 * P_2) \rfloor$ is supplied.

LEMMA 6.9.　$r_2 \in rs(x).\mathcal{I} \wedge r_2 \overset{1/2}{\Longmapsto} P * \lfloor \lceil P \rceil \twoheadrightarrow (P_1 * P_2) \rfloor * \mathsf{chainres}(rs)$

$$\sqsubseteq \quad \exists rs', r_3, r_4. \, r_3, r_4 \notin rs \wedge rs' = rs \blacktriangleleft_x \blacktriangleleft_{\mathcal{I}} ((\bullet \setminus r_2) \uplus \{r_3, r_4\}) \wedge$$

$$r_3 \overset{1/2}{\Longmapsto} P_1 * r_4 \overset{1/2}{\Longmapsto} P_2 * \mathsf{chainres}(rs')$$

1 $\left\{ \mathsf{recv}(x, P) * \mathsf{send}(y, Q) * x \prec y \right\}$

2 $\langle \mathsf{skip} \rangle$; // *Definition of* recv, send, $\prec$. *Use Lemma* 6.7 *to establish*

3   // *that all three predicates refer to the same region.*

4 $\left\{ \begin{array}{l} \exists r, r_2, r_3.\, \mathsf{oreg}(r, \{\mathsf{Chain}(rs) \mid (y, x) \in \mathsf{ord}(rs)\}) * \\ \left(\begin{array}{c} \mathsf{oreg}(r, \{\mathsf{Chain}(rs) \mid rs(x) = (\_, \mathcal{I}, \_, \_) \wedge r_2 \in \mathcal{I} \wedge \mathsf{wf}(rs)\}) \\ * r_2 \stackrel{1/2}{\Longmapsto} P * [\mathsf{change}(x, r_2)]_1^r \end{array} \right) * \\ \left(\begin{array}{c} \mathsf{oreg}\left(r, \{\mathsf{Chain}(rs) \mid \mathsf{wf}(rs) \wedge rs(y) = (r_3, \_, 0, \_)\}\right) \\ * r_3 \stackrel{1/2}{\Longmapsto} Q * [\mathsf{send}(y)]_1^r \end{array} \right) \end{array} \right\}$

5 // *Conjunction on regions. Stable because of the permissions held.*

6 $\left\{ \begin{array}{l} \exists r, r_2, r_3.\, r_2 \stackrel{1/2}{\Longmapsto} P * [\mathsf{change}(x, r_2)]_1^r * r_3 \stackrel{1/2}{\Longmapsto} Q * [\mathsf{send}(y)]_1^r * \\ \mathsf{oreg}\left(r, \left\{ \mathsf{Chain}(rs) \,\middle|\, \begin{array}{l} rs = rs_1 \cdot e \cdot rs_2 \cdot e' \cdot rs_3 \wedge \mathsf{wf}(rs) \wedge \\ e = (y, r_3, \_, 0, \_) \wedge e' = (x, \_, \mathcal{I}, \_, \_) \wedge r_2 \in \mathcal{I} \end{array} \right\} \right) \end{array} \right\}$

7 $\langle \mathsf{skip} \rangle$; // *Open the region, drop* $\triangleright$ *using AFrame.*

8 $\left\langle \begin{array}{l} \exists r, r_2, r_3.\, r_2 \stackrel{1/2}{\Longmapsto} P * [\mathsf{change}(x, r_2)]_1^r * r_3 \stackrel{1/2}{\Longmapsto} Q * [\mathsf{send}(y)]_1^r * \\ \mathsf{chainds}(rs) * \mathsf{chainres}(rs) * \mathsf{unused}(r, rs) \wedge rs = rs_1 \cdot e \cdot rs_2 \cdot e' \cdot rs_3 \wedge \\ \mathsf{wf}(rs) \wedge e = (y, r_3, \mathcal{I}', 0, \mathcal{W}') \wedge e' = (x, r', \mathcal{I}, f, \mathcal{W}) \wedge r_2 \in \mathcal{I} \end{array} \right\rangle$

9 // *Pull out node.*

10 $\left\langle \begin{array}{l} \exists r, r_2, r_3.\, r_2 \stackrel{1/2}{\Longmapsto} P * [\mathsf{change}(x, r_2)]_1^r * r_3 \stackrel{1/2}{\Longmapsto} Q * [\mathsf{send}(y)]_1^r * \\ \mathsf{chainds}(rs) * \mathsf{chainres}(rs_1) * \mathsf{resource}(\mathcal{I}', \mathcal{W}' \uplus \{r_3\}) * \mathsf{chainres}(rs_2 \cdot e' \cdot rs_3) * \mathsf{unused}(r, rs) \wedge \\ rs = rs_1 \cdot e \cdot rs_2 \cdot e' \cdot rs_3 \wedge \mathsf{wf}(rs) \wedge e = (y, r_3, \mathcal{I}', 0, \mathcal{W}') \wedge e' = (x, r', \mathcal{I}, f, \mathcal{W}) \wedge r_2 \in \mathcal{I} \end{array} \right\rangle$

11 // *Apply Lemma* 6.8 *to supply the resource*

12 $\left\langle \begin{array}{l} \exists r, r_2, r_3.\, [\mathsf{change}(x, r_2)]_1^r * r_4 \stackrel{1/2}{\Longmapsto} (P \twoheadrightarrow Q) * [\mathsf{send}(y)]_1^r * \\ \mathsf{chainds}(rs) * \mathsf{chainres}(rs_1) * \mathsf{resource}(\mathcal{I}', \mathcal{W}' \uplus \{r_2, r_4\}) * \mathsf{chainres}(rs_2 \cdot e' \cdot rs_3) * \mathsf{unused}(r, rs) \wedge \\ rs = rs_1 \cdot e \cdot rs_2 \cdot e' \cdot rs_3 \wedge \mathsf{wf}(rs) \wedge e = (y, r_3, \mathcal{I}', 0, \mathcal{W}') \wedge e' = (x, r', \mathcal{I}, f, \mathcal{W}) \wedge r_2 \in \mathcal{I} \end{array} \right\rangle$

13 // *Close the chain.*

14 $\left\langle \begin{array}{l} \exists r, r_2, r_4.\, r_4 \stackrel{1/2}{\Longmapsto} (P \twoheadrightarrow Q) * [\mathsf{send}(y)]_1^r * \mathsf{chainds}(rs') * \mathsf{chainres}(rs') * \mathsf{unused}(r, rs') \wedge \\ rs = rs_1 \cdot e \cdot rs_2 \cdot e' \cdot rs_3 \wedge rs' = rs_1 \cdot (y, r_4, \mathcal{I}', 0, \mathcal{W}' \uplus \{r_2\}) \cdot rs_2 \cdot e' \cdot rs_3 \wedge \\ \mathsf{wf}(rs) \wedge e = (y, \_, \mathcal{I}', 0, \mathcal{W}') \wedge e' = (x, r', \mathcal{I}, f, \mathcal{W}) \wedge r_2 \in \mathcal{I} \end{array} \right\rangle$

15 // *Close region using the* renun *transition.*

16 $\left\{ \exists r, r_4.\, r_4 \stackrel{1/2}{\Longmapsto} (P \twoheadrightarrow Q) * [\mathsf{send}(y)]_1^r * \mathsf{oreg}\left(r, \{\mathsf{Chain}(rs) \mid \mathsf{wf}(rs) \wedge rs(y) = (r_4, \_, 0, \_)\}\right) \right\}$

17 // *Definition of* send.

18 $\left\{ \mathsf{send}(y, P \twoheadrightarrow Q) \right\}$

Fig. 22.   Proof of the renunciation axiom for out-of-order implementation.

$1 \quad \left\{ \mathsf{recv}(x, P) * \lfloor \lceil P \rceil \twoheadrightarrow (P_1 * P_2) \rfloor \right\}$

$2 \quad$ *// Definition of* recv*.*

$3 \quad \left\{ \begin{array}{c} \exists r_1, r_2.\, \mathsf{oreg}\,(r_1, \{\mathsf{Chain}(rs) \mid rs(x) = (\_, \mathcal{I}, \_, \_) \wedge r_2 \in \mathcal{I} \wedge \mathsf{wf}(rs)\}) \\ * \, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(x, r_2)]_1^{r_1} * \lfloor \lceil P \rceil \twoheadrightarrow (P_1 * P_2) \rfloor \end{array} \right\}$

$4 \quad$ *// Enter the region.*

$5 \quad \left\langle \begin{array}{c} \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(x, r_2)]_1^{r_1} * \lfloor \lceil P \rceil \twoheadrightarrow (P_1 * P_2) \rfloor * \\ \triangleright \left(\begin{array}{c} \mathsf{chainds}(rs) * \mathsf{chainres}(rs_1) * \mathsf{resource}(\mathcal{I} \uplus \{r_2\}, \mathcal{W}) * \mathsf{chainres}(rs_2) \wedge \\ rs = rs_1 \cdot (x, r, \mathcal{I} \uplus \{r_2\}, f, \mathcal{W}) \cdot rs_2 \wedge \mathsf{unused}(r_1, rs) \end{array} \right) \end{array} \right\rangle$

$6 \quad \langle \mathtt{skip} \rangle$ *// AFrame to drop* $\triangleright$*, Lemma* 6.9 *to split the region.*

$7 \quad \left\langle \begin{array}{c} \exists r_1, r_2, r_3, r_4.\, r_3, r_4 \notin rs \wedge r_3 \overset{1/2}{\longmapsto} P_1 * r_4 \overset{1/2}{\longmapsto} P_2 * [\mathsf{change}(x, r_2)]_1^{r_1} * \\ \mathsf{chainds}(rs) * \mathsf{chainres}(rs_1) * \mathsf{resource}(\mathcal{I} \uplus \{r_3, r_4\}, \mathcal{W}) * \mathsf{chainres}(rs_2) * \\ \mathsf{unused}(r_1, rs) \wedge rs = rs_1 \cdot (x, r, \mathcal{I} \uplus \{r_2\}, f, \mathcal{W}) \cdot rs_2 \end{array} \right\rangle$

$8 \quad$ *// Close chain, pull* [change] *perms out of* unused*.*

$9 \quad \left\langle \begin{array}{c} \exists r_1, r_2, r_3, r_4.\, r_3 \overset{1/2}{\longmapsto} P_1 * r_4 \overset{1/2}{\longmapsto} P_2 * [\mathsf{change}(x, r_3)]_1^{r_1} * [\mathsf{change}(x, r_4)]_1^{r_1} * \\ \mathsf{chainds}(rs) * \mathsf{chainres}(rs_1) * \mathsf{resource}(\mathcal{I} \uplus \{r_3, r_4\}, \mathcal{W}) * \mathsf{chainres}(rs_2) * \mathsf{unused}(r_1, rs') \\ \wedge rs = rs_1 \cdot (x, r, \mathcal{I} \uplus \{r_2\}, f, \mathcal{W}) \cdot rs_2 \wedge rs' = rs_1 \cdot (x, r, \mathcal{I} \uplus \{r_3, r_4\}, f, \mathcal{W}) \cdot rs_2 \end{array} \right\rangle$

$10 \quad$ *// Close region using the* split *transition.*

$11 \quad \left\{ \begin{array}{c} \exists r_1, r_2, r_3, r_4.\, r_3 \overset{1/2}{\longmapsto} P_1 * r_4 \overset{1/2}{\longmapsto} P_2 * [\mathsf{change}(x, r_3)]_1^{r_1} * [\mathsf{change}(x, r_4)]_1^{r_1} * \\ \mathsf{oreg}\,(r_1, \{\mathsf{Chain}(rs) \mid rs(x) = (\_, \mathcal{I}, \_, \_) \wedge r_3, r_4 \in \mathcal{I} \wedge \mathsf{wf}(rs)\}) \end{array} \right\}$

$12 \quad$ *// Definition of* recv*.*

$13 \quad \left\{ \mathsf{recv}(x, P_1) * \mathsf{recv}(x, P_2) \right\}$

Fig. 23. Proof of the splitting axiom for the out-of-order implementation.

PROOF. Given in Appendix B. This proof is very similar in structure to the proof without chaining given in §5.3. □

7. OUT-OF-ORDER SUMMARIZATION

In this section we consider our third, most complex, channel implementation. This implementation satisfies our abstract specification, but internally, signals propagate up a tree structure towards a shared root rather than along a linear chain. This barrier implementation corresponds closely to the one in [Navabi et al. 2008] – the summarisation process is how it achieves its efficiency. Verifying this implementation demonstrates that our approach scales to custom synchronisation constructs developed for performance-sensitive concurrency applications.

7.1. Implementation approach

Figure 24 shows the implementation. The data-structure is an inverted tree – that is, nodes point upwards towards a single common root. An instance of the data-structure is illustrated in Figure 25.

All of the nodes in the tree are of type `chan_addr`, and each contains a fixed-size boolean array. Booleans in an array correspond to channel flags, either Low or High. Flags that are leaves, i.e. that do not have a child subtree, represent the channels in the chain. Scanning these leaf flags in order gives the sequence of flags in the abstract chain.

```
 1 typedef struct chan_hdr chan_hdr;      35 signal(chan_addr x){
 2                                        36   int i;
 3 typedef struct chan_addr {             37   bool ret = FALSE;
 4   chan_hdr *hdr;                        38   chan_addr a = x;
 5   int off;                             39   while (a.hdr != NULL && !ret){
 6 } chan_addr;                            40     a.hdr->flags[a.off] = 1;
 7                                        41     for(i=0; i<=a.hdr->loff; i++){
 8 typedef struct chan_hdr {              42       if (a.hdr->flags[i] != 1)
 9   chan_addr up;                         43         ret = TRUE;
10   int loff;                            44     }
11   bool flags[MAX];                     45     a = a.hdr->up;
12 } chan_hdr;                             46   }
13                                        47 }
14 (chan_addr, chan_addr)                 48
15 extend(chan_addr x){                   49
16   chan_addr nx,r;                      50
17   if(x.off == x.hdr->loff              51
18             && x.off < MAX){           52
19     x.hdr->flags[x.loff+1] = 0;        53
20     x.hdr->loff++;                     54
21     r.hdr = x.hdr;                     55
22     r.off = x.off + 1;                 56
23     nx = x;                            57 wait(chan_addr x){
24   } else {                             58   chan_addr a = x;
25     nh = malloc(chan_hdr);             59   while (a.hdr != NULL) {
26     nh->up = x;                        60     for(skip; a.off<=a.hdr->loff;
27     nh->loff = 1;                      61                         a.off++){
28     nx.hdr = nh;                       62       while(a.hdr->flags[a.off] != 1)
29     nx.off = 0;                        63         skip;
30     r.hdr = nh;                        64     }
31     r.off = 1;                         65     a = a.hdr->up;
32   }                                    66     a.off++;
33   return (r,nx);                       67   }
34 }                                      68 }
```

Fig. 24. Channel implementation with out-of-order signalling and summarization.

Non-leaf flags in the tree summarise the states of multiple flags in the chain. Nodes are equipped with an `up` field containing an address and an offset – the address is the parent node, while the offset refers to a particular 'summary' flag in its array. If all the flags in a node are set, then its parent summary flag can also be set. This means that if a summary flag is set, all the flags in the subtree must also be set. Thus `wait` only needs to read from a single flag to know that all the flags in the child subtree are set.

To `signal` a channel, the function reads the array location from the array header, and writes 1 into the flag at the appropriate offset (Fig. 24, line 40). `signal` then reads all the sibling flags in the same array. If any of them are unset, it exits. Otherwise it retrieves the address to the next level in the tree and loops if it is not at the root (lines 41–45). In this way, if a flag's siblings are all set, then `signal` will set the parent summary flag. If all the siblings of the summary flag are also set, it will set its parent, and in this way iterate up the tree.

The `wait` function exploits summaries to reduce the number of flags it must test. Rather than examining all the preceding flags in the chain, `wait` just examines the summary closer
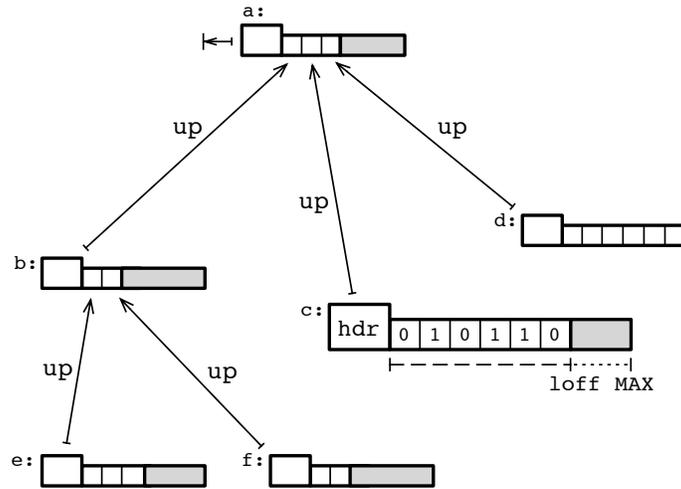
Fig. 25. Example of the summarising channel structure.

to the root. The function waits on each preceding flag in the current array (Fig. 24, lines 60–64 – note that increasing the index moves logically earlier in the chain). It then reads the **up** address stored in the array header and loops if it is not at the root. Because it only ever climbs the tree, the function avoids the cost of iterating over the entire chain. The following diagram shows the nodes accessed when traversing up from **x**:
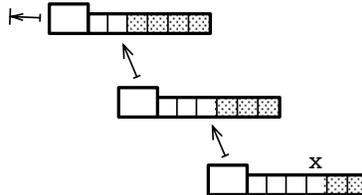
Arrays are allocated with a fixed maximum size **MAX**, and are gradually filled when extending the chain. The current number of flags active in the array is stored in the header field **loff**. Thus **extend** has two cases depending whether there is room in the current array for another leaf (checked on line 17).

— If there is space, then the new leaf is inserted immediately following the current one in the array (lines 20–22).
— If no space remains in the array, then a new array is allocated, and the current flag is used as a summary. The first and second locations in the new array represent the current and newly-created channels (lines 25–31).

As a result of extension, **wait** may be passed the address of a leaf flag, which is then silently converted into a summary. However, this process is sound: if this summary is set, then the appropriate child flag must also be set (this subtlety is what necessitates the predicate **finalleaf** in the proof below).

7.2. Proof Strategy

Abstract state. At level of specification, the behaviour of the algorithm is unchanged from §6 – as a result, we can reuse some of the reasoning from there. The algorithm's abstract

state has two parts: an *abstract chain* and a *heap map*. The former dictates the resources promised and renounced at each point in the chain, while the latter defines the underlying inverted-tree pointer structure.

An *abstract chain* is a sequence in CNode^+, defined exactly as in §6 (see p.38). The only difference is that each node location *loc* is now a pair $\langle h, o \rangle$ consisting of a physical address and an offset – this corresponds to the `chan_addr` type. We reuse abstract chains so that we can reuse reasoning about how resources move through the chain. However, for this implementation an abstract chain alone is insufficient, because it does not represent the underlying inverted tree data-structure.

A *heap map* represents the inverted-tree structure which controls signal propagation and summarisation. Heap maps are finite, partial functions from physical addresses – each address represents a `chan_hdr`-typed object forming the tree.

$$d \in \mathsf{HMap} : \mathsf{Addr} \xrightarrow{fin} \{uhdr : \mathsf{Addr}; \; uoff : \mathsf{Int}; \; loff : \mathsf{Int}; \; flags : \{0..\mathsf{MAX}\} \to \mathsf{Bool}\}$$

(To simplify the representation, we flatten the `chan_adr`-typed field 'up' into the two sub-fields *uhdr* and *uoff*.)

To ensure a heap map represents a correct inverted tree we require several well-formedness properties:

— Each address-offset pair $\langle x, i \rangle$ has at most one child in d which points upward to it. This ensures leaves are uniquely identified.
— Chains of upward pointers are non-cyclic and point towards a common root. This ensures that the data-structure is tree-shaped.
— Flags are only set in non-leaf nodes if all the corresponding leaf nodes are set. This ensures that signals are properly summarised and propagated up the tree.

Given a chain rs and heap map d, well-formedness also requires that the two portions of the abstract state correspond. Loosely, this means that the leaves of the inverted tree defined by d are exactly the addresses in the chain rs.

Abstract state well-formedness. In order to define formally that rs and d are well-formed and correspond correctly, we require several auxiliary notions:

$$
\begin{aligned}
\mathsf{child}_d(x, i, y) &\triangleq d(y).uhdr = x \wedge d(y).uoff = i \wedge 0 \le i \le d(x).loff \\
\mathsf{leaf}_d(\langle x, i \rangle) &\triangleq \langle x, i \rangle \in d \wedge \nexists y. \, d(y).uhdr = x \wedge d(y).uoff = i \\
\mathsf{descend}_d(\langle x, i \rangle, \langle y, j \rangle) &\triangleq \langle x, i \rangle = \langle y, j \rangle \vee \mathsf{descend}_d(\langle x, i \rangle, \langle d(y).uhdr, d(y).uoff \rangle) \\
\mathsf{isset}_d(\langle x, i \rangle) &\triangleq d(x).flags(i) = 1 \\
\mathsf{allset}_d(x) &\triangleq \forall i. \, 0 \le i \le d(x).loff \implies \mathsf{isset}_d(\langle x, i \rangle)
\end{aligned}
$$

(Note we say that a location-offset pair $\langle x, i \rangle$ is in d if $x \in \mathrm{dom}(d) \wedge i \le d.loff$.)

The child, $\mathsf{descend}$, and leaf predicates record corresponding structural facts about relationships in the tree. Well-formedness on d (defined below) requires that paths through *uhdr* are finite, which suffices to ensure that $\mathsf{descend}$ is well-defined. isset and allset respectively assert that a single address and a whole array have their flags set.

$$
\begin{aligned}
\langle x, i \rangle <_d \langle y, j \rangle &\triangleq \exists z, ix, iy. \, \mathsf{descend}_d(\langle z, ix \rangle, \langle x, i \rangle) \wedge \mathsf{descend}_d(\langle z, iy \rangle, \langle y, j \rangle) \wedge ix > iy \\
a <_d^{\mathbb{1}} b &\triangleq a <_d b \vee \mathsf{descend}_d(a, b) \vee \mathsf{descend}_d(b, a) \\
\mathsf{finalleaf}_d(x, y) &\triangleq \mathsf{descend}_d(x, y) \wedge \mathsf{leaf}_d(y) \wedge (\forall z. \, \mathsf{descend}_d(x, z) \wedge \mathsf{leaf}_d(z) \Rightarrow z <_d^{\mathbb{1}} y)
\end{aligned}
$$

The order predicate $<_d$ says that two addresses are ordered in the tree, meaning that they share a common ancestor array in which they are also ordered. This defines a transitive and irreflexive order. $<_d^{\mathbb{1}}$ says that either two addresses are related by $<_d$, or that one is the descendant of the other (i.e. they are on the same path and one summarises the other).

$\mathsf{finalleaf}_d(x, y)$ indicates that y is the maximal leaf according to $<^1_d$ that is summarised by x. This is useful because applications of `extend` may mean clients `wait` on x when the actual leaf has been superceded by y. For each x there exists at most one y satisfying $\mathsf{finalleaf}$, so we generally use it as a partial function, i.e. $\mathsf{finalleaf}_d(x)$ stands for the unique y such that $\mathsf{finalleaf}_d(x, y)$.

We can now define $\mathsf{wf}(rs, d)$, which requires that rs and d are independently well-formed, and that they are correctly tied together. As abstract chains are unchanged from §6, we can reuse the prior definition of $\mathsf{wf}(rs)$ (§6.2). For heap-maps, well-formedness is defined as follows:

$$
\begin{aligned}
\mathsf{wf}(d) \quad \triangleq \quad & \exists r \colon \mathsf{Addr}.\, \exists \tau \colon \mathrm{dom}(d) \to \mathbb{N}. \\
& \forall x, i, y, z.\, (\langle x, i \rangle \in d \implies \exists j.\, \mathsf{descend}(\langle r, j \rangle, \langle x, i \rangle)) \,\wedge \\
& \quad ((\mathsf{child}_d(x, i, y) \wedge \mathsf{child}_d(x, i, z)) \implies y = z) \,\wedge \\
& \quad ((\mathsf{child}_d(x, i, y) \wedge \mathsf{isset}_d(x, i)) \implies \mathsf{allset}_d(y)) \,\wedge \\
& \quad (d(x).uhdr = y \wedge y \neq \mathsf{NULL} \implies \\
& \qquad d(y) \text{ defined} \wedge d(x).uoff \leq d(y).loff \wedge \tau(y) < \tau(x))
\end{aligned}
$$

Here the address r is the location of the tree root, while the function τ records the distance from the current node to the root – this enforces the absence of cycles. The first clause of the definition ensures all elements in the tree share a common root. The second ensures that children are uniquely identified by address and offset. The third ensures that setting a flag summarises all descendants. The final clause guarantees the existence of non-NULL parents to a node, and enforces the distance function τ.

$$
\begin{aligned}
\mathsf{wf}(rs, d) \quad \triangleq \quad & \mathsf{wf}(rs) \wedge \mathsf{wf}(d) \,\wedge \\
& \forall r \in rs.\ r.loc = \langle l, o \rangle \implies r.flg = d(l).flags[o] \,\wedge \\
& \forall r \in rs.\, \mathsf{leaf}_d(r.loc) \,\wedge \\
& \forall r_1, r_2.\, (rs = \_ \cdot r_1 \cdot \_ \cdot r_2 \cdot \_) \implies r_2.loc <_d r_1.loc
\end{aligned}
$$

As well as requiring that rs and d are well-formed on their own, this requires (1) that flags in the chain are correctly set in the heap map; (2) locations in the chain are leaves in the heap map; and (3) order in the chain is reflected in the heap-map order $<_d$.

Transition map. We define a new transition map T_s to capture changes to the heap map (Figure 26). Many of the transitions are inherited from T_c, the transition relation for the chained implementation (§6.2).

`set` and `ext` alter the underlying inverted-tree data-structure, and so are defined to allow this. `set` just updates the tree-map flag appropriately using our lens notation. `ext` has two cases, reflecting the conditional in the implementation. Either there is enough space in the array to fit another flag, or another array must be added to the heap map. In the latter case, note that the address of the existing channel a is shifted into the new array.

Furthermore, `signal` can mark summary flags, i.e. flags that are not leaves in the tree; this is allowed by the transition `mark`. Finally, note $\mathsf{finalleaf}$ in the definition of `get`: this is needed because the real flag may shift its position due to `extend`, with `wait` left reading from a summary. Here x is the original flag which has been superceded, and $\mathsf{finalleaf}_d(x)$ is the current position of the flag. It is safe to pass an address to `wait` which may have been converted into a summary because extension ensures the original flag must be one of those summarised by x.

Interpretation function. For a given map d, `chainds` maps down to the corresponding data-structure definition. Heap maps are intentionally close to the underlying heap: each element

$$
\begin{aligned}
\mathsf{renun}_s(x, (rs, d), (rs', d')) &\triangleq d' = d \wedge \mathsf{renun}_c(x, rs, rs') \\
\mathsf{set}_s(x, (rs, d), (rs', d')) &\triangleq x = \langle h, o \rangle \wedge d' = d \blacktriangleleft_h \blacktriangleleft_{flags[o]} (1) \wedge \mathsf{set}_c(\langle h, o \rangle, rs, rs') \\
\mathsf{ext}_s(x, (rs, d), (rs', d')) &\triangleq rs = (rs_1 \cdot a \cdot rs_2) \wedge rs' = (rs_1 \cdot b \cdot c \cdot rs_2) \wedge a.flg = 0 \wedge \\
&\qquad \exists y.\, b = a[loc \mapsto y] \wedge c.flg = 0 \wedge c.\mathcal{W} = \emptyset \wedge x = \langle h, o \rangle = a.loc \wedge \\
&\qquad \left(\begin{pmatrix} b.off = \langle h', 0 \rangle \wedge c.off = \langle h', 1 \rangle \wedge \\ d' = d \uplus h' \mapsto \{uhdr = h; uoff = o; loff = 1; flags = \lambda_\_.\, 0\} \end{pmatrix} \right. \\
&\qquad \left. \vee \begin{pmatrix} o = d(h).loff \wedge b.loc = a.loc \wedge c.loc = \langle h, o+1 \rangle \wedge \\ d' = (d \blacktriangleleft_h \blacktriangleleft_{loff} (o+1)) \blacktriangleleft_h \blacktriangleleft_{flags[o+1]} (0) \end{pmatrix} \right) \\
\mathsf{split}_s(x, r, (rs, d), (rs', d')) &\triangleq d = d' \wedge \mathsf{split}_c(\mathsf{finalleaf}_d(x), r, rs, rs') \\
\mathsf{sat}_s((rs, d), (rs', d')) &\triangleq d = d' \wedge \mathsf{sat}_c(rs, rs') \\
\mathsf{get}_s(x, r, (rs, d), (rs', d')) &\triangleq d = d' \wedge \mathsf{get}_c(\mathsf{finalleaf}_d(x), r, rs, rs') \\
T_\mathsf{s}(\mathsf{send}(x)) &\triangleq \{(a, b) \mid \mathsf{wf}(b) \wedge (\mathsf{renun}_s(x, a, b) \vee \mathsf{set}_s(x, a, b) \vee \mathsf{ext}_s(x, a, b))\} \\
T_\mathsf{s}(\mathsf{change}(x, r)) &\triangleq \{(a, b) \mid \mathsf{wf}(b) \wedge (\mathsf{split}_s(x, r, a, b) \vee \mathsf{sat}_s(a, b) \vee \mathsf{get}_s(x, r, a, b))\} \\
T_\mathsf{s}(\mathsf{mark}) &\triangleq \left\{ ((rs, d), (rs, d')) \;\middle|\; \begin{aligned} &\mathsf{wf}(rs, d') \wedge d' = d \blacktriangleleft_h \blacktriangleleft_{flags[o]} (1) \\ &\wedge \neg \mathsf{leaf}_d(\langle h, o \rangle)) \end{aligned} \right\}
\end{aligned}
$$

Fig. 26. Definition of the transition relation T_s for the summarising implementation.

in the domain maps to a distinct `chan_addr` object in memory.

$$
\begin{aligned}
\mathsf{chainds}(d) \quad \triangleq \quad &\circledast_{x \in \mathrm{dom}(d)}. \\
&x.\mathrm{up} \Mapsto \{\mathrm{hdr} = d(x).uhdr; \mathrm{off} = d(x).uoff\} * x.\mathrm{loff} \mapsto (d(x).loff) \\
&\qquad * \circledast_{i \in \{0..\mathsf{MAX}\}}. x.\mathrm{flags}[i] \mapsto (d(x).flags(i))
\end{aligned}
$$

Once a `chan_addr` object has been allocated, its up-address cannot be modified. To represent this in the definition, we write $x \Mapsto v$ to indicate that x is immutable – shorthand for $\exists f.\, x \xmapsto{f} v$. This is useful as immutable locations can be freely shared between threads: $x \Mapsto v \implies x \Mapsto v * x \Mapsto v$.

The interpretation function I_s converts an abstract state $\mathsf{Chain}(rs, d)$ into a data-structure:

$$
I_s(r)(\mathsf{Chain}(rs, d)) \quad \triangleq \quad \mathsf{chainds}(d) * \mathsf{chainres}(rs) * \mathsf{unused}(r, rs, d)
$$

As $\mathsf{chainds}$ defines the concrete heap structure, it only requires a heap-map d. Conversely, $\mathsf{chainres}$ defines the pattern of splits, promises, and renunciations, and so only requires an abstract chain rs. Indeed, $\mathsf{chainres}$ is defined identically to §6.

The predicate unused, representing the 'library' of unused permissions, requires both d and rs. This is because the position of a flag may move as a result of extension. As a result, `wait` may be passed a summary flag x – before extension, the passed node would have been a leaf. The flag x will be represented in d, but not represented in the chain rs, and thus unused requires both in order to keep track of permissions.

To define unused, the set $\mathsf{C}(rs, d)$ represents possible targets of `wait`; a permission is missing from the set of unused `change` permissions, $\mathsf{uC}(rs, d)$, only if it targets one of these

nodes.

$$\mathsf{C}(rs,d) \triangleq \{\langle y,i\rangle \mid r \in rs \wedge \mathsf{finalleaf}_d(\langle y,i\rangle, r.loc) \wedge \neg\mathsf{finalleaf}_d(\langle d(y).uhdr, d(y).uoff\rangle, r.loc)\}$$

$$\mathsf{uC}(rs,d) \triangleq \{(x,r) \mid r \in rs(\mathsf{finalleaf}_d(x)).\mathcal{I} \wedge x \in \mathsf{C}(rs,d) \wedge \neg\exists y.\, r \in rs(y).\mathcal{W}\}$$

$$\mathsf{unused}(r,rs,d) \triangleq (\bigcircledast\, x \notin \mathsf{uS}(rs).\, [\mathsf{send}(x)]_1^r) * (\bigcircledast (x,r') \notin \mathsf{uC}(rs,d).\, [\mathsf{change}(x,r')]_1^r)$$

Predicate definitions. We can now define the send and recv predicates. These largely follow the definitions in §6: the differences in underlying data-structures are abstracted by the interpretation function and transition relation.

$$\mathsf{oreg}(r,S) \triangleq \mathsf{region}(S, T_\mathsf{s}, I_\mathsf{s}(r), r)$$

$$\mathsf{send}(x,P) \triangleq \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{send}(x)]_1^{r_1} * [\mathsf{mark}]_-^{r_1} *$$
$$\mathsf{oreg}(r_1, \{\mathsf{Chain}(rs,d) \mid \mathsf{wf}(rs,d) \wedge rs(x) = (r_2, \_, 0, \_)\})$$

$$\mathsf{recv}(x,P) \triangleq \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(x,r_2)]_1^{r_1} *$$
$$\mathsf{oreg}(r_1, \{\mathsf{Chain}(rs,d) \mid \mathsf{wf}(rs,d) \wedge r_2 \in rs(\mathsf{finalleaf}_d(x)).\mathcal{I}\})$$

$$x \prec y \triangleq \exists r.\, \mathsf{oreg}(r, \{\mathsf{Chain}(rs,d) \mid \mathsf{wf}(rs,d) \wedge (\mathsf{finalleaf}_d(y), \mathsf{finalleaf}_d(x)) \in \mathsf{ord}(rs)\})$$

We use x to stand for $\langle x.\mathtt{hdr}, x.\mathtt{off}\rangle$ if x is a $\mathtt{chan\_addr}$ struct and we use $[\mathsf{mark}]_-^r$ as notation for $\exists \pi.\, [\mathsf{mark}]_\pi^r$ to represent non-exclusive ownership of the mark action.

The stability of most of the predicates is obvious; however, the fact that finalleaf can change means we prove stability explicitly for recv.

LEMMA 7.1. $\mathsf{recv}(x,P)$ *is stable.*

PROOF. Assume the initial abstract state of the chain is $\mathsf{Chain}(rs,d)$ and that ext takes the step $(x', (rs,d), (rs', d'))$. The case where $x \neq x'$ is trivial, so assume $x = x'$. We now need to show $\mathsf{wf}(rs', d') \wedge r_2 \in rs'(\mathsf{finalleaf}_{d'}(x)).\mathcal{I}$. Assume a, b and c are chain nodes as used in the definition of ext.

The requirement $\mathsf{wf}(rs', d')$ holds as a constraint on the transition relation. It remains to show the second clause. By the definition of ext, $\mathsf{finalleaf}_d(x) = a.loc$ and $a.\mathcal{I} = b.\mathcal{I}$. There are now two cases: either ext generates a new array, or it adds an element to the existing array. In the latter case, d' only changes by adding an element at a higher index in the array. Thus $\mathsf{finalleaf}_{d'}(x) = b.loc$. In the former case, ext adds a new array which must also descend from x. If $\mathsf{finalleaf}_{d'}(x) \neq b.loc$, there must be leaf z such that $b.loc <_d z$, but the only new leaf c is at the next index in the new array, meaning $c.loc <_d b.loc$. Thus the result follows by contradiction. □

7.3. Verifying `wait`, `signal`, `extend`

Proving `signal`. A sketch-proof for `signal` is given in Fig. 27. The algorithm begins by setting the flag at the appropriate address (line 9). Abstractly the reasoning here is the same as when setting a flag in the non-summarising implementation (§6.3) so we omit it. The algorithm then climbs up the tree. If all the flags have been set in a given array, the summary flag is also set (line 9). Well-formedness allows summary nodes to be set if all their children are set. If a flag is discovered which is not set, or the loop climbs to the top of the tree, the algorithm exits.

The assignment on line 9 applies the transition relation step set or mark, depending on whether the node is a leaf or a summary. The following lemma ensures that the library of unused permissions is preserved after each such transition relation step.

```
1  {send(x, P) * ⌊P⌋}
2  signal(chan_addr x){
3    int i; ret = FALSE;
4    chan_addr a = x;
```

$$\left\{ \begin{array}{l} \mathtt{a} = \mathtt{x} \land \lfloor P \rfloor * \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{send}(\mathtt{x})]_1^{r_1} * \\ \mathsf{oreg}(r_1, \{\mathsf{Chain}(rs, d) \mid \mathsf{wf}(rs, d) \land \mathsf{leaf}_d(\mathtt{x}) \land d(\mathtt{x.hdr}).\mathit{flags}[\mathtt{x.off}] = 0\}) \end{array} \right\}$$

```
6    while (a.hdr != NULL && !ret){
7     // Loop invariant.
```

$$\left\{ \begin{array}{l} (\mathtt{ret} \land \mathsf{emp}) \lor \\ \left(\begin{array}{l} \mathtt{a} = \mathtt{x} \land \lfloor P \rfloor * \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{send}(\mathtt{x})]_1^{r_1} * [\mathsf{mark}]_-^{r_1} * \\ \mathsf{oreg}(r_1, \{\mathsf{Chain}(rs, d) \mid \mathsf{wf}(rs, d) \land \mathsf{leaf}_d(\mathtt{x}) \land d(\mathtt{x.hdr}).\mathit{flags}[\mathtt{x.off}] = 0\}) \end{array} \right) \lor \\ \left(\begin{array}{l} \exists r_1.\, [\mathsf{mark}]_-^{r_1} * \\ \mathsf{oreg}(r_1, \{\mathsf{Chain}(rs, d) \mid \mathsf{wf}(rs, d) \land \neg\mathsf{leaf}_d(\mathtt{a}) \land (\forall l.\, (\mathsf{descend}_d(\mathtt{a}, l) \land \mathtt{a} \neq l) \Rightarrow \mathsf{isset}_d(l))\}) \end{array} \right) \end{array} \right\}$$

```
9     a.hdr->flags[a.off] = 1; // Transition relation step set / mark.
10    for(i=0; i<=a.hdr->loff; i++){
```

$$\left\{ \left(\begin{array}{l} \exists r_1.\, [\mathsf{mark}]_-^{r_1} * \\ \mathsf{oreg}\left(r_1, \left\{ \mathsf{Chain}(rs, d) \,\middle|\, \begin{array}{l} \mathsf{wf}(rs, d) \land (\forall l.\, \mathsf{descend}_d(\mathtt{a}, l) \Rightarrow \mathsf{isset}_d(l)) \\ \land\, \forall 0 \leq n < \mathtt{i}.\, \mathsf{isset}_d(\langle \mathtt{a.hdr}, n \rangle) \end{array} \right\} \right) \end{array} \right) \right\}$$

```
12     if (a.hdr->flags[i] != 1)
13        ret = TRUE;
14    }
```

$$\left\{ \begin{array}{l} (\mathtt{ret} \land \mathsf{oreg}(r_1, \{\mathsf{Chain}(rs, d) \mid \mathsf{wf}(rs, d) \land \mathtt{a} \in \mathrm{dom}(d)\})) \lor \\ \left(\begin{array}{l} \exists r_1.\, [\mathsf{mark}]_-^{r_1} * \\ \mathsf{oreg}(r_1, \{\mathsf{Chain}(rs, d) \mid \mathsf{wf}(rs, d) \land (\forall l.\, \mathsf{descend}_d(\mathtt{a}, l) \Rightarrow \mathsf{isset}_d(l)) \land \mathsf{allset}_d(\mathtt{a.hdr})\}) \end{array} \right) \end{array} \right\}$$

```
16    a = a.hdr->up; // If exit condition holds, delete assertions using weakening.
17   }
18  }
19  {emp}
```

Fig. 27. Sketch-proof of `signal` with summarization.

LEMMA 7.2.

$$\mathsf{set}_s(x, (rs, d), (rs', d')) \wedge \mathsf{unused}(r, rs, d) * [\mathsf{send}(x)]_1^r \implies \mathsf{unused}(r, rs', d')$$

$$((rs, d), (rs', d')) \in T_s(\mathsf{mark}) \wedge \mathsf{unused}(r, rs, d) \implies \mathsf{unused}(r, rs', d')$$

PROOF. Trivial from the definition of T_s and unused. □

Proving `wait`. A sketch-proof for `wait` is given in Fig. 28 / 29. This proof just deals with the part of the code establishing that all the flags in the chain have been set. In the proof, we use $\mathsf{last}_d(\mathtt{a})$ to stand for the last address in the array associated with $\mathtt{a}$, i.e. $\langle \mathtt{a.hdr}, d(\mathtt{a.hdr}).loff \rangle$.

The loop starting at line 5 checks the flags in the current array. In line 8 the algorithm waits for the current node's flag. This may not be a leaf – it may be a summary node somewhere inside the tree. Once this passes, by the second clause of well-formedness (page 53) we can conclude that all the flags in the subsequence rs_3 have been set. Then the algorithm increments the offset – as $\mathtt{a}$ is not at the last offset for the array, there must exist an adjacent channel address at this position.

To prove this algorithm correct, we need several sub-lemmas. The first states that once the algorithm reaches the root of the tree, there are no locations in d that are earlier according to $<_d$. This ensures that searching the tree covers all the preceding channels in the chain.

LEMMA 7.3. $\mathsf{wf}(d) \wedge d(a).loff = o \wedge d(a).uhdr = \mathsf{NULL} \implies \neg \exists x.\, x <_d \langle a, o \rangle$

PROOF. Assume such an x exists. Then by the definition of $<_d$, there must exist an address y such that x and $\langle a, o \rangle$ are both descended from y. As the $uhdr$ field is NULL, the only possibility is that both addresses are in the object at a. By the definition of $<_d$, x must be further right in the flag array, but o is the right-most address. This contradicts the assumption and completes the proof. □

The second lemma states that examining the elements reachable through the heap map suffices to show that the corresponding elements in the abstract chain have been set. This lemma justifies our splitting of the invariant into a separate heap map and abstract chain structure.

LEMMA 7.4. $\mathsf{wf}(rs, d) \wedge z \in \mathsf{dom}(rs) \wedge \mathsf{descend}_d(x, z) \wedge \mathsf{leaf}_d(z) \wedge$
$\qquad (\forall l.\, l <_d^{\mathbb{1}} x \wedge \mathsf{leaf}_d(l) \Rightarrow \mathsf{isset}_d(l))$
$\qquad \implies \quad \exists rs_1, rs_2.\, rs = rs_1 \cdot rs(z) \cdot rs_2 \wedge \mathsf{ctrue}(rs(z) \cdot rs_2)$

PROOF. As $z \in \mathsf{dom}(rs)$, we can easily divide up rs into $rs_1 \cdot rs(z) \cdot rs_2$. Now pick an arbitrary element y in $dom(rs(z) \cdot rs_2)$ and suppose that $rs(y).flg$ is not set. By well-formedness, it must be true that $y <_d^{\mathbb{1}} z$. Now we show that $y <_d^{\mathbb{1}} x$. The contrary, $x <_d y$, would imply that $z <_d y$, contradicting our assumption. Therefore by the premise the associated flag must be set. However, well-formedness requires that flags are mirrored in rs and d, contradicting our assumption and completing the proof. □

The final lemma shows that shifting left from the current maximal node reaches a node earlier in the order. Note that the existence of the new node is shown on left of the implication because this lemma applied in a negative position in the proof.

LEMMA 7.5. $0 \le d(a).off < d(a).loff \wedge \mathsf{wf}(d) \wedge a \blacktriangleleft_{off} (\bullet + 1) <_d k <_d b \implies a <_d^{\mathbb{1}} k <_d b$

PROOF. The result follows from the structure of the heap map and the definition of $<_d$. □

Proving `extend`. A sketch-proof of `extend` is given in Fig. 30. There are two cases for extending the chain: either the node is the last element in the current array and there is

```
1  {recv(x, P)}
2  wait(chan_addr x){
3    chan_addr a = x;
```

$$\left\{ \begin{array}{l} \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(\mathbf{x}, r_2)]_1^{r_1} * \mathsf{a} = \mathbf{x} * \\ \mathsf{oreg}(r_1, \{\mathsf{Chain}(rs, d) \mid \mathsf{wf}(rs, d) \wedge r_2 \in rs(\mathsf{finalleaf}_d(\mathbf{x})).\mathcal{I}\}) \end{array} \right\}$$

```
5    while (a.hdr != NULL) {
```

$$\left\{ \begin{array}{l} \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(\mathbf{x}, r_2)]_1^{r_1} * \\ \mathsf{oreg}\left(r_1, \left\{ \mathsf{Chain}(rs, d) \left| \begin{array}{l} \mathsf{wf}(rs, d) \wedge r_2 \in rs(\mathsf{finalleaf}_d(\mathbf{x})).\mathcal{I} \wedge \\ \forall l.\, (\mathsf{a} <_d l <_d^{1\!\!1} \mathbf{x}) \wedge \mathsf{leaf}_d(l) \Rightarrow \mathsf{isset}_d(l) \end{array} \right. \right\} \right) \end{array} \right\}$$

```
7      for(skip; a.off<=a.hdr->loff; a.off++){
8        while(a.hdr->flags[a.off] != 1) skip;
```

$$\left\{ \begin{array}{l} \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(\mathbf{x}, r_2)]_1^{r_1} * \\ \mathsf{oreg}\left(r_1, \left\{ \mathsf{Chain}(rs, d) \left| \begin{array}{l} \mathsf{wf}(rs, d) \wedge r_2 \in rs(\mathsf{finalleaf}_d(\mathbf{x})).\mathcal{I} \wedge \mathsf{isset}_d(\mathsf{a}) \wedge \\ \forall l.\, (\mathsf{a} <_d l <_d^{1\!\!1} \mathbf{x}) \wedge \mathsf{leaf}_d(l) \Rightarrow \mathsf{isset}_d(l) \end{array} \right. \right\} \right) \end{array} \right\}$$

```
10       // Appeal to well-formedness to add 'a' to checked interval.
```

$$\left\{ \begin{array}{l} \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(\mathbf{x}, r_2)]_1^{r_1} * \\ \mathsf{oreg}\left(r_1, \left\{ \mathsf{Chain}(rs, d) \left| \begin{array}{l} \mathsf{wf}(rs, d) \wedge r_2 \in rs(\mathsf{finalleaf}_d(\mathbf{x})).\mathcal{I} \wedge \\ \forall l.\, (\mathsf{a} <_d^{1\!\!1} l <_d^{1\!\!1} \mathbf{x}) \wedge \mathsf{leaf}_d(l) \Rightarrow \mathsf{isset}_d(l) \end{array} \right. \right\} \right) \end{array} \right\}$$

```
12       // Apply Lemma 7.5 to show preceding node exists.
13     }
14     // Stable because chain cannot be extended once all the flags have been set.
```

$$\left\{ \begin{array}{l} \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(\mathbf{x}, r_2)]_1^{r_1} * \\ \mathsf{oreg}\left(r_1, \left\{ \mathsf{Chain}(rs, d) \left| \begin{array}{l} \mathsf{wf}(rs, d) \wedge r_2 \in rs(\mathsf{finalleaf}_d(\mathbf{x})).\mathcal{I} \wedge \\ \forall l.\, (\mathsf{last}_d(\mathsf{a}) <_d^{1\!\!1} l <_d^{1\!\!1} \mathbf{x}) \wedge \mathsf{leaf}_d(l) \Rightarrow \mathsf{isset}_d(l) \end{array} \right. \right\} \right) \end{array} \right\}$$

```
16     // Case-split on whether d(a.hdr).uhdr = NULL, if so, apply Lemma 7.3.
17     // As ∀x. ¬x <_d last_d(a), well-formedness gives us ∀x. last_d(a) <_d^1 x.
```

$$\left\{ \begin{array}{l} \exists r_1, r_2.\, r_2 \overset{1/2}{\longmapsto} P * [\mathsf{change}(\mathbf{x}, r_2)]_1^{r_1} * \\ \mathsf{oreg}\left(r_1, \left\{ \mathsf{Chain}(rs, d) \left| \begin{array}{l} \mathsf{wf}(rs, d) \wedge r_2 \in rs(\mathsf{finalleaf}_d(\mathbf{x})).\mathcal{I} \wedge \\ \left(\begin{array}{l} d(\mathsf{a.hdr}).uhdr \neq \mathsf{NULL} \implies \\ \quad \forall l.\, (\mathsf{last}_d(\mathsf{a}) <_d^{1\!\!1} l <_d^{1\!\!1} \mathbf{x}) \wedge \mathsf{leaf}_d(l) \Rightarrow \mathsf{isset}_d(l) \\ \vee\ d(\mathsf{a.hdr}).uhdr = \mathsf{NULL} \implies \\ \quad \forall l.\, l <_d^{1\!\!1} \mathbf{x} \wedge \mathsf{leaf}_d(l) \Rightarrow \mathsf{isset}_d(l) \end{array} \right) \end{array} \right. \right\} \right) \end{array} \right\}$$

Fig. 28.   Sketch-proof of `wait` with summarization (completed in Fig. 29).

```
19        a = a.hdr->up;
```

$$\left\{\begin{array}{l} \exists r_1, r_2.\, r_2 \stackrel{1/2}{\longmapsto} P * [\mathsf{change}(\mathrm{x}, r_2)]_1^{r_1} * \\ \left(\mathsf{oreg}\left(r_1, \left\{\mathsf{Chain}(rs, d) \middle| \begin{array}{l} \mathsf{wf}(rs, d) \wedge r_2 \in rs(\mathsf{finalleaf}_d(\mathrm{x})).\mathcal{I} \wedge \\ \begin{pmatrix} \mathsf{a.hdr} \neq \mathsf{NULL} \implies \\ \forall l.\, \mathrm{a} <_d^1 l <_d^1 \mathrm{x} \wedge \mathsf{leaf}_d(l) \Rightarrow \mathsf{isset}_d(l) \\ \vee\ \mathsf{a.hdr} = \mathsf{NULL} \implies \\ \forall l.\, l <_d^1 \mathrm{x} \wedge \mathsf{leaf}_d(l) \Rightarrow \mathsf{isset}_d(l) \end{pmatrix} \end{array}\right\}\right)\right) \end{array}\right\}$$

20

```
21        a.off++; // Apply Lemma 7.5 to non-NULL case.
```

$$\left\{\begin{array}{l} \exists r_1, r_2.\, r_2 \stackrel{1/2}{\longmapsto} P * [\mathsf{change}(\mathrm{x}, r_2)]_1^{r_1} * \\ \left(\mathsf{oreg}\left(r_1, \left\{\mathsf{Chain}(rs, d) \middle| \begin{array}{l} \mathsf{wf}(rs, d) \wedge r_2 \in rs(\mathsf{finalleaf}_d(\mathrm{x})).\mathcal{I} \wedge \\ \begin{pmatrix} \mathsf{a.hdr} \neq \mathsf{NULL} \implies \\ \forall l.\, \mathrm{a} <_d l <_d^1 \mathrm{x} \wedge \mathsf{leaf}_d(l) \Rightarrow \mathsf{isset}_d(l) \\ \vee\ \mathsf{a.hdr} = \mathsf{NULL} \implies \\ \forall l.\, l <_d^1 \mathrm{x} \wedge \mathsf{leaf}_d(l) \Rightarrow \mathsf{isset}_d(l) \end{pmatrix} \end{array}\right\}\right)\right) \end{array}\right\}$$

22

```
23    }
```

$$\left\{\begin{array}{l} \exists r_1, r_2.\, r_2 \stackrel{1/2}{\longmapsto} P * [\mathsf{change}(\mathrm{x}, r_2)]_1^{r_1} * \\ \mathsf{oreg}\left(r_1, \left\{\mathsf{Chain}(rs, d) \middle| \begin{array}{l} \mathsf{wf}(rs, d) \wedge r_2 \in rs(\mathsf{finalleaf}_d(\mathrm{x})).\mathcal{I} \wedge \\ \forall l.\, l <_d^1 \mathrm{x} \wedge \mathsf{leaf}_d(l) \Rightarrow \mathsf{isset}_d(l) \end{array}\right\}\right) \end{array}\right\}$$

24

```
25    // Apply Lemma 7.4 to show all nodes in chain are set.
```

$$\left\{\begin{array}{l} \exists r_1, r_2.\, r_2 \stackrel{1/2}{\longmapsto} P * [\mathsf{change}(\mathrm{x}, r_2)]_1^{r_1} * \\ \left(\mathsf{oreg}\left(r_1, \left\{\mathsf{Chain}(rs, d) \middle| \begin{array}{l} \mathsf{wf}(rs, d) \wedge r_2 \in rs(\mathsf{finalleaf}_d(\mathrm{x})).\mathcal{I} \wedge \\ \forall z.\, \mathsf{leaf}_d(z) \wedge z \in \mathsf{dom}(rs) \wedge \mathsf{descend}_d(\mathrm{x}, z) \implies \\ \exists rs_1, rs_2.\, rs = rs_1 \cdot rs(z) \cdot rs_2 \wedge \mathsf{ctrue}(rs(z) \cdot rs_2) \end{array}\right\}\right)\right) \end{array}\right\}$$

26

```
27    // Identical reasoning to chained implementation (§6.3)
28    ...
29  }
```

Fig. 29.   Sketch-proof of `wait` with summarization (continued from Fig. 28).

space to add an extra node; or there is no space and the algorithm allocates a fresh array. This choice is made by the conditional in line 6.

   The proof needs the following lemmas to show that the unused predicate representing unused permissions is preserved by extending the chain.

LEMMA 7.6.

$$\begin{pmatrix} \mathsf{ext}_s(\langle x, i\rangle, (rs, d), (rs', d')) \wedge \mathsf{dom}(d') = \mathsf{dom}(d) \wedge r' \notin rs \wedge \\ rs'(\langle x, i+1\rangle).\mathcal{I} = \{r'\} \wedge \mathsf{wf}(rs, d) \wedge \mathsf{wf}(rs', d') \wedge \mathsf{unused}(r, rs, d) \end{pmatrix} \implies$$
$$\mathsf{unused}(r, rs', d') * [\mathsf{send}(\langle x, i+1\rangle)]_1^r * [\mathsf{change}(\langle x, i+1\rangle, r')]_1^r$$

PROOF. Begin by observing that, by the definition of ext_s, $\langle x, i+1\rangle \notin rs$, and thus that:

$$\mathsf{unused}(r, rs, d) \implies [\mathsf{send}(\langle x, i+1\rangle)]_1^r * [\mathsf{change}(\langle x, i+1\rangle, r')]_1^r * \mathsf{true}$$

As $\langle x, i+1\rangle \in rs'$, it holds immediately that $\langle x, i+1\rangle \in \mathsf{uS}(rs')$. As $\langle x, i+1\rangle$ is a leaf, $\mathsf{finalleaf}_{d'}(\langle x, i+1\rangle) = \langle x, i+1\rangle$. As it is not the first leaf in the array x, it cannot have a

1 $\left\{\mathsf{send}(\mathbf{x}, P) * \circledast_{e \in E}\, e \prec \mathbf{x} * \circledast_{l \in L}\, \mathbf{x} \prec l\right\}$

2 `extend(chan_addr x){`

3   *// Frame off order predicates and unfold* `send`*.*

4 $\left\{\begin{aligned} &\exists r_1, r_2.\, r_2 \overset{1/2}{\Longmapsto} P * [\mathsf{send}(\mathbf{x})]_1^{r_1} * [\mathsf{mark}]_-^{r_1} * \\ &\mathsf{oreg}(r_1, \{\mathsf{Chain}(rs, d) \mid \mathsf{wf}(rs, d) \wedge \mathsf{leaf}_d(\mathbf{x}) \wedge rs(\mathbf{x}).\mathit{flg} = 0\}) \end{aligned}\right\}$

5   `chan_addr nx,r;`

6   `if(x.off == x.hdr->loff && x.off < MAX){`

7 $\left\{\begin{aligned} &\exists r_1, r_2.\, r_2 \overset{1/2}{\Longmapsto} P * [\mathsf{mark}]_-^{r_1} * [\mathsf{send}(\mathbf{x})]_1^{r_1} * r' \overset{1}{\Longmapsto} Q * r'' \overset{1}{\Longmapsto} Q * \\ &\mathsf{oreg}\left(r_1, \left\{\mathsf{Chain}(rs, d) \,\middle|\, \begin{aligned} &\mathsf{wf}(rs, d) \wedge \mathsf{leaf}_d(\mathbf{x}) \wedge rs(\mathbf{x}).\mathit{flg} = 0 \\ &\wedge\, d(\mathbf{x}.\mathsf{hdr}).\mathit{loff} = \mathbf{x}.\mathsf{off} \wedge \mathbf{x}.\mathsf{off} < \mathsf{MAX} \end{aligned}\right\}\right) \end{aligned}\right\}$

8   `x.hdr->flags[x.loff+1] = 0;`

9   `x.hdr->loff++;`

10   *// Transition step* ext. *Use Lemmas* 7.6 *and* 7.8 *to show perms available.*

11 $\left\{\begin{aligned} &\exists r_1, r_2.\, r_2 \overset{1/2}{\Longmapsto} P * r' \overset{1/2}{\Longmapsto} Q * r'' \overset{1/2}{\Longmapsto} Q * \\ &[\mathsf{mark}]_-^{r_1} * [\mathsf{send}(\mathbf{x})]_1^{r_1} * [\mathsf{send}(\langle \mathbf{x}.\mathsf{hdr}, \mathbf{x}.\mathsf{off} + 1\rangle)]_1^{r_1} * [\mathsf{change}(\langle \mathbf{x}.\mathsf{hdr}, \mathbf{x}.\mathsf{off} + 1\rangle, r'')]_1^{r_1} * \\ &\mathsf{oreg}\left(r_1, \left\{\mathsf{Chain}(rs, d) \,\middle|\, \begin{aligned} &rs = rs_1 \cdot rs(\mathbf{x}) \cdot (\langle \mathbf{x}.\mathsf{hdr}, \mathbf{x}.\mathsf{off} + 1\rangle, r', \{r''\}, 0, \emptyset) \cdot rs_2 \wedge \\ &\mathsf{wf}(rs, d) \wedge \mathsf{leaf}_d(\mathbf{x}) \wedge rs(\mathbf{x}).\mathit{flg} = 0 \wedge \mathsf{leaf}_d(\langle \mathbf{x}.\mathsf{hdr}, \mathbf{x}.\mathsf{off} + 1\rangle) \\ &\wedge\, d(\mathbf{x}.\mathsf{hdr}).\mathit{loff} = \mathbf{x}.\mathsf{off} + 1 \wedge \mathbf{x}.\mathsf{off} < \mathsf{MAX} \end{aligned}\right\}\right) \end{aligned}\right\}$

12   `r.hdr = x.hdr; r.off = x.off + 1;`

13   `nx = x;`

14 $\left\{\begin{aligned} &\exists r_1, r_2.\, r_2 \overset{1/2}{\Longmapsto} P * r' \overset{1/2}{\Longmapsto} Q * r'' \overset{1/2}{\Longmapsto} Q * \\ &[\mathsf{mark}]_-^{r_1} * [\mathsf{send}(\mathbf{nx})]_1^{r_1} * [\mathsf{send}(\mathbf{r})]_1^{r_1} * [\mathsf{change}(\mathbf{r}, r'')]_1^{r_1} * \\ &\mathsf{oreg}\left(r_1, \left\{\mathsf{Chain}(rs, d) \,\middle|\, \begin{aligned} &rs = rs_1 \cdot rs(\mathbf{nx}) \cdot (\mathbf{r}, r', \{r''\}, 0, \emptyset) \cdot rs_2 \wedge \mathsf{wf}(rs, d) \wedge \\ &\mathsf{leaf}_d(\mathbf{nx}) \wedge rs(\mathbf{nx}).\mathit{flg} = 0 \wedge \mathsf{leaf}_d(\mathbf{r}) \wedge rs(\mathbf{r}).\mathit{flg} = 0 \end{aligned}\right\}\right) \end{aligned}\right\}$

15   `} else {`

16 $\left\{\begin{aligned} &\exists r_1, r_2.\, r_2 \overset{1/2}{\Longmapsto} P * r' \overset{1}{\Longmapsto} Q * r'' \overset{1}{\Longmapsto} Q * [\mathsf{mark}]_-^{r_1} * [\mathsf{send}(\mathbf{x})]_1^{r_1} * \\ &\mathsf{oreg}(r_1, \{\mathsf{Chain}(rs, d) \mid \mathsf{wf}(rs, d) \wedge \mathsf{leaf}_d(\mathbf{x}) \wedge rs(\mathbf{x}).\mathit{flg} = 0\}) \end{aligned}\right\}$

17   `nh = malloc(chan_hdr);`

18   `nh->up = x; nh->loff = 1;`

19   *// Transition step* ext. *Use Lemmas* 7.7 *and* 7.8 *to show perms available.*

20 $\left\{\begin{aligned} &\exists r_1, r_2.\, r_2 \overset{1/2}{\Longmapsto} P * r' \overset{1/2}{\Longmapsto} Q * r'' \overset{1/2}{\Longmapsto} Q * \\ &[\mathsf{mark}]_-^{r_1} * [\mathsf{send}(\langle \mathbf{nh}, 0\rangle)]_1^{r_1} * [\mathsf{send}(\langle \mathbf{nh}, 1\rangle)]_1^{r_1} * [\mathsf{change}(\langle \mathbf{nh}, 1\rangle, r'')]_1^{r_1} * \\ &\mathsf{oreg}\left(r_1, \left\{\mathsf{Chain}(rs, d) \,\middle|\, \begin{aligned} &rs = rs_1 \cdot (rs(\mathbf{x}) \blacktriangleleft_{loc} \langle \mathbf{nh}, 0\rangle) \cdot (\langle \mathbf{nh}, 1\rangle, r', \{r''\}, 0, \emptyset) \cdot rs_2 \wedge \\ &\mathsf{wf}(rs, d) \wedge \mathsf{leaf}_d(\mathbf{x}) \wedge rs(\mathbf{x}).\mathit{flg} = 0 \wedge \mathsf{leaf}_d(\langle \mathbf{nh}, 0\rangle) \wedge \mathsf{leaf}_d(\langle \mathbf{nh}, 1\rangle) \\ &\wedge\, rs(\langle \mathbf{nh}, 0\rangle).\mathit{flg} = 0 \wedge rs(\langle \mathbf{nh}, 1\rangle).\mathit{flg} = 0 \end{aligned}\right\}\right) \end{aligned}\right\}$

21   `nx.hdr = nh; nx.off = 0;`

22   `r.hdr = nh; r.off = 1;`

23   `}`

24   `return (r,nx);`

25 `}`

26 $\left\{\mathsf{send}(\mathbf{r}, Q) * \mathsf{recv}(\mathbf{r}, Q) * \mathsf{send}(\mathbf{nx}, P) * \mathbf{r} \prec \mathbf{nx} * \circledast_{e \in E}\, e \prec \mathbf{r} * \circledast_{l \in L}\, \mathbf{nx} \prec l\right\}$

Fig. 30. Sketch-proof of `extend` with summarization.

finalleaf parent, meaning it must be in $\mathsf{C}(rs', d')$. Thus $\langle x, i+1 \rangle \in \mathsf{uC}(rs', d')$. This suffices to show that $[\mathsf{send}(\langle x, i+1 \rangle)]$ and $[\mathsf{change}(\langle x, i+1 \rangle, r')]$ can be safely removed from unused. $\square$

LEMMA 7.7.

$$\begin{pmatrix} \mathsf{ext}_s(\langle x, i \rangle, (rs, d), (rs', d')) \wedge \mathrm{dom}(d') = \mathrm{dom}(d) \uplus \{l\} \wedge r' \notin rs \wedge \\ rs'(\langle l, 1 \rangle).\mathcal{I} = \{r'\} \wedge \mathsf{wf}(rs, d) \wedge \mathsf{wf}(rs', d') \wedge \mathsf{unused}(r, rs, d) * [\mathsf{send}(\langle x, i \rangle)]_1^r \end{pmatrix} \implies$$
$$\mathsf{unused}(r, rs', d') * [\mathsf{send}(\langle l, 0 \rangle)]_1^r * [\mathsf{send}(\langle l, 1 \rangle)]_1^r * [\mathsf{change}(\langle l, 1 \rangle, r')]_1^r$$

PROOF. By the structure of ext, $\langle l, 0 \rangle$ and $\langle l, 1 \rangle$ are not in rs, but are in rs'. The ability to retrieve $[\mathsf{send}(\langle l, 0 \rangle)]_1^r * [\mathsf{send}(\langle l, 1 \rangle)]_1^r$ follows immediately. As $\langle l, 0 \rangle$ is leftmost in the array l, its parent $\langle x, i \rangle$ is the maximal final-leaf in $\mathsf{C}(rs', d')$. However, $\langle l, 1 \rangle$ is not leftmost, and thus is in $\mathsf{C}(rs', d')$. By the same argument used in the previous lemma, $[\mathsf{change}(\langle l, 1 \rangle, r')]_1^r$ can be removed from the unused. $\square$

LEMMA 7.8. $\mathsf{ext}(x, (rs, d), (rs', d')) \wedge \mathsf{wf}(rs, d) \wedge \mathsf{wf}(rs', d') \implies \mathsf{uC}(rs, d) \subseteq \mathsf{uC}(rs', d')$

PROOF. There are two cases for extension: in-place extension in the array, or creation of a new array. In the former case, finalleaf is preserved for existing nodes because the only new node is less than all existing nodes in the array. In the latter case, the parent of the new array is a finalleaf to the new array, and all other finalleaf relationships are preserved.

Now pick a pair $(x, i) \in \mathsf{uC}(rs, d)$, the set of used send permissions. Extending the chain can't stop x from satisfying finalleaf or make any node higher than x satisfy finalleaf. Therefore $x \in \mathsf{C}(rs', d')$ after extension. The only alteration to renounced sets $\mathcal{W}$ in rs' is to add a new empty set. Thus $\neg \exists y. \, r \in rs'(y).\mathcal{W}$. Finally, both cases of extension preserve the promise sets $\mathcal{I}$, ensuring that $r \in rs'(\mathsf{finalleaf}_{d'}(x)).\mathcal{I}$. $\square$

7.4. Verifying splitting and renunciation axioms

The splitting and renunciation axioms do not depend on the underlying data-structure representation, and therefore are largely identical to the ones given in §6.4. The main difference is the new definition of unused. The renunciation case is straightforward, but for splitting we need to show that we can pull the appropriate change permissions out of the 'library' predicate unused. This is captured by the following lemma:

LEMMA 7.9.

$$\mathsf{unused}(r_1, rs, d) * [\mathsf{change}(x, r_2)]_1^{r_1} \wedge \mathsf{split}(x, (rs, d), (rs', d')) \wedge$$
$$rs(\mathsf{finalleaf}_d(x)) = (r, \mathcal{I} \uplus \{r_2\}, f, \mathcal{W}) \wedge rs' = rs \blacktriangleleft_{\mathsf{finalleaf}_{d'}(x)} \blacktriangleleft_{\mathcal{I}} ((\bullet \setminus \{r_2\}) \uplus \{r_3, r_4\})$$
$$\implies \mathsf{unused}(r_1, rs', d') * [\mathsf{change}(x, r_3)]_1^{r_1} * [\mathsf{change}(x, r_4)]_1^{r_1}$$

PROOF. By the definition of split, $d = d'$, and element locations in rs are unchanged in rs'. Thus it holds that $\mathsf{C}(rs, d) = \mathsf{C}(rs', d')$ and $\mathsf{finalleaf}_d(x) = \mathsf{finalleaf}_{d'}(x)$. From the definition of uC the available change permissions are controlled by the set $rs(\mathsf{finalleaf}_d(x)).\mathcal{I}$. This set is correctly updated by the transition, which completes the proof. $\square$

8. COMPARISON TO CONFERENCE PAPER

This paper substantially expands and revises the proofs of correctness given in our conference paper [Dodds et al. 2011]. All the proofs have been restructured, and the proof of the summarising implementation (§7) is entirely new. This paper also fixes a subtle logical error which rendered some of the reasoning in our conference paper unsound. In this section, we describe how this problem arose, and how we have fixed it.

Our specifications rely crucially on higher-order quantification to abstract over the resources transferred through channels. To support this, in [Dodds et al. 2011] we extended the

original concurrent abstract predicates logic [Dinsdale-Young et al. 2010] with higher-order assertions and quantification.

In concurrent abstract predicates, resources describe not only the current state of shared regions but also the protocols that govern these shared regions. In the case of higher-order shared resources, these protocols are themselves expressed in terms of assertion variables that might be instantiated with shared resources. Support for such higher-order shared resources thus require a semantic domain of protocols that include assertions over (among other things) protocols. This results in a circularity and the resulting equation ($protocol \cong \mathcal{P}(... \times protocol)$) has no solution in set-theory, by a simple cardinality argument.

The logic and model presented in [Dodds et al. 2011] broke this circularity by ignoring protocol assertions when interpreting protocols. As a consequence, many of the properties we relied on when reasoning about the higher-order resources $\mathsf{box}(i, P, \pi)$ and $\mathsf{fut}(i, P)$ are unsound. (In that paper, fut played a similar role to recv in this paper, while box was used in verifying the splitting axiom.) For instance, $\mathsf{fut}(i, P)$ is generally not stable when P is instantiated with an assertion that includes a protocol assertion, because $\mathsf{fut}(i, P)$ asserts the existence of a shared region whose protocol is defined in terms of P. A more detailed discussion of this class of problem is given in [Svendsen et al. 2013].

The program *logic* itself presented in [Dodds et al. 2011] still appears sound. However, many steps in the *proofs of programs* depend on unsound auxiliary entailment steps. These steps are common in separation logic proofs, but in most earlier work entailments generally capture comparatively simple properties. We failed to appreciate how deeply the proofs in [Dodds et al. 2011] relied on very subtle entailments between shared regions that were broken in our modified model. The problem came to light a year later when Svendsen attempted to use our logic to verify the Joins library [Svendsen et al. 2013]. Resolving this kind of problem motivated the development of iCAP, which is the proof technique we use in this paper.

iCAP uses step-indexing to stratify the construction of the semantic domain of protocols. The resulting logic does support higher-order shared resources, but requires $\triangleright$ operators to ensure that protocols are properly stratified. Thus the problematic circularity in [Dodds et al. 2011] is appropriately resolved in the rules of the logic. At the level of human process, we have been much more meticulous in this paper in identifying and checking entailment steps used in program proofs.

REFERENCES

Christian J. Bell, Andrew Appel, and David Walker. 2009. Concurrent Separation Logic for Pipelined Parallelization. In *SAS*.

Emery D. Berger, Ting Yang, Tongping Liu, and Gene Novark. 2010. Grace: Safe Multithreaded Programming for C/C++. In *OOPSLA*.

Lars Birkedal, Rasmus Ejlers Møgelberg, Jan Schwinghammer, and Kristian Støvring. 2012. First steps in synthetic guarded domain theory: step-indexing in the topos of trees. *Logical Methods in Computer Science* 8, 4 (2012).

Robert L. Bocchino, Jr., Vikram S. Adve, Danny Dig, Sarita V. Adve, Stephen Heumann, Rakesh Komuravelli, Jeffrey Overbey, Patrick Simmons, Hyojin Sung, and Mohsen Vakilian. 2009. A type and effect system for deterministic parallel Java. In *OOPSLA '09*. ACM, 97–116.

Richard Bornat, Cristiano Calcagno, Peter O'Hearn, and Matthew Parkinson. 2005. Permission accounting in separation logic. In *POPL*. 259–270.

M. Botinčan, M. Dodds, and S. Jagannathan. 2013. Resource-Sensitive Synchronization Inference by Abduction. In *ACM Trans. on Programming Languages and Systems*.

P. da Rocha Pinto, T. Dinsdale-Young, M. Dodds, P. Gardner, and M. Wheelhouse. 2011. A Simple Abstraction for Complex Concurrent Indexes. In *OOPSLA*.

P. da Rocha Pinto, T. Dinsdale-Young, and P. Gardner. 2014. TaDA: A Logic for Time and Data Abstraction. In *ECOOP*. To appear.

Thomas Dinsdale-Young, Mike Dodds, Philippa Gardner, Matthew J Parkinson, and Viktor Vafeiadis. 2010. Concurrent Abstract Predicates. In *ECOOP*.

Mike Dodds, Xinyu Feng, Matthew J. Parkinson, and Viktor Vafeiadis. 2009. Deny-Guarantee Reasoning. In *ESOP*.

M. Dodds, S. Jagannathan, and M. J. Parkinson. 2011. Modular Reasoning for Deterministic Parallelism. In *POPL*.

Xinyu Feng, Rodrigo Ferreira, and Zhong Shao. 2007. On the Relationship between Concurrent Separation Logic and Assume-Guarantee Reasoning. In *ESOP*.

Alexey Gotsman, Josh Berdine, Byron Cook, Noam Rinetzky, and Mooly Sagiv. 2007. Local Reasoning for Storable Locks and Threads. In *APLAS*.

Christian Haack, Marieke Huisman, and Clément Hurlin. 2008. Reasoning about Java's Reentrant Locks. In *APLAS*. 171–187.

C. A. R. Hoare and Peter W. O'Hearn. 2008. Separation Logic Semantics for Communicating Processes. *ENTCS* 212 (2008), 3–25.

Aquinas Hobor, Andrew W. Appel, and Francesco Zappa Nardelli. 2008. Oracle Semantics for Concurrent Separation Logic. In *ESOP*.

Bart Jacobs and Frank Piessens. 2009. *Modular Full Functional Specification and Verification of Lock-Free Data Structures*. Technical Report CW 551. Katholieke Universiteit Leuven, Dept. of Computer Science.

Cliff B. Jones. 1983. Tentative steps toward a development method for interfering programs. *TOPLAS* 5, 4 (1983), 596–619.

Neel R. Krishnaswami, Lars Birkedal, and Jonathan Aldrich. 2010. Verifying event-driven programs using ramified frame properties. In *TLDI*.

K. R. M. Leino, P. Müller, and J. Smans. 2010. Deadlock-free Channels and Locks. In *ESOP*.

A. Navabi, X. Zhang, and S. Jagannathan. 2008. Quasi-static Scheduling for Safe Futures. In *PPoPP*. ACM, 23–32.

Peter W. O'Hearn. 2007. Resources, Concurrency and Local Reasoning. *TCS* (2007).

M. J. Parkinson and G. M. Bierman. 2005. Separation logic and abstraction. In *POPL*. 247–258.

Martin C. Rinard and Monica S. Lam. 1992. Semantic Foundations of Jade. In *POPL*. ACM, 105–118.

Kasper Svendsen and Lars Birkedal. 2014a. Impredicative Concurrent Abstract Predicates. In *ESOP*.

Kasper Svendsen and Lars Birkedal. 2014b. *Impredicative Concurrent Abstract Predicates*. Technical Report. Aarhus University. https://bitbucket.org/logsem/public/src/master/icap/esop2014-tr.pdf

Kasper Svendsen, Lars Birkedal, and Matthew J. Parkinson. 2013. Joins: A Case Study in Modular Specification of a Concurrent Reentrant Higher-Order Library. In *ECOOP*.

Aaron Turon, Derek Dreyer, and Lars Birkedal. 2013. Unifying refinement and hoare-style reasoning in a logic for higher-order concurrency. In *ICFP*.

Viktor Vafeiadis. 2007. *Modular Fine-Grained Concurrency Verification*. Ph.D. Dissertation. University of Cambridge.

Viktor Vafeiadis and Matthew J Parkinson. 2007. A marriage of Rely/Guarantee and Separation Logic. In *CONCUR*. 256–271.

Jules Villard, Étienne Lozes, and Cristiano Calcagno. 2010. Tracking Heaps That Hop with Heap-Hop. In *TACAS*. 275–279.

A. Welc, S. Jagannathan, and A. Hosking. 2005. Safe Futures for Java. In *OOPSLA*. 439–435.

John Wickerson, Mike Dodds, and Matthew Parkinson. 2010. Explicit Stabilisation for Modular Rely-Guarantee Reasoning. In *ESOP*.

A. SAVED PROPOSITIONS

Lemma A.1. *Property (3) of saved propositions implies property (4).*

Proof.

$$r \overset{\pi_1}{\Longmapsto} P * r \overset{\pi_2}{\Longmapsto} Q * (X \mathrel{-\!*} \triangleright(Q * Y)) * (P \mathrel{-\!*} Z)$$

Property (3), frame off saved propositions.

$$\implies \quad (X \mathrel{-\!*} \triangleright(Q * Y)) * (P \mathrel{-\!*} Z) * (\triangleright Q \Rightarrow \triangleright P)$$

Apply SMono, $\triangleright(P \mathrel{-\!*} Q) \implies (\triangleright P \mathrel{-\!*} \triangleright Q)$, *and* $(P \Rightarrow Q) \implies (P \mathrel{-\!*} Q)$.

$$\implies \quad (X \mathrel{-\!*} \triangleright(Q * Y)) * (\triangleright P \mathrel{-\!*} \triangleright Z) * (\triangleright Q \mathrel{-\!*} \triangleright P)$$

Merge $\mathrel{-\!*}$.

$$\implies \quad (X \mathrel{-\!*} \triangleright(Q * Y)) * (\triangleright Q \mathrel{-\!*} \triangleright Z)$$

Add frame $\triangleright Y$ *to* $\mathrel{-\!*}$, *rearrange* $\triangleright$.

$$\implies \quad (X \mathrel{-\!*} \triangleright(Q * Y)) * (\triangleright(Q * Y) \mathrel{-\!*} \triangleright(Z * Y))$$

Merge $\mathrel{-\!*}$, *frame saved propositions back on.*

$$\implies \quad r \overset{\pi_1}{\Longmapsto} P * r \overset{\pi_2}{\Longmapsto} Q * (X \mathrel{-\!*} \triangleright(Z * Y))$$

□

A.1. Encoding in iCAP

A saved proposition $r \overset{\pi}{\Longmapsto} P$ is encoded as a normal iCAP predicate with a structure guaranteeing the properties we want. Intuitively, this predicate consists of a shared region with identifier r, with the proposition P encoded into its transition relation. Linearity comes from a permission with fractional argument π.

More formally, we assume two transition-system states $\{\mathsf{1st}, \mathsf{2nd}\}$ and a single token tok, and invariant map I_{prp} and transition relation T_{prp} defined as follows:

$$I_{\mathsf{prp}}(Q)(\mathsf{1st}) \triangleq \mathsf{emp}$$
$$I_{\mathsf{prp}}(Q)(\mathsf{2nd}) \triangleq Q$$
$$T_{\mathsf{prp}}(\mathsf{tok}) \triangleq \{(\mathsf{1st}, \mathsf{2nd})\}$$

We then define the saved proposition $r \overset{\pi}{\Longmapsto} Q$ as follows:

$$r \overset{\pi}{\Longmapsto} Q \quad \triangleq \quad \mathsf{region}(\{\mathsf{1st}, \mathsf{2nd}\}, T_{\mathsf{prp}}, I_{\mathsf{prp}}(Q), r) * [\mathsf{tok}]^\pi_r$$

The fact that the representation transition state $\mathsf{1st}$ is emp means that we are not obliged to supply P when creating the saved proposition. The second state encodes the value of the saved proposition.

The linearity property (property 2) holds trivially from the linearity of permissions. Two saved propositions with arguments π_1 and π_2 must contain tok permissions with fractional arguments π_1 and π_2. Combining these gives the required result.

For the unification property (property 3) we need to reason more deeply about the iCAP model. The following facts about regions and invariant maps hold in iCAP – for proofs see [Svendsen and Birkedal 2014b].

$$\mathsf{region}(S, T, I, r) * \mathsf{region}(S', T', J, r) \quad \implies \quad (\triangleright I(s) \Rightarrow \triangleright J(s)) \tag{5}$$

$$\mathsf{region}(S, T, I, r) * \mathsf{region}(S', T', J, r) \quad \implies \quad (\triangleright I(s) \mathrel{-\!*} \triangleright J(s)) \tag{6}$$

The later modality, $\triangleright$, is needed in these properties because we are reasoning about the contents of a shared region – albeit one that will not contain any resource. We can then prove the unification property as follows:

PROOF (PROPERTY 3).

$$r \overset{\pi_1}{\Longmapsto} P * r \overset{\pi_2}{\Longmapsto} Q$$

$\Rightarrow\ r \overset{\pi_1}{\Longmapsto} P * \mathsf{region}(\{\mathsf{1st}, \mathsf{2nd}\}, T_{\mathsf{prp}}, I_{\mathsf{prp}}(P), r)$ region *is duplicable.*

    $* \, r \overset{\pi_2}{\Longmapsto} Q * \mathsf{region}(\{\mathsf{1st}, \mathsf{2nd}\}, T_{\mathsf{prp}}, I_{\mathsf{prp}}(Q), r)$

$\Rightarrow\ r \overset{\pi_1}{\Longmapsto} P * r \overset{\pi_2}{\Longmapsto} Q * (\triangleright(I_{\mathsf{prp}}(P)(\mathsf{2nd})) \Rightarrow \triangleright(I_{\mathsf{prp}}(Q)(\mathsf{2nd})))$ *Property 5.*

$\Rightarrow\ r \overset{\pi_1}{\Longmapsto} P * r \overset{\pi_2}{\Longmapsto} Q * (\triangleright P \Rightarrow \triangleright Q)$ *defn of* I_{prp}

$\square$

For unification inside separating implication we reason as follows:

PROOF (PROPERTY 4). Using property 6 and the same proof technique as above, we can derive a slightly different version of the unification property:

$$r \overset{\pi_1}{\Longmapsto} P * r \overset{\pi_2}{\Longmapsto} Q \quad \implies \quad (\triangleright P) \mathrel{-\!\!*} (\triangleright Q) \tag{7}$$

The proof of Property (4) then goes as follows:

$r \overset{\pi_1}{\Longmapsto} P * r \overset{\pi_2}{\Longmapsto} Q * (X \mathrel{-\!\!*} \triangleright(Q * Y)) * (P \mathrel{-\!\!*} Z)$

  SMono

$\Rightarrow\ r \overset{\pi_1}{\Longmapsto} P * r \overset{\pi_2}{\Longmapsto} Q * (X \mathrel{-\!\!*} \triangleright(Q * Y)) * \triangleright(P \mathrel{-\!\!*} Z)$

  LBin, assume $\triangleright$ *distributes over* $\mathrel{-\!\!*}$

$\Rightarrow\ r \overset{\pi_1}{\Longmapsto} P * r \overset{\pi_2}{\Longmapsto} Q * (X \mathrel{-\!\!*} (\triangleright Q * \triangleright Y)) * ((\triangleright P) \mathrel{-\!\!*} \triangleright Z)$

  Assume property 7

$\Rightarrow\ r \overset{\pi_1}{\Longmapsto} P * r \overset{\pi_2}{\Longmapsto} Q * (X \mathrel{-\!\!*} (\triangleright Q * \triangleright Y)) * ((\triangleright P) \mathrel{-\!\!*} \triangleright Z) * ((\triangleright Q) \mathrel{-\!\!*} (\triangleright P))$

  Transitivity of $\mathrel{-\!\!*}$

$\Rightarrow\ r \overset{\pi_1}{\Longmapsto} P * r \overset{\pi_2}{\Longmapsto} Q * (X \mathrel{-\!\!*} (\triangleright Q * \triangleright Y)) * ((\triangleright Q) \mathrel{-\!\!*} \triangleright Z)$

  Framing of $\mathrel{-\!\!*}$

$\Rightarrow\ r \overset{\pi_1}{\Longmapsto} P * r \overset{\pi_2}{\Longmapsto} Q * (X \mathrel{-\!\!*} (\triangleright Q * \triangleright Y)) * ((\triangleright Q) * (\triangleright Y)) \mathrel{-\!\!*} ((\triangleright Z) * (\triangleright Y))$

  Transitivity of $\mathrel{-\!\!*}$

$\Rightarrow\ r \overset{\pi_1}{\Longmapsto} P * r \overset{\pi_2}{\Longmapsto} Q * (X \mathrel{-\!\!*} ((\triangleright Z) * (\triangleright Y)))$

$\square$

B. PROOFS FOR OUT-OF-ORDER SIGNALLING

This appendix gives proofs for some of the lemmas stated in §6.

LEMMA 6.1.

$$\mathsf{resource}(\mathcal{I}, \emptyset) \sqsubseteq \underset{i \in \mathcal{I}}{\circledast}.\, \exists Q \colon \mathsf{Prop}.\, i \overset{1/2}{\Longmapsto} Q(i) * \lceil \triangleright Q(i) \rceil$$

PROOF.

$$\mathcal{W} = \emptyset \wedge \exists Q\colon \mathcal{I} \to \mathsf{Prop}, \ R\colon \mathcal{W} \to \mathsf{Prop}.$$
$$\circledast_{i \in \mathcal{I}}.\, i \overset{1/2}{\Longmapsto} Q(i) * \circledast_{w \in \mathcal{W}}.\, w \overset{1/2}{\Longmapsto} R(w)$$
$$* \lfloor (\triangleright \circledast_{w \in \mathcal{W}}.\, R(w)) \rightarrow\!\!* \triangleright \circledast_{i \in \mathcal{I}}.\, \lceil Q(i) \rceil \rfloor$$

Simplify using $\mathcal{W} = \emptyset$, weakening.

$$\sqsubseteq \quad \exists Q\colon \mathcal{I} \to \mathsf{Prop}.\, \circledast_{i \in \mathcal{I}}.\, i \overset{1/2}{\Longmapsto} Q(i) * \lfloor \triangleright \circledast_{i \in \mathcal{I}}.\, Q(i) \rfloor$$

Switch from $\lfloor - \rfloor$ to $\lceil - \rceil$, pull out $\circledast$.

$$\sqsubseteq \quad \exists Q\colon \mathcal{I} \to \mathsf{Prop}.\, \circledast_{i \in \mathcal{I}}.\, i \overset{1/2}{\Longmapsto} Q(i) * \circledast_{i \in \mathcal{I}}.\, \lceil \triangleright Q(i) \rceil$$

Push in the existential.

$$\sqsubseteq \quad \circledast_{i \in \mathcal{I}}.\, \exists Q\colon \mathsf{Prop}.\, i \overset{1/2}{\Longmapsto} Q(i) * \lceil \triangleright Q(i) \rceil$$

$\square$

LEMMA 6.2. $\quad r \overset{1/2}{\Longmapsto} P * \lfloor P \rfloor * \mathsf{resource}(\mathcal{I}, \mathcal{W} \uplus \{r\}) \quad \sqsubseteq \quad \mathsf{resource}(\mathcal{I}, \mathcal{W})$

PROOF.

$$r \overset{1/2}{\Longmapsto} R * \lfloor R \rfloor * \mathsf{resource}(\mathcal{I}, \mathcal{W} \uplus \{r\})$$

Definition of $\mathsf{resource}$.

$$\sqsubseteq \quad r \overset{1/2}{\Longmapsto} R * \lfloor R \rfloor *$$
$$\left(\begin{array}{c} \exists Q\colon \mathcal{I} \to \mathsf{Prop}, \ R\colon \mathcal{W} \uplus \{r\} \to \mathsf{Prop}. \\ \circledast_{i \in \mathcal{I}}.\, i \overset{1/2}{\Longmapsto} Q(i) * \circledast_{w \in \mathcal{W} \uplus \{r\}}.\, w \overset{1/2}{\Longmapsto} R(w) \\ * \lfloor (\triangleright \circledast_{w \in \mathcal{W} \uplus \{r\}}.\, R(w)) \rightarrow\!\!* \triangleright \circledast_{i \in \mathcal{I}}.\, \lceil Q(i) \rceil \rfloor \end{array} \right)$$

Property (3), monotonicity of $\triangleright$, monotonicity of $\lfloor - \rfloor$.

$$\sqsubseteq \quad \exists P\colon \mathsf{Prop}.\, r \overset{1/2}{\Longmapsto} R * \lfloor \triangleright R \rfloor * (\lfloor \triangleright R \rfloor \Rightarrow \lfloor \triangleright P \rfloor)$$
$$\left(\begin{array}{c} \exists Q\colon \mathcal{I} \to \mathsf{Prop}, \ R\colon \mathcal{W} \to \mathsf{Prop}. \\ r \overset{1/2}{\Longmapsto} P * \circledast_{i \in \mathcal{I}}.\, i \overset{1/2}{\Longmapsto} Q(i) * \circledast_{w \in \mathcal{W}}.\, w \overset{1/2}{\Longmapsto} R(w) \\ * \lfloor (\triangleright P * \triangleright \circledast_{w \in \mathcal{W}}.\, R(w)) \rightarrow\!\!* \triangleright \circledast_{i \in \mathcal{I}}.\, \lceil Q(i) \rceil \rfloor \end{array} \right)$$

Modus ponens.

$$\sqsubseteq \quad \exists P\colon \mathsf{Prop}.\, r \overset{1/2}{\Longmapsto} R * \lfloor \triangleright P \rfloor *$$
$$\left(\begin{array}{c} \exists Q\colon \mathcal{I} \to \mathsf{Prop}, \ R\colon \mathcal{W} \to \mathsf{Prop}. \\ r \overset{1/2}{\Longmapsto} P * \circledast_{i \in \mathcal{I}}.\, i \overset{1/2}{\Longmapsto} Q(i) * \circledast_{w \in \mathcal{W}}.\, w \overset{1/2}{\Longmapsto} R(w) \\ * \lfloor (\triangleright P * \triangleright \circledast_{w \in \mathcal{W}}.\, R(w)) \rightarrow\!\!* \triangleright \circledast_{i \in \mathcal{I}}.\, \lceil Q(i) \rceil \rfloor \end{array} \right)$$

Combine $\lfloor - \rfloor$, modus ponens for $\rightarrow\!\!*$, weakening.

$$\sqsubseteq \quad \exists Q\colon \mathcal{I} \to \mathsf{Prop}, \ R\colon \mathcal{W} \to \mathsf{Prop}.$$
$$\circledast_{i \in \mathcal{I}}.\, i \overset{1/2}{\Longmapsto} Q(i) * \circledast_{w \in \mathcal{W}}.\, w \overset{1/2}{\Longmapsto} R(w)$$
$$* \lfloor (\triangleright \circledast_{w \in \mathcal{W}}.\, R(w)) \rightarrow\!\!* \triangleright \circledast_{i \in \mathcal{I}}.\, \lceil Q(i) \rceil \rfloor$$

$$\sqsubseteq \quad \mathsf{resource}(\mathcal{I}, \mathcal{W})$$

$\square$

LEMMA 6.3. $\quad \mathsf{chainres}(rs) \wedge \mathsf{wf}(rs) \wedge \mathsf{ctrue}(rs)$
$$\sqsubseteq \quad \exists rs'.\, \mathsf{chainres}(rs') \wedge \mathsf{cconf}(rs') \wedge rs \overset{pr}{\to}{}^* rs' \wedge \mathsf{wf}(rs')$$

PROOF. We perform a sequence of smaller view-shifts corresponding to converting each cnode in turn, starting with the earliest element in the chain:

$$P_0 \sqsubseteq P_1 \sqsubseteq P_2 \sqsubseteq \ldots \sqsubseteq P_n$$

Here n is the length of rs and the subscript $1, 2, 3 \ldots$ denotes the length of suffix of rs which has been checked. We write $rs[a, b]$ for the subsequence of rs from element a to element b, inclusive of both. Thus the inductive invariant is:

$$
\begin{aligned}
P_i \quad \triangleq \quad & \exists rs'. \, \mathsf{chainres}(rs[0, n-i] \cdot rs') \wedge \mathsf{cconf}(rs') \\
& \wedge \, rs[(n-i)+1, n] \xrightarrow{pr}_* rs' \wedge \mathsf{wf}(rs[0, (n-i)] \cdot rs')
\end{aligned}
$$

The base case of the proof is simple. If $i = 0$ take rs' to be empty and the invariant follows trivially from the premise. Let us assume that $i > 0$. We reason as follows to pull out the intermediate chain node:

$$\exists rs'. \, \mathsf{chainres}(rs[0, n-i] \cdot rs') \wedge \mathsf{cconf}(rs') \wedge rs[(n-i)+1, n] \xrightarrow{pr}_* rs' \wedge \mathsf{wf}(rs[0, n-i] \cdot rs')$$

$$\sqsubseteq$$

$$\exists rs', s. \, \mathsf{chainres}(rs[0, n-(i+1)]) * \mathsf{resource}(s.\mathcal{I}, s.\mathcal{W}) * \mathsf{chainres}(rs')$$
$$\wedge \, \mathsf{cconf}(rs') \wedge rs[n-i, n] \xrightarrow{pr}_* s \cdot rs' \wedge \mathsf{wf}(rs[0, n-(i+1))] \cdot s \cdot rs')$$

If $s.\mathcal{W} = \emptyset$ then we are done. Otherwise, we induct on the size of $\mathcal{W}$, showing that it can be reduced to $\emptyset$ by classical entailment. (Recall that by definition $\mathcal{W}$ is finite.)

Pick an element $w \in \mathcal{W}$. Since the chain is well-formed, the region identifier w must also be a member of some set $s'.\mathcal{I}$ for an earlier element $s' \in rs'$. If w is a member of $s'.\mathcal{I}$ for multiple s', we pick the first such $s' \in rs'$. Since cconf holds for rs' it follows that $s'.\mathcal{W} = \emptyset$. Thus, there exists rs_1', rs_2' and s' such that $w \in s'.\mathcal{I}$, $s'.\mathcal{W} = \emptyset$, $rs' = rs_1' \cdot s' \cdot rs_2'$ and $\forall x \in rs_1'. \, w \notin x.\mathcal{I}$. By the definition of resource, there exists a saved proposition $w \xmapsto{1/2} R$ inside $\mathsf{resource}(s.\mathcal{I}, s.\mathcal{W})$. By the same definition, the saved proposition $w \xmapsto{1/2} P$ and $\mathsf{resource} \triangleright P$ must be included in $\mathsf{resource}(s'.\mathcal{I}, \emptyset)$. We move the resource from one node to the other, and delete the saved proposition using weakening:

$$\mathsf{resource}(s.\mathcal{I}, s.\mathcal{W}) * \mathsf{chainres}(rs_1') * \mathsf{resource}(s'.\mathcal{I}, \emptyset) * \mathsf{chainres}(rs_2')$$

$$\sqsubseteq$$

$$\mathsf{resource}(s.\mathcal{I}, s.\mathcal{W} \setminus \{w\}) * \mathsf{chainres}(rs_1') * \mathsf{resource}(s'.\mathcal{I} \setminus \{w\}, \emptyset) * \mathsf{chainres}(rs_2')$$

Let rs'' denote $rs_1' \cdot s'[\mathcal{I} \mapsto s'.\mathcal{I} \setminus \{w\}] \cdot rs_2'$. By definition of $\xrightarrow{\mathcal{W}}$ it thus follows that

$$rs[n-i, n] \xrightarrow{pr}_* s \cdot rs' \xrightarrow{\mathcal{W}} s[\mathcal{W} \mapsto s.\mathcal{W} \setminus \{w\}] \cdot rs''$$

Hence, by Lemma B.1 it follows that

$$\mathsf{wf}(rs[0, n-(i+1))] \cdot s[\mathcal{W} \mapsto s.\mathcal{W} \setminus \{w\}] \cdot rs'')$$

The result is that the assertion is rewritten as follows:

$$\sqsubseteq \quad \exists rs', s. \, \mathsf{chainres}(rs[0, n-(i+1)]) * \mathsf{resource}(s.\mathcal{I}, s.\mathcal{W} \setminus \{w\}) * \mathsf{chainres}(rs')$$
$$\wedge \, \mathsf{cconf}(rs') \wedge rs[n-i, n] \xrightarrow{pr}_* s \cdot rs' \wedge \mathsf{wf}(rs[0, n-(i+1))] \cdot s \cdot rs')$$

Thus we have rewritten $s.\mathcal{W}$ into a smaller set. By inducting on the size of this set we can get to the point where $\mathcal{W}' = \emptyset$. This allows us to complete one step of the outer induction, which completes the inductive proof. $\qquad\square$

LEMMA B.1.

$$\mathsf{wf}(rs) \wedge rs \xrightarrow{\mathcal{W}} rs' \Rightarrow (\mathsf{available}(rs) = \mathsf{available}(rs') \wedge \mathsf{wf}(rs'))$$

PROOF. By definition of $rs \xrightarrow{\mathcal{W}} rs'$ there exists $rs_1, rs_2, rs_3, s_1, s_2$ and w such that

$$rs = rs_1 \cdot s_1 \cdot rs_2 \cdot s_2 \cdot rs_3, \qquad w \in s_2.\mathcal{I}, \qquad w \in s_1.\mathcal{W},$$
$$rs' = rs_1 \cdot (s_1 \blacktriangleleft_\mathcal{W} (\bullet \setminus w)) \cdot rs_2 \cdot (s_2 \blacktriangleleft_\mathcal{I} (\bullet \setminus w)) \cdot rs_3$$

Since $s_2.\mathcal{W} \cap s_2.\mathcal{I} = \emptyset$ and $\mathsf{wf}(s_2 \cdot rs_3)$ it follows that $w \notin (s_2.\mathcal{W} \cup \mathsf{available}(rs_3))$ and $\forall s \in rs_2.\ w \notin s.\mathcal{I}$. Thus,

$$\mathsf{available}((s_2 \blacktriangleleft_\mathcal{I} (\bullet \setminus w)) \cdot rs_3) = (\mathsf{available}(rs_3) \setminus s_2.\mathcal{W}) \uplus (s_2.\mathcal{I} \setminus \{w\})$$
$$= \mathsf{available}(s_2 \cdot rs_3) \setminus \{w\}$$

Since $w \in s_1.\mathcal{W}$ and $\mathsf{wf}(s_1 \cdot rs_2 \cdot s_2 \cdot rs_3)$ it follows that $w \in \mathsf{available}(rs_2 \cdot s_2 \cdot rs_3)$. Hence, since $\forall s \in rs_2.\ w \notin s.\mathcal{I}$ it follows that $\forall s \in rs_2.\ w \notin s.\mathcal{W}$. Thus,

$$\mathsf{available}((s_1 \blacktriangleleft_\mathcal{W} (\bullet \setminus w)) \cdot rs_2 \cdot (s_2 \blacktriangleleft_\mathcal{I} (\bullet \setminus w)) \cdot rs_3)$$
$$= (\mathsf{available}(rs_2 \cdot (s_2 \blacktriangleleft_\mathcal{I} (\bullet \setminus w)) \cdot rs_3) \setminus (s_1.\mathcal{W} \setminus \{w\})) \uplus s_1.\mathcal{I}$$
$$= ((\mathsf{available}(rs_2 \cdot s_2 \cdot rs_3) \setminus \{w\}) \setminus (s_1.\mathcal{W} \setminus \{w\})) \uplus s_1.\mathcal{I}$$
$$= (\mathsf{available}(rs_2 \cdot s_2 \cdot rs_3) \setminus s_1.\mathcal{W}) \uplus s_1.\mathcal{I}$$
$$= \mathsf{available}(s_1 \cdot rs_2 \cdot s_2 \cdot rs_3)$$

from which it follows easily that $\mathsf{available}(rs) = \mathsf{available}(rs')$. To show that $\mathsf{wf}(rs')$ we must also show that

$$s_1.\mathcal{W} \setminus \{w\} \subseteq \mathsf{available}(rs_2 \cdot (s_2 \blacktriangleleft_\mathcal{I} (\bullet \setminus w)) \cdot rs_3)$$
$$= \mathsf{available}(rs_2 \cdot s_2 \cdot rs_3) \setminus \{w\}$$

and $s_2.\mathcal{W} \subseteq \mathsf{available}(rs_3)$ both of which follow easily from the assumption that $\mathsf{wf}(rs)$. It remains to show that all sets $\mathcal{I}$ for the chain are pairwise disjoint, and likewise for all sets $\mathcal{W}$. As we have only removed identifiers, this is satisfied trivially. $\square$

LEMMA 6.4. $\mathsf{resource}(\mathcal{I} \uplus r_2, \emptyset) * r_2 \overset{1/2}{\Longmapsto} P \ \sqsubseteq \ \mathsf{resource}(\mathcal{I}, \emptyset) * \lceil \rhd P \rceil$

PROOF.

$\mathsf{resource}(\mathcal{I} \uplus r_2, \emptyset) * r_2 \overset{1/2}{\Longmapsto} P$

   Lemma 6.1.

$\sqsubseteq \quad r_2 \overset{1/2}{\Longmapsto} P * \circledast_{i \in \mathcal{I} \uplus r_2}. \exists R.\ i \overset{1/2}{\Longmapsto} R * \lceil \rhd R \rceil$

   Pull out resource for identifier r_2.

$\sqsubseteq \quad r_2 \overset{1/2}{\Longmapsto} P * \exists R'.\ r_2 \overset{1/2}{\Longmapsto} R' * \lceil \rhd R' \rceil * \circledast_{i \in \mathcal{I}}. \exists R.\ i \overset{1/2}{\Longmapsto} R * \lceil \rhd R \rceil$

   Property (3), monotonicity of $\lceil - \rceil$.

$\sqsubseteq \quad r_2 \overset{1/2}{\Longmapsto} P * \exists R'.\ (\lceil \rhd R' \rceil \Rightarrow \lceil \rhd P \rceil) * r_2 \overset{1/2}{\Longmapsto} R' * \lceil \rhd R' \rceil * \circledast_{i \in \mathcal{I}}. \exists R.\ i \overset{1/2}{\Longmapsto} R * \lceil \rhd R \rceil$

   Modus ponens, weakening.

$\sqsubseteq \quad \lceil \rhd P \rceil * \circledast_{i \in \mathcal{I}}. \exists R.\ i \overset{1/2}{\Longmapsto} R * \lceil \rhd R \rceil$

   Definition of resource.

$\sqsubseteq \quad \lceil \rhd P \rceil * \mathsf{resource}(\mathcal{I}, \emptyset)$

$\square$

LEMMA 6.7.

$\{\mathsf{oreg}(r, \{\mathsf{Chain}(rs) \mid x \in rs\}) * \mathsf{oreg}(r', \{\mathsf{Chain}(rs') \mid x \in rs'\})\} \quad \langle \mathtt{skip} \rangle \quad \{r = r'\}$

PROOF. Prove this by case-splitting on whether the two regions are equal. Suppose the two are equal – then the specification is proved. If they are unequal, we prove this leads to a contradiction by opening both regions and examining their contents. Each region asserts exclusive ownership of heap cell $x.loc$ which leads to a contradiction. Therefore the post-condition is false, allowing us to prove any post-condition. □

LEMMA 6.8. $\mathsf{resource}(\mathcal{I}, \mathcal{W} \uplus \{r\}) * r \overset{1/2}{\Longmapsto} S * r' \overset{1/2}{\Longmapsto} T_1$

$\qquad\sqsubseteq\quad \exists r''.\, \mathsf{resource}(\mathcal{I}, \mathcal{W} \uplus \{r', r''\}) * r'' \overset{1/2}{\Longmapsto} (T_1 \mathbin{-\!\!*} S)$

PROOF. First we construct a new saved proposition r'' such that $r'' \Rrightarrow (T_1 \mathbin{-\!\!*} S)$. Now it suffices to prove

$\qquad \mathsf{resource}(\mathcal{I}, \mathcal{W} \uplus \{r\}) * r \overset{1/2}{\Longmapsto} S * r' \overset{1/2}{\Longmapsto} T_1 * r'' \overset{1/2}{\Longmapsto} T_2 \wedge \mathsf{valid}(T_1 * T_2 \Rightarrow S)$
$\qquad \sqsubseteq\quad \mathsf{resource}(\mathcal{I}, \mathcal{W} \uplus \{r', r''\})$

$r \overset{1/2}{\Longmapsto} S * r' \overset{1/2}{\Longmapsto} T_1 * r'' \overset{1/2}{\Longmapsto} T_2 \wedge \mathsf{valid}(T_1 * T_2 \Rightarrow S) *$
$\begin{pmatrix} \exists Q \colon \mathcal{I} \to \mathsf{Prop},\ R \colon \mathcal{W} \uplus \{r\} \to \mathsf{Prop}. \\ \circledast_{i \in \mathcal{I}}.\, i \overset{1/2}{\Longmapsto} Q(i) * \circledast_{w \in \mathcal{W} \uplus \{r\}}.\, w \overset{1/2}{\Longmapsto} R(w) \\ * \lfloor (\triangleright \circledast_{w \in \mathcal{W} \uplus \{r\}}.\, R(w)) \mathbin{-\!\!*} \triangleright \circledast_{i \in \mathcal{I}}.\, \lceil Q(i) \rceil \rfloor \end{pmatrix}$

$\quad$ *Rearrange*
$\sqsubseteq\ \exists P \colon \mathsf{Prop},\ Q \colon \mathcal{I} \to \mathsf{Prop},\ R \colon \mathcal{W} \uplus \{r', r''\} \to \mathsf{Prop}.$
$\qquad\quad \mathsf{valid}(R(r') * R(r'') \Rightarrow S) \wedge r \overset{1/2}{\Longmapsto} S * r \overset{1/2}{\Longmapsto} P$
$\qquad\quad * \circledast_{i \in \mathcal{I}}.\, i \overset{1/2}{\Longmapsto} Q(i) * \circledast_{w \in \mathcal{W} \uplus \{r', r''\}}.\, w \overset{1/2}{\Longmapsto} R(w)$
$\qquad\qquad * \lfloor \triangleright P * \triangleright \circledast_{w \in \mathcal{W}}.\, R(w)) \mathbin{-\!\!*} \triangleright \circledast_{i \in \mathcal{I}}.\, \lceil Q(i) \rceil \rfloor$

$\quad$ *Property (3), SMono and distributing $\triangleright$ over $\implies$.*
$\sqsubseteq\qquad\qquad \exists P \colon \mathsf{Prop},\ Q \colon \mathcal{I} \to \mathsf{Prop},\ R \colon \mathcal{W} \uplus \{r', r''\} \to \mathsf{Prop}.$
$\mathsf{valid}((\triangleright R(r')) * (\triangleright R(r'')) \Rightarrow \triangleright P) \wedge \circledast_{i \in \mathcal{I}}.\, i \overset{1/2}{\Longmapsto} Q(i) * \circledast_{w \in \mathcal{W} \uplus \{r', r''\}}.\, w \overset{1/2}{\Longmapsto} R(w)$
$\qquad\qquad\qquad * \lfloor (\triangleright P * \triangleright \circledast_{w \in \mathcal{W}}.\, R(w)) \mathbin{-\!\!*} \triangleright \circledast_{i \in \mathcal{I}}.\, \lceil Q(i) \rceil \rfloor$

$\quad$ *By mono of $\lfloor \_ \rfloor$, $\mathbin{-\!\!*}$ and $*$.*
$\sqsubseteq\qquad\qquad \exists Q \colon \mathcal{I} \to \mathsf{Prop},\ R \colon \mathcal{W} \uplus \{r', r''\} \to \mathsf{Prop}.$
$\qquad\qquad\quad \circledast_{i \in \mathcal{I}}.\, i \overset{1/2}{\Longmapsto} Q(i) * \circledast_{w \in \mathcal{W} \uplus \{r', r''\}}.\, w \overset{1/2}{\Longmapsto} R(w)$
$\qquad * \lfloor ((\triangleright R(r')) * (\triangleright R(r''))) \mathbin{-\!\!*} \triangleright \circledast_{w \in \mathcal{W}}.\, R(w)) \mathbin{-\!\!*} \triangleright \circledast_{i \in \mathcal{I}}.\, \lceil Q(i) \rceil \rfloor$

$\quad$ *Rearrange*
$\sqsubseteq\ \exists P \colon \mathsf{Prop},\ Q \colon \mathcal{I} \to \mathsf{Prop},\ R \colon \mathcal{W} \uplus \{r', r''\} \to \mathsf{Prop}.$
$\qquad\quad \circledast_{i \in \mathcal{I}}.\, i \overset{1/2}{\Longmapsto} Q(i) * \circledast_{w \in \mathcal{W} \uplus \{r', r''\}}.\, w \overset{1/2}{\Longmapsto} R(w)$
$\qquad\quad * \lfloor \triangleright \circledast_{w \in \mathcal{W} \uplus \{r', r''\}}.\, R(w)) \mathbin{-\!\!*} \triangleright \circledast_{i \in \mathcal{I}}.\, \lceil Q(i) \rceil \rfloor$

$\hfill\square$

LEMMA B.2.

$r_2 \overset{1/2}{\Longmapsto} P * \lfloor \lceil P \rceil \mathbin{-\!\!*} (P_1 * P_2) \rfloor * r_3 \overset{1/2}{\Longmapsto} P_1 * r_4 \overset{1/2}{\Longmapsto} P_2 * \mathsf{resource}(\mathcal{I} \uplus \{r_2\}, \mathcal{W})$

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad \sqsubseteq\quad \mathsf{resource}(\mathcal{I} \uplus \{r_3, r_4\}, \mathcal{W})$

PROOF.

$$r_2 \overset{1/2}{\longmapsto} P * \lfloor \lceil P \rceil \twoheadrightarrow (P_1 * P_2) \rfloor * r_3 \overset{1/2}{\longmapsto} P_1 * r_4 \overset{1/2}{\longmapsto} P_2 * \mathsf{resource}(\mathcal{I} \uplus \{r_2\}, \mathcal{W})$$

Definition of resource.

$$\sqsubseteq \quad r_2 \overset{1/2}{\longmapsto} P * \lfloor \lceil P \rceil \twoheadrightarrow (P_1 * P_2) \rfloor * r_3 \overset{1/2}{\longmapsto} P_1 * r_4 \overset{1/2}{\longmapsto} P_2$$

$$\begin{pmatrix} \exists Q \colon \mathcal{I} \uplus \{r_2\} \to \mathsf{Prop}, \ R \colon \mathcal{W} \to \mathsf{Prop}. \\ \circledast_{i \in \mathcal{I} \uplus \{r_2\}}. \ i \overset{1/2}{\longmapsto} Q(i) * \circledast_{w \in \mathcal{W}}. \ w \overset{1/2}{\longmapsto} R(w) \\ * \left(\lfloor (\triangleright \circledast_{w \in \mathcal{W}}. \ R(w)) \twoheadrightarrow \triangleright \circledast_{i \in \mathcal{I} \uplus \{r_2\}}. \ \lceil Q(i) \rceil \rfloor \right) \end{pmatrix}$$

Pull out r_2, $\triangleright$ *mono w.r.t.* $\twoheadrightarrow$, *property* (4).

$$\sqsubseteq \qquad\qquad r_2 \overset{1/2}{\longmapsto} P * r_3 \overset{1/2}{\longmapsto} P_1 * r_4 \overset{1/2}{\longmapsto} P_2$$

$$\begin{pmatrix} \exists Q \colon \mathcal{I} \uplus \{r_2\} \to \mathsf{Prop}, \ R \colon \mathcal{W} \to \mathsf{Prop}. \\ r_2 \overset{1/2}{\longmapsto} Q(r_2) * \circledast_{i \in \mathcal{I}}. \ i \overset{1/2}{\longmapsto} Q(i) * \circledast_{w \in \mathcal{W}}. \ w \overset{1/2}{\longmapsto} R(w) \\ * \lfloor (\triangleright \circledast_{w \in \mathcal{W}}. \ R(w)) \twoheadrightarrow \triangleright (P_1 * P_2 * \circledast_{i \in \mathcal{I}}. \ \lceil Q(i) \rceil) \rfloor \end{pmatrix}$$

Fold r_3, r_4 *into* $\mathcal{I}$, *weaken with* $\lceil \_ \rceil$, *weakening.*

$$\sqsubseteq \qquad \exists Q \colon \mathcal{I} \uplus r_2 \to \mathsf{Prop}, \ R \colon \mathcal{W} \to \mathsf{Prop}.$$

$$\circledast_{i \in \mathcal{I} \uplus \{r_3, r_4\}}. \ i \overset{1/2}{\longmapsto} Q(i) * \circledast_{w \in \mathcal{W}}. \ w \overset{1/2}{\longmapsto} R(w)$$

$$* \lfloor (\triangleright \circledast_{w \in \mathcal{W}}. \ R(w)) \twoheadrightarrow \triangleright \circledast_{i \in \mathcal{I} \uplus \{r_3, r_4\}}. \ \lceil Q(i) \rceil \rfloor)$$

Definition of resource.

$$\sqsubseteq \quad \mathsf{resource}(\mathcal{I} \uplus \{r_3, r_4\}, \mathcal{W})$$

$\square$

LEMMA 6.9. $r_2 \in rs(x).\mathcal{I} \wedge r_2 \overset{1/2}{\longmapsto} P * \lfloor \lceil P \rceil \twoheadrightarrow (P_1 * P_2) \rfloor * \mathsf{chainres}(rs)$
$$\sqsubseteq \quad \exists rs', r_3, r_4. \ r_3, r_4 \notin rs \wedge rs' = rs \blacktriangleleft_x \blacktriangleleft_{\mathcal{I}} (\bullet \setminus r_2) \uplus \{r_3, r_4\} \wedge$$
$$r_3 \overset{1/2}{\longmapsto} P_1 * r_4 \overset{1/2}{\longmapsto} P_2 * \mathsf{chainres}(rs')$$

Proof.

$$r_2 \in rs(x).\mathcal{I} \wedge r_2 \overset{1/2}{\Longmapsto} P * \lfloor \lceil P \rceil \mathbin{-\!\!*} (P_1 * P_2) \rfloor * \mathsf{chainres}(rs)$$

Make saved propositions, fresh by construction.

$$\sqsubseteq \quad r_2 \in rs(x).\mathcal{I} \wedge r_2 \overset{1/2}{\Longmapsto} P * \lfloor \lceil P \rceil \mathbin{-\!\!*} (P_1 * P_2) \rfloor * \mathsf{chainres}(rs) *$$
$$\exists r_3, r_4.\, r_3 \overset{1}{\Longmapsto} P_1 * r_4 \overset{1}{\Longmapsto} P_2 \wedge r_3, r_4 \notin rs$$

Pull out resource *predicate for* x.

$$\sqsubseteq \quad \exists rs_1, rs_2.\, r_2 \in rs(x).\mathcal{I} \wedge r_2 \overset{1/2}{\Longmapsto} P * \lfloor \lceil P \rceil \mathbin{-\!\!*} (P_1 * P_2) \rfloor *$$
$$\mathsf{chainres}(rs_1) * \mathsf{resource}(rs(x).\mathcal{I}, rs(x).\mathcal{W} \uplus \{rs(x).res \mid rs(x).\mathit{flg} = 0\}) * \mathsf{chainres}(rs_2) *$$
$$\exists r_3, r_4.\, r_3 \overset{1}{\Longmapsto} P_1 * r_4 \overset{1}{\Longmapsto} P_2 \wedge r_3, r_4 \notin rs \wedge rs = rs_1 \cdot rs(x) \cdot rs_2$$

Apply Lemma B.2.

$$\sqsubseteq \quad \exists rs_1, rs_2, r_3, r_4.\, r_2 \in rs(x).\mathcal{I} \wedge r_2 \overset{1/2}{\Longmapsto} P * \lfloor \lceil P \rceil \mathbin{-\!\!*} (P_1 * P_2) \rfloor *$$
$$\mathsf{chainres}(rs_1) * \mathsf{resource}((rs(x).\mathcal{I} \setminus r_2) \uplus \{r_3, r_4\}, rs(x).\mathcal{W} \uplus \{rs(x).res \mid rs(x).\mathit{flg} = 0\}) *$$
$$\mathsf{chainres}(rs_2) * r_3 \overset{1/2}{\Longmapsto} P_1 * r_4 \overset{1/2}{\Longmapsto} P_2 \wedge r_3, r_4 \notin rs \wedge rs = rs_1 \cdot rs(x) \cdot rs_2$$

Definition of chainres.

$$\sqsubseteq \quad \exists rs', r_3, r_4.\, r_3, r_4 \notin rs \wedge rs' = rs \blacktriangleleft_x \blacktriangleleft_{\mathcal{I}} (\bullet \setminus r_2) \uplus \{r_3, r_4\} \wedge$$
$$r_3 \overset{1/2}{\Longmapsto} P_1 * r_4 \overset{1/2}{\Longmapsto} P_2 * \mathsf{chainres}(rs')$$

$\square$

C. ICAP PROOF SYSTEM

In this Appendix we introduce the formal iCAP proof system. The introduction is self-contained, but does not cover the full iCAP proof system. In particular, certain iCAP features, such as guarded recursive predicates and phantom state are not necessary for the present paper, and have been dropped from the proof system.

C.1. Syntax

The proof system consists of two logics, an assertion logic and a specification logic, over a common simply-typed term language generated by the following grammar:

$$\mathsf{M}, \mathsf{N}, \mathsf{P}, \mathsf{Q}, \mathsf{S}, \mathsf{T}, \mathsf{F}, \mathsf{R} ::=$$
$$\mid \lambda \mathsf{x} : \tau.\, \mathsf{M} \mid \mathsf{M}\,\mathsf{N} \mid \mathsf{x}$$
$$\mid \perp \mid \top \mid \mathsf{M} \vee \mathsf{N} \mid \mathsf{M} \wedge \mathsf{M} \mid \mathsf{M} \Rightarrow \mathsf{N} \mid \forall \mathsf{x} : \tau.\, \mathsf{P} \mid \exists \mathsf{x} : \tau.\, \mathsf{P} \mid \mathsf{M} =_\tau \mathsf{N}$$
$$\mid \mathsf{P} * \mathsf{Q} \mid \mathsf{P} \mathbin{-\!\!*} \mathsf{Q} \mid \mathsf{emp} \mid \mathsf{M}.\mathsf{F} \mapsto \mathsf{N} \mid \mathsf{M}{:}\mathsf{N} \mid \mathsf{region}(\mathsf{R}, \mathsf{M}, \mathsf{N}) \mid [\mathsf{M}]_\mathsf{N}^\mathsf{R} \mid \mathsf{stable}(\mathsf{P})$$
$$\mid \mathsf{P} \sqsubseteq^\mathsf{R} \mathsf{Q} \mid (\Delta).\{\mathsf{P}\}\overline{\mathsf{s}}\{\mathsf{Q}\} \mid (\Delta).\langle \mathsf{P} \rangle \overline{\mathsf{s}} \langle \mathsf{Q} \rangle^\mathsf{R} \mid \mathsf{M}.\mathsf{N} : (\Delta).\{\mathsf{P}\}\{\mathsf{x}.\mathsf{Q}\} \mid \mathsf{M} : (\Delta).\{\mathsf{P}\}\{\mathsf{x}.\mathsf{Q}\}$$
$$\mid \rhd \mathsf{M} \mid \mathsf{valid}(\mathsf{P}) \mid \mathsf{spec}(\mathsf{S}) \mid \Delta_X(x)$$

Here X is an arbitrary set and x an arbitrary element of X. The $\Delta_X(x)$ gives a shallow embedding of the meta-theory into iCAP. Correspondingly, the grammar of types (given below) features a type constructor $\Delta(X)$ for injecting arbirary sets into iCAP.

$$\text{Types } \tau, \sigma \ ::= \ 1 \mid \tau \to \sigma \mid \tau \times \sigma \mid \tau + \sigma \mid \mathcal{P}(\tau) \mid \Delta(X) \mid \mathsf{Prop} \mid \mathsf{Spec}$$

In addition to the usual type-constructors, iCAP includes two proposition types – one for each logic – the Prop type for the assertion logic and the Spec type for the specification logic. Base-types for values, Val, state identifiers, SId, action identifiers, AId, region identifiers, RId, class names, Class, and field names, Field are just syntactic suger for injections of the

corresponding set (e.g., Val is syntactic suger for $\Delta(\mathit{Val})$). We will usually leave out the explicit injection for elements when reasoning about elements of these injected types.

The typing rules of the logic are given below. The rules have been split into standard higher-order logic typing rules, followed by iCAP specific typing rules. Terms are typed in a logical variable context, Γ, and program variable context, Δ. Logical variables are used purely for specification purposes and may not appear free in the code of Hoare triples. Program variables may appear free in both the pre- and postcondition of Hoare triples and in the code snippet. The logical variable context, Γ, maps variables to types, while all variables in the program variable context, Δ, have the type Val.

$$\frac{(x : \tau) \in \Gamma}{\Gamma; \Delta \vdash x : \tau} \qquad \frac{(x : \mathsf{Val}) \in \Delta}{\Gamma; \Delta \vdash x : \mathsf{Val}} \qquad \frac{\Gamma, x : \tau; \Delta \vdash \mathsf{M} : \sigma}{\Gamma; \Delta \vdash \lambda x : \tau.\ \mathsf{M} : \tau \to \sigma} \qquad \frac{\Xi \vdash \mathsf{M} : \tau \to \sigma \qquad \Xi \vdash \mathsf{N} : \tau}{\Xi \vdash \mathsf{M}\ \mathsf{N} : \sigma}$$

$$\frac{p \in \{\bot, \top, \mathsf{emp}\}}{\Xi \vdash p : \mathsf{Prop}} \qquad \frac{\Xi \vdash \mathsf{P} : \mathsf{Prop} \qquad \Xi \vdash \mathsf{Q} : \mathsf{Prop} \qquad op \in \{\vee, \wedge, \Rightarrow, *, \twoheadrightarrow\}}{\Xi \vdash \mathsf{P}\ op\ \mathsf{Q} : \mathsf{Prop}}$$

$$\frac{\Xi \vdash \mathsf{M} : \tau \qquad \Xi \vdash \mathsf{N} : \tau}{\Xi \vdash \mathsf{M} =_\tau \mathsf{N} : \mathsf{Prop}} \qquad \frac{\Gamma, x : \tau; \Delta \vdash \mathsf{P} : \mathsf{Prop} \qquad Q \in \{\exists, \forall\}}{\Gamma; \Delta \vdash Qx : \tau.\ \mathsf{P} : \mathsf{Prop}} \qquad \frac{p \in \{\bot, \top\}}{\Gamma; - \vdash p : \mathsf{Spec}}$$

$$\frac{\Gamma; - \vdash \mathsf{S} : \mathsf{Spec} \qquad \Gamma; - \vdash \mathsf{T} : \mathsf{Spec} \qquad op \in \{\vee, \wedge, \Rightarrow\}}{\Gamma; - \vdash \mathsf{S}\ op\ \mathsf{T} : \mathsf{Spec}} \qquad \frac{\Gamma; - \vdash \mathsf{M} : \tau \qquad \Gamma; - \vdash \mathsf{N} : \tau}{\Gamma; - \vdash \mathsf{M} =_\tau \mathsf{N} : \mathsf{Spec}}$$

$$\frac{\Gamma, x : \tau; - \vdash \mathsf{S} : \mathsf{Spec} \qquad Q \in \{\forall, \exists\}}{\Gamma; - \vdash Qx : \tau.\ \mathsf{S} : \mathsf{Spec}}$$

$$\frac{\Xi \vdash \mathsf{M} : \mathcal{P}(\mathsf{SId}) \qquad \Xi \vdash \mathsf{I} : \mathsf{SId} \to \mathsf{Prop} \qquad \Xi \vdash \mathsf{T} : \mathsf{AId} \to \mathcal{P}(\mathsf{SId} \times \mathsf{SId}) \qquad \Xi \vdash \mathsf{R} : \mathsf{RId}}{\Xi \vdash \mathsf{region}(\mathsf{M}, \mathsf{I}, \mathsf{T}, \mathsf{R}) : \mathsf{Prop}}$$

$$\frac{\Xi \vdash \mathsf{A} : \mathsf{AId} \qquad \Xi \vdash \mathsf{R} : \mathsf{RId} \qquad \Xi \vdash \mathsf{P} : \mathsf{Perm}}{\Xi \vdash [\mathsf{A}]^\mathsf{R}_\mathsf{P} : \mathsf{Prop}} \qquad \frac{x \in X \qquad X \in obj(Sets)}{\Xi \vdash \Delta_X(x) : \Delta(X)}$$

$$\frac{\Xi \vdash \mathsf{M} : \mathsf{Val} \qquad \Xi \vdash \mathsf{C} : \mathsf{Class}}{\Xi \vdash \mathsf{M} : \mathsf{C} : \mathsf{Prop}} \qquad \frac{\Xi \vdash \mathsf{M} : \mathsf{Val} \qquad \Xi \vdash \mathsf{F} : \mathsf{Field} \qquad \Xi \vdash \mathsf{N} : \mathsf{Val}}{\Xi \vdash \mathsf{M}.\mathsf{F} \mapsto \mathsf{N} : \mathsf{Prop}}$$

$$\frac{\Gamma; - \vdash \mathsf{S} : \mathsf{Spec}}{\Gamma; \Delta \vdash \mathsf{spec}(\mathsf{S}) : \mathsf{Prop}} \qquad \frac{\Gamma; - \vdash \mathsf{P} : \mathsf{Prop}}{\Gamma; - \vdash \mathsf{valid}(\mathsf{P}) : \mathsf{Spec}} \qquad \frac{\Xi \vdash \mathsf{P} : \mathsf{Prop}}{\Xi \vdash \triangleright\mathsf{P} : \mathsf{Prop}} \qquad \frac{\Gamma; - \vdash \mathsf{P} : \mathsf{Prop}}{\Gamma; - \vdash \mathsf{stable}(\mathsf{P}) : \mathsf{Spec}}$$

$$\frac{\Gamma; - \vdash \mathsf{R} : \mathcal{P}(\mathsf{RId}) \qquad \Gamma; - \vdash \mathsf{P} : \mathsf{Prop} \qquad \Gamma; - \vdash \mathsf{Q} : \mathsf{Prop}}{\Gamma; - \vdash \mathsf{P} \sqsubseteq^\mathsf{R} \mathsf{Q} : \mathsf{Spec}} \qquad \frac{\Gamma; - \vdash \mathsf{S} : \mathsf{Spec}}{\Gamma; - \vdash \triangleright\mathsf{S} : \mathsf{Spec}}$$

$$\frac{\Gamma; \Delta \vdash \mathsf{P} : \mathsf{Prop} \qquad \Gamma; \Delta \vdash \mathsf{Q} : \mathsf{Prop} \qquad FV(\mathsf{s}) \subseteq \mathsf{vars}(\Delta)}{\Gamma; - \vdash (\Delta).\{\mathsf{P}\}\mathsf{s}\{\mathsf{Q}\} : \mathsf{Spec}}$$

$$\frac{\Gamma; \Delta \vdash \mathsf{P} : \mathsf{Prop} \qquad \Gamma; \Delta \vdash \mathsf{Q} : \mathsf{Prop} \qquad \Gamma; - \vdash \mathsf{R} : \mathcal{P}(\mathsf{RId}) \qquad FV(\mathsf{s}) \subseteq \mathsf{vars}(\Delta)}{\Gamma; - \vdash (\Delta).\langle\mathsf{P}\rangle\mathsf{s}\langle\mathsf{Q}\rangle^\mathsf{R} : \mathsf{Spec}}$$

C.2. Logics

The iCAP proof system consists of two logics: an assertion logic for reasoning about the current state and a specification logic for reasoning about the behavior of programs. The specification logic is given by the specification entailment judgment $\Gamma \mid \Phi \vdash S$ where S is a specification and Φ is a specification context. The assertion logic is given by the assertion entailment judgment $\Gamma; \Delta \mid \Phi \mid P \vdash Q$ where P and Q are assertions and Φ is a specification context. The assertion entailment includes the specification context Φ, to allow the use of assertion assumptions embedded in specifications.

C.2.1. Specification and assertion embeddings. The valid embedding is used to export axioms about abstract resources in library specifications to clients. For instance, the axioms about the channel order (duplication and transitivity) are implicitly expressed as validities about the channel order resource. The introduction and elimination rule for valid specifications is given below.

$$\frac{\Gamma; - \mid \Phi \mid \top \vdash P}{\Gamma \mid \Phi \vdash \mathsf{valid}(P)} \qquad\qquad \frac{\Gamma \mid \Phi \vdash \mathsf{valid}(P)}{\Gamma; - \mid \Phi \mid \top \vdash P}$$

C.2.2. Later operator. The SLOB rule internalizes induction on step-indicies in the logic and is implicitly used when verifying mutually recursive methods. In particular, to verify a method call, it suffices to know that the body of the called method satisfies a given specification one step later, to know that the call satisfies the given specification now. Intuitively, calling a method uses one step in the operational semantics, before the method body starts executing.

SLOB
$$\frac{\Gamma \mid \Phi, \triangleright S \vdash S}{\Gamma \mid \Phi \vdash S}$$

SMONO
$$\frac{}{\Gamma \mid \Phi \vdash S \Rightarrow \triangleright S}$$

LPOINTS
$$\frac{\Gamma \mid \Phi \vdash \langle x.f \mapsto v \rangle \; c \; \langle Q \rangle^{\mathcal{E}}}{\Gamma \mid \Phi \vdash \langle \triangleright(x.f \mapsto v) \rangle \; c \; \langle Q \rangle^{\mathcal{E}}}$$

LBIN
$$\frac{op \in \{\wedge, \vee, *\}}{\Gamma; \Delta \mid \Phi \mid \triangleright(P \; op \; Q) \dashv\vdash (\triangleright P) \; op \; (\triangleright Q)}$$

LIMPL
$$\frac{}{\Gamma; \Delta \mid \Phi \mid \triangleright(P \Rightarrow Q) \vdash (\triangleright P) \Rightarrow (\triangleright Q)}$$

LWAND
$$\frac{}{\Gamma; \Delta \mid \Phi \mid \triangleright(P \mathbin{-\!\!*} Q) \vdash (\triangleright P) \mathbin{-\!\!*} (\triangleright Q)}$$

LCEIL
$$\frac{}{\Gamma; \Delta \mid \Phi \mid \triangleright\lceil P \rceil \dashv\vdash \lceil \triangleright P \rceil}$$

LFLOOR
$$\frac{}{\Gamma; \Delta \mid \Phi \mid \triangleright\lfloor P \rfloor \dashv\vdash \lfloor \triangleright P \rfloor}$$

LQUANT
$$\frac{Q \in \{\forall, \exists\}}{\Gamma; \Delta \mid \Phi \mid \triangleright(Qx : \tau. \; P(x)) \dashv\vdash Qx : \tau. \; \triangleright P(x)}$$

$$\frac{}{\Gamma; \Delta \mid \Phi \mid \mathsf{region}(X, T, I_1, r) * \mathsf{region}(Y, T, I_2, r) \vdash I_1(s) \Rightarrow I_2(s)}$$

C.2.3. View-shifts. The view-shift relation includes standard assertion implication. In addition it is transitive and supports framing of stable frames. There is no implicit assumption

that the pre- and postcondition of view-shifts is stable.

$$\text{VTRANS}$$
$$\dfrac{\Gamma \mid \Phi \vdash \mathsf{P} \sqsubseteq^{\mathcal{E}} \mathsf{Q} \qquad \Gamma \mid \Phi \vdash \mathsf{Q} \sqsubseteq^{\mathcal{E}} \mathsf{R}}{\Gamma \mid \Phi \vdash \mathsf{Q} \sqsubseteq^{\mathcal{E}} \mathsf{R}}$$

$$\text{VIMPL}$$
$$\dfrac{\Gamma \mid \Phi \mid \mathsf{P} \vdash \mathsf{Q}}{\Gamma \mid \Phi \vdash \mathsf{P} \sqsubseteq^{\mathcal{E}} \mathsf{Q}}$$

$$\text{VALLOC}$$
$$\dfrac{\begin{array}{c}\Gamma \mid \Phi \vdash \Gamma \mid \Phi \vdash \forall \alpha \in \mathsf{AId}.\ \forall x \in \mathsf{SId} \times \mathsf{SId}.\ (\triangleright T(\alpha)(x)) \Rightarrow T(\alpha)(x) \vee \triangleright \bot \\ \Gamma \mid \Phi \vdash \mathcal{E} \text{ is infinite} \qquad \Gamma \mid \Phi \vdash \forall n \in \mathcal{E}.\ \mathsf{P} * \circledast_{\alpha \in A}[\alpha]_1^n \Rightarrow \triangleright I(n)(x) \\ \Gamma \mid \Phi \vdash \forall n \in \mathcal{E}.\ \forall s.\ \mathsf{stable}(I(n)(s)) \qquad \Gamma \mid \Phi \vdash A \cap B = \emptyset \end{array}}{\Gamma \mid \Phi \vdash \mathsf{P} \sqsubseteq^{\mathcal{E}} \exists n \in \mathcal{E}.\ \mathsf{region}(\{x\}, T, I(n), n) * \circledast_{\alpha \in B}[\alpha]_1^n}$$

$$\text{VOPEN}$$
$$\dfrac{\begin{array}{c}\Gamma \mid \Phi \vdash \mathsf{stable}(\mathsf{P}) \qquad \Gamma \mid \Phi \vdash \mathsf{stable}(\mathsf{Q}) \\ \Gamma \mid \Phi \vdash \forall x \in X.\ f(x) \in Y \\ \Gamma \mid \Phi \vdash \forall x \in X.\ (x, f(x)) \in T(\alpha) \vee f(x) = x \\ \Gamma \mid \Phi \vdash \forall x \in X.\ \mathsf{P} * \triangleright I(x) * [\alpha]_\pi^n \sqsubseteq^{\mathcal{E}} \mathsf{Q} * \triangleright I(f(x)) \end{array}}{\Gamma \mid \Phi \vdash \mathsf{region}(X, T, I, n) * \mathsf{P} * [\alpha]_\pi^n \sqsubseteq^{\mathcal{E} \uplus \{n\}} \mathsf{region}(Y, T, I, n) * \mathsf{Q}}$$

The VALLOC rule presented above generalizes the region allocation rule presented in 4, by allowing the newly allocated region to immediately take ownership of action permissions on the newly allocated region ($\circledast_{\alpha \in A}[\alpha]_1^n$). In addition, it allows the allocator to pick an infinite set $\mathcal{E}$ of region identifiers, from which the region identifier of the newly allocated region will be chosen. This is used to reason about inequality of region idenfiers, which is necessary when reasoning about nested region openings.

C.2.4. Atomic commands. All pre- and postconditions that appear in non-atomic Hoare triples are implicitly required to be stable.

$$\dfrac{\begin{array}{c}\Gamma, \Delta \vdash \mathsf{P}, \mathsf{Q} : \mathsf{Prop} \qquad \mathsf{atomic}(\mathsf{s}) \\ \Gamma, \Delta \mid \Phi \vdash \mathsf{stable}(\mathsf{P}) \qquad \Gamma, \Delta \mid \Phi \vdash \mathsf{stable}(\mathsf{Q}) \\ \Gamma \mid \Phi \vdash (\Delta).\langle \mathsf{P} \rangle \mathsf{s} \langle \mathsf{Q} \rangle^{\mathcal{E}} \end{array}}{\Gamma \mid \Phi \vdash (\Delta).\{\mathsf{P}\}\mathsf{s}\{\mathsf{Q}\}}$$

$$\dfrac{\begin{array}{c}\Gamma, \Delta \vdash \mathsf{P}, \mathsf{Q} : \mathsf{Prop} \qquad \Gamma, \Delta \vdash \mathcal{E}_1, \mathcal{E}_2 : \mathcal{P}(\mathsf{RId}) \\ \Gamma \mid \Phi \vdash (\Delta).\langle \mathsf{P} \rangle \mathsf{s} \langle \mathsf{Q} \rangle^{\mathcal{E}_1 \backslash \mathcal{E}_2} \end{array}}{\Gamma \mid \Phi \vdash (\Delta).\langle \mathsf{P} \rangle \mathsf{s} \langle \mathsf{Q} \rangle^{\mathcal{E}_1}}$$

$$\text{ATOMIC}$$
$$\dfrac{\begin{array}{c}\Gamma, \Delta \mid \Phi \vdash \mathsf{stable}(\mathsf{P}) \qquad \Gamma, \Delta \mid \Phi \vdash \forall y.\ \mathsf{stable}(\mathsf{Q}(y)) \\ \Gamma, \Delta \mid \Phi \vdash \forall x \in X.\ (x, f(x)) \in \overline{T(A)} \vee f(x) = x \\ \Gamma \mid \Phi \vdash \forall x \in X.\ (\Delta).\langle \mathsf{P} * \circledast_{\alpha \in A}[\alpha]_{g(\alpha)}^n * \triangleright I(x) \rangle\ c\ \langle \mathsf{Q}(x) * \triangleright I(f(x)) \rangle^{\mathcal{E}} \end{array}}{\Gamma \mid \Phi \vdash (\Delta).\ \langle \mathsf{P} * \circledast_{\alpha \in A}[\alpha]_{g(\alpha)}^n * \mathsf{region}(X, T, I, n) \rangle}$$

$$c$$

$$\langle \exists x.\ \mathsf{Q}(x) * \mathsf{region}(\{f(x)\}, T, I, n) \rangle^{\mathcal{E} \uplus \{n\}}$$

C.2.5. Stability.

$$\frac{\Gamma \mid \Phi \vdash \forall \alpha \notin A.\ \forall x \in X.\ T(\alpha)(x) \subseteq X}{\Gamma \mid \Phi \vdash \mathsf{stable}(\mathsf{region}(X,T,I,n) * \circledast_{\alpha \in A}[\alpha]_1^n)} \qquad \overline{\Gamma \mid \Phi \vdash \mathsf{stable}(\top)} \qquad \overline{\Gamma \mid \Phi \vdash \mathsf{stable}(\bot)}$$

$$\overline{\Gamma \mid \Phi \vdash \mathsf{stable}(\mathsf{emp})} \qquad \overline{\Gamma \mid \Phi \vdash \mathsf{stable}(\mathsf{M.F} \mapsto \mathsf{N})} \qquad \overline{\Gamma \mid \Phi \vdash \mathsf{stable}(\mathsf{M : N})}$$

$$\frac{\Gamma \mid \Phi \vdash \mathsf{stable}(\mathsf{P}) \qquad \Gamma \mid \Phi \vdash \mathsf{stable}(\mathsf{Q}) \qquad op \in \{\vee, \wedge, *\}}{\Gamma \mid \Phi \vdash \mathsf{stable}(\mathsf{P}\ op\ \mathsf{Q})} \qquad \overline{\Gamma \mid \Phi \vdash \mathsf{stable}(\mathsf{M} =_\tau \mathsf{N})}$$

$$\frac{\Gamma \mid \Phi \vdash \forall x : \tau.\ \mathsf{stable}(\mathsf{P}(x)) \qquad Q \in \{\forall, \exists\}}{\Gamma \mid \Phi \vdash \mathsf{stable}(Qx : \tau.\ \mathsf{P}(x))} \qquad \frac{\Gamma \vdash \mathsf{S} : \mathsf{Spec}}{\Gamma \mid \Phi \vdash \mathsf{stable}(\mathsf{spec}(\mathsf{S}))} \qquad \frac{\Gamma \mid \Phi \vdash \mathsf{stable}(\mathsf{P})}{\Gamma \mid \Phi \vdash \mathsf{stable}(\triangleright\mathsf{P})}$$

C.2.6. Structural rules.

VFRAME
$$\frac{\Gamma \mid \Phi \vdash \mathsf{P} \sqsubseteq^{\mathcal{E}} \mathsf{Q} \qquad \Gamma \mid \Phi \mid \mathsf{stable}(\mathsf{R})}{\Gamma \mid \Phi \vdash \mathsf{P} * \mathsf{R} \sqsubseteq^{\mathcal{E}} \mathsf{Q} * \mathsf{R}}$$

AFRAME
$$\frac{\Gamma \mid \Phi \vdash (\Delta).\langle \mathsf{P}\rangle\ c\ \langle \mathsf{Q}\rangle^{\mathcal{E}} \qquad \Gamma, \Delta \mid \Phi \vdash \mathsf{stable}(\mathsf{R})}{\Gamma \mid \Phi \vdash (\Delta).\langle \mathsf{P} * \triangleright\mathsf{R}\rangle\ c\ \langle \mathsf{Q} * \mathsf{R}\rangle^{\mathcal{E}}}$$

FRAME
$$\frac{\Gamma \mid \Phi \vdash (\Delta).\{\mathsf{P}\}\ c\ \{\mathsf{Q}\} \qquad \Gamma, \Delta \mid \Phi \vdash \mathsf{stable}(\mathsf{R})}{\Gamma \mid \Phi \vdash (\Delta).\{\mathsf{P} * \mathsf{R}\}\ c\ \{\mathsf{Q} * \mathsf{R}\}}$$

CONSEQ
$$\frac{\Gamma, \Delta \mid \Phi \vdash \mathsf{P}_1 \sqsubseteq^{\mathcal{E}} \mathsf{P}_2 \qquad \Gamma \mid \Phi \vdash (\Delta).\{\mathsf{P}_2\}\ c\ \{\mathsf{Q}_2\} \qquad \Gamma, \Delta \mid \Phi \vdash \mathsf{Q}_2 \sqsubseteq^{\mathcal{E}} \mathsf{Q}_1}{\Gamma \mid \Phi \vdash (\Delta).\{\mathsf{P}_1\}\ c\ \{\mathsf{Q}_1\}}$$

ACONSEQ
$$\frac{\Gamma, \Delta \mid \Phi \vdash \mathsf{P}_1 \sqsubseteq^{\mathcal{E}} \mathsf{P}_2 \qquad \Gamma \mid \Phi \vdash (\Delta).\langle \mathsf{P}_2\rangle\ c\ \langle \mathsf{Q}_2\rangle^{\mathcal{E}} \qquad \Gamma, \Delta \mid \Phi \vdash \mathsf{Q}_2 \sqsubseteq^{\mathcal{E}} \mathsf{Q}_1}{\Gamma \mid \Phi \vdash (\Delta).\langle \mathsf{P}_1\rangle\ c\ \langle \mathsf{Q}_1\rangle^{\mathcal{E}}}$$

SEQ
$$\frac{\Gamma \mid \Phi \vdash (\Delta).\{\mathsf{P}\}\mathsf{s}_1\{\mathsf{Q}\} \qquad \Gamma \mid \Phi \vdash (\Delta).\{\mathsf{Q}\}\mathsf{s}_2\{\mathsf{R}\}}{\Gamma \mid \Phi \vdash (\Delta).\{\mathsf{P}\}\mathsf{s}_1; \mathsf{s}_2\{\mathsf{R}\}}$$