

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

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 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.

This work is supported TODO.

Author's addresses: TODO

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 0000-0000/YYYY/01-ARTA \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

For standard synchronisation constructs – e.g. locks, queues, stacks – abstract specifications are well-understood. 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 enforcing the original sequential program dependencies on shared resources. A resource could be any program variable, data structure, memory region, lock, etc. for which 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 supports 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.
- 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 *explicit stabilization* [Wickerson et al. 2010] to reason about the stability of complex separation logic assertions (see §3.5).

This paper also corrects a subtle logical problem in [Dodds et al. 2011], discovered by Svendsen a year after publication. As is often true in logic, this problem arose as a result of self-reference – in this case, a circularity in the model of higher-order propositions 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. §4 gives a very simple barrier implementation, and shows how it can be verified with respect to the core of the specification. §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.

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. We share two co-authors with [Svendsen et al. 2013], and iCAP was largely developed as a improvement on the HOCAP logic it uses.

The most 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.

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 specifications 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 ensure 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) {
    if(*x < 10) {
        *y = *y + v;
        *x = *x + v;
        sleep(rand());
    } else {
        sleep(rand());
    }
}
```

```
*x = 0;
*y = 0;

f(x,y,5);
f(x,y,11);
```

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:

```
f1(x,y,v,i) {
    if(*x < 10) {
        *y = *y + v;
        *x = *x + v;
        signal(i);
        sleep(rand());
    } else {
        signal(i);
        sleep(rand());
    }
}

f2(x,y,v,i) {
    wait(i);
    if(*x < 10) {
        *y = *y + v;
        *x = *x + v;
        sleep(rand());
    } else {
        sleep(rand());
    }
}
```

PROGRAM BODY:

```
*x = 0; *y = 0;
chan *i = newchan();

f1(x,y,5,i) || f2(x,y,11,i);
```

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'*: $\mathbf{x} \mapsto v * \mathbf{y} \mapsto v'$. 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 $\mathbf{x} \mapsto v * \mathbf{y} \mapsto v'$ 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]:

$$\frac{\{P_1\} C_1 \{Q_1\} \quad \{P_2\} C_2 \{Q_2\}}{\{P_1 * P_2\} C_1 \parallel C_2 \{Q_1 * Q_2\}} \text{ PAR}$$

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:

- recv(i, P)** – By calling **wait** on **i**, the thread will acquire a resource satisfying the assertion *P*.
- send(i, P)** – By calling **signal** on **i** when holding a resource satisfying *P*, the thread will lose the resource *P*.

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

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

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:

{ emp }	i = newchan()	{ send(i, P) * recv(i, P) }
{ send(i, P) * [P] }	signal(i)	{ emp }
{ recv(i, P) }	wait(i)	{ [P] }

The specification of `newchan` is implicitly universally quantified for all assertions P , meaning that we can construct a channel for any assertion.¹ 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).

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

{ x ↦ 0 * y ↦ 0 * send(i, x ↦ 5 * y ↦ 5) }	f1(x, y, 5, i)	{ emp }
{ recv(i, x ↦ 5 * y ↦ 5) }	f2(x, y, 11, i)	{ x ↦ 16 * y ↦ 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:

¹In iCAP, assertions can be shared between multiple threads. In this case we need to establish that each assertion is *stable*, i.e. invariant under changes performed by other threads. The *explicit stabilisation* operators $\lfloor - \rfloor$ and $\lceil - \rceil$ in the specification are needed because P might contain assertions about shared state. These are discussed further at the end of the section, in §3.5. For the moment, note that if P is a thread-local assertion such as $x \mapsto v$, then $\lfloor P \rfloor \iff P \iff \lceil P \rceil$.

$$\begin{array}{c}
\{x \mapsto _ * y \mapsto _ \} \\
*x = 0; *y = 0; \text{chan } *i = \text{newchan}(); \\
\left\{ \begin{array}{l}
\{x \mapsto 0 * y \mapsto 0 * \text{send}(i, x \mapsto 5 * y \mapsto 5) * \text{recv}(i, x \mapsto 5 * y \mapsto 5)\} \\
\{x \mapsto 0 * y \mapsto 0 * \text{send}(i, x \mapsto 5 * y \mapsto 5)\} \parallel \left\{ \begin{array}{l} \text{recv}(i, x \mapsto 5 * y \mapsto 5) \end{array} \right\} \\
\text{f1}(x, y, 5, i) \qquad \qquad \qquad \text{f2}(x, y, 11, i) \\
\{ \text{emp} \} \qquad \qquad \qquad \{x \mapsto 16 * y \mapsto 16\}
\end{array} \right\} \text{PAR rule application.} \\
\{x \mapsto 16 * y \mapsto 16\}
\end{array}$$

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:

$$\begin{array}{l}
*x = *y + 1; \quad \parallel \quad \text{wait}(i); \quad \parallel \quad \text{wait}(i); \\
\text{signal}(i) \quad \parallel \quad z := *x \quad \parallel \quad *y := 4
\end{array}$$

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:

$$\{\text{recv}(a, P) * \llbracket P \rrbracket \multimap (P_1 * P_2)\} \quad \langle \text{skip} \rangle \quad \{\text{recv}(a, P_1) * \text{recv}(a, P_2)\}$$

This axiom states that when a thread has been promised a resource P , access can be split between two threads, potentially before the resource is available. The specification uses separation logic's separating implication operator \multimap (separating implication and separating conjunction have a similar relationship shared by classical implication, \Rightarrow , and classical conjunction, \wedge , i.e. they are adjoint). The assertion $P \multimap (P_1 * P_2)$ asserts ownership of a resource that when combined with P can be split into P_1 and P_2 . As the resource P is promised by the recv predicate, the result of applying this property is two new recv predicates, one of which promises P_1 , and the other P_2 .

Note that this property is *not* a logical entailment – applying it requires an operational step, **skip**. This is because the property manipulates a shared higher-order resource: $\text{recv}(a, P)$. 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. To use this property, we have to assume that every application in a proof 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, if we take the second branch of the conditional in **f**, we do not need the heap location **y**. It is safe to notify future threads that **y** is available, conditional on all logically prior threads releasing it, even though the thread itself never acquired access to the resource. Without renunciation the thread would have to **wait** for all the logically earlier threads to **signal** before it could proceed.

Chains. To support chains, 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:

$$\begin{aligned} x \prec y &\implies x \prec y * x \prec y && (\text{duplication}) \\ x \prec y * y \prec z &\implies x \prec z && (\text{transitivity}) \end{aligned}$$

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 (\mathbf{a}, \mathbf{b}) , that are later than all the channels before \mathbf{x} and before all the channels after \mathbf{x} , and \mathbf{a} is before \mathbf{b} in the chain. It also creates **recv**, **send**, and order predicates representing the new channel.

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

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

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

Using this specification, we can partially discharge a **send** using the preceding **recv** predicate. The assertion $P * Q$ stands for the resource which gives Q if combined with a resource satisfying P . Thus, combining it with the predicate $\text{recv}(x, P)$ provides Q . In other words, the thread gives up the ability to ever acquire the resource, and instead forwards this ability to logically later threads. When the resource becomes available from prior threads, the *next* thread in the order will receive it (unless it, too, renounces the resource).

3.4. Full abstract specification

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

Example parallelisation. Suppose we want to run many copies of the function **f** in sequence, for example over an array of values **vs**.

```
for(j = 0; j < max; j++){
    f(x,y,vs[j]);
}
```

To parallelize this program, we define **fp** (left of Fig. 3) a version of **f** which is safe to run in parallel with many copies of itself. To do this, **f** is modified to call both **signal** and **wait**. Each call to **f** receives the resource from logically earlier threads (those invoked in earlier loop iterations) with **wait**, then releases it to logically later threads (those invoked in later loop iterations) using **signal**.

In the original transformation involving **f1** and **f2**, we did not distinguish between the resources **x** and **y**. However, we need to gain access to **y** only if we take the first branch

SPECIFICATIONS:

$$\begin{array}{lll}
\{\text{emp}\} & \mathbf{i} = \text{newchan}() & \{\text{recv}(\mathbf{i}, P) * \text{send}(\mathbf{i}, P)\} \\
\{\text{send}(\mathbf{i}, P) * \lfloor P \rfloor\} & \text{signal}(\mathbf{i}) & \{\text{emp}\} \\
\{\text{recv}(\mathbf{i}, P)\} & \text{wait}(\mathbf{i}) & \{\lfloor P \rfloor\}
\end{array}$$

$$\left\{ \text{send}(\mathbf{x}, P) * \bigotimes_{e \in E} e \prec \mathbf{x} * \bigotimes_{l \in L} \mathbf{x} \prec l \right\} (\mathbf{a}, \mathbf{b}) = \text{extend}(\mathbf{x}) \left\{ \text{send}(\mathbf{a}, Q) * \text{recv}(\mathbf{a}, Q) * \text{send}(\mathbf{b}, P) \right\}$$

$$\left\{ * \mathbf{a} \prec \mathbf{b} * \bigotimes_{e \in E} e \prec \mathbf{a} * \bigotimes_{l \in L} \mathbf{b} \prec l \right\}$$

AXIOMS:

$$\begin{array}{lll}
x \prec y & \implies & x \prec y * x \prec y \\
x \prec y * y \prec z & \implies & x \prec z \\
\{\text{recv}(x, P) * \text{send}(y, Q) * x \prec y\} & \langle \text{skip} \rangle & \{\text{send}(y, P * Q)\} \\
\{\text{recv}(a, P) * \lfloor P \rfloor * (P_1 * P_2)\} & \langle \text{skip} \rangle & \{\text{recv}(a, P_1) * \text{recv}(a, P_2)\}
\end{array}$$

Fig. 2. Full abstract specification for deterministic parallelism.

FUNCTION DEFINITION:

```

fp(x,y,v,ix,iy,ixp,iyp) {
  wait(ixp);
  if (*x < 10) {
    wait(iyp);
    *y = *y + v;
    signal(iy);
    *x = *x + v;
    signal(ix);
  } else {
    signal(ix);
    signal(iy);
    sleep(rand());
  }
}

```

PROGRAM BODY:

```

ixf = newchan();
(ixn,ixf) = extend(ixf);
signal(ixn);
iyf = newchan();
(iyn,iyf) = extend(iyf);
signal(iyn);
for(j = 0; j < max; j++){
  v = vs[j];
  ixl = ixn;
  (ixn,ixf) = extend(ixf);
  iyl = iyn;
  (iyn,iyf) = extend(iyf);
  fork( fp(x,y,v,ixn,iyn,ixl,iyl) );
}
wait(ixn);
wait(iyn);

```

Fig. 3. Example parallelization of **f** and a client.

```

1  {recv(ixp, x ↦ -) * send(ix, x ↦ -) * recv(iyp, y ↦ -) * send(iy, y ↦ -) * iyp < iy}
2  fp(x, y, v, ix, iy, ixp, iyp) {
3    wait(ixp);
4    {x ↦ - * send(ix, x ↦ -) * recv(iyp, y ↦ -) * send(iy, y ↦ -) * iyp < iy}
5    if (*x < 10) {
6      wait(iyp);
7      {x ↦ - * y ↦ - * send(ix, x ↦ -) * send(iy, y ↦ -)}
8      *y = *y + v;
9      signal(iy);
10     {x ↦ - * send(ix, x ↦ -)}
11     *x = *x + v;
12     signal(ix);
13   } else {
14     signal(ix);
15     {recv(iyp, y ↦ -) * send(iy, y ↦ -) * iyp < iy}
16     signal(iy);
17     {emp}
18     sleep(rand());
19   }
20 }
21 {emp}

```

Fig. 4. Proof for parallelized program `fp`.

of the conditional. Otherwise we can release `y` to logically future threads. To realise this parallelism in the new version of `f`, we use two chains of channels: one for `x`, and one for `y`. The function takes arguments `ix` and `iy` representing the next points in the two chains, and `ixp` and `iyp` representing the previous points.

In Fig. 4, we verify `fp` against the following specification:

$$\left\{ \begin{array}{l} \text{recv}(\text{ixp}, x \mapsto -) * \text{send}(ix, x \mapsto -) * \\ \text{recv}(iyp, y \mapsto -) * \text{send}(iy, y \mapsto -) * iyp < iy \end{array} \right\} \text{fp}(\dots) \quad \{\text{emp}\}$$

Note that we only prove basic memory safety, but we could easily verify stronger invariants.

Line 16 in Fig. 4 is noteworthy. There, the precondition *does not* assert that the thread has access to `y ↦ -`; rather, it asserts it can acquire access by calling `wait`. Instead of doing this, the thread renounces access to the resource without ever holding it.

The parallelized version of the client program is given on the right of Fig. 3; a proof of correctness is given in Fig. 5. The predicates `send(ixf, true)` and `send(iyf, true)` represent the logically last element of the chain, while `array()` stands for the array holding processed values.

Thus, using our abstract specification, we have shown that our parallelized version of the program is memory-safe. With a little more effort, we could verify the behaviour of the program. 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. While the example

```

1  {  $x \mapsto \_ * y \mapsto \_ * \text{array}(\text{vs}, \text{max})$  }
2  ixf = newchan();
3  (ixn,ixf) = extend(ixf);
4  signal(ixn);
5  {  $y \mapsto \_ * \text{array}(\text{vs}, \text{max}) * \text{recv}(\text{ixn}, x \mapsto \_) * \text{send}(\text{ixf}, \text{true}) * \text{ixn} \prec \text{ixf}$  }
6  iyf = newchan();
7  (iyn,iyf) = extend(iyf);
8  signal(iyn);
9  {  $\text{array}(\text{vs}, \text{max}) * \text{recv}(\text{ixn}, x \mapsto \_) * \text{send}(\text{ixf}, \text{true}) * \text{ixn} \prec \text{ixf}$ 
   *  $\text{recv}(\text{iyn}, y \mapsto \_) * \text{send}(\text{iyf}, \text{true}) * \text{iyn} \prec \text{iyf}$  }
10 for(j = 0; j < max; j++){
11   ixl = ixn;
12   (ixn,ixf) = extend(ixf);
13   {  $\text{array}(\text{vs}, \text{max}) * \text{recv}(\text{ixl}, x \mapsto \_) * \text{send}(\text{ixn}, x \mapsto \_)$ 
   *  $\text{recv}(\text{ixn}, x \mapsto \_) * \text{send}(\text{ixf}, \text{true}) * \text{ixn} \prec \text{ixf}$ 
   *  $\text{recv}(\text{iyn}, y \mapsto \_) * \text{send}(\text{iyf}, \text{true}) * \text{iyn} \prec \text{iyf}$  }
14   iyl = iyn;
15   (iyn,iyf) = extend(iyf);
16   {  $\text{array}(\text{vs}, \text{max}) * \text{recv}(\text{ixl}, x \mapsto \_) * \text{send}(\text{ixn}, x \mapsto \_)$ 
   *  $\text{recv}(\text{ixn}, x \mapsto \_) * \text{send}(\text{ixf}, \text{true}) * \text{ixn} \prec \text{ixf}$ 
   *  $\text{recv}(\text{iyl}, y \mapsto \_) * \text{send}(\text{iyn}, y \mapsto \_) * \text{iyl} \prec \text{iyn}$ 
   *  $\text{recv}(\text{iyn}, y \mapsto \_) * \text{send}(\text{iyf}, \text{true}) * \text{iyn} \prec \text{iyf}$  }
17   fork( fp(x,y,vs[j],ixn,iyn,ixl,iyl) ); // apply function spec.
18 }
19 wait(ixn);
20 wait(iyn);
21 {  $\text{array}(\text{vs}, \text{max}) * x \mapsto \_ * y \mapsto \_$  }

```

Fig. 5. Proof of fp client program.

given here is trivial, in [Botinčan et al. 2013] we use the same mechanism as the basis for a general parallelisation analysis.

3.5. Explicit Stabilisation

Our abstract specification includes the assertion $[P]$ in the precondition to `signal()`, and $[P]$ in the postcondition to `wait()`, as well as similar assertions in the axiom for splitting. $[-]$ and $[-]$ are *explicit stabilisation* operators [Wickerson et al. 2010].

In iCAP assertions can refer to resources shared between threads. Any such assertion in the proof of one thread could potentially be invalidated by interference from other threads. Therefore, any assertion in iCAP must be *stable*, meaning invariant under the worst-case interference from other threads (we inherit this terminology from Jones’s rely-guarantee logic). In our abstract specification, higher-order arguments such as P could include assertions about shared state. Thus, P and similar assertions must be provably stable.

In our conference paper [Dodds et al. 2011] we used a predicate `stable(P)` to explicitly assert stability of P the specification. However, this predicate is hard to manipulate in proofs

```

struct chan {
    int flag;
}

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

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

wait(chan *x) {
    while(x->flag == 0)
        skip;
}

```

Fig. 6. Implementation of the barrier library.

because it satisfies few distribution properties. This is especially problematic for assertions containing separating implication, \ast . In this paper, we instead use explicit stabilization operators, which strengthen or weaken an assertion until it is stable:

- $\lfloor P \rfloor$ stands for the weakest assertion stronger than P that is stable.
- $\lceil P \rceil$ stands for the strongest assertion weaker than P that is stable.

Two properties of explicit stabilization operators make them easy to remove if not needed:

- If P is already stable, then $\lfloor P \rfloor \iff P \iff \lceil P \rceil$. This holds by default if P contains no shared assertions nor predicate variables.
- $\lfloor P \rfloor \implies P \implies \lceil P \rceil$.

Furthermore, explicit stabilisation operators are semi-distributive over separating conjunction. As a result, they are easy to move around in proofs.

- $\lfloor P \rfloor \ast \lfloor Q \rfloor \implies \lfloor P \ast Q \rfloor$.
- $\lceil P \ast Q \rceil \implies \lceil P \rceil \ast \lceil Q \rceil$.

The predicate **stable** is much harder to use in proofs, because the only one of these properties that holds is the first one, elimination if the assertion contains no shared assertions or predicate variables. The properties are especially useful when combined with separating implication. We use \ast extensively, but are aware of no other approach that would allow reasoning about the stability of complex assertions that include separating implication. The following Lemma provides the common pattern of our usage:

LEMMA 3.1. *If P stable and $P \ast Q \Rightarrow R$, then $P \Rightarrow \lfloor Q \ast R \rfloor$.*

4. PROOF STRATEGY FOR THE SIMPLIFIED SPECIFICATION

This section provides an intuitive introduction to our general proof approach, and to iCAP, the reasoning system our proofs are based on. To motivate this, we define a very simple channel implementation and verify it with respect to just the first three axioms of our abstract specification:

```

{emp} i = newchan() {recv(i, P) * send(i, P)}
{send(i, P) * ⌊P⌋}   signal(i)   {emp}
{recv(i, P)}         wait(i)     {⌈P⌉}

```

By avoiding splitting, chain extension, and renunciation, we can illustrate the basic features of iCAP in a straightforward manner. In §5 and §6 we reintroduce the necessary extra reasoning to verify our full abstract specification.

Figure 6 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.

4.1. Introduction to iCAP

Before we can present our verification, we need to sketch the key details of the logic we use: Impredicative Concurrent Abstract Predicates (iCAP). iCAP is a separation logic variant intended for verifying concurrent higher-order programs. Full details are given in [Svendsen and Birkedal 2014a].

To handle concurrency, iCAP extends separation logic with *regions* containing resources shared between threads. Resources stored in a region must be modified atomically in a way that satisfies the expectations of the other threads. An assertion about a region has the following form:

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

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 the shared region owns in the 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 owned by the given thread. Ownership of actions is tracked using tokens. For instance, the following assertion asserts ownership of a **set** and **change** token: $[\text{set}]_i^{r_1} * [\text{change}]_j^{r_2}$. 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.

Tokens are linear objects created at the same time as a region – by issuing threads with different tokens, we grant them different abilities over the shared region. 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. Tokens can be split arbitrarily: $[\alpha]_{i+j}^r \Leftrightarrow [\alpha]_i^r * [\alpha]_j^r$. In contrast, region assertions can be duplicated arbitrarily:

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

Thus, each region assertion should be interpreted as one view on a resource shared between many threads. Such a region assertion, $\text{region}(R, T, I, r)$ is *stable* (i.e., closed under possible inference from the environment) if the set R is closed under all transitions in T labelled with actions potentially owned by the environment.

As well as information about the underlying memory cells, the abstract state of region captures auxiliary state representing the way threads can interact. To manipulate this auxiliary state, iCAP uses the *view-shift* operator \sqsubseteq . View-shift generalises standard implication to allow creation, destruction and manipulation of regions.

As a convenience, we assume that we can garbage collect unwanted portions of assertions. As we mostly apply this to logical constructs, this assumption loses little generality, and in any case could easily be lifted at the cost of larger proofs.

Higher-order assertions and ‘later’. The features discussed above could equally be expressed in prior logics such as RGSep [Vafeiadis and Parkinson 2007]. 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.

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. To capture this stratification in the logic, iCAP introduces a ‘later’ modality, written \triangleright . Intuitively, $\triangleright P$ expresses that the assertion P holds after one step of execution. To ensure that protocols defined in iCAP are well-defined they are implicitly interpreted one step later. The region assertion $\text{region}(\{x\}, T, I, r)$ thus expresses that the shared region r currently owns the resources described by $\triangleright I(x)$.

If an assertion P holds now, then it also holds after one step of execution. This is expressed by the (SMono) rule given below. In general, if $\triangleright P$ holds now, it is not the case that P also holds now. Instead, \triangleright operators can be eliminated by taking an execution step, as expressed by the frame rule for atomic commands (AFrame). Using (SMono) and the rule of consequence, one can derive the standard CAP frame rule.

$$\begin{aligned} P &\Longrightarrow \triangleright P && \text{(SMONO)} \\ \text{atomic}(C) \wedge \text{stable}(R) \wedge \{P\} C \{Q\} &\Longrightarrow \{P * \triangleright R\} C \{Q * R\} && \text{(AFRAME)} \end{aligned}$$

The effect of this 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] 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.

Properties of ‘later’. The later operator commutes over conjunction, disjunction and separating conjunction and semi-commutes over implication and separating implication:

$$\begin{aligned} \triangleright(P \wedge Q) &\iff \triangleright P \wedge \triangleright Q, & \triangleright(P \vee Q) &\iff \triangleright P \vee \triangleright Q, & \triangleright(P * Q) &\iff \triangleright P * \triangleright Q && \text{(LBIN)} \\ \triangleright(P \Rightarrow Q) &\implies \triangleright P \Rightarrow \triangleright Q, & \triangleright(P \multimap Q) &\implies \triangleright P \multimap \triangleright Q && \text{(LIMPL/LWAND)} \end{aligned}$$

Later also commutes over existential and universal quantification over non-empty types τ :

$$\triangleright \exists x : \tau. P(x) \iff \exists x : \tau. \triangleright P(x), \quad \triangleright \forall x : \tau. P(x) \iff \forall x : \tau. \triangleright P(x) \quad \text{(LQUANT)}$$

In proof outlines we often apply these properties silently. Lastly, later can be moved freely in and out of stabilization brackets.

$$\triangleright[P] \iff [\triangleright P], \quad \triangleright[P] \iff [\triangleright P] \quad \text{(LCEIL/LFLOOR)}$$

Reasoning about shared regions. All statements that access resources owned by shared regions must be atomic and obey the protocols of the regions in question. To illustrate how this is expressed formally, we consider a simplified proof rule for reasoning about shared regions. We refer to the iCAP technical report [Svendsen and Birkedal 2014b] for the general proof rules.

We reason about shared regions using rules that allow us to treat the shared resources 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 in proof outlines. To illustrate, consider the verification of an atomic statement C that changes the abstract state of a

shared region r from x to y :

$$\{\text{region}(\{x\}, T, I, r) * [\alpha]_\pi^r * P\} C \{\text{region}(\{y\}, T, I, r) * Q\}$$

To prove that this triple holds, we must first prove that we own sufficient permissions to change the abstract state from x to y . In this case we own non-exclusive permission to the α action. We thus have to prove that there is an α -labelled path from x to y in the transition system T . Furthermore, we have to prove that C does indeed transform the resources associated with abstract state x to those associated with abstract state y , according to I . Since protocols are implicitly interpreted one step later, we have to prove the following triple:

$$\{\triangleright I(x) * [\alpha]_\pi^r * P\} C \{\triangleright I(y) * Q\}$$

Here P and Q can be used to transfer local resources in and out of the shared region as appropriate.

Imagine that $I = \lambda v. l \mapsto v$ and C is the atomic statement $l := y$. Then we would be left with the following proof obligation (after instantiating P and Q with emp).

$$\{\triangleright(l \mapsto x) * [\alpha]_\pi^r\} l := y \{\triangleright l \mapsto y\}$$

Conceptually, this is provable because primitive points-to assertions, such as $l \mapsto x$, are independent of the step-indexing. This is capture by the structural (LPOINTS) rule given below, which allows us to remove a later from the pre-condition of a points-to assertion.

$$\text{atomic}(C) \wedge \{x \mapsto y\} C \{Q\} \implies \{\triangleright x \mapsto y\} C \{Q\} \quad (\text{LPOINTS})$$

4.2. Verifying the Simple Implementation

Verifying the implementation shown in Figure 6 with respect to the abstract specification amounts to giving concrete definitions to the abstract predicates **send** and **recv**, and then using iCAP to prove that the resulting concrete specification is satisfied.

We begin by defining the structure of the shared region. There are three possible abstract states: either the flag is low, or the flag is high and the sent resource is available, or the flag is high and the resource has been taken (we call this ‘done’). We define three corresponding states: $\{\text{Low}, \text{High}, \text{Done}\}$. Each abstract state is associated with an invariant by I_b , defined as follows:

$$\begin{aligned} I_b(x, P)(\text{Low}) &\triangleq x.\text{flag} \mapsto 0 \\ I_b(x, P)(\text{High}) &\triangleq x.\text{flag} \mapsto 1 * \lceil P \rceil \\ I_b(x, P)(\text{Done}) &\triangleq x.\text{flag} \mapsto 1 \end{aligned}$$

Here x is the location of the channel, while P is the resource controlled by the channel. Note that moving from **Low** to **High** requires the resource $\lceil P \rceil$ and setting the flag, while moving from **High** to **Done** removes $\lceil P \rceil$.

To ensure that these regions are invariant in concurrent contexts, we need to show that the interpretation of each abstract state, **High**, **Low** and **Done**, is stable. For these definitions, this is simple. Stability distributes over separating conjunction, so we can consider each conjunct separately. Points-to assertions are automatically stable. Furthermore, the higher-order parameter P is wrapped in explicit stabilization operators $\lceil - \rceil$, ensuring that it is also stable.

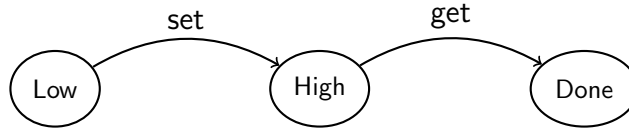
The transition relation T_b has the following form. Each transition corresponds to an operation that can be performed on the channel.


```

1  {send(x, P) * [P]}
2  signal(chan *x) {
3    // definition of send.
4    {region({Low}, Tb, Ib(x, P), r) * [set]1r * [P]}
5    // enter the region.
6    {▷(x.flag ↦ 0) * [set]1r * [P]}
7    // drop ▷ using LPoints
8    {x.flag ↦ 0 * [set]1r * [P]}
9    x->flag := 1;
10   {x.flag ↦ 1 * [set]1r * [P]}
11   // flip explicit stabilization with [P] ⇒ [P].
12   {x.flag ↦ 1 * [set]1r * [P]}
13   // close region and delete set token.
14   {region({High}, Tb, Ib(x, P), r)}
15   // stabilise the region's abstract state.
16   {region({High, Done}, Tb, Ib(x, P), r)}
17 }
18 {emp}

```

Fig. 7. Proof of `signal()` using simple predicate definitions.



Using these definitions, we can give the predicates `send` and `recv` an interpretation:

$$\begin{aligned}
\text{send}(x, P) &\triangleq \exists r. \text{region}(\{\text{Low}\}, T_b, I_b(x, P), r) * [\text{set}]_1^r \\
\text{recv}(x, P) &\triangleq \exists r. \text{region}(\{\text{Low}, \text{High}\}, T_b, I_b(x, P), r) * [\text{get}]_1^r
\end{aligned}$$

The `send` predicate asserts that the shared region is `Low` – otherwise we could not send the resource. On the other hand, `recv` asserts it is either `Low` or `High`, but not `Done`. $[\text{set}]_1^r$ and $[\text{get}]_1^r$ are tokens allowing the thread to take particular transitions in T_b . The `send` predicate allows the thread to set the flag and supply the resource, while `recv` allows the thread to retrieve the supplied resource.

Figures 7 and 8 show sketch proofs for `signal()` and `wait()` respectively. Both proofs involve accessing resources stored in the shared region – for example, `signal` enters the region between lines 5–13 of Fig. 7. This whole section corresponds to one atomic action. On entering the region at line 6, note the region invariant is wrapped in a \triangleright operator, but this can be dropped immediately using the `LPOINTS` rule.

When closing the region (Fig. 7, line 13), the resulting invariant must satisfy one of the abstract states. By examining the definition of I_b above, observe that `High` is satisfied.

PROOF BODY:

```

1  {recv(x, P)}
2  wait(chan *x) {
3    int b;
4    do {
5      {region({Low, High}, Tb, Ib(x, P), r) * [get]1r}
6      // case-split on Low / High.
7      {(region({Low}, Tb, Ib(x, P), r) ∨ region({High}, Tb, Ib(x, P), r)) * [get]1r}
8      b = x->flag; // Low & High cases given below.
9      {(region({Low, High}, Tb, Ib(x, P), r) * [get]1r ∧ b = 0) ∨
       (region({Done}, Tb, Ib(x, P), r) * [get]1r * [P] ∧ b = 1)}
10     } while (b == 0)
11   } // Garbage collect the shared region assertions.
12   {[P]}

```

LOW CASE:

```

1  {region({Low}, Tb, Ib(x, P), r) * [get]1r}
2  // open region
3  {▷(x.flag ↦ 0) * [get]1r}
4  // drop ▷ and read value.
5  b = x->flag;
6  {x.flag ↦ 0 * [get]1r ∧ b = 0}
7  // close region, and stabilise
8  {region({Low, High}, Tb, Ib(x, P), r)
   * [get]1r ∧ b = 0}

```

HIGH CASE:

```

1  {region({High}, Tb, Ib(x, P), r) * [get]1r}
2  // open region
3  {▷(x.flag ↦ 1 * [P]) * [get]1r}
4  // drop ▷ and read value.
5  b = x->flag;
6  {x.flag ↦ 1 * [P] * [get]1r ∧ b = 1}
7  // close region.
8  {region({Done}, Tb, Ib(x, P), r)
   * [P] * [get]1r ∧ b = 1}

```

Fig. 8. Proof of wait() using simple predicate definitions.

Furthermore, because the thread holds the **set** token, according to the transition relation T_b it is allowed to move from the **Low** (line 4) to the **High** (line 14) abstract state.

The final step of the proof is to ensure that all assertions are stable, i.e. invariant under interference from other threads. By itself, the assertion $\text{region}(\{\text{High}\}, T_b, I_b(x, P), r)$ is unstable, because according to T_b , some thread could move the region from **High** to **Done**. On line 15 we weaken the set of abstract states to add **Done**. This assertion is stable.

The proof of **wait** (Fig. 8) uses similar reasoning. 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. 8. 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.

5. SPLITTING CHANNELS

In this section, we extend our proof to cover *splitting*, expressed by the following axiom:

$$\{\text{recv}(a, P) * \llbracket P \rrbracket * (P_1 * P_2)\} \quad \langle \text{skip} \rangle \quad \{\text{recv}(a, P_1) * \text{recv}(a, P_2)\}$$

Verifying this axiom requires us to introduce some extra logical machinery. (As the simple implementation sequentialises signalling, we leave extension and renunciation to §6.)

5.1. Saved Propositions

The splitting axiom requires the proof to manipulate offered resource: the resource parameter to `recv` has to be split into two separate offered resources. For example, in the axiom, the offered resource P is split into resources P_1 and P_2 . As the splitting axiom is exposed to the client, P , P_1 and P_2 cannot be fixed in the proof; they are chosen by the client.

In iCAP, modifications to a shared region must be represented in the transition system for the region. An obvious (but wrong) way to represent splitting is to parameterise transition system states with sets of propositions representing the splitting. Then we would have transition system states $\text{Low}(\mathcal{I})$ and $\text{High}(\mathcal{I})$, with $\mathcal{I} \in \mathcal{P}_{\text{fin}}(\text{Prop})$ representing the splittings. In the case of the axiom, we would then have the transition:

$$\begin{aligned} \text{Low}(\mathcal{I} \uplus \{P\}) &\rightsquigarrow \text{Low}(\mathcal{I} \uplus \{P_1, P_2\}) \\ \text{High}(\mathcal{I} \uplus \{P\}) &\rightsquigarrow \text{High}(\mathcal{I} \uplus \{P_1, P_2\}) \end{aligned}$$

However, iCAP does not support states and actions indexed by step-indexed sets, such as `Prop` (the type of assertions). Our solution is to introduce a new logical construct we call a *saved proposition*. One can see saved propositions as capturing some the features of higher-order propositions stored in the heap. However, unlike true HO stored propositions, we can encode and verify saved propositions as predicates in iCAP, avoiding the need to develop a new logic. In our proofs, saved propositions have all the properties we need, and we are uncertain whether there would be any benefit in extending iCAP with true HO stored propositions. Our encoding of saved propositions is given in Appendix A.

A saved proposition, written $r \stackrel{\pi}{\Rightarrow} P$, associates an identifier r with a proposition P . As r comes from a non step-indexed set (in fact RID, the set of region names) we can parameterise transition system states by them. We thus represent each splitting by a set of identifiers combined with a collection of saved propositions. In the case of the axiom, we would assume saved propositions $r \stackrel{\pi}{\Rightarrow} P * r_1 \stackrel{\pi_1}{\Rightarrow} P_1 * r_2 \stackrel{\pi_2}{\Rightarrow} P_2$, and give the transition as:

$$\begin{aligned} \text{Low}(\mathcal{I} \uplus \{r\}) &\rightsquigarrow \text{Low}(\mathcal{I} \uplus \{r_1, r_2\}) \\ \text{High}(\mathcal{I} \uplus \{r\}) &\rightsquigarrow \text{High}(\mathcal{I} \uplus \{r_1, r_2\}) \end{aligned}$$

The difference is subtle but important. By introducing the indirection from identifiers to propositions we lose some properties. Most importantly, we cannot easily unify saved propositions: given $r \stackrel{\pi_1}{\Rightarrow} P$ and $r \stackrel{\pi_2}{\Rightarrow} Q$, 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 four properties:

$$\text{emp} \implies \exists r. r \stackrel{1}{\vdash} P \quad (1)$$

$$r \stackrel{\pi_1}{\vdash} P * r \stackrel{\pi_2}{\vdash} P \iff \begin{cases} r \stackrel{\pi_1 + \pi_2}{\vdash} P & \text{if } (\pi_1 + \pi_2) \leq 1 \\ \text{false} & \text{otherwise} \end{cases} \quad (2)$$

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

$$\left(\begin{array}{c} r \stackrel{\pi_1}{\vdash} P * r \stackrel{\pi_2}{\vdash} Q * \\ (X * \triangleright (Q * Y)) * (P * Z) \end{array} \right) \implies r \stackrel{\pi_1}{\vdash} P * r \stackrel{\pi_2}{\vdash} Q * X * \triangleright (Z * Y) \quad (4)$$

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 fixed 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. *Property* (4) says that we can apply property (3) inside separating implications. This is useful in the proof when modifying a resource embedded into a larger assertion.

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 $\text{Low}(\mathcal{I})$ and $\text{High}(\mathcal{I})$, where \mathcal{I} is a finite set of region identifiers in $\mathcal{P}_{fin}(\text{RID})$. 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 $\text{Low}(\mathcal{I})$ or $\text{High}(\mathcal{I})$, then each $i \in \mathcal{I}$ corresponds to a thread which expects to receive a resource. To find out *what* resource P is expected, we examine the associated saved proposition $i \stackrel{\pi}{\vdash} P$.

Actions in the transition relation T_m are of the form **send** and **change**(r), where r is a region identifier. The first action, **send**, sets the flag, and simply moves from Low to High . The second, **change**(r), is both the action of taking the resource associated to region r when the flag is high, and splitting the resource from region r to the resource required by r_1 and r_2 .

$$T_m(\text{send}) \triangleq \{(\text{Low}(\mathcal{I}), \text{High}(\mathcal{I}))\}$$

$$\begin{aligned} T_m(\text{change}(r)) \triangleq & \{(\text{High}(\mathcal{I} \uplus \{r\}), \text{High}(\mathcal{I}))\} \cup \{(\text{Low}(\mathcal{I} \uplus \{r\}), \text{Low}(\mathcal{I} \uplus \{r_1, r_2\}))\} \\ & \cup \{(\text{High}(\mathcal{I} \uplus \{r\}), \text{High}(\mathcal{I} \uplus \{r_1, r_2\}))\} \end{aligned}$$

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

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

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

where

$$\text{waiting}(P, \mathcal{I}) \triangleq \exists Q: \mathcal{I} \rightarrow \text{Prop}. \lfloor (\triangleright P) \multimap \bigotimes_{i \in \mathcal{I}} Q(i) \rfloor * \bigotimes_{i \in \mathcal{I}} i \xRightarrow{1/2} Q(i)$$

$$\text{ress}(\mathcal{I}) \triangleq \bigotimes_{i \in \mathcal{I}} \exists R. i \xRightarrow{1/2} R * \lceil \triangleright R \rceil$$

$$\text{changes}_r(\mathcal{I}) \triangleq \bigotimes_{r' \notin \mathcal{I}} [\text{change}(r')]_1^r$$

The definitions here use three auxiliary predicates: **waiting**, standing for resources that have been promised but not supplied; **ress**, standing for resources once they have been supplied; and **changes**, standing for change permissions on unused regions – this allows new shared propositions to be added when splitting.

The representation of **Low** consists of the flag, change tokens, and the **waiting** predicate. $\text{waiting}(P, \mathcal{I})$ 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 shared propositions $i \xRightarrow{1/2} Q(i)$. The assertion $\lfloor (\triangleright P) \multimap \bigotimes_{i \in \mathcal{I}} Q(i) \rfloor$ 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. $\text{ress}(\mathcal{I})$ pairs together fractional saved propositions, $i \xRightarrow{1/2} R$ 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:

$$\text{creg}(x, r, P, S) \triangleq \text{region}(S, T_m, I_m(x, r, P), r)$$

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.

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

The definition of $\text{recv}(x, Q)$ predicate is more complex. It includes $r' \xRightarrow{1/2} Q$, 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).

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

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

Proving newchan. The proof of **newchan** (Fig. 9) works by allocating an abstract region containing the concrete state for the channel. Most steps in the proof are straightforward. The challenging ones are 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:

$$\text{emp} \implies \lfloor (\triangleright P) \multimap (\triangleright P) \rfloor$$

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

```

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  $\xRightarrow{1/2}$  P * x.flag ↦ 0 * [(▷P) ↯ (▷P)]}
10   // Create channel region
11   {∃r1, r2. r1  $\xRightarrow{1/2}$  P * creg(x, r2, P, {Low({r1})}) * [send]1r2 * [change(r1)]1r2}
12   // Satisfy the send predicate
13   {∃r1, r2. send(x, P) * r1  $\xRightarrow{1/2}$  P * creg(x, r2, P, {Low({r1})}) * [change(r1)]1r2}
14 } // Satisfy the recv predicate
15 {send(x, P) * recv(x, P)}

```

Fig. 9. Proof of `newchan()` w.r.t. the full specification.

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

$$\begin{aligned}
& r_1 \xRightarrow{1/2} P * x \mapsto 0 * [(\triangleright P) \multimap (\triangleright P)] \\
& \sqsubseteq \exists r_2. \text{creg}(x, r_2, P, \{\text{Low}(\{r_1\}, \emptyset)\}) * [\text{send}]_1^{r_2} * [\text{change}(r_1)]_1^{r_2}
\end{aligned}$$

To prove this, we appeal to iCAP's VALLOC rule, which controls construction of new regions.

Proving signal. The proof of `signal` (Fig. 10) works by opening the channel region (line 4), merging in the supplied resource $[P]$ 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 $\text{Low}(\mathcal{I})$ to $\text{High}(\mathcal{I})$ 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. $[P] * \text{waiting}(P, \mathcal{I}) \sqsubseteq \text{ress}(\mathcal{I})$

PROOF. $[P] * \text{waiting}(P, \mathcal{I})$

$$\begin{aligned}
& \sqsubseteq [\triangleright P] * \exists Q: \mathcal{I} \rightarrow \text{Prop}. [\triangleright P \multimap \bigotimes_{i \in \mathcal{I}} Q(i)] * (\bigotimes_{i \in \mathcal{I}} i \xRightarrow{1/2} Q(i)) \\
& \sqsubseteq \exists Q: \mathcal{I} \rightarrow \text{Prop}. [\triangleright \bigotimes_{i \in \mathcal{I}} Q(i)] * (\bigotimes_{i \in \mathcal{I}} i \xRightarrow{1/2} Q(i)) \\
& \sqsubseteq \exists Q: \mathcal{I} \rightarrow \text{Prop}. (\bigotimes_{i \in \mathcal{I}} [\triangleright Q(i)]) * (\bigotimes_{i \in \mathcal{I}} i \xRightarrow{1/2} Q(i)) \\
& \sqsubseteq (\bigotimes_{i \in \mathcal{I}} \exists R. i \xRightarrow{1/2} R * [\triangleright R]) \\
& \sqsubseteq \text{ress}(\mathcal{I})
\end{aligned}$$

To prove the second step, we appeal to the fact that $[-]$ is semi-distributive over separating conjunction, $[A] * [B] \implies [A * B]$, and *modus ponens* for separating implication, $A * (A \multimap B) \implies B$. The third step follows from the fact that $[-]$ is weaker than $[-]$, and $[-]$ is semi-distributive over separating conjunction, $[A * B] \implies [A] * [B]$. \square

```

1  {send(x, P) * [P]}
2  signal(chan *x) {
3    {creg(x, r, P, {Low(I') | true}) * [send]1r * [P]}
4    // open the region. Low(I) is an arbitrary member of {Low(I') | true}.
5    {▷Im(x, r, P)(Low(I)) * [send]1r * [P]}
6    // invariant definition.
7    {[send]1r * [P] * ▷(x.flag ↦ 0 * waiting(P, I) * changesr(I))}
8    // pull out points-to using LBin and LPoints.
9    {[send]1r * [P] * x.flag ↦ 0 * ▷(waiting(P, I) * changesr(I))}
10   x->flag := 1; // rewrite flag, drop ▷ using AFrame
11   {[send]1r * [P] * x.flag ↦ 1 * waiting(P, I) * changesr(I)}
12   // Lemma 5.1, add a ▷ using SMono.
13   {[send]1r * ▷(x.flag ↦ 1 * ressource(I) * changesr(I))}
14   // Invariant definition.
15   {[send]1r * ▷Im(x, r, P)(High(I))}
16   // close the region.
17   {[send]1r * creg(x, r, P, {High(I') | true})}
18 } // delete unneeded predicates.
19 {emp}

```

Fig. 10. Proof of `signal()` w.r.t. the full specification.

Proving wait. In the proof of `wait` (Fig. 11) 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 \mathcal{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.

LEMMA 5.2. $r \xRightarrow{1/2} P * \text{ress}(\mathcal{I} \uplus \{r\}) \sqsubseteq \text{ress}(\mathcal{I}) * \triangleright[P]$

PROOF. $r \xRightarrow{1/2} P * \text{ress}(\mathcal{I} \uplus \{r\})$

$$\begin{aligned}
& \sqsubseteq r \xRightarrow{1/2} P * \text{ress}(\mathcal{I}) * \exists R. r \xRightarrow{1/2} R * \triangleright[R] && \text{(Rearrange)} \\
& \sqsubseteq r \xRightarrow{1/2} P * \text{ress}(\mathcal{I}) * \exists R. r \xRightarrow{1/2} R * \triangleright[P] && \text{(Property 3, mono of } \triangleright[-] \text{)} \\
& \sqsubseteq \text{ress}(\mathcal{I}) * \triangleright[P] && \text{(LCEIL/LFLOOR)}
\end{aligned}$$

□

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 12. The core of the proof is two lemmas which express splitting in the `Low` and `High` cases.

```

1  {recv(x, P)}
2  wait(chan *x) {
3    { $\exists R, r, r'. r' \xRightarrow{1/2} P * [\text{change}(r')]_1^r * \text{creg}(x, r, R, \{\text{LoHi}(\mathcal{I}') \mid r' \in \mathcal{I}'\})$ }
4    // Open the region. Only consider the High case.
5    { $r' \in \mathcal{I} \wedge r' \xRightarrow{1/2} P * [\text{change}(r')]_1^r * \triangleright I_m(x, r, R)(\text{High}(\mathcal{I}))$ }
6    // Apply invariant definition
7    // Pull out points-to using LBin and LPoints
8    { $r' \in \mathcal{I} \wedge r' \xRightarrow{1/2} P * [\text{change}(r')]_1^r * x.\text{flag} \mapsto 1 * \triangleright (\text{ress}(\mathcal{I}) * \text{changes}_r(\mathcal{I}))$ }
9    assume(x->flag == 1) // drop  $\triangleright$  using AFrame, push in [change(r')] perm.
10   { $r' \in \mathcal{I} \wedge r' \xRightarrow{1/2} P * x.\text{flag} \mapsto 1 * \text{ress}(\mathcal{I}) * \text{changes}_r(\mathcal{I} \uplus \{r'\})$ }
11   // Lemma 5.2, add  $\triangleright$  using SMono
12   { $\triangleright [P] * \triangleright (x.\text{flag} \mapsto 1 * \text{ress}(\mathcal{I} \setminus \{r'\}) * \text{changes}_r(\mathcal{I} \setminus \{r'\}))$ }
13   // Invariant definition
14   { $\triangleright [P] * \triangleright I_m(x, r, R)(\text{High}(\mathcal{I} \setminus \{r'\}))$ }
15   // close the region
16   { $\triangleright [P] * \exists R, r. \text{creg}(x, r, R, \{\text{High}(\mathcal{I}) \mid \text{true}\})$ }
17   // abstract garbage collect
18   { $\triangleright [P]$ }
19 }

```

Fig. 11. Proof of wait() w.r.t. the full specification.

```

1  {recv(x, P) *  $\llbracket [P] \multimap P_1 * P_2 \rrbracket$ }
2  { $\exists R, r, r'. r' \xRightarrow{1/2} P * [\text{change}(r')]_1^r * \text{creg}(x, r, R, \{\text{LoHi}(\mathcal{I}') \mid r' \in \mathcal{I}'\}) * \llbracket [P] \multimap P_1 * P_2 \rrbracket$ }
3  // Open the region using an arbitrary state containing r'
4  { $\exists R, r, r'. r' \xRightarrow{1/2} P * [\text{change}(r')]_1^r * \triangleright I_m(x, r, R)(\text{LoHi}(\mathcal{I} \uplus \{r'\})) * \llbracket [P] \multimap P_1 * P_2 \rrbracket$ }
5  skip // Use AFrame to remove  $\triangleright$ 
6  { $\exists R, r, r'. r' \xRightarrow{1/2} P * [\text{change}(r')]_1^r * I_m(x, r, R)(\text{LoHi}(\mathcal{I} \uplus \{r'\})) * \llbracket [P] \multimap P_1 * P_2 \rrbracket$ }
7  // Lemma 5.5
8  { $\exists R, r, r_1, r_2. r_1 \xRightarrow{1/2} P_1 * r_1 \xRightarrow{1/2} P_2 * [\text{change}(r_1)]_1^r * [\text{change}(r_2)]_1^r$   

   { $* I_m(x, r, R)(\text{LoHi}(\mathcal{I} \uplus \{r_1, r_2\}))$ }
9  // Close the region
10 { $\exists R, r, r_1, r_2. r_1 \xRightarrow{1/2} P_1 * r_1 \xRightarrow{1/2} P_2 * [\text{change}(r_1)]_1^r * [\text{change}(r_2)]_1^r$   

   { $* \text{creg}(x, r, R, \{\text{LoHi}(\mathcal{I}') \mid r_1, r_2 \in \mathcal{I}'\})$ }
11 {recv(x, P1) * recv(x, P2)}

```

Fig. 12. Proof outline for splitting axiom.

$$\begin{aligned}
\text{LEMMA 5.3. } r &\xRightarrow{1/2} P * \llbracket [P] \multimap P_1 * P_2 \rrbracket * \text{waiting}(R, \mathcal{I} \uplus r) \\
&\sqsubseteq \exists r_1, r_2. r_1 \xRightarrow{1/2} P_1 * r_2 \xRightarrow{1/2} P_2 * \text{waiting}(R, \mathcal{I} \uplus \{r_1, r_2\})
\end{aligned}$$

PROOF.

$$\begin{aligned}
r &\xRightarrow{1/2} P * \llbracket [P] \multimap P_1 * P_2 \rrbracket * \text{waiting}(R, \mathcal{I} \uplus r) \\
&\sqsubseteq r \xRightarrow{1/2} P * \llbracket [P] \multimap P_1 * P_2 \rrbracket * \\
&\quad \exists Q: \mathcal{I} \uplus \{r\} \rightarrow \text{Prop}. \llbracket (\triangleright R) \multimap \bigotimes_{i \in \mathcal{I} \uplus r}. Q(i) \rrbracket * (\bigotimes_{i \in \mathcal{I} \uplus r}. i \xRightarrow{1/2} Q(i)) \\
&\sqsubseteq \exists Q'. r \xRightarrow{1/2} P * r \xRightarrow{1/2} Q' * \llbracket [P] \multimap P_1 * P_2 \rrbracket * \\
&\quad \exists Q: \mathcal{I} \rightarrow \text{Prop}. \llbracket (\triangleright R) \multimap Q' * \bigotimes_{i \in \mathcal{I}}. Q(i) \rrbracket * (\bigotimes_{i \in \mathcal{I}}. i \xRightarrow{1/2} Q(i)) \\
&\sqsubseteq \llbracket [P] \multimap P_1 * P_2 \rrbracket * \\
&\quad \exists Q: \mathcal{I} \rightarrow \text{Prop}. \llbracket (\triangleright R) \multimap \triangleright P * \bigotimes_{i \in \mathcal{I}}. Q(i) \rrbracket * (\bigotimes_{i \in \mathcal{I}}. i \xRightarrow{1/2} Q(i)) \\
&\sqsubseteq \llbracket \triangleright [P] \multimap (\triangleright P_1) * (\triangleright P_2) \rrbracket * \\
&\quad \exists Q: \mathcal{I} \rightarrow \text{Prop}. \llbracket (\triangleright R) \multimap \triangleright [P] * \bigotimes_{i \in \mathcal{I}}. Q(i) \rrbracket * (\bigotimes_{i \in \mathcal{I}}. i \xRightarrow{1/2} Q(i)) \\
&\sqsubseteq \exists Q: \mathcal{I} \rightarrow \text{Prop}. \llbracket (\triangleright R) \multimap (\triangleright P_1) * (\triangleright P_2) * \bigotimes_{i \in \mathcal{I}}. Q(i) \rrbracket * (\bigotimes_{i \in \mathcal{I}}. i \xRightarrow{1/2} Q(i)) \\
&\sqsubseteq \exists r_1, r_2. r_1 \xRightarrow{1} P_1 * r_2 \xRightarrow{1} P_2 \wedge r_1, r_2 \notin \mathcal{I} * \\
&\quad \exists Q: \mathcal{I} \rightarrow \text{Prop}. \llbracket (\triangleright R) \multimap (\triangleright P_1) * (\triangleright P_2) * \bigotimes_{i \in \mathcal{I}}. Q(i) \rrbracket * (\bigotimes_{i \in \mathcal{I}}. i \xRightarrow{1/2} Q(i)) \\
&\sqsubseteq \exists r_1, r_2. r_1 \xRightarrow{1/2} P_1 * r_2 \xRightarrow{\frac{1}{2}} P_2 * \\
&\quad \exists Q: \mathcal{I} \uplus \{r_1, r_2\} \rightarrow \text{Prop}. \llbracket (\triangleright R) \multimap \bigotimes_{i \in \mathcal{I} \uplus \{r_1, r_2\}}. Q(i) \rrbracket * (\bigotimes_{i \in \mathcal{I} \uplus \{r_1, r_2\}}. i \xRightarrow{1/2} Q(i)) \\
&\sqsubseteq \exists r_1, r_2. r_1 \xRightarrow{1/2} P_1 * r_2 \xRightarrow{\frac{1}{2}} P_2 * \text{waiting}(R, \mathcal{I} \uplus \{r_1, r_2\})
\end{aligned}$$

□

$$\begin{aligned}
\text{LEMMA 5.4. } r &\xRightarrow{1/2} P * \llbracket [P] \multimap P_1 * P_2 \rrbracket * \text{ress}(\mathcal{I} \uplus r) \\
&\sqsubseteq \exists r_1, r_2. r_1 \xRightarrow{1/2} P_1 * r_2 \xRightarrow{1/2} P_2 * \text{ress}(\mathcal{I} \uplus \{r_1, r_2\})
\end{aligned}$$

PROOF.

$$\begin{aligned}
r &\xRightarrow{1/2} P * \llbracket [P] \multimap P_1 * P_2 \rrbracket * \text{ress}(\mathcal{I} \uplus r) \\
&\sqsubseteq (\llbracket [P] \multimap P_1 * P_2 \rrbracket * \text{ress}(\mathcal{I}) * \triangleright [P]) && (\text{Lemma 5.2}) \\
&\sqsubseteq ((\triangleright [P]) \multimap (\triangleright P_1) * (\triangleright P_2)) * \text{ress}(\mathcal{I}) * \triangleright [P] && (\text{SMono, dist } \triangleright \text{ over } \multimap) \\
&\sqsubseteq \text{ress}(\mathcal{I}) * (\triangleright P_1) * (\triangleright P_2) && (\text{Modus poens}) \\
&\sqsubseteq \text{ress}(\mathcal{I}) * (\triangleright P_1) * (\triangleright P_2) * \exists r_1, r_2. r_1 \xRightarrow{1} P_1 * r_2 \xRightarrow{1} P_2 && (\text{Property 1}) \\
&\sqsubseteq \text{ress}(\mathcal{I} \uplus \{r_1, r_2\}) * \exists r_1, r_2. r_1 \xRightarrow{1/2} P_1 * r_2 \xRightarrow{1/2} P_2 && (\text{Sub-lemma})
\end{aligned}$$

The last step consists of two applications of the following sub-lemma:

$$\begin{aligned}
&\text{ress}(\mathcal{I}) * \triangleright P * r \xRightarrow{1} P \\
&\sqsubseteq (\bigotimes_{i \in \mathcal{I}} \exists Q. i \xRightarrow{1/2} Q * \llbracket \triangleright Q \rrbracket) * \llbracket \triangleright P \rrbracket * r \xRightarrow{1} P \\
&\sqsubseteq (\bigotimes_{i \in \mathcal{I} \uplus \{r\}} \exists Q. i \xRightarrow{1/2} Q * \llbracket \triangleright Q \rrbracket) * r \xRightarrow{1/2} P \\
&\sqsubseteq \text{ress}(\mathcal{I} \uplus \{r\}) * r \xRightarrow{1/2} P
\end{aligned}$$

□

These lemmas are combined as follows.

$$\begin{aligned}
\text{LEMMA 5.5. } r' &\stackrel{1/2}{\Longrightarrow} P * I_m(x, r, R)(\text{LoHi}(\mathcal{I} \uplus \{r'\})) * [\text{change}(r')]_r^1 * \llbracket P \rrbracket \multimap P_1 * P_2 \\
&\sqsubseteq \\
&\exists r_1, r_2. [\text{change}(r_1)]_r^1 * r_1 \stackrel{1/2}{\Longrightarrow} P_1 * [\text{change}(r_2)]_r^1 * r_2 \stackrel{1/2}{\Longrightarrow} P_2 * \\
&\quad \triangleright I_m(x, r, R)(\text{LoHi}(\mathcal{I} \uplus \{r_1, r_2\}))
\end{aligned}$$

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. \square

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. Calls to **signal** can complete in any order consistent with the specification. To ensure renounced resources are available, **wait** checks all predecessors in the chain. We prove that this implementation implements our abstract specification.

This new implementation uses a linked list data-structure. The implementations of **extend**, **newchan**, **signal** and **wait** operations are defined as follows:

```

struct chan {
    int flag;
    chan *prev;
}

signal(chan *x) {
    x->flag = 1;
}

wait(chan *x) {
    chan *c = x;
    while(c != NULL) {
        while(c->flag == 0) skip;
        c = c->prev;
    }
}

chan *newchan() {
    chan *x = new(chan);
    x->flag = 0;
    x->prev = NULL;
    return x;
}

extend(chan *x) {
    chan *z = new(chan);
    z->flag = 0;
    z->prev = x->prev;
    x->prev = z;
    return (z,x);
}

```

Calling **signal** sets the 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.

6.1. Proof Structure

In our proof of the simple barrier implementation, each shared region contained a single channel. In the current implementation, **wait** traverses the whole preceding chain. To allow this, the proof stores the whole chain in a single region, represented by a single invariant. This makes it more complex, but the core approach remains the same:

- **send**, **recv**, and ordering predicates are reified as constraints on the shared region.
- Resource obligations are represented by sets of identifiers, which are tied to saved propositions.

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.

Abstract chain type. To represent the state of channels in the chain, we define a type of abstract chain nodes, **CNode**:

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

Each **CNode** represents one channel in the chain, while an *abstract chain* is a finite sequence in \mathbf{CNode}^+ . Given an element $(loc, res, \mathcal{I}, flag, \mathcal{W}) \in \mathbf{CNode}$, loc is the physical address, res is the identifier used when sending resources, \mathcal{I} is the set of identifiers for resources promised to other predicates, $flag$ is the current state of the flag, and \mathcal{W} is the set of identifiers for resources earlier in the chain promised to this channel through renunciation. (Given a **CNode** s , we sometimes write $s.flag$, $s.\mathcal{I}$ etc. to identify the appropriate components of the tuple. We use 0 and 1 to represent the **Low** and **High** flag state, respectively.)

Given an abstract chain $x \cdot xs$, it is important that identifiers in the set $x.\mathcal{W}$ are in the set of resources promised in xs (i.e. in some set \mathcal{I}). To ensure this, we define well-formedness of abstract chains as follows:

$$\begin{aligned} \mathbf{available}([\]) &\triangleq \emptyset, & \mathbf{available}(s \cdot xs) &\triangleq (\mathbf{available}(xs) \setminus s.\mathcal{W}) \uplus s.\mathcal{I} \\ \mathbf{wf}([\]) &\triangleq \mathbf{true}, & \mathbf{wf}(s \cdot xs) &\triangleq \mathbf{wf}(xs) \wedge s.\mathcal{W} \subseteq \mathbf{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 \end{aligned}$$

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 promised resources are empty.

$$\begin{aligned} \mathbf{ctrue}(rs) &\triangleq \forall e \in rs. e.flag = 1 \\ \mathbf{cconf}(rs) &\triangleq \forall e \in rs. e.flag = 1 \wedge e.\mathcal{W} = \emptyset \end{aligned}$$

Finally, we use $rs_1 \xrightarrow{\mathcal{W}}^* rs_2$ to denote that the abstract chain rs_2 follows from rs_1 by cancelling out renounced resources with the corresponding promise. Formally, $\xrightarrow{\mathcal{W}}^*$ is defined as the transitive-reflexive closure of the following relation, which cancels a single promise to a later node using a renounced resource from an earlier node.

$$\begin{aligned} rs \xrightarrow{\mathcal{W}} rs' &\triangleq \exists x, y, r. \\ & r \in rs(x).\mathcal{W} \wedge r \in rs(y).\mathcal{I} \wedge (x, y) \in \mathbf{ord}(rs) \wedge \\ & rs' = (rs \triangleleft_x \triangleleft_{\mathcal{W}} (\bullet \setminus \{r\})) \triangleleft_y \triangleleft_{\mathcal{I}} (\bullet \setminus \{r\}) \end{aligned}$$

The notation \triangleleft is a *lens* allowing a single field of a chain to be updated without modifying the remainder of the chain.

$$x \triangleleft_i f \triangleq x[i \mapsto (f[x(i)/\bullet])] \qquad x \triangleleft_i \triangleleft_j f \triangleq x[i \mapsto (x(i) \triangleleft_j f)]$$

$$\begin{aligned}
\text{renun}_c(x, rs, rs') &\triangleq rs(x).flg = 0 \wedge rs' = (rs \blacktriangleleft_x \blacktriangleleft_{res} r') \blacktriangleleft_x \blacktriangleleft_{\mathcal{W}} (\bullet \uplus w) \\
\text{set}_c(x, rs, rs') &\triangleq rs(x).flg = 0 \wedge rs' = rs \blacktriangleleft_x \blacktriangleleft_{flg} (1) \\
\text{ext}_c(x, rs, rs') &\triangleq rs = (rs_1 \cdot a \cdot rs_2) \wedge rs' = (rs_1 \cdot a \cdot b \cdot rs_2) \wedge \\
&\quad a.loc = x \wedge a.flg = 0 \wedge b.flg = 0 \wedge b.\mathcal{W} = \emptyset \\
\text{split}_c(x, r, rs, rs') &\triangleq r \in rs(x).\mathcal{I} \wedge rs' = rs \blacktriangleleft_x \blacktriangleleft_{\mathcal{I}} ((\bullet \setminus \{r\}) \uplus \{r_2, r_3\}) \\
\text{sat}_c(x, rs, rs') &\triangleq rs = (rs_1 \cdot rs_2) \wedge rs' = (rs_1 \cdot rs'_2) \wedge \text{cconf}(rs'_2) \wedge rs_2 \xrightarrow{\mathcal{W}^*} rs'_2 \\
\text{get}_c(x, r, rs, rs') &\triangleq rs = (rs_1 \cdot rs(x) \cdot rs_2) \wedge \text{cconf}(rs(x) \cdot rs_2) \wedge rs' = rs \blacktriangleleft_x \blacktriangleleft_{\mathcal{I}} (\bullet \setminus \{r\}) \\
\\
T_c(\text{send}(x)) &\triangleq \{(a, b) \mid \text{wf}(b) \wedge (\text{renun}_c(x, a, b) \vee \text{set}_c(x, a, b) \vee \text{ext}_c(x, a, b))\} \\
T_c(\text{change}(x, r)) &\triangleq \{(a, b) \mid \text{wf}(b) \wedge (\text{split}_c(x, r, a, b) \vee \text{sat}_c(x, a, b) \vee \text{get}_c(x, r, a, b))\}
\end{aligned}$$

Fig. 13. Definition of T_c , the transition relation for the chained channel implementation.

$$\begin{aligned}
\text{chainds}(x \cdot y \cdot rs) &\triangleq x.loc \mapsto \{\text{prev} = y.loc; \text{flag} = x.flg\} * \text{chainds}(y \cdot rs) \\
\text{chainds}(x \cdot \text{null}) &\triangleq x.loc \mapsto \{\text{prev} = \text{NULL}; \text{flag} = x.flg\} \\
\\
\text{resource}(\mathcal{I}, \mathcal{W}) &\triangleq \exists Q: \mathcal{I} \rightarrow \text{Prop}, R: \mathcal{W} \rightarrow \text{Prop}. \\
&\quad \bigotimes_{i \in \mathcal{I}} i \xrightarrow{1/2} Q(i) * \bigotimes_{w \in \mathcal{W}} w \xrightarrow{1/2} R(w) \\
&\quad * \lfloor (\triangleright \bigotimes_{w \in \mathcal{W}} R(w)) * \triangleright \bigotimes_{i \in \mathcal{I}} \lceil Q(i) \rceil \rfloor \\
\text{chainres}(x \cdot rs) &\triangleq \text{resource}(x.\mathcal{I}, x.\mathcal{W} \uplus \{x.res \mid \neg x.flg\}) * \text{chainres}(rs) \\
\text{chainres}(\text{null}) &\triangleq \text{emp} \\
\\
\text{uS}(rs) &\triangleq \{x \mid (x, \neg, \neg, 0, \neg) \in rs\} \\
\text{uC}(rs) &\triangleq \{(x, i) \mid (x, \neg, \mathcal{I}, \neg, \neg) \in rs \wedge i \in \mathcal{I} \wedge \neg \exists (y, \neg, \neg, \mathcal{W}) \in rs. i \in \mathcal{W}\} \\
\text{unused}(r, rs) &\triangleq (\bigotimes x \notin \text{uS}(rs). [\text{send}(x)]_1^r) * (\bigotimes (x, r') \notin \text{uC}(rs). [\text{change}(x, r')]_1^r)
\end{aligned}$$

Fig. 14. Predicates used in defining the state of a region.

Predicate definitions. As usual, we begin by defining the structure of the shared region. Abstract states have the form $\text{Chain}(rs)$, where rs is an abstract chain. Actions have the form $\text{send}(x)$ and $\text{change}(x, r)$, where x is an address, and r a region identifier. The transition relation T_c is defined in Fig. 13. We assume that physical addresses are used uniquely, so where convenient we use chains as finite functions of type

$$\text{Addr} \xrightarrow{fin} (\text{Rld} \times \mathcal{P}(\text{Rld}) \times \{\text{High}, \text{Low}\} \times \mathcal{P}(\text{Rld}))$$

The transition relation defines six kinds of transitions in Fig. 13. For send we have renunciation, which adds an element to \mathcal{W} ; setting the flag; and extending the chain, which creates a new $\text{CNode } b$. 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. 14). The predicate chainds represents the list

data-structure underpinning the implementation. Each link in the chain has the appropriate prev and flag values set.

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 renounced resources that are being waited for \mathcal{W} . When there are no waiting resources the predicate can be drastically simplified to just the promised resources:

$$\text{LEMMA 6.1. } \text{resource}(\mathcal{I}, \emptyset) \sqsubseteq \bigotimes_{i \in \mathcal{I}} \exists Q: \text{Prop. } i \xrightarrow{1/2} Q(i) * [\triangleright Q(i)]$$

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:

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

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

$$\text{oreg}(r, S) \triangleq \text{region}(S, T_c, I_c(r), r)$$

We can now define the **send**, **recv**, and ordering predicates. We write $(x, y) \in \text{seq}(rs)$ to say that the two addresses x and y appear adjacent in the sequence rs , and $(x, y) \in \text{ord}(rs)$ to say just that they are ordered in rs .

$$\begin{aligned} \text{send}(x, P) \triangleq \exists r_1, r_2. \text{oreg}(r_1, \{\text{Chain}(rs) \mid \text{wf}(rs) \wedge rs(x) = (r_2, -, 0, -)\}) \\ * r_2 \xrightarrow{1/2} P * [\text{send}(x)]_1^{r_1} \end{aligned}$$

$$\begin{aligned} \text{recv}(x, P) \triangleq \exists r_1, r_2. \text{oreg}(r_1, \{\text{Chain}(rs) \mid rs(x) = (-, \mathcal{I}, -, -) \wedge r_2 \in \mathcal{I} \wedge \text{wf}(rs)\}) \\ * r_2 \xrightarrow{1/2} P * [\text{change}(x, r_2)]_1^{r_1} \end{aligned}$$

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

6.2. Proving signal, wait and extend

Proving signal. The sketch-proof is shown in Fig. 15 – 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 is sufficient to allow the flag to be set.

$$\text{LEMMA 6.2. } r \xrightarrow{1/2} P * \lfloor P \rfloor * \text{resource}(\mathcal{I}, \mathcal{W} \uplus \{r\}) \sqsubseteq \text{resource}(\mathcal{I}, \mathcal{W})$$

PROOF. Given in Appendix B. □

In the proof of **signal**, this lemma is used to show that the appropriate resource has been supplied (Fig. 15, 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. 16. The three most important steps are checking that all preceding bits are set (lines 5–11), checking that renounced resources have been supplied (line 14), and retrieving the resource from the chain (line 18). The last two of these require helper lemmas, given below.

```

1  {send(x, P) * [P]}
2  // Definition of send.
3  {∃ r1, r2. r2  $\xRightarrow{1/2}$  P * [send(x)]1r1 * [P] *
   oreg(r1, {Chain(rs) | wf(rs) ∧ rs(x) = (r2, I, 0, W)})}
4  // Enter the region.
5  {∃ r1, r2. r2  $\xRightarrow{1/2}$  P * [send(x)]1r1 * [P] *
   ▷ (chainds(rs) * chainres(rs) * unused(r1, rs) ∧ wf(rs) ∧ rs(x) = (r2, I, 0, W))}
6  // Split up using chain axioms.
7  {∃ r1, r2. r2  $\xRightarrow{1/2}$  P * [send(x)]1r1 * [P] *
   ▷ (∃ rs1, rs2, a. chainds(rs1) * x ↦ {prev = hd(rs2).loc; flag = 0} * chainds(rs2) *
     chainres(rs) * unused(r1, rs) ∧ rs = rs1 · a · rs2 ∧ wf(rs) ∧ a = (x, r2, I, 0, W))}
8  x->flag = 1; // Set the flag, drop the ▷ using Aframe.
9  {∃ r1, r2. r2  $\xRightarrow{1/2}$  P * [send(x)]1r1 * [P] *
   ∃ rs1, rs2, a. chainds(rs1) * x ↦ {prev = hd(rs2).loc; flag = 1} * chainds(rs2) *
   chainres(rs) * unused(r1, rs) ∧ rs = rs1 · a · rs2 ∧ wf(rs) ∧ a = (x, r2, I, 0, W)}
10 // Pull out resource predicate.
11 {∃ r1, r2. r2  $\xRightarrow{1/2}$  P * [send(x)]1r1 * [P] *
   ∃ rs1, rs2, a. chainds(rs1) * x ↦ {prev = hd(rs2).loc; flag = 1} * chainds(rs2)
   * chainres(rs1) * resource(I, W ⊔ {r2}) * chainres(rs2)
   * unused(r1, rs) ∧ rs = rs1 · a · rs2 ∧ wf(rs) ∧ a = (x, r2, I, 0, W)}
12 // Apply Lemma 6.2.
13 {∃ r1, r2. [send(x)]1r1 *
   ∃ rs1, rs2, a. chainds(rs1) * x ↦ {prev = hd(rs2).loc; flag = 1} * chainds(rs2)
   * chainres(rs1) * resource(I, W) * chainres(rs2)
   * unused(r1, rs) ∧ rs = rs1 · a · rs2 ∧ wf(rs) ∧ a = (x, r2, I, 0, W)}
14 // Collapse the chain, add ▷.
15 {∃ r1, r2.
   ▷ (chainds(rs') * chainres(rs') * unused(r1, rs') ∧ rs' = rs[x ↦ (r2, I, 1, W)])}
16 // Close the region using transition set.
17 // Well-formed structure of chain hasn't changed.
18 {∃ r1, r2. oreg(r1, {Chain(rs) | wf(rs) ∧ rs(x) = (r2, I, 1, W)})}
19 // Garbage collect.
20 {emp}

```

Fig. 15. Sketch-proof for `signal` with out-of-order signalling.

```

1  {recv(x, P)}
2  wait(x){
3    chan *c = x;
4    // Definition of recv.
5    {
6       $\exists r_1, r_2. \text{oreg}(r_1, \{\text{Chain}(rs) \mid \exists \mathcal{I}'. rs(\mathbf{x}) = (-, \mathcal{I}', -, -) \wedge r_2 \in \mathcal{I}' \wedge \text{wf}(rs)\})$ 
7       $* r_2 \xRightarrow{1/2} P * [\text{change}(\mathbf{x}, r_2)]_1^{r_1} \wedge c = \mathbf{x}$ 
8    }
9    while(c != NULL) {
10     while(!c->flag) skip;
11     c = c->prev;
12   }
13   // Flags up to x are now set. Stable because flags cannot get unset.
14   {
15      $\exists r_1, r_2. r_2 \xRightarrow{1/2} P * [\text{change}(\mathbf{x}, r_2)]_1^{r_1} *$ 
16      $\text{oreg}\left(r_1, \left\{ \text{Chain}(rs) \mid \begin{array}{l} \exists \mathcal{I}'. \text{wf}(rs) \wedge \exists i. rs[i] = (\mathbf{x}, -, \mathcal{I}' \uplus r_2, -, -) \wedge \\ \forall j. i \leq j < \text{len}(rs) \implies rs[j] = (-, -, -, 1, -) \end{array} \right\}\right)$ 
17   }
18   <skip>; // Enter the region and drop  $\triangleright$  using AFrame.
19   {
20      $\exists rs_1, rs_2, a. r_2 \xRightarrow{1/2} P * [\text{change}(\mathbf{x}, r_2)]_1^{r_1} * \text{chainds}(rs) * \text{chainres}(rs) * \text{unused}(r_1, rs) \wedge$ 
21      $\text{wf}(rs) \wedge rs = rs_1 \cdot a \cdot rs_2 \wedge a = (\mathbf{x}, r, \mathcal{I}' \uplus r_2, -, \mathcal{W}) \wedge \text{ctrue}(a \cdot rs_2)$ 
22   }
23   // Apply Lemma 6.3 to convert from ctrue to cconf..
24   {
25      $\exists rs_1, rs_2, rs'_2, a. r_2 \xRightarrow{1/2} P * [\text{change}(\mathbf{x}, r_2)]_1^{r_1} *$ 
26      $\text{chainds}(rs) * \text{chainres}(rs_1 \cdot rs'_2) * \text{unused}(r_1, rs) \wedge \text{wf}(rs) \wedge rs = rs_1 \cdot a \cdot rs_2 \wedge$ 
27      $a = (\mathbf{x}, r, \mathcal{I}' \uplus r_2, -, \mathcal{W}) \wedge \text{ccconf}(rs'_2) \wedge a \cdot rs_2 \xrightarrow{\mathcal{W}}^* rs'_2$ 
28   }
29   // Pull out the resource predicate for x.
30   {
31      $\exists rs_1, rs_2, rs'_2, a. r_2 \xRightarrow{1/2} P * [\text{change}(\mathbf{x}, r_2)]_1^{r_1} *$ 
32      $\text{chainds}(rs) * \text{chainres}(rs_1) * \text{resource}(\mathcal{I}' \uplus r_2, \emptyset) * \text{chainres}(rs'_2) * \text{unused}(r_1, rs) \wedge$ 
33      $\text{wf}(rs) \wedge rs = rs_1 \cdot a \cdot rs_2 \wedge a = (\mathbf{x}, r, \mathcal{I}' \uplus r_2, -, \mathcal{W}) \wedge \text{ccconf}(rs'_2) \wedge a \cdot rs_2 \xrightarrow{\mathcal{W}}^* rs'_2$ 
34   }
35   // Apply Lemma 6.4.
36   {
37      $\lceil \triangleright P \rceil * \exists rs_1, rs_2, rs'_2, a. r_2 \xRightarrow{1/2} P * [\text{change}(\mathbf{x}, r_2)]_1^{r_1} *$ 
38      $\text{chainds}(rs) * \text{chainres}(rs_1) * \text{resource}(\mathcal{I}', \emptyset) * \text{chainres}(rs'_2) * \text{unused}(r_1, rs) \wedge$ 
39      $\text{wf}(rs) \wedge rs = rs_1 \cdot a \cdot rs_2 \wedge a = (\mathbf{x}, r, \mathcal{I}' \uplus r_2, -, \mathcal{W}) \wedge \text{ccconf}(rs'_2) \wedge a \cdot rs_2 \xrightarrow{\mathcal{W}}^* rs'_2$ 
40   }
41   // Close the region, use transitions sat and get.
42   {
43      $\exists r_1, r_2. \lceil \triangleright P \rceil *$ 
44      $\text{oreg}\left(r_1, \left\{ \text{Chain}(rs') \mid \begin{array}{l} \text{wf}(rs') \wedge rs = rs_1 \cdot a \cdot rs_2 \wedge a \cdot rs_2 \xrightarrow{\mathcal{W}}^* rs'_2 \wedge \\ \text{ccconf}(rs'_2) \wedge rs' = (rs_1 \cdot rs'_2) \triangleleft_x \triangleleft_{\mathcal{I}} (\bullet \setminus \{r_2\}) \end{array} \right\}\right)$ 
45   }
46   } // Use  $\lceil \triangleright P \rceil \implies \triangleright [P]$  and clean up garbage.
47   { $\triangleright [P]$ }

```

Fig. 16. Sketch-proof for wait with out-of-order signalling.

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.

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).

$$\begin{aligned} \text{LEMMA 6.3. } & \text{chainres}(rs) \wedge \text{wf}(rs) \wedge \text{ctrue}(rs) \\ & \sqsubseteq \exists rs'. \text{chainres}(rs') \wedge \text{cconf}(rs') \wedge rs \xrightarrow{\mathcal{W}^*} rs' \wedge \text{wf}(rs') \end{aligned}$$

PROOF. Given in Appendix B. \square

We apply this lemma on line 14 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:

$$\text{LEMMA 6.4. } \text{resource}(\mathcal{I} \uplus r_2, \emptyset) * r_2 \xrightarrow{1/2} P \sqsubseteq \text{resource}(\mathcal{I}, \emptyset) * [\triangleright P]$$

PROOF. Given in Appendix B \square

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 18 of the sketch-proof.

Proving extend. The sketch-proof is given in Fig. 17. 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.

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

PROOF. Each identifier r'' used in rs is associated with a fractional saved proposition $r'' \xrightarrow{1/2} P$. We case-split on the finite set of possible equalities and appeal to the linearity of saved propositions (Property 2). \square

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

LEMMA 6.6.

$$\begin{aligned} rs &= rs_1 \cdot (x, r_2, \mathcal{I}, 0, \mathcal{W}) \cdot rs_2 \wedge \text{unused}(r_1, rs) \wedge z, r', r'' \notin rs \\ &\sqsubseteq \text{unused}(r_1, rs_1 \cdot (x, r_2, \mathcal{I}, 0, \mathcal{W}) \cdot (z, r', \{r''\}, 0, \emptyset) \cdot rs_2) * [\text{send}(z)]_1^{r_1} * [\text{change}(z, r'')]_1^{r_1} \end{aligned}$$

PROOF. Appeal to the definition of **unused**. \square

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


```

1  { send(x, P) *  $\bigotimes_{e \in E} e \prec x$  *  $\bigotimes_{l \in L} x \prec l$  }
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 [send(x)] permission.
5    {  $\exists r_1, r_2, x', r_2 \xRightarrow{1/2} P * [\text{send}(x)]_1^{r_1} * z \mapsto \{\text{prev} = x'; \text{flag} = 0\} *$ 
       $\text{oreg}(r_1, \{\text{Chain}(rs) \mid \text{wf}(rs) \wedge rs(x) = (x, r_2, -, 0, -) \wedge (x, x') \in \text{seq}(rs) \wedge z \notin rs\})$  }
6    skip; skip; // Make a pair of stored propositions, appeal to Lemma 6.5.
7    {  $\exists r_1, r_2, x', r', r'', r_2 \xRightarrow{1/2} P * [\text{send}(x)]_1^{r_1} * z \mapsto \{\text{prev} = x'; \text{flag} = 0\} * r' \xRightarrow{1} Q * r'' \xRightarrow{1} Q *$ 
       $\text{oreg}(r_1, \{\text{Chain}(rs) \mid \text{wf}(rs) \wedge rs(x) = (x, r_2, -, 0, -) \wedge (x, x') \in \text{seq}(rs) \wedge z, r', r'' \notin rs\})$  }
8    // Create a node in the chain.
9    {  $\exists r_1, r_2, x', r', r'', r_2 \xRightarrow{1/2} P * [\text{send}(x)]_1^{r_1} *$ 
       $\text{oreg}(r_1, \{\text{Chain}(rs) \mid \text{wf}(rs) \wedge rs(x) = (x, r_2, -, 0, -) \wedge (x, x') \in \text{seq}(rs) \wedge z, r', r'' \notin rs\}) *$ 
       $r' \xRightarrow{1/2} Q * r'' \xRightarrow{1/2} Q * z \mapsto \{\text{prev} = x'; \text{flag} = 0\} * \text{resource}(\{r''\}, \{r'\})$  }
10   // Enter the region, split the chain up.
11   {  $\exists r_1, r_2, x', r', r'', h, rs_1, rs_2. r_2 \xRightarrow{1/2} P * [\text{send}(x)]_1^{r_1} *$ 
       $\triangleright \left( \text{chainds}(rs_1) * x \mapsto \{\text{prev} = x'; \text{flag} = 0\} * \text{chainds}(rs_2) * \text{chainres}(rs) * \right.$ 
       $\left. \text{unused}(r_1, rs) \wedge rs = rs_1 \cdot e \cdot rs_2 \wedge \text{wf}(rs) \wedge e = (x, r_2, \mathcal{I}, 0, \mathcal{W}) \wedge z, r', r'' \notin rs \right) *$ 
       $r' \xRightarrow{1/2} Q * r'' \xRightarrow{1/2} Q * z \mapsto \{\text{prev} = \text{hd}(rs_2).loc; \text{flag} = 0\} * \text{resource}(\{r''\}, \{r'\})$  }
12   x->prev = z; // Stitch the cnode into the chain using LPoints / AFrame.
13   {  $\exists r_1, r_2, x', r', r'', h, rs_1, rs_2. r_2 \xRightarrow{1/2} P * [\text{send}(x)]_1^{r_1} *$ 
       $\text{chainds}(rs_1) * x \mapsto \{\text{prev} = z; \text{flag} = 0\} * \text{chainds}(rs_2) * \text{chainres}(rs) *$ 
       $\text{unused}(r_1, rs) \wedge rs = rs_1 \cdot e \cdot rs_2 \wedge \text{wf}(rs) \wedge e = (x, r_2, \mathcal{I}, 0, \mathcal{W}) \wedge z, r', r'' \notin rs *$ 
       $r' \xRightarrow{1/2} Q * r'' \xRightarrow{1/2} Q * z \mapsto \{\text{prev} = \text{hd}(rs_2).loc; \text{flag} = 0\} * \text{resource}(\{r''\}, \{r'\})$  }
14   // Fold the chainds predicate back up.
15   {  $\exists r_1, r_2, x', r', r'', h, rs_1, rs_2. r_2 \xRightarrow{1/2} P * [\text{send}(x)]_1^{r_1} * r' \xRightarrow{1/2} Q * r'' \xRightarrow{1/2} Q$ 
       $\text{chainds}(rs_1 \cdot e \cdot e' \cdot rs_2) * \text{chainres}(rs_1 \cdot e \cdot e' \cdot rs_2) * \text{unused}(r_1, rs) \wedge$ 
       $rs = rs_1 \cdot e \cdot rs_2 \wedge \text{wf}(rs) \wedge e = (x, r_2, \mathcal{I}, 0, \mathcal{W}) \wedge e' = (z, r', \{r''\}, 0, \emptyset) \wedge z, r', r'' \notin rs$  }
16   // Get [send(z)] * [change(z, r'')] from unused(r1, rs) using Lemma 6.6.
17   {  $\exists r_1, r_2, r', r'', h, rs_1, rs_2.$ 
       $\text{chainds}(rs_1 \cdot e \cdot e' \cdot rs_2) * \text{chainres}(rs_1 \cdot e \cdot e' \cdot rs_2) * \text{unused}(r_1, rs_1 \cdot e \cdot e' \cdot rs_2)$ 
       $\wedge rs = rs_1 \cdot e \cdot rs_2 \wedge \text{wf}(rs) \wedge e = (x, r_2, \mathcal{I}, 0, \mathcal{W}) \wedge e' = (z, r', \{r''\}, 0, \emptyset)$ 
       $\wedge r_2 \xRightarrow{1/2} P * [\text{send}(x)]_1^{r_1} * r' \xRightarrow{1/2} Q * r'' \xRightarrow{1/2} Q * [\text{send}(z)]_1^{r_1} * [\text{change}(z, r'')]_1^{r_1}$  }
18   // Close the region using the ext transition.
19   {  $\exists r_1, r_2, r', r'', h.$ 
       $\text{oreg}(r_1, \{\text{Chain}(rs') \mid rs' = rs_1 \cdot (x, r_2, -, 0, -) \cdot (z, r', \mathcal{I}', 0, -) \cdot rs_2 \wedge r'' \in \mathcal{I}' \wedge \text{wf}(rs')\})$ 
       $\wedge r_2 \xRightarrow{1/2} P * [\text{send}(x)]_1^{r_1} * r' \xRightarrow{1/2} Q * r'' \xRightarrow{1/2} Q * [\text{send}(z)]_1^{r_1} * [\text{change}(z, r'')]_1^{r_1}$  }
20   return (z, x);
21 } // Frame on order predicates. Fold invariant into predicates.
22 { send(z, Q) * recv(z, Q) * send(x, P) * z < x *  $\bigotimes_{e \in E} e \prec z$  *  $\bigotimes_{l \in L} x \prec l$  }

```

Fig. 17. Sketch-proof of extend with out-of-order signalling.

is stable because we hold the **send** tokens on x and z , meaning these channels cannot be extended or renounced.

Proving newchan. Omitted: this proof is similar to **extend**, but simpler.

6.3. Proving Renunciation and Splitting Axioms

Renunciation. The axiom is defined as follows:

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

The sketch-proof is given in Fig. 18. The important steps are conjoining the three predicates to give a single stable view on the shared structure (line 2), supplying the renounced resource to the shared region (line 12), and closing the region to give a new **send** predicate (line 14).

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.

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

PROOF. Given in Appendix B. \square

We use this lemma on line 2, Fig. 18. The conjoined region that arises from this lemma (line 5) 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 **resource** predicate (line 12) we use the following lemma to show that the renunciation set \mathcal{W} is updated appropriately:

$$\begin{aligned} \text{LEMMA 6.8.} \quad & \text{resource}(\mathcal{I}, \mathcal{W} \uplus \{r\}) * r \xRightarrow{1/2} Q * r' \xRightarrow{1/2} P \\ & \sqsubseteq \exists r''. \text{resource}(\mathcal{I}, \mathcal{W} \uplus \{r', r''\}) * r'' \xRightarrow{1/2} (P * Q) \end{aligned}$$

PROOF. Given in Appendix B. \square

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 $P * Q$. Internally this corresponds to garbage collecting one saved proposition and creating another.

On line 14 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]** token is held.

Splitting. The axiom is defined as follows:

$$\{\text{recv}(a, P) * \llbracket P \rrbracket * (P_1 * P_2)\} \quad \langle \text{skip} \rangle \quad \{\text{recv}(a, P_1) * \text{recv}(a, P_2)\}$$

A sketch-proof is given in Fig. 19. The key step is splitting the promised resource set \mathcal{I} for a node (line 6). To do this, we use the following lemma.

$$\begin{aligned} \text{LEMMA 6.9.} \quad & r_2 \in rs(x). \mathcal{I} \wedge r_2 \xRightarrow{1/2} P * \llbracket P \rrbracket * (P_1 * P_2) * \text{chainres}(rs) \\ & \sqsubseteq \exists rs', r_3, r_4. r_3, r_4 \notin rs \wedge rs' = rs \blacktriangleleft_x \blacktriangleleft_{\mathcal{I}} (\bullet \setminus r_2) \uplus \{r_3, r_4\} \wedge \\ & \quad r_3 \xRightarrow{1/2} P_1 * r_4 \xRightarrow{1/2} P_2 * \text{chainres}(rs') \end{aligned}$$

$$\begin{array}{l}
1 \quad \{ \text{recv}(x, P) * \text{send}(y, Q) * x \prec y \} \\
2 \quad \langle \text{skip} \rangle; \text{ // Definition of } \text{recv}, \text{ send}, \text{ Lemma 6.7.} \\
3 \quad \left\{ \begin{array}{l} \exists r, r_2, r_3. \text{oreg}(r, \{ \text{Chain}(rs) \mid (y, x) \in \text{ord}(rs) \}) * \\ \left(\text{oreg}(r, \{ \text{Chain}(rs) \mid rs(x) = (_, \mathcal{I}, _, _) \wedge r_2 \in \mathcal{I} \wedge \text{wf}(rs) \}) \right) * \\ * r_2 \xRightarrow{1/2} P * [\text{change}(x, r_2)]_1^r \\ \left(\text{oreg}(r, \{ \text{Chain}(rs) \mid \text{wf}(rs) \wedge rs(y) = (r_3, _, 0, _) \}) \right) \\ * r_3 \xRightarrow{1/2} Q * [\text{send}(y)]_1^r \end{array} \right\} \\
4 \quad \text{ // Conjunction on regions. Stable because of the permissions held.} \\
5 \quad \left\{ \begin{array}{l} \exists r, r_2, r_3. r_2 \xRightarrow{1/2} P * [\text{change}(x, r_2)]_1^r * r_3 \xRightarrow{1/2} Q * [\text{send}(y)]_1^r * \\ \text{oreg} \left(r, \left\{ \text{Chain}(rs) \mid \begin{array}{l} rs = rs_1 \cdot e \cdot rs_2 \cdot e' \cdot rs_3 \wedge \text{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\} \\
6 \quad \langle \text{skip} \rangle; \text{ // Open the region, drop } \triangleright \text{ using } \text{AFrame.} \\
7 \quad \left\{ \begin{array}{l} \exists r, r_2, r_3. r_2 \xRightarrow{1/2} P * [\text{change}(x, r_2)]_1^r * r_3 \xRightarrow{1/2} Q * [\text{send}(y)]_1^r * \\ \text{chainds}(rs) * \text{chainres}(rs) * \text{unused}(r, rs) \wedge rs = rs_1 \cdot e \cdot rs_2 \cdot e' \cdot rs_3 \wedge \\ \text{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\} \\
8 \quad \text{ // Pull out node.} \\
9 \quad \left\{ \begin{array}{l} \exists r, r_2, r_3. r_2 \xRightarrow{1/2} P * [\text{change}(x, r_2)]_1^r * r_3 \xRightarrow{1/2} Q * [\text{send}(y)]_1^r * \\ \text{chainds}(rs) * \text{chainres}(rs_1) * \text{resource}(\mathcal{I}', \mathcal{W}' \uplus \{r_3\}) * \text{chainres}(rs_2 \cdot e' \cdot rs_3) * \text{unused}(r, rs) \wedge \\ rs = rs_1 \cdot e \cdot rs_2 \cdot e' \cdot rs_3 \wedge \text{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\} \\
10 \quad \text{ // Apply Lemma 6.8 to supply the resource} \\
11 \quad \left\{ \begin{array}{l} \exists r, r_2, r_3. [\text{change}(x, r_2)]_1^r * r_4 \xRightarrow{1/2} (P \multimap Q) * [\text{send}(y)]_1^r * \\ \text{chainds}(rs) * \text{chainres}(rs_1) * \text{resource}(\mathcal{I}', \mathcal{W}' \uplus \{r_2, r_4\}) * \text{chainres}(rs_2 \cdot e' \cdot rs_3) * \text{unused}(r, rs) \wedge \\ rs = rs_1 \cdot e \cdot rs_2 \cdot e' \cdot rs_3 \wedge \text{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\} \\
12 \quad \text{ // Close the chain.} \\
13 \quad \left\{ \begin{array}{l} \exists r, r_2, r_4. r_4 \xRightarrow{1/2} (P \multimap Q) * [\text{send}(y)]_1^r * \text{chainds}(rs') * \text{chainres}(rs') * \text{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 \\ \text{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\} \\
14 \quad \text{ // Close region using the } \text{renun} \text{ transition.} \\
15 \quad \left\{ \exists r, r_4. r_4 \xRightarrow{1/2} (P \multimap Q) * [\text{send}(y)]_1^r * \text{oreg}(r, \{ \text{Chain}(rs) \mid \text{wf}(rs) \wedge rs(y) = (r_4, _, 0, _) \}) \right\} \\
16 \quad \text{ // Definition of } \text{send.} \\
17 \quad \{ \text{send}(y, P \multimap Q) \}
\end{array}$$

Fig. 18. Proof of the renunciation axiom for out-of-order implementation.

```

1  {  $\text{recv}(x, P) * \llbracket [P] \multimap (P_1 * P_2) \rrbracket$  }
2  // Definition of  $\text{recv}$ .
3  {  $\exists r_1, r_2. \text{oreg}(r_1, \{\text{Chain}(rs) \mid rs(x) = (-, \mathcal{I}, -, -) \wedge r_2 \in \mathcal{I} \wedge \text{wf}(rs)\})$  }
   {  $* r_2 \xRightarrow{1/2} P * [\text{change}(x, r_2)]_1^{r_1} * \llbracket [P] \multimap (P_1 * P_2) \rrbracket$  }
4  // Enter the region.
5  {  $\exists r_1, r_2. r_2 \xRightarrow{1/2} P * [\text{change}(x, r_2)]_1^{r_1} * \llbracket [P] \multimap (P_1 * P_2) \rrbracket * \triangleright$ 
   {  $\left( \text{chainds}(rs) * \text{chainres}(rs_1) * \text{resource}(\mathcal{I} \uplus \{r_2\}, \mathcal{W}) * \text{chainres}(rs_2) \wedge \right.$ 
   {  $\left. rs = rs_1 \cdot (x, r, \mathcal{I} \uplus \{r_2\}, f, \mathcal{W}) \cdot rs_2 \wedge \text{unused}(r_1, rs) \right)$  } }
6  {  $\langle \text{skip} \rangle$  // AFrame to drop  $\triangleright$ , Lemma 6.9 to split the region.
7  {  $\exists r_1, r_2, r_3, r_4. r_3, r_4 \notin rs \wedge r_3 \xRightarrow{1/2} P_1 * r_4 \xRightarrow{1/2} P_2 * [\text{change}(x, r_2)]_1^{r_1} * \left\{ \begin{array}{l} \text{chainds}(rs) * \text{chainres}(rs_1) * \text{resource}(\mathcal{I} \uplus \{r_3, r_4\}, \mathcal{W}) * \text{chainres}(rs_2) * \\ \text{unused}(r_1, rs) \wedge rs = rs_1 \cdot (x, r, \mathcal{I} \uplus \{r_2\}, f, \mathcal{W}) \cdot rs_2 \end{array} \right\}$  }
8  // Close chain, pull  $[\text{change}]$  perms out of unused.
9  {  $\exists r_1, r_2, r_3, r_4. r_3 \xRightarrow{1/2} P_1 * r_4 \xRightarrow{1/2} P_2 * [\text{change}(x, r_3)]_1^{r_1} * [\text{change}(x, r_4)]_1^{r_1} * \left\{ \begin{array}{l} \text{chainds}(rs) * \text{chainres}(rs_1) * \text{resource}(\mathcal{I} \uplus \{r_3, r_4\}, \mathcal{W}) * \text{chainres}(rs_2) * \text{unused}(r_1, rs') \\ \wedge rs = rs_1 \cdot (x, r, \mathcal{I} \uplus \{r_2\}, f, \mathcal{W}) \cdot rs_2 \wedge rs' = rs_1 \cdot (x, r, \mathcal{I} \uplus \{r_3, r_4\}, f, \mathcal{W}) \cdot rs_2 \end{array} \right\}$  }
10 // Close region using the split transition.
11 {  $\exists r_1, r_2, r_3, r_4. r_3 \xRightarrow{1/2} P_1 * r_4 \xRightarrow{1/2} P_2 * [\text{change}(x, r_3)]_1^{r_1} * [\text{change}(x, r_4)]_1^{r_1} * \left\{ \begin{array}{l} \text{oreg}(r_1, \{\text{Chain}(rs) \mid rs(x) = (-, \mathcal{I}, -, -) \wedge r_3, r_4 \in \mathcal{I} \wedge \text{wf}(rs)\}) \end{array} \right\}$  }
12 // Definition of  $\text{recv}$ .
13 {  $\text{recv}(x, P_1) * \text{recv}(x, P_2)$  }

```

Fig. 19. 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. \square

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 20 shows the implementation source-code. The data-structure is organised as follows. Flags are stored in fixed-size arrays of boolean values. These arrays are organised into a tree, with arrays closer to the root summarising arrays closer to the leaves. The header of each non-root array stores the address of a parent flag up the tree (in the up field). If one of these non-leaf flags is set, it denotes that all the leaves below it are also set. Thus wait

```

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

```

Fig. 20. Channel implementation with out-of-order signalling and summarization.

can check these summary flags rather than the entire chain. An instance of the structure is illustrated in Fig. 21.

Signals can be sent out-of-order, as with our previous implementation. 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. 20, 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, a call to **signal** summarises its siblings. If a flag’s siblings are 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. If a summary flag is set, the algorithm can conclude that all the summarised leaf flags are also set.

The **wait** function exploits summaries to reduce the number of flags it must test. Rather than examining all the preceding elements in the whole chain, **wait** just examines summary flags that precede the current node. Because it only ever climbs the tree, the function avoids the cost of iterating over the entire chain. The function spin-waits on each preceding flag in

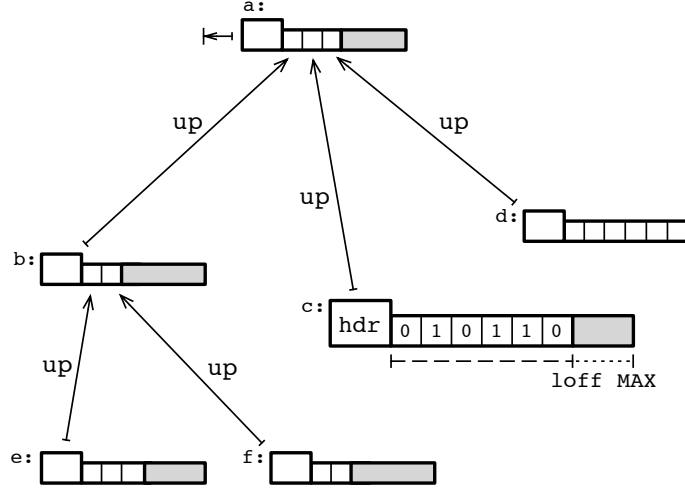
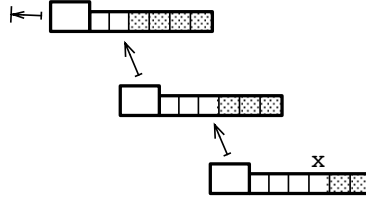


Fig. 21. Example of the summarising channel structure.

the current array (Fig. 20, lines 60–64 – note that increasing the index goes 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. The following diagram shows the nodes accessed when traversing up from `x`:



Extending the chain to insert a new link in a channel is a subtle process – array elements begin as leaves, but may become summary nodes. 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 existing and new channels are moved to the start of this array (lines 25–31).

7.2. Proof Strategy

At the abstract level, the behaviour of the implementation is unchanged from §6. Channels are arranged into a logical chain, with later channels waiting on earlier ones. Before acquiring resources using `wait`, all the preceding flags in the chain must be set using `signal`. The added summarization mechanism just changes the flags that must be examined to determine all flags in an interval are set. As a result, we can reuse much of the reasoning from §6.

Abstract state and basic predicates. The abstract state of a chain is once again an element of CNode^+ , identical to the one used in the previous section, except a 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 in the implementation). To represent the contents of the heap, we define a *heap map* d , consisting of a finite, partial function from physical addresses:

$$d: \text{Addr} \xrightarrow{\text{fin}} \{ \text{uhdr}: \text{Addr}; \text{uoff}: \text{Int}; \text{loff}: \text{Int}; \text{flags}: \{0..\text{MAX}\} \rightarrow \text{Bool} \}$$

A heap map records the `chan_hdr`-typed objects forming the tree. To simplify the representation, we flatten the `chan_hdr`-typed field `up` into the two fields `uhdr` and `uoff`. A location $\langle x, i \rangle$ is in d (written $\langle x, i \rangle \in d$) if $x \in \text{dom}(d) \wedge i \leq d.\text{loff}$.

For a given map d , `chainds` defines the corresponding data-structure. In the following definition, we write $x \mapsto v$ to indicate that x is immutable – shorthand for $\exists f. x \xrightarrow{f} v$.

$$\begin{aligned} \text{chainds}(d) &\triangleq \bigotimes_{x \in \text{dom}(d)} \cdot \\ x.\text{up} &\mapsto \{ \text{chan_hdr} = d(x).\text{uhdr}; \text{off} = d(x).\text{uoff} \} * x.\text{loff} \mapsto (d(x).\text{loff}) \\ &\quad * \bigotimes_{i \in \{0.. \text{MAX}\}} \cdot x.\text{flags}[i] \mapsto (d(x).\text{flags}(i)) \end{aligned}$$

In order to ensure that rs and d are well-formed and correspond correctly, we require several auxiliary notions:

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

Intuitively, the `child`, `descend`, and `leaf` predicates record the 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 `descend` is well-defined. `isset` and `allset` respectively assert that a single address and a whole array have their flags set.

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

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. `finalleaf` $_d(x, y)$ indicates that y is the right-most leaf according to $<_d$ 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 , so we generally use `finalleaf` as a partial function, i.e. `finalleaf` $_d(x)$ stands for the unique y such that `finalleaf` $_d(x, y)$.

Well-formedness ensures that rs and d are independently well-formed, and that they are correctly tied together into an inverted tree structure.

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

The address r is the location of the common tree root. The function τ records the distance from the current node to the tree root: this ensures that all paths up through the tree are finite. The first clause 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 .

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

The region state $\text{Chain}(rs, d)$ now takes as its argument both an abstract chain rs and a heap map d . Note that the **chainds** predicate now takes only a heap map – this captures all the structure it requires.

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

Because of this, the predicate **unused** takes a heap map d as an argument in addition to a chain rs . uS , the set of unused $[\text{send}(x)]$ permissions, remains as in the previous section.

Well-formedness ensures several properties of $<_d$. However, note that $<_d$ is not truly an order, as it is only transitive over leaf nodes.

$$\begin{aligned} \text{LEMMA 7.1. } \text{wf}(d) \implies \forall a, b \in d. & (a <_d^1 b \vee b < a) \wedge \neg(a <_d b <_d^1 a) \wedge \\ & (a <_d^1 b \wedge \text{leaf}_d(b) \implies a <_d b \vee \text{descend}_d(a, b)) \end{aligned}$$

PROOF. The first property follows from the existence of a single root node. For the second, as $a <_d b$, there must exist a common array where the ancestors of a and b are ordered. This contradicts both $\text{descend}_d(a, b)$ and $\text{descend}_d(b, a)$, and irreflexivity precludes $b <_d a$. For the third, either $a <_d b$, $\text{descend}_d(b, a)$ or $\text{descend}_d(a, b)$. All three cases trivially satisfy the conclusion. \square

Predicate definitions. As noted above, chain extension may result in calls to **wait** targeting summary nodes – before extension, these nodes would have been leaves. This is sound because these nodes will be set as part of summarization. For a given leaf l , the node that **wait** will be watching is the maximal node x such that $\text{finalleaf}(x, l)$ holds. Intuitively, this is because new arrays created through extension use the first element in the new array to stand for the extended chain element. The set $\text{C}(rs, d)$ represents these maximal nodes; a permission is missing from the set of unused **change** permissions, $\text{uC}(rs, d)$, only if it targets one of these nodes.

$$\begin{aligned} \text{C}(rs, d) &\triangleq \{ \langle y, i \rangle \mid r \in rs \wedge \text{finalleaf}_d(\langle y, i \rangle, r.loc) \wedge \neg \text{finalleaf}_d(\langle d(y).uhdr, d(y).uoff \rangle, r.loc) \} \\ \text{uC}(rs, d) &\triangleq \{ \langle x, r \rangle \mid r \in rs(\text{finalleaf}_d(x)).\mathcal{I} \wedge x \in \text{C}(rs, d) \wedge \neg \exists y. r \in rs(y).\mathcal{W} \} \\ \text{unused}(r, rs, d) &\triangleq (\otimes x \notin \text{uS}(rs). [\text{send}(x)]_1^{r_1}) * (\otimes (x, r') \notin \text{uC}(rs, d). [\text{change}(x, r')]_1^{r_1}) \end{aligned}$$

We can now define the **send** and **recv** predicates (we also use x to refer to the pair $\langle x.hdr, x.off \rangle$ when x is a **chan_addr** struct):

$$\begin{aligned} \text{oreg}(r, S) &\triangleq \text{region}(S, T_s, I_s(r), r) \\ \text{send}(x, P) &\triangleq \exists r_1, r_2. r_2 \xrightarrow{1/2} P * [\text{send}(x)]_1^{r_1} * [\text{mark}]_-^{r_1} * \\ &\quad \text{oreg}(r_1, \{ \text{Chain}(rs, d) \mid \text{wf}(rs, d) \wedge rs(x) = (r_2, -, 0, -) \}) \\ \text{recv}(x, P) &\triangleq \exists r_1, r_2. r_2 \xrightarrow{1/2} P * [\text{change}(x, r_2)]_1^{r_1} * \\ &\quad \text{oreg}(r_1, \{ \text{Chain}(rs, d) \mid \text{wf}(rs, d) \wedge r_2 \in rs(\text{finalleaf}_d(x)).\mathcal{I} \}) \\ x < y &\triangleq \exists r. \text{oreg}(r, \{ \text{Chain}(rs, d) \mid \text{wf}(rs, d) \wedge (y, x) \in \text{ord}(rs) \}) \end{aligned}$$

$$\begin{aligned}
\text{renun}_s(x, (rs, d), (rs', d')) &\triangleq d' = d \wedge \text{renun}_c(x, rs, rs') \\
\text{set}_s(x, (rs, d), (rs', d')) &\triangleq d' = d \triangleleft_h \triangleleft_{\text{flags}[o]} (1) \wedge \text{set}_c(\langle h, o \rangle, rs, rs') \\
\text{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.\text{flg} = 0 \wedge a.\text{loc} = x \wedge \\
&\quad \exists y. b = a[\text{loc} \mapsto y] \wedge c.\text{flg} = 0 \wedge c.\mathcal{W} = \emptyset \wedge \langle h, o \rangle = a.\text{loc} \wedge \\
&\quad \left(\begin{aligned} &b.\text{off} = \langle h', 0 \rangle \wedge c.\text{off} = \langle h', 1 \rangle \wedge \\ &d' = d \uplus h' \mapsto \{u\text{hdr} = h; u\text{off} = o; l\text{off} = 1; \text{flags} = \lambda_. 0\} \wedge \\ &\vee \left(\begin{aligned} &o = d(h).\text{loff} \wedge b.\text{loc} = a.\text{loc} \wedge c.\text{loc} = \langle h, o + 1 \rangle \wedge \\ &d' = (d \triangleleft_h \triangleleft_{\text{loff}} (o + 1)) \triangleleft_h \triangleleft_{\text{flags}[o+1]} (0) \end{aligned} \right) \end{aligned} \right) \\
\text{split}_s(x, r, (rs, d), (rs', d')) &\triangleq d = d' \wedge \text{split}_c(\text{finalleaf}_d(x), r, rs, rs') \\
\text{sat}_s(x, (rs, d), (rs', d')) &\triangleq d = d' \wedge \text{sat}_c(\text{finalleaf}_d(x), rs, rs') \\
\text{get}_s(x, r, (rs, d), (rs', d')) &\triangleq d = d' \wedge \text{get}_c(\text{finalleaf}_d(x), r, rs, rs') \\
T_s(\text{send}(x)) &\triangleq \{(a, b) \mid \text{wf}(b) \wedge (\text{renun}_s(x, a, b) \vee \text{set}_s(x, a, b) \vee \text{ext}_s(x, a, b))\} \\
T_s(\text{change}(x, r)) &\triangleq \{(a, b) \mid \text{wf}(b) \wedge (\text{split}_s(x, r, a, b) \vee \text{sat}_s(x, a, b) \vee \text{get}_s(x, r, a, b))\} \\
T_s(\text{mark}) &\triangleq \left\{ ((rs, d), (rs, d')) \mid \begin{aligned} &\text{wf}(rs, d') \wedge d' = d \triangleleft_h \triangleleft_{\text{flags}[o]} (1) \\ &\wedge \neg \text{leaf}_d(\langle h, o \rangle) \end{aligned} \right\}
\end{aligned}$$

Fig. 22. Definition of the transition relation T_s for the summarising implementation.

We define a new transition map T_s to capture changes to the heap map. The definition is given in Figure 22. Most of the transitions are inherited from T_c , the transition relation for the chained implementation. (We modify the type of T_c slightly by assuming it is defined on address-offset pairs, rather than just addresses). However, **set** and **extend** correspond to changes to the underlying heap, and they thus have to be extensively altered. Furthermore, **signal** can mark summary bits that are not leaves in the tree; this is allowed by the transition **mark**. We use $[\text{mark}]_\pi^r$ as notation for $\exists \pi. [\text{mark}]_\pi^r$ to represent non-exclusive ownership of the **mark** action. Finally, note **finalleaf** in the definition of **change**: this is needed because a leaf x may be converted by **extend** into a summary node representing $\text{finalleaf}_d(x)$.

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

LEMMA 7.2. $\text{recv}(x, P)$ is stable.

PROOF. Assume the initial state of the chain is $\text{Chain}(rs, d)$ and that **ext** takes the step $(x', (rs, d), (rs', d'))$. The case where $x \neq x'$ is simple, so assume $x = x'$. We now need to show $\text{wf}(rs', d') \wedge r_2 \in rs'(\text{finalleaf}_{d'}(x)).\mathcal{I}$.

The requirement $\text{wf}(rs', d')$ holds as a constraint on the transition relation. It remains to show the second clause. By the definition of **ext**, $\text{finalleaf}_d(x) = a.\text{loc}$ 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 $\text{finalleaf}_{d'}(x) = b.\text{loc}$. In the former case, **ext** adds a new array which must also descend from x . If $\text{finalleaf}_{d'}(x) \neq b.\text{loc}$, there must be leaf z such that $b.\text{loc} <_d z$, but the only new leaf c is at the next index in the new array, meaning $c.\text{loc} <_d b.\text{loc}$. Thus the result follows by contradiction. \square

7.3. Verifying wait, signal, extend

```

1  {send(x, P) * [P]}
2  signal(chan_addr x){
3    int i; ret = FALSE;
4    chan_addr a = x;
5    {a = x ∧ [P] * ∃r1, r2. r2  $\xRightarrow{1/2}$  P * [send(x)]1r1 *
   oreg(r1, {Chain(rs, d) | wf(rs, d) ∧ leafd(x) ∧ d(x.hdr).flags[x.off] = 0})}
6    while (a.hdr != NULL && !ret){
7      // Loop invariant.
8      {
9        (ret ∧ emp) ∨
10       {
11         (a = x ∧ [P] * ∃r1, r2. r2  $\xRightarrow{1/2}$  P * [send(x)]1r1 * [mark]-r1 *
12         oreg(r1, {Chain(rs, d) | wf(rs, d) ∧ leafd(x) ∧ d(x.hdr).flags[x.off] = 0})) ∨
13         (∃r1. [mark]-r1 *
14         oreg(r1, {Chain(rs, d) | wf(rs, d) ∧ ¬leafd(a) ∧ (∀l. (descendd(a, l) ∧ a ≠ l) ⇒ issetd(l))}))
15       }
16       a.hdr->flags[a.off] = 1; // Transition relation step set / mark.
17       for(i=0; i<=a.hdr->loff; i++){
18         {
19           (∃r1. [mark]-r1 *
20           oreg(r1, {Chain(rs, d) | wf(rs, d) ∧ (∀l. descendd(a, l) ⇒ issetd(l))
21           ∧ ∀0 ≤ n < i. issetd(⟨a.hdr, n⟩)}))
22         }
23         if (a.hdr->flags[i] != 1)
24           ret = TRUE;
25       }
26       {
27         (ret ∧ oreg(r1, {Chain(rs, d) | wf(rs, d) ∧ a ∈ dom(d)})) ∨
28         (∃r1. [mark]-r1 *
29         oreg(r1, {Chain(rs, d) | wf(rs, d) ∧ (∀l. descendd(a, l) ⇒ issetd(l)) ∧ allsetd(a.hdr)}))
30       }
31       a = a.hdr->up; // Abstract garbage collect.
32     }
33   }
34   {emp}

```

Fig. 23. Sketch-proof of signal with summarization.

Proving signal. A sketch-proof for **signal** is given in Fig. 23. 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 (§6.2) 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, so doing this does not change the abstract representation. 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.

LEMMA 7.3.

$$\begin{aligned} \text{set}_s(x, (rs, d), (rs', d')) \wedge \text{unused}(r, rs, d) * [\text{send}(x)]_1^r &\implies \text{unused}(r, rs', d') \\ ((rs, d), (rs', d')) \in T_s(\text{mark}) \wedge \text{unused}(r, rs, d) &\implies \text{unused}(r, rs', d') \end{aligned}$$

Proving wait. A sketch-proof for **wait** is given in Fig. 24 / 25. 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 $\text{last}_d(\mathbf{a})$ to stand for the last address in the array associated with \mathbf{a} , i.e. $\langle \mathbf{a}.\text{hdr}, d(\mathbf{a}.\text{hdr}).\text{loff} \rangle$.

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 39) we can conclude that all the flags in the subsequence rs_3 have been set. Then the algorithm increments the offset – as \mathbf{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, it has examined all the addresses greater than the starting address.

$$\text{LEMMA 7.4. } \text{wf}(d) \wedge d(\mathbf{a}).\text{loff} = o \wedge d(\mathbf{a}).\text{uhdr} = \text{NULL} \implies \neg \exists x. x <_d \langle \mathbf{a}, o \rangle$$

PROOF. Assume such an x exists. Then by the definition of $<_d$, there must exist an address y such that x and $\langle \mathbf{a}, o \rangle$ are both descended from y . As the uhdr field is **NULL**, the only possibility is that both addresses are in the object at \mathbf{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. \square

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.

$$\begin{aligned} \text{LEMMA 7.5. } \text{wf}(rs, d) \wedge z \in \text{dom}(rs) \wedge \text{descend}_d(x, z) \wedge \text{leaf}_d(z) \wedge \\ (\forall l. l <_d^{\parallel} x \wedge \text{leaf}_d(l) \implies \text{isset}_d(l)) \\ \implies \exists rs_1, rs_2. rs = rs_1 \cdot rs(z) \cdot rs_2 \wedge \text{ctrue}(rs(z) \cdot rs_2) \end{aligned}$$

PROOF. As $z \in \text{dom}(rs)$, we can easily divide up rs into $rs_1 \cdot rs(z) \cdot rs_2$. Now pick an arbitrary element y in $\text{dom}(rs(z) \cdot rs_2)$ and suppose that $rs(y).\text{flag}$ is not set. By well-formedness, it must be true that $y <_d^{\parallel} z$. Now we show that $y <_d^{\parallel} x$. 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. \square

```

1  {recv(x, P)}
2  wait(chan_addr x){
3    chan_addr a = x;
4    {
5       $\exists r_1, r_2. r_2 \xRightarrow{1/2} P * [\text{change}(x, r_2)]_1^{r_1} * a = x * \text{oreg}(r_1, \{\text{Chain}(rs, d) \mid \text{wf}(rs, d) \wedge r_2 \in rs(\text{finalleaf}_d(x)).\mathcal{I}\})\}$ 
6      while (a.hdr != NULL) {
7        {
8           $\exists r_1, r_2. r_2 \xRightarrow{1/2} P * [\text{change}(x, r_2)]_1^{r_1} * \text{oreg}\left(r_1, \left\{\text{Chain}(rs, d) \mid \begin{array}{l} \text{wf}(rs, d) \wedge r_2 \in rs(\text{finalleaf}_d(x)).\mathcal{I} \wedge \\ \forall l. a <_d l <_d^{\mathbb{I}} x \wedge \text{leaf}_d(l) \Rightarrow \text{isset}_d(l) \end{array} \right\}\right)$ 
9          for(skip; a.off <= a.hdr->loff; a.off++){
10            while(a.hdr->flags[a.off] != 1) skip;
11            {
12               $\exists r_1, r_2. r_2 \xRightarrow{1/2} P * [\text{change}(x, r_2)]_1^{r_1} * \text{oreg}\left(r_1, \left\{\text{Chain}(rs, d) \mid \begin{array}{l} \text{wf}(rs, d) \wedge r_2 \in rs(\text{finalleaf}_d(x)).\mathcal{I} \wedge \text{isset}_d(a) \wedge \\ \forall l. a <_d l <_d^{\mathbb{I}} x \wedge \text{leaf}_d(l) \Rightarrow \text{isset}_d(l) \end{array} \right\}\right)$ 
13              // Appeal to well-formedness.
14              {
15                 $\exists r_1, r_2. r_2 \xRightarrow{1/2} P * [\text{change}(x, r_2)]_1^{r_1} * \text{oreg}\left(r_1, \left\{\text{Chain}(rs, d) \mid \begin{array}{l} \text{wf}(rs, d) \wedge r_2 \in rs(\text{finalleaf}_d(x)).\mathcal{I} \wedge \\ \forall l. a <_d^{\mathbb{I}} l <_d^{\mathbb{I}} x \wedge \text{leaf}_d(l) \Rightarrow \text{isset}_d(l) \end{array} \right\}\right)$ 
16                // Apply Lemma 7.6 when incrementing a.off.
17              }
18              // Stable because chain cannot be extended once all the flags have been set.
19              {
20                 $\exists r_1, r_2. r_2 \xRightarrow{1/2} P * [\text{change}(x, r_2)]_1^{r_1} * \text{oreg}\left(r_1, \left\{\text{Chain}(rs, d) \mid \begin{array}{l} \text{wf}(rs, d) \wedge r_2 \in rs(\text{finalleaf}_d(x)).\mathcal{I} \wedge \\ \forall l. \text{last}_d(a) <_d^{\mathbb{I}} l <_d^{\mathbb{I}} x \wedge \text{leaf}_d(l) \Rightarrow \text{isset}_d(l) \end{array} \right\}\right)$ 
21                // Case-split on whether d(a.hdr).uhdr = NULL, if so, apply Lemma 7.4.
22                // As  $\forall x. \neg x <_d \text{last}_d(a)$ , well-formedness gives us  $\forall x. \text{last}_d(a) <_d^{\mathbb{I}} x$ .
23                {
24                   $\exists r_1, r_2. r_2 \xRightarrow{1/2} P * [\text{change}(x, r_2)]_1^{r_1} * \text{oreg}\left(r_1, \left\{\text{Chain}(rs, d) \mid \begin{array}{l} \text{wf}(rs, d) \wedge r_2 \in rs(\text{finalleaf}_d(x)).\mathcal{I} \wedge \\ \left( \begin{array}{l} d(a.hdr).uhdr \neq \text{NULL} \Rightarrow \\ \forall l. \text{last}_d(a) <_d^{\mathbb{I}} l <_d^{\mathbb{I}} x \wedge \text{leaf}_d(l) \Rightarrow \text{isset}_d(l) \\ \vee d(a.hdr).uhdr = \text{NULL} \Rightarrow \\ \forall l. l <_d^{\mathbb{I}} x \wedge \text{leaf}_d(l) \Rightarrow \text{isset}_d(l) \end{array} \right) \end{array} \right\}\right)$ 

```

Fig. 24. Sketch-proof of wait with summarization (completed in Fig. 25).

```

19   a = a.hdr->up;
    {
20     {
21       {
22         {
23           {
24             {
25               {
26                 {
27                   {
28                     {
29                       {

```

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

```

21   a.off++; // Apply Lemma 7.6 again.
    {
22     {
23       {
24         {
25           {
26             {
27               {
28                 {
29                   {

```

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

```

23   }
    {
24     {
25       {
26         {
27           {
28             {
29               {

```

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

```

25   // Apply Lemma 7.5.
    {
26     {
27       {
28         {
29           {

```

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

```

27   // Identical reasoning to chained implementation.
28   ...
29   }

```

Fig. 25. Sketch-proof of `wait` with summarization (continued from Fig. 24).

The final lemma shows that shifting left from the current maximal node reaches a node earlier in the order.

LEMMA 7.6. $0 \leq d(a).off < d(a).loff \wedge \text{wf}(d) \wedge a \blacktriangleleft_{off} (\bullet+1) <_d k <_d b \implies a <_d^{\mathbb{1}} k <_d b$

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

Proving extend. A sketch-proof of `extend` is given in Fig. 26. There are two cases for extending the chain: either the node is the last element in the current array and there is 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 pair of lemmas to show that the `unused` predicate representing unused permissions is preserved by extending the chain.

```

1  { send(x, P) *  $\bigotimes_{e \in E} e \prec x$  *  $\bigotimes_{l \in L} x \prec l$  }
2  extend(chan_addr x){
3    // Frame off order predicates and unfold send.
4    {  $\exists r_1, r_2. r_2 \xRightarrow{1/2} P * [\text{send}(x)]_1^{r_1} * [\text{mark}]_-^{r_1} * \text{oreg}(r_1, \{\text{Chain}(rs, d) \mid \text{wf}(rs, d) \wedge \text{leaf}_d(x) \wedge rs(x).flg = 0\})$  }
5    chan_addr nx, r;
6    if(x.off == x.hdr->loff && x.off < MAX){
7      {  $\exists r_1, r_2. r_2 \xRightarrow{1/2} P * [\text{mark}]_-^{r_1} * [\text{send}(x)]_1^{r_1} * r' \xRightarrow{1} Q * r'' \xRightarrow{1} Q * \text{oreg}\left(r_1, \left\{ \text{Chain}(rs, d) \mid \begin{array}{l} \text{wf}(rs, d) \wedge \text{leaf}_d(x) \wedge rs(x).flg = 0 \\ \wedge d(x.hdr).loff = x.off \wedge x.off < MAX \end{array} \right\} \right)$  }
8      x.hdr->flags[x.loff+1] = 0;
9      x.hdr->loff++; // Transition relation step ext.
10     {  $\exists r_1, r_2. r_2 \xRightarrow{1/2} P * r' \xRightarrow{1/2} Q * r'' \xRightarrow{1/2} Q * [\text{mark}]_-^{r_1} * [\text{send}(x)]_1^{r_1} * [\text{send}(\langle x.hdr, x.off + 1 \rangle)]_1^{r_1} * [\text{change}(\langle x.hdr, x.off + 1 \rangle, r'')]_1^{r_1} * \text{oreg}\left(r_1, \left\{ \text{Chain}(rs, d) \mid \begin{array}{l} rs = rs_1 \cdot rs(x) \cdot (\langle x.hdr, x.off + 1 \rangle, r', \{r''\}, 0, \emptyset) \cdot rs_2 \wedge \\ \text{wf}(rs, d) \wedge \text{leaf}_d(x) \wedge rs(x).flg = 0 \wedge \text{leaf}_d(\langle x.hdr, x.off + 1 \rangle) \\ \wedge d(x.hdr).loff = x.off + 1 \wedge x.off < MAX \end{array} \right\} \right)$  }
11     r.hdr = x.hdr; r.off = x.off + 1;
12     nx = x;
13     {  $\exists r_1, r_2. r_2 \xRightarrow{1/2} P * r' \xRightarrow{1/2} Q * r'' \xRightarrow{1/2} Q * [\text{mark}]_-^{r_1} * [\text{send}(nx)]_1^{r_1} * [\text{send}(r)]_1^{r_1} * [\text{change}(r, r'')]_1^{r_1} * \text{oreg}\left(r_1, \left\{ \text{Chain}(rs, d) \mid \begin{array}{l} rs = rs_1 \cdot rs(nx) \cdot (r, r', \{r''\}, 0, \emptyset) \cdot rs_2 \wedge \text{wf}(rs, d) \wedge \\ \text{leaf}_d(nx) \wedge rs(nx).flg = 0 \wedge \text{leaf}_d(r) \wedge rs(r).flg = 0 \end{array} \right\} \right)$  }
14   } else {
15     {  $\exists r_1, r_2. r_2 \xRightarrow{1/2} P * r' \xRightarrow{1} Q * r'' \xRightarrow{1} Q * [\text{mark}]_-^{r_1} * [\text{send}(x)]_1^{r_1} * \text{oreg}(r_1, \{\text{Chain}(rs, d) \mid \text{wf}(rs, d) \wedge \text{leaf}_d(x) \wedge rs(x).flg = 0\})$  }
16     nh = malloc(chan_hdr);
17     nh->up = x; nh->loff = 1;
18     // Transition relation step ext.
19     {  $\exists r_1, r_2. r_2 \xRightarrow{1/2} P * r' \xRightarrow{1/2} Q * r'' \xRightarrow{1/2} Q * [\text{mark}]_-^{r_1} * [\text{send}(\langle nh, 0 \rangle)]_1^{r_1} * [\text{send}(\langle nh, 1 \rangle)]_1^{r_1} * [\text{change}(\langle nh, 1 \rangle, r'')]_1^{r_1} * \text{oreg}\left(r_1, \left\{ \text{Chain}(rs, d) \mid \begin{array}{l} rs = rs_1 \cdot (rs(x) \triangleleft_{loc} \langle nh, 0 \rangle) \cdot (\langle nh, 1 \rangle, r', \{r''\}, 0, \emptyset) \cdot rs_2 \wedge \\ \text{wf}(rs, d) \wedge \text{leaf}_d(x) \wedge rs(x).flg = 0 \wedge \text{leaf}_d(\langle nh, 0 \rangle) \wedge \text{leaf}_d(\langle nh, 1 \rangle) \\ \wedge rs(\langle nh, 0 \rangle).flg = 0 \wedge rs(\langle nh, 1 \rangle).flg = 0 \end{array} \right\} \right)$  }
20     nx.hdr = nh; nx.off = 0;
21     r.hdr = nh; r.off = 1;
22   }
23   return (r, nx);
24 }
25 { send(r, Q) * rcv(r, Q) * send(nx, P) * r < nx *  $\bigotimes_{e \in E} e \prec r$  *  $\bigotimes_{l \in L} nx \prec l$  }

```

Fig. 26. Sketch-proof of extend with summarization.

LEMMA 7.7.

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

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

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

As $\langle x, i+1 \rangle \in rs'$, it holds immediately that $\langle x, i+1 \rangle \in \text{uS}(rs')$. As $\langle x, i+1 \rangle$ is a leaf, $\text{finalleaf}_{d'}(\langle x, i+1 \rangle) = \langle x, i+1 \rangle$. As it is not the first leaf in the array x , it cannot have a finalleaf parent, meaning it must be in $\text{C}(rs', d')$. Thus $\langle x, i+1 \rangle \in \text{uC}(rs', d')$. This suffices to show that $[\text{send}(\langle x, i+1 \rangle)]$ and $[\text{change}(\langle x, i+1 \rangle, r')]$ can be safely removed from unused . \square

LEMMA 7.8.

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

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 $[\text{send}(\langle l, 0 \rangle)]_1^r * [\text{send}(\langle l, 1 \rangle)]_1^r$ 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 $\text{C}(rs', d')$. However, $\langle l, 1 \rangle$ is not leftmost, and thus is in $\text{C}(rs', d')$. By the same argument used in the previous lemma, $[\text{change}(\langle l, 1 \rangle, r')]_1^r$ can be removed from the unused . \square

LEMMA 7.9. $\text{ext}(x, (rs, d), (rs', d')) \wedge \text{wf}(rs, d) \wedge \text{wf}(rs', d') \implies \text{uC}(rs, d) \subseteq \text{uC}(rs', d')$

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 $(x, i) \in \text{uC}(rs, d)$. Extending the chain can't stop x from satisfying finalleaf or make any node higher than x satisfy finalleaf . Therefore $x \in \text{C}(rs', d')$ 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 $r \in rs'(\text{finalleaf}_{d'}(x)). \mathcal{I}$. \square

7.4. Verifying 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 §6.3. The main difference is the new definition of unused . The renunciation case is straightforward, but we need to show that we can pull the appropriate change permissions out of the region. This is captured by the following lemma:

LEMMA 7.10.

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

PROOF. By the definition of split , $d = d'$, and element locations in rs are unchanged in rs' . Thus it holds that $\text{C}(rs, d) = \text{C}(rs', d')$ and $\text{finalleaf}_d(x) = \text{finalleaf}_{d'}(x)$. From the

definition of uC the available **change** permissions are controlled by the set $rs(\text{finalleaf}_a(x)).\mathcal{I}$. 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 ($\text{protocol} \cong \mathcal{P}(\dots \times \text{protocol})$) 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 $\text{box}(i, P, \pi)$ and $\text{fut}(i, P)$ 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, $\text{fut}(i, P)$ is generally not stable when P is instantiated with an assertion that includes a protocol assertion, because $\text{fut}(i, P)$ asserts the existence of a shared region whose protocol is defined in terms of P .

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 Rinetzkzy, 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

A saved proposition $r \models^\pi P$ 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 $\{1st, 2nd\}$ and a single token state tok , and invariant map I_{prp} and transition relation T_{prp} defined as follows:

$$\begin{aligned} I_{prp}(Q)(1st) &\triangleq \text{emp} \\ I_{prp}(Q)(2nd) &\triangleq Q \\ T_{prp}(tok) &\triangleq \{(1st, 2nd)\} \end{aligned}$$

We then define the saved proposition $r \models^\pi Q$ as follows:

$$r \models^\pi Q \triangleq \text{region}(\{1st, 2nd\}, T_{prp}, I_{prp}(Q), r) * [tok]_r^\pi$$

The fact that the representation transition state $1st$ 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].

$$\text{region}(S, T, I, r) * \text{region}(S', T', J, r) \implies (\triangleright I(s) \Rightarrow \triangleright J(s)) \quad (5)$$

$$\text{region}(S, T, I, r) * \text{region}(S', T', J, r) \implies (\triangleright I(s) \multimap \triangleright J(s)) \quad (6)$$

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).

$$\begin{aligned} &r \models^{\pi_1} P * r \models^{\pi_2} Q \\ \Rightarrow &r \models^{\pi_1} P * \text{region}(\{1st, 2nd\}, T_{prp}, I_{prp}(P), r) && \text{region is duplicable.} \\ &\quad * r \models^{\pi_2} Q * \text{region}(\{1st, 2nd\}, T_{prp}, I_{prp}(Q), r) \\ \Rightarrow &r \models^{\pi_1} P * r \models^{\pi_2} Q * (\triangleright(I_{prp}(P)(2nd)) \Rightarrow \triangleright(I_{prp}(Q)(2nd))) && \text{Property 5.} \\ \Rightarrow &r \models^{\pi_1} P * r \models^{\pi_2} Q * (\triangleright P \Rightarrow \triangleright Q) && \text{defn of } I_{prp} \end{aligned}$$

□

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:

$$r \models^{\pi_1} P * r \models^{\pi_2} Q \implies (\triangleright P) \multimap (\triangleright Q) \quad (7)$$

The proof of Property (4) then goes as follows:

$$\begin{aligned}
& r \xRightarrow{\pi_1} P * r \xRightarrow{\pi_2} Q * (X \multimap \triangleright(Q * Y)) * (P \multimap Z) \\
& \quad \text{SMono} \\
& \Rightarrow r \xRightarrow{\pi_1} P * r \xRightarrow{\pi_2} Q * (X \multimap \triangleright(Q * Y)) * \triangleright(P \multimap Z) \\
& \quad \text{LBin, assume } \triangleright \text{ distributes over } \multimap \\
& \Rightarrow r \xRightarrow{\pi_1} P * r \xRightarrow{\pi_2} Q * (X \multimap (\triangleright Q * \triangleright Y)) * ((\triangleright P) \multimap \triangleright Z) \\
& \quad \text{Assume property } \gamma \\
& \Rightarrow r \xRightarrow{\pi_1} P * r \xRightarrow{\pi_2} Q * (X \multimap (\triangleright Q * \triangleright Y)) * ((\triangleright P) \multimap \triangleright Z) * ((\triangleright Q) \multimap (\triangleright P)) \\
& \quad \text{Transitivity of } \multimap \\
& \Rightarrow r \xRightarrow{\pi_1} P * r \xRightarrow{\pi_2} Q * (X \multimap (\triangleright Q * \triangleright Y)) * ((\triangleright Q) \multimap \triangleright Z) \\
& \quad \text{Framing of } \multimap \\
& \Rightarrow r \xRightarrow{\pi_1} P * r \xRightarrow{\pi_2} Q * (X \multimap (\triangleright Q * \triangleright Y)) * ((\triangleright Q) * (\triangleright Y)) \multimap ((\triangleright Z) * (\triangleright Y)) \\
& \quad \text{Transitivity of } \multimap \\
& \Rightarrow r \xRightarrow{\pi_1} P * r \xRightarrow{\pi_2} Q * (X \multimap ((\triangleright Z) * (\triangleright Y)))
\end{aligned}$$

□

B. PROOFS FOR OUT-OF-ORDER SIGNALLING

This appendix gives proofs for some of the lemmas stated in §6.

LEMMA 6.1.

$$\text{resource}(\mathcal{I}, \emptyset) \sqsubseteq \bigotimes_{i \in \mathcal{I}}. \exists Q: \text{Prop}. i \xRightarrow{1/2} Q(i) * \lceil \triangleright Q(i) \rceil$$

PROOF.

$$\begin{aligned}
& \mathcal{W} = \emptyset \wedge \exists Q: \mathcal{I} \rightarrow \text{Prop}, R: \mathcal{W} \rightarrow \text{Prop}. \\
& \quad \bigotimes_{i \in \mathcal{I}}. i \xRightarrow{1/2} Q(i) * \bigotimes_{w \in \mathcal{W}}. w \xRightarrow{1/2} R(w) \\
& \quad * \lfloor (\triangleright \bigotimes_{w \in \mathcal{W}}. R(w)) \multimap \triangleright \bigotimes_{i \in \mathcal{I}}. \lceil Q(i) \rceil \rfloor \\
& \quad \text{Simplify using } \mathcal{W} = \emptyset, \text{ garbage collect.} \\
& \sqsubseteq \exists Q: \mathcal{I} \rightarrow \text{Prop}. \bigotimes_{i \in \mathcal{I}}. i \xRightarrow{1/2} Q(i) * \lfloor \triangleright \bigotimes_{i \in \mathcal{I}}. Q(i) \rfloor \\
& \quad \text{Switch from } \lfloor - \rfloor \text{ to } \lceil - \rceil, \text{ pull out } \bigotimes. \\
& \sqsubseteq \exists Q: \mathcal{I} \rightarrow \text{Prop}. \bigotimes_{i \in \mathcal{I}}. i \xRightarrow{1/2} Q(i) * \bigotimes_{i \in \mathcal{I}}. \lceil \triangleright Q(i) \rceil \\
& \quad \text{Push in the existential.} \\
& \sqsubseteq \bigotimes_{i \in \mathcal{I}}. \exists Q: \text{Prop}. i \xRightarrow{1/2} Q(i) * \lceil \triangleright Q(i) \rceil
\end{aligned}$$

□

$$\text{LEMMA 6.2. } r \xRightarrow{1/2} P * \lfloor P \rfloor * \text{resource}(\mathcal{I}, \mathcal{W} \uplus \{r\}) \sqsubseteq \text{resource}(\mathcal{I}, \mathcal{W})$$

PROOF.

$$r \xRightarrow{1/2} R * [R] * \text{resource}(\mathcal{I}, \mathcal{W} \uplus \{r\})$$

Definition of resource.

$$\sqsubseteq \quad r \xRightarrow{1/2} R * [R] * \left(\begin{array}{l} \exists Q: \mathcal{I} \rightarrow \text{Prop}, R: \mathcal{W} \uplus \{r\} \rightarrow \text{Prop}. \\ \otimes_{i \in \mathcal{I}} i \xRightarrow{1/2} Q(i) * \otimes_{w \in \mathcal{W} \uplus \{r\}} w \xRightarrow{1/2} R(w) \\ * \lfloor (\triangleright \otimes_{w \in \mathcal{W} \uplus \{r\}} R(w)) \multimap \triangleright \otimes_{i \in \mathcal{I}} [Q(i)] \rfloor \end{array} \right)$$

Property (3), monotonicity of \triangleright , monotonicity of $\lfloor - \rfloor$.

$$\sqsubseteq \quad \exists P: \text{Prop}. r \xRightarrow{1/2} R * [\triangleright R] * ([\triangleright R] \Rightarrow [\triangleright P]) \left(\begin{array}{l} \exists Q: \mathcal{I} \rightarrow \text{Prop}, R: \mathcal{W} \rightarrow \text{Prop}. \\ r \xRightarrow{1/2} P * \otimes_{i \in \mathcal{I}} i \xRightarrow{1/2} Q(i) * \otimes_{w \in \mathcal{W}} w \xRightarrow{1/2} R(w) \\ * \lfloor (\triangleright P * \triangleright \otimes_{w \in \mathcal{W}} R(w)) \multimap \triangleright \otimes_{i \in \mathcal{I}} [Q(i)] \rfloor \end{array} \right)$$

Modus ponens.

$$\sqsubseteq \quad \exists P: \text{Prop}. r \xRightarrow{1/2} R * [\triangleright P] * \left(\begin{array}{l} \exists Q: \mathcal{I} \rightarrow \text{Prop}, R: \mathcal{W} \rightarrow \text{Prop}. \\ r \xRightarrow{1/2} P * \otimes_{i \in \mathcal{I}} i \xRightarrow{1/2} Q(i) * \otimes_{w \in \mathcal{W}} w \xRightarrow{1/2} R(w) \\ * \lfloor (\triangleright P * \triangleright \otimes_{w \in \mathcal{W}} R(w)) \multimap \triangleright \otimes_{i \in \mathcal{I}} [Q(i)] \rfloor \end{array} \right)$$

Combine $\lfloor - \rfloor$, modus ponens for \multimap , garbage collect.

$$\sqsubseteq \quad \begin{array}{l} \exists Q: \mathcal{I} \rightarrow \text{Prop}, R: \mathcal{W} \rightarrow \text{Prop}. \\ \otimes_{i \in \mathcal{I}} i \xRightarrow{1/2} Q(i) * \otimes_{w \in \mathcal{W}} w \xRightarrow{1/2} R(w) \\ * \lfloor (\triangleright \otimes_{w \in \mathcal{W}} R(w)) \multimap \triangleright \otimes_{i \in \mathcal{I}} [Q(i)] \rfloor \\ \sqsubseteq \quad \text{resource}(\mathcal{I}, \mathcal{W}) \end{array}$$

□

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

$$\sqsubseteq \quad \exists rs'. \text{chainres}(rs') \wedge \text{cconf}(rs') \wedge rs \xrightarrow{\mathcal{W}^*} rs' \wedge \text{wf}(rs')$$

PROOF. We perform a sequence of smaller view-shifts corresponding to converting each *cnode* in turn, starting with the earliest element in the chain:

$$P_0 \sqsubseteq P_1 \sqsubseteq P_2 \sqsubseteq \dots \sqsubseteq P_n$$

Here n is the length of rs and the subscript $1, 2, 3 \dots$ 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:

$$\begin{aligned} P_i &\triangleq \quad \exists rs'. \text{chainres}(rs[0, n-i] \cdot rs') \wedge \text{cconf}(rs') \\ &\quad \wedge rs[(n-i) + 1, n] \xrightarrow{\mathcal{W}^*} rs' \wedge \text{wf}(rs[0, (n-i)] \cdot rs') \end{aligned}$$

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:

$$\begin{aligned}
& \exists rs'. \text{chainres}(rs[0, n-i] \cdot rs') \wedge \text{cconf}(rs') \wedge rs[(n-i)+1, n] \xrightarrow{\mathcal{W}^*} rs' \wedge \text{wf}(rs[0, n-i] \cdot rs') \\
& \sqsubseteq \\
& \exists rs', s. \text{chainres}(rs[0, n-(i+1)]) * \text{resource}(s.\mathcal{I}, s.\mathcal{W}) * \text{chainres}(rs') \\
& \wedge \text{cconf}(rs') \wedge rs[n-i, n] \xrightarrow{\mathcal{W}^*} s \cdot rs' \wedge \text{wf}(rs[0, n-(i+1)]) \cdot s \cdot rs'
\end{aligned}$$

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 $w \xrightarrow{1/2} R$ inside $\text{resource}(s.\mathcal{I}, s.\mathcal{W})$. By the same definition, the saved proposition $w \xrightarrow{1/2} P$ and $\text{resource} \triangleright P$ must be included in $\text{resource}(s'.\mathcal{I}, \emptyset)$. We move the resource from one node to the other, and garbage collect the saved proposition:

$$\begin{aligned}
& \text{resource}(s.\mathcal{I}, s.\mathcal{W}) * \text{chainres}(rs'_1) * \text{resource}(s'.\mathcal{I}, \emptyset) * \text{chainres}(rs'_2) \\
& \sqsubseteq \\
& \text{resource}(s.\mathcal{I}, s.\mathcal{W} \setminus \{w\}) * \text{chainres}(rs'_1) * \text{resource}(s'.\mathcal{I} \setminus \{w\}, \emptyset) * \text{chainres}(rs'_2)
\end{aligned}$$

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

$$rs[n-i, n] \xrightarrow{\mathcal{W}^*} s \cdot rs' \xrightarrow{\mathcal{W}} s[\mathcal{W} \mapsto s.\mathcal{W} \setminus \{w\}] \cdot rs''$$

Hence, by Lemma B.1 it follows that

$$\text{wf}(rs[0, n-(i+1)]) \cdot s[\mathcal{W} \mapsto s.\mathcal{W} \setminus \{w\}] \cdot rs''$$

The result is that the assertion is rewritten as follows:

$$\begin{aligned}
& \sqsubseteq \exists rs', s. \text{chainres}(rs[0, n-(i+1)]) * \text{resource}(s.\mathcal{I}, s.\mathcal{W} \setminus \{w\}) * \text{chainres}(rs') \\
& \wedge \text{cconf}(rs') \wedge rs[n-i, n] \xrightarrow{\mathcal{W}^*} s \cdot rs' \wedge \text{wf}(rs[0, n-(i+1)]) \cdot s \cdot rs'
\end{aligned}$$

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. \square

LEMMA B.1.

$$\text{wf}(rs) \wedge rs \xrightarrow{\mathcal{W}} rs' \Rightarrow (\text{available}(rs) = \text{available}(rs') \wedge \text{wf}(rs'))$$

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

$$\begin{aligned}
rs &= rs_1 \cdot s_1 \cdot rs_2 \cdot s_2 \cdot rs_3, & w &\in s_2.\mathcal{I}, & w &\in s_1.\mathcal{W}, \\
rs' &= rs_1 \cdot (s_1 \blacktriangleleft_{\mathcal{W}} (\bullet \setminus w)) \cdot rs_2 \cdot (s_2 \blacktriangleleft_{\mathcal{I}} (\bullet \setminus w)) \cdot rs_3
\end{aligned}$$

Since $s_2.\mathcal{W} \cap s_2.\mathcal{I} = \emptyset$ and $\text{wf}(s_2 \cdot rs_3)$ it follows that $w \notin (s_2.\mathcal{W} \cup \text{available}(rs_3))$ and $\forall s \in rs_2. w \notin s.\mathcal{I}$. Thus,

$$\begin{aligned}
\text{available}((s_2 \blacktriangleleft_{\mathcal{I}} (\bullet \setminus w)) \cdot rs_3) &= (\text{available}(rs_3) \setminus s_2.\mathcal{W}) \uplus (s_2.\mathcal{I} \setminus \{w\}) \\
&= \text{available}(s_2 \cdot rs_3) \setminus \{w\}
\end{aligned}$$

Since $w \in s_1 \cdot \mathcal{W}$ and $\text{wf}(s_1 \cdot rs_2 \cdot s_2 \cdot rs_3)$ it follows that $w \in \text{available}(rs_2 \cdot s_2 \cdot rs_3)$. Hence, since $\forall s \in rs_2. w \notin s \cdot \mathcal{I}$ it follows that $\forall s \in rs_2. w \notin s \cdot \mathcal{W}$. Thus,

$$\begin{aligned}
& \text{available}((s_1 \blacktriangleleft_{\mathcal{W}} (\bullet \setminus w)) \cdot rs_2 \cdot (s_2 \blacktriangleleft_{\mathcal{I}} (\bullet \setminus w)) \cdot rs_3) \\
&= (\text{available}(rs_2 \cdot (s_2 \blacktriangleleft_{\mathcal{I}} (\bullet \setminus w)) \cdot rs_3) \setminus (s_1 \cdot \mathcal{W} \setminus \{w\})) \uplus s_1 \cdot \mathcal{I} \\
&= ((\text{available}(rs_2 \cdot s_2 \cdot rs_3) \setminus \{w\}) \setminus (s_1 \cdot \mathcal{W} \setminus \{w\})) \uplus s_1 \cdot \mathcal{I} \\
&= (\text{available}(rs_2 \cdot s_2 \cdot rs_3) \setminus s_1 \cdot \mathcal{W}) \uplus s_1 \cdot \mathcal{I} \\
&= \text{available}(s_1 \cdot rs_2 \cdot s_2 \cdot rs_3)
\end{aligned}$$

from which it follows easily that $\text{available}(rs) = \text{available}(rs')$. To show that $\text{wf}(rs')$ we must also show that

$$\begin{aligned}
s_1 \cdot \mathcal{W} \setminus \{w\} &\subseteq \text{available}(rs_2 \cdot (s_2 \blacktriangleleft_{\mathcal{I}} (\bullet \setminus w)) \cdot rs_3) \\
&= \text{available}(rs_2 \cdot s_2 \cdot rs_3) \setminus \{w\}
\end{aligned}$$

and $s_2 \cdot \mathcal{W} \subseteq \text{available}(rs_3)$ both of which follow easily from the assumption that $\text{wf}(rs)$. 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. $\text{resource}(\mathcal{I} \uplus r_2, \emptyset) * r_2 \xRightarrow{1/2} P \sqsubseteq \text{resource}(\mathcal{I}, \emptyset) * \lceil \triangleright P \rceil$

PROOF.

$$\text{resource}(\mathcal{I} \uplus r_2, \emptyset) * r_2 \xRightarrow{1/2} P$$

Lemma 6.1.

$$\sqsubseteq r_2 \xRightarrow{1/2} P * \bigotimes_{i \in \mathcal{I} \uplus r_2} \exists R. i \xRightarrow{1/2} R * \lceil \triangleright R \rceil$$

Pull out resource for identifier r_2 .

$$\sqsubseteq r_2 \xRightarrow{1/2} P * \exists R'. r_2 \xRightarrow{1/2} R' * \lceil \triangleright R' \rceil * \bigotimes_{i \in \mathcal{I}} \exists R. i \xRightarrow{1/2} R * \lceil \triangleright R \rceil$$

Property (3), monotonicity of $\lceil - \rceil$.

$$\sqsubseteq r_2 \xRightarrow{1/2} P * \exists R'. (\lceil \triangleright R' \rceil \Rightarrow \lceil \triangleright P \rceil) * r_2 \xRightarrow{1/2} R' * \lceil \triangleright R' \rceil * \bigotimes_{i \in \mathcal{I}} \exists R. i \xRightarrow{1/2} R * \lceil \triangleright R \rceil$$

Modus ponens, garbage collect.

$$\sqsubseteq \lceil \triangleright P \rceil * \bigotimes_{i \in \mathcal{I}} \exists R. i \xRightarrow{1/2} R * \lceil \triangleright R \rceil$$

Definition of resource.

$$\sqsubseteq \lceil \triangleright P \rceil * \text{resource}(\mathcal{I}, \emptyset)$$

\square

LEMMA 6.7.

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

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 $x.\text{loc}$ which leads to a contradiction. Therefore the post-condition is **false**, allowing us to prove any post-condition. \square

LEMMA 6.8. $\text{resource}(\mathcal{I}, \mathcal{W} \uplus \{r\}) * r \xRightarrow{1/2} S * r' \xRightarrow{1/2} T_1$
 $\sqsubseteq \exists r''. \text{resource}(\mathcal{I}, \mathcal{W} \uplus \{r', r''\}) * r'' \xRightarrow{1/2} (T_1 \multimap S)$

PROOF. First we construct a new saved proposition r'' such that $r'' \Rightarrow (T_1 \multimap S)$. Now it suffices to prove

$$\begin{aligned} & \text{resource}(\mathcal{I}, \mathcal{W} \uplus \{r\}) * r \xRightarrow{1/2} S * r' \xRightarrow{1/2} T_1 * r'' \xRightarrow{1/2} T_2 \wedge \text{valid}(T_1 * T_2 \Rightarrow S) \\ & \sqsubseteq \text{resource}(\mathcal{I}, \mathcal{W} \uplus \{r', r''\}) \end{aligned}$$

$$\begin{aligned} & r \xRightarrow{1/2} S * r' \xRightarrow{1/2} T_1 * r'' \xRightarrow{1/2} T_2 \wedge \text{valid}(T_1 * T_2 \Rightarrow S) * \\ & \left(\begin{array}{l} \exists Q: \mathcal{I} \rightarrow \text{Prop}, R: \mathcal{W} \uplus \{r\} \rightarrow \text{Prop}. \\ \bigotimes_{i \in \mathcal{I}} i \xRightarrow{1/2} Q(i) * \bigotimes_{w \in \mathcal{W} \uplus \{r\}} w \xRightarrow{1/2} R(w) \\ * \lfloor \triangleright \bigotimes_{w \in \mathcal{W} \uplus \{r\}} R(w) \rfloor \multimap \triangleright \bigotimes_{i \in \mathcal{I}} \lceil Q(i) \rceil \end{array} \right) \\ & \text{Rearrange} \\ & \sqsubseteq \exists P: \text{Prop}, Q: \mathcal{I} \rightarrow \text{Prop}, R: \mathcal{W} \uplus \{r', r''\} \rightarrow \text{Prop}. \\ & \quad \text{valid}(R(r') * R(r'') \Rightarrow S) \wedge r \xRightarrow{1/2} S * r \xRightarrow{1/2} P \\ & \quad * \bigotimes_{i \in \mathcal{I}} i \xRightarrow{1/2} Q(i) * \bigotimes_{w \in \mathcal{W} \uplus \{r', r''\}} w \xRightarrow{1/2} R(w) \\ & \quad * \lfloor \triangleright P * \triangleright \bigotimes_{w \in \mathcal{W}} R(w) \rfloor \multimap \triangleright \bigotimes_{i \in \mathcal{I}} \lceil Q(i) \rceil \\ & \text{Property (3), SMono and distributing } \triangleright \text{ over } \implies . \\ & \sqsubseteq \exists P: \text{Prop}, Q: \mathcal{I} \rightarrow \text{Prop}, R: \mathcal{W} \uplus \{r', r''\} \rightarrow \text{Prop}. \\ & \quad \text{valid}((\triangleright R(r')) * (\triangleright R(r'')) \Rightarrow \triangleright P) \wedge \bigotimes_{i \in \mathcal{I}} i \xRightarrow{1/2} Q(i) * \bigotimes_{w \in \mathcal{W} \uplus \{r', r''\}} w \xRightarrow{1/2} R(w) \\ & \quad * \lfloor \triangleright P * \triangleright \bigotimes_{w \in \mathcal{W}} R(w) \rfloor \multimap \triangleright \bigotimes_{i \in \mathcal{I}} \lceil Q(i) \rceil \\ & \text{By mono of } \lfloor _ \rfloor, \multimap \text{ and } *. \\ & \sqsubseteq \exists Q: \mathcal{I} \rightarrow \text{Prop}, R: \mathcal{W} \uplus \{r', r''\} \rightarrow \text{Prop}. \\ & \quad \bigotimes_{i \in \mathcal{I}} i \xRightarrow{1/2} Q(i) * \bigotimes_{w \in \mathcal{W} \uplus \{r', r''\}} w \xRightarrow{1/2} R(w) \\ & \quad * \lfloor ((\triangleright R(r')) * (\triangleright R(r''))) * \triangleright \bigotimes_{w \in \mathcal{W}} R(w) \rfloor \multimap \triangleright \bigotimes_{i \in \mathcal{I}} \lceil Q(i) \rceil \\ & \text{Rearrange} \\ & \sqsubseteq \exists P: \text{Prop}, Q: \mathcal{I} \rightarrow \text{Prop}, R: \mathcal{W} \uplus \{r', r''\} \rightarrow \text{Prop}. \\ & \quad \bigotimes_{i \in \mathcal{I}} i \xRightarrow{1/2} Q(i) * \bigotimes_{w \in \mathcal{W} \uplus \{r', r''\}} w \xRightarrow{1/2} R(w) \\ & \quad * \lfloor \triangleright \bigotimes_{w \in \mathcal{W} \uplus \{r', r''\}} R(w) \rfloor \multimap \triangleright \bigotimes_{i \in \mathcal{I}} \lceil Q(i) \rceil \end{aligned}$$

□

LEMMA B.2.

$$\begin{aligned} & r_2 \xRightarrow{1/2} P * \lfloor \lceil P \rceil \multimap (P_1 * P_2) \rfloor * r_3 \xRightarrow{1/2} P_1 * r_4 \xRightarrow{1/2} P_2 * \text{resource}(\mathcal{I} \uplus \{r_2\}, \mathcal{W}) \\ & \sqsubseteq \text{resource}(\mathcal{I} \uplus \{r_3, r_4\}, \mathcal{W}) \end{aligned}$$

PROOF.

$$r_2 \xRightarrow{1/2} P * \llbracket [P] \multimap (P_1 * P_2) \rrbracket * r_3 \xRightarrow{1/2} P_1 * r_4 \xRightarrow{1/2} P_2 * \text{resource}(\mathcal{I} \uplus \{r_2\}, \mathcal{W})$$

Definition of resource.

$$\sqsubseteq r_2 \xRightarrow{1/2} P * \llbracket [P] \multimap (P_1 * P_2) \rrbracket * r_3 \xRightarrow{1/2} P_1 * r_4 \xRightarrow{1/2} P_2$$

$$\left(\begin{array}{l} \exists Q: \mathcal{I} \uplus \{r_2\} \rightarrow \text{Prop}, R: \mathcal{W} \rightarrow \text{Prop}. \\ \bigotimes_{i \in \mathcal{I} \uplus \{r_2\}} i \xRightarrow{1/2} Q(i) * \bigotimes_{w \in \mathcal{W}} w \xRightarrow{1/2} R(w) \\ * (\llbracket \triangleright \bigotimes_{w \in \mathcal{W}} R(w) \rrbracket \multimap \triangleright \bigotimes_{i \in \mathcal{I} \uplus \{r_2\}} \llbracket Q(i) \rrbracket \rrbracket) \end{array} \right)$$

Pull out r_2 , \triangleright mono w.r.t. \multimap , property (4).

$$\sqsubseteq r_2 \xRightarrow{1/2} P * r_3 \xRightarrow{1/2} P_1 * r_4 \xRightarrow{1/2} P_2$$

$$\left(\begin{array}{l} \exists Q: \mathcal{I} \uplus \{r_2\} \rightarrow \text{Prop}, R: \mathcal{W} \rightarrow \text{Prop}. \\ r_2 \xRightarrow{1/2} Q(r_2) * \bigotimes_{i \in \mathcal{I}} i \xRightarrow{1/2} Q(i) * \bigotimes_{w \in \mathcal{W}} w \xRightarrow{1/2} R(w) \\ * (\llbracket \triangleright \bigotimes_{w \in \mathcal{W}} R(w) \rrbracket \multimap \triangleright (P_1 * P_2 * \bigotimes_{i \in \mathcal{I}} \llbracket Q(i) \rrbracket) \rrbracket) \end{array} \right)$$

Fold r_3, r_4 into \mathcal{I} , weaken with $\llbracket - \rrbracket$, garbage collect.

$$\sqsubseteq \exists Q: \mathcal{I} \uplus r_2 \rightarrow \text{Prop}, R: \mathcal{W} \rightarrow \text{Prop}.$$

$$\bigotimes_{i \in \mathcal{I} \uplus \{r_3, r_4\}} i \xRightarrow{1/2} Q(i) * \bigotimes_{w \in \mathcal{W}} w \xRightarrow{1/2} R(w)$$

$$* (\llbracket \triangleright \bigotimes_{w \in \mathcal{W}} R(w) \rrbracket \multimap \triangleright \bigotimes_{i \in \mathcal{I} \uplus \{r_3, r_4\}} \llbracket Q(i) \rrbracket \rrbracket)$$

Definition of resource.

$$\sqsubseteq \text{resource}(\mathcal{I} \uplus \{r_3, r_4\}, \mathcal{W})$$

□

LEMMA 6.9. $r_2 \in rs(x). \mathcal{I} \wedge r_2 \xRightarrow{1/2} P * \llbracket [P] \multimap (P_1 * P_2) \rrbracket * \text{chainres}(rs)$

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

$$r_3 \xRightarrow{1/2} P_1 * r_4 \xRightarrow{1/2} P_2 * \text{chainres}(rs')$$

PROOF.

$$r_2 \in rs(x).\mathcal{I} \wedge r_2 \xRightarrow{1/2} P * \llbracket [P] \multimap (P_1 * P_2) \rrbracket * \text{chainres}(rs)$$

Make saved propositions, fresh by construction.

$$\sqsubseteq r_2 \in rs(x).\mathcal{I} \wedge r_2 \xRightarrow{1/2} P * \llbracket [P] \multimap (P_1 * P_2) \rrbracket * \text{chainres}(rs) * \\ \exists r_3, r_4. r_3 \xRightarrow{1} P_1 * r_4 \xRightarrow{1} P_2 \wedge r_3, r_4 \notin rs$$

Pull out resource predicate for x .

$$\sqsubseteq \exists rs_1, rs_2. r_2 \in rs(x).\mathcal{I} \wedge r_2 \xRightarrow{1/2} P * \llbracket [P] \multimap (P_1 * P_2) \rrbracket * \\ \text{chainres}(rs_1) * \text{resource}(rs(x).\mathcal{I}, rs(x).\mathcal{W} \uplus \{rs(x).res \mid rs(x).flg = 0\}) * \text{chainres}(rs_2) * \\ \exists r_3, r_4. r_3 \xRightarrow{1} P_1 * r_4 \xRightarrow{1} P_2 \wedge r_3, r_4 \notin rs \wedge rs = rs_1 \cdot rs(x) \cdot rs_2$$

Apply Lemma B.2.

$$\sqsubseteq \exists rs_1, rs_2, r_3, r_4. r_2 \in rs(x).\mathcal{I} \wedge r_2 \xRightarrow{1/2} P * \llbracket [P] \multimap (P_1 * P_2) \rrbracket * \\ \text{chainres}(rs_1) * \text{resource}((rs(x).\mathcal{I} \setminus r_2) \uplus \{r_3, r_4\}, rs(x).\mathcal{W} \uplus \{rs(x).res \mid rs(x).flg = 0\}) * \\ \text{chainres}(rs_2) * r_3 \xRightarrow{1/2} P_1 * r_4 \xRightarrow{1/2} P_2 \wedge r_3, r_4 \notin rs \wedge rs = rs_1 \cdot rs(x) \cdot rs_2$$

Definition of chainres.

$$\sqsubseteq \exists rs', r_3, r_4. r_3, r_4 \notin rs \wedge rs' = rs \blacktriangleleft_x \blacktriangleleft_{\mathcal{I}} (\bullet \setminus r_2) \uplus \{r_3, r_4\} \wedge \\ r_3 \xRightarrow{1/2} P_1 * r_4 \xRightarrow{1/2} P_2 * \text{chainres}(rs')$$

□