



Reasoning over Permissions Regions in Concurrent Separation Logic

James Brotherston^{1(✉)}, Diana Costa¹, Aquinas Hobor², and John Wickerson³

¹ University College London, London, UK

J.Brotherston@ucl.ac.uk

² National University of Singapore, Singapore, Singapore

³ Imperial College London, London, UK

Abstract. We propose an extension of separation logic with fractional permissions, aimed at reasoning about concurrent programs that share arbitrary *regions* or data structures in memory. In existing formalisms, such reasoning typically either fails or is subject to stringent side conditions on formulas (notably *precision*) that significantly impair automation. We suggest two formal syntactic additions that collectively remove the need for such side conditions: first, the use of both “weak” and “strong” forms of separating conjunction, and second, the use of nominal labels from hybrid logic. We contend that our suggested alterations bring formal reasoning with fractional permissions in separation logic considerably closer to common pen-and-paper intuition, while imposing only a modest bureaucratic overhead.

Keywords: Separation logic · Permissions · Concurrency · Verification

1 Introduction

Concurrent separation logic (CSL) is a version of separation logic designed to enable compositional reasoning about concurrent programs that manipulate memory possibly shared between threads [6, 26]. Like standard separation logic [28], CSL is based on *Hoare triples* $\{A\}C\{B\}$, where C is a program and A and B are formulas (called the *precondition* and *postcondition* of the code respectively). The heart of the formalism is the following *concurrency rule*:

$$\frac{\{A_1\}C_1\{B_1\} \quad \{A_2\}C_2\{B_2\}}{\{A_1 \otimes A_2\}C_1 \parallel C_2\{B_1 \otimes B_2\}}$$

where \otimes is a so-called *separating conjunction*. This rule says that if two threads C_1 and C_2 are run on spatially separated resources $A_1 \otimes A_2$ then the result will be the spatially separated result, $B_1 \otimes B_2$, of running the two threads individually.

However, since many or perhaps even most interesting concurrent programs do share some resources, \otimes typically does not denote strict disjoint separation of memories, as it does in standard separation logic (where it is usually written as $*$).

Instead, it usually denotes a weaker sort of “separation” designed to ensure that the two threads at least cannot interfere with each others’ data. This gives rise to the idea of *fractional permissions*, which allow us to divide writeable memory into multiple read-only copies by adding a permission value to each location in heap memory. In the usual model, due to Boyland [5], permissions are rational numbers in the half-open interval $(0, 1]$, with 1 denoting the write permission, and values in $(0, 1)$ denoting read-only permissions. We write the formula A^π , where π is a permission, to denote a “ π share” of the formula A . For example, $(x \mapsto a)^{0.5}$ (typically written as $x \xrightarrow{0.5} a$ for convenience) denotes a “half share” of a single heap cell, with address x and value a . The separating conjunction $A \otimes B$ then denotes heaps realising A and B that are “compatible”, rather than disjoint: where the heaps overlap, they must agree on the data value, and one adds the permissions at the overlapping locations [4]. E.g., at the logical level, we have the entailment:

$$x \xrightarrow{0.5} a \otimes x \xrightarrow{0.5} b \models a = b \wedge x \mapsto a. \quad (1)$$

Happily, the concurrency rule of CSL is still sound in this setting (see e.g. [29]).

However, the use of this weaker notion of separation \otimes causes complications for formal reasoning in separation logic, especially if one wishes to reason over arbitrary regions of memory rather than individual pointers. There are two particular difficulties, as identified by Le and Hobor [24]. The first is that, since \otimes denotes possibly-overlapping memories, one loses the main useful feature of separation logic: its nonambiguity about separation, which means that desirable entailments such as $A^{0.5} \otimes B^{0.5} \models (A \otimes B)^{0.5}$ turn out to be false. E.g.:

$$x \xrightarrow{0.5} a \otimes y \xrightarrow{0.5} b \not\models (x \mapsto a \otimes y \mapsto b)^{0.5}.$$

Here, the two “half-pointers” on the LHS might be aliased ($x = y$ and $a = b$), meaning they are two halves of the same pointer, whereas on the RHS they must be non-aliased (because we cannot combine two “whole” pointers). This ambiguity becomes quite annoying when one adds arbitrary predicate symbols to the logic, e.g. to support inductively defined data structures.

The second difficulty is that although recombining single pointers is straightforward, as indicated by Eq. (1), recombining the shares of arbitrary formulae is challenging. E.g., $A^{0.5} \otimes A^{0.5} \not\models A$, as shown by the counterexample

$$(x \mapsto 1 \vee y \mapsto 2)^{0.5} \otimes (x \mapsto 1 \vee y \mapsto 2)^{0.5} \not\models x \mapsto 1 \vee y \mapsto 2.$$

The LHS can be satisfied by a heap with a 0.5-share of x and a 0.5-share of y , whereas the RHS requires a full (1) share of either x or y .

Le et al. [24] address these problems by a combination of the use of *tree shares* (essentially Boolean binary trees) rather than rational numbers as permissions, and semantic restrictions on when the above sorts of permissions reasoning can be applied. For example, recombining permissions ($A^{0.5} \otimes A^{0.5} \models A$) is permitted only when the formula is *precise* in the usual separation logic sense (cf. [28]). The chief drawback with this approach is the need to repeatedly check these side

conditions on formulas when reasoning, as well as that said reasoning cannot be performed on imprecise formulas.

Instead, we propose to resolve these difficulties by a different, two-pronged extension to the syntax of the logic. First, we propose that the usual “strong” separating conjunction $*$, which enforces the strict disjointness of memory, *should be retained* in the formalism in addition to the weaker \otimes . The stronger $*$ supports entailments such as $A^{0.5} * B^{0.5} \models (A * B)^{0.5}$, which does not hold when \otimes is used instead. Second, we introduce *nominal labels* from hybrid logic (cf. [3, 10]) to remember that two copies of a formula have the same origin. We write a nominal α to denote a unique heap, in which case entailments such as $(\alpha \wedge A)^{0.5} \otimes (\alpha \wedge A)^{0.5} \models \alpha \wedge A$ become valid. We remark that labels have been adopted for similar “tracking” purposes in several other separation logic proof systems [10, 21, 23, 25].

The remainder of this paper aims to demonstrate that our proposed extensions are (i) weakly *necessary*, in that expected reasoning patterns fail under the usual formalism, (ii) *correct*, in that they recover the desired logical principles, and (iii) *sufficient* to verify typical concurrent programming patterns that use sharing. Section 2 gives some simple examples that motivate our extensions. Section 3 then formally introduces the syntax and semantics of our extended formalism. In Sect. 4 we show that our logic obeys the logical principles that enable us to reason smoothly with fractional permissions over arbitrary formulas, and in Sect. 5 we give some longer worked examples. Finally, in Sect. 6 we conclude and discuss directions for future work.

2 Motivating Examples

In this section, we aim to motivate our extensions to separation logic with permissions by showing, firstly, how the failures of the logical principles described in the introduction actually arise in program verification examples and, secondly, how these failures are remedied by our proposed changes.

The overall context of our work is reasoning about concurrent programs that share some data structure or region in memory, which can be described as a formula in the assertion language. If A is such a formula then we write A^π to denote a “ π share” of the formula A , meaning informally that all of the pointers in the heap memory satisfying A are owned with share π . The main question then becomes how this notion interacts with the separating conjunction \otimes . There are two key desirable logical equivalences:

$$(A \otimes B)^\pi \equiv A^\pi \otimes B^\pi \tag{I}$$

$$A^{\pi \oplus \sigma} \equiv A^\pi \otimes A^\sigma \tag{II}$$

Equivalence (I) describes distributing a fractional share over a separating conjunction, whereas equivalence (II) describes combining two pieces of a previously split resource. Both equivalences are true in the \models direction but, as we have seen in the Introduction, false in the \models one. Generally speaking, \otimes is like Humpty Dumpty: easy to break apart, but not so easy to put back together again.

The key to understanding the difficulty is the following equivalence:

$$x \overset{\pi}{\mapsto} a \circledast y \overset{\sigma}{\mapsto} b \equiv (x \overset{\pi}{\mapsto} a * y \overset{\sigma}{\mapsto} b) \vee (x = y \wedge a = b \wedge x \overset{\pi \oplus \sigma}{\mapsto} a)$$

In other words, either x and y are not aliased, or they *are* aliased and the permissions combine (the additive operation \oplus on rational shares is simply normal addition when the sum is ≤ 1 and undefined otherwise). This disjunction undermines the notational economies that have led to separation logic's great successes in scalable verification [11]; in particular, (I) fails because the left disjunct might be true, and (II) fails because the right disjunct might be. At a high level, \circledast is a bit too easy to introduce, and therefore also a bit too hard to eliminate.

2.1 Weak vs. Strong Separation and the Distribution Principle

One of the challenges of the weak separating conjunction \circledast is that it interacts poorly with inductively defined predicates. Consider porting the usual separation logic definition of a possibly-cyclic linked list segment from x to y from a sequential setting to a concurrent one by a simple substitution of \circledast for $*$:

$$\text{ls } x \ y \ =_{\text{def}} \ (x = y \wedge \text{emp}) \vee (\exists z. x \mapsto z \circledast \text{ls } z \ y).$$

Now consider a simple recursive procedure $\text{foo}(x, y)$ that traverses a linked list segment from x to y :

```
foo(x,y) { if x=y then return; else foo([x],y); }
```

It is easy to see that foo leaves the list segment unchanged, and therefore satisfies the following Hoare triple:

$$\{(\text{ls } x \ y)^{0.5}\} \text{foo}(x, y); \{(\text{ls } x \ y)^{0.5}\}.$$

The intuitive proof of this fact would run approximately as follows:

$$\begin{aligned} & \{(\text{ls } x \ y)^{0.5}\} \text{foo}(x, y) \{ \\ & \quad \text{if } x=y \text{ then return; } \{(\text{ls } x \ y)^{0.5}\} \\ & \quad \text{else} \quad \{x \neq y \wedge (x \mapsto z \circledast \text{ls } z \ y)^{0.5}\} \\ & \quad \quad \{x \overset{0.5}{\mapsto} z \circledast (\text{ls } z \ y)^{0.5}\} \\ & \quad \text{foo}([x], y); \quad \{x \overset{0.5}{\mapsto} z \circledast (\text{ls } z \ y)^{0.5}\} \\ & \quad \quad \{x \mapsto z \circledast \text{ls } z \ y\}^{0.5} \\ & \quad \quad \{(\text{ls } x \ y)^{0.5}\} \\ & \quad \} \quad \{(\text{ls } x \ y)^{0.5}\} \end{aligned}$$

However, because of the use of \circledast , the highlighted inference step is not sound:

$$x \overset{0.5}{\mapsto} z \circledast (\text{ls } z \ y)^{0.5} \not\equiv (x \mapsto z \circledast \text{ls } z \ y)^{0.5}. \quad (2)$$

To see this, consider a heap with the following structure, viewed in two ways:

$$x \overset{0.5}{\mapsto} z \otimes z \overset{0.5}{\mapsto} x \otimes x \overset{0.5}{\mapsto} z = x \mapsto z \otimes z \overset{0.5}{\mapsto} x$$

This heap satisfies the LHS of the entailment in (2), as it is the \otimes -composition of a 0.5-share of $x \mapsto z$ and a 0.5-share of $\text{ls } z z$, a cyclic list segment from z back to itself (note that here $z = y$). However, it does not satisfy the RHS, since it is not a 0.5-share of the \otimes -composition of $x \mapsto z$ with $\text{ls } z z$, which would require the pointer to be disjoint from the list segment.

The underlying reason for the failure of this example is that, in going from $(x \mapsto z \otimes \text{ls } z z)^{0.5}$ to $x \overset{0.5}{\mapsto} z \otimes (\text{ls } z z)^{0.5}$, we have lost the information that the pointer and the list segment are actually disjoint. This is reflected in the general failure of the distribution principle $A^\pi \otimes B^\pi \models (A \otimes B)^\pi$, of which the above is just one instance. Accordingly, our proposal is that the “strong” separating conjunction $*$ from standard separation logic, which forces disjointness of the heaps satisfying its conjuncts, should *also* be retained in the logic alongside \otimes , on the grounds that (II) *is* true for the stronger connective:

$$(A * B)^\pi \equiv A^\pi * B^\pi. \quad (3)$$

If we then define our list segments using $*$ in the traditional way, namely

$$\text{ls } x y \stackrel{\text{def}}{=} (x = y \wedge \text{emp}) \vee (\exists z. x \mapsto z * \text{ls } z y),$$

then we can observe that this second definition of ls is identical to the first on permission-free formulas, since \otimes and $*$ coincide in that case. However, when we replay the verification proof above with the new definition of ls , every \otimes in the proof above becomes a $*$, and the proof then becomes sound. Nevertheless, we can still use \otimes to describe permission-decomposition of list segments at a higher level; e.g., $\text{ls } x y$ can still be decomposed as $(\text{ls } x y)^{0.5} \otimes (\text{ls } x y)^{0.5}$.

2.2 Nominal Labelling and the Combination Principle

Unfortunately, even when we use the strong separating conjunction $*$ to define list segments ls , a further difficulty still remains. Consider a simple concurrent program that runs two copies of `foo` in parallel on the same list segment:

$$\text{foo}(x, y); \parallel \text{foo}(x, y);$$

Since `foo` only reads from its input list segment, and satisfies the specification $\{(\text{ls } x y)^{0.5}\} \text{foo}(x, y); \{(\text{ls } x y)^{0.5}\}$, this program satisfies the specification

$$\{\text{ls } x y\} \text{foo}(x, y); \parallel \text{foo}(x, y); \{\text{ls } x y\}.$$

Now consider constructing a proof of this specification in CSL. First we view the list segment $\text{ls } x y$ as the \otimes -composition of two read-only copies, with permission 0.5 each; then we use CSL’s concurrency rule (see Sect. 1) to compose the

specifications of the two threads; last we recombine the two read-only copies to obtain the original list segment. The proof diagram is as follows:

$$\begin{array}{c}
 \{\text{ls } xy\} \\
 \{(\text{ls } xy)^{0.5} \otimes (\text{ls } xy)^{0.5}\} \\
 \{\text{ls } xy\}^{0.5} \quad \parallel \quad \{\text{ls } xy\}^{0.5} \\
 \text{foo}(x, y); \quad \parallel \quad \text{foo}(x, y); \\
 \{\text{ls } xy\}^{0.5} \quad \parallel \quad \{\text{ls } xy\}^{0.5} \\
 \{(\text{ls } xy)^{0.5} \otimes (\text{ls } xy)^{0.5}\} \\
 \quad \quad \quad \xrightarrow{\text{red}} \{\text{ls } xy\}
 \end{array}$$

However, again, the highlighted inference step in this proof is not correct:

$$(\text{ls } xy)^{0.5} \otimes (\text{ls } xy)^{0.5} \not\models \text{ls } xy. \tag{4}$$

A countermodel is a heap with the following structure, again viewed in two ways:

$$(x \xrightarrow{0.5} y \otimes y \xrightarrow{0.5} y) \otimes x \xrightarrow{0.5} y = x \mapsto y \otimes y \xrightarrow{0.5} y$$

According to the first view of such a heap, it satisfies the LHS of (4), as it is the \otimes -composition of two 0.5-shares of $\text{ls } xy$ (one of two cells, and one of a single cell). However, it does not satisfy $\text{ls } xy$, since that would require every cell in the heap to be owned with permission 1.

Like in our previous example, the reason for the failure of this example is that we have lost information. In going from $\text{ls } xy$ to $(\text{ls } xy)^{0.5} \otimes (\text{ls } xy)^{0.5}$, we have forgotten that the two formulas $(\text{ls } xy)^{0.5}$ are in fact *copies of the same region*. For formulas A that are *precise* in that they uniquely describe part of any given heap [12], e.g. formulas $x \mapsto a$, this loss of information does not happen and we do have $A^{0.5} \otimes A^{0.5} \models A$; but for non-precise formulas such as $\text{ls } xy$, this principle fails.

However, we regard this primarily as a technical shortcoming of the formalism, rather than a failure of our intuition. It *ought* to be true that we can take any region of memory, split it into two read-only copies, and then later merge the two copies to re-obtain the original region. Were we conducting the above proof on pen and paper, we would very likely explain the difficulty away by adopting some kind of labelling convention, allowing us to remember that two formulas have been obtained from the same memory region by dividing permissions.

In fact, that is almost exactly our proposed remedy to the situation. We introduce *nominals*, or *labels*, from hybrid logic, where a nominal α is interpreted as denoting a unique heap. Any formula of the form $\alpha \wedge A$ is then precise (in the above sense), and so obeys the combination principle

$$(\alpha \wedge A)^\pi \otimes (\alpha \wedge A)^\sigma \models (\alpha \wedge A)^{\sigma \oplus \pi}, \tag{5}$$

where \oplus is addition on permissions. Thus we can repair the faulty CSL proof above by replacing every instance of the formula $\text{ls } xy$ by the “labelled” formula $\alpha \wedge \text{ls } xy$ (and adding an initial step in which we introduce the fresh label α).

2.3 The Jump Modality

However, this is not quite the end of the story. Readers may have noticed that replacing $\text{ls } xy$ by the “labelled” version $\alpha \wedge \text{ls } xy$ also entails establishing a slightly stronger specification for the function `foo`, namely:

$$\{(\alpha \wedge \text{ls } xy)^{0.5}\} \text{foo}(x, y); \{(\alpha \wedge \text{ls } xy)^{0.5}\}.$$

This introduces an extra difficulty in the proof (cf. Sect. 2.1); at the recursive call to `foo`($[x], y$), the precondition now becomes $\alpha^{0.5} \wedge (x \xrightarrow{0.5} z * (\text{ls } zy)^{0.5})$, which means that we cannot apply separation logic’s *frame rule* [32] to the pointer formula without first weakening away the label-share $\alpha^{0.5}$.

For this reason, we shall also employ hybrid logic’s “jump” modality $@_-$, where the formula $@_\alpha A$ means that A is true of the heap denoted by the label α . In the above, we can introduce labels β and γ for the list components $x \mapsto z$ and $\text{ls } zy$ respectively, whereby we can represent the decomposition of the list by the assertion $@_\alpha(\beta * \gamma)$. Since this is a *pure* assertion that does not depend on the heap, it can be safely maintained when applying the frame rule, and used after the function call to restore the label α , using the easily verifiable fact that

$$@_\alpha(\beta * \gamma) \wedge (\beta * \gamma) \models \alpha.$$

Similar reasoning over labelled decompositions of data structures is seemingly necessary whenever treating recursion; we return to it in more detail in Sect. 5.

3 Separation Logic with Labels and Permissions (SL_{LP})

Following the motivation given in the previous section, here we give the syntax and semantics of a separation logic, SL_{LP} , with permissions over arbitrary formulas, making use of both strong *and* weak separating conjunctions, and nominal labels (from hybrid logic [3, 10]). First, we define a suitable notion of permissions and associated operations.

Definition 3.1. *A permissions algebra is a tuple $(\text{Perm}, \oplus, \otimes, 1)$, where Perm is a set (of “permissions”), $1 \in \text{Perm}$ is called the write permission, and \oplus and \otimes are respectively partial and total binary functions on Perm , satisfying associativity, commutativity, cancellativity and the following additional axioms:*

$$\begin{array}{ll} \pi_1 \oplus \pi_2 \neq \pi_2 & (\text{non-zero}) \\ \forall \pi. \pi \oplus 1 \text{ is undefined} & (\text{top}) \\ \forall \pi. \exists \pi_1, \pi_2. \pi = \pi_1 \oplus \pi_2 & (\text{divisibility}) \\ (\pi_1 \oplus \pi_2) \otimes \pi = (\pi_1 \otimes \pi) \oplus (\pi_2 \otimes \pi) & (\text{left-dist}) \end{array}$$

The most common example of a permissions algebra is the Boyland fractional permission model $\langle (0, 1] \cap \mathbb{Q}, \oplus, \times, 1 \rangle$, where permissions are rational numbers in $(0, 1]$, \times is standard multiplication, and \oplus is standard addition but undefined if $p + p' > 1$. From now on, we assume a fixed but arbitrary permissions algebra.

With the permissions structure in place, we can now define the syntax of our logic. We assume disjoint, countably infinite sets \mathbf{Var} of variables, \mathbf{Pred} of predicate symbols (with associated arities) and \mathbf{Label} of labels.

Definition 3.2. *We define formulas of SL_{LP} by the grammar:*

$$\begin{aligned} A ::= & x = y \mid \neg A \mid A \wedge A \mid A \vee A \mid A \rightarrow A && (\text{pure}) \\ & \mid \mathbf{emp} \mid x \mapsto y \mid P(\mathbf{x}) \mid A * A \mid A \otimes A \mid A \multimap A \mid A \multimap^* A && (\text{spatial}) \\ & \mid A^\pi \mid \alpha \mid @_\alpha A && (\text{perms/labels}) \end{aligned}$$

where x, y range over \mathbf{Var} , π ranges over \mathbf{Perm} , P ranges over \mathbf{Pred} , α ranges over \mathbf{Label} and \mathbf{x} ranges over tuples of variables of length matching the arity of the predicate symbol P . We write $x \xrightarrow{\pi} y$ for $(x \mapsto y)^\pi$, and $x \neq y$ for $\neg(x = y)$.

The “magic wands” \multimap and \multimap^* are the implications adjoint to $*$ and \otimes , as usual in separation logic. We include them for completeness, but we use \multimap only for fairly complex examples (see Sect. 5.3) and in fact do not use \multimap^* at all.

Semantics. We interpret formulas in a standard model of stacks and heaps-with-permissions (cf. [4]), except that our models also incorporate a valuation of nominal labels. We assume an infinite set \mathbf{Val} of *values* of which an infinite subset $\mathbf{Loc} \subset \mathbf{Val}$ are considered addressable *locations*. A *stack* is as usual a map $s : \mathbf{Var} \rightarrow \mathbf{Val}$. A *heap-with-permissions*, which we call a *p-heap* for short, is a finite partial function $h : \mathbf{Loc} \rightarrow_{\text{fin}} \mathbf{Val} \times \mathbf{Perm}$ from locations to value-permission pairs. We write $\text{dom}(h)$ for the *domain* of h , i.e. the set of locations on which h is defined. Two p-heaps h_1 and h_2 are called *disjoint* if $\text{dom}(h_1) \cap \text{dom}(h_2) = \emptyset$, and *compatible* if, for all $\ell \in \text{dom}(h_1) \cap \text{dom}(h_2)$, we have $h_1(\ell) = (v, \pi_1)$ and $h_2(\ell) = (v, \pi_2)$ and $\pi_1 \oplus \pi_2$ is defined. (Thus, trivially, disjoint heaps are also compatible.) We define the multiplication $\pi \cdot h$ of a p-heap h by permission π by extending \otimes pointwise:

$$(\pi \cdot h)(\ell) = (v, \pi \otimes \pi') \Leftrightarrow h(\ell) = (v, \pi').$$

We also assume that each predicate symbol P of arity k is given a fixed interpretation $\llbracket P \rrbracket \in (\mathbf{Val}^k \times \mathbf{PHeaps})$, where \mathbf{PHeaps} is the set of all p-heaps. Here we allow an essentially free interpretation of predicate symbols, but they could also be given by a suitable inductive definition schema, as is done in many papers on separation logic (e.g. [7, 8]). Finally, a *valuation* is a function $\rho : \mathbf{Label} \rightarrow \mathbf{PHeaps}$ assigning a single p-heap $\rho(\alpha)$ to each label α .

Definition 3.3 (Strong and weak heap composition). *The strong composition $h_1 \circ h_2$ of two disjoint p-heaps h_1 and h_2 is defined as their union:*

$$(h_1 \circ h_2)(\ell) = \begin{cases} h_1(\ell) & \text{if } \ell \notin \text{dom}(h_2) \\ h_2(\ell) & \text{if } \ell \notin \text{dom}(h_1) \end{cases}$$

$$\begin{array}{ll}
s, h, \rho \models x = y & \Leftrightarrow s(x) = s(y) \\
s, h, \rho \models \neg A & \Leftrightarrow s, h, \rho \not\models A \\
s, h, \rho \models A \wedge B & \Leftrightarrow s, h, \rho \models A \text{ and } s, h, \rho \models B \\
s, h, \rho \models A \vee B & \Leftrightarrow s, h, \rho \models A \text{ or } s, h, \rho \models B \\
s, h, \rho \models A \rightarrow B & \Leftrightarrow s, h, \rho \models A \text{ implies } s, h, \rho \models B \\
s, h, \rho \models \text{emp} & \Leftrightarrow \text{dom}(h) = \emptyset \\
s, h, \rho \models x \mapsto y & \Leftrightarrow \text{dom}(h) = \{s(x)\} \text{ and } h(s(x)) = (s(y), 1) \\
s, h, \rho \models P(\mathbf{x}) & \Leftrightarrow (s(\mathbf{x}), h) \in \llbracket P \rrbracket \\
s, h, \rho \models A * B & \Leftrightarrow \exists h_1, h_2. h = h_1 \circ h_2 \text{ and } s, h_1, \rho \models A \text{ and } s, h_2, \rho \models B \\
s, h, \rho \models A \otimes B & \Leftrightarrow \exists h_1, h_2. h = h_1 \bar{\circ} h_2 \text{ and } s, h_1, \rho \models A \text{ and } s, h_2, \rho \models B \\
s, h, \rho \models A \multimap B & \Leftrightarrow \forall h'. \text{ if } h \circ h' \text{ defined and } s, h', \rho \models A \text{ then } s, h \circ h', \rho \models B \\
s, h, \rho \models A \multimap\!\!\!\multimap B & \Leftrightarrow \forall h'. \text{ if } h \bar{\circ} h' \text{ defined and } s, h', \rho \models A \text{ then } s, h \bar{\circ} h', \rho \models B \\
s, h, \rho \models A^\pi & \Leftrightarrow \exists h'. h = \pi \cdot h' \text{ and } s, h', \rho \models A \\
s, h, \rho \models \alpha & \Leftrightarrow h = \rho(\alpha) \\
s, h, \rho \models @_\alpha A & \Leftrightarrow s, \rho(\alpha), \rho \models A
\end{array}$$

Fig. 1. Definition of the satisfaction relation $s, h, \rho \models A$ for SL_{LP} .

If h_1 and h_2 are not disjoint then $h_1 \circ h_2$ is undefined.

The weak composition $h_1 \bar{\circ} h_2$ of two compatible p -heaps h_1 and h_2 is defined as their union, adding permissions at overlapping locations:

$$(h_1 \bar{\circ} h_2)(\ell) = \begin{cases} (v, \pi_1 \oplus \pi_2) & \text{if } h_1(\ell) = (v, \pi_1) \text{ and } h_2(\ell) = (v, \pi_2) \\ h_1(\ell) & \text{if } \ell \notin \text{dom}(h_2) \\ h_2(\ell) & \text{if } \ell \notin \text{dom}(h_1) \end{cases}$$

If h_1 and h_2 are not compatible then $h_1 \bar{\circ} h_2$ is undefined.

Definition 3.4. The satisfaction relation $s, h, \rho \models A$, where s is a stack, h a p -heap, ρ a valuation and A a formula, is defined by structural induction on A in Fig. 1. We write the entailment $A \models B$, where A and B are formulas, to mean that if $s, h, \rho \models A$ then $s, h, \rho \models B$. We write the equivalence $A \equiv B$ to mean that $A \models B$ and $B \models A$.

4 Logical Principles of SL_{LP}

In this section, we establish the main logical entailments and equivalences of SL_{LP} that capture the various interactions between the separating conjunctions \otimes and $*$, permissions and labels. As well as being of interest in their own right, many of these principles will be essential in treating the practical verification examples in Sect. 5. In particular, the permission distribution principle for $*$ (cf. (3), Sect. 2) is given in Lemma 4.3, and the permission combination principle for labelled formulas (cf. (5), Sect. 2) is given in Lemma 4.4.

Proposition 4.1. *The following equivalences all hold in SL_{LP} :*

$$\begin{array}{ll} A \otimes B \equiv B \otimes A & A * B \equiv B * A \\ A \otimes (B \otimes C) \equiv (A \otimes B) \otimes C & A * (B * C) \equiv (A * B) * C \\ A \otimes \text{emp} \equiv A & A * \text{emp} \equiv A \end{array}$$

Additionally, the following residuation laws hold:

$$A \models B \multimap C \Leftrightarrow A \otimes B \models C \quad \text{and} \quad A \models B \multimap C \Leftrightarrow A * B \models C.$$

In addition, we can always weaken $$ to \otimes : $A * B \models A \otimes B$.*

Next, we establish an additional connection between the two separating conjunctions \otimes and $*$.

Lemma 4.2 (\otimes / $*$ distribution). *For all formulas A, B, C and D ,*

$$(A \otimes B) * (C \otimes D) \models (A * C) \otimes (B * D). \tag{*/\otimes}$$

Proof. First we show a corresponding model-theoretic property: for any p-heaps h_1, h_2, h_3 and h_4 such that $(h_1 \bar{\circ} h_2) \circ (h_3 \bar{\circ} h_4)$ is defined,

$$(h_1 \bar{\circ} h_2) \circ (h_3 \bar{\circ} h_4) = (h_1 \circ h_3) \bar{\circ} (h_2 \circ h_4) \tag{6}$$

Since $(h_1 \bar{\circ} h_2) \circ (h_3 \bar{\circ} h_4)$ is defined by assumption, we have that $h_1 \bar{\circ} h_2$ and $h_3 \bar{\circ} h_4$ are disjoint and that h_1 and h_2 , as well as h_3 and h_4 are compatible. In particular, h_1 and h_3 are disjoint, so $h_1 \circ h_3$ is defined; the same reasoning applies to h_2 and h_4 . Moreover, since h_1 and h_2 are compatible, $h_1 \circ h_3$ and $h_2 \circ h_4$ must be compatible and so $(h_1 \circ h_3) \bar{\circ} (h_2 \circ h_4)$ is defined.

Now, writing h for $(h_1 \bar{\circ} h_2) \circ (h_3 \bar{\circ} h_4)$, and letting $\ell \in \text{dom}(h)$, we have

$$h(\ell) = \begin{cases} h_1(\ell) & \text{if } \ell \notin \text{dom}(h_3), \ell \notin \text{dom}(h_4) \text{ and } \ell \notin \text{dom}(h_2) \\ h_2(\ell) & \text{if } \ell \notin \text{dom}(h_3), \ell \notin \text{dom}(h_4) \text{ and } \ell \notin \text{dom}(h_1) \\ (v, \pi_1 \oplus \pi_2) & \text{if } \ell \notin \text{dom}(h_3), \ell \notin \text{dom}(h_4) \text{ and } h_1(\ell) = (v, \pi_1) \\ & \text{and } h_2(\ell) = (v, \pi_2) \\ h_3(\ell) & \text{if } \ell \notin \text{dom}(h_1), \ell \notin \text{dom}(h_2) \text{ and } \ell \notin \text{dom}(h_4) \\ h_4(\ell) & \text{if } \ell \notin \text{dom}(h_1), \ell \notin \text{dom}(h_2) \text{ and } \ell \notin \text{dom}(h_3) \\ (u, \pi_3 \oplus \pi_4) & \text{if } \ell \notin \text{dom}(h_1), \ell \notin \text{dom}(h_2) \text{ and } h_3(\ell) = (u, \pi_3) \\ & \text{and } h_4(\ell) = (u, \pi_4) \end{cases}$$

We can merge the first and fourth cases by noting that $h(\ell) = (h_1 \circ h_3)(\ell)$ if $\ell \notin \text{dom}(h_2 \circ h_4)$, and similarly for the second and fifth cases. We can also rewrite the last two cases by observing that $\ell \notin \text{dom}(h_3)$ implies $h_1(\ell) = (h_1 \circ h_3)(\ell)$, and so on, resulting in

$$\begin{aligned} h(\ell) &= \begin{cases} (h_1 \circ h_3)(\ell) & \text{if } \ell \notin \text{dom}(h_2 \circ h_4) \\ (h_2 \circ h_4)(\ell) & \text{if } \ell \notin \text{dom}(h_1 \circ h_3) \\ (w, \sigma_1 \oplus \sigma_2) & \text{if } (h_1 \circ h_3)(\ell) = (w, \sigma_1) \text{ and } (h_2 \circ h_4)(\ell) = (w, \sigma_2) \end{cases} \\ &= ((h_1 \circ h_3) \bar{\circ} (h_2 \circ h_4))(\ell). \end{aligned}$$

Now we show the main result. Suppose $s, h, \rho \models (A \otimes B) * (C \otimes D)$. This gives us $h = (h_1 \bar{\circ} h_2) \circ (h_3 \bar{\circ} h_4)$, where $s, h_1, \rho \models A$ and $s, h_2, \rho \models B$ and $s, h_3, \rho \models C$ and $s, h_4, \rho \models D$. By Eq. (6), we have $h = (h_1 \circ h_3) \bar{\circ} (h_2 \circ h_4)$, which gives us exactly that $s, h, \rho \models (A * C) \otimes (B * D)$, as required. \square

Next, we establish principles for distributing permissions over various connectives, in particular over the strong $*$, stated earlier as (3) in Sect. 2.

Lemma 4.3 (Permission distribution). *The following equivalences hold for all formulas A and B , and permissions π and σ :*

$$\begin{aligned} (A^\sigma)^\pi &\equiv A^{\sigma \otimes \pi} & (\otimes) \\ (A \vee B)^\pi &\equiv A^\pi \vee B^\pi & (\vee^\pi) \\ (A \wedge B)^\pi &\equiv A^\pi \wedge B^\pi & (\wedge^\pi) \\ (A * B)^\pi &\equiv A^\pi * B^\pi & (*^\pi) \end{aligned}$$

Proof. We just show the most interesting case, $(*^\pi)$. First of all, we establish a corresponding model-theoretic property: for any permission π and disjoint p-heaps h_1 and h_2 , meaning $h_1 \circ h_2$ is defined,

$$\pi \cdot (h_1 \circ h_2) = (\pi \cdot h_1) \circ (\pi \cdot h_2). \quad (7)$$

To see this, we first observe that for any $\ell \in \text{dom}(h_1 \circ h_2)$, we have that either $\ell \in \text{dom}(h_1)$ or $\ell \in \text{dom}(h_2)$. We just show the case $\ell \in \text{dom}(h_1)$, since the other is symmetric. Writing $h_1(\ell) = (v_1, \pi_1)$, and using the fact that $\ell \notin \text{dom}(h_2)$,

$$\pi \cdot (h_1 \circ h_2)(\ell) = (v_1, \pi \otimes \pi_1) = (\pi \cdot h_1)(\ell) = ((\pi \cdot h_1) \circ (\pi \cdot h_2))(\ell).$$

Now for the main result, let s, h and ρ be given. We have

$$\begin{aligned} s, h, \rho &\models (A * B)^\pi \\ \Leftrightarrow h &= \pi \cdot h' \text{ and } s, h', \rho \models A * B \\ \Leftrightarrow h &= \pi \cdot h' \text{ and } h' = h_1 \circ h_2 \text{ and } s, h_1, \rho \models A \text{ and } s, h_2, \rho \models B \\ \Leftrightarrow h &= \pi \cdot (h_1 \circ h_2) \text{ and } s, h_1, \rho \models A \text{ and } s, h_2, \rho \models B \\ \Leftrightarrow h &= (\pi \cdot h_1) \circ (\pi \cdot h_2) \text{ and } s, h_1, \rho \models A \text{ and } s, h_2, \rho \models B & \text{by (7)} \\ \Leftrightarrow h &= h'_1 \circ h'_2 \text{ and } s, h'_1, \rho \models A^\pi \text{ and } s, h'_2, \rho \models B^\pi \\ \Leftrightarrow s, h, \rho &\models A^\pi * B^\pi. \end{aligned} \quad \square$$

We now establish the main principles for dividing and combining permissions formulas using \otimes . As foreshadowed in Sect. 2, the combination principle holds only for formulas that are conjoined with a nominal label (cf. Eq. (5)).

Lemma 4.4 (Permission division and combination). *For all formulas A , nominals α , and permissions π_1, π_2 such that $\pi_1 \oplus \pi_2$ is defined:*

$$\begin{aligned} A^{\pi_1 \oplus \pi_2} &\models A^{\pi_1} \otimes A^{\pi_2} & (\text{Split } \otimes) \\ (\alpha \wedge A)^{\pi_1} \otimes (\alpha \wedge A)^{\pi_2} &\models (\alpha \wedge A)^{\pi_1 \oplus \pi_2} & (\text{Join } \otimes) \end{aligned}$$

Proof. **Case (Split \otimes):** Suppose that $s, h, \rho \models A^{\pi_1 \oplus \pi_2}$. We have $h = (\pi_1 \oplus \pi_2) \cdot h'$, where $s, h', \rho \models A$. That is, for any $\ell \in \text{dom}(h)$, we have $h'(\ell) = (v, \pi)$ say and, using the permissions algebra axiom (left-dist) from Definition 3.1,

$$h(\ell) = (v, (\pi_1 \oplus \pi_2) \otimes \pi) = (v, (\pi_1 \otimes \pi) \oplus (\pi_2 \otimes \pi)).$$

Now we define p-heaps h_1 and h_2 , both with domain exactly $\text{dom}(h)$, by

$$h_i(\ell) = (v, \pi_i \otimes \pi) \Leftrightarrow h'(\ell) = (v, \pi) \quad \text{for } i \in \{1, 2\}.$$

By construction, $h_1 = \pi_1 \cdot h'$ and $h_2 = \pi_2 \cdot h'$. Since $s, h', \rho \models A$, this gives us $s, h_1, \rho \models A^{\pi_1}$ and $s, h_2, \rho \models A^{\pi_2}$. Furthermore, also by construction, h_1 and h_2 are compatible, with $h = h_1 \bar{\circ} h_2$. Thus $s, h, \rho \models A^{\pi_1} \otimes A^{\pi_2}$, as required.

Case (Join \otimes): First of all, we show that for any p-heap h ,

$$(\pi_1 \cdot h) \bar{\circ} (\pi_2 \cdot h) = (\pi_1 \oplus \pi_2) \cdot h. \quad (8)$$

To see this, we observe that for any $\ell \in \text{dom}(h)$, writing $h(\ell) = (v, \pi)$ say,

$$\begin{aligned} & ((\pi_1 \oplus \pi_2) \cdot h)(\ell) \\ &= (v, (\pi_1 \oplus \pi_2) \otimes \pi) \\ &= (v, (\pi_1 \otimes \pi) \oplus (\pi_2 \otimes \pi)) \quad \text{by (left-dist)} \\ &= (h_1 \oplus h_2)(\ell) \text{ where } h_1(\ell) = (v, \pi_1 \otimes \pi) \text{ and } h_2 = (v, \pi_2 \otimes \pi) \\ &= ((\pi_1 \cdot h) \bar{\circ} (\pi_2 \cdot h))(\ell). \end{aligned}$$

Now, for the main result, suppose $s, h, \rho \models (\alpha \wedge A)^{\pi_1} \otimes (\alpha \wedge A)^{\pi_2}$. We have $h = h_1 \bar{\circ} h_2$ where $s, h_1, \rho \models (\alpha \wedge A)^{\pi_1}$ and $s, h_2, \rho \models (\alpha \wedge A)^{\pi_2}$. That is, $h = (\pi_1 \cdot h'_1) \bar{\circ} (\pi_2 \cdot h'_2)$, where $s, h'_1, \rho \models \alpha \wedge A$ and $s, h'_2, \rho \models \alpha \wedge A$. Thus $h'_1 = h'_2 = \rho(\alpha)$ and so, by (8), we have $h = (\pi_1 \oplus \pi_2) \cdot h'_1$, where $s, h'_1, \rho \models \alpha \wedge A$. This gives us $s, h, \rho \models (\alpha \wedge A)^{\pi_1 \oplus \pi_2}$, as required.

Lastly, we state some useful principles for labels and the “jump” modality.

Lemma 4.5 (Labelling and jump). *For all formulas A and labels α ,*

$$\begin{aligned} @_{\alpha} A \wedge \alpha^{\pi} &\models A^{\pi} && (@ \text{Elim}) \\ (\alpha \wedge A)^{\pi} &\models @_{\alpha} A && (@ \text{Intro}) \\ @_{\alpha} (\beta_1^{\pi} * \beta_2^{\sigma}) \wedge (\beta_1^{\pi} \otimes \beta_2^{\sigma}) &\models \alpha \wedge (\beta_1^{\pi} * \beta_2^{\sigma}) && (@ / * / \otimes) \end{aligned}$$

Proof. We just show the case ($@ / * / \otimes$), the others being easy. Suppose $s, h, \rho \models @_{\alpha} (\beta_1^{\pi} * \beta_2^{\sigma}) \wedge (\beta_1^{\pi} \otimes \beta_2^{\sigma})$, meaning that $s, \rho(\alpha), \rho \models \beta_1^{\pi} * \beta_2^{\sigma}$ and $s, h, \rho \models \beta_1^{\pi} \otimes \beta_2^{\sigma}$. Then we have $\rho(\alpha) = (\pi \cdot \rho(\beta_1)) \circ (\sigma \cdot \rho(\beta_2))$, while $h = (\pi \cdot \rho(\beta_1)) \bar{\circ} (\sigma \cdot \rho(\beta_2))$. Since \circ is defined only when its arguments are disjoint p-heaps, we obtain that $h = \rho(\alpha) = (\pi \cdot \rho(\beta_1)) \circ (\sigma \cdot \rho(\beta_2))$. Thus $s, h, \rho \models \alpha \wedge (\beta_1^{\pi} * \beta_2^{\sigma})$. \square

$$\begin{array}{c}
\frac{\{A_1\} C_1 \{B_1\} \quad \{A_2\} C_2 \{B_2\}}{\{A_1 \otimes A_2\} C_1 \parallel C_2 \{B_1 \otimes B_2\}} \text{ (Par) } \quad \frac{\{\alpha \wedge A\} C \{B\}}{\{A\} C \{B\}} \text{ (Label) } \\
\\
\frac{\{A\} C \{B\}}{\{A * F\} C \{B * F\}} \text{ (Frame *) } \quad \frac{\{A\} C \{B\}}{\{A \otimes F\} C \{B \otimes F\}} \text{ (Frame } \otimes \text{) } \\
\\
\text{(Par) } \text{ModVars}(C_2) \cap \text{FreeVars}(A_1, B_1) = \text{ModVars}(C_1) \cap \text{FreeVars}(A_2, B_2) = \emptyset \\
\text{(Label) } \alpha \text{ fresh} \quad \text{(Frame *) } \text{ModVars}(C) \cap \text{FreeVars}(F) = \emptyset \quad \text{(Frame } \otimes \text{) } \text{ see } \S 5.3
\end{array}$$

Fig. 2. The key CSL proof rules used in our examples; not shown are standard rules for consequence, conditionals, load/store, etc. The fresh-labelling rule (Label) and combination of both weak (Frame \otimes) and strong (Frame $*$) frame rules are novel to our approach. We require weak conjunction \otimes for the parallel rule (Par).

5 Concurrent Program Verification Examples

In this section, we demonstrate how SL_{LP} can be used in conjunction with the usual principles of CSL to construct verification proofs of concurrent programs, taking three examples of increasing complexity.

Our examples all operate on *binary trees* in memory, defined as usual in separation logic (again note the use of $*$ rather than \otimes):

$$\text{tree}(x) =_{\text{def}} (x = \text{null} \wedge \text{emp}) \vee (\exists d, l, r. x \mapsto (d, l, r) * \text{tree}(l) * \text{tree}(r)).$$

Our proofs employ (a subset of) the standard rules of CSL—with the most important being the concurrency rule from the Introduction, the separation logic *frame rules* for both $*$ and \otimes , and a new rule enabling us to introduce fresh labels into the precondition of a triple (similar to the way Hoare logic usually handles existential quantifiers). These key rules are shown in Fig. 2. We simplify our Hoare triple to remove elements to handle function call/return and furthermore omit the presentation of the standard collection of rules for consequence, load, store, if-then-else, assignment, etc.; readers interested in such aspects can consult [1]. Both of our frame rules have the usual side condition on modified program variables. The strong frame rule (Frame $*$) has an additional side condition that will be discussed in Sect. 5.3; until then it is trivially satisfied.

5.1 Parallel Read

Consider the following program:

```

check(x) {
  if (x == null) { return; }
  read(x); || read(x);
}

```

This is intended to be a straightforward example where we take a tree rooted at x and, if x is non-null, split into parallel threads that run the program `read` on x , and whose specification is $\{\alpha^\pi \wedge \text{tree}(x)^\sigma\} \text{read}(x) \{\alpha^\pi \wedge \text{tree}(x)^\sigma\}$. We prove that `check` satisfies the specification $\{\text{tree}(x)^\pi\} \text{check}(x) \{\text{tree}(x)^\pi\}$; the verification proof is in Fig. 3. The proof makes use of the basic operations of our theory: labelling, splitting and joining. The example follows precisely these steps, starting by labelling the formula $\text{tree}(x)^\pi \wedge x \neq \text{null}$ with α . The concurrency rule (Par) allows us to put formulas back together after the parallel call, and the two copies $(\alpha \wedge \text{tree}(x)^\pi)^{0.5}$ that were obtained are glued back together to yield $\text{tree}(x)^\pi$, since they have the same label.

$$\begin{array}{l}
\{\text{tree}(x)^\pi\} \\
\text{check}(x) \{ \\
\{(\text{tree}(x)^\pi \wedge x = \text{null}) \vee (\text{tree}(x)^\pi \wedge x \neq \text{null})\} \\
\text{if } (x == \text{null}) \{ \{x = \text{null} \wedge \text{tree}(x)^\pi\} \\
\text{return;} \\
\{\text{tree}(x)^\pi\} \\
\} \\
\{\alpha \wedge \text{tree}(x)^\pi \wedge x \neq \text{null}\} \qquad \text{by (Label)} \\
\{(\alpha \wedge \text{tree}(x)^\pi)^{0.5} \otimes (\alpha \wedge \text{tree}(x)^\pi)^{0.5}\} \qquad \text{by (Split } \otimes) \\
\left. \begin{array}{l}
\{(\alpha \wedge \text{tree}(x)^\pi)^{0.5}\} \\
\{\alpha^{0.5} \wedge \text{tree}(x)^{\pi \otimes 0.5}\} \\
\text{read}(x); \\
\{\alpha^{0.5} \wedge \text{tree}(x)^{\pi \otimes 0.5}\} \\
\{(\alpha \wedge \text{tree}(x)^\pi)^{0.5}\}
\end{array} \right\| \dots \qquad \text{by } (\wedge^\pi), (\otimes) \\
\{(\alpha \wedge \text{tree}(x)^\pi)^{0.5} \otimes (\alpha \wedge \text{tree}(x)^\pi)^{0.5}\} \qquad \text{by } (\wedge^\pi), (\otimes) \\
\{\alpha \wedge \text{tree}(x)^\pi\} \qquad \text{by (Join } \otimes) \\
\} \\
\{\text{tree}(x)^\pi\}
\end{array}$$

Fig. 3. Verification proof of program `check` in Example 5.1.

5.2 Parallel Tree Processing (Le and Hobor [24])

Consider the following program, which was also employed as an example in [24]:

```

proc(x) {
  if (x == null) { return; }
  print(x->d);    || print(x->d);
  proc(x->l);     || proc(x->l);
  proc(x->r);     || proc(x->r);
}

```

This code takes a tree rooted at x and, if x is non-null, splits into parallel threads that call `proc` recursively on its left and right

branches. We prove, in Fig. 4, that `proc` satisfies the specification $\{\alpha \wedge \text{tree}(x)^\pi\} \text{proc}(x) \{\alpha \wedge \text{tree}(x)^\pi\}$. First we unroll the definition of `tree`(x) and distribute the permission over Boolean connectives and $*$. If the tree is empty the process stops. Otherwise, we label each component with a new label and introduce the “jump” statement $@_\alpha(\beta_1 * \beta_2 * \beta_3)$, recording the decomposition of the tree into its three components. Since such statements are *pure*, i.e. independent of the heap, we can “carry” this formula along our computation without interfering with the frame rule(s). Now that every subregion is labelled, we split the formula into two copies, each with half share, but after distributing 0.5 over $*$ and \wedge we end up with half shares in the labels as well. We relabel each subregion with new “whole” labels, and again introduce pure $@$ -formulas that record the relation between the old and the new labels. At this moment we enter the parallel threads and recursively apply `proc` to the left and right subtrees of x . Assuming the specification of `proc` for subtrees of x , we then retrieve the original label α from the trail of crumbs left by the $@$ -formulas. We can then recombine the α -labelled threads using (`Join` \otimes) to arrive at the desired postcondition.

5.3 Cross-thread Data Transfer

Our previous examples involve only “isolated tank” concurrency: a program has some resources and splits them into parallel threads that do not communicate with each other before—remembering Humpty Dumpty!—ultimately re-merging. For our last example, we will show that our technique is expressive enough to handle more sophisticated kinds of sharing, in particular inter-thread coarse-grained communication. We will show that we can not only share read-only data, but in fact prove that one thread has acquired the full ownership of a structure, even when the associated root pointers are not easily exposed.

To do so, we add some communication primitives to our language, together with their associated Hoare rules. Coarse-grained concurrency such as locks, channels, and barriers have been well-investigated in various flavours of concurrent separation logic [19, 26, 31]. We will use a channel for our example in this section but with simplified rules: the Hoare rule for a channel c to send message number i whose message invariant is R_i^c is $\{R_i^c(x)\} \text{send}(c, x) \{\text{emp}\}$, while the corresponding rule to receive is $\{\text{emp}\} \text{receive}(c) \{\lambda \text{ret}. R_i^c(\text{ret})\}$. We ignore details such as identifying which party is allowed to send/receive at a given time [14] or the resource ownership of the channel itself [18].

These rules interact poorly with the strong frame rule from Fig. 2:

$$\frac{\{A\} C \{B\}}{\{A * F\} C \{B * F\}} \quad (\dagger, \ddagger) \text{ (Frame } *) \quad \begin{array}{l} (\dagger) \text{ ModVars}(C) \cap \text{FreeVars}(F) = \emptyset \\ (\ddagger) C \text{ does not receive resources} \end{array}$$

The revealed side condition (\ddagger) means that C does not contain any subcommands that “transfer in” resources, such as `unlock`, `receive`, etc.; this side condition is a bit stronger than necessary but has a simple definition and can be checked syntactically. Without (\ddagger) , we can reach a contradiction. Assume that the current

```

{ $\alpha \wedge \text{tree}(x)^\pi$ }
proc(x) {
  {( $\alpha \wedge (x = \text{null} \wedge \text{emp})^\pi$ )  $\vee$  ( $\alpha \wedge (x \mapsto (d, l, r) * \text{tree}(l) * \text{tree}(r))^\pi$ )} by ( $\wedge^\pi$ ), ( $\vee^\pi$ )
  {( $\alpha \wedge x = \text{null} \wedge \text{emp}$ )  $\vee$  ( $\alpha \wedge (x \mapsto (d, l, r) * \text{tree}(l) * \text{tree}(r))^\pi$ )}
  if (x == null) { { $\alpha \wedge x = \text{null} \wedge \text{emp}$ }
  return;
  { $\alpha \wedge (x = \text{null} \wedge \text{emp})^\pi$ }
  { $\alpha \wedge \text{tree}(x)^\pi$ }
}
{ $\alpha \wedge (x \xrightarrow{\pi} (d, l, r) * \text{tree}(l)^\pi * \text{tree}(r)^\pi)$ } by ( $*^\pi$ )
{(( $\beta_1 \wedge x \xrightarrow{\pi} (d, l, r)$ ) * ( $\beta_2 \wedge \text{tree}(l)^\pi$ ) * ( $\beta_3 \wedge \text{tree}(r)^\pi$ ))  $\wedge$ 
  @ $_\alpha$ ( $\beta_1 * \beta_2 * \beta_3$ )} by (Label), (@ Intro)
{(( $\beta_1^{0.5} \wedge x \xrightarrow{\pi \otimes 0.5} (d, l, r)$ ) * ( $\beta_2^{0.5} \wedge \text{tree}(l)^\pi \otimes 0.5$ ) *
  ( $\beta_3^{0.5} \wedge \text{tree}(r)^\pi \otimes 0.5$ ))  $\wedge$  (@ $_\alpha$ ( $\beta_1 * \beta_2 * \beta_3$ )) $^{0.5}$ ) $\otimes$ 
  (( $\beta_1^{0.5} \wedge x \xrightarrow{\pi \otimes 0.5} (d, l, r)$ ) * ( $\beta_2^{0.5} \wedge \text{tree}(l)^\pi \otimes 0.5$ ) *
  ( $\beta_3^{0.5} \wedge \text{tree}(r)^\pi \otimes 0.5$ ))  $\wedge$  (@ $_\alpha$ ( $\beta_1 * \beta_2 * \beta_3$ )) $^{0.5}$ )} by (Split  $\otimes$ ),
  ( $*^\pi$ ), ( $\wedge^\pi$ )
{((( $\gamma_1 \wedge x \xrightarrow{\pi \otimes 0.5} (d, l, r) \wedge @_{\gamma_1} \beta_1^{0.5}$ ) * ( $\gamma_2 \wedge \text{tree}(l)^\pi \otimes 0.5 \wedge @_{\gamma_2} \beta_2^{0.5}$ ) *
  ( $\gamma_3 \wedge \text{tree}(r)^\pi \otimes 0.5 \wedge @_{\gamma_3} \beta_3^{0.5}$ ))  $\wedge$  @ $_\alpha$ ( $\beta_1 * \beta_2 * \beta_3$ )) $\otimes$ 
  ((( $\gamma_4 \wedge x \xrightarrow{\pi \otimes 0.5} (d, l, r) \wedge @_{\gamma_4} \beta_1^{0.5}$ ) * ( $\gamma_5 \wedge \text{tree}(l)^\pi \otimes 0.5 \wedge @_{\gamma_5} \beta_2^{0.5}$ ) *
  ( $\gamma_6 \wedge \text{tree}(r)^\pi \otimes 0.5 \wedge @_{\gamma_6} \beta_3^{0.5}$ ))  $\wedge$  @ $_\alpha$ ( $\beta_1 * \beta_2 * \beta_3$ ))} by (Label),
  (@ Intro)
{(( $\gamma_1 \wedge x \xrightarrow{\pi \otimes 0.5} (d, l, r) \wedge @_{\gamma_1} \beta_1^{0.5}$ ) *
  ( $\gamma_2 \wedge \text{tree}(l)^\pi \otimes 0.5 \wedge @_{\gamma_2} \beta_2^{0.5}$ ) *
  ( $\gamma_3 \wedge \text{tree}(r)^\pi \otimes 0.5 \wedge @_{\gamma_3} \beta_3^{0.5}$ ))  $\wedge$  @ $_\alpha$ ( $\beta_1 * \beta_2 * \beta_3$ )}
print(x->d);
{(( $\gamma_1 \wedge x \xrightarrow{\pi \otimes 0.5} (d, l, r) \wedge @_{\gamma_1} \beta_1^{0.5}$ ) *
  ( $\gamma_2 \wedge \text{tree}(l)^\pi \otimes 0.5 \wedge @_{\gamma_2} \beta_2^{0.5}$ ) *
  ( $\gamma_3 \wedge \text{tree}(r)^\pi \otimes 0.5 \wedge @_{\gamma_3} \beta_3^{0.5}$ ))  $\wedge$  @ $_\alpha$ ( $\beta_1 * \beta_2 * \beta_3$ )}
proc(x->l);
{(( $\gamma_1 \wedge x \xrightarrow{\pi \otimes 0.5} (d, l, r) \wedge @_{\gamma_1} \beta_1^{0.5}$ ) *
  ( $\gamma_2 \wedge \text{tree}(l)^\pi \otimes 0.5 \wedge @_{\gamma_2} \beta_2^{0.5}$ ) *
  ( $\gamma_3 \wedge \text{tree}(r)^\pi \otimes 0.5 \wedge @_{\gamma_3} \beta_3^{0.5}$ ))  $\wedge$  @ $_\alpha$ ( $\beta_1 * \beta_2 * \beta_3$ )}
proc(x->r);
{(( $\gamma_1 \wedge x \xrightarrow{\pi \otimes 0.5} (d, l, r) \wedge @_{\gamma_1} \beta_1^{0.5}$ ) *
  ( $\gamma_2 \wedge \text{tree}(l)^\pi \otimes 0.5 \wedge @_{\gamma_2} \beta_2^{0.5}$ ) *
  ( $\gamma_3 \wedge \text{tree}(r)^\pi \otimes 0.5 \wedge @_{\gamma_3} \beta_3^{0.5}$ ))  $\wedge$  @ $_\alpha$ ( $\beta_1 * \beta_2 * \beta_3$ )}
{(( $\beta_1^{0.5} \wedge x \xrightarrow{\pi \otimes 0.5} (d, l, r)$ ) * ( $\beta_2^{0.5} \wedge \text{tree}(l)^\pi \otimes 0.5$ ) *
  ( $\beta_3^{0.5} \wedge \text{tree}(r)^\pi \otimes 0.5$ ))  $\wedge$  (@ $_\alpha$ ( $\beta_1 * \beta_2 * \beta_3$ )) $^{0.5}$ }
{((( $\beta_1 \wedge x \xrightarrow{\pi} (d, l, r)$ ) * ( $\beta_2 \wedge \text{tree}(l)^\pi$ ) *
  ( $\beta_3 \wedge \text{tree}(r)^\pi$ ))  $\wedge$  @ $_\alpha$ ( $\beta_1 * \beta_2 * \beta_3$ )) $^{0.5}$ } by ( $\wedge^\pi$ ), ( $*^\pi$ )
{ $\alpha \wedge (x \xrightarrow{\pi} (d, l, r) * \text{tree}(l)^\pi * \text{tree}(r)^\pi)$ } by (@ / * /  $\otimes$ )
{ $\alpha \wedge (x \xrightarrow{\pi} (d, l, r) * \text{tree}(l)^\pi * \text{tree}(r)^\pi)$ } $^{0.5}$ 
  ( $\alpha \wedge (x \xrightarrow{\pi} (d, l, r) * \text{tree}(l)^\pi * \text{tree}(r)^\pi)$ ) $^{0.5}$ } by (Par)
{ $\alpha \wedge (x \xrightarrow{\pi} (d, l, r) * \text{tree}(l)^\pi * \text{tree}(r)^\pi)$ } by (Join  $\otimes$ )
}
{ $\alpha \wedge \text{tree}(x)^\pi$ }

```

Fig. 4. Verification proof of Le and Hobor’s program from [24] in Example 5.2.

```

100 void transfer(int key) { { emp }
101                          rt* = make_tree();
102                          { tree(rt) }
    { (α ∧ tree(rt))0.5 } || { (α ∧ tree(rt))0.5 }
    tree* sub = find(rt, key) || ... ;
    send(ch, sub) || tree* sub = receive(ch) ;
    ... || modify(sub) ;
    receive(ch) || send(ch, ()) ;
    { (ε ∧ tree(rt))0.5 } || { (ε ∧ tree(rt))0.5 }
400                          { tree(rt) }
401                          delete_tree(rt); { emp } }

```

Fig. 5. Verification proof of the top and bottom of `transfer` in Example 5.3.

message invariant R_i^c is $x \overset{0.5}{\mapsto} a$, which has been sent by thread B. Now thread A, which had the other half of $x \overset{0.5}{\mapsto} a$, can reason as follows:

$$\frac{\{\text{emp}\} \text{receive}(c) \{x \overset{0.5}{\mapsto} a\}}{\{\text{emp} * x \overset{0.5}{\mapsto} a\} \text{receive}(c) \{x \overset{0.5}{\mapsto} a * x \overset{0.5}{\mapsto} a\}} \text{ (Frame } *), \text{ without } (\ddagger)$$

The postcondition is a contradiction as no location strongly separates from itself. However, given (\ddagger) the strong frame rule can be proven by induction.

The consequence of (\ddagger) , from a verification point of view, is that when resources are transferred in they arrive *weakly separated*, by \otimes , since we must use the weak frame rule around the receiving command. The troublesome issue is that this newly “arriving” state can thus \otimes -overlap awkwardly with the existing state. Fortunately, judicious use of labels can sort things out.

Consider the code in Fig. 5. The basic idea is simple: we create some data at the top (line 101) and then split its ownership 50-50 to two threads. The left thread finds a subtree, and passes its half of that subtree to the right via a channel. The right thread receives the root of that subtree, and thus has full ownership of that subtree along with half-ownership of the rest of the tree. Accordingly, the right thread can modify that subtree before notifying the left subtree and passing half of the modified subtree back. After merging, full ownership of the entire tree is restored and so on line 401 the program can delete it. Figure 5 only contains the proof and line numbers for the top and bottom shared portions. The left and the right thread’s proofs appear in Fig. 6.

By this point the top and bottom portions of the verification are straightforward. After creating the tree `tree(rt)` at line 102, we introduce the label α , split the formula using (`Split` \otimes), and then pass $(\alpha \wedge \text{tree}(rt))^{0.5}$ to both threads. After the parallel execution, due to the call to `modify(sub)` in the right thread, the tree has changed in memory. Accordingly, the label for the tree must also change as indicated by the $(\epsilon \wedge \text{tree}(rt))^{0.5}$ in both threads after parallel processing. These are then recombined on line 400 using the re-combination principle (`Join` \otimes), before the tree is deallocated via standard sequential techniques.

```

200 {  $(\alpha \wedge \text{tree}(\text{rt}))^{0.5}$  }
201 tree* sub = find(rt, key);
202 {  $(\alpha^{0.5} \wedge \text{tree}(\text{sub}) * (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt})))^{0.5}$  }
203 {  $(\alpha^{0.5} \wedge (\beta \wedge \text{tree}(\text{sub})) * (\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt}))))^{0.5}$  }
204 {  $\alpha^{0.5} \wedge ((\beta \wedge \text{tree}(\text{sub})) * (\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt}))))^{0.5} \wedge$ 
  {  $(\text{@}_{\alpha}^{0.5}((\beta \wedge \text{tree}(\text{sub})) * (\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt}))))^{0.5})$  }
205 {  $(\beta \wedge \text{tree}(\text{sub}))^{0.5} * (\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt})))^{0.5} \wedge$ 
  {  $(\text{@}_{\alpha}^{0.5}((\beta \wedge \text{tree}(\text{sub})) * (\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt}))))^{0.5})$  }
206 {  $(\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt})))^{0.5} \otimes (\beta \wedge \text{tree}(\text{sub}))^{0.5} \wedge$ 
  {  $(\text{@}_{\alpha}^{0.5}((\beta \wedge \text{tree}(\text{sub})) * (\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt}))))^{0.5})$  }
207 send(ch, sub);
208 {  $(\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt})))^{0.5}$  }
209 ...
210 {  $(\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt})))^{0.5}$  }
211 receive(ch);
212 {  $(\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt})))^{0.5} \otimes ((\text{@}_{\gamma}^{0.5}(\delta \wedge \text{tree}(\text{sub}) \multimap \epsilon \wedge \text{tree}(\text{rt}))^{0.5}) \wedge$ 
  {  $\gamma \perp \delta \wedge \delta \wedge \text{tree}(\text{sub})^{0.5})$  }
213 {  $\gamma \wedge (\delta \wedge \text{tree}(\text{sub}) \multimap \epsilon \wedge \text{tree}(\text{rt}))^{0.5} \otimes \delta \wedge \text{tree}(\text{sub})^{0.5} \wedge \gamma \perp \delta$  }
214 {  $(\delta \wedge \text{tree}(\text{sub}) \multimap \epsilon \wedge \text{tree}(\text{rt}))^{0.5} * \delta \wedge \text{tree}(\text{sub})^{0.5}$  }
215 {  $(\epsilon \wedge \text{tree}(\text{rt}))^{0.5}$  }

```

```

300 {  $(\alpha \wedge \text{tree}(\text{rt}))^{0.5}$  }
301 ...
302 {  $(\alpha \wedge \text{tree}(\text{rt}))^{0.5}$  }
303 tree* sub = receive(ch);
304 {  $(\alpha \wedge \text{tree}(\text{rt}))^{0.5} \otimes (\beta \wedge \text{tree}(\text{sub}))^{0.5} \wedge$ 
  {  $(\text{@}_{\alpha}^{0.5}((\beta \wedge \text{tree}(\text{sub})) * (\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt}))))^{0.5})$  }
305 {  $((\beta \wedge \text{tree}(\text{sub})) * (\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt}))))^{0.5} \otimes (\beta \wedge \text{tree}(\text{sub}))^{0.5}$  }
306 {  $((\beta \wedge \text{tree}(\text{sub}))^{0.5} \otimes (\beta \wedge \text{tree}(\text{sub}))^{0.5}) * (\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt})))^{0.5}$  }
307 {  $\text{tree}(\text{sub}) * (\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt})))^{0.5}$  }
308 modify(sub);
309 {  $\text{tree}(\text{sub}) * (\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt})))^{0.5}$  }
310 {  $(\delta \wedge \text{tree}(\text{sub})) * (\gamma \wedge ((\delta \wedge \text{tree}(\text{sub})) \multimap (\epsilon \wedge \text{tree}(\text{rt}))))^{0.5} \wedge \gamma \perp \delta$  }
311 {  $((\delta \wedge \text{tree}(\text{sub}))^{0.5} \otimes (\delta \wedge \text{tree}(\text{sub}))^{0.5}) * (\gamma \wedge ((\delta \wedge \text{tree}(\text{sub})) \multimap (\epsilon \wedge \text{tree}(\text{rt}))))^{0.5} \wedge$ 
  {  $\gamma \perp \delta \wedge (\text{@}_{\gamma}^{0.5}((\delta \wedge \text{tree}(\text{sub})) \multimap (\epsilon \wedge \text{tree}(\text{rt}))))^{0.5})$  }
312 {  $((\delta \wedge \text{tree}(\text{sub}))^{0.5} * (\gamma \wedge ((\delta \wedge \text{tree}(\text{sub})) \multimap (\epsilon \wedge \text{tree}(\text{rt}))))^{0.5}) \otimes$ 
  {  $(\delta \wedge \text{tree}(\text{sub}))^{0.5} \wedge \gamma \perp \delta \wedge (\text{@}_{\gamma}^{0.5}((\delta \wedge \text{tree}(\text{sub})) \multimap (\epsilon \wedge \text{tree}(\text{rt}))))^{0.5})$  }
313 {  $(\epsilon \wedge \text{tree}(\text{rt}))^{0.5} \otimes$ 
  {  $(\delta \wedge \text{tree}(\text{sub}))^{0.5} \wedge \gamma \perp \delta \wedge (\text{@}_{\gamma}^{0.5}((\delta \wedge \text{tree}(\text{sub})) \multimap (\epsilon \wedge \text{tree}(\text{rt}))))^{0.5})$  }
314 send(ch, ());
315 {  $(\epsilon \wedge \text{tree}(\text{rt}))^{0.5}$  }

```

Fig. 6. Verifications of the left (top) and right (bottom) threads of **transfer**.

Let us now examine the more interesting proofs of the individual threads in Fig. 6. Line 201 calls the **find** function, which searches a binary tree for a subtree rooted with key *key*. Following Cao *et al.* [13] we specify **find** as follows:

$$\{ \text{tree}(x)^\pi \} \text{find}(x) \{ \lambda \text{ret}. (\text{tree}(\text{ret}) * (\text{tree}(\text{ret}) \multimap \text{tree}(x)))^\pi \}$$

Here *ret* is bound to the return value of `find`, and the postcondition can be considered to represent the returned subtree $\text{tree}(\text{ret})$ separately from the tree-with-a-hole $\text{tree}(\text{ret}) \multimap \text{tree}(x)$, using a \multimap style to represent replacement as per Hobor and Villard [20]. This is the invariant on line 202.

Line 203 then attaches the fresh labels β and γ to the \ast -separated subparts, and line 204 snapshots the formula current at label α using the @ operator; $\text{@}_\alpha^\pi P$ should be read as “when one has a π -fraction of α , P holds”; it is definable using @ and an existential quantifier over labels. On line 205 we forget (in the left thread) the label α for the current heap for housekeeping purposes, and then on line 206 we weaken the strong separating conjunction \ast to the weak one @ before sending the root of the subtree `sub` on line 207.

In the `transfer` program, the invariant for the first channel message is

$$(\beta \wedge \text{tree}(\text{sub}))^{0.5} \wedge (\text{@}_\alpha^{0.5}((\beta \wedge \text{tree}(\text{sub})) \ast (\gamma \wedge (\text{tree}(\text{sub}) \multimap \text{tree}(\text{rt}))))^{0.5})$$

In other words, half of the ownership of the tree rooted at `sub` plus the (pure) @ -fact about the shape of the heap labeled by α . Comparing lines 206 and 208 we can see that this information has been shipped over the wire (the @ -information has been dropped since no longer needed). The left thread then continues to process until synchronizing again with the `receive` in line 211.

Before we consider the second synchronization, however, let us instead jump to the corresponding receive in the right thread at line 303. After the receive, the invariant on line 304 has the (weakly separated) resources sent from the left thread on line 206. We then “jump” label α using the @ -information to reach line 305. We can redistribute the β inside the \ast on line 306 since we already know that β and γ are disjoint. On line 307 we reach the payoff by combining both halves of the subtree `sub`, enabling the modification of the subtree in line 308.

On line 310 we label the two subheaps, and specialize the magic wand so that given the specific heap δ it will yield the specific heap ϵ ; we also record the pure fact that γ and δ are disjoint, written $\gamma \perp \delta$. On line 311 we snapshot γ and split the tree `sub` 50-50; then on line 312 we push half of `sub` out of the strong \ast . On line 313 we combine the subtree and the tree-with-hole to reach the final tree ϵ . We then send on line 314 with the channel’s second resource invariant:

$$(\delta \wedge \text{tree}(\text{sub}))^{0.5} \wedge \gamma \perp \delta \wedge (\text{@}_\gamma^{0.5}((\delta \wedge \text{tree}(\text{sub})) \multimap (\epsilon \wedge \text{tree}(\text{rt}))))^{0.5})$$

After the send, on line 315 we have reached the final fractional tree ϵ .

Back in the left-hand thread, the second send is received in line 211, leading to the weakly-separated postcondition in line 212. In line 213 we “jump” label γ , and then in line 214 we use the known disjointness of γ and δ to change the @ to \ast . Finally in line 215 we apply the magic wand to reach the postcondition.

6 Conclusions and Future Work

We propose an extension of separation logic with fractional permissions [4] in order to reason about sharing over arbitrary regions of memory. We identify two

fundamental logical principles that fail when the “weak” separating conjunction \otimes is used in place of the usual “strong” $*$, the first being distribution of permissions— $A^\pi \otimes B^\pi \not\models (A \otimes B)^\pi$ —and the second being the re-combination of permission-divided formulas, $A^\pi \otimes A^\sigma \not\models A^{\pi \oplus \sigma}$. We avoid the former difficulty by *retaining* the strong $*$ in the formalism alongside \otimes , and the latter by using nominal *labels*, from hybrid logic, to record exact aliasing between read-only copies of a formula.

The main previous work addressing these issues, by Le and Hobor [24], uses a combination of permissions based on *tree shares* [17] and semantic side conditions on formulas to overcome the aforementioned problems. The *rely-guarantee* separation logic in [30] similarly restricts concurrent reasoning to structures described by precise formulas only. In contrast, our logic is a little more complex, but we can use permissions of any kind, and do not require side conditions. In addition, our use of labelling enables us to handle examples involving the transfer of data structures between concurrent threads.

On the other hand, we think it probable that the kind of examples we consider in this paper could also be proven by hand in at least some of the verification formalisms derived from CSL (e.g. [16, 22, 27]). For example, using the “concurrent abstract predicates” in [16], one can explicitly declare shared regions of memory in a fairly ad-hoc way. However, such program logics are typically very complicated and, we believe, quite unlikely to be amenable to automation.

We feel that the main appeal of the present work lies in its relative simplicity—we build on standard CSL with permissions and invoke only a modest amount of extra syntax—which bodes well for its potential automation (at least for simpler examples). In practical terms, an obvious way to proceed would be to develop a prototype verifier for concurrent programs based on our logic SL_{LP} . An important challenge in this area is to develop heuristics—e.g., for splitting, labelling and combining formulas—that work acceptably well in practice.

An even greater challenge is to move from *verifying* user-provided specifications to *inferring* them automatically, as is done e.g. by Facebook INFER. In separation logic, this crucially depends on solving the *biabduction* problem, which aims to discover “best fit” solutions for applications of the frame rule [9, 11]. In the CSL setting, a further problem seems to lie in deciding how applications of the concurrency rule should divide resources between threads.

Finally, automating the verification approach set out in this paper will likely necessitate restricting our full logic to some suitably tractable fragment, e.g. one analogous to the well-known *symbolic heaps* in standard separation logic (cf. [2, 15]). The identification of such tractable fragments is another important theoretical problem in this area. It is our hope that this paper will serve to stimulate interest in the automation of concurrent separation logic in particular, and permission-sensitive reasoning in general.

References

1. Appel, A.W., et al.: Program Logics for Certified Compilers. Cambridge University Press, New York (2014)

2. Berdine, J., Calcagno, C., O'Hearn, P.W.: A decidable fragment of separation logic. In: Lodaya, K., Mahajan, M. (eds.) FSTTCS 2004. LNCS, vol. 3328, pp. 97–109. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30538-5_9
3. Blackburn, P., de Rijke, M., Venema, Y.: *Modal Logic*. Cambridge University Press, Cambridge (2001)
4. Bornat, R., Calcagno, C., O'Hearn, P., Parkinson, M.: Permission accounting in separation logic. In: *Proceedings of POPL-32*, pp. 59–70. ACM (2005)
5. Boyland, J.: Checking interference with fractional permissions. In: Cousot, R. (ed.) SAS 2003. LNCS, vol. 2694, pp. 55–72. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-44898-5_4
6. Brookes, S.: A semantics for concurrent separation logic. *Theoret. Comput. Sci.* **375**(1–3), 227–270 (2007)
7. Brotherston, J.: Formalised inductive reasoning in the logic of bunched implications. In: Nielson, H.R., Filé, G. (eds.) SAS 2007. LNCS, vol. 4634, pp. 87–103. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74061-2_6
8. Brotherston, J., Fuhs, C., Gorogiannis, N., Navarro Pérez, J.: A decision procedure for satisfiability in separation logic with inductive predicates. In: *Proceedings of CSL-LICS*, pp. 25:1–25:10. ACM (2014)
9. Brotherston, J., Gorogiannis, N., Kanovich, M.: Biabduction (and related problems) in array separation logic. In: de Moura, L. (ed.) CADE 2017. LNCS (LNAI), vol. 10395, pp. 472–490. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63046-5_29
10. Brotherston, J., Villard, J.: Parametric completeness for separation theories. In: *Proceedings of POPL-41*, pp. 453–464. ACM (2014)
11. Calcagno, C., Distefano, D., O'Hearn, P., Yang, H.: Compositional shape analysis by means of bi-abduction. *J. ACM* **58**(6), 1–66 (2011)
12. Calcagno, C., O'Hearn, P., Yang, H.: Local action and abstract separation logic. In: *Proceedings of LICS-22*, pp. 366–378. IEEE Computer Society (2007)
13. Cao, Q., Wang, S., Hobor, A., Appel, A.W.: Proof pearl: magic wand as frame (2019)
14. Costea, A., Chin, W.-N., Qin, S., Craciun, F.: Automated modular verification for relaxed communication protocols. In: Ryu, S. (ed.) APLAS 2018. LNCS, vol. 11275, pp. 284–305. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-02768-1_16
15. Demri, S., Lozes, E., Lugiez, D.: On symbolic heaps modulo permission theories. In: *Proceedings of FSTTCS-37*, pp. 25:1–25:13. Dagstuhl (2017)
16. Dinsdale-Young, T., Dodds, M., Gardner, P., Parkinson, M.J., Vafeiadis, V.: Concurrent abstract predicates. In: D'Hondt, T. (ed.) ECOOP 2010. LNCS, vol. 6183, pp. 504–528. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14107-2_24
17. Dockins, R., Hobor, A., Appel, A.W.: A fresh look at separation algebras and share accounting. In: Hu, Z. (ed.) APLAS 2009. LNCS, vol. 5904, pp. 161–177. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-10672-9_13
18. Hobor, A., Appel, A.W., Nardelli, F.Z.: Oracle semantics for concurrent separation logic. In: Drossopoulou, S. (ed.) ESOP 2008. LNCS, vol. 4960, pp. 353–367. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78739-6_27
19. Hobor, A., Gherghina, C.: Barriers in concurrent separation logic: now with tool support!. *Logical Methods Comput. Sci.* **8**, 1–36 (2012)
20. Hobor, A., Villard, J.: The ramifications of sharing in data structures. In: *Proceedings of POPL-40*, pp. 523–536. ACM (2013)

21. Hóu, Z., Clouston, R., Goré, R., Tiu, A.: Proof search for propositional abstract separation logics via labelled sequents. In: Proceedings of POPL-41, pp. 465–476. ACM (2014)
22. Krebbers, R., Jung, R., Bizjak, A., Jourdan, J.-H., Dreyer, D., Birkedal, L.: The essence of higher-order concurrent separation logic. In: Yang, H. (ed.) ESOP 2017. LNCS, vol. 10201, pp. 696–723. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54434-1_26
23. Larchey-Wendling, D., Galmiche, D.: Exploring the relation between intuitionistic BI and Boolean BI: an unexpected embedding. *Math. Struct. Comput. Sci.* **19**, 1–66 (2009)
24. Le, X.-B., Hobor, A.: Logical reasoning for disjoint permissions. In: Ahmed, A. (ed.) ESOP 2018. LNCS, vol. 10801, pp. 385–414. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-89884-1_14
25. Lee, W., Park, S.: A proof system for separation logic with magic wand. In: Proceedings of POPL-41, pp. 477–490. ACM (2014)
26. O’Hearn, P.W.: Resources, concurrency and local reasoning. *Theoret. Comput. Sci.* **375**(1–3), 271–307 (2007)
27. Raad, A., Villard, J., Gardner, P.: CoLoSL: concurrent local subjective logic. In: Vitek, J. (ed.) ESOP 2015. LNCS, vol. 9032, pp. 710–735. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46669-8_29
28. Reynolds, J.C.: Separation logic: a logic for shared mutable data structures. In: Proceedings of LICS-17, pp. 55–74. IEEE Computer Society (2002)
29. Vafeiadis, V.: Concurrent separation logic and operational semantics. In: Proceedings of MFPS-27, pp. 335–351. Elsevier (2011)
30. Vafeiadis, V., Parkinson, M.: A marriage of rely/guarantee and separation logic. In: Caires, L., Vasconcelos, V.T. (eds.) CONCUR 2007. LNCS, vol. 4703, pp. 256–271. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74407-8_18
31. Villard, J., Lozes, É., Calcagno, C.: Tracking heaps that hop with heap-hop. In: Esparza, J., Majumdar, R. (eds.) TACAS 2010. LNCS, vol. 6015, pp. 275–279. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-12002-2_23
32. Yang, H., O’Hearn, P.: A semantic basis for local reasoning. In: Nielsen, M., Engberg, U. (eds.) FoSSaCS 2002. LNCS, vol. 2303, pp. 402–416. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45931-6_28

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

