

# Sound Automation of Magic Wands

No Author Given

No Institute Given

**Abstract.** The magic wand  $\text{-*}$  (also called separating implication) is a separation logic connective commonly used to specify properties of partial data structures, for instance during iterative traversals. A *footprint* of a magic wand formula  $A \text{-*} B$  is a state that, combined with any state in which  $A$  holds, yields a state in which  $B$  holds. The key challenge of proving a magic wand (also called *packaging* a wand) is to find such a footprint. Existing package algorithms either have a high annotation overhead or, as we show in this paper, are unsound.

We present a formal framework that precisely characterises a wide design space of possible package algorithms applicable to a large class of separation logics. We prove in Isabelle/HOL that our formal framework is sound and complete, and use it to develop a novel package algorithm that offers competitive automation and is sound. Moreover, we present a novel, restricted definition of wands and prove in Isabelle/HOL that it is possible to soundly combine fractions of such wands, which is not the case for arbitrary wands. We have implemented our techniques for the Viper language, and demonstrate that they are effective in practice.

## 1 Introduction

Separation logic [34] (SL hereafter) is a program logic that has been widely used to prove complex properties of heap-manipulating programs. The two main logical connectives that enable such reasoning are the *separating conjunction*  $*$  and the *separating implication* (more commonly known as the *magic wand*)  $\text{-*}$ , in combination with *resource assertions* which represent e.g. exclusive ownership of (and permission to access) particular heap locations. The separating conjunction expresses that two assertions prescribe ownership of disjoint parts of the heap, useful, for instance, to reason about aliasing or race conditions. More precisely, the assertion  $A * B$  holds in a program state  $\sigma$  if and only if  $\sigma$  can be split into two *compatible* program states  $\sigma_A$  and  $\sigma_B$  such that  $A$  and  $B$  hold in  $\sigma_A$  and  $\sigma_B$ , respectively. In SL, heaps of program states are *partial* maps from locations to values; their domains represent heap locations exclusively owned. Two program states are compatible if (the domains of) their heaps are disjoint.

Intuitively, a magic wand  $A \text{-*} B$  can be used to express the difference between the heap locations that  $B$  and  $A$  provide permission to access. The magic wand is useful, for instance, to specify partial data structures, where  $B$  specifies the entire data structure and  $A$  specifies a part that is missing [37, 29].  $A \text{-*} B$  holds in a state  $\sigma_w$ , if and only if for *any* program state  $\sigma_A$  in which  $A$  holds and that

is compatible with  $\sigma_w$ ,  $B$  holds in the state obtained by combining the heaps of  $\sigma_A$  and  $\sigma_w$ . Thus, if  $A * (A \multimap B)$  holds in a state, then so does  $B$ , analogously to the *modus ponens* inference rule in propositional logic.

The magic wand has been shown to enable or greatly simplify proofs in many different cases [38, 24, 17, 37, 29, 16, 11, 1]. For instance, Yang [38] uses the magic wand to prove the Schorr-Waite graph marking algorithm. Dodds *et al.* [16] employ the wand to specify synchronisation barriers for deterministic parallelism. Examples using magic wands to specify partial data structures include tracking ongoing traversals of a data structure [37, 29], where the left-hand side of the wand specifies the part of the data structure yet to be traversed, or for specifying protocols that enforce orderly modification of data structures [24, 17, 21] (e.g. the protocol governing Java iterators). More recently, wands have been used for formal reasoning about borrowed references in the Rust programming language, which employs an ownership type system. Magic wands can concisely specify the data structure from which the borrowed reference was taken, in terms of modifications to the partial data structure accessible via the reference [1].

The complexity of SL proofs has given rise to a variety of automatic SL verifiers that reduce the required proof effort. Given the usefulness of magic wands, it is important that such verifiers also provide automatic support for wands. However, reasoning about a magic wand requires reasoning about *all* states in which the left-hand side holds, which is challenging. It has been shown that a separation logic even without the separating conjunction (but with the magic wand) is as expressive as a variant of second-order logic and, thus, undecidable [8].

Two different approaches [35, 4] that provide partially-automated support are implemented in the verifiers Viper [30] and VerCors [3]. However, the approach implemented in VerCors [4] incurs significant annotation overhead, and the approach in Viper [35] suffers from a fundamental, previously undiscovered flaw that renders the approach unsound. Both approaches require user-provided *package operations* to direct the verifier’s proof search. *Packaging* a wand  $A \multimap B$  expresses that the verifier should prove and subsequently record  $A \multimap B$ . To package  $A \multimap B$  the verifier must split the current state into two compatible states  $\sigma'$  and  $\sigma_w$  such that  $A \multimap B$  holds in  $\sigma_w$ . We call  $\sigma_w$  a *footprint* of the wand. After successfully packaging a wand, the verifier must disallow changes to  $\sigma_w$  to preserve the wand’s validity: the verifier *packages the footprint into the wand*.

The key challenge for supporting magic wands in automatic verifiers is to define a *package algorithm* that packages a wand. In VerCors’s package algorithm, a user must manually specify a footprint for the wand and the algorithm checks whether the wand holds in the specified footprint. This leads to a lot of annotation overhead. Viper’s current package algorithm reduces this overhead significantly by defining an algorithm to automatically infer a suitable footprint. Unfortunately, as we show in this paper, Viper’s current algorithm has a fundamental flaw that causes the algorithm to infer an *incorrect* footprint in certain cases, i.e. a state in which the wand does *not* hold. As a result, one can prove results that do not hold, which leads to unsound reasoning. We will explain the fundamental flaw in Sec. 2; it illustrates the subtlety of supporting this important connective.

**Approach and Contributions.** In this paper, we present a formal foundation for sound package algorithms (key to automating wands in automatic SL verifiers), and we implement a novel such algorithm based on these foundations. Our algorithm requires the same annotation overhead as the prior, flawed Viper algorithm, which is (to our knowledge) the most automatic existing approach. We introduce a formal framework expressed via a novel *package logic* that defines the design space for package algorithms. The soundness of a package algorithm can be justified by showing that the algorithm finds a proof in our package logic. The design space for package algorithms is large since there are various aspects that affect how one expresses the algorithm including (1) which footprint an algorithm infers or checks for a wand (there are often multiple options as we will show in Sec. 3), (2) the state model (which differs between different SL verifiers), and (3) restricted definitions of wands (for instance, to ensure each wand has a unique minimal footprint). Our package logic deals with (1) by capturing all sound derivations for the same wand. To deal with (2) and (3), our logic is parametric along multiple dimensions. For instance, the state model can be any separation algebra to support different SL extensions (e.g. fractional permissions [5]).

Our logic also supports parameters to restrict the allowed footprints for wands in systematic ways. Such restrictions are useful, for instance, in a logic supporting *fractional permissions*. Fractional permissions permit splitting ownership/resources into shared fragments which typically permit read access to the underlying data. However, as we show in Sec. 4, fractional parts of general magic wands cannot always be soundly recombined. Existing solutions for other connectives impose side conditions to enable sound recombinations [25], which are often hard to check automatically. We instead introduce a novel restriction of magic wands to avoid such side conditions and develop a corresponding second package algorithm again based on the formal framework provided by our package logic. We make the following contributions:

- We formalise a *package logic* that can be used as a basis for a wide range of package algorithms (Sec. 3). The logic has multiple parameters including: a separation algebra to model the states and a parameter to restrict the definition of a wand in a systematic way. We formally prove the logic sound and complete for any instantiation of the parameters in Isabelle/HOL.
- We develop a novel, restricted definition of a wand (Sec. 4) and prove in Isabelle/HOL that this wand can always be recombined.
- We implement sound package algorithms for both the standard and the restricted wand in the Viper verifier and justify their soundness directly via our package logic (Sec. 5). We evaluate both algorithms on the Viper test suite. Our evaluation shows that (1) our algorithms perform similarly well to prior work and correctly reject examples where prior work is unsound, and (2) our restricted wand definition is expressive enough for most examples.

We will make our Isabelle formalisation and the implementation of our new package algorithms available as part of our artefact.

## 2 Background and Motivation

In this section, we present the necessary background for this paper. We use *implicit dynamic frames* [36] to represent SL assertions, since both existing automatic verifiers that support wands (VerCors and Viper) are based on it. There is a known strong correspondence between SL and implicit dynamic frames [32].

### 2.1 Implicit Dynamic Frames

Just like SL assertions, implicit dynamic frames (IDF hereafter) assertions specify not only value information, but also *permissions* to heap locations that are allowed to be accessed. To justify dereferencing a heap location, the corresponding permission is required, ensuring memory safety. IDF assertions specify permissions to locations and value information separately. An assertion  $\text{acc}(\mathbf{x.val})$  (an *accessibility predicate*) denotes permission to the heap *location*  $\mathbf{x.val}$ , while  $\mathbf{x.val} = v$  expresses that  $\mathbf{x.val}$  contains *value*  $v$ . The separating conjunction in IDF enforces disjointness (formally: acts multiplicatively) with respect to resource assertions such as accessibility predicates; in particular, if  $\text{acc}(\mathbf{x.val}) * \text{acc}(\mathbf{y.val})$  holds in a state, then  $\mathbf{x}$  and  $\mathbf{y}$  must be different (analogously to SL).

The main difference between IDF and SL is that SL does not allow general heap-dependent expressions such as  $\mathbf{x.val} = v$  or  $\mathbf{x.left.right}$  [36] to be specified separately from the permissions to the heap locations they depend on. The IDF assertion  $\text{acc}(\mathbf{x.val}) * \mathbf{x.val} = v$  must be expressed in SL via the *points-to assertion*  $\mathbf{x.val} \mapsto v$ , which also conveys exclusive permission to the location  $\mathbf{x.val}$ . IDF supports heap dependent expressions within *self-framing* assertions: those which require permissions to all the heap locations on whose values they depend (e.g.  $\text{acc}(\mathbf{x.val}) * \mathbf{x.val} = v$  is self-framing but  $\mathbf{x.val} = v$  is not) [36].

### 2.2 A Typical Example Using Magic Wands

Fig. 1 shows a variation of an example from the VerifyThis verification competition [18]. The method `leftLeaf` computes the leftmost leaf of a binary tree in an iterative fashion (`package` and `apply` operations, shown in blue, should be ignored for now). The pre- and postconditions of `leftLeaf` are both  $\text{Tree}(\mathbf{x})$ , which is a *predicate instance* used to specify all permissions to `val`, `left`, and `right` fields in the tree rooted at  $\mathbf{x}$  (the recursive definition of this predicate is on the right of Fig. 1). Proving this specification amounts to proving that `leftLeaf` is memory-safe and that the permissions to the input tree are preserved, which is typically necessary to enable further such methods to be called on the same tree.

The key challenge when verifying `leftLeaf` is specifying an appropriate loop invariant. The loop invariant must track the permissions to the subtree rooted at  $\mathbf{y}$  that still needs to be traversed, since otherwise dereferencing  $\mathbf{y.left}$  in the loop body is not allowed. Additionally, the invariant must track all of the remaining permissions in the input tree rooted at  $\mathbf{x}$  (the permissions to the nodes already traversed and others unreachable from  $\mathbf{y}$ ), since otherwise the postcondition cannot be satisfied. The former can be easily expressed with  $\text{Tree}(\mathbf{y})$ . The latter can be

```

1 method leftLeaf(x: Ref) : (y: Ref)
2   requires Tree(x)
3   ensures Tree(x) {
4     y := x
5     package Tree(x) -* Tree(x)
6
7     while(y.left != null)
8       inv Tree(y) * (Tree(y) -* Tree(x)) {
9         y := y.left
10        package Tree(y) -* Tree(x)
11         // { hints for package}
12      }
13    apply Tree(y) -* Tree(x)
14  }

```

$\text{Tree}(x: \text{Ref}) \triangleq$   
 $\text{acc}(x.\text{val}) * \text{acc}(x.\text{left}) * \text{acc}(x.\text{right})$   
 $(x.\text{left} \neq \text{null} \Rightarrow \text{Tree}(x.\text{left})) * (x.\text{right} \neq \text{null} \Rightarrow \text{Tree}(x.\text{right}))$

**Fig. 1.** The code on the left computes the leftmost leaf of a binary tree and includes specifications to prove memory safety. The predicate describing the permissions of a tree is defined on the right. The loop invariant uses a wand to summarise the permissions of the input tree excluding the tree yet to be traversed. The blue operations are ghost operations to guide the verifier. We omit ghost operations specific to predicates. The `package` on lines 10 and 11 requires further hints in existing approaches (the hints vary based on the approach) as we discuss in App. J. Here, for instance, one must specify that in order to package the wand, the wand in the loop invariant must be applied.

elegantly achieved with a magic wand  $\text{Tree}(y) -* \text{Tree}(x)$ . This wand promises  $\text{Tree}(x)$  if one combines the wand with  $\text{Tree}(y)$ . That is, the wand represents (at least) the difference between the permissions making up the two trees. Using SL’s modus-ponens-like inference rule (directed by the `apply` operation on line 13, explained next), one can show that the loop invariant entails the postcondition.

### 2.3 Wand Ghost Operations

Automatic SL verifiers such as GRASShopper [33], VeriFast [20], VerCors, and Viper generally represent permissions owned by a program state in two ways: by recording predicate instances (such as  $\text{Tree}(x)$  in Fig. 1) and *direct* permissions to heap locations. Magic wand instances provide a third way to represent permissions and are recorded analogously. Verifiers that support them require two wand-specific *ghost operations*, which instruct the verifiers when to prove a wand and when to apply a recorded wand instance using SL’s modus-ponens-like rule.

A `package` ghost operation expresses that a verifier should prove a new wand instance in the current state and report an error if the proof attempt fails. To prove a new wand instance, the verifier must split the current state into two states  $\sigma'$  and  $\sigma_w$  such that the wand holds in the *footprint* state  $\sigma_w$ ; on success, permissions in the footprint are effectively exchanged for the resulting magic wand instance. We call a procedure that selects a footprint by splitting the

current state a *package algorithm*. On lines 5 and 10 of Fig. 1, new wands are packaged to establish and preserve the invariant, respectively.

The **apply** operation *applies* a wand  $A \multimap B$  using SL’s modus-ponens-like rule if the verifier records a wand instance of  $A \multimap B$  and  $A$  holds in the current state (and otherwise fails), exchanging these for the assertion  $B$ . The **apply** operation is directly justified by the wand’s semantics: Combining a wand’s footprint with *any* state in which  $A$  holds is guaranteed to yield a state in which  $B$  holds. For the **apply** operation on line 13 of Fig. 1, the verifier removes the applied wand instance and  $\text{Tree}(y)$ , in exchange for the predicate instance  $\text{Tree}(x)$ .

## 2.4 The Footprint Inference Attempt (FIA)

Package algorithms differ in how a footprint for the specified magic wand is selected. In one existing approach [4], the user must manually provide the footprint and the algorithm checks whether the specified footprint is correct. In the other approach [35], a footprint is inferred. We explain and compare to the latter approach since it is the more automatic of the two; hereafter, we refer to its package algorithm as *the Footprint Inference Attempt (FIA)*. Inferring a correct footprint is challenging due the complexity of the wand connective. In particular, we have discovered that, in certain cases, the FIA infers *incorrect* footprints, leading to unsound reasoning. The goal of this subsection is to understand the FIA’s key ideas, which our solution will build on, and why it is unsound.

In general, there may be multiple valid are many footprints for a magic wand  $A \multimap B$ . The FIA attempts to infer a footprint which is as close as possible to the *difference* between the permissions required by  $B$  and  $A$ , taking as few permissions as possible while aiming for a footprint compatible with  $A$  (so that the resulting wand can be later applied) [35]. That is, the FIA includes only permissions in the footprint it infers that are specified by  $B$  and *not* guaranteed by  $A$ .

For a wand  $A \multimap B$ , the FIA constructs an arbitrary state  $\sigma_A$  that satisfies  $A$  (representing  $\sigma_A$  symbolically). Then, the FIA tries to construct a state  $\sigma_B$  in which  $B$  holds by taking permissions (and copying corresponding heap values) from  $\sigma_A$  if possible and the current state otherwise. If this algorithm succeeds, the (implicit) inferred footprint consists of the permissions that were taken from the current state. The FIA constructs  $\sigma_B$  by iterating over the permissions and logical constraints in  $B$ . For each permission, the FIA checks whether  $\sigma_A$  owns the permission. If so, the FIA adds the permission to  $\sigma_B$  and removes the permission from  $\sigma_A$ . Otherwise, the FIA removes the permission from the current state or fails if the current state does not have the permission. For each logical constraint, the FIA checks that the constraint holds in  $\sigma_B$  as constructed so far. We show an example of the FIA correctly packaging a wand in App. A.

**Unsoundness of the FIA.** We have discovered that for some wands  $A \multimap B$ , the FIA determines an *incorrect* footprint for the magic wand. This unsoundness can arise when the FIA performs a case split on the content of the arbitrary

state  $\sigma_A$  satisfying  $A$ . In such situations, the FIA infers a footprint for each case *separately*, making use of properties that hold in that case. For certain wands, this leads to different footprints being selected for each case, while *none* of the inferred footprints can be used to justify  $B$  in *all* cases, i.e. for *all* states  $\sigma_A$  that satisfy  $A$ . As a result, the packaged wand does *not* hold in any of the inferred footprints, which can make verification unsound, as we illustrate below.

The wand  $w := \text{acc}(\mathbf{x.f}) * (\mathbf{x.f} = \mathbf{y} \vee \mathbf{x.f} = \mathbf{z}) -* \text{acc}(\mathbf{x.f}) * \text{acc}(\mathbf{x.f.g})$  illustrates the problem. For this wand, every state  $\sigma_A$  satisfying the left-hand side must have permission to  $\mathbf{x.f}$ . However  $\mathbf{x.f}$  may either point to  $\mathbf{y}$  or  $\mathbf{z}$ . If  $\mathbf{x.f}$  points to  $\mathbf{y}$  in  $\sigma_A$ , then to justify the right-hand side's second conjunct, the footprint must contain permission to  $\mathbf{y.g}$ . Analogously, if  $\mathbf{x.f}$  points to  $\mathbf{z}$  in  $\sigma_A$ , then the footprint must contain permission to  $\mathbf{z.g}$ . The wand's semantics requires a footprint to justify the wand's right-hand side for all states in which the left-hand side holds, and thus, a correct footprint must be able to justify *both* cases. Hence, the footprint must have permission to *both*  $\mathbf{y.g}$  and  $\mathbf{z.g}$ . However, the FIA's inferred footprint is in effect the disjunction of these two permissions.

Packaging the above wand  $w$  using the FIA leads to unsound reasoning. After the incorrect package described above, the assertion  $(\text{acc}(\mathbf{y.g}) \vee \text{acc}(\mathbf{z.g})) * w$  can be proved since the FIA removes permission to either  $\mathbf{y.g}$  or  $\mathbf{z.g}$  from the current state, but not both. However, this assertion does not actually hold! According to the semantics of wands,  $w$ 's footprint must include permission to both  $\mathbf{y.g}$  and  $\mathbf{z.g}$ ; the separating conjunction means that the assertion is equivalent to false.

The unsoundness of the FIA shows the subtlety and challenge of developing sound package algorithms. Algorithms that soundly infer a single footprint for all states in which the wand's left-hand side holds must be more involved than the FIA. Ensuring their soundness requires a *formal* framework to construct them and justify their correctness. We introduce such a framework in the next section.

### 3 A Logical Framework for Packaging Wands

In this section, we present a new logical framework that defines the design space for (sound) package algorithms. The core of this framework is our *package logic*, which defines the space of potential algorithmic choices of a footprint for a particular magic wand. Successfully packaging a wand in a given state is (as we will show) equivalent to finding a derivation in our package logic, and any actual package algorithm must correspond to a proof search in our logic (if it is sound). In particular, we provide soundness (Thm. 1) and completeness (Thm. 2) results for our logic. We define a specific package algorithm with this logic at its foundation, inspired by the FIA package algorithm [35] (described in Sec. 2.4) but amending its unsoundness, resulting in (to the best of our knowledge) the first sound and relatively automatic package algorithm.

All definitions and results in this section have been fully mechanised in Isabelle/HOL. Our mechanised definitions are parametric with the underlying verification logic in various senses: the underlying separation algebra is a parameter, the syntax of assertions is defined in a way which allows simple extension

with different base cases and connectives, and the semantics of magic wands itself can be restricted if only particular kinds of footprint are desired in practice. As a specific example of the latter parameter, in Sec. 4 we define a novel restriction of magic wand footprints which guarantees better properties in combination with certain usages of fractional permissions; this is seamlessly supported by the general package logic presented here. Nonetheless, to simplify the exposition of this section, we will assume that any magic wand footprint satisfying the connective’s standard semantics is an acceptable result.

### 3.1 Footprint Selection Strategies

As we explained in Sec. 1, there is a wide design space for package algorithms; in particular, many potential strategies for finding a magic wand’s footprint exist and none is clearly optimal. Recall that a footprint is a state, and thus consists of permissions to certain heap locations as well as storing their corresponding values; for simplicity we identify a footprint by the permissions it contains.

For example, consider the following magic wand (using fractional permissions)  $\text{acc}(\mathbf{x.b}, 1/2) \multimap \text{acc}(\mathbf{x.b}, 1/2) * (\mathbf{x.b} \Rightarrow \text{acc}(\mathbf{x.f}))$ . Suppose this magic wand is to be packaged in a state where full permissions to both  $\mathbf{x.b}$  and  $\mathbf{x.f}$  are held, and the value of  $\mathbf{x.b}$  is currently false. Two valid potential footprints are:

1. Full permission to  $\mathbf{x.f}$ . This is sufficient to guarantee the right-hand side will hold regardless of the value that  $\mathbf{x.b}$  has by the time the wand is applied.
2. Half permission to  $\mathbf{x.b}$ . By including this permission, the fact that  $\mathbf{x.b}$  is currently false is also included, and thus permission to  $\mathbf{x.f}$  is not needed.

There is no clear reason to prefer one choice over the other: different package algorithms (or manual choices) might choose either. Our package logic allows either choice along with any of many less optimal choices, such as taking both permissions. On the other hand, as motivated earlier in Sec. 3.1, our package logic must (and does) enforce that a single valid footprint is chosen for a wand that works for each and every potential state satisfying its left-hand side.

### 3.2 Package Logic: Preliminaries

In order to capture different state models and different flavours of separation logic, our package logic is parameterised by a separation algebra. For space reasons, we present here a simplified overview of this separation algebra, but all definitions are given in App. C and have been mechanised in Isabelle/HOL. We consider a separation algebra [10, 15] where  $\Sigma$  is the set of states,  $\oplus : \Sigma \times \Sigma \rightarrow \Sigma$  is a partial operation that is commutative and associative, and  $e \in \Sigma$ , which corresponds to the empty state, is a neutral element for  $\oplus$ . We write  $\succeq$  for the induced partial order of the resulting partial commutative monoid, and  $\sigma_1 \# \sigma_2$  iff  $\sigma_1 \oplus \sigma_2$  is defined (i.e.  $\sigma_1$  and  $\sigma_2$  are compatible). Finally, if  $\sigma_2 \succeq \sigma_1$ , we define the subtraction  $\sigma_2 \ominus \sigma_1$  to be the  $\succeq$ -largest state  $\sigma_r$  such that  $\sigma_2 = \sigma_1 \oplus \sigma_r$ .

We define our package logic for an assertion language with the following grammar:  $A = A * A \mid \mathcal{B} \Rightarrow A \mid \mathcal{B}$ , where  $A$  ranges over assertions and  $\mathcal{B}$  over *semantic assertions*. The semantics of this assertion language is formally defined



in App. C. To allow our package logic to be applied to a variety of underlying assertion logics, we distinguish only the two most-relevant connectives: the separating conjunction and an implication (for expressing conditional assertions). To support additional connectives and base-cases of the assertion logic, the third type of assertion we consider is a general kind of *semantic assertion*, i.e. function from  $\Sigma$  to Booleans. This third type can be instantiated to represent logical assertions that don't match the first two cases. In particular, assertions such as  $x.f = 5$ ,  $\text{acc}(x.f)$ , abstract predicates (such as  $\text{Tree}(x)$ ) or magic wands can be represented as semantic assertions. This core assertion language can also be easily extended with native support for e.g. the logical conjunction and disjunction connectives; we explain in App. D how to extend the rules of the logic accordingly.

### 3.3 The Package Logic

We define our package logic to prescribe the design space of algorithms for deciding how, in an initial state  $\sigma_0$ , to select a valid footprint (or fail) for a magic wand  $A \multimap B$ . The aim is to infer states  $\sigma_w$  and  $\sigma_1$  that partition  $\sigma_0$  (i.e.  $\sigma_0 = \sigma_1 \oplus \sigma_w$ ) such that  $\sigma_w$  is a valid footprint for  $A \multimap B$  (when combined with any compatible state satisfying  $A$ , the resulting state satisfies  $B$ ). In particular, all permissions (and logical facts) required by the assertion  $B$  must either come from the footprint or be guaranteed to be provided by any compatible state satisfying  $A$ .

Recall from Sec. 2.4 that the mistake underlying the FIA approach ultimately resulted from allowing multiple different footprints to be selected conditionally on state satisfying  $A$ , rather than a single footprint which works for all such states. Our package logic addresses this concern by defining judgements in terms of the *set* of all states satisfying  $A$ ; whenever *any* of these tracked states is insufficient to provide a permission required by  $B$ , our logic will force this permission to be added *in general* to the wand's footprint (taken from the current state).

A *witness set*  $S$  is a set of pairs of states  $(\sigma_A, \sigma_B)$ ; conceptually, the first represents the state available for trying to prove  $B$  *in addition* to the current state; this is initially a state satisfying the wand's left-hand side  $A$ . The second represents the state assembled (so-far) to attempt to satisfy the right-hand side  $B$ . We write  $S^1$  for the set of first elements of all pairs in a witness set  $S$ . A *context*  $\Delta$  is a pair  $(\sigma, S)$  of a state and a witness set; here,  $\sigma$  represents the (as-yet unused remainder of the) current state in which the wand is being packaged.

The basic idea behind a derivation in our logic is to show how to assemble a witness set in which *all* second elements are states satisfying  $B$ , via some combinations of: (1) moving a part of the first element of a pair in the witness set into the second, and (2) moving a part of the outer state  $\sigma$  into *all* first elements of the pairs (this becomes a part of the wand's footprint). The actual judgements of the logic are a little more complex, to correctly record any hypotheses (called *path-conditions*) that result from deconstructing conditional assertions in  $B$ .

**Configurations and Reductions.** A *configuration* represents a current objective in our package logic: the part of the wand's right-hand side still to be

satisfied as well as the current state of a footprint computation. A configuration is a triple  $\langle B, pc, (\sigma, S) \rangle$ , where  $B$  is an assertion,  $pc$  is a *path condition* (a function from  $\Sigma$  to Booleans), and  $(\sigma, S)$  is a context. Conceptually,  $B$  is the assertion still to be satisfied,  $pc$  represents hypotheses we are currently working under, and the context  $(\sigma, S)$  tracks the current state and witness set, as described above.

A *reduction* is a judgement  $\langle B, pc, (\sigma_0, S_0) \rangle \rightsquigarrow (\sigma_1, S_1)$ , representing the achievement of the objective described via the configuration on the left, resulting in the final context on the right;  $\sigma_1$  is the new version of the outer state (and becomes the new current state after the **package** operation); whatever was removed from the initial outer state is implicitly the selected footprint state  $\sigma_w$ . If a reduction is derivable in our package logic, this footprint  $\sigma_w$  guarantees that for all  $(\sigma_A, \sigma_B) \in S_0$ , if  $(\sigma_A \oplus \sigma_B) \# \sigma_w$ , then  $\sigma_A \oplus \sigma_w$  satisfies  $pc \Rightarrow B$ . The condition  $(\sigma_A \oplus \sigma_B) \# \sigma_w$  ensures that the pair  $(\sigma_A, \sigma_B)$  actually corresponds to a state in which the wand can be applied given the chosen footprint  $\sigma_w$ , as we explain later. The package logic defines the steps an algorithm may take to achieve this goal.

We represent packaging a wand  $A \multimap B$  in state  $\sigma_0$  by the derivation of a reduction  $\langle B, \lambda\sigma. \top, (\sigma_0, \{(\sigma_A, e) \mid \sigma_A \models A\}) \rangle \rightsquigarrow (\sigma_1, S_1)$ , for some state  $\sigma_1$  and witness set  $S_1$ . The path condition is initially true (we are not yet under any particular hypotheses). The initial witness set contains all pairs of a state  $\sigma_A$  that satisfies  $A$  and the empty state  $e$ , to which a successful reduction will add permissions in order to satisfy  $B$ <sup>1</sup>. Note that an actual algorithm based on our package logic need not explicitly compute this set (which could well be infinite), but can instead track it symbolically. If the algorithm finds a derivation of this reduction, it has proven that the difference between  $\sigma_0$  and  $\sigma_1$  is a valid footprint of the wand  $A \multimap B$ , since the logic is sound (Thm. 1 below).

**Rules.** Fig. 2 presents the four rules of our logic, defining (via derivable reductions) how a configuration can be reduced to a context. There is a rule for each type of assertion  $B$ : *Implication* for an implication, *Star* for a separating conjunction, and *Atom* for a semantic assertion. The logic also includes the rule *Extract*, which represents a choice to extract permissions from the outer state and adds them to all pairs of states in the witness set. In the following, we informally write *reducing an assertion* to refer to the process of deriving (in the logic) that the relevant configuration containing this assertion reduces to some context.

To reduce an implication  $B \Rightarrow A$ , the rule *Implication* conjoins the hypothesis  $B$  with the previous path condition, leaving  $A$  to be reduced. Informally, this expresses that satisfying  $pc \Rightarrow (b \Rightarrow A)$  is equivalent to satisfying  $(pc \wedge b) \Rightarrow A$ .

For a separating conjunction  $A_1 * A_2$ , the *Star* rule expresses that both  $A_1$  and  $A_2$  must be reduced, in order to reduce  $A_1 * A_2$ ; permissions used in the reduction of the first conjunct must not be used again, which is reflected by the threading-through of the intermediate context  $\Delta_1$ <sup>2</sup>.

<sup>1</sup> If  $B$  is intuitionistic, this can be simplified to only the  $\succeq$ -minimal states that satisfy  $A$ .  $B$  is intuitionistic [34] iff, if  $B$  holds in a state  $\sigma$ , then  $B$  holds in any state  $\sigma'$  such that  $\sigma' \succeq \sigma$ . In intuitionistic SL or in IDF, all assertions are intuitionistic.

<sup>2</sup> The order in the premises is unimportant since  $A_1 * A_2$  and  $A_2 * A_1$  are equivalent.

$$\begin{array}{c}
\frac{\langle A, \lambda\sigma. pc(\sigma) \wedge b(\sigma), \Delta \rangle \rightsquigarrow \Delta'}{\langle b \Rightarrow A, pc, \Delta \rangle \rightsquigarrow \Delta'} \textit{Implication} \qquad \frac{\langle A_1, pc, \Delta_0 \rangle \rightsquigarrow \Delta_1 \quad \langle A_2, pc, \Delta_1 \rangle \rightsquigarrow \Delta_2}{\langle A_1 * A_2, pc, \Delta_0 \rangle \rightsquigarrow \Delta_2} \textit{Star} \\
\\
\frac{\begin{array}{l} \forall (\sigma_A, \sigma_B) \in S. pc(\sigma_A) \implies \sigma_A \succeq choice(\sigma_A, \sigma_B) \wedge \mathcal{B}(choice(\sigma_A, \sigma_B)) \\ S_{\top} = \{(\sigma_A \ominus choice(\sigma_A, \sigma_B), \sigma_B \oplus choice(\sigma_A, \sigma_B)) \mid (\sigma_A, \sigma_B) \in S \wedge pc(\sigma_A)\} \\ S_{\perp} = \{(\sigma_A, \sigma_B) \mid (\sigma_A, \sigma_B) \in S \wedge \neg pc(\sigma_A)\} \end{array}}{\langle \mathcal{B}, pc, (\sigma, S) \rangle \rightsquigarrow (\sigma, S_{\top} \cup S_{\perp})} \textit{Atom} \\
\\
\frac{\begin{array}{l} \sigma_0 = \sigma_1 \oplus \sigma_w \quad \textit{stable}(\sigma_w) \quad \langle A, pc, (\sigma_1, S_1) \rangle \rightsquigarrow \Delta \\ S_1 = \{(\sigma_A \oplus \sigma_w, \sigma_B) \mid (\sigma_A, \sigma_B) \in S_0 \wedge (\sigma_A \oplus \sigma_B) \# \sigma_w\} \end{array}}{\langle A, pc, (\sigma_0, S_0) \rangle \rightsquigarrow \Delta} \textit{Extract}
\end{array}$$

**Fig. 2.** Rules of the package logic.

The *Atom* rule specifies how to prove that all states in  $S^1$  (where  $S$  is the witness set) satisfy the assertion  $pc \Rightarrow \mathcal{B}$ . To understand the premises, consider a pair  $(\sigma_A, \sigma_B) \in S$ . If  $\sigma_A$  does not satisfy the path condition, i.e.  $\neg pc(\sigma_A)$ , then  $\sigma_A$  *does not* have to justify  $\mathcal{B}$ , and thus the pair  $(\sigma_A, \sigma_B)$  is left unchanged; this case corresponds to the set  $S_{\perp}$ . Conversely, if  $\sigma_A$  satisfies the path condition, i.e.  $pc(\sigma_A)$ , then  $\sigma_A$  must satisfy  $\mathcal{B}$ , and the corresponding permissions must be transferred from  $\sigma_A$  to  $\sigma_B$ . Since some assertions may be satisfied in different ways, such as disjunctions, the algorithm has a choice in how to satisfy  $\mathcal{B}$ , which might be different for each pair  $(\sigma_A, \sigma_B)$ . This choice is represented by  $choice(\sigma_A, \sigma_B)$ , which must satisfy  $\mathcal{B}$  and be smaller or equal to  $\sigma_A$ . We update the witness set by transferring  $choice(\sigma_A, \sigma_B)$  from  $\sigma_A$  to  $\sigma_B$ . This second case corresponds to the set  $S_{\top}$ . Note that the *Atom* rule can be applied only if  $\sigma_A$  satisfies  $\mathcal{B}$ , for all pairs  $(\sigma_A, \sigma_B) \in S$  such that  $pc(\sigma_A)$ . If not, a package algorithm must either first extract more permissions from the outer state with the *Extract* rule, or fail.

The *Extract* rule (applicable at any step of a derivation), expresses that we can extract permissions (the state<sup>3</sup>  $\sigma_w$ ) from the outer state  $\sigma_0$ , and combine them with the first element of each pair of states in the witness set. Note that  $(\sigma_A, \sigma_B)$  is removed from the witness set if  $\sigma_A \oplus \sigma_B$  is not compatible with  $\sigma_w$ . In such cases, adding  $\sigma_w$  to  $\sigma_A$  would create a pair in the witness set representing a state in which the wand cannot be applied. Consequently, there is no need to establish the right-hand side of the wand for this pair and our logic correspondingly removes it. Finally, the rule requires that we reduce the assertion  $A$  in the new context.

The strategy of a package algorithm is mostly reflected by when and how it uses the *Extract* rule. To package the wand  $\text{acc}(\mathbf{x.b}, 1/2) * \text{acc}(\mathbf{x.b}, 1/2) * (\mathbf{x.b} \Rightarrow \text{acc}(\mathbf{x.f}))$  from Sec. 3.1. one algorithm might use this rule to extract permission

<sup>3</sup> We explain formally in App. C the notion of a stable state, which is a technicality of our general state model; in standard SL and many variants, all states are stable.

to  $x.f$ ; another might use it to extract permission to  $x.b$  (if  $x.b$  had value false in the original state). App. E shows a full derivation of a reduction in our logic.

### 3.4 Soundness and Completeness

We write  $\vdash \langle B, pc, \Delta \rangle \rightsquigarrow \Delta'$  to express that a reduction can be derived in the logic. As explained above, the goal of a package algorithm is to find a derivation of  $\langle B, \lambda_{\cdot}. \top, (\sigma, \{(\sigma_A, e) \mid \sigma_A \in S_A\}) \rangle \rightsquigarrow (\sigma', S')$ . If it succeeds, then the difference between  $\sigma'$  and  $\sigma$  is a valid footprint of  $A \multimap B$ , since our package logic is sound. In particular, we have proven the following soundness result in Isabelle/HOL:

**Theorem 1. Soundness.** *Let  $B$  be a well-formed<sup>4</sup> assertion. If*

1. *the set  $S_A$  contains all states that satisfy  $A$ . i.e.  $\forall \sigma_A. \sigma_A \models A \Rightarrow \sigma_A \in S_A$ ,*
2.  *$\vdash \langle B, \lambda_{\cdot}. \top, (\sigma, \{(\sigma_A, e) \mid \sigma_A \in S_A\}) \rangle \rightsquigarrow (\sigma', S')$ , and*
3. *at least one of the following conditions holds:*

(a)  *$B$  is intuitionistic*

(b) *For all  $(\sigma_A, \sigma_B) \in S'$ ,  $\sigma_A$  contains no permission (i.e.  $\sigma_A \oplus \sigma_A = \sigma_A$ )*

*then there exists a stable state  $\sigma_w$  s.t.  $\sigma = \sigma' \oplus \sigma_w$  and  $\sigma_w$  is a footprint of  $A \multimap B$ .*

The third premise shows that, in an intuitionistic SL or in IDF, the correspondence between a derivation in the logic and a valid footprint of a wand is straightforward (case (a)). However, in classical SL, one must in general additionally check that all permissions in the witness set have been consumed (case (b)).

We have also proved in Isabelle/HOL that our package logic is complete, i.e. any valid footprint can be computed via a derivation in our package logic:

**Theorem 2. Completeness.** *Let  $B$  be a well-formed<sup>4</sup> assertion. If  $\sigma_w$  is a stable footprint of  $A \multimap B$ , and  $\sigma = \sigma' \oplus \sigma_w$ , then there exists a witness set  $S'$  such that  $\vdash \langle B, \lambda_{\cdot}. \top, (\sigma, \{(\sigma_A, e) \mid \sigma_A \in S_A\}) \rangle \rightsquigarrow (\sigma', S')$ .*

### 3.5 A Sound Package Algorithm

We now describe an automatic package algorithm that corresponds to a proof search strategy in our package logic, and which is thus sound. To convey the main ideas, consider packaging a wand of the shape  $A \multimap B_1 * \dots * B_n$ .<sup>5</sup> Our algorithm traverses the assertion  $B_1 * \dots * B_n$  from left to right, similarly to the FIA approach; this traversal is justified by repeated applications of the rule *Star*. Assume at some point during this traversal that the current context is  $(\sigma_0, S)$ . When we encounter the assertion  $B_i$ , we have two possible cases:

1. All states  $\sigma_A \in S^1$  satisfy  $B_i$ , which means that the permissions (or values) required by  $B_i$  are provided by the left-hand side of the wand. In this case, for each pair  $(\sigma_A, \sigma_B) \in S$ , we transfer permissions (and the corresponding

<sup>4</sup> We formally define well-formedness in App. C. Intuitively, a well-formed assertion roughly corresponds to a self-framing assertion as defined in Sec. 2.1.

<sup>5</sup> In App. I, we additionally show how our package algorithm handles implications.

values) to satisfy  $B_i$  from  $\sigma_A$  to  $\sigma_B$ , using the rule *Atom*. Note that the transferred permissions might be different for each pair  $(\sigma_A, \sigma_B)$ . This gives us a new witness set  $S'$ , while the outer state  $\sigma_0$  is left unchanged. We must then handle the next assertion  $B_{i+1}$  in the context  $(\sigma_0, S')$ .

2. There is at least one pair  $(\sigma_A, \sigma_B) \in S$  such that  $B_i$  does not hold in  $\sigma_A$ . In this case, the algorithm fails if combining the permissions (and values) contained in the outer state with each  $\sigma_A \in S^1$  is not sufficient to satisfy  $B_i$ . Otherwise, we apply the rule *Extract* to transfer permissions from the outer state  $\sigma_0$  to each state  $\sigma_A$  in  $S^1$  such that  $B_i$  holds in  $\sigma_A$ . This gives us a new context  $(\sigma'_0, S')$ . We can now apply the first case with the context  $(\sigma'_0, S')$ .

We show, as an example, how the above algorithm soundly packages the wand  $\text{acc}(\mathbf{x.f}) * (\mathbf{x.f} = \mathbf{y} \vee \mathbf{x.f} = \mathbf{z}) \multimap \text{acc}(\mathbf{x.f}) * \text{acc}(\mathbf{x.f.g})$  in App. B.

## 4 Using the Logic with Combinable Wands

Extending SL with fractional permissions [6] is well-known to be useful for reasoning about heap-manipulating concurrent programs with shared state. In this setting, permission amounts are generalised to fractions  $0 \leq p \leq 1$ . Reading a heap location is permitted if  $p > 0$ , and writing if  $p = 1$ , which permits concurrent reads and ensures exclusive writes. The assertion  $\text{acc}(\mathbf{x.f}, p)$  holds in a state that has *at least*  $p$  permission to  $\mathbf{x.f}$ . A permission amount  $p + q$  to a heap location  $\mathbf{x.f}$  can be split into a permission amount  $p$  and a permission amount  $q$ , i.e.  $\text{acc}(\mathbf{x.f}, p + q) \models \text{acc}(\mathbf{x.f}, p) * \text{acc}(\mathbf{x.f}, q)$ , and these two permissions can be recombined, i.e.  $\text{acc}(\mathbf{x.f}, p) * \text{acc}(\mathbf{x.f}, q) \models \text{acc}(\mathbf{x.f}, p + q)$ .

This concept has been generalised [7, 19, 13, 25, 9] to *fractional assertions*  $A^p$ , representing a fraction  $p$  of  $A$ .  $A^p$  holds in a state  $\sigma$  iff there exists a state  $\sigma_A$  in which  $A$  holds and  $\sigma$  is obtained from  $\sigma_A$  by multiplying all permission amounts held by  $p$  [25, 9]; in this case, we write  $\sigma = p \cdot \sigma_A$ . For example,  $\text{acc}(\mathbf{x.f})^p \equiv \text{acc}(\mathbf{x.f}, p)$ , and  $\text{Tree}(\mathbf{x})^p$  (where  $\text{Tree}$  is the predicate defined in Fig. 1) expresses  $p$  permission to all nodes of the tree rooted in  $\mathbf{x}$ .

Using fractional assertions, one might specify a function `find`, which searches a binary tree and yields a subtree whose root contains key `key`, as follows [9]:  $\{ \text{Tree}(\mathbf{x})^p \} \text{find}(\mathbf{x}, \text{key}) \{ \lambda \text{ret}. (\text{Tree}(\text{ret}) * (\text{Tree}(\text{ret}) \multimap \text{Tree}(\mathbf{x}))^p) \}$ , in which `ret` corresponds to the return value of `find`. This postcondition is similar to the loop invariant in Fig. 1, except that it needs only a fraction  $p$  of  $\text{Tree}(\mathbf{x})$ . A number of automatic SL verifiers, such as Caper [14], Chalice [27], VerCors [2], VeriFast [20], and Viper [30], support fractional assertions in some form.

**Combinable Assertions.** While it is always possible to split an assertion  $A^{p+q}$  into  $A^p * A^q$ , recombining  $A^p * A^q$  into  $A^{p+q}$  is sound only under some conditions, for example [25] if  $A$  is *precise* (in the usual SL sense [34]). We say that  $A$  is *combinable* iff the entailment  $A^p * A^q \models A^{p+q}$  holds for any two positive fractions  $p$  and  $q$  such that  $p + q \leq 1$ . As an example,  $\text{acc}(\mathbf{x.f})$  is combinable, but  $\text{acc}(\mathbf{x.f}) \vee \text{acc}(\mathbf{x.g})$  is not because a state containing half permission to both  $\mathbf{x.f}$  and  $\mathbf{x.g}$  satisfies  $(\text{acc}(\mathbf{x.f}) \vee \text{acc}(\mathbf{x.g}))^{0.5} * (\text{acc}(\mathbf{x.f}) \vee \text{acc}(\mathbf{x.g}))^{0.5}$ ,

but not  $\text{acc}(\mathbf{x.f}) \vee \text{acc}(\mathbf{x.g})$ . Combinable assertions are particularly useful to reason about concurrent programs, for instance, to combine the postconditions of parallel branches when they terminate [9].

However, a magic wand is in general *not* combinable, as we show below. This is problematic for SL verifiers; they cannot soundly combine wands, nor predicates that could possibly contain wands in their bodies. One way to prevent the latter is to forbid magic wands in predicate bodies entirely, but this limits the common usage of predicates to abstract over general assertions in specifications [31]. Another solution is to disallow combining fractional instances of a predicate if its body contains a wand, which means requiring additional annotations to “taint” such predicates transitively. This is overly restrictive for wands which are actually combinable and complicates reasoning about abstract predicate families [31].

To address this issue, we propose a novel restriction of the wand, called *combinable wand* (we use *standard wand* to refer to the usual, unrestricted connective). Unlike standard wands in general, a combinable wand is always combinable if its right-hand side is combinable. Thus, by only using combinable wands instead of standard wands, all assertions in logics such as those employed by VerCors and Viper can be made combinable without any of the other aforementioned restrictions regarding predicates. Sec. 5 shows that the restriction combinable wands impose is sufficiently weak for practical purposes. Finally, footprints of combinable wands can be automatically inferred by package algorithms built on our package logic. All results in this section have been proven in Isabelle/HOL.

**Standard Wands are not Combinable in General.** Even if  $B$  is combinable, the standard wand  $A \ast B$  is, in general, not. As an example, the wand  $w := \text{acc}(\mathbf{x.f}, 1/2) \ast \text{acc}(\mathbf{x.g})$  is not combinable, because  $w^{0.5} \ast w^{0.5} \not\vdash w$ . To see this, consider two states  $\sigma_f$  and  $\sigma_g$ , containing full permissions to only  $\mathbf{x.f}$  and  $\mathbf{x.g}$ , respectively. Both states are valid footprints of  $w$ , i.e.  $\sigma_f \models w$  (because  $\sigma_f$  is incompatible with all states that satisfy the left-hand side) and  $\sigma_g \models w$  (because  $\sigma_g$  entails the right-hand side). Thus, by definition,  $0.5 \cdot \sigma_f \models w^{0.5}$  and  $0.5 \cdot \sigma_g \models w^{0.5}$ . However,  $0.5 \cdot \sigma_f \oplus 0.5 \cdot \sigma_g$ , i.e. a state with half permission to both  $\mathbf{x.f}$  and  $\mathbf{x.g}$ , is *not* a valid footprint of  $w$ , and thus  $w^{0.5} \ast w^{0.5} \not\vdash w$ .

Intuitively,  $w$  is not combinable because one of its footprints,  $\sigma_f$ , is incompatible with the left-hand side of the wand, but becomes compatible when the footprint is scaled down to a fraction.<sup>6</sup> After scaling, the wand no longer holds trivially, and the footprint does not necessarily establish the right-hand side.

To make this intuition more precise, we introduce the notion of *scalable footprints*. For a state  $\sigma$ , we define  $\text{scaled}(\sigma)$  to be the set of copies of  $\sigma$  multiplied by any fraction  $0 < \alpha \leq 1$ , i.e.  $\text{scaled}(\sigma) := \{\alpha \cdot \sigma \mid 0 < \alpha \leq 1\}$ . A footprint  $\sigma_w$  is *scalable w.r.t. a state  $\sigma_A$*  iff either (1)  $\sigma_A$  is compatible with *all* states from  $\text{scaled}(\sigma_w)$ , or (2)  $\sigma_A$  is compatible with *no* state in  $\text{scaled}(\sigma_w)$ . A footprint is *scalable for a wand  $A \ast B$*  iff it is scalable w.r.t. all states that satisfy  $A$ .

<sup>6</sup> In this example,  $\sigma_f$  is incompatible with *all* states that satisfy the left-hand side. However, footprints that are incompatible with only *some* states that satisfy the left-hand side might also render wands not combinable, as we show in App. F.

Intuitively, this means that the footprint does not “jump” between satisfying the wand trivially and having to satisfy the right-hand side. In the above example,  $\sigma_g$  is a scalable footprint for  $w$ , but  $\sigma_f$  is not.

**Making Wands Combinable.** The previous paragraphs show that, even if  $B$  is combinable, the standard wand  $A \multimap B$  is in general not combinable because it can be satisfied by non-scalable footprints. Therefore, we define a novel restricted interpretation for wands that *forces* footprints to be scalable, in the following sense. The restricted interpretation of a wand accepts all scalable footprints, and transforms non-scalable footprints before checking whether they actually satisfy the wand. We call a wand with this restricted interpretation a *combinable wand*, and write  $A \multimap_c B$  to differentiate it from the standard wand  $A \multimap B$ .

For standard wands, *any* state  $\sigma_w$  is a footprint of  $A \multimap B$  iff, for all states  $\sigma_A$  that satisfy  $A$ ,  $\sigma_A \# \sigma_w \Rightarrow \sigma_A \oplus \sigma_w \models B$ . We obtain the definition of combinable wands by replacing  $\sigma_w$  with a (possibly smaller) state  $\mathcal{R}(\sigma_A, \sigma_w)$  that is scalable w.r.t.  $\sigma_A$ .  $\mathcal{R}(\sigma_A, \sigma_w)$  is defined as  $\sigma_w$  if *no* state in  $\text{scaled}(\sigma_w)$  is compatible with any  $\sigma_A$ ; in that case, condition (2) of scalable footprints holds for  $\mathcal{R}(\sigma_A, \sigma_w)$  w.r.t.  $\sigma_A$ . Otherwise,  $\mathcal{R}(\sigma_A, \sigma_w)$  is obtained by removing just enough permissions from  $\sigma_w$  to ensure that *all* states in  $\text{scaled}(\mathcal{R}(\sigma_A, \sigma_w))$  are compatible with  $\sigma_A$ , which ensures that condition (1) holds for  $\mathcal{R}(\sigma_A, \sigma_w)$  w.r.t.  $\sigma_A$ .

To formally define  $\mathcal{R}(\sigma_A, \sigma_w)$ , we fix a concrete separation algebra (formally defined in App. G), whose states are pairs  $(\pi, h)$  of a *permission mask*  $\pi$ , which maps heap locations to fractional permissions, and a partial heap  $h$ , which maps heap locations to values.

**Definition 1.** Let  $(\pi_A, h_A)$  and  $(\pi_w, h_w)$  be two states, and let  $\pi'_w$  be the permission mask such that  $\forall l. \pi'_w(l) = \min(\pi_w(l), 1 - \pi_A(l))$ . Then

$$\mathcal{R}((\pi_A, h_A), (\pi_w, h_w)) = \begin{cases} (\pi_w, h_w) & \text{if } \forall \sigma \in \text{scaled}((\pi_w, h_w)). \neg(\pi_A, h_A) \# \sigma \\ (\pi'_w, h_w) & \text{otherwise} \end{cases}$$

The combinable wand  $A \multimap_c B$  is then interpreted as follows:

$$\sigma_w \models A \multimap_c B \iff (\forall \sigma_A. \sigma_A \models A \wedge \sigma_A \# \mathcal{R}(\sigma_A, \sigma_w) \implies \sigma_A \oplus \mathcal{R}(\sigma_A, \sigma_w) \models B)$$

The following theorem (proved in Isabelle/HOL) shows some key properties of combinable wands.

**Theorem 3.** Let  $B$  be an intuitionistic assertion.

1. If  $B$  is combinable, then  $A \multimap_c B$  is combinable.
2.  $A \multimap_c B \models A \multimap B$ .
3. If  $A$  is a binary assertion, then  $A \multimap_c B$  and  $A \multimap B$  are equivalent.

Property 1 expresses that combinable wands constructed from combinable assertions are combinable, which enables verification methodologies underlying tools such as VerCors and Viper to support flexible combinations of wands and

predicates (as motivated at the start of this section). Property 2 implies that  $A * (A \multimap_c B) \models B$ , that is, combinable wands can be applied like standard wands. Property 3 states that combinable wands pose no restrictions if the left-hand side is binary, that is, if it can be expressed without fractional permissions (formally defined in App. G). For example, the predicate  $\text{Tree}(x)$  from Fig. 1 is binary, which implies that the wands  $\text{Tree}(y) \multimap_c \text{Tree}(x)$  and  $\text{Tree}(y) \multimap \text{Tree}(x)$  are equivalent. This property is an important reason for why combinable wands are expressive enough for practical purposes, as we further evidence in Sec. 5.

Footprints of combinable wands can be automatically inferred by package algorithms built on our package logic. We explain in App. H how to lift the package logic presented in Sec. 3 to handle alternative definitions of allowable wand footprints such as the restrictions imposed by Def. 1.

## 5 Evaluation

We have implemented package algorithms for the standard wands and combinable wands in a custom branch of Viper’s [30] verification condition generator (VCG). Both are based on the package logic described in Sec. 3, adapted to the fractional permission setting. Both algorithms automate the proof search strategy outlined in Sec. 3.5. Viper’s VCG translates Viper programs to Boogie [28] programs. It uses a total-heap semantics of IDF [32], where Viper states include a heap and a permission mask (tracking fractional permission amounts). The heap and mask are represented in Boogie as maps; we also represent witness sets as Boogie maps.

We evaluate our implementations of the package algorithms on Viper’s test suite and compare them to Viper’s implementation of the FIA as presented in Sec. 2.4. Our key findings are that our algorithms (1) enable the verification of almost all correct **package** operations. (2) correctly report **package** operations that are supposed to fail (in contrast to the FIA), and (3) have an acceptable performance overhead compared to the FIA. Moreover, interpreting wands as combinable wands as explained in Sec. 4 has only a minor effect on the results, but correctly rejects attempts to package a non-combinable wand. This finding suggests that verifiers could improve their expressiveness by allowing flexible combinations of wands and predicates with only a minor completeness penalty.

For our evaluation, we considered all 85 files in the test suite for Viper’s VCG that contain at least one **package** operation. From these 85 files, we removed 29 files containing features that our implementation does not yet support. 28 of these 29 files require proof scripts to guide the footprint inference, which are orthogonal to the concerns of this paper (see the appendix App. J for details).

Table 1 gives an overview of our results. These confirm that our algorithms for standard and combinable wands (S-Alg and C-Alg) do not produce false negatives, that is, are sound. In contrast, the FIA does verify an incorrect program (which is similar to the example in Sec. 2.4). While this is only a single unsound example, it is worth emphasizing that (a) it comes from the pre-existing test suite of the tool itself, (b) the unsoundness was not known of until our work, and (c) soundness issues in a program verifier are critical to address; we show how to achieve this.



Algorithm	Expected result	Incorrectly verified	Spurious errors
FIA	55	1	0
S-Alg	51	0	5
C-Alg	48	0	8

**Table 1.** Verification results on our 56 benchmarks with the FIA, our algorithm for standard wands (S-Alg), and for combinable wands (C-Alg). For each algorithm, we report the number of correct verification results, false negatives, and false positives.

Compared with the FIA, our implementation reports a handful of false positives (spurious errors). For S-Alg, 3 out of 5 false positives are caused by missing features of our implementation (such as remembering a subset of the permissions that are inside predicate instances when manipulating predicates); these features could be straightforwardly added in the future. The other 2 false positives are caused by S-Alg’s strategy. In one, the only potential footprint prevents the wand from ever being applied; although technically a false positive, it seems useful to reject the wand and alert the user. The other case is due to a coarse-grained heuristic applied by S-Alg that can be improved.

C-Alg reports the expected result in 48 benchmarks. Importantly, it correctly rejects one wand that indeed does not hold as a combinable wand. 5 of the 8 false positives are identical to those for S-Alg. In the other three benchmarks, the wands still do hold as combinable wands, but further extensions to C-Alg are required to handle them due to technical challenges regarding predicate instances. Once these extensions have been implemented, C-Alg will be as precise as S-Alg, indicating that comparable program verifiers could switch to combinable wands to simply enable sound, flexible combinations with predicates.

To evaluate performance, we ran each of the three implementations 5 times on each of the 56 benchmarks on a Lenovo T480 with 32 GB of RAM and a i7-8550U 1.8 GHz CPU, running on Windows 10. We removed the slowest and fastest time, and then took the mean of the remaining 3 runs. The FIA takes between 1 and 11 seconds per benchmark. On average, S-Alg is 21% slower than the FIA. For 46 of the 56 examples, the increase is less than 30%, and for 3 examples S-Alg is between a factor 2 and 3.4 slower. The overhead is most likely due to the increased complexity of our algorithms, which track more states explicitly and require more quantified axioms in the Boogie encoding. C-Alg is on average 10% slower than S-Alg. We consider the performance overhead of our algorithms to be acceptable, especially since wands occur much more frequently in our benchmarks than in average Viper projects, as judged by existing tests and examples. More representative projects will, thus, incur a much smaller slow-down.

## 6 Related Work

VerCors [3] and Viper [30] are to the best of our knowledge the only automatic SL verifiers that support magic wands. Both employ `package` and `apply` ghost operations. VerCors’ package algorithm requires a user to manually specify a

footprint whereas Viper infers footprints using the FIA, which is unsound as we show in Sec. 2.4. Our package algorithm is as automatic as the FIA but is sound.

Lee and Park [26] develop a sound and complete proof system for SL including the magic wand. Moreover, they derive a decision procedure from their completeness proof for propositional SL. However, more expressive versions of SL (that include e.g. predicates and quantifiers) are undecidable [8] and so this decision procedure cannot be directly applied in the logics employed by program verifiers.

Iris [22] provides a custom proof mode [23] for interactive SL proofs in Coq [12]. Separation logics expressed in Iris support wands and are more expressive than those of automatic SL verifiers at the cost of requiring more user guidance. Packaging a wand in the proof mode requires manually specifying a footprint and proving that the footprint is correct. While tactics can be used in principle to automate parts of this process, there are no specific tactics to infer footprints.

Fractional assertions have been used in various forms [7, 19, 13, 25, 9]. Le and Hobor [25] allow combining two fractional assertions  $A^p$  and  $A^q$  only if  $A$  is *precise* in the SL sense (i.e.  $A$  describes the contents of the heaps in which it holds precisely). To avoid imposing a side condition requiring  $A$  to be precise, Brotherston *et al.* [9] introduce *nominal labels* that can be associated with assertions. If an assertion is split into two fractional assertions, then the same fresh label can be associated with both fractional parts to indicate that they were split from the same assertion. Brotherston *et al.* allow combining two fractional assertions if both assertions are associated with the same label. However, their solution has not been implemented in a verifier and their work does not deal with packaging wands. Our solution also avoids requiring that an assertion is precise and allows combining assertions even if they were not split from the same assertion. Instead of introducing nominal labels, we introduce a light restriction that ensures that wands are always combinable. As a result, general assertions containing combinable wands but no other potentially imprecise connectives (such as disjunction) are combinable. In particular, all assertions employed in verifiers such as VerCors and Viper can be made combinable thanks to our work.

## 7 Conclusion


We presented a package logic that precisely characterises sound package algorithms for automated reasoning about magic wands. Based on this logic, we developed a novel package algorithm that is inspired by an existing approach, but is sound. Moreover, we identified a sufficient criterion for wands to be combinable, such that they can be used flexibly in logics with fractional permissions, and presented a package algorithm for combinable wands. We implemented our solutions in Viper and demonstrated their practical usefulness. The soundness and completeness of our package logic, as well as key properties of combinable wands are all proved in Isabelle/HOL. As future work, we plan to extend the implementation of the two package algorithms described in Sec. 5 by porting various features of the pre-existing FIA implementation. Moreover, we will use our package logic to develop another algorithm for Viper’s symbolic-execution verifier.

## References

1. Astrauskas, V., Müller, P., Poli, F., Summers, A.J.: Leveraging Rust types for modular specification and verification. In: Object-Oriented Programming Systems, Languages, and Applications (OOPSLA). vol. 3, pp. 147:1–147:30. ACM (2019). <https://doi.org/10.1145/3360573>
2. Blom, S., Darabi, S., Huisman, M., Oortwijn, W.: The VerCors tool set: Verification of parallel and concurrent software. In: Polikarpova, N., Schneider, S. (eds.) Integrated Formal Methods (IFM). Lecture Notes in Computer Science, vol. 10510, pp. 102–110. Springer (2007)
3. Blom, S., Darabi, S., Huisman, M., Oortwijn, W.: The VerCors tool set: Verification of parallel and concurrent software. In: Polikarpova, N., Schneider, S. (eds.) Integrated Formal Methods. pp. 102–110. Springer International Publishing, Cham (2017)
4. Blom, S., Huisman, M.: Witnessing the elimination of magic wands. *International Journal on Software Tools for Technology Transfer (STTT)* **17**(6), 757–781 (2015). <https://doi.org/10.1007/s10009-015-0372-3>
5. Boyland, J.: Checking interference with fractional permissions. In: Static Analysis (SAS). Lecture Notes in Computer Science, vol. 2694, pp. 55–72. Springer (2003)
6. Boyland, J.: Checking interference with fractional permissions. In: Cousot, R. (ed.) Static Analysis (SAS). pp. 55–72 (2003)
7. Boyland, J.T.: Semantics of fractional permissions with nesting. *Transactions on Programming Languages and Systems (TOPLAS)* **32**(6), 22:1–22:33 (2010). <https://doi.org/10.1145/1749608.1749611>
8. Brochenin, R., Demri, S., Lozes, E.: On the almighty wand. *Information and Computation* **211**, 106–137 (2012). <https://doi.org/https://doi.org/10.1016/j.ic.2011.12.003>
9. Brotherston, J., Costa, D., Hobor, A., Wickerson, J.: Reasoning over permissions regions in concurrent separation logic. In: Lahiri, S.K., Wang, C. (eds.) Computer Aided Verification (CAV) (2020)
10. Calcagno, C., O’Hearn, P.W., Yang, H.: Local action and abstract separation logic. In: Logic in Computer Science (LICS). pp. 366–375 (2007)
11. Cao, Q., Wang, S., Hobor, A., Appel, A.W.: Proof pearl: Magic wand as frame (2019)
12. Coq Development Team, T.: The Coq Reference Manual, version 8.10 (2019), available electronically at <http://coq.inria.fr/documentation>
13. Dinsdale-Young, T., Dodds, M., Gardner, P., Parkinson, M.J., Vafeiadis, V.: Concurrent abstract predicates. In: European Conference on Object-Oriented Programming (ECOOP). pp. 504–528 (2010). [https://doi.org/10.1007/978-3-642-14107-2\\_24](https://doi.org/10.1007/978-3-642-14107-2_24)
14. Dinsdale-Young, T., da Rocha Pinto, P., Andersen, K.J., Birkedal, L.: Caper - automatic verification for fine-grained concurrency. In: Yang, H. (ed.) European Symposium on Programming (ESOP). Lecture Notes in Computer Science, vol. 10201, pp. 420–447. Springer (2017)
15. Dockins, R., Hobor, A., Appel, A.W.: A fresh look at separation algebras and share accounting. In: Hu, Z. (ed.) Asian Symposium on Programming Languages and Systems (APLAS). pp. 161–177 (2009)
16. Dodds, M., Jagannathan, S., Parkinson, M.J.: Modular reasoning for deterministic parallelism. In: Principles of Programming Languages (POPL) (2011). <https://doi.org/10.1145/1926385.1926416>

17. Haack, C., Hurlin, C.: Resource usage protocols for iterators. *Journal of Object Technology (JOT)* **8**(4), 55–83 (June 2009)
18. Huisman, M., Klebanov, V., Monahan, R.: Verifythis 2012 - A program verification competition. *Int. J. Softw. Tools Technol. Transf.* **17**(6), 647–657 (2015). <https://doi.org/10.1007/s10009-015-0396-8>
19. Jacobs, B., Piessens, F.: Expressive modular fine-grained concurrency specification. In: *Principles of Programming Languages (POPL)*. pp. 271–282 (2011). <https://doi.org/10.1145/1926385.1926417>
20. Jacobs, B., Smans, J., Philippaerts, P., Vogels, F., Penninckx, W., Piessens, F.: VeriFast: A powerful, sound, predictable, fast verifier for C and Java. In: Bobaru, M.G., Havelund, K., Holzmann, G.J., Joshi, R. (eds.) *NASA Formal Methods (NFM)*. *Lecture Notes in Computer Science*, vol. 6617, pp. 41–55. Springer (2011)
21. Jensen, J., Birkedal, L., Sestoft, P.: Modular verification of linked lists with views via separation logic. *Journal of Object Technology (JOT)* **10**, 2: 1–20 (January 2011). <https://doi.org/10.1145/1924520.1924524>
22. Jung, R., Krebbers, R., Jourdan, J., Bizjak, A., Birkedal, L., Dreyer, D.: Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.* **28**, e20 (2018). <https://doi.org/10.1017/S0956796818000151>
23. Krebbers, R., Jourdan, J.H., Jung, R., Tassarotti, J., Kaiser, J.O., Timany, A., Charguéraud, A., Dreyer, D.: Mosel: A general, extensible modal framework for interactive proofs in separation logic **2** (2018). <https://doi.org/10.1145/3236772>
24. Krishnaswami, N.R.: Reasoning about iterators with separation logic. In: *Specification and Verification of Component-Based Systems (SAVCBS)* (2006). <https://doi.org/10.1145/1181195.1181213>
25. Le, X.B., Hobor, A.: Logical reasoning for disjoint permissions. In: Ahmed, A. (ed.) *European Symposium on Programming (ESOP)* (2018)
26. Lee, W., Park, S.: A proof system for separation logic with magic wand. In: *Principles of Programming Languages (POPL)* (2014). <https://doi.org/10.1145/2535838.2535871>
27. Leino, K.R.M., Müller, P., Smans, J.: Verification of concurrent programs with Chalice. In: Aldini, A., Barthe, G., Gorrieri, R. (eds.) *Foundations of Security Analysis and Design V*. *Lecture Notes in Computer Science*, vol. 5705, pp. 195–222. Springer (2009), <http://www.springerlink.com>
28. Leino, K.R.M.: This is Boogie 2, available from <http://research.microsoft.com/en-us/um/people/leino/papers/krml178.pdf>
29. Maeda, T., Sato, H., Yonezawa, A.: Extended alias type system using separating implication. *Workshop on Types in Language Design and Implementation (TLDI)* (2011). <https://doi.org/10.1145/1929553.1929559>
30. Müller, P., Schwerhoff, M., Summers, A.J.: Viper: A verification infrastructure for permission-based reasoning. In: Jobstmann, B., Leino, K.R.M. (eds.) *Verification, Model Checking, and Abstract Interpretation (VMCAI)*. *Lecture Notes in Computer Science*, vol. 9583, pp. 41–62. Springer (2016)
31. Parkinson, M., Bierman, G.: Separation logic and abstraction. In: Palsberg, J., Abadi, M. (eds.) *Principle of Programming Languages (POPL)*. pp. 247–258. ACM (2005)
32. Parkinson, M.J., Summers, A.J.: The relationship between separation logic and implicit dynamic frames. *Logical Methods in Computer Science* **8**(3:01), 1–54 (2012)
33. Piskac, R., Wies, T., Zufferey, D.: GRASShopper—complete heap verification with mixed specifications. In: Ábrahám, E., Havelund, K. (eds.) *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*. *Lecture Notes in Computer Science*, vol. 8413, pp. 124–139. Springer (2014)

34. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: Logic in Computer Science (LICS). pp. 55–74. IEEE (2002)
35. Schwerhoff, M., Summers, A.J.: Lightweight Support for Magic Wands in an Automatic Verifier. In: Boyland, J.T. (ed.) European Conference on Object-Oriented Programming (ECOOP). LIPIcs, vol. 37, pp. 614–638. Schloss Dagstuhl (2015)
36. Smans, J., Jacobs, B., Piessens, F.: Implicit dynamic frames: Combining dynamic frames and separation logic. In: European Conference on Object-Oriented Programming (ECOOP). Lecture Notes in Computer Science, vol. 5653, pp. 148–172. Springer (2009)
37. Tuerk, T.: Local reasoning about while-loops. In: Verified Software: Theories, Tools and Experiments - Theory Workshop (VS-Theory) (2010)
38. Yang, H.: An example of local reasoning in bi pointer logic: the Schorr-Waite graph marking algorithm. *SPACE* **1** (2001)



**Fig. 3.** A visualisation of the footprint inference baseline traversing the right-hand side of  $\text{acc}(x.f) \multimap \text{acc}(x.f) * \text{acc}(y.g)$  in a state  $\sigma$ .  $A$  holds in  $\sigma_A$ . Each small grey box represents a permission to a heap location.


## A Example for the Footprint Inference Attempt

Fig. 3 visualises the FIA for the right-hand-side of the wand  $\text{acc}(x.f) \multimap \text{acc}(x.f) * \text{acc}(y.g)$  in a current state  $\sigma$  that has permissions to  $x.f$  and  $y.g$ .  $\sigma_A$  is an arbitrary state in which the wand’s left-hand-side holds, and thus has permission to  $x.f$ . When traversing the first conjunct  $\text{acc}(x.f)$ , the FIA removes permission to  $x.f$  from  $\sigma_A$  and adds it to  $\sigma_B$ . For the second conjunct  $\text{acc}(y.g)$ , the FIA removes permission to  $y.g$  from  $\sigma$ , since  $\sigma_A$  does not contain any permission to  $y.g$ . Since these were all the requirements from the wand’s right-hand side, the FIA succeeds and the footprint (the permissions taken from  $\sigma$ ) is a state with permission to  $y.g$ , which is a correct footprint.

## B Example for a Package Algorithm Based on the Package Logic

We illustrate in Fig. 4 the algorithm described in Sec. 3.5. Our goal is to package the wand  $\text{acc}(x.f) * (x.f = y \vee x.f = z) \multimap \text{acc}(x.f) * \text{acc}(x.f.g)$ . Recall that the FIA infers an incorrect footprint for this wand, as explained in Sec. 2.4. We assume that the initial outer state  $\sigma_0$  contains permissions to  $y.g$ ,  $z.g$ , and some other resources. The initial context is  $(\sigma_0, S_0)$ , where  $S_0^1$  contains two minimal states that satisfy  $\text{acc}(x.f) * (x.f = y \vee x.f = z)$  (the left-hand side of the wand). The two states in  $S_0^1$  have permission to  $x.f$ , and  $x.f$  contains value  $y$  in one state and  $z$  in the other one. The second element of both pairs in  $S_0$  is the empty state.

We first handle the first conjunct of the right-hand side,  $\text{acc}(x.f)$ . Since both states in  $S_0^1$  satisfy  $\text{acc}(x.f)$ , the first case applies. For the two pairs  $(\sigma_A, \sigma_B) \in S_0$ , we transfer permission to  $x.f$  from  $\sigma_A$  to  $\sigma_B$ , and we obtain the new witness set  $S_1$ . We then handle the second conjunct of the right-hand side,  $\text{acc}(x.f.g)$ . Since no state  $\sigma_A \in S_1^1$  satisfies  $\text{acc}(x.f.g)$  (since the two states in  $S_1^1$  have no permissions anymore), the second case applies. We need to transfer permissions from the outer state  $\sigma_0$  to all states of  $S_1^1$ . Moreover, since  $x.f$  evaluates to  $y$  in one state of  $S_1^1$  and to  $z$  in the other one, we transfer permissions to both  $y.g$  and  $z.g$  to the states of  $S_1^1$ , and we obtain the new



**Fig. 4.** Illustration of how the algorithm described in Sec. 3.5 packages the wand  $\text{acc}(x.f) * (x.f = y \vee x.f = z) \multimap \text{acc}(x.f) * \text{acc}(x.f.g)$ .  $\sigma_0$  represents the program state before the package, and  $\sigma_1$  is the outer state after permissions to  $y.g$  and  $z.g$  have been extracted (which correspond to a footprint of the wand inferred by the algorithm). The  $S_i$  represent witness sets, i.e. sets of pairs of states. Pairs of states are represented as stacks of two states. Finally, permissions to heap locations are represented with light grey rectangles.

context  $(\sigma_1, S_2)$  (where  $\sigma_1$  is  $\sigma_0$  without permissions to  $y.g$  and  $z.g$ ). For each pair of states  $(\sigma_A, \sigma_B) \in S_2$ , we can now transfer permissions from  $\sigma_A$  to  $\sigma_B$  to satisfy  $\text{acc}(x.f.g)$ .

In the end, the footprint inferred is the difference between the final outer state  $\sigma_1$  and the initial outer state  $\sigma_0$ , i.e. the state that contains permissions to (and values of) exactly  $y.g$  and  $z.g$ . Therefore, the algorithm presented above finds the correct footprint for this wand, contrary to the FIA as we explained in Sec. 2.4. After the package the new state of the program is  $\sigma_1$ , in which we record the wand instance and proceed with the verification of the subsequent statements.

## C Separation Algebra and Assertions

In this section, we formally define the separation algebra and the assertion language that our package logic builds on.

**Definition 2.** A separation algebra is a quintuple  $(\Sigma, \oplus, e, |\_|, \text{stable})$ , where:

1.  $\Sigma$  is a set of states,  $\oplus$  is a partial addition on  $\Sigma$  that is commutative and associative, and  $e$  is the neutral element of  $\oplus$ .
2.  $|\_|$  (the core) is a function from  $\Sigma$  to  $\Sigma$ .

3. *stable* is a function from  $\Sigma$  to Booleans.
4. The following axioms are satisfied:
  - (a)  $\forall x \in \Sigma. x = x \oplus |x| \wedge |x| = |x| \oplus |x|$
  - (b)  $\forall x, c \in \Sigma. x = x \oplus c \implies (\exists r \in \Sigma. |x| = c \oplus r)$
  - (c)  $\forall a, b, c \in \Sigma. c = a \oplus b \implies |c| = |a| \oplus |b|$
  - (d)  $\mathbf{stable}(e) \wedge (\forall a, b, c \in \Sigma. c = a \oplus b \wedge \mathbf{stable}(a) \wedge \mathbf{stable}(b) \implies \mathbf{stable}(c))$
  - (e)  $\forall a, b, c \in \Sigma. c = a \oplus b \wedge c = c \oplus c \implies a = a \oplus a$  (positivity)
  - (f)  $\forall a, b, x, y \in \Sigma. a = b \oplus x \wedge a = b \oplus y \wedge |x| = |y| \implies x = y$  (cancellativity)

$|\sigma|$  represents the pure (duplicable) resources contained in the state  $\sigma$ . A state  $\sigma$  is *pure* iff  $\sigma = \sigma \oplus \sigma$ , that is, iff it contains only pure information. In an implicit dynamic frame setting, the values stored in the heap are considered pure resources (and thus duplicable), but ownership is not (since it cannot be duplicated). Pure resources can also include local variables, which can be used to represent SL assertions of the form  $\exists v. x.f \mapsto v * A$  as separating conjunctions without existentials (where the value of  $v$  is represented as a duplicable resource).

In an implicit dynamic frame setting, a state might contain a value of a heap location even if it does not have permission to this heap location. This is necessary to define the evaluation of the separating conjunction, as explained in Sec. 2.1. A state  $\sigma$  is *stable*, written  $\mathbf{stable}(\sigma)$ , iff  $\sigma$  only contains values of heaps locations to which it has some permission. When packaging a wand in IDF, we only consider stable footprints. Not doing so would give a different semantics to IDF wands and SL wands. Consider for example the SL assertion  $A := x.f \mapsto 5 * (x.f \mapsto \_ * x.f \mapsto 5)$ , which can also be expressed in IDF [32]. In SL,  $A$  is equivalent to false. In IDF,  $A$  is equivalent to false only if footprints are enforced to be stable. Indeed, if we would allow non-stable states to be footprints of wands, then a state with no permission at all but in which  $x.f$  contains the value 5 would be a valid footprint of  $x.f \mapsto \_ * x.f \mapsto 5$ . Thus, a state in which  $x.f \mapsto 5$  holds would satisfy  $A$ .

Axioms (a) and (b) state that the core of a state is its maximal pure part, while axiom (c) requires the function  $|\_ |$  to be linear. Axiom (d) requires that the unit is stable, and that the sum of two stable states is also stable. Finally, the positivity axiom states that any state smaller than a pure state has to be pure, and the cancellativity axiom states that the algebra is cancellative for non-pure resources.

Using this separation algebra, we define the following partial order on elements of  $\Sigma$ : A state  $\sigma_2 \in \Sigma$  is larger than another state  $\sigma_1 \in \Sigma$ , written  $\sigma_2 \succeq \sigma_1$ , iff  $\exists \sigma_r \in \Sigma. \sigma_2 = \sigma_1 \oplus \sigma_r$ . Moreover, we write  $\sigma_1 \# \sigma_2$  iff  $\sigma_1 \oplus \sigma_2$  is defined. Finally, we define a subtraction operator,  $\sigma_A \ominus \sigma_B$ , which corresponds to the largest state  $\sigma_r$  such that  $\sigma_A = \sigma_B \oplus \sigma_r$  if  $\sigma_A \succeq \sigma_B$  (the other case is not relevant).

In order to enable a package algorithm to deconstruct the separating conjunctions and implications of the right-hand side of a wand, and to extract the footprint piecewise as illustrated in App. B, we consider an assertion language that contains the separating conjunction and the implication connectives. This allows us to write logical rules that only apply to a separating conjunction and to an implication, respectively. Moreover, to be as general as possible, we do



not fix the other connectives of the assertion language. Thus, the third type of assertion we consider in our assertion language is the general type of *semantic assertions*, i.e. functions from  $\Sigma$  to Booleans. This third type represents any SL assertion that is neither a separating conjunction nor an implication. In particular, assertions such as  $\mathbf{x.f} = 5$ ,  $\mathbf{acc}(\mathbf{x.f})$ , abstract predicates (such as  $\mathbf{Tree}(\mathbf{x})$ ) or magic wands are represented as semantic assertions.

**Definition 3.** *Let  $\mathcal{B}$  range over semantic assertions, i.e. functions from  $\Sigma$  to booleans. Assertions (ranged over by  $A$ ) are defined as follows:*

$$A = A * A \mid \mathcal{B} \Rightarrow A \mid \mathcal{B}$$

For a state  $\sigma \in \Sigma$  and an assertion  $A$ , we write  $\sigma \models A$  to say that  $\sigma$  satisfies  $A$ , and define it as follows:

$$\begin{aligned} \sigma \models A_1 * A_2 &\iff (\exists \sigma_1, \sigma_2. \sigma = \sigma_1 \oplus \sigma_2 \wedge \sigma_1 \models A_1 \wedge \sigma_2 \models A_2) \\ \sigma \models \mathcal{B} \Rightarrow A &\iff (\mathcal{B}(\sigma) \implies \sigma \models A) \\ \sigma \models \mathcal{B} &\iff \mathcal{B}(\sigma) \end{aligned}$$

App. D explains how this extend this assertion language and the logic to handle other connectives, such as the disjunction or the normal conjunction.

This assertion language is too permissive for our purpose. In particular, we only want to consider assertions that are well-formed, that is, if an assertion is well-defined and holds in a state, then adding pure resources to this state should not render the assertion false. Informally, for an IDF assertion, well-formed corresponds to being self-framing. We achieve this with monotonicity constraints: A semantic assertion  $\mathcal{B}$  which appears on the left-hand side of an implication should stay false if we add pure resources to a state in which it is false, and semantic assertions which are not on the left-hand side of an implication should behave in the opposite way.

**Definition 4.** *A semantic assertion  $\mathcal{B}$  is monotonically pure, written  $\mathit{monoPure}(\mathcal{B})$ , iff  $(\forall \sigma, \sigma_p \in \Sigma. \sigma_p \text{ is pure} \wedge \mathcal{B}(\sigma) \wedge \sigma \# \sigma_p \implies \mathcal{B}(\sigma \oplus \sigma_p))$ .*

We write  $\mathit{wf}(A)$  to say that the assertion  $A$  is well-formed. It is defined as follows:

$$\begin{aligned} \mathit{wf}(A_1 * A_2) &\iff \mathit{wf}(A_1) \wedge \mathit{wf}(A_2) \\ \mathit{wf}(\mathcal{B} \Rightarrow A) &\iff \mathit{monoPure}(\neg \mathcal{B}) \wedge \mathit{wf}(A) \\ \mathit{wf}(\mathcal{B}) &\iff \mathit{monoPure}(\mathcal{B}) \end{aligned}$$

## D Extending the Logic

The framework and the logic presented in Sec. 3 operate only on a simple language for assertions: An assertion is either a separating conjunction (star) of two assertions, an implication of a pure semantic assertion on the left-hand side and an assertion on the right-hand side, or a semantic assertion. Note that

any assertion can be represented in this framework, since any assertion can be represented as a semantic assertion. The star and the implication connectives that our framework provides enables a package algorithm to (1) deconstruct an assertion with these connectives, and (2) apply the rule *Extract* with some heuristic at the “leaves” of this assertion.

Thus, while other connectives such as the disjunction or the non-separating conjunction can still be handled using semantic assertions, one might want to extend the assertion language along with the logic such that the algorithm can deconstruct these assertions even deeper. In the following, we describe how one can extend the set of rules from Fig. 2 to handle disjunctions and non-separating conjunctions.

*Disjunctions.* To satisfy the disjunction  $A \vee B$ , a state must either satisfy  $A$  or satisfy  $B$ . When dealing with a set of extended states, a package algorithm can choose which extended states must satisfy  $A$ , and which ones must satisfy  $B$ . More precisely, a rule to handle disjunctions could proceed in five steps:

1. Separate the witness set into the set  $S_0^1$  of extended states that must prove  $A$ , and the set  $S_0^B$  of extended states that must prove  $B$ .
2. Use the rules to handle the assertion  $A$  with the witness set  $S_0^1$ . This gives a new witness set  $S_1^1$ .
3. In step 2, the algorithm might have added a partial footprint to the witness set  $S_0^1$  to get the new witness set  $S_1^1$ . Thus, this partial footprint should be added to  $S_0^B$ , which gives a new witness  $S_1^B$ .
4. Use the rules to handle the assertion  $B$  with the witness set  $S_1^B$ , which gives a new witness set  $S_2^B$ .
5. In step 4, the algorithm might have added a partial footprint to the witness set  $S_1^B$  to get the new witness set  $S_2^B$ . Thus, this partial footprint should also be added to  $S_1^1$ , which gives a new witness  $S_2^1$ .
6. The final witness set is  $S_2^1 \cup S_2^B$ .

*Non-separating conjunctions.* To satisfy the non-separating conjunction  $A \wedge B$ , a state must satisfy  $A$  and  $B$ . Thus, the idea in this case is to first use the rules to satisfy  $A$ , then “reset” the states and use the rules to satisfy  $B$ , and finally take the “union” of these states. More precisely,

1. Use the rules to handle the assertion  $A$  with the initial witness set  $S_0$ . This gives a new witness set  $S_1$ .
2. Record, for each extended state, the resources which have been added to its second element to go from  $S_0$  to  $S_1$ , and then transfer back these resources from the second element to the first element of the extended state. This gives a new witness set  $S_2$ .
3. Use the rules to handle the assertion  $B$  with the initial witness set  $S_2$ . This gives a new witness set  $S_3$ .
4. For each extended state of  $S_3$ , consider the state  $\sigma_r$  which has been added to its second element to go from  $S_2$  to  $S_3$ . If  $\sigma_r \models A$ , then do not modify this extended state. If  $\sigma_r \not\models A$ , then transfer another state  $\sigma'_r$  from the first

element to the second element, such that  $\sigma_r \oplus \sigma'_r \models A$ . These transformations yield the final witness set.

## E Example of a Derivation Using the Package Logic

We illustrate how these rules can be used to package the wand from Sec. 3.1,  $w := \text{acc}(\mathbf{x.f}) * (\mathbf{x.f} = \mathbf{y} \vee \mathbf{x.f} = \mathbf{z}) \multimap \text{acc}(\mathbf{x.f}) * \text{acc}(\mathbf{x.f.g})$ . This proof corresponds to the strategy illustrated in App. B. We omit the path condition since it is always the trivial condition  $(\lambda\sigma. \top)$ . Assume that the outer state  $\sigma_0$  is the addition of  $\sigma_{yz}$ , a state that contains permission to  $\mathbf{y.g}$  and  $\mathbf{z.g}$ , and  $\sigma_1$ .  $S_0 := \{(\sigma_A, e) \mid \sigma_A \in \Sigma \wedge \sigma_A \models \text{acc}(\mathbf{x.f}) * (\mathbf{x.f} = \mathbf{y} \vee \mathbf{x.f} = \mathbf{z})\}$  is the initial witness set. We show below a part of a proof that  $\langle \text{acc}(\mathbf{x.f}) * \text{acc}(\mathbf{x.f.g}), (\sigma_0, S_0) \rangle \rightsquigarrow (\sigma_1, S_3)$  is correct, and thus that  $\sigma_{yz}$  is a correct footprint of the wand  $w$  (since  $\sigma_0 = \sigma_1 \oplus \sigma_{yz}$ ):

$$\frac{\frac{\dots}{\langle \text{acc}(\mathbf{x.f}), (\sigma_0, S_0) \rangle \rightsquigarrow (\sigma, S_1)} \textit{Atom} \quad \frac{\frac{\dots}{\langle \text{acc}(\mathbf{x.f.g}), (\sigma_1, S_2) \rangle \rightsquigarrow (\sigma_1, S_3)} \textit{Atom} \quad \sigma_0 = \sigma_{yz} \oplus \sigma_1 \quad \text{stable}(\sigma_{yz}) \quad S_2 = \dots}{\langle \text{acc}(\mathbf{x.f.g}), (\sigma_0, S_1) \rangle \rightsquigarrow (\sigma_1, S_3)} \textit{Extract}}{\langle \text{acc}(\mathbf{x.f}) * \text{acc}(\mathbf{x.f.g}), (\sigma_0, S_0) \rangle \rightsquigarrow (\sigma_1, S_3)} \textit{Star}$$

The derivation can be read from bottom to top and from left to right. Using the rule *Star*, we first split the assertion  $\text{acc}(\mathbf{x.f}) * \text{acc}(\mathbf{x.f.g})$  into its two conjuncts,  $\text{acc}(\mathbf{x.f})$  (on the left) and  $\text{acc}(\mathbf{x.f.g})$  (on the right). We then handle  $\text{acc}(\mathbf{x.f})$  using the rule *Atom*.  $\text{acc}(\mathbf{x.f})$  holds in the first element of each pair of  $S_0$ , since any state that satisfies the left-hand side of the wand has permission to  $\mathbf{x.f}$ . Therefore, we use the rule *Atom* with a *choice* function that always chooses the relevant state with exactly full permission to  $\mathbf{x.f}$ .  $S_1$  is the updated witness set where this permission to  $\mathbf{x.f}$  has been transferred from the first to the second element of each pair of states. Next, we want to handle  $\text{acc}(\mathbf{x.f.g})$  using the rule *Atom*. However, we cannot do this directly from the witness set  $S_1$ : We know that, for each  $(\sigma_A, \sigma_B) \in S_1$ ,  $\mathbf{x.f.g}$  evaluated in  $\sigma_A$  is either  $\mathbf{y}$  or  $\mathbf{z}$ , but  $\sigma_A$  does not have any permission to either  $\mathbf{y.g}$  or  $\mathbf{z.g}$ . Thus, we transfer the permissions to both  $\mathbf{y.g}$  and  $\mathbf{z.g}$  from the outer state  $\sigma_0$  to all states of  $S_1^1$ , using the rule *Extract*, which results in the new context  $(\sigma_1, S_2)$ . Finally, we apply the rule *Atom* to prove  $\langle \text{acc}(\mathbf{x.f.g}), (\sigma_1, S_2) \rangle \rightsquigarrow (\sigma_1, S_3)$ , where the *choice* function chooses for each pair the corresponding state that contains full permission to  $\mathbf{x.f.g}$ .

This example follows a general proof search strategy, which deconstructs the right-hand side using the rules *Star* and *Implication*. When we arrive at a semantic assertion, we first try to prove it directly from the witness set (using the rule *Atom*), as we did to prove  $\text{acc}(\mathbf{x.f})$ . If we cannot prove this semantic assertion directly, we use the rule *Extract* to extract the necessary resources from the outer state before using the rule *Atom*, as we did to prove  $\text{acc}(\mathbf{x.f.g})$ . This general proof search strategy corresponds to the package algorithm described in Sec. 3.5. Different heuristics in when and how to use the rule *Extract* lead to

different proof strategies, and thus to different package algorithms, as explained in Sec. 3.1. Given that our logic is sound, any package algorithm that corresponds to a proof search strategy in the package logic is sound.

## F Example of a Wand that is not Combinable

In Sec. 4, we show that the wand  $\text{acc}(\mathbf{x.f}, 1/2) \multimap \text{acc}(\mathbf{x.g})$  is not combinable, because of a footprint that is incompatible with *all* states that satisfy  $\text{acc}(\mathbf{x.f}, 1/2)$ , but that becomes compatible with some when scaled down (by half). In this section, we show a wand that is incompatible with *some* states that satisfy the left-hand side of the wand, to illustrate that this is still an issue.

Consider, the wand  $w' := \text{acc}(\mathbf{x.f}) * (\mathbf{x.f} = \mathbf{y} \vee \mathbf{x.f} = \mathbf{z}) * \text{acc}(\mathbf{x.f.g}, 1/2) \multimap \text{acc}(\mathbf{y.g})$ .  $w'$  is not combinable. It is straightforward to see that  $\text{acc}(\mathbf{y.g}) \models w'$ . Moreover,  $\text{acc}(\mathbf{y.g}, 1/2) * \text{acc}(\mathbf{z.g}) \models w'$ . Indeed,  $\text{acc}(\mathbf{z.g})$  combined with  $(\mathbf{x.f} = \mathbf{y} \vee \mathbf{x.f} = \mathbf{z}) * \text{acc}(\mathbf{x.f.g}, 1/2)$  implies that  $\mathbf{x.f} = \mathbf{y}$ , and  $\text{acc}(\mathbf{y.g}, 1/2)$  combined with  $\mathbf{x.f} = \mathbf{y}$  and  $\text{acc}(\mathbf{x.f.g}, 1/2)$  entails the right-hand side  $\text{acc}(\mathbf{y.g})$ . However,  $\text{acc}(\mathbf{y.g})^{0.5} * (\text{acc}(\mathbf{y.g}, 1/2) * \text{acc}(\mathbf{z.g}))^{0.5} \equiv \text{acc}(\mathbf{y.g}, 3/4) * \text{acc}(\mathbf{z.g}, 1/2) \not\models w'$ , since  $\mathbf{x.f} = \mathbf{z}$  is now possible. Footprints satisfying  $\text{acc}(\mathbf{y.g})$  are scalable. However, footprints that only satisfy  $\text{acc}(\mathbf{y.g}, 1/2) * \text{acc}(\mathbf{z.g})$  are not scalable.

## G A State Model for Fractional Permissions

We define in this section an implicit dynamic frame state model with fractional permissions, to instantiate the separation algebra as described in Def. 2. Moreover, we define the meaning of a *binary* assertion.

### Definition 5. State model.

Let  $L$  be a set of heap locations which contains a special element *null*, and let  $V$  be a set of values.

- A state is a pair  $(\pi, h)$  of a permission mask  $\pi$  and a partial heap  $h$ , where
- $\pi : L \rightarrow \mathbb{Q} \cap [0, 1]$  maps each heap location to a fractional permission between 0 and 1 included, and
- $h : L \rightarrow V$  is a partial mapping from heap locations to values.

A state  $(\pi, h)$  is valid iff (1)  $\pi(\text{null}) = 0$  and  $\forall l \in L. \pi(l) > 0 \implies h(l)$  is defined. (1) ensures that having ownership of a heap location implies that this heap location is not null, while (2) ensures that the values of all heap locations owned are defined.  $\Sigma$  is defined as the set of all valid states.

This state model corresponds to a separation algebra:

**Definition 6.** Given two valid states  $(\pi_1, h_1)$  and  $(\pi_2, h_2)$ , the addition  $(\pi_1, h_1) \oplus (\pi_2, h_2)$  is defined iff (1)  $h_1$  and  $h_2$  agree on heap locations where they are both defined and (2)  $\forall l \in L. \pi_1(l) + \pi_2(l) \leq 1$ . In this case,  $(\pi_1, h_1) \oplus (\pi_2, h_2) = (\pi_1 + \pi_2, h_1 \cup h_2)$  is a valid state.

The empty state  $e$  is defined as  $(\lambda_{\cdot}.0, \emptyset)$ . The core of a state  $(\pi, h)$  is defined as  $|\!(\pi, h)\!| = (\lambda_{\cdot}.0, h)$ . A state is stable iff  $\forall l \in L. \pi(l) > 0 \iff h(l)$  is defined.  $(\Sigma, \oplus, e, |\!_{\cdot}\!|, \text{stable})$  defines a separation algebra.

We define the partial multiplication of a state by a positive rational as follows:

**Definition 7.** Let  $\alpha \in \mathbb{Q}^+$  be a positive rational, and  $(\pi, h) \in \Sigma$  be a valid state. The product  $\alpha \odot (\pi, h)$  is defined iff  $\forall l \in L. \alpha \times \pi(l) \leq 1$ . In this case,  $\alpha \odot (\pi, h) := (\lambda l. \alpha \times \pi(l), h)$  (which is a valid state).

**Definition 8.** We define the binary restriction of a permission mask  $\pi$  as follows:

$$\text{bin}(\pi)(l) = \begin{cases} 1 & \text{if } \pi(l) = 1 \\ 0 & \text{otherwise} \end{cases}$$

An assertion  $A$  is binary iff  $\forall (\pi, h) \in \langle A \rangle. (\text{bin}(\pi), h) \in \langle A \rangle$

## H Leveraging the Logic

The definition of the combinable wand  $A \multimap_c B$  in Sec. 4 corresponds to the normal definition of a magic wand, except that the footprint is transformed before being combined with states that satisfy  $A$ . We generalise this pattern with the concept of *monotonic transformers*. A *transformer* is a function  $t$  that transforms a state  $\sigma$  into the state  $t(\sigma)$ . It is monotonic iff  $\forall \sigma_1, \sigma_2. \sigma_2 \succeq \sigma_1 \implies t(\sigma_2) \succeq t(\sigma_1)$ . In the case of combinable wands, the function  $\lambda \sigma. \mathcal{R}(\sigma_A, \sigma)$  is a monotonic transformer, for each  $\sigma_A$  that satisfies  $A$ . In the following, we explain how to lift the package logic such that it is sound and complete w.r.t. to the following wand's definition:

$$\sigma_w \models A \multimap_{\mathcal{T}} B \iff (\forall \sigma_A. \sigma_A \models A \wedge \sigma_A \# \mathcal{T}(\sigma_A, \sigma_w) \Rightarrow \sigma_A \oplus \mathcal{T}(\sigma_A, \sigma_w) \models B)$$

where  $\lambda \sigma. \mathcal{T}(\sigma_A, \sigma)$  is a monotonic transformer for each  $\sigma_A$  that satisfies  $A$ . Note that we get the definition of the usual wand by setting  $\mathcal{T}(\sigma_A, \sigma) = \sigma$  for all  $\sigma_A$  and  $\sigma$ .

The witness set is lifted from a set of pairs of states  $(\sigma_A, \sigma_B)$  to a set of tuples  $(\sigma_A, \sigma_B, t)$ , where  $t$  is the monotonic transformer associated to  $\sigma_A$ . The initial witness set is thus  $\{(\sigma_A, e, \lambda \sigma. \mathcal{T}(\sigma_A, \sigma)) \mid \sigma_A \models A\}$ . We also need to modify the rule *Extract*, such that we combine a *transformed* version of  $\sigma_w$  to elements of the witness set (recall that  $\sigma_w$  represents the permissions we extract from the outer state). Consider a triple  $(\sigma_A, \sigma_B, t)$  from the witness set. We cannot simply combine  $\sigma_A$  with  $t(\sigma_w)$ , because the footprint might be extracted piecewise in the package logic, and the transformer  $t$  is only applied to the *complete* footprint in the above definition. Therefore, we need to compute the *part* of the transformed footprint that we need to combine with  $\sigma_A$ . To do this, we need to keep track of the current footprint that has been extracted so far. If  $\sigma_f$  is the footprint extracted so far, and  $\sigma_w$  is the additional part we want to extract from the outer state, the state  $\sigma_A$  must be combined with  $t(\sigma_f \oplus \sigma_w) \ominus t(\sigma_f)$ . We subtract  $t(\sigma_f)$  from  $t(\sigma_f \oplus \sigma_w)$ , since  $t(\sigma_f)$  has already been added to this tuple. In order to keep track of the footprint  $\sigma_f$  extracted so far, we extend contexts from a

```

package( $\sigma_0, A \multimap B, ps$ ) =
 $S_0 \leftarrow \{(\sigma_A, e) \mid \sigma_A \in \text{consleft} - \text{handside}(T_0, \bar{A}, \bar{A})\} * A_2 \rightarrow$ 
 $(\sigma_1, S_1) \leftarrow \text{handleProofScript}(\sigma_0, S_0, ps)$ 
 $(\sigma_2, S_2) \leftarrow \text{proveright} - \text{handside}((\sigma_1, S_1), \top, \bar{B})$ 
 $\text{return } \sigma_2$ 

consleft - handside( $T, pc, A$ ) =
 $\text{return } \text{consleft} - \text{handside}(\text{consleft} - \text{handside}(T, pc, A_1), pc, A_2)$ 
 $\text{return } \text{consleft} - \text{handside}(T, pc \wedge b, A)$ 
 $A \text{ is } b \text{ where } b \text{ is pure} \rightarrow$ 
 $\text{return } \{\sigma_A \mid \sigma_A \in T \wedge \sigma_A \models b\}$ 
 $A \text{ is } r \text{ where } r \text{ is a resource} \rightarrow$ 
 $\text{return } \{\sigma_A \oplus R(\sigma_A, r) \mid \sigma_A \in T \wedge \sigma_A \# R(\sigma_A, r)\}$ 

proveright - handside( $(\sigma, S), pc, B$ ) =
 $B \text{ is } B_1 * B_2 \rightarrow \text{return } \text{proveright} - \text{handside}(\text{proveright} - \text{handside}((\sigma, S), pc, B_1), pc, B_2)$ 
 $B \text{ is } b \Rightarrow B \rightarrow \text{return } \text{proveright} - \text{handside}(T, pc \wedge b, B)$ 
 $B \text{ is } b \text{ where } b \text{ is pure} \rightarrow \text{assert}(\forall \sigma_A \in S^1. pc(\sigma_A) \Rightarrow (\sigma_A \models b)); \text{return } (\sigma, S)$ 
 $B \text{ is } r \text{ where } r \text{ is a resource} \rightarrow$ 
 $\text{assert}(\forall \sigma_A \in S^1. \sigma_A \# \sigma \Rightarrow \sigma_A \oplus \sigma \succeq R(\sigma_A, r)); (\sigma', S') \leftarrow (\sigma, S)$ 
 $\text{if } \neg(\forall \sigma_A \in S^1. \sigma_A \succeq R(\sigma_A, r)) \text{ then}$ 
  Compute a minimal  $\sigma_w$  s.t.  $\sigma \succeq \sigma_w$  and  $\forall \sigma_A \in S^1. \sigma_A \# \sigma_w \Rightarrow \sigma_A \oplus \sigma_w \succeq R(\sigma_A, r)$ 
 $(\sigma', S') \leftarrow (\sigma \ominus \sigma_w, \{(\sigma_A \oplus \sigma_w, \sigma_B) \mid (\sigma_A, \sigma_B) \in S \wedge (\sigma_A \oplus \sigma_B) \# \sigma_w\})$ 
 $\text{fi}$ 
 $\text{return } (\sigma', \{(\sigma_A \ominus R(\sigma_A, r), \sigma_B \oplus R(\sigma_A, r)) \mid (\sigma_A, \sigma_B) \in S'\})$ 

```

**Fig. 5.** High-level representation of the package algorithm we have implemented in Viper’s VCG to compute standard wands. *package* is the main function, *consleft* - *handside* constructs a set of minimal states that satisfy an assertion, and *proveright* - *handside* automates a proof search in the package logic.

pair of a program state  $\sigma$  and a witness set  $S$  to tuples  $(\sigma, S, \sigma_f)$ . Finally, the current footprint is updated to be  $\sigma_f \oplus \sigma_w$  in the rule *Extract*. We have proven in Isabelle/HOL that this lifted logic is sound and complete for the above wand’s definition.

## I Automation

For the sake of presentation, we only discuss here the algorithm that packages standard wands. It is straightforward to adapt it to compute combinable wands, by following the approach described in App. H. Fig. 5 presents, on a high-level, the algorithm we have implemented in Viper’s VCG. Viper’s VCG uses a total-heap semantics of IDF [32], where Viper states (ignoring local variables) consist of a heap and a permission mask (mapping resources to the held ownership amounts). The heap and the mask are represented in Boogie with maps. Based on this representation of Viper states, we can represent sets of states (in the case of *consleft* - *handside* below) and witness sets (in the case of *handleProofScript* and *proveright* - *handside* below) with Boogie maps, which allows us to flexibly manipulate these sets.

The main function, *package*, takes as input a program state  $\sigma_0$ , a wand  $A \multimap B$ , and a *proof script*  $ps$ . We ignore proof scripts here since they are orthogonal to the automation of the proof search, but we explain what they are in App. J. The *package* function calls the function *consleft* - *handside*, which creates a minimal set of states that satisfy  $A$ , to create the initial witness set  $S_0$ . It then calls the function *proveright* - *handside*, which automates a proof search in the package logic, to extract a footprint of  $A \multimap B$  from  $\sigma_0$ . The *package* function finally returns

the program state  $\sigma_2$ , which corresponds to the state  $\sigma_0$  to which a footprint of  $A * B$  has been subtracted. After the package algorithm has successfully executed, the verifier can record an instance of the wand  $A * B$  in  $\sigma_2$  to get the new program state.

The call  $\text{consleft} - \text{handside}(T_0, \top, A)$  constructs a set  $T$  of minimal states that satisfy the assertion  $A$ .  $T_0$  represents a set of “empty” states, i.e. states with no permissions but a total heap. The functions  $\text{consleft} - \text{handside}$  and  $\text{proveright} - \text{handside}$  work similarly to each other, traversing the assertion they receive as input. In particular, both  $\text{consleft} - \text{handside}$  and  $\text{proveright} - \text{handside}$  pattern match the assertion, which gives rise to four cases. If the assertion is a separating conjunction, both functions handle first the first conjunct and then the second conjunct. If it is an implication, both functions syntactically conjoin the left-hand side of the implication to their path condition  $pc$ . In the case of  $\text{proveright} - \text{handside}$ , the separating conjunction case corresponds to the rule *Star* from the package logic, and the implication case to the rule *Implication*.

Finally, both functions distinguish pure assertions from assertions that correspond to resources. Resource assertions in Viper correspond to permissions to heap locations (e.g.  $\text{acc}(\mathbf{x.f})$ ), to predicates (e.g.  $\text{Tree}(\mathbf{x})$ ), or to magic wands. Pure assertions are assertions that do not contain resources, such as  $\mathbf{x.f} = 5$ . In the case of a pure assertion  $b$ ,  $\text{consleft} - \text{handside}$  filters out the states that do not satisfy  $b$ , while  $\text{proveright} - \text{handside}$  asserts that all elements of  $S^1$  satisfy  $b$ . The latter corresponds to an application of the rule *Atom*.<sup>7</sup>

To handle resource assertions, we use the notation  $R(\sigma_A, r)$ , which corresponds to a minimal state that satisfies the resource assertion  $r$  in the state  $\sigma_A$ . We need to evaluate  $r$  in  $\sigma_A$  because  $r$  might be heap-dependent, for example  $r$  could be  $\text{acc}(\mathbf{x.f.g})$  or  $\text{Tree}(\mathbf{x.left})$ . In the case of a resource assertion  $r$ ,  $\text{consleft} - \text{handside}$  combines all states of the set  $T$  with a minimal state that satisfies  $r$ , while  $\text{proveright} - \text{handside}$  applies the following strategy: If all states of  $S^1$  satisfy  $r$  (in which case the if-branch is not entered),  $\text{proveright} - \text{handside}$  directly applies the rule *Atom*. In this is not the case, then  $\text{proveright} - \text{handside}$  first applies the rule *Extract* by computing a minimal state  $\sigma_w$  to extract from the outer state  $\sigma$ , which corresponds to the if-branch, and then the rule *Atom*. Finally, the initial **assert** statement checks that all states of  $S^1$  combined with the outer state satisfy  $r$ , which ensures that the following code corresponds to correct applications of rules from the package logic.

## J Proof scripts

A proof script is a program statement that helps the package algorithm infer or check a footprint. They are mainly useful when one must manipulate predicate instances or magic wand instances in order to infer a footprint, since complete

<sup>7</sup> In this case, the witness set  $S$  is not modified because, for each  $(\sigma_A, \sigma_B) \in S$ ,  $\text{choice}(\sigma_A, \sigma_B)$  (recall that  $\text{choice}$  is a parameter of the rule *Atom*) corresponds to pure resources (as defined in App. G) that are already present in  $\sigma_A$  and  $\sigma_B$ , and thus  $\sigma_A \ominus \text{choice}(\sigma_A, \sigma_B) = \sigma_A$ , and  $\sigma_B \oplus \text{choice}(\sigma_A, \sigma_B) = \sigma_B$ .

automation in such cases is infeasible. Both Viper and VerCors support proof scripts. Since Viper’s and our package algorithm infer footprints, one must provide less elaborate proof scripts than for VerCors.

A package algorithm executes all program statements in a proof script before considering the wand’s right-hand side. Executing a proof script is similar to justifying the right-hand side: permission from the wand’s left-hand side or the current state must be potentially used to do so.

In our implementation, we support the following inductively defined proof scripts:

$$P = \mathbf{assert} A \mid \mathbf{fold} Q(x) \mid \mathbf{unfold} Q(x) \mid \mathbf{apply} A \multimap A \mid \\ P; P \mid \mathbf{if}(b) \{P\} \mathbf{else} \{P\}$$

where  $Q(x)$  is a predicate instance,  $b$  is a boolean expression, and  $A$  is an assertion.

The proof script  $\mathbf{assert} A$  forces the algorithm to justify assertion  $A$ , which can be used to direct the algorithm towards a specific footprint.  $\mathbf{fold} Q(x)$  forces the algorithm to justify the permissions in the body of  $Q(x)$  and to exchange them for the predicate instance  $Q(x)$ , while  $\mathbf{unfold} Q(x)$  does the opposite.  $\mathbf{apply} A \multimap B$  forces the algorithm to apply the wand. Finally, proof scripts can be composed sequentially or one can be put under a conditional.

The **package** in line 10 of Fig. 1 requires the following proof script for Viper and for our package algorithm

$$\mathbf{fold} \text{Tree}(y0); \mathbf{apply} \text{Tree}(y0) \multimap \text{Tree}(x)$$

to successfully infer a footprint, where  $y0$  is an auxiliary variable that contains the value of  $y$  at the beginning of the loop iteration. Note that this proof script does not explicitly specify the footprint. To execute the two operations, the package algorithms still must remove the necessary permissions from the left-hand side or the current state. In this case, the first **fold** statement forces the algorithm to select the permissions for  $y0$  as part of the footprint and the **apply** statement forces the algorithm to select the applied wand instance as part of the footprint (where the left-hand side is obtained after executing the **fold** statement).

In VerCors, the proof script for the **package** in line 10 additionally requires **assert** statements (called **use** statements in VerCors) for all the permissions in the footprint, since VerCors does not infer footprints. Moreover, one needs to add the statement  $\mathbf{assert} y0.\text{left} = y$ , since VerCors cannot infer that  $y$  must be the left node of  $y0$ .