



A Decision Procedure for Guarded Separation Logic: Complete Entailment Checking for Separation Logic with Inductive Definitions

Matheja, Christoph; Pagel, Jens; Zuleger, Florian

Published in:
ACM Transactions on Computational Logic

Link to article, DOI:
[10.1145/3534927](https://doi.org/10.1145/3534927)

Publication date:
2023

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Matheja, C., Pagel, J., & Zuleger, F. (2023). A Decision Procedure for Guarded Separation Logic: Complete Entailment Checking for Separation Logic with Inductive Definitions. *ACM Transactions on Computational Logic*, 24(1), Article 1. <https://doi.org/10.1145/3534927>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

A Decision Procedure for Guarded Separation Logic

Complete Entailment Checking for Separation Logic with Inductive Definitions

CHRISTOPH MATHEJA, Technical University of Denmark, Denmark, and ETH Zurich, Switzerland

JENS PAGEL, TU Wien, Austria

FLORIAN ZULEGER, TU Wien, Austria

We develop a doubly-exponential decision procedure for the satisfiability problem of *guarded separation logic*—a novel fragment of separation logic featuring user-supplied inductive predicates, Boolean connectives, and separating connectives, including restricted (guarded) versions of negation, magic wand, and septraction. Moreover, we show that dropping the guards for any of the above connectives leads to an undecidable fragment.

We further apply our decision procedure to reason about *entailments* in the popular symbolic heap fragment of separation logic. In particular, we obtain a doubly-exponential decision procedure for entailments between (quantifier-free) symbolic heaps with inductive predicate definitions of bounded treewidth (SL_{btw})—one of the most expressive decidable fragments of separation logic. Together with the recently shown $2EXPTIME$ -hardness for entailments in said fragment, we conclude that the entailment problem for SL_{btw} is $2EXPTIME$ -complete—thereby closing a previously open complexity gap.

CCS Concepts: • **Theory of computation** → **Automated reasoning; Separation logic; Abstraction; Logic and verification; Problems, reductions and completeness.**

Additional Key Words and Phrases: decision procedures, entailment, magic wands, inductive predicates

1 INTRODUCTION

Separation Logic (SL) [Ishtiaq and O’Hearn 2001; Reynolds 2002] is a popular formalism for Hoare-style verification of imperative, heap-manipulating programs. At its core, SL extends first-order logic with two connectives—the *separating conjunction* \star and the *separating implication* \rightarrow (aka magic wand)—for concisely specifying how resources, such as program memory, can be split-up and extended, respectively. Based on these connectives, SL enables *local reasoning*, i.e., sound verification of program parts in isolation, about the resources employed by a program—a key property responsible for SL’s broad adoption in static analysis [Berdine et al. 2007; Calcagno and Distefano 2011; Calcagno et al. 2015, 2011; Gotsman et al. 2007], automated verification [Berdine et al. 2005a, 2011; Chin et al. 2012; Jacobs et al. 2011; Müller et al. 2017; Piskac et al. 2014b; Ta et al. 2018], and interactive theorem proving [Appel 2014; Jung et al. 2018].

Regardless of the flavor of formal reasoning, any *automated* approach based on SL ultimately relies on a solver for discharging either the *satisfiability problem*—does the SL formula ϕ have a model?—or the *entailment problem*—is every model of ϕ also a model of ψ , or, equivalently, is $\phi \wedge \neg\psi$ unsatisfiable? While both problems are undecidable (and equivalent) in general [Calcagno et al. 2001], various decidable SL fragments, in which entailments cannot be reduced to the (un)satisfiability problem because negation is forbidden, have been proposed in the literature, e.g., [Berdine et al. 2004; Cook et al. 2011; Echenim et al. 2020a; Iosif et al. 2013].

Authors’ addresses: Christoph Matheja, Technical University of Denmark, Denmark, and ETH Zurich, Switzerland, cmatheja@inf.ethz.ch; Jens Pagel, pagel@forsyte.at, TU Wien, Vienna, Austria; Florian Zuleger, TU Wien, Vienna, Austria, zuleger@forsyte.at.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, or to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

1529-3785/2022/9-ART \$15.00

<https://doi.org/10.1145/3534927>

$$\begin{aligned} \text{tll}(x_1, x_2, x_3) &\Leftarrow (x_1 \mapsto \langle \text{nil}, \text{nil}, x_3 \rangle) \star (x_1 \approx x_2) \\ \text{tll}(x_1, x_2, x_3) &\Leftarrow \exists \langle l, r, m \rangle . (x_1 \mapsto \langle l, r, \text{nil} \rangle) \\ &\quad \star \text{tll}(l, x_2, m) \\ &\quad \star \text{tll}(r, m, x_3) \end{aligned}$$

(a) An SID specifying trees with linked leaves.

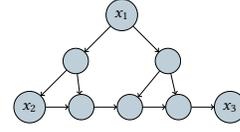
(b) A model of $\text{tll}(x_1, x_2, x_3)$.

Fig. 1. The SID of Iosif et al. [2013] defining trees with linked leaves and an illustration of a model.

In particular, the *symbolic heap* fragment—an idiomatic form of SL formulas with \star but without \rightarrow that is often encountered when manually writing program proofs [Berdine et al. 2005b]—has received a lot of attention. Symbolic heaps appear, for instance, in the automated tools INFER [Calcagno and Distefano 2011], SLEEK [Chin et al. 2012], SONGBIRD [Ta et al. 2016], GRASSHOPPER [Piskac et al. 2014b], VERIFAST [Jacobs et al. 2011], SLS [Ta et al. 2018], and SPEN [Enea et al. 2017]. To support complex data structure specifications, symbolic heaps are often enriched with *systems of inductive predicate definitions* (SIDs). For example, Fig. 1 depicts an SID specifying trees with linked leaves as well as an illustration of a model (*nil*-pointers have been omitted for readability).

The precise form of permitted SIDs has a significant impact on the decidability and complexity of reasoning about symbolic heaps: Brotherston et al. [2014] showed that *satisfiability* is EXPTIME-complete for symbolic heaps over arbitrary SIDs, whereas the *entailment problem* is undecidable in general (cf. [Antonopoulos et al. 2014; Iosif et al. 2014]). To deal with entailments, tools rely on specialized methods for fixed predicates [Berdine et al. 2004; Cook et al. 2011; Piskac et al. 2013, 2014a], decision procedures for restricted classes of SIDs [Iosif et al. 2013, 2014], or incomplete approaches, e.g., fold/unfold reasoning [Chin et al. 2012] or cyclic proofs [Brotherston et al. 2011].

Among the largest decidable classes of symbolic heaps with user-supplied SIDs is the fragment of *symbolic heaps with bounded treewidth* (SL_{btw}) developed by Iosif et al. [2013], which supports rich data structure definitions, such as the one in Fig. 1. Further examples include doubly-linked lists and binary trees with parent pointers. Decidability is achieved by imposing three syntactic conditions on SIDs, which allow reducing the entailment problem for SL_{btw} to the (decidable) satisfiability problem for monadic second-order logic (MSO) over graphs of bounded treewidth (cf. [Courcelle and Engelfriet 2012]). This reduction yields an elementary decision procedure (by analyzing the resulting quantifier depth, it is in 4EXPTIME). However, it is infeasible in practice. Furthermore, there is a “complexity gap” between the above decision procedure and a recent result proving that the entailment problem for SL_{btw} is at least 2EXPTIME-hard [Echenim et al. 2020b].

The goal of this article is twofold: First, we look beyond symbolic heaps and study *guarded separation logic* (GSL)—a novel SL fragment featuring both standard Boolean and separating connectives (including restricted forms of negation and magic wand) as well as SIDs supported by SL_{btw} . In particular, we develop a doubly-exponential decision procedure for the *satisfiability problem* of GSL. Second, we show that the *entailment problem* for SL_{btw} can be reduced to the satisfiability problem for GSL because an SL_{btw} formula ϕ entails an SL_{btw} formula ψ iff the GSL formula $\phi \wedge \neg\psi$ is unsatisfiable. Consequently, we close the aforementioned complexity gap and conclude that the entailment problem for SL_{btw} is 2EXPTIME-complete.

Guarded separation logic. Inspired by work on first-order logic with *guarded negation*, we propose the fragment GSL of *guarded separation logic*. GSL supports negation \neg , magic wand \star and septraction \oplus [Brochenin et al. 2012], but requires each of these connectives to appear in conjunction with another GSL formula ϕ , i.e., $\phi \wedge \neg\psi$, $\phi \wedge (\psi \star \vartheta)$, or $\phi \wedge (\psi \oplus \vartheta)$, acting as its *guard*; hence, the name. By construction, a guard is never equivalent to true and thus cannot be dropped.

While we consider the satisfiability problem of quantifier-free GSL formulas, we admit arbitrary inductive predicates as long as they can be defined in SL_{btw} , which supports existential quantifiers. Hence, the formulas

below belong to GSL and are thus covered by our decision procedure.

$$\begin{aligned} \text{tll}(x, y, z) \wedge \neg x \mapsto \langle \text{nil}, \text{nil}, z \rangle & \quad (\text{a tree with linked leaves and at least three nodes}) \\ (x \mapsto \langle y, z \rangle \star \text{tll}(y, \ell, r) \star \text{tll}(z, r, \text{nil})) \wedge \neg \text{tll}(x, \ell, \text{nil}) & \quad (\text{encoding of an entailment in } \mathbf{SL}_{\text{btw}}) \\ (x \mapsto \langle y, z \rangle \star \text{tll}(z, r, \text{nil})) \wedge (\text{tll}(y, \ell, r) \rightarrow \star \text{tll}(x, \ell, \text{nil})) & \quad (\text{the root's left subtree is missing}) \end{aligned}$$

Abstraction-based satisfiability checking. Our decision procedure for GSL satisfiability—and thus also for \mathbf{SL}_{btw} entailments—is based on the *compositional computation* of an *abstraction* of program states, i.e., the universe of potential models, that *refines* the satisfaction relation \models of GSL.¹ That is, we will develop an abstract domain \mathbb{A} and an abstraction function $\text{abst} : \mathbf{States} \rightarrow \mathbb{A}$ with the following three key properties:

- (1) *Refinement.* Whenever $\text{abst}(\sigma) = \text{abst}(\sigma')$ holds for two states σ and σ' , then σ and σ' satisfy the same GSL formulas, i.e., the equivalence relation induced by our abstraction function,

$$\sigma \equiv_{\text{abst}} \sigma' \quad \text{iff} \quad \text{abst}(\sigma) = \text{abst}(\sigma'),$$

refines the satisfaction relation \models of GSL.

- (2) *Compositionality.* For each logical connective supported by GSL, the abstraction function abst can be computed compositionally from already known abstractions. For example, for the separating conjunction \star , this means that there exists an effectively computable operation $\bullet : \mathbb{A} \times \mathbb{A} \rightarrow \mathbb{A}$ such that, for all states σ and σ' ,

$$\text{abst}(\sigma \uplus \sigma') = \text{abst}(\sigma) \bullet \text{abst}(\sigma'),$$

where $\sigma \uplus \sigma'$ denotes the “disjoint union” of two states used to assign semantics to the separating conjunction.

- (3) *Finiteness.* The abstract domain \mathbb{A} has only finitely many elements after fixing the number of free variables.

Put together, refinement and compositionality allow lifting the abstraction function over states

$$\text{abst} : \mathbf{States} \rightarrow \mathbb{A}$$

to a function over models of GSL formulas

$$\text{abst}_{\text{GSL}} : \text{GSL} \rightarrow \mathbb{A}, \quad \phi \mapsto \{\text{abst}(\sigma) \mid \sigma \in \mathbf{States}, \sigma \models \phi\},$$

Provided we can compute the abstraction abst_{GSL} of every atomic GSL formula, we can then use abst_{GSL} for satisfiability checking: The GSL formula ϕ is satisfiable iff $\text{abst}_{\text{GSL}}(\phi) \neq \emptyset$. Finiteness then ensures that the set $\text{abst}_{\text{GSL}}(\phi)$ is finite; it can thus be computed and checked for emptiness.

We will provide a more detailed overview in Section 6 of our abstraction once we have precisely defined the semantics of guarded separation logic formulas.

Contributions. The main contributions of this article can be summarized as follows:

- We study the decidability of (quantifier-free) *guarded* separation logic (GSL)—a novel separation logic fragment that goes beyond symbolic heaps with user-defined inductive definitions by featuring restricted (guarded) versions of the magic wand, septraction, and negation.
- We show that omitting the guards for *any* of the three operators — magic wand, septraction, and negation — leads to an undecidable logic. Together with our decidability results, this yields an almost tight decidability delineation between separation logics that admit user-defined inductive predicate definitions.
- We present a decision procedure for the satisfiability problem of GSL based on the compositional computation of finite abstractions, called Φ -types, of potential models.
- We analyze the complexity of the above decision procedure and show that satisfiability of GSL is decidable in 2ExpTime .

¹We will properly formalize all notions mentioned in this section in the remainder of this article.

- We apply our decision procedure for GSL to decide, again in 2ExpTime , the *entailment* problem for (quantifier-free) symbolic heaps with user-defined inductive definitions of bounded-tree; in light of the recently shown 2ExpTime -hardness by Echenim et al. [2020b], we obtain that said entailment problem is 2ExpTime -complete—thereby closing an existing complexity gap.

This article unifies and revises the results of two conference papers [Katelaan et al. 2019; Katelaan and Zuleger 2020]. We note that both papers only sketch the main ideas and most of the proofs were omitted. In this article, we dedicate a whole section to the careful motivation of our abstraction (see Section 6) and present all the proofs (an early version of this article was put on arXiv [Pagel et al. 2020] in order to convince the reviewers of [Katelaan and Zuleger 2020] of the correctness of our results). We remark that we have significantly improved the presentation and reworked all the technical details in comparison to our earlier technical report [Pagel et al. 2020].

Organization of the article. After agreeing on basic notational conventions in Section 2, we briefly recap separation logic in Section 3. In particular, we consider user-defined inductive definitions and the bounded treewidth fragment upon which our own SL fragments are based.

We introduce the novel fragment of guarded separation logic in Section 4. Section 5 shows that even small extensions of guarded separation logic lead to an unsatisfiable satisfiability problem. The remainder of this article is concerned with developing a decision procedure for guarded separation logic and, by extension, the entailment problem for symbolic heaps with inductive definitions of bounded treewidth. Section 6 informally discusses the main ideas underlying our decision procedure. The formal details are worked out in Sections 7 to 9. In particular, in Section 9, we present the decision procedure itself and analyze its complexity. Finally, we conclude in Section 10. To improve readability, some technical proofs have been moved to the Appendix at the end of this article.

Acknowledgments. We thank Mnacho Echenim, Radu Iosif, and Nicolas Peltier for their outstandingly thorough study of [Katelaan et al. 2019], which presented the originally proposed abstraction-based decision procedure, and their help in discovering an incompleteness issue, which we were able to fix in our follow-up work [Katelaan and Zuleger 2020; Pagel et al. 2020].

2 NOTATION

Throughout this article, we adhere to the following basic notational conventions.

Sequences. Finite sequences are denoted either in boldface, e.g., \mathbf{x} , or by explicitly listing their elements, e.g., $\langle x_1, \dots, x_k \rangle$; the *empty sequence* is $\langle \rangle$. The *length* of the sequence \mathbf{x} is $|\mathbf{x}|$. We call \mathbf{x} *repetition-free* if its elements are pairwise different. To reduce notational clutter, we often omit the brackets around sequences of length one, i.e., we write x instead of $\langle x \rangle$. The sequence $\mathbf{x} \cdot \mathbf{y}$ is obtained from *concatenating* the sequences \mathbf{x} and \mathbf{y} . A^* is the *set of all finite sequences* over some set A ; A^+ is the set of all *non-empty* finite sequences over A .

Sets from sequences. We frequently treat sequences as sets if the ordering of elements is irrelevant. For example, $x \in \mathbf{x}$ states that the sequence \mathbf{x} contains the element x , $\mathbf{x} \cup \mathbf{y}$ is the set consisting of all elements of the sequences \mathbf{x} and \mathbf{y} , etc.

Partial functions. We denote by $f: A \rightarrow B$ a (*partial*) *function* with *domain* $\text{dom}(f) \triangleq A$ and *image* $\text{img}(f) \triangleq B$. If f is undefined on x , i.e., $x \notin \text{dom}(f)$, we write $f(x) = \perp$. Moreover, $f \circ g$ is the *composition* of the functions f and g mapping every x to $f(g(x))$. We interpret the *size* $|f|$ of a partial function f as the cardinality of its domain, i.e., $|f| \triangleq |\text{dom}(f)|$; f is *finite* if $|f|$ is finite.

We often describe finite partial functions as sets of mappings. The set $\{x_1 \mapsto y_1, \dots, x_k \mapsto y_k\}$, for example, represents the partial function that, for every $i \in [1, k]$, maps x_i to y_i ; it is undefined for all other values. Furthermore, $f \cup g$ denotes the (not necessarily disjoint) union of f and g ; it is defined iff $f(x) = g(x)$ holds for

$$\begin{aligned}
& x \in \mathbf{Var}, u, v \in \mathbf{Var} \cup \mathbf{Loc}, \mathbf{w} \in (\mathbf{Var} \cup \mathbf{Loc})^+, \text{pred} \in \mathbf{Preds} \\
\phi_{\text{atom}} & ::= \mathbf{emp} \mid u \approx v \mid u \not\approx v \mid u \mapsto \mathbf{w} \mid \text{pred}(\mathbf{w}) \\
\phi & ::= \phi_{\text{atom}} \mid \phi \star \phi \mid \phi \neg \star \phi \mid \phi \wedge \phi \mid \phi \vee \phi \mid \neg \phi \mid \exists x. \phi \mid \forall x. \phi
\end{aligned}$$

Fig. 2. The syntax of a first-order separation logic with user-defined predicates (SL)

all $x \in \text{dom}(f) \cap \text{dom}(g)$. Formally:

$$(f \cup g)(x) \triangleq \begin{cases} f(x), & \text{if } x \in \text{dom}(f), \\ g(x), & \text{if } x \in \text{dom}(g) \setminus \text{dom}(f), \\ \perp, & \text{otherwise.} \end{cases}$$

We write $f \uplus g$ instead of $f \cup g$ whenever we additionally require that $\text{dom}(f) \cap \text{dom}(g) = \emptyset$. Furthermore, we denote by $f[x/v]$ the *updated* partial function in which x maps to v , i.e.,

$$f[x/v](y) \triangleq \begin{cases} v & \text{if } y = x, \\ f(y), & \text{otherwise.} \end{cases}$$

In particular, if $f(x)$ is undefined, $f[x/v]$ adds x to the domain of the resulting function. By slight abuse of notation, we write $f[x/\perp]$ to denote the function in which x is removed from the domain of f . To compare partial functions, we say that function g *subsumes* function f , written $f \subseteq g$, if (1) g is at least as defined as g and (2) g agrees with f on their common domain. Formally:

$$f \subseteq g \quad \text{iff} \quad \text{dom}(f) \subseteq \text{dom}(g) \quad \text{and} \quad \forall x \in \text{dom}(f): f(x) = g(x).$$

Whenever f and g map to sets or functions, we also use a weaker ordering. The relation $f \sqsubseteq g$ is defined as $f \subseteq g$ but only requires that $g(x)$ subsumes $f(x)$ if both are defined on x , i.e.,

$$f \sqsubseteq g \quad \text{iff} \quad \text{dom}(f) \subseteq \text{dom}(g) \quad \text{and} \quad \forall x \in \text{dom}(f): f(x) \subseteq g(x).$$

Functions over sequences. We implicitly lift partial functions $f: A \rightarrow B$ to functions $f: A^* \rightarrow B^*$ over sequences by pointwise application. That is, for a sequence $\langle a_1, \dots, a_k \rangle \in A^*$, we define

$$f(\langle a_1, \dots, a_k \rangle) \triangleq f(a_1, \dots, a_k) \triangleq \langle f(a_1), \dots, f(a_k) \rangle,$$

where, as indicated above, we omit the brackets indicating sequences to improve readability.

Finally, we lift the update $f[x/v]$ of a single value to sequences of values $\mathbf{x} = \langle x_1, \dots, x_k \rangle$ and $\mathbf{v} = \langle v_1, \dots, v_k \rangle$ by setting $f[\mathbf{x}/\mathbf{v}] \triangleq f[x_1/v_1][x_2/v_2] \dots [x_k/v_k]$.

3 SEPARATION LOGIC WITH INDUCTIVE DEFINITIONS

We briefly recapitulate the basics of first-order separation logic with user-defined predicates. That is, we introduce the syntax and semantics of both separation logic and systems of inductive definitions, the symbolic heap fragment, and the bounded treewidth fragment originally studied by Iosif et al. [2013]. Most of the presented material is fairly standard (cf., among others, [Brotherston et al. 2014; Iosif et al. 2014; Ishtiaq and O’Hearn 2001; Reynolds 2002]) with the notable exception that *our semantics of pure formulas enforces the heap to be empty*. A reader familiar with separation logic may skim over this section to familiarize herself with our notation.

3.1 The Syntax of Separation Logic

Figure 2 defines the syntax of first-order separation logic with user-defined predicates (SL for short), where x is drawn from a countably infinite set \mathbf{Var} of *variables*, u and v are either variables or *locations* drawn from the countably infinite set \mathbf{Loc} , and \mathbf{w} is a finite sequence whose elements can be both variables and locations.

In particular, notice that any location in \mathbf{Loc} may appear as a constant in formulas. Moreover, pred is taken from a finite set \mathbf{Preds} of *predicate identifiers*; each predicate pred is equipped with an *arity* $\text{ar}(\text{pred}) \in \mathbb{N}$ that determines its number of parameters.

Informally, the meaning of the atomic formulas is as follows:

- The *empty-heap predicate* \mathbf{emp} denotes the empty heap.
- The *equality* $u \approx v$ and the *disequality* $u \neq v$ express that u and v alias and that they do not alias in the current program state (whose heap needs to be empty), respectively.
- The *points-to assertion* $u \mapsto \mathbf{w}$ states that the address u points to a heap-allocated object consisting of $|\mathbf{w}| > 0$ fields, where the i -th field stores the i -th location of the sequence \mathbf{w} .
- The *predicate call* $\text{pred}(\mathbf{w})$ allows to refer to user-defined data structures, e.g., lists and trees.

The SL formulas \mathbf{emp} and $u \mapsto v$ are called *spatial atoms* because they describe the spatial layout of the heap, whereas (dis-)equalities are called *pure atoms* [Ishtiaq and O’Hearn 2001] because they do not depend on the heap. Apart from atoms, SL formulas are built from

- classical propositional connectives, i.e., conjunction (\wedge), disjunction (\vee), and negation (\neg),
- existential (\exists) and universal (\forall) quantifiers, and
- *separating connectives*, i.e., *separating conjunction* \star and *implication* (or magic wand) \multimap .

As usual, one can derive additional operators such as standard implication $\phi \Rightarrow \psi \triangleq \neg\phi \vee \psi$ and *septraction* $\phi \oplus \psi \triangleq \neg(\phi \star \neg\psi)$ (cf. [Brochenin et al. 2012; Thakur et al. 2014]).

The semantics of the classical connectives is standard. Let us briefly compare their meaning with the intuition underlying the separating connectives. While $\phi \wedge \psi$ means that the program state satisfies both ϕ and ψ simultaneously, $\phi \star \psi$ denotes that (the heap component of) the program state can be split into two disjoint parts which separately satisfy ϕ and ψ . Similarly, while $\phi \Rightarrow \psi$ means that every program state satisfying ϕ also satisfies ψ , $\phi \multimap \psi$ means that the *extension* of the program state with any program state that satisfies ϕ yields a program state that satisfies ψ .

The magic wand is useful for weakest-precondition reasoning, e.g., to express memory allocation [Batz et al. 2019; Ishtiaq and O’Hearn 2001; Reynolds 2002]. However, automated verification tools often do not or only partially support the magic wand, because its inclusion quickly leads to undecidability [Appel 2014; Blom and Huisman 2015; Schwerhoff and Summers 2015].

3.1.1 Substitution. Various constructions in this article involve syntactically replacing variables and locations—we thus give a generic definition that allows performing multiple substitutions at once. Let $\mathbf{y}, \mathbf{z} \in (\mathbf{Var} \cup \mathbf{Loc})^*$ be sequences of the same length, where \mathbf{y} is repetition-free. We denote by $\phi[\mathbf{y}/\mathbf{z}]$ the formula obtained from ϕ through (simultaneous) *substitution* of each element in \mathbf{y} by the element in \mathbf{z} at the same position; Appendix A.1 provides a formal definition. E.g.,

$$(\exists x. x \mapsto \langle y, 7 \rangle \star \text{ls}(y, x))[\langle y, 7 \rangle / \langle 3, x \rangle] = \exists x. (x \mapsto \langle 3, x \rangle) \star \text{ls}(3, x).$$

Moreover, we write $\phi(\mathbf{z})$ as a shortcut for the substitution $\phi[\text{fvars}(\phi)/\mathbf{z}]$.

3.2 The Stack-Heap Model

We interpret SL in terms of the widely-used *stack-heap model*, which already appears in the seminal papers of Ishtiaq and O’Hearn [2001] and Reynolds [2002]. A stack-heap pair $\langle \mathfrak{s}, \mathfrak{h} \rangle$ consists of a stack \mathfrak{s} assigning *values* in \mathbf{Val} to variables and a heap \mathfrak{h} assigning values in \mathbf{Val} to allocated memory *locations* taken from $\mathbf{Loc} \subseteq \mathbf{Val}$.

We fix the set of locations $\mathbf{Loc} \triangleq \mathbb{N}_{>0}$. To simplify the technical development, we will work with the sets of values $\mathbf{Val} = \mathbf{Loc}$ and $\mathbf{Val} = \mathbf{Loc} \cup \{\text{nil}\}$, where the latter extends the former with the null pointer *nil*.

We will mostly work with $\mathbf{Val} = \mathbf{Loc}$ (and generally prefer the term locations over values), which simplifies the presentation of our decision procedure. However, we will use $\mathbf{Val} \triangleq \mathbf{Loc} \cup \{\text{nil}\}$ in examples and in the

undecidability proofs (in these examples and proofs one could also replace *nil* by a fresh variable that refers to a location which is never allocated; hence, the use of *nil* in these places has been a decision of taste and presentational convenience). We later present a reduction that extends our decision procedure to the extended set of values that includes the null pointer (see Cor. 9.4). We note, however, that it would not be difficult to extend our decision procedure to directly deal with $\mathbf{Val} \triangleq \mathbf{Loc} \cup \{\mathit{nil}\}$. We summarize that everywhere, where *nil* is not explicitly mentioned, the reader should identify values with locations, that is, assume $\mathbf{Val} = \mathbf{Loc} \triangleq \mathbb{N}_{>0}$; in such cases, we will prefer the term locations over values. The set **Stacks** of *stacks* then consists of all finite partial functions mapping variables to locations, i.e.,

$$\mathbf{Stacks} \triangleq \{\mathfrak{s} \mid \mathfrak{s}: V \rightarrow \mathbf{Val}, V \subseteq \mathbf{Var}, |V| < \infty, \},$$

In order to treat both evaluations of variables and constant locations uniformly, we slightly abuse notation and set $\mathfrak{s}(v) \triangleq v$ for all values $v \in \mathbf{Val}$.² The set **Heaps** of *heaps* consists of all finite partial functions mapping allocated memory locations to sequences of locations, i.e.,

$$\mathbf{Heaps} \triangleq \{\mathfrak{h} \mid \mathfrak{h}: L \rightarrow \mathbf{Val}^+, L \subseteq \mathbf{Loc}, |L| < \infty\}.$$

By mapping locations to sequences rather than single locations, the heap assigns every allocated memory location to the entire structure allocated at this location. This is a fairly standard—but far from ubiquitous [Calcagno et al. 2006; Reynolds 2002]—abstraction of the actual memory layout; it simplifies the memory model without losing precision as long as we do not use pointer arithmetic.

We frequently refer to stack-heap pairs $\langle \mathfrak{s}, \mathfrak{h} \rangle$ as (program) *states*.

3.2.1 Location terminology. We denote by $\mathit{locs}(\phi)$ the set of all locations in \mathbf{Loc} that explicitly appear as constants symbols in SL formula ϕ . Similarly, the set of all locations appearing in heap \mathfrak{h} is $\mathit{locs}(\mathfrak{h}) \triangleq \mathit{dom}(\mathfrak{h}) \cup \bigcup_{v \in \mathit{img}(\mathfrak{h})} v$; we lift this set to states $\langle \mathfrak{s}, \mathfrak{h} \rangle$ by setting $\mathit{locs}(\langle \mathfrak{s}, \mathfrak{h} \rangle) \triangleq \mathit{img}(\mathfrak{s}) \cup \mathit{locs}(\mathfrak{h})$.

We often distinguish between allocated, referenced, and dangling locations $\ell \in \mathbf{Loc}$: ℓ is *allocated* in heap \mathfrak{h} if $\ell \in \mathit{dom}(\mathfrak{h})$; it is *referenced* if $\ell \in \mathit{img}(\mathfrak{h})$. Finally, ℓ is *dangling* if it appears in \mathfrak{h} but is not allocated.

We call a variable $x \in \mathbf{Var}$ allocated, referenced, or dangling if the location $\mathfrak{s}(x)$ is allocated, referenced, or dangling, respectively. We collect all allocated variables and all referenced variables in state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ in the sets $\mathit{all}(\langle \mathfrak{s}, \mathfrak{h} \rangle) \triangleq \{x \mid \mathfrak{s}(x) \in \mathit{dom}(\mathfrak{h})\}$ and $\mathit{ref}(\langle \mathfrak{s}, \mathfrak{h} \rangle) \triangleq \{x \mid \mathfrak{s}(x) \in \mathit{img}(\mathfrak{h})\}$.

3.3 The Semantics of Separation Logic

Figure 3 defines the semantics of SL in terms of a satisfaction relation \models_{Φ} , where the sole purpose of Φ —explained in Section 3.3.1 below—is to assign semantics to user-defined predicate calls. A state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ that satisfies an SL formula ϕ , i.e., $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$, is called a *model* of ϕ .

The empty-heap predicate **emp** holds iff the heap is empty; equalities and disequalities between variables hold iff the stack maps the variables to identical and different locations, respectively. For (dis-)equalities, we additionally require that the heap is empty. This is non-standard, but not unprecedented [Piskac et al. 2013], and will simplify the technical development.

A points-to assertion $x \mapsto \langle y_1, \dots, y_k \rangle$ holds in the singleton heap that allocates exactly the location $\mathfrak{s}(x)$ and stores the locations $\mathfrak{s}(y_1), \dots, \mathfrak{s}(y_k)$ at this location. This interpretation of points-to assertions is often called a *precise* [Calcagno et al. 2007; Yang 2001] semantics, because the heap contains precisely the object described by the points-to assertion, and nothing else.

For the separating conjunction, $\langle \mathfrak{s}, \mathfrak{h} \rangle \models \psi \star \theta$ holds if and only if there exist domain-disjoint heaps $\mathfrak{h}_1, \mathfrak{h}_2$ such that their union (\uplus , see Section 2) is \mathfrak{h} and both $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle \models_{\Phi} \psi$ and $\langle \mathfrak{s}, \mathfrak{h}_2 \rangle \models_{\Phi} \theta$ hold. While the separating

²This convention does not affect the formal definition of stacks; in particular, their domain and image remains unchanged.

ϕ	$\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$ iff
emp	$\mathfrak{h} = \emptyset$
$u \approx v$	$\mathfrak{s}(u) = \mathfrak{s}(v)$ and $\mathfrak{h} = \emptyset$
$u \not\approx v$	$\mathfrak{s}(u) \neq \mathfrak{s}(v)$ and $\mathfrak{h} = \emptyset$
$u \mapsto y$	$\mathfrak{h} = \{\mathfrak{s}(u) \mapsto \mathfrak{s}(y)\}$
pred (y)	$\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \psi(y)$ for some $(\text{pred}(x) \leftarrow \psi) \in \Phi$
$\psi \star \theta$	exists $\mathfrak{h}_1, \mathfrak{h}_2$ such that $\mathfrak{h} = \mathfrak{h}_1 \uplus \mathfrak{h}_2$, $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle \models_{\Phi} \psi$, and $\langle \mathfrak{s}, \mathfrak{h}_2 \rangle \models_{\Phi} \theta$
$\psi \star \theta$	for all \mathfrak{h}_0 , if $\text{dom}(\mathfrak{h}_0) \cap \text{dom}(\mathfrak{h}) = \emptyset$ and $\langle \mathfrak{s}, \mathfrak{h}_0 \rangle \models_{\Phi} \psi$ then $\langle \mathfrak{s}, \mathfrak{h}_0 \uplus \mathfrak{h} \rangle \models_{\Phi} \theta$
$\psi \wedge \theta$	$\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \psi$ and $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \theta$
$\psi \vee \theta$	$\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \psi$ or $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \theta$
$\neg \psi$	not $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \psi$
$\exists x. \psi$	exists $v \in \mathbf{Loc}$ such that $\langle \mathfrak{s}[x/v], \mathfrak{h} \rangle \models_{\Phi} \psi$
$\forall x. \psi$	for all $v \in \mathbf{Loc}$, $\langle \mathfrak{s}[x/v], \mathfrak{h} \rangle \models_{\Phi} \psi$

Fig. 3. Semantics of SL

conjunction is about splitting the heap, the magic wand is about extending it: $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi \star \psi$ holds iff all ways to extend \mathfrak{h} with a disjoint model of ϕ yields a model of ψ .

The semantics of the Boolean connectives and the quantifiers is standard. In particular, as justified by the lemma below, the semantics of quantifiers can also be interpreted in terms of syntactic substitution rather than updating the stack (which is formally defined in Section 2).

LEMMA 3.1 (SUBSTITUTION LEMMA). *For all SL formulas ϕ , states $\langle \mathfrak{s}, \mathfrak{h} \rangle$, variables $x \in \text{fvvars}(\phi)$, and locations $\ell \in \mathbf{Loc}$, we have $\langle \mathfrak{s}[x/v], \mathfrak{h} \rangle \models_{\Phi} \phi$ iff $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi[x/v]$.*

PROOF. By induction on the structure of SL formulas. □

EXAMPLE 3.2. (1) $(x \mapsto y) \star (y \mapsto \text{nil})$ states that the heap consists of exactly two objects, one pointed to by x , the other pointed to by y ; that the object pointed to by x contains a pointer to the object pointed to by y ; and that the object pointed to by y contains a null pointer. Put less precisely but more concisely, x points to y , y points to nil , and x and y are separate objects on the heap. The precise semantics of assertions guarantees that there are no other objects in the heap.

(2) $(x \mapsto y) \wedge (z \mapsto y)$ states that (a) the heap consists of a single object x that points to y and that simultaneously (b) the heap consists of a single object z that points to y . This formula is only satisfiable for stacks \mathfrak{s} with $\mathfrak{s}(x) = \mathfrak{s}(z)$.

(3) $(x \mapsto y) \star (z \mapsto y)$ states that after adding a pointer from x to y to the heap, we obtain a heap that contains a single pointer from z to y . This formula is only satisfiable for the empty heap and for stacks \mathfrak{s} with $\mathfrak{s}(x) = \mathfrak{s}(z)$.

(4) $\forall x. (x \mapsto \text{nil}) \star ((\neg \mathbf{emp}) \star (\neg \mathbf{emp}))$ states that the heap contains at least one pointer: no matter which variable we additionally allocate, the resulting heap can be split into two nonempty parts, so the original heap must itself have been nonempty—the formula is equivalent to $\neg \mathbf{emp}$.

3.3.1 Systems of inductive definitions. Predicates are interpreted in terms of a user-supplied system of inductive definitions (SID). An SID is a finite set Φ of rules of the form $\text{pred}(\mathbf{x}) \leftarrow \phi(\mathbf{x})$, where $\text{pred} \in \mathbf{Preds}$ is a predicate symbol, $\text{fvvars}(\text{pred}) = \mathbf{x} \in \mathbf{Var}^*$ are the formal parameters of pred with $|\mathbf{x}| = \text{ar}(\text{pred})$, and ϕ is an SL formula with free variables \mathbf{x} ; the size $|\Phi|$ of Φ is the sum of the sizes of the formulas in its rules. We collect all predicates that occur in Φ in the set $\mathbf{Preds}(\Phi)$. Moreover, we assume that all rules with the same predicate pred on the left-hand side have the same (repetition-free) sequence of parameters \mathbf{x} .

A stack-heap pair $\langle s, h \rangle$ satisfies the predicate call $\text{pred}(y)$ with respect to SID Φ iff Φ contains a rule $\text{pred}(x) \Leftarrow \phi$ such that $\langle s, h \rangle$ satisfies the rule's right-hand side once we instantiate its formal parameters with the arguments passed to the predicate call, i.e., $\langle s, h \rangle \models_{\Phi} \phi(y)$.

Notice that rules involving arbitrary SL formulas—e.g., $\text{pred}(x) \Leftarrow \neg \text{pred}(x)$ —do *not* necessarily lead to a well-defined semantics of predicate calls. We will restrict the formulas allowed to appear in SIDs in Section 3.4.2 to ensure that our semantics is always well-defined.

EXAMPLE 3.3 (INDUCTIVE DEFINITIONS). (1) Let Φ_{ls} be the SID given by the following rules:

$$\begin{array}{ll} \text{lseg}(x_1, x_2) \Leftarrow x_1 \mapsto x_2 & \text{ls}(x_1) \Leftarrow x_1 \mapsto \text{nil} \\ \text{lseg}(x_1, x_2) \Leftarrow \exists y. x_1 \mapsto y \star \text{lseg}(y, x_2) & \text{ls}(x_1) \Leftarrow \exists y. (x_1 \mapsto y) \star \text{ls}(y) \end{array}$$

The predicate $\text{lseg}(x_1, x_2)$ describes non-empty singly-linked list segments with head x_1 and tail x_2 ; the predicate $\text{ls}(x_1)$ describes those list segments that are terminated by a null pointer. Hence, the formulas $\text{lseg}(x_1, \text{nil})$ and $\text{ls}(x_1)$ are equivalent with respect to the SID Φ_{ls} .

(2) The SID $\Phi_{\text{odd/even}}$ below defines all non-empty list segments of odd and even length, respectively.

$$\begin{array}{ll} \text{odd}(x_1, x_2) \Leftarrow x_1 \mapsto x_2 & \text{even}(x_1, x_2) \Leftarrow \exists y. (x_1 \mapsto y) \star \text{odd}(y, x_2) \\ \text{odd}(x_1, x_2) \Leftarrow \exists y. (x_1 \mapsto y) \star \text{even}(y, x_2) & \end{array}$$

(3) The SID Φ_{tree} below defines null-terminated binary trees with root x_1 .

$$\text{tree}(x_1) \Leftarrow x_1 \mapsto \langle \text{nil}, \text{nil} \rangle \quad \text{tree}(x_1) \Leftarrow \exists \langle l, r \rangle. (x_1 \mapsto \langle l, r \rangle) \star \text{tree}(l) \star \text{tree}(r)$$

3.3.2 *Satisfiability and entailment.* An SL formula ϕ is *satisfiable* with respect to Φ iff there exists a state $\langle s, h \rangle$ such that $\langle s, h \rangle \models_{\Phi} \phi$. Moreover, the SL formula ϕ *entails* the SL formula ψ given SID Φ , written $\phi \models_{\Phi} \psi$, iff for all states $\langle s, h \rangle$, we have $\langle s, h \rangle \models_{\Phi} \phi$ implies $\langle s, h \rangle \models_{\Phi} \psi$.

3.3.3 *Isomorphic states.* Our decision procedure will exploit that SL formulas cannot distinguish between individual locations—as long as formulas do not explicitly use constant locations. More formally, formulas cannot distinguish isomorphic states:

DEFINITION 3.4 (ISOMORPHIC STATES). Two states $\langle s, h \rangle$ and $\langle s', h' \rangle$ are isomorphic, written $\langle s, h \rangle \cong \langle s', h' \rangle$, iff there exists a bijection $\sigma: \text{locs}(\langle s, h \rangle) \rightarrow \text{locs}(\langle s', h' \rangle)$ such that

- (1) for all x , $s'(x) = \sigma(s(x))$, and
- (2) $h' = \{\sigma(l) \mapsto \sigma(h(l)) \mid l \in \text{dom}(h)\}$.

LEMMA 3.5. Let ϕ be an SL formula with $\text{locs}(\phi) = \emptyset$. Then, for all states $\langle s, h \rangle$ and $\langle s', h' \rangle$,

$$\langle s, h \rangle \cong \langle s', h' \rangle \quad \text{implies} \quad \langle s, h \rangle \models_{\Phi} \phi \quad \text{iff} \quad \langle s', h' \rangle \models_{\Phi} \phi.$$

PROOF. By induction on the structure of SL formulas. □

3.4 The Bounded Treewidth Fragment

Our main goal is to develop a decision procedure for entailments in an SL fragment that extends the so-called bounded treewidth fragment (SL_{btw}) of Iosif et al. [2013]. In this section, we briefly recapitulate that fragment as we will rely on similar restrictions for SIDs.

3.4.1 *Symbolic heaps.* Formulas in SL_{btw} are restricted to symbolic heaps with user-supplied predicates—a popular fragment of SL that is both expressive enough for specifying complex heap shapes and suitable to serve

as an abstract domain for program analyses (cf. [Berdine et al. 2007, 2005b; Calcagno et al. 2011]). A *symbolic heap* is a formula of the form

$$\underbrace{\exists x_1, \dots, x_k}_{k \geq 0} \cdot \underbrace{\phi_{\text{atom}} \star \dots \star \phi_{\text{atom}}}_{1 \text{ or more atoms}}$$

Notice that negation, disjunction, universal quantifiers, and magic wands are *not* allowed in symbolic heaps. In particular, this means—since pure formulas are evaluated in the empty heap—that there is *no* symbolic heap that is always satisfied, i.e., equivalent to true.

When working with symbolic heaps, it is convenient to group the atoms into (1) a spatial part collecting all points-to assertions, (2) a part collecting all predicate calls, and (3) a pure part collecting all equalities and disequalities (in that order). Hence, the set SH^\exists of *symbolic heaps* ϕ_{sh} is given by

$$\phi_{\text{sh}} ::= \underbrace{\exists e}_{e \in \text{Var}^*} \cdot \underbrace{(u_1 \mapsto v_1) \star \dots \star (u_k \mapsto v_k)}_{\text{spatial part, emp for } k=0} \star \underbrace{\text{pred}_1(w_1) \star \dots \star \text{pred}_l(w_l)}_{\text{predicate calls, emp for } l=0} \\ \star \underbrace{u_1 \approx v_1 \star \dots \star u_m \approx v_m \star u'_1 \not\approx v'_1 \star \dots \star u'_n \not\approx v'_n}_{\text{pure part, emp for } m=0, n=0 \text{ respectively}}$$

3.4.2 Symbolic heap SIDs. Our semantics of predicate calls (see Figure 3) is well-defined as long as all formulas appearing in the underlying SID are symbolic heaps—a requirement that we impose throughout the remainder of this article. For instance, all SIDs in Example 3.3 only use symbolic heaps in their rules. The restriction of SID rules to symbolic heaps is standard. In fact, our semantics coincides with other semantics from the separation logic literature that—instead of replacing predicates by rules step by step—are based on least fixed points [Brotherston 2007; Brotherston et al. 2014] or derivation trees [Iosif et al. 2013, 2014; Jansen et al. 2017; Matheja 2020].

3.4.3 The bounded treewidth fragment. Since negation is not available, the entailment problem for symbolic heaps is genuinely different from the satisfiability problem: it is impossible to solve an entailment $\phi \models_{\Phi} \psi$ by checking the unsatisfiability of $\phi \wedge \neg\psi$, because the latter formula is not a symbolic heap. In fact, the satisfiability problem for symbolic heaps is decidable in general [Brotherston et al. 2014], whereas the entailment problem is not [Antonopoulos et al. 2014; Iosif et al. 2014]. However, various subclasses of symbolic heaps with a decidable—and even tractable [Cook et al. 2011]—entailment problem have been studied in the literature (e.g., [Berdine et al. 2004; Iosif et al. 2013, 2014; Le et al. 2017]); as such, the symbolic heap fragment has been the main focus of a recent competition of entailment solvers (SL-COMP) [et al. 2019]. The largest of these fragments has been developed by Iosif et al. [2013]; it achieves decidability by imposing three restrictions—progress, connectivity, and establishment—on SIDs to ensure that all models of predicates are of *bounded treewidth*.³

Local allocation and references. To formalize the above three assumptions for SIDs, we need two auxiliary definitions: we collect all variables and locations that appear on the left-hand side of points-to assertions in formula ϕ in the *local allocation* set $\text{lalloc}(\phi)$; analogously, the *local references* set $\text{lref}(\phi)$ collects all variables and locations appearing on the right-hand side of points-to assertions.

We now present the three aforementioned conditions imposed on SIDs Φ to ensure decidability of the entailment problem for symbolic heaps.

Progress. A predicate pred satisfies *progress* iff there exists a free variable $x \in \text{fv}(\text{pred})$ such that, for all rules $(\text{pred}(x) \leftarrow \phi) \in \Phi$, (1) ϕ contains exactly one point-to assertion, and (2) x is allocated in ϕ , i.e., $\text{lalloc}(\phi) = \{x\}$.

³More precisely, when viewed as graphs, all models of the SIDs satisfying the three aforementioned restrictions have bounded treewidth. For a formal definition of treewidth, we refer to [Diestel 2016].

In this case, we call x the *root* of pred . Moreover, if the i -th parameter of pred , say x_i , is its root, then we set $\text{predroot}(\text{pred}(x)) \triangleq x_i$.

Connectivity. A predicate pred satisfies *connectivity* iff for all rules of pred , all variables that are allocated in the recursive calls of the rule are also referenced in the rule. Formally, for all rules $(\text{pred} \Leftarrow \phi) \in \Phi$ and for all calls $\text{pred}'(y)$ appearing in ϕ , we have $\text{predroot}(\text{pred}'(y)) \subseteq \text{lref}(\phi)$.

Establishment. A predicate pred is *established* iff all existentially quantified variables across all rules of pred are eventually allocated, or equal to a parameter. Formally, for all rules $(\text{pred}(x) \Leftarrow \exists y. \phi) \in \Phi$ and for all states $\langle s, \mathfrak{h} \rangle$, if $\langle s, \mathfrak{h} \rangle \models_{\Phi} \phi$ then $s(y) \subseteq \text{dom}(\mathfrak{h}) \cup s(x)$.

3.4.4 SIDs of bounded treewidth. We denote by ID_{btw} the set of all SIDs in which all predicates satisfy *progress*, *connectivity*, and *establishment*. E.g., the SIDs given in Example 3.3 belong to ID_{btw} .

THEOREM 3.6 ([ECHENIM ET AL. 2020B; IOSIF ET AL. 2013]). *The entailment problem for symbolic heaps over SIDs in ID_{btw} is decidable, of elementary complexity, and 2-EXPTIME hard.*

In the remainder of this article, we strengthen the above theorem in two ways: First, we give a larger decidable SL fragment, and, second, we develop a 2-EXPTIME decision procedure which, by the above lower bounds, is of optimal asymptotic complexity. Moreover, we show that even small extensions of our fragments lead to an undecidable entailment problem.

3.4.5 Global Assumptions about SIDs. Unless stated otherwise, we assume that all SIDs considered in this article belong to ID_{btw} . Moreover, to avoid notational clutter, Φ always refers to an arbitrary, but fixed SID in ID_{btw} unless it is explicitly given. Without loss of generality, we make two further assumptions about the rules in SIDs to simplify the technical development.

First, we assume that non-recursive rules do *not* contain existential quantifiers because they can always be eliminated: due to progress and establishment, all existentially-quantified variables in a non-recursive rule must be provably equal to either a constant or a parameter of the predicate.

Second, to avoid dedicated reasoning about points-to assertions, we may add dedicated predicates simulating points-to assertions to every SID; we call the resulting SIDs *pointer-closed*:

DEFINITION 3.7 (POINTER-CLOSED SID). *An SID Φ is pointer-closed w.r.t. ϕ iff it contains a predicate ptr_k and a single rule $\text{ptr}_k(\langle x_1, \dots, x_{k+1} \rangle) \Leftarrow x_1 \mapsto \langle x_2, \dots, x_{k+1} \rangle$ for all points-to assertions mapping to structures of length k in ϕ .*

Since all predicates introduced by transforming an SID into a pointer-closed one satisfy progress, connectivity, and establishment, we can safely assume that SIDs in ID_{btw} are pointer-closed. As a consequence of this assumption, we consider the number of formal parameters of predicates to be at least as large as the number of fields of points-to assertions whenever we analyze complexities.

4 THE GUARDED FRAGMENT OF SEPARATION LOGIC

To obtain fragments of SL with both support for complex data structure predicates and a decidable entailment problem, we rely on the same restrictions on user-defined predicates as Iosif et al. [2013]: the semantics of predicate calls needs to be determined by SIDs taken from the bounded treewidth fragment ID_{btw} . In contrast to Iosif et al. [2013], our work is, however, not limited to entailments between symbolic heaps over the predicates at hand. Rather, we additionally consider reasoning about a novel (quantifier-free) *guarded* fragment of separation logic (GSL) featuring restricted variants of negation \neg , magic wand \star , and septraction \otimes .

Intuitively, the guarded fragment enforces that the aforementioned connectives \neg , \star , and \otimes only appear in conjunction with another formula restricting the possible shapes of the heap. We will show in Section 5 that this

restriction is crucial: lifting it for any of these connectives yields an undecidable entailment problem—even if the remaining connectives are removed.

4.1 Guarded Formulas

The set **GSL** of formulas in (quantifier-free) *guarded separation logic* is given by the grammar

$$\begin{aligned} \phi ::= \phi_{\text{atom}} \quad (::= \mathbf{emp} \mid u \approx v \mid u \not\approx v \mid u \mapsto \mathbf{w} \mid \text{pred}(\mathbf{w})) & \quad (\text{same atoms as SL}) \\ & \mid \phi \star \phi \mid \phi \wedge \phi \mid \phi \vee \phi & \quad (\text{standard connectives}) \\ & \mid \phi \wedge \neg\phi \mid \phi \wedge (\phi \star \phi) \mid \phi \wedge (\phi \oplus \phi). & \quad (\text{guarded connectives}) \end{aligned}$$

The atoms as well as the connectives \star , \wedge , and \vee are the same as for the full logic **SL** introduced in Section 3.1. Moreover, negation \neg , magic wand \star , and septraction \oplus may only appear in *guarded* form, i.e., in conjunction with another guarded formula ϕ . Since **GSL** is a syntactic fragment of **SL**, the semantics of **GSL** is given by the semantics of **SL** presented in Section 3.3.

EXAMPLE 4.1. Assume the predicate $\text{lseg}(x_1, x_2)$ represents all non-empty list segments from x_1 to x_2 ; a formal definition is found in Example 3.3. Moreover, consider the following guarded formulas:

- (1) $\text{lseg}(x, y) \wedge \neg x \mapsto y$ states that the heap consists of a list of length at least two.
- (2) $\text{lseg}(x, y) \wedge (\text{lseg}(y, z) \oplus \text{lseg}(x, x))$ states that the heap consists of a list segment from x to y that can be extended to a cyclic list by adding a list from y to z ; it entails that x and z are aliases.

In contrast to variants of separation logic in the literature (cf. [Calcagno et al. 2011; Reynolds 2002]), our separation logic **SL** does not contain an atom true , which is always satisfied. While true is, of course, definable in **SL**, e.g., $\mathbf{emp} \vee \neg\mathbf{emp}$, it is *not* definable in **GSL**. In particular, $x \approx x$ is not equivalent to true as our semantics of equalities and disequalities requires the heap to be empty. This is crucial: If true were definable, **GSL** would coincide with the set of all quantifier-free **SL** formulas, because we could choose true for all guards.

4.2 Guarded States and Dangling Pointers

The decision procedure developed in this article exploits that all models of guarded formulas are themselves guarded in the sense that they have only a limited amount of dangling pointers. We recall from Section 3.2.1 that a dangling pointer is a location that is not allocated; the set of all dangling locations in heap \mathfrak{h} is thus given by $\text{dangling}(\mathfrak{h}) \triangleq \text{locs}(\mathfrak{h}) \setminus \text{dom}(\mathfrak{h})$.

In the following, we first define guarded states, then show that establishment implies a models of atomic predicates to be guarded, and finally lift this result to arbitrary guarded formulas.

DEFINITION 4.2 (GUARDED STATE). The set **GStates** of guarded states is given by

$$\mathbf{GStates} \triangleq \{ \langle \mathfrak{s}, \mathfrak{h} \rangle \mid \mathfrak{s} \in \mathbf{Stacks}, \mathfrak{h} \in \mathbf{Heaps}, \text{dangling}(\mathfrak{h}) \subseteq \text{img}(\mathfrak{s}) \}.$$

Guarded states are well-behaved with regard to taking the union of heaps:

LEMMA 4.3. Let $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle, \langle \mathfrak{s}, \mathfrak{h}_2 \rangle \in \mathbf{GStates}$ with $\mathfrak{h}_1 \uplus \mathfrak{h}_2 \neq \perp$. Then, $\langle \mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2 \rangle \in \mathbf{GStates}$.

PROOF. We observe that $\text{dangling}(\mathfrak{h}_1 \uplus \mathfrak{h}_2) \subseteq \text{dangling}(\mathfrak{h}_1) \cup \text{dangling}(\mathfrak{h}_2) \subseteq \text{img}(\mathfrak{s})$ □

Furthermore, due to establishment (cf. Section 3.4.5), models of predicate calls are guarded:

LEMMA 4.4. For all predicates $\text{pred} \in \mathbf{Preds}(\Phi)$ and all states $\langle \mathfrak{s}, \mathfrak{h} \rangle$, we have

$$\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{pred}(\mathbf{x}) \quad \text{implies} \quad \langle \mathfrak{s}, \mathfrak{h} \rangle \in \mathbf{GStates}.$$

PROOF. By induction on the number of rule applications needed to establish $\langle s, h \rangle \models_{\Phi} \text{pred}(x)$; a detailed proof is found in Appendix A.2. \square

We now lift the result from Lemma 4.4 from atomic predicates to arbitrary guarded formulas. We will use the following result that every model of a guarded formula satisfies a finite number of predicates conjoined by the separating conjunction:

LEMMA 4.5. *Let $\phi \in \text{GSL}$ be a guarded formula with $\text{fvars}(\phi) = \mathbf{x}$. Then, for every state $\langle s, h \rangle \models_{\Phi} \phi$, there are predicates $\text{pred}_i \in \text{Preds}(\Phi)$ and variables $z_i \subseteq \mathbf{x}$ such that $\langle s, h \rangle \models_{\Phi} \star_{1 \leq i \leq k} \text{pred}_i(z_i)$.*

PROOF. By structural induction on ϕ ; see Appendix A.4 for details. \square

COROLLARY 4.6. *For all $\phi \in \text{GSL}$ and all states $\langle s, h \rangle \models_{\Phi} \phi$, we have $\langle s, h \rangle \in \text{GStates}$.*

PROOF. Immediate from Lemma 4.4 and Lemma 4.5. \square

On the importance of guardedness. The fact that all appearances of negation, magic wand, and sepraction in GSL are guarded by a conjunction with another GSL formula is crucial for limiting the number of dangling pointers in the above lemma.

For the negation and the magic wand this is straightforward: without guards, both can be used to define true, e.g., $\mathbf{emp} \vee \neg \mathbf{emp}$ and $((x \mapsto \text{nil}) \star (x \mapsto \text{nil})) \rightarrow \mathbf{emp}$. Since true is satisfied by all states, the number of dangling locations is unbounded. For the sepraction \oplus , consider the following SID:

$$\begin{aligned} \text{tll}(r, l, t) &\Leftarrow l \mapsto t \star r \approx l & \text{lseg}(l, t) &\Leftarrow \exists n. (l \mapsto n) \star \text{lseg}(n, t) \\ \text{tll}(r, l, t) &\Leftarrow \exists \langle u, v, m \rangle. (r \mapsto \langle u, v \rangle) \star \text{tll}(u, l, m) \star \text{tll}(s_2, m, r) & \text{lseg}(l, t) &\Leftarrow l \mapsto t \end{aligned}$$

The tll predicate encodes a binary tree with root r and leftmost leaf l overlaid with a singly-linked list segment from l to t whose nodes are the leaves of the tree. Now, assume a state $\langle s, h \rangle$ satisfying the unguarded formula $\text{lseg}(l, t) \oplus \text{tll}(r, l, t)$. In other words, there exists a heap h_1 with $\langle s, h_1 \rangle \models_{\Phi} \text{lseg}(l, t)$ and $\langle s, h \cup h_1 \rangle \models_{\Phi} \text{tll}(r, l, t)$. Since each list element is a leaf of the tree in heap h , we have $\text{dangling}(h) = \text{dom}(h_1)$ —a finite, but unbounded set of dangling pointers.

5 BEYOND GUARDED SEPARATION LOGIC: UNDECIDABILITY PROOFS

Before we develop our decision procedure for the fragment GSL of guarded separation logic with inductive definitions of bounded treewidth, we further justify the need for guarding negation, magic wand, and sepraction. More precisely, we show in this section that omitting the guards for *any* of the above three operators leads to an undecidable logic. Together with our decidability results presented afterwards, this yields an almost tight delineation between undecidability of separation logics that allow arbitrary SIDs in ID_{btw} .

5.1 Encoding Context-Free Language in SIDs

All of our undecidability results, which are presented in Section 5.2, rely on a novel encoding of the language-intersection problem for context-free grammars—a well-known undecidable problem.

DEFINITION 5.1 (CONTEXT-FREE GRAMMAR). *A context-free grammar (CFG) in Chomsky normal form is a 4-tuple $G = \langle \mathbf{N}, \mathbf{T}, \mathbf{R}, S \rangle$, where \mathbf{N} is a finite set of nonterminals; \mathbf{T} is a finite set of terminals, which is disjoint from \mathbf{N} ; $\mathbf{R} \subseteq \mathbf{N} \times (\mathbf{N}^2 \cup \mathbf{T})$ is a finite set of production rules mapping nonterminals to two nonterminals or a single terminal; and $S \in \mathbf{N}$ is the start symbol. CFG is the set of all CFGs.*

We often denote production rules $\langle a, b \rangle$ by $a \rightarrow b$ to improve readability. Since we assume all CFGs to be in Chomsky normal form, all rules are either of the form $N \rightarrow AB$ or $N \rightarrow a$, where N, A, B are nonterminals in \mathbf{N} and a is a terminal in \mathbf{T} .

$$\begin{aligned}
\text{letter}_i(a) &\Leftarrow a \mapsto \underbrace{\langle \text{nil}, \dots, \text{nil} \rangle}_{\text{length } i}, \quad 1 \leq i \leq n \\
N(x_1, x_2, x_3) &\Leftarrow \exists l, r, m. (x_1 \mapsto \langle l, r \rangle) \star A(l, x_2, m) \star B(r, m, x_3), \quad j \in \{1, 2\}, (N \rightarrow AB) \in \mathbf{R}_j \\
N(x_1, x_2, x_3) &\Leftarrow \exists a. (x_1 \mapsto \langle x_3, a \rangle) \star \text{letter}_k(a) \star x_1 \approx x_2, \quad j \in \{1, 2\}, (N \rightarrow a_k) \in \mathbf{R}_j \\
\text{word}(x, y) &\Leftarrow \exists a. (x \mapsto \langle y, a \rangle) \star \text{letter}_i(a), \quad 1 \leq i \leq n \\
\text{word}(x, y) &\Leftarrow \exists \langle n, a \rangle. (x \mapsto \langle n, a \rangle) \star \text{letter}_i(a) \star \text{word}(n, y), \quad 1 \leq i \leq n
\end{aligned}$$

Fig. 4. The SID Φ encoding derivations of the CFGs $G_1 = \langle N_1, T, R_1, S_1 \rangle$ and $G_2 = \langle N_2, T, R_2, S_2 \rangle$.

DEFINITION 5.2. Let $G = \langle N, T, R, S \rangle \in \text{CFG}$ and let $v, w \in (N \cup T)^*$. We write $v \Rightarrow w$ if there exist strings $u_1, u_2 \in (N \cup T)^*$ and a rule $a \rightarrow b$ such that $v = u_1 \cdot a \cdot u_2$ and $w = u_1 \cdot b \cdot u_2$. We write \Rightarrow^+ for the transitive closure of \Rightarrow . The language of G is given by $\mathcal{L}(G) \triangleq \{w \in T^* \mid S \Rightarrow^+ w\}$.

In the following, we exploit the following classic undecidability result (cf. [Bar-Hillel et al. 1961]):

THEOREM 5.3 (UNDECIDABILITY OF LANGUAGE INTERSECTION). Given two CFGs G_1 and G_2 , it is undecidable whether $\mathcal{L}(G_1) \cap \mathcal{L}(G_2) \neq \emptyset$ holds—even if neither G_1 nor G_2 accept the empty string $\langle \rangle$.

Encoding CFGs as SIDs. Throughout the remainder of this section, we fix a set $T = \{a_1, \dots, a_n\}$ of terminals and two CFGs G_1 and G_2 , where we assume that their sets of nonterminals do not overlap, i.e., $N_1 \cap N_2 = \emptyset$. Fig. 4 depicts the SID Φ encoding G_1 and G_2 : For each terminal symbol a_i , we introduce a predicate $\text{letter}_i(a)$.⁴ Moreover, for each nonterminal $N \in N_1 \cup N_2$, there is a corresponding predicate encoding the derivations of G_1 and G_2 as trees with linked leaves (TLL), similar to the SID in Fig. 1 on p. 2. The predicate word overapproximates the possible *front*, i.e., the list of linked leaves of the TLL; we will need it later to prove undecidability of individual fragments.

By construction, every word in $\mathcal{L}(G_i)$ corresponds to at least one state $\langle s, h \rangle$ with $\langle s, h \rangle \models_{\Phi} S_i(x_1, x_2, x_3)$. Furthermore, every model $\langle s, h \rangle$ of $S_i(x_1, x_2, x_3)$ corresponds to both a *derivation tree*⁵ and a word in $\mathcal{L}(G_i)$, where the inner nodes of the TLL correspond to the derivation tree and its front corresponds to the word in $\mathcal{L}(G_i)$.

EXAMPLE 5.4. Fig. 5 illustrates both a derivation tree (Figure 5b) and a model of our encoding (Fig. 5c) for the CFG $G = \langle \{S, A, B, C\}, \{a_1, a_2\}, R, S \rangle$ whose rules are provided in Figure 5a. We observe that the depicted model encodes the aforementioned derivation tree: Every nonterminal is translated to a node in a binary tree (blue). The leaves of the tree are linked. They each have a successor that encodes a terminal symbol of the derivation (orange): The node contains k pointers to nil to represent terminal a_k . The list of linked leaves and orange nodes together form the induced word, i.e., $a_2 a_2 a_1 a_1 a_1$.

To show that our encoding is correct, i.e., it adequately captures the language of a given CFG, we need to refer to the word induced by a given model. We first define the terminals of such a word, which are given by the letter predicates in the models' list of linked leaves.

⁴While it is convenient to model each terminal a_i through a single points-to assertion mapping a to i null pointers, it is noteworthy, that, in principle, points-to assertions mapping to at most two locations suffice.

⁵We do not formally define derivation trees for CFGs as they are not required for the formal development; we refer to [Hopcroft et al. 2007] for a thorough introduction of CFGs.

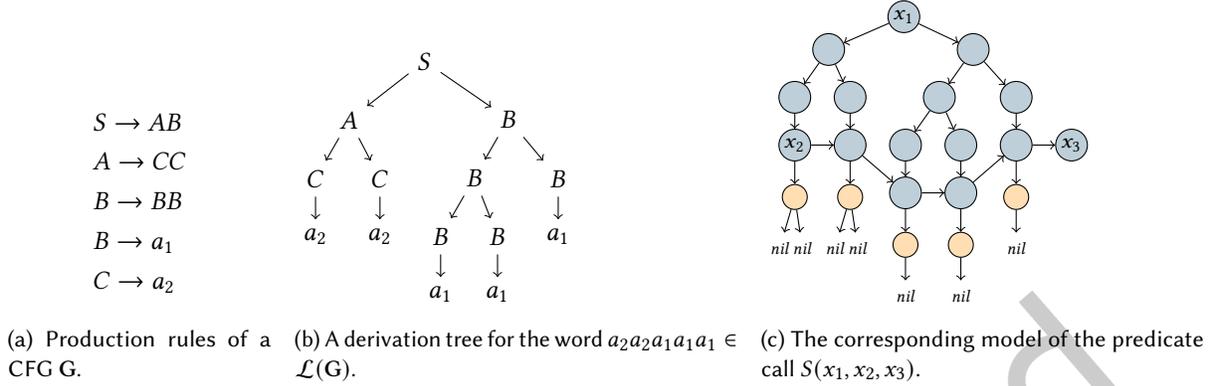


Fig. 5. Encoding a derivation of a context-free grammar as a stack-heap model.

DEFINITION 5.5 (INDUCED LETTERS). Let $G = \langle N, T, R, S \rangle$ and let Φ be the corresponding SID encoding. Let $\langle s, h \rangle \models_{\Phi} \text{word}(x, y)$ and let $j_1, \dots, j_m \in \{1, \dots, n\}$ be such that

$$\begin{aligned} \langle s, h \rangle \models_{\Phi} \exists n_1, \dots, n_{m-1}, b_1, \dots, b_m. & ((x \mapsto \langle n_1, b_1 \rangle) \star \text{letter}_{j_1}(b_1)) \\ & \star ((n_1 \mapsto \langle n_2, b_2 \rangle) \star \text{letter}_{j_2}(b_2)) \\ & \star \dots \star ((n_{m-1} \mapsto \langle y, b_m \rangle) \star \text{letter}_{j_m}(b_m)). \end{aligned}$$

We define the induced letters of $\langle s, h \rangle$ and x, y as $\text{letters}(s, h, x, y) \triangleq a_{j_1} a_{j_2} \dots a_{j_m}$.

Every model of the predicate $N(x_1, x_2, x_3)$ contains a sub-heap satisfying the word predicate:

LEMMA 5.6. Let $G = \langle N, T, R, S \rangle$ and let Φ be its SID encoding. Let $x_1, x_2, x_3 \in \text{Var}$, $N \in N$ and let $\langle s, h \rangle \models_{\Phi} N(x_1, x_2, x_3) \star t$. Then there exists a unique heap $h_w \subseteq h$ with $\langle s, h_w \rangle \models_{\Phi} \text{word}(x_2, x_3)$.

PROOF. A straightforward induction shows that the models of the predicate call $N(x_1, x_2, x_3)$ are trees with linked leaves with root $s(x_1)$, leftmost leaf $s(x_2)$, and successor of rightmost leaf $s(x_3)$. We pick as h_w the heap that contains $s(x_2)$ as well as all locations reachable from $s(x_2)$ in h . This gives us precisely the list from $s(x_2)$ to $s(x_3)$. Moreover, every leaf satisfies a formula of the form $\exists a. (y \mapsto \langle z, a \rangle) \star \text{letter}_k(a)$. Consequently, $\langle s, h_w \rangle \models_{\Phi} \text{word}(x_2, x_3)$. \square

Lemma 5.6 ensures that models of our encoding of CFGs induce a word over the given alphabet.

DEFINITION 5.7 (INDUCED WORD). Let $G = \langle N, T, R, S \rangle$ and let Φ be its SID encoding. Let $x_1, x_2, x_3 \in \text{Var}$, $N \in N$ and let $\langle s, h \rangle \models_{\Phi} N(x_1, x_2, x_3)$. Let $h_w \subseteq h$ be the unique heap with $\langle s, h_w \rangle \models_{\Phi} \text{word}(x_2, x_3)$. The induced word of $\langle s, h \rangle$ and N is $\text{wordof}_N(s, h, x_2, x_3) \triangleq \text{letters}(s, h_w, x_2, x_3)$.

Every word $w \in \mathcal{L}(G)$ is then the induced word of a model of the corresponding SID encoding.

LEMMA 5.8 (COMPLETENESS OF THE ENCODING). Let $G = \langle N, T, R, S \rangle$ and let Φ be the corresponding SID encoding. Let $x_1, x_2, x_3 \in \text{Var}$ and let $w \in \mathcal{L}(G)$. Then there exists a model $\langle s, h \rangle$ of $S(x_1, x_2, x_3)$ with $\text{wordof}_S(s, h, x_2, x_3) = w$.

PROOF. By mathematical induction on the number of \Rightarrow steps; see Appendix A.5. \square

Likewise, every induced word of a model of the corresponding SID encoding is in $\mathcal{L}(G)$.

LEMMA 5.9 (SOUNDNESS OF THE ENCODING). Let $G = \langle N, T, R, S \rangle$ and let Φ be its SID encoding. Let $x_1, x_2, x_3 \in \text{Var}$ and let $\langle s, h \rangle \models_{\Phi} S(x_1, x_2, x_3)$. Then $\text{wordof}_S(s, h, x_2, x_3) \in \mathcal{L}(G)$.

PROOF. By induction on the height h of the tree contained in \mathfrak{h} ; see Appendix A.6. \square

5.2 Undecidability of Unguarded Fragments

We are now ready to prove that omitting guards leads to undecidable SL fragments. To conveniently describe these fragments, we write $\mathbf{SL}_{\text{btw}}(\cdot_1, \dots, \cdot_k)$ for the restriction of quantifier-free formulas in \mathbf{SL}_{btw} to formulas built from all atoms as well as the additional symbols and connectives \cdot_1, \dots, \cdot_k . For example, formulas in $\mathbf{SL}_{\text{btw}}(\wedge, \star, \text{t})$ are built from atomic predicates, the predicate t (true), and the binary connectives \star, \wedge . As usual, t holds in all models, i.e., $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{t}$ for all states $\langle \mathfrak{s}, \mathfrak{h} \rangle$.

First, we show that allowing t as well as both standard and separating conjunction \wedge and \star (i.e., **GSL** without disjunction or any of the guarded connectives) immediately leads to undecidability.

THEOREM 5.10. *The satisfiability problem for the fragment $\mathbf{SL}_{\text{btw}}(\wedge, \star, \text{t})$ is undecidable.*

PROOF. Let Φ be the encoding of the CFGs $\mathbf{G}_1 = \langle \mathbf{N}_1, \mathbf{T}, \mathbf{R}_1, \mathbf{S}_1 \rangle$ and $\mathbf{G}_2 = \langle \mathbf{N}_2, \mathbf{T}, \mathbf{R}_2, \mathbf{S}_2 \rangle$ as described in Section 5.1. Moreover, consider the $\mathbf{SL}_{\text{btw}}(\wedge, \star, \text{t})$ formula $\phi \triangleq (\mathbf{S}_1(a, x, y) \star \text{t}) \wedge (\mathbf{S}_2(b, x, y) \star \text{t})$. Then ϕ is satisfiable iff $\mathcal{L}(\mathbf{G}_1) \cap \mathcal{L}(\mathbf{G}_2) \neq \emptyset$; see Appendix A.7 for details. \square

COROLLARY 5.11. *The satisfiability problem of $\mathbf{SL}_{\text{btw}}(\wedge, \star, \neg)$ is undecidable.*

PROOF. This follows directly from the undecidability of $\mathbf{SL}_{\text{btw}}(\wedge, \star, \text{t})$ (Theorem 5.10), because t is definable in $\mathbf{SL}_{\text{btw}}(\wedge, \star, \neg)$; for example, $\text{t} \triangleq \neg(\mathbf{emp} \wedge \neg \mathbf{emp})$. \square

COROLLARY 5.12. *The satisfiability problem of $\mathbf{SL}_{\text{btw}}(\wedge, \star, \rightarrow)$ is undecidable.*

PROOF. Follows directly from the undecidability of $\mathbf{SL}_{\text{btw}}(\wedge, \star, \text{t})$ (Theorem 5.10), because t is definable in $\mathbf{SL}_{\text{btw}}(\wedge, \star, \rightarrow)$; for example, $\text{t} \triangleq (x \neq x) \rightarrow \mathbf{emp}$. \square

Our final undecidability proof concerns septractions. We need one more auxiliary result before we can prove this result.

LEMMA 5.13. *Let $\mathbf{G}_2 = \langle \mathbf{N}_2, \mathbf{T}, \mathbf{R}_2, \mathbf{S}_2 \rangle$ be the CFG fixed in Section 5.1. Moreover, let Φ be the corresponding SID encoding, $\text{word}_2(x, y) \triangleq (\text{word}(x, y) \oplus \mathbf{S}_2(a, x, y)) \oplus \mathbf{S}_2(a, x, y)$, and let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state. Then $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{word}_2(x, y)$ iff $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{word}(x, y)$ and $\text{letters}(\mathfrak{s}, \mathfrak{h}, x, y) \in \mathcal{L}(\mathbf{G}_2)$.*

PROOF. See Appendix A.8. \square

To prove the undecidability of separation logic in the presence of unguarded septractions, we show that $\psi \triangleq \text{word}_2(x, y) \oplus \mathbf{S}_1(a, x, y)$ is satisfiable iff $\mathcal{L}(\mathbf{G}_1) \cap \mathcal{L}(\mathbf{G}_2) \neq \emptyset$. Intuitively, this holds because ψ is satisfiable iff it is possible to replace the “word part” of models of $\mathbf{S}_1(a, x, y)$ with the “word part” of models of $\mathbf{S}_2(b, x, y)$.

THEOREM 5.14. *The satisfiability problem of $\mathbf{SL}_{\text{btw}}(\oplus)$ is undecidable.*

PROOF. $\psi \triangleq \text{word}_2(x, y) \oplus \mathbf{S}_1(a, x, y)$ is satisfiable iff $\mathcal{L}(\mathbf{G}_1) \cap \mathcal{L}(\mathbf{G}_2) \neq \emptyset$; see Appendix A.9. \square

We have shown that all extensions of the guarded fragment **GSL**, in which one of the guards is dropped, lead to an undecidable satisfiability problem. In the remainder of this paper, we develop a decision procedure for **GSL**; keeping all guards thus indeed ensures decidability.

6 TOWARDS A COMPOSITIONAL ABSTRACTION FOR GSL

In Section 1, we sketched our goal of using a finite *compositional abstraction* that *refines* the satisfaction relation in order to decide the satisfiability problem for the separation logic fragment GSL. The same procedure then also allows deciding entailments between (quantifier-free) symbolic heaps in SH^\exists with user-defined predicates (defined by rules that may, of course, contain quantifiers). The key challenge is to develop an abstraction mechanism that can deal with arbitrary user-defined predicates from the ID_{btw} fragment. To get an abstraction that satisfies refinement, we need to be able to deduce from the abstraction which predicate calls hold in the underlying model. To this end, we will abstract every state by a set of formulas that relates the state to predicates of the SID.

In the following we introduce our abstraction, starting with a simple but insufficient idea, and then incrementally improve on it.

Purpose of this section. This section serves as a roadmap; it outlines the main concepts underlying our decision procedure and explains them informally by means of examples. We will formalize all of these concepts in follow-up sections—references to the formal details are provided where appropriate. Similarly, the remaining sections will frequently refer back to this section to either give further details on the examples, or to pin-point the progress of our formalization.

6.1 First Attempt: Abstracting States by Symbolic Heaps

Our first idea is to abstract a state by the quantifier-free symbolic heaps that it satisfies:

$$\text{abst}_1(\mathfrak{s}, \mathfrak{h}) \triangleq \{\phi \in \text{SH}^\exists \mid \phi \text{ quantifier-free, } \langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi\}.$$

Let us analyze the properties of this abstraction function. (For the moment, we ignore whether we can actually compute this abstraction.)

6.1.1 Finiteness. The abstraction abst_1 is finite (for a fixed number of variables, given by the domain of the stack \mathfrak{s}), because there are only finitely many quantifier-free symbolic heaps up to logical equivalence: Since $\Phi \in \text{ID}_{\text{btw}}$, every predicate call in a symbolic heap ϕ has to allocate at least one free variable due to the *progress* property (cf. Section 3.4.3); the same holds for every points-to assertion. Consequently, every satisfiable quantifier-free formula can contain at most $|\text{fvvars}(\phi)|$ many predicate calls and points-to assertions. In principle, we can, of course, “blow up” a satisfiable formula ϕ to arbitrary size by adding **emp** atoms and (dis-)equalities, but any fixed *aliasing constraint* over $\text{fvvars}(\phi)$, i.e., any fixed relationship between the free variables of formula ϕ , can be expressed with at most $|\text{fvvars}(\phi)|^2$ such atoms. Hence, we are able to obtain a bound on the size of the symbolic heaps that need to be considered for constructing the abstraction.

6.1.2 Refinement. Recall that our abstraction satisfies refinement iff states leading to the same abstraction satisfy the same formulas. This immediately holds for abst_1 —at least on the quantifier-free symbolic-heap fragment of GSL.

6.1.3 Compositionality. Can we also *compose* abstractions, i.e., can we find a (computable) operator \bullet with $\text{abst}_1(\mathfrak{s}, \mathfrak{h}_1) \bullet \text{abst}_1(\mathfrak{s}, \mathfrak{h}_2) = \text{abst}_1(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2)$? Unfortunately, the example below demonstrates that finding such an operator is quite challenging. Assume that Φ defines the list-segment predicate lseg (cf. Example 3.3). Moreover, consider a state $(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2)$ such that

- $\mathfrak{s}(u) \neq \mathfrak{s}(v)$ for all $u, v \in \{x, y, z\}$ with $u \neq v$,
- $\text{abst}_1(\mathfrak{s}, \mathfrak{h}_1) = \{\text{lseg}(x, y)\}$, $\text{abst}_1(\mathfrak{s}, \mathfrak{h}_2) = \{\text{lseg}(y, z)\}$, and $\text{abst}_1(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2) = \{\text{lseg}(x, z)\}$,

where we omit pure constraints in the sets $\text{abst}_1(\dots)$ for readability. We highlight that it is a-priori unclear how to infer that $\text{lseg}(x, z)$ holds in the composed state, i.e., that $\text{abst}_1(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2) = \{\text{lseg}(x, z)\}$ —at least by relying solely on the assumptions $\text{abst}_1(\mathfrak{s}, \mathfrak{h}_1) = \{\text{lseg}(x, y)\}$ and $\text{abst}_1(\mathfrak{s}, \mathfrak{h}_2) = \{\text{lseg}(y, z)\}$. In particular, it is unclear

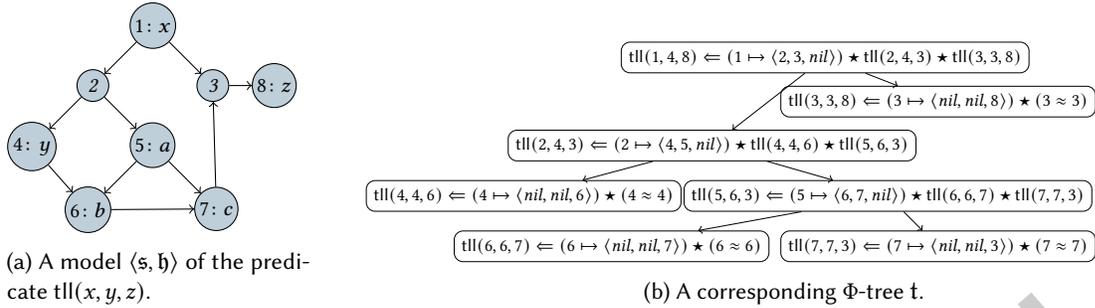


Fig. 6. A model $\langle s, h \rangle \models_{\Phi} tll(x, y, z)$ and the Φ -tree t corresponding to this model.

how to derive this fact by a syntactic argument. One might turn towards an argument based on the semantics, i.e., considering the definition of $lseg$ in Φ . However, then the above composition operation \bullet boils down to an entailment check $lseg(x, y) \star lseg(y, z) \models_{\Phi} lseg(x, z)$. Hence, we end up with a chicken-and-egg problem: we need an entailment checker to implement the composition operator \bullet that we would like to use in the implementation of our abstraction-based (satisfiability and) entailment checker.

6.2 Second Attempt: Unfolding Predicates into Forests

Next, we attempt to extend the abstraction $abst_1$ to get a “more syntactic” composition operation. To this end, we need to take a step back and reflect on the semantics of SIDs.

6.2.1 Unfolding predicate calls. According to the SL semantics (Section 3.3), $\langle s, h \rangle \models_{\Phi} pred(z)$ holds iff there exists a rule $pred(x) \Leftarrow \psi(x) \in \Phi$ such that $\langle s, h \rangle \models_{\Phi} \psi(z)$. We say that we have *unfolded* the predicate $pred$ by the above rule. In general, ψ may itself contain predicate calls. To prove $\langle s, h \rangle \models_{\Phi} \psi(z)$, we must continue unfolding the remaining predicate calls according to rules of the respective predicates until no predicate calls remain.

It is natural to visualize an unfolding process as a tree. In fact, defining the semantics of inductive predicates based on such *unfolding trees* is a common approach in the literature (cf. [Iosif et al. 2013, 2014; Jansen et al. 2017; Matheja 2020]). In this article, we use a variant of unfolding trees, called Φ -trees, which we will formally introduce in Definition 7.1.

EXAMPLE 6.1 (Φ -TREE). Recall the SID Φ from Fig. 1, which defines the predicate tll . Figure 6a depicts a state $\langle s, h \rangle$ with $\langle s, h \rangle \models_{\Phi} tll(x, y, z)$. Each node is labeled with a location and the stack variable evaluating to the location (if any). The depicted state $\langle s, h \rangle$ thus corresponds to

$$\begin{aligned} s &= \{x \mapsto 1, y \mapsto 4, z \mapsto 8, a \mapsto 5, b \mapsto 6, c \mapsto 7\}, \text{ and} \\ h &= \{1 \mapsto \langle 2, 3, nil \rangle, 2 \mapsto \langle 4, 5, nil \rangle, 5 \mapsto \langle 6, 7, nil \rangle, 4 \mapsto \langle nil, nil, 6 \rangle, \\ &\quad 6 \mapsto \langle nil, nil, 7 \rangle, 7 \mapsto \langle nil, nil, 3 \rangle, 3 \mapsto \langle nil, nil, 8 \rangle\}. \end{aligned}$$

Figure 6b shows a Φ -tree t corresponding to this state. Each node of t is labeled with a rule instance, i.e., a rule of the SID in which all variables—both formal parameters and existentially-quantified variables—have been instantiated with the locations of the state. This is different from other notions of unfolding trees, e.g., [Iosif et al. 2013], in which nodes are labeled by rules, not rule instances. Note that t induces the heap h : h is the union of all the points-to assertions that occur in the node labels of t . We denote this heap by $heap(t)$.

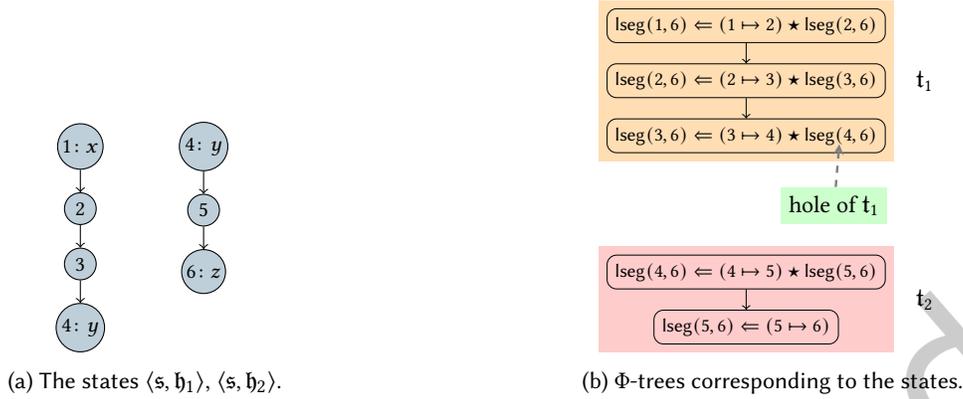


Fig. 7. States $\langle s, h_1 \rangle \models_{\Phi} \text{lseg}(x, y)$ and $\langle s, h_2 \rangle \models_{\Phi} \text{lseg}(y, z)$ and Φ -trees t_1, t_2 corresponding to the states. The tree t_1 contains one predicate call that is not unfolded, $\text{lseg}(4, 6)$. We say that 4, the root of this folded predicate call, is a hole of the tree.

Φ -trees enable an alternative reading of the entailment $\text{lseg}(x, y) \star \text{lseg}(y, z) \models_{\Phi} \text{lseg}(x, z)$: The entailment is valid iff whenever $\langle s, h \rangle \models_{\Phi} \text{lseg}(x, y) \star \text{lseg}(y, z)$ holds, it is possible to find a Φ -tree t with root $\text{lseg}(s(x), s(z))$ such that $\text{heap}(t) = h$.

6.2.2 Abstracting states by forests. Our next abstraction attempt is to encode the existence of a suitable Φ -tree. More precisely, we encode that the models of $\text{lseg}(x, y)$ and $\text{lseg}(y, z)$ each correspond to *partial* Φ -trees of $\text{lseg}(x, z)$ that can be combined into an unfolding tree of $\text{lseg}(x, z)$. A partial Φ -tree is obtained by prematurely stopping the unfolding process. Consequently, such a tree may contain *holes*—predicate calls that have not been unfolded.

EXAMPLE 6.2 (Φ -TREES WITH HOLES). Recall the entailment $\text{lseg}(x, y) \star \text{lseg}(y, z) \models_{\Phi} \text{lseg}(x, z)$ from above. Figure 7a shows states $\langle s, h_1 \rangle, \langle s, h_2 \rangle$ with $\langle s, h_1 \rangle \models_{\Phi} \text{lseg}(x, y)$ and $\langle s, h_2 \rangle \models_{\Phi} \text{lseg}(y, z)$. By the semantics of \star , it holds for $h \triangleq h_1 \uplus h_2$ that $\langle s, h \rangle \models_{\Phi} \text{lseg}(x, y) \star \text{lseg}(y, z)$.

How can Φ -trees be used to argue that $\langle s, h \rangle \models_{\Phi} \text{lseg}(x, z)$? Fig. 7b shows a Φ -forest—a set of Φ -trees—consisting of the partial Φ -trees t_1 and t_2 with $\text{heap}(t_1) = h_1$ and $\text{heap}(t_2) = h_2$, respectively. Notice that t_1 contains a hole: Since the predicate call $\text{lseg}(4, 6)$ is not unfolded, the hole, location 4, is not allocated in the tree, even though it is the root parameter of the hole predicate $\text{lseg}(4, 6)$. We can merge t_1 and t_2 into a larger tree by plugging t_2 into the hole of t_1 , i.e., we add an edge from the hole of t_1 to the root of t_2 . This is possible because the root of t_2 is labeled with the aforementioned hole predicate, $\text{lseg}(4, 6)$. The resulting tree is a Φ -tree for $\text{lseg}(x, z)$. That is, it is a tree without holes whose root is labeled with a rule instance of the predicate call $\text{lseg}(s(x), s(z))$. By merging the two trees, we verified the above entailment: every model of $\text{lseg}(x, y) \star \text{lseg}(y, z)$ is also a model of $\text{lseg}(x, z)$.

In fact, we go one step further and consider Φ -forests (cf. Definition 7.4), i.e., sets of partial Φ -trees. For example, the set $\{t_1, t_2\}$ illustrated in Fig. 7 is a Φ -forest.

EXAMPLE 6.3 (Φ -FOREST). Continuing Example 6.1, Fig. 8 depicts a Φ -forest $\mathfrak{f} = \{t_1, t_2, t_3, t_4\}$ that encodes one way to obtain the state $\langle s, h \rangle$ through iterative unfolding of predicate calls. Both t_1 and t_2 only partially unfold the predicates at their roots, leaving locations 5 resp. 6 and 7 as holes. By merging the four trees, we get the tree t from Example 6.1.

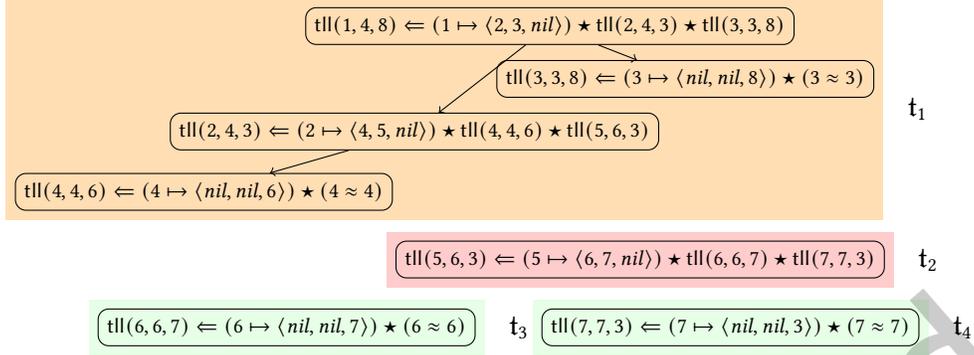


Fig. 8. A Φ -forest $\mathfrak{f} = \{t_1, t_2, t_3, t_4\}$ for the state from Example 6.1, used in Example 6.3.

Our second idea towards a suitable abstraction is to abstract a state by computing all Φ -forests consisting of trees whose combined heap matches the heap of the state:

$$\text{abst}_2(\mathfrak{s}, \mathfrak{h}) \triangleq \left\{ \mathfrak{f} \mid \mathfrak{f} \text{ is a } \Phi\text{-forest of } (\mathfrak{s}, \mathfrak{h}) \text{ with } \mathfrak{h} = \bigcup_{t \in \mathfrak{f}} \text{heap}(t) \right\}.$$

6.2.3 Compositionality. We can define a suitable composition operation $\text{abst}_2(\mathfrak{s}, \mathfrak{h}_1) \bullet \text{abst}_2(\mathfrak{s}, \mathfrak{h}_2)$ by computing all ways to merge the Φ -forests of $\text{abst}_2(\mathfrak{s}, \mathfrak{h}_1)$ and $\text{abst}_2(\mathfrak{s}, \mathfrak{h}_2)$. This approach yields precisely the set of all Φ -forests of $\langle \mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2 \rangle$, i.e., the set $\text{abst}_2(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2)$ from above, as required.

6.2.4 Finiteness. Unfortunately, abst_2 results in an *infinite* abstraction due to two main issues:

Issue 1 The tree nodes are labeled with concrete locations, so the result of abst_2 differs even if the states are isomorphic, i.e., identical up to renaming of locations. Apart from leading to an infinite abstraction, distinguishing such states is undesirable as they satisfy the same GSL formulas as long as we do not explicitly use constant locations.

Issue 2 If we keep track of all Φ -forests, the size of $\text{abst}_2(\mathfrak{s}, \mathfrak{h})$ grows with the size of \mathfrak{h} —it is unbounded. For example, the abstraction of a list-segment of size n contains the forest that consist of a single tree with n nodes, the forest that consists of n one-node trees as well as all possibilities in between.

6.3 Third Attempt: Forest Projections

Our first attempt yields a finite abstraction that is not compositional, whereas our second attempt is compositional but not finite. We now construct a finite *and* compositional abstraction by considering an additional abstraction—called the *projection*—on top of Φ -forests with holes. To this end, we denote by $\text{rootpred}(t)$ the root and by $\text{allholepreds}(t)$ the hole predicates of a Φ -tree t .

EXAMPLE 6.4. (1) Let t_1, t_2 be the Φ -trees from Example 6.2, which are illustrated in Fig. 7b. Then,

- $\text{rootpred}(t_1) = \text{lseg}(1, 6)$, $\text{allholepreds}(t_1) = \{\text{lseg}(4, 6)\}$, and
- $\text{rootpred}(t_2) = \text{lseg}(4, 6)$, $\text{allholepreds}(t_2) = \emptyset$.

(2) Let t_1, t_2, t_3, t_4 be the Φ -trees from Example 6.3, which are illustrated in Fig. 8. Then,

- $\text{rootpred}(t_1) = \text{tll}(1, 4, 8)$, $\text{allholepreds}(t_1) = \{\text{tll}(5, 6, 3)\}$,
- $\text{rootpred}(t_2) = \text{tll}(5, 6, 3)$, $\text{allholepreds}(t_2) = \{\text{tll}(6, 6, 7), \text{tll}(7, 7, 3)\}$,
- $\text{rootpred}(t_3) = \text{tll}(6, 6, 7)$, $\text{allholepreds}(t_3) = \emptyset$, and

- $\text{rootpred}(t_4) = \text{tll}(7, 7, 3)$, $\text{allholepreds}(t_4) = \emptyset$.

The *projection* of a Φ -forest \mathfrak{f} can be viewed as a GSL formula encoding, for each tree $t \in \mathfrak{f}$, a model of $\text{rootpred}(t)$ from which models of $\text{allholepreds}(t)$ have been subtracted, i.e.,⁶

$$\star_{t \in \mathfrak{f}} [(\star \text{allholepreds}(t)) \rightarrow \text{rootpred}(t)]. \quad (\dagger)$$

The goal of the projection operation is to combat Issue 2 identified above: to restore finiteness, we keep only limited information about each unfolding tree, and remember only its root predicate and its hole predicates. The magic wand introduced by the projection operation in the formula (\dagger) allows us to maintain the compositionality of the abstraction.

EXAMPLE 6.5 (FOREST PROJECTION—WITH LOCATIONS). *Recall from Example 6.2, the states $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle \models_{\Phi} \text{lseg}(x, y)$, $\langle \mathfrak{s}, \mathfrak{h}_2 \rangle \models_{\Phi} \text{lseg}(y, z)$, and the corresponding Φ -trees t_1, t_2 . The projection of stack \mathfrak{s} and Φ -forest $\{t_1, t_2\}$ is then the formula $(\text{lseg}(4, 6) \rightarrow \text{lseg}(1, 6)) \star (\text{emp} \rightarrow \text{lseg}(4, 6))$.*

6.3.1 Abstracting from locations. Our goal, which we will soon complete, has been to define a compositional abstraction over states. To this end, we introduced (partial) unfolding trees. These trees are naturally defined through the instantiation of SID rules with locations. Unfortunately, locations present an obstacle towards obtaining a finite abstraction (Issue 1): we get a different abstraction even for Φ -forests that encode the same model up to isomorphism! However, after having projected unfolding trees to formulas, we are able to reverse the instantiation of variables with locations. We in fact define the projection operation (\dagger) to output variables instead of locations: Say t is a Φ -tree of the state $\langle \mathfrak{s}, \mathfrak{h} \rangle$. Then we replace the locations in the formula (\dagger) as follows:

- (1) Every location $v \in \text{img}(\mathfrak{s})$ is replaced by a variable x satisfying $\mathfrak{s}(x) = v$.
- (2) Every location in $\text{dom}(\mathfrak{h}) \setminus \text{img}(\mathfrak{s})$ is replaced by an *existentially-quantified* variable, because there exists a location in the heap \mathfrak{h} that corresponds to the location in the formula (\dagger) .
- (3) All other locations are replaced by *universally-quantified* variables, because these locations do not occur in the heap \mathfrak{h} (this holds because we will always assume $\langle \mathfrak{s}, \text{heap}(t) \rangle \in \text{GStates}$) and can thus be picked in an arbitrary way.

We remark that the formal definition of projection uses non-standard quantifiers \exists and \forall in projections (further discussed below). For the moment, this difference does not matter; it is safe to replace them with the usual first-order quantifiers \exists and \forall for intuition.

EXAMPLE 6.6 (FOREST PROJECTION—WITH VARIABLES (WITHOUT QUANTIFIERS)). *Continuing Example 6.5, the projection of \mathfrak{s} and Φ -forest $\{t_1, t_2\}$ using the above replacement is*

$$(\text{lseg}(y, z) \rightarrow \text{lseg}(x, z)) \star (\text{emp} \rightarrow \text{lseg}(y, z)).$$

We will later prove that the projection operation is sound (Lemma 7.25), i.e., for the given example,

$$\langle \mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2 \rangle \models_{\Phi} (\text{lseg}(y, z) \rightarrow \text{lseg}(x, z)) \star (\text{emp} \rightarrow \text{lseg}(y, z)).$$

We further note that the idea of connecting holes with corresponding roots in Φ -trees is mirrored on the level of formulas: since the magic wand \rightarrow is the left-adjoint of the separating conjunction \star , an application of modus ponens⁷ (formalized in Lemma 7.18) suffices to establish that

$$\langle \mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2 \rangle \models_{\Phi} \text{emp} \rightarrow \text{lseg}(x, z).$$

⁶Recall that $\star \{\phi_1, \dots, \phi_n\}$ is a shortcut for $\phi_1 \star \dots \star \phi_n$.

⁷i.e., $\phi \star (\phi \rightarrow \psi) \Rightarrow \psi$

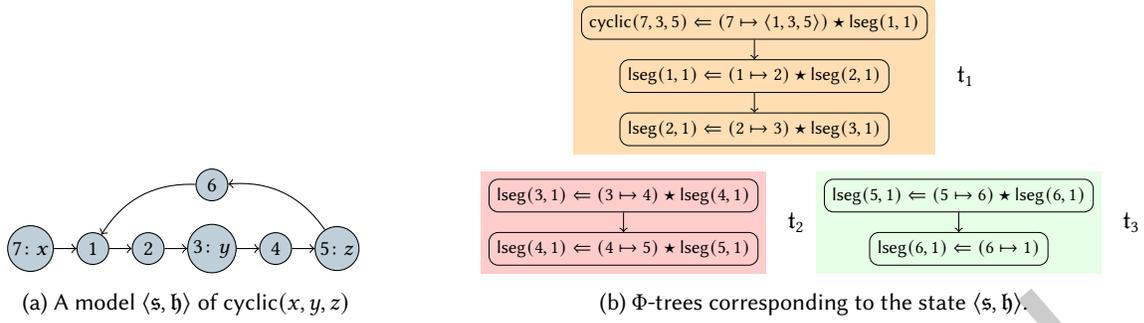


Fig. 9. A state $\langle s, h \rangle \models_{\Phi} \text{cyclic}(x, y, z)$ and Φ -trees corresponding to this state.

EXAMPLE 6.7 (FOREST PROJECTION—WITH VARIABLES (AND QUANTIFIERS)). We consider a Φ , consisting of the list segment predicate lseg (see Example 3.3) and the following additional predicate:

$$\text{cyclic}(x, y, z) \leftarrow \exists a. x \mapsto \langle a, y, z \rangle \star \text{ls}(a, a)$$

Fig. 9b depicts a model $\langle s, h \rangle$ of $\text{cyclic}(x, y, z)$ and Φ -trees t_1, t_2, t_3 with $h = \text{heap}(t_1) \cup \text{heap}(t_2) \cup \text{heap}(t_3)$. The projection of stack s and Φ -forest $\{t_1, t_2, t_3\}$ is

$$\exists a. (\text{lseg}(y, a) \star \text{cyclic}(x, y, z)) \star (\text{lseg}(z, a) \star \text{lseg}(y, a)) \star \text{lseg}(z, a).$$

We will later prove that the projection operation is sound (Lemma 7.25), i.e., for the given example,

$$\langle s, h \rangle \models_{\Phi} \exists a. (\text{lseg}(y, a) \star \text{cyclic}(x, y, z)) \star (\text{lseg}(z, a) \star \text{lseg}(y, a)) \star \text{lseg}(z, a).$$

Equipped with this extended projection operation, we are now in a position to specify the third (and almost final) abstraction function:

$$\text{abst}_3(s, h) \triangleq \left\{ \phi \mid \phi \text{ is the projection of a } \Phi \text{-forest } f \text{ of } (s, h) \text{ with } h = \bigcup_{t \in f} \text{heap}(t) \right\}.$$

6.3.2 *Compositionality.* As already hinted at in Example 6.6, the projection of formulas allows us to define a (computable) operator \bullet with $\text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2) = \text{abst}_3(s, h_1 \uplus h_2)$ such that:

- (1) We have $\phi \star \psi \in \text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2)$ for all $\phi \in \text{abst}_3(s, h_1)$ and $\psi \in \text{abst}_3(s, h_2)$.
- (2) The set of formulas $\text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2)$ is closed under application of modus ponens.
- (3) The set $\text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2)$ is closed under certain rules for manipulating quantifiers.

EXAMPLE 6.8 (COMPOSITION OPERATION ON PROJECTIONS). (1) We recall the states $\langle s, h_1 \rangle \models_{\Phi} \text{lseg}(x, y)$ and $\langle s, h_2 \rangle \models_{\Phi} \text{lseg}(y, z)$ from Example 6.2, and the Φ -trees t_1, t_2 . The projection of s and $\{t_1\}$ is $\text{lseg}(y, z) \star \text{lseg}(x, z)$, and the projection of s and $\{t_2\}$ is $\text{emp} \star \text{lseg}(y, z)$. Hence,

$$(\text{lseg}(y, z) \star \text{lseg}(x, z)) \star (\text{emp} \star \text{lseg}(y, z)) \in \text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2).$$

By applying modus ponens, we get $\text{emp} \star \text{lseg}(x, z) \in \text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2)$. The above reasoning approach will lead to a compositional proof of the entailment

$$\text{lseg}(x, y) \star \text{lseg}(y, z) \models_{\Phi} \text{lseg}(x, z).$$

- (2) Let $\langle s, h \rangle$ be the model and let t_1, t_2, t_3 be the Φ -trees from Example 6.7. We set $h_1 = \text{heap}(t_1) \cup \text{heap}(t_3)$ and $h_2 = \text{heap}(t_2)$. The projection of s and $\{t_1, t_3\}$ is

$$\exists a. (\text{lseg}(y, a) \rightarrow \text{cyclic}(x, y, z)) \star \text{lseg}(z, a),$$

and the projection of s and $\{t_2\}$ is $\forall a'. \text{lseg}(z, a') \rightarrow \text{lseg}(y, a')$. Hence, we have

$$\begin{aligned} & [\exists a. (\text{lseg}(y, a) \rightarrow \text{cyclic}(x, y, z)) \star \text{lseg}(z, a)] \star \\ & [\forall a'. \text{lseg}(z, a') \rightarrow \text{lseg}(y, a')] \in \text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2). \end{aligned}$$

By instantiating a' with a and moving $\exists a.$ to the front of the formula, we get that

$$\begin{aligned} & \exists a. (\text{lseg}(y, a) \rightarrow \text{cyclic}(x, y, z)) \star \text{lseg}(z, a) \star \\ & (\text{lseg}(z, a) \rightarrow \text{lseg}(y, a)) \in \text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2). \end{aligned}$$

By applying *modus ponens* (twice), we get $\exists a. \text{cyclic}(x, y, z) \in \text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2)$. As the variable a does not appear free anymore, the quantifier can be dropped, and we get

$$\text{cyclic}(x, y, z) \in \text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2).$$

The above reasoning approach will lead to a compositional proof of the entailment

$$\text{fork}(x, y, z) \star \text{lseg}(y, z) \models_{\Phi} \text{cyclic}(x, y, z),$$

where Φ extends the SID from Example 6.7 by the predicate

$$\text{fork}(x, y, z) \Leftarrow \exists a. (x \mapsto \langle a, y, z \rangle) \star \text{ls}(a, y) \star \text{ls}(z, a).$$

- (3) Let t_1, t_2, t_3, t_4 be the Φ -trees from Example 6.3 for the state $\langle s, h \rangle$ of Example 6.1. We set $h_1 = \text{heap}(t_1) \cup \text{heap}(t_3) \cup \text{heap}(t_4)$ and $h_2 = \text{heap}(t_2)$. The projection of s and $\{t_1, t_3, t_4\}$ is

$$\exists r. (\text{tll}(a, b, c) \rightarrow \text{tll}(x, y, z)) \star (b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star (c \mapsto \langle \text{nil}, \text{nil}, r \rangle),$$

and the projection of s and $\{t_2\}$ is $\forall r'. ((b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star (c \mapsto \langle \text{nil}, \text{nil}, r' \rangle)) \rightarrow \text{tll}(a, b, c)$. Hence,

$$\begin{aligned} & [\exists r. (\text{tll}(a, b, c) \rightarrow \text{tll}(x, y, z)) \star (b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star (c \mapsto \langle \text{nil}, \text{nil}, r \rangle)] \star \\ & [\forall r'. ((b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star (c \mapsto \langle \text{nil}, \text{nil}, r' \rangle)) \rightarrow \text{tll}(a, b, c)] \in \text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2). \end{aligned}$$

By instantiating r' with r and moving $\exists r.$ to the front of the formula, we get that

$$\begin{aligned} & \exists r. (\text{tll}(a, b, c) \rightarrow \text{tll}(x, y, z)) \star (b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star (c \mapsto \langle \text{nil}, \text{nil}, r \rangle) \star \\ & (((b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star (c \mapsto \langle \text{nil}, \text{nil}, r \rangle)) \rightarrow \text{tll}(a, b, c)) \in \text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2). \end{aligned}$$

By applying *modus ponens* for the magic wand (twice), we get that

$$\exists r. \text{tll}(x, y, z) \in \text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2).$$

As the variable r does not appear free anymore, the quantifier can be dropped and we get

$$\text{tll}(x, y, z) \in \text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2).$$

The above reasoning approach will lead to a compositional proof of the entailment

$$\text{tllHole}(x, y, z, a, b, c) \star (a \mapsto \langle b, c, \text{nil} \rangle) \models_{\Phi} \text{tll}(x, y, z),$$

where Φ extends the TLL SID from Fig. 1 by the predicates

$$\begin{aligned}
\text{tllHole}(x, y, z, a, b, c) &\Leftarrow \exists l, r. (x \mapsto \langle l, r, \text{nil} \rangle) \star \text{helper}(l, r, y, z, a, b, c) \\
\text{helper}(l, r, y, z, a, b, c) &\Leftarrow (l \mapsto \langle y, a, \text{nil} \rangle) \star \text{list4}(y, b, c, r, z) \\
\text{list4}(y, b, c, r, z) &\Leftarrow (y \mapsto \langle \text{nil}, \text{nil}, b \rangle) \star \text{list3}(b, c, r, z) \\
\text{list3}(b, c, r, z) &\Leftarrow (b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star \text{list2}(c, r, z) \\
\text{list2}(c, r, z) &\Leftarrow (c \mapsto \langle \text{nil}, \text{nil}, r \rangle) \star \text{ptr}(r, z).
\end{aligned}$$

6.3.3 Guarded Quantifiers. We now discuss the semantics of the special quantifiers \exists and \forall used in the projection operation. We rely on a non-standard semantics because we want our approach to directly support SIDs with (dis-)equalities. If one would disallow (dis-)equalities in SIDs, one could use the usual \exists and \forall quantifiers instead of \exists and \forall (which is sufficient for the SIDs in Example 6.8). We motivate our non-standard semantics with the SID Φ given by the following predicates:

$$p(x, a, b) \Leftarrow \exists y. (x \mapsto y) \star q(y, a) \star x \neq a \star a \neq b \quad q(y, a) \Leftarrow (y \mapsto \text{nil}) \star y \neq a$$

We consider the stack $\mathfrak{s} = \{x \mapsto 1\}$ and the heap $\mathfrak{h} = \{1 \mapsto 2, 2 \mapsto \text{nil}\}$. We further consider the unfolding tree t consisting of the root $p(1, 4, 5) \Leftarrow (1 \mapsto 2) \star q(2, 4) \star 1 \neq 4 \star 4 \neq 5$ with the single child $q(2, 4) \Leftarrow (2 \mapsto \text{nil}) \star 2 \neq 4$; note that $\text{heap}(t) = \mathfrak{h}$. The projection of this unfolding tree is the formula $\forall a, b. p(x, a, b)$. As discussed earlier, we want that the projection operation is sound, i.e., $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \forall a, b. p(x, a, b)$. However, using the standard quantifier \forall instead of \forall does not work:

$$(1) \langle \mathfrak{s}, \mathfrak{h} \rangle \not\models_{\Phi} p(1, 5, 5) \quad (2) \langle \mathfrak{s}, \mathfrak{h} \rangle \not\models_{\Phi} p(1, 1, 5) \quad (3) \langle \mathfrak{s}, \mathfrak{h} \rangle \not\models_{\Phi} p(1, 2, 5)$$

The above example shows that we need to prevent instantiating universally quantified variables with

- (1) identical locations, see $\langle \mathfrak{s}, \mathfrak{h} \rangle \not\models_{\Phi} p(1, 5, 5)$,
- (2) locations that are in the image of the stack, see $\langle \mathfrak{s}, \mathfrak{h} \rangle \not\models_{\Phi} p(1, 1, 5)$, and
- (3) locations that are existentially quantified, see $\langle \mathfrak{s}, \mathfrak{h} \rangle \not\models_{\Phi} p(1, 2, 5)$.

For the semantics of \forall we use that in SL_{btw} all existentially quantified variables (that are not equal to a parameter) are allocated because of the establishment requirement, and set

$$\begin{aligned}
\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \forall \langle a_1, \dots, a_k \rangle. \phi \text{ iff for all pairwise different locations} \\
v_1, \dots, v_k \in \text{Loc} \setminus (\text{dom}(\mathfrak{h}) \cup \text{img}(\mathfrak{s})) \text{ it holds that } \langle \mathfrak{s} \cup \{a_1 \mapsto v_1, \dots, a_k \mapsto v_k\}, \mathfrak{h} \rangle \models_{\Phi} \phi
\end{aligned}$$

Our main requirement for giving semantics to the \exists quantifier is the correctness of the following entailment, which we already used in Example 6.8: $(\exists e. \phi) \star (\forall a. \psi) \models_{\Phi} \exists e. \phi \star \psi[a/e]$ (\dagger)

This is ensured by the following semantics for \exists :

$$\begin{aligned}
\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \exists \langle e_1, \dots, e_k \rangle. \phi \text{ iff for all pairwise different locations} \\
v_1, \dots, v_k \in \text{dom}(\mathfrak{h}) \setminus \text{img}(\mathfrak{s}) \text{ such that } \langle \mathfrak{s} \cup \{e_1 \mapsto v_1, \dots, e_k \mapsto v_k\}, \mathfrak{h} \rangle \models_{\Phi} \phi.
\end{aligned}$$

We call our quantifiers \exists and \forall *guarded* because they exclude the instantiation of variables with certain locations. We note that our quantifiers are *not* dual, i.e., $\exists x. \phi$ is *not* equivalent to $\neg \forall x. \neg \phi$. However, we believe that our semantics is sufficiently motivated by our considerations on the soundness of the projection and the entailment (\dagger) .

6.3.4 Finiteness. Did we solve Issues 1 and 2 from the second attempt? Unfortunately not completely. However, one additional restriction on unfolding forests will be sufficient to guarantee the finiteness of the abstraction. We first explain the issue by means of an example:

EXAMPLE 6.9. Let Φ be an SID that defines the list-segment predicate lseg . Let $\langle s, h \rangle \models_{\Phi} \text{lseg}(x, \text{nil})$ with $|h| > n$. Then, there exists a forest \mathfrak{f} such that $h = \bigcup_{t \in \mathfrak{f}} \text{heap}(t)$ and whose projection is

$$\begin{aligned} & \exists y_1, \dots, y_n. \text{lseg}(y_n, \text{nil}) \star (\text{lseg}(y_n, \text{nil}) \rightarrow \text{lseg}(y_{n-1}, \text{nil})) \\ & \star \dots \star (\text{lseg}(y_2, \text{nil}) \rightarrow \text{lseg}(y_1, \text{nil})) \star (\text{lseg}(y_1, \text{nil}) \rightarrow \text{lseg}(x, \text{nil})) \end{aligned}$$

As there exist such models $\langle s, h \rangle$ for arbitrary $n \in \mathbb{N}$, there are infinitely many (non-equivalent) formulas resulting from projections of unfolding forests.

Fortunately, we do not need to consider all unfolding forests for deciding the satisfiability of the considered separation logic GSL. We recall that our goal is to define a compositional abstraction:

$$\text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2) = \text{abst}_3(s, h_1 \uplus h_2)$$

Hence, we need to ensure that every unfolding tree of $\langle s, h_1 \uplus h_2 \rangle$ can be composed via \bullet from unfolding trees of $\langle s, h_1 \rangle$ and $\langle s, h_2 \rangle$. In our approach we will have the guarantee that $\langle s, h_1 \rangle$ and $\langle s, h_2 \rangle$ are guarded (cf. Corollary 8.19). With this in mind, let us consider an unfolding tree of $\langle s, h_1 \uplus h_2 \rangle$ that is composed of some trees of $\langle s, h_1 \rangle$ and $\langle s, h_2 \rangle$, i.e., without loss of generality there is a pointer that is allocated in $\langle s, h_1 \rangle$ and points to a location in $\langle s, h_2 \rangle$. Then, this pointer is dangling for the state $\langle s, h_1 \rangle$ and the target of this pointer is in the image of the stack. Now, we recall from the definition of composition that the target of the pointer is a hole of an unfolding tree of $\langle s, h_1 \rangle$ and the root of an unfolding tree of $\langle s, h_2 \rangle$. Thus, we can restrict our attention to unfolding trees whose roots and holes are in the image of the stack! This motivates the following definition:

An unfolding tree t is s -delimited, if the root and holes of t are in the image of the stack s .

Equipped with this definition (which we will formalize in Definition 8.10), we restrict the abstraction function abst_3 to forests of delimited unfolding trees. This guarantees the finiteness of the abstraction: The formulas resulting from the projection of such forests have the property that (1) all root parameters of predicate calls are free variables and every variable occurs at most once as a root parameter, and (2) all root parameters of predicates calls on the left-hand side of a magic wand are free variables and every variable occurs at most once as a root parameter. Because the number of free variables is bounded, finiteness easily follows.

6.4 Summary of Overview

To sum up, we propose abstracting the state $\langle s, h \rangle$ in the following way:

- (1) We compute all s -delimited Φ -forests of $\langle s, h \rangle$.
- (2) We project these forests onto formulas.
- (3) The abstraction of $\langle s, h \rangle$ is the set of all these formulas; we call this set the *type* of $\langle s, h \rangle$.

The resulting abstraction is (1) finite (the set of types is finite), (2) compositional (we have $\text{abst}_3(s, h_1) \bullet \text{abst}_3(s, h_2) = \text{abst}_3(s, h_1 \uplus h_2)$), and (3) computable (we only need to apply rules for modus ponens and for manipulating quantifiers as illustrated in Example 6.8).

Outline of the following sections. In the remainder of this article, we give the technical details for the material overviewed in this section. In Section 7, we formalize Φ -forests (Section 7.1), their projections (Section 7.2), and how to compose forest projections (Section 7.3). The type abstraction is introduced in Section 8. We discuss how the satisfiability problem for guarded SL formulas can be reduced to computing types in Section 8.1. We formalize s -delimited forests in Section 8.3 and discuss how types can be computed compositionally in Sections 8.4 and 8.5. Finally, in Section 9, we present algorithms for computing the types of GSL formulas, summarize our overall decision procedure, and discuss our decidability and complexity results.

7 FORESTS AND THEIR PROJECTIONS

We now start formalizing the concepts that have been informally introduced in Section 6: Φ -forests (Section 7.1), their projection onto formulas (Section 7.2), and how to compose them (Section 7.3).

7.1 Forests

Our main objects of study in this section are Φ -forests (Definition 7.4) made up of Φ -trees (Definition 7.1). As motivated in Section 6, a Φ -tree encodes one fixed way to unfold a predicate call by means of the rules of the SID Φ . The differences between the unfolding trees of Iosif et al. [2013, 2014]; Jansen et al. [2017] and our Φ -trees are that (1) we instantiate variables with locations, and (2) Φ -trees can have *holes*, i.e., we allow that one or more of the predicate calls introduced (by means of recursive rules) in the unfolding process remain folded.

7.1.1 Rule instances. We annotate every node of a Φ -tree with a *rule instance* of the SID Φ , i.e., a formula obtained from a rule of the SID by instantiating both the formal arguments of the predicates and the existentially quantified variables of the rule with locations:

$$\begin{aligned} \mathbf{RuleInst}(\Phi) \triangleq \{ & \text{pred}(\mathbf{v}) \Leftarrow \phi[\mathbf{x} \cdot \mathbf{y}/\mathbf{v} \cdot \mathbf{w}] \mid (\text{pred}(\mathbf{x}) \Leftarrow \exists \mathbf{y}. \phi) \in \Phi, \\ & \mathbf{v} \in \mathbf{Loc}^{\text{ar}(\text{pred})}, \mathbf{w} \in \mathbf{Loc}^{|\mathbf{y}|}, \text{ and all (dis-)equalities in } \phi[\mathbf{x} \cdot \mathbf{y}/\mathbf{v} \cdot \mathbf{w}] \text{ are valid} \} \end{aligned}$$

In the above definition, we refer only to those (dis-)equalities that occur explicitly in the formula, not those implied by recursive calls or by the separating conjunction. Validity of these (dis-)equalities is straightforward to check because all variables have been instantiated with concrete locations.

The notion of a rule instance is motivated as follows: whenever $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{pred}(\mathbf{v})$, there is at least one rule instance $(\text{pred}(\mathbf{v}) \Leftarrow \psi) \in \mathbf{RuleInst}(\Phi)$ such that $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \psi$.

7.1.2 Φ -trees. We represent a Φ -tree as a partial function $t: \mathbf{Loc} \rightarrow (2^{\mathbf{Loc}} \times \mathbf{RuleInst}(\Phi))$, where the set \mathbf{Loc} of locations serves as the nodes of the tree; every node is mapped to its successors in the (directed) tree and to its label, a rule instance. Moreover, for t to be a Φ -tree, it must satisfy additional consistency criteria. To formalize these criteria, we fix some SID Φ and a node

$$t(l) = \langle \mathbf{v}, (\text{pred}(\mathbf{v}) \Leftarrow (a \mapsto \mathbf{b}) \star \text{pred}_1(\mathbf{v}_1) \star \dots \star \text{pred}_m(\mathbf{v}_m) \star \Pi) \rangle,$$

where Π is a set of equalities and disequalities. Notice that all rule instances are of the above form because—by our global assumptions in Section 3.4.5— Φ satisfies the progress property. We introduce the following shortcuts for the node at location l to simplify working with Φ -trees:

$$\begin{aligned} \text{succ}_t(l) &\triangleq \mathbf{v} && \text{(locations corresponding to the successors of node } l\text{)} \\ \text{head}_t(l) &\triangleq \text{pred}(\mathbf{v}) && \text{(the predicate on the lhs of the rule instance)} \\ \text{heap}_t(l) &\triangleq \{a \mapsto \mathbf{b}\} && \text{(the unique heap satisfying the points-to assertion in the rule instance)} \\ \text{calls}_t(l) &\triangleq \{\text{pred}_1(\mathbf{v}_1), \dots, \text{pred}_m(\mathbf{v}_m)\} && \text{(the predicate calls in the rule instance)} \\ \text{rule}_t(l) &\triangleq \text{pred}(\mathbf{v}) \Leftarrow (a \mapsto \mathbf{b}) \star \text{pred}_1(\mathbf{v}_1) \star \dots \star \text{pred}_m(\mathbf{v}_m) \star \Pi && \text{(the rule instance)} \end{aligned}$$

Moreover, we define the *hole predicates* of l as those predicate calls in $\text{calls}_t(l)$ whose root does not occur in $\text{succ}_t(l)$; the *holes* of l are the corresponding locations:

- $\text{holepreds}_t(l) \triangleq \{\text{pred}'(\mathbf{z}') \in \text{calls}_t(l) \mid \forall c \in \text{succ}_t(l). \text{head}_t(c) \neq \text{pred}'(\mathbf{z}')\}$, and
- $\text{holes}_t(l) \triangleq \{\text{predroot}(\text{pred}'(\mathbf{z}')) \mid \text{pred}'(\mathbf{z}') \in \text{holepreds}_t(l)\}$.

We lift some of the above definitions from individual locations to entire trees \mathfrak{t} :

$$\begin{aligned} \text{heap}(\mathfrak{t}) &\triangleq \bigcup_{c \in \text{dom}(\mathfrak{t})} \text{heap}_{\mathfrak{t}}(c) && \text{(the heap satisfying exactly the points-to assertions in } \mathfrak{t} \text{)} \\ \text{ptrlocs}(\mathfrak{t}) &\triangleq \bigcup_{(c \mapsto \mathbf{d}) \in \text{heap}(\mathfrak{t})} \{c\} \cup \mathbf{d} && \text{(all locations that appear in points-to assertions in } \mathfrak{t} \text{)} \\ \text{allholepreds}(\mathfrak{t}) &\triangleq \bigcup_{c \in \text{dom}(\mathfrak{t})} \text{holepreds}_{\mathfrak{t}}(c) && \text{(all hole predicates in } \mathfrak{t} \text{)} \\ \text{allholes}(\mathfrak{t}) &\triangleq \bigcup_{l \in \text{dom}(\mathfrak{t})} \text{holes}_{\mathfrak{t}}(l) && \text{(all holes in } \mathfrak{t} \text{)} \end{aligned}$$

We denote by $\text{graph}(\mathfrak{t})$ the directed graph induced by the successors of locations in \mathfrak{t} . That is,

$$\text{graph}(\mathfrak{t}) \triangleq \langle \text{dom}(\mathfrak{t}), \{ \langle x, y \rangle \mid x \in \text{dom}(\mathfrak{t}), y \in \text{succ}_{\mathfrak{t}}(x) \} \rangle.$$

The *height* of \mathfrak{t} is the length of the longest path in the directed graph $\text{graph}(\mathfrak{t})$.

DEFINITION 7.1 (Φ -TREE). A partial function $\mathfrak{t}: \text{Loc} \rightarrow (2^{\text{Loc}} \times \text{RuleInst}(\Phi))$ is a Φ -tree iff

- (1) Φ is in the fragment of SIDs of bounded treewidth, i.e., $\Phi \in \text{ID}_{\text{btw}}$,
- (2) $\text{graph}(\mathfrak{t})$ is a directed tree, and
- (3) \mathfrak{t} is Φ -consistent, i.e., for all locations $l \in \text{dom}(\mathfrak{t})$, we have:
 - l is the single allocated location in its rule instance, i.e., $\text{heap}_{\mathfrak{t}}(l) = \{l \mapsto \dots\}$,
 - l points to its successors in \mathfrak{t} , i.e., $\text{heap}_{\mathfrak{t}}(l) = \{l \mapsto \mathbf{b}\}$ implies $\text{succ}_{\mathfrak{t}}(l) \subseteq \mathbf{b}$, and
 - the predicate calls associated with the successors $\text{succ}_{\mathfrak{t}}(l) = \langle v_1, \dots, v_k \rangle$, of l appear in the rule instance at location l , i.e., $\{\text{head}_{\mathfrak{t}}(v_1), \dots, \text{head}_{\mathfrak{t}}(v_k)\} \subseteq \text{calls}_{\mathfrak{t}}(l)$.

Since every Φ -tree \mathfrak{t} is a directed tree, it has a root, which we denote by $\text{root}(\mathfrak{t})$; the corresponding predicate call is $\text{rootpred}(\mathfrak{t}) \triangleq \text{head}_{\mathfrak{t}}(\text{root}(\mathfrak{t}))$.

EXAMPLE 7.2 (Φ -TREE). (1) A Φ -tree over the SID $\Phi_{\text{odd/even}}$ (cf. Example 3.3) is given by

$$\mathfrak{t}(l) \triangleq \begin{cases} \langle b, \text{even}(l_1, a) \Leftarrow (l_1 \mapsto b) \star \text{odd}(b, a) \rangle & \text{if } l = l_1 \\ \langle \emptyset, \text{odd}(b, a) \Leftarrow (b \mapsto l_2) \star \text{even}(l_2, a) \rangle & \text{if } l = b \\ \perp & \text{otherwise.} \end{cases}$$

Formally, \mathfrak{t} is defined over the locations $\text{dom}(\mathfrak{t}) = \{l_1, b\}$. Moreover, we have $\text{succ}_{\mathfrak{t}}(l_1) = \{b\}$, $\text{head}_{\mathfrak{t}}(l_1) = \text{even}(l_1, a)$, $\text{calls}_{\mathfrak{t}}(l_1) = \{\text{odd}(b, a)\}$, $\text{heap}(\mathfrak{t}) = \{l_1 \mapsto b, b \mapsto l_2\}$, $\text{heap}_{\mathfrak{t}}(l_1) = \{l_1 \mapsto b\}$, $\text{ptrlocs}(\mathfrak{t}) = \{l_1, b, l_2\}$, $\text{allholes}(\mathfrak{t}) = \{l_2\}$, and $\text{allholepreds}(\mathfrak{t}) = \{\text{even}(l_2, a)\}$.

(2) All of the trees considered in Section 6 are Φ -trees.

We remark that the above definition of Φ -trees does not account for rule instances in which the same predicate call appears multiple times. Similarly, we do not account for multiple predicate calls with the same root parameter. As we will see in Section 8.3, such cases do not need to be considered. We can thus ignore these cases in favor of a simpler formalization.

Our main motivation for considering Φ -trees is that they give a more structured view on models of predicate calls. In particular, every such model corresponds to (at least one) Φ -tree without holes:

LEMMA 7.3. Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state and $\text{pred} \in \text{Preds}(\Phi)$. Then, $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{pred}(z_1, \dots, z_k)$ iff there exists a Φ -tree \mathfrak{t} with $\text{rootpred}(\mathfrak{t}) = \text{pred}(\mathfrak{s}(z_1), \dots, \mathfrak{s}(z_k))$, $\text{allholepreds}(\mathfrak{t}) = \emptyset$, and $\text{heap}(\{\mathfrak{t}\}) = \mathfrak{h}$.

PROOF. The statement directly follows by induction on the number of rules applied to derive $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{pred}(z_1, \dots, z_k)$ resp. the height of the tree \mathfrak{t} . \square

7.1.3 Φ -Forests. We combine zero or more Φ -trees into Φ -forests.

DEFINITION 7.4 (Φ -FOREST). A Φ -forest \mathfrak{f} is a finite set of Φ -trees $\mathfrak{f} = \{t_1, \dots, t_k\}$ with pairwise disjoint locations, i.e., $\text{dom}(t_i) \cap \text{dom}(t_j) = \emptyset$ for $i \neq j$.

We assume that all definitions are lifted from Φ -trees to Φ -forests, i.e., for $\mathfrak{f} = \{t_1, \dots, t_k\}$, we define

- the *induced heap* of \mathfrak{f} as $\text{heap}(\mathfrak{f}) \triangleq \bigcup_{t \in \mathfrak{f}} \text{heap}(t)$; if $l \in \text{dom}(t_i)$ then $\text{rule}_{\mathfrak{f}}(l) = \text{rule}_{t_i}(l)$;
- $\text{graph}(\mathfrak{f}) \triangleq \langle \text{dom}(\mathfrak{f}), \{(x, y) \mid 1 \leq i \leq k, x \in \text{dom}(t_i), y \in \text{succ}_{t_i}(x)\} \rangle$; and
- $\text{dom}(\mathfrak{f}) \triangleq \bigcup_i \text{dom}(t_i)$; $\text{roots}(\mathfrak{f}) \triangleq \{\text{root}(t_i) \mid 1 \leq i \leq k\}$; $\text{allholes}(\mathfrak{f}) \triangleq \bigcup_{1 \leq i \leq k} \text{allholes}(t_i)$.

EXAMPLE 7.5 (Φ -FOREST). Both Example 6.3 and Example 6.6 define a Φ -forest.

7.1.4 *Composing Forests*. As motivated in Section 6.2.3, Φ -forests are composed by (1) taking their disjoint union and (2) optionally merging pairs of trees of the resulting forest by identifying the root of one tree with a hole of another tree.

Disjoint union of forests. The union of two Φ -forests corresponds to ordinary set union, provided no location is in the domain of both forests; otherwise, it is undefined.

DEFINITION 7.6 (UNION OF Φ -FORESTS). Let $\mathfrak{f}_1, \mathfrak{f}_2$ be Φ -forests. The union of $\mathfrak{f}_1, \mathfrak{f}_2$ is given by

$$\mathfrak{f}_1 \uplus \mathfrak{f}_2 \triangleq \begin{cases} \mathfrak{f}_1 \cup \mathfrak{f}_2 & \text{if } \text{dom}(\mathfrak{f}_1) \cap \text{dom}(\mathfrak{f}_2) = \emptyset, \\ \perp, & \text{otherwise.} \end{cases}$$

LEMMA 7.7. Let $\mathfrak{f} = \mathfrak{f}_1 \uplus \mathfrak{f}_2$. Then $\text{heap}(\mathfrak{f}) = \text{heap}(\mathfrak{f}_1) \uplus \text{heap}(\mathfrak{f}_2)$.

PROOF. $\text{heap}(\mathfrak{f}) = \bigcup_{t \in \mathfrak{f}} \text{heap}(t) = (\bigcup_{t \in \mathfrak{f}_1} \text{heap}(t)) \cup (\bigcup_{t \in \mathfrak{f}_2} \text{heap}(t)) = \text{heap}(\mathfrak{f}_1) \uplus \text{heap}(\mathfrak{f}_2)$. (Where we have \uplus rather than \cup because $\mathfrak{f}_1 \uplus \mathfrak{f}_2$ is defined.) \square

Splitting forests. We formalize the process of merging Φ -trees in a roundabout way: we first define a way to *split* the trees of a forest into sub-trees at a fixed set of locations—the inverse of merging forests. This may seem like an arbitrary choice, but will simplify the technical development in follow-up sections. We first consider two examples of splitting before formalizing it in Definition 7.9.

EXAMPLE 7.8 (SPLITTING FORESTS). (1) Let t be the Φ -tree from Example 7.2. The $\{b\}$ -split of $\{t\}$ is given by $\{t_1, t_2\}$, for the trees $t_1 = \{l_1 \mapsto \langle \emptyset, \text{even}(l_1, a) \leftarrow (l_1 \mapsto b) \star \text{odd}(b, a) \rangle\}$ and $t_2 = \{b \mapsto \langle \emptyset, \text{odd}(b, a) \leftarrow (b \mapsto l_2) \star \text{even}(l_2, a) \rangle\}$. $\{t_1, t_2\}$ is the \mathbf{l} -split of $\{t\}$ for all $\mathbf{l} \supseteq \{b\}$: in our definition of \mathbf{l} -split we will not require for the locations in \mathbf{l} to actually occur in the forest.

(2) Recall the forest $\mathfrak{f} = \{t_1, t_2, t_3\}$ from Example 6.3 and the tree t from Example 6.1. Then \mathfrak{f} is the $\{2, 4\}$ -split of $\{t\}$. Likewise, \mathfrak{f} is the $\{1, 2, 4, 7\}$ -split of $\{t\}$, because 1 is the root of a tree and 7 does not occur in the forest. In contrast, \mathfrak{f} is not the $\{1, 2, 5\}$ -split of $\{t\}$, because $5 \in \text{dom}(\mathfrak{f}) \setminus \text{roots}(\mathfrak{f})$.

DEFINITION 7.9 (\mathbf{l} -SPLIT). Let $\mathfrak{f}, \bar{\mathfrak{f}}$ be Φ -forests and $\mathbf{l} \subseteq \text{Loc}$. Then $\bar{\mathfrak{f}}$ is an \mathbf{l} -split of \mathfrak{f} if

- (1) both forests cover the same locations, i.e., $\text{dom}(\mathfrak{f}) = \text{dom}(\bar{\mathfrak{f}})$,
- (2) both forests contain the same rule instances, i.e., $\text{rule}_{\mathfrak{f}}(d) = \text{rule}_{\bar{\mathfrak{f}}}(d)$ for all $d \in \text{dom}(\mathfrak{f})$, and
- (3) the graph of $\bar{\mathfrak{f}}$ is obtained from the graph of \mathfrak{f} by removing edges leading to locations in \mathbf{l} , i.e., $\text{graph}(\bar{\mathfrak{f}}) = \text{graph}(\mathfrak{f}) \setminus \{(a, b) \mid a \in \text{Loc}, b \in \mathbf{l}\}$.

LEMMA 7.10 (UNIQUENESS OF \mathbf{l} -SPLIT). For all $\mathbf{l} \subseteq \text{Loc}$, every Φ -forest has a unique \mathbf{l} -split $\text{split}(\mathfrak{f}, \mathbf{l})$.

PROOF. See Appendix A.10. \square

To formalize how we merge trees, we define a derivation relation \blacktriangleright^* between forests in which we iteratively split trees at suitable locations. Intuitively, $\bar{f}_1 \blacktriangleright^* \bar{f}_2$ holds if splitting the trees in \bar{f}_2 at zero or more locations yields \bar{f}_1 ; or, equivalently, if “merging” zero or more trees of \bar{f}_1 yields \bar{f}_2 .

DEFINITION 7.11 (FOREST DERIVATION). *The forest \bar{f}_2 is one-step derivable from the forest \bar{f}_1 , denoted $\bar{f}_1 \blacktriangleright \bar{f}_2$ iff there exists a location $l \in \text{dom}(\bar{f})$ such that $\bar{f}_1 = \text{split}(\bar{f}_2, \{l\})$.*

The reflexive-transitive closure of \blacktriangleright is denoted by \blacktriangleright^ .*

EXAMPLE 7.12. *Let \perp denote the everywhere undefined partial function. Then, consider the forests $\bar{f} \triangleq \{t_1, t_2\}$ and $\bar{\bar{f}} \triangleq \{\bar{t}\}$ given by the trees below. Then $\bar{f} \blacktriangleright \bar{\bar{f}}$ because $\bar{f} = \text{split}(\bar{\bar{f}}, l_2)$.*

$$\begin{aligned} t_1 &\triangleq \{l_1 \mapsto \langle \perp, \text{odd}(l_1, l_4) \Leftarrow (l_1 \mapsto l_2) \star \text{even}(l_2, l_4) \rangle\} \\ t_2 &\triangleq \{l_2 \mapsto \langle l_3, \text{even}(l_2, l_4) \Leftarrow (l_2 \mapsto l_3) \star \text{odd}(l_3, l_4) \rangle, \\ &\quad l_3 \mapsto \langle \perp, \text{odd}(l_3, l_4) \Leftarrow (l_3 \mapsto l_4) \rangle\} \\ \bar{t} &\triangleq \{l_1 \mapsto \langle l_2, \text{odd}(l_1, l_4) \Leftarrow (l_1 \mapsto l_2) \star \text{even}(l_2, l_4) \rangle, \\ &\quad l_2 \mapsto \langle l_3, \text{even}(l_2, l_4) \Leftarrow (l_2 \mapsto l_3) \star \text{odd}(l_3, l_4) \rangle, \\ &\quad l_3 \mapsto \langle \perp, \text{odd}(l_3, l_4) \Leftarrow (l_3 \mapsto l_4) \rangle\}. \end{aligned}$$

We note that multiple steps of \blacktriangleright correspond to splitting at multiple locations, because

$$\text{split}(\bar{f}, \{l_1, \dots, l_k\}) = \text{split}(\dots \text{split}(\text{split}(\bar{f}, \{l_1\}), \{l_2\}), \dots, \{l_k\}).$$

LEMMA 7.13. $\bar{f}_1 \blacktriangleright^* \bar{f}_2$ iff there exists a set of locations \mathbf{l} with $\bar{f}_1 = \text{split}(\bar{f}_2, \mathbf{l})$.

Moreover, forests in the \blacktriangleright^* relation describe the same states:

LEMMA 7.14. *Let \bar{f} be a Φ -forest and $\bar{\bar{f}} \blacktriangleright^* \bar{f}$. Then $\text{heap}(\bar{\bar{f}}) = \text{heap}(\bar{f})$.*

PROOF. Since $\bar{\bar{f}} \blacktriangleright^* \bar{f}$, there exists—by Lemma 7.13—a set of locations \mathbf{l} with $\bar{\bar{f}} = \text{split}(\bar{f}, \mathbf{l})$. By definition of \mathbf{l} -splits, we have (1) $\text{dom}(\bar{\bar{f}}) = \text{dom}(\bar{f})$ and (2) $\text{rule}_{\bar{f}}(l) = \text{rule}_{\bar{\bar{f}}}(l)$ for every location $l \in \text{dom}(\bar{\bar{f}})$. Consequently, $\text{heap}(\bar{\bar{f}}) = \text{heap}(\bar{f})$. \square

Based on the \blacktriangleright^* , we define the *composition* operation on pairs of forests as motivated in Section 6.2:

DEFINITION 7.15 (FOREST COMPOSITION). *The composition of \bar{f}_1 and \bar{f}_2 is $\bar{f}_1 \bullet_{\mathbf{F}} \bar{f}_2 \triangleq \{\bar{f} \mid \bar{f}_1 \cup \bar{f}_2 \blacktriangleright^* \bar{f}\}$.*

7.2 Forest Projections

In Section 6.3, we informally presented the *projection* of Φ -forests onto GSL formulas, and discussed the need for using *guarded quantifiers*. As a reminder, we repeat here the informal definition of the projection: Given a stack \mathfrak{s} and a Φ -forest $\bar{t} = \{t_1, \dots, t_k\}$,

- (1) we compute the formula $\phi \triangleq \star_{1 \leq i \leq k} (\star \text{allholepreds}(t_i)) \rightarrow \text{rootpred}(t_i)$, in which all parameters of all predicate calls are locations;
- (2) we replace in ϕ every location $v \in \text{img}(\mathfrak{s})$ by an arbitrary but fixed variable x with $\mathfrak{s}(x) = v$ holds;
- (3) we replace every location $v \in \text{dom}(\bar{t}) \setminus \text{img}(\mathfrak{s})$ by a *guarded existential*;
- (4) we replace every other location by a *guarded universal*.

(We point out that any occurrence of *nil* in ϕ would not be replaced during step (2)-(4) by the projection operation because it is not a location.) Now we make the above definitions precise. First, we introduce the *projection* of trees and forests (Section 7.2.1). Then, we state the definition of *guarded quantifiers* (in Section 7.2.2); Finally, we introduce the *stack-projection* (in Section 7.2.3).

7.2.1 Tree and Forest Projections. We are now ready to define the forest projection outlined in Section 6.3. We begin with defining the projection of a tree:

DEFINITION 7.16 (PROJECTION OF A TREE). *The projection $\text{project}^{\text{Loc}}(\mathfrak{t})$ of a Φ -tree \mathfrak{t} is given by*

$$\text{project}^{\text{Loc}}(\mathfrak{t}) \triangleq (\star \text{allholepreds}(\mathfrak{t})) \star \text{rootpred}(\mathfrak{t}).$$

EXAMPLE 7.17. *Recall from Example 7.2 the Φ -tree \mathfrak{t} over an SID describing lists of even and odd length. This tree admits the tree projection $\text{project}^{\text{Loc}}(\mathfrak{t}) = \text{even}(l_2, a) \star \text{even}(l_1, a)$.*

Tree projections are sound in the sense that the induced heap of a tree satisfies its tree projection. To prove this result, we need the following variant of modus ponens (cf. [Reynolds 2002]):

LEMMA 7.18 (GENERALIZED MODUS PONENS).

$$((\text{pred}_2(x_2) \star \psi) \star \text{pred}_1(x_1)) \star (\psi' \star \text{pred}_2(x_2)) \text{ implies } (\psi \star \psi') \star \text{pred}_1(x_1).$$

LEMMA 7.19 (SOUNDNESS OF TREE PROJECTIONS). *Let \mathfrak{t} be a Φ -tree. Then, we have $\langle _, \text{heap}(\mathfrak{t}) \rangle \models_{\Phi} \text{project}^{\text{Loc}}(\mathfrak{t})$ (where $_$ denotes an arbitrary stack).*

PROOF. By mathematical induction on the height of \mathfrak{t} ; see Appendix A.12 for details. \square

7.2.2 Guarded Quantifiers. As motivated in Section 6.3.3, we introduce *guarded* quantifiers, which we denote by \exists and \forall , respectively. Specifically, we consider formulas $\exists e. (\forall a. (\phi_{\text{qf}} \star \dots \star \phi_{\text{qf}}))$, where ϕ_{qf} denotes *quantifier-free* SL formulas (cf. Section 3.1). We collect all formulas of the above form in the set $\text{SL}_{\text{btw}}^{\exists\forall}$. Our guarded quantifiers have the following semantics:

- $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \exists \langle e_1, \dots, e_k \rangle . \phi$ iff there exist pairwise different locations $v_1, \dots, v_k \in \text{dom}(\mathfrak{h}) \setminus \text{img}(\mathfrak{s})$ such that $\langle \mathfrak{s} \cup \{e_1 \mapsto v_1, \dots, e_k \mapsto v_k\}, \mathfrak{h} \rangle \models_{\Phi} \phi$.
- $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \forall \langle a_1, \dots, a_k \rangle . \phi$ iff for all pairwise different locations $v_1, \dots, v_k \in \text{Loc} \setminus (\text{dom}(\mathfrak{h}) \cup \text{img}(\mathfrak{s}))$, we have $\langle \mathfrak{s} \cup \{a_1 \mapsto v_1, \dots, a_k \mapsto v_k\}, \mathfrak{h} \rangle \models_{\Phi} \phi$.

Notice that our guarded quantifiers differ from the standard ones in three aspects: first, guarded quantifiers cannot be instantiated with locations that are already known, i.e., not with any location that is already in the stack. Second, we require that the quantified locations are pairwise different. Third, our quantifiers are *not* dual, i.e., $\exists x. \phi$ is *not* equivalent to $\neg \forall x. \neg \phi$. For guarded states, location terms in a formula can be replaced by a guarded universal quantifier as long as they do not appear in the state in question:

LEMMA 7.20. *Let $\langle \mathfrak{s}, \mathfrak{h} \rangle \in \text{GStates}$ and ϕ be a quantifier free SL formula with $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$. Moreover, let $\mathbf{v} \in (\text{Loc} \setminus (\text{dom}(\mathfrak{h}) \cup \text{img}(\mathfrak{s})))^*$ be a repetition-free sequence of locations. Then, for every set $\mathbf{a} \triangleq \{a_1, \dots, a_{|\mathbf{v}|}\}$ of fresh variables (i.e., $\mathbf{a} \cap \text{dom}(\mathfrak{s}) = \emptyset$), we have $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \forall \mathbf{a}. \phi[\mathbf{v}/\mathbf{a}]$.*

PROOF. See Appendix A.11. \square

Many standard equivalences of separation logic continue to hold for formulas with guarded quantifiers; we list corresponding *rewriting rules* in Fig. 10. These rules establish the *rewriting equivalence* \equiv , which preserves logical equivalence—we will only consider formulas up to \equiv .

LEMMA 7.21 (SOUNDNESS OF REWRITING EQUIVALENCE). *If $\phi_1 \equiv \phi_2$ then $\phi_1 \models_{\Phi} \phi_2$.*

7.2.3 Stack-Projection. We now abstract from locations in projections (cf. Section 6.3.1), replacing every location l in the projection of a forest \mathfrak{f} by a variable: a stack variable, if l is in the image of the stack, an existentially-quantified one if $l \in \text{dom}(\mathfrak{f})$, and a universally-quantified one otherwise.

$$\begin{array}{c}
\frac{\phi_1 \equiv \phi_2}{\phi_1 \star \psi \equiv \phi_2 \star \psi} \text{ (mono)} \quad \frac{}{\phi_1 \star \mathbf{emp} \equiv \phi_1} \text{ (emp)} \quad \frac{\phi_1 \equiv \phi_2}{\phi_2 \star \psi \equiv \phi_1 \star \psi} \text{ (anti)} \quad \frac{}{\phi_1 \equiv \phi_1} \text{ (id)} \\
\\
\frac{}{\phi_1 \star (\phi_2 \star \phi_3) \equiv (\phi_1 \star \phi_2) \star \phi_3} \text{ (assoc)} \quad \frac{}{\phi_1 \star \phi_2 \equiv \phi_2 \star \phi_1} \text{ (comm)} \quad \frac{\phi_1 \equiv \phi_2}{\phi_2 \equiv \phi_1} \text{ (sym)} \\
\\
\frac{\mathbb{Q} \in \{\forall, \exists\} \quad z \notin \text{vars}(\phi)}{\mathbb{Q}y. \phi \equiv \mathbb{Q}z. \phi[y/z]} \text{ (ren)} \quad \frac{\phi_1 \equiv \phi_2}{\exists y. \phi_1 \equiv \exists y. \phi_2} \text{ (\exists-intro)} \quad \frac{\phi_1 \equiv \phi_2}{\forall y. \phi_1 \equiv \forall y. \phi_2} \text{ (\forall-intro)} \\
\\
\frac{\phi_1 \equiv \phi_3 \quad \phi_3 \equiv \phi_2}{\phi_1 \equiv \phi_2} \text{ (trans)} \quad \frac{\mathbb{Q} \in \{\forall, \exists\} \quad z \notin \text{vars}(\phi)}{\mathbb{Q}z. \phi \equiv \phi} \text{ (drop)}
\end{array}$$

Fig. 10. A set of rules for rewriting $\text{SL}_{\text{btw}}^{\exists\forall}$ formulas into equivalent formulas.

Aliasing and Variable Order. In case of aliasing, i.e., if there are multiple variables that are mapped to the same location l , there are multiple choices for replacing l by a stack variable x with $\mathfrak{s}(x) = l$. This has the consequence that the projection would not be unique. In order to avoid this problem, we assume an arbitrary, but fixed, linear ordering of the variables Var . We then choose the variable among all the aliases of a variable that is maximal according to this variable ordering. Formally:

DEFINITION 7.22 (STACK-CHOICE FUNCTION $\mathfrak{s}_{\text{max}}^{-1}$). *Let \mathfrak{s} be a stack. Then, the stack-choice function of \mathfrak{s} maps a location $l \in \text{img}(\mathfrak{s})$ to $\mathfrak{s}_{\text{max}}^{-1}(l) = \max\{x \in \text{dom}(\mathfrak{s}) \mid \mathfrak{s}(x) = l\}$.*

Quantified Variables. We (mostly) maintain the convention that we denote (guarded) universally resp. existentially quantified variables by a_1, a_2, \dots resp. e_1, e_2, \dots . We will always assume that $\{a_1, a_2, \dots\} \cap \{e_1, e_2, \dots\} = \emptyset$ and that $\text{dom}(\mathfrak{s}) \cap (\{a_1, a_2, \dots\} \cup \{e_1, e_2, \dots\}) = \emptyset$ for any stack \mathfrak{s} . We are now ready to give the main definition of this subsection:

DEFINITION 7.23 (STACK-PROJECTION). *Let $\mathfrak{f} = \{t_1, \dots, t_k\}$ be a Φ -forest, \mathfrak{s} be a stack, and*

- *let $\phi = \star_{1 \leq i \leq k} \text{project}^{\text{Loc}}(t_i)$ be the projection of trees of \mathfrak{f} conjoined by \star ,*
- *let $\mathbf{w} = \text{locs}(\phi) \cap (\text{dom}(\mathfrak{f}) \setminus \text{img}(\mathfrak{s}))$ be some (arbitrarily ordered) sequence of locations that occur in the formula ϕ and are allocated in $\text{heap}(\mathfrak{f})$ but are not the value of any stack variable,*
- *and let $\mathbf{v} = \text{locs}(\phi) \setminus (\text{img}(\mathfrak{s}) \cup \text{dom}(\mathfrak{f}))$ be some (arbitrarily ordered) sequence of locations that occur in the formula ϕ and are neither allocated nor the value of any stack variable.*

Then, we define the stack-projection of \mathfrak{s} and \mathfrak{f} as

$$\text{project}(\mathfrak{s}, \mathfrak{f}) \triangleq \exists \mathbf{e}. \forall \mathbf{a}. \phi[\text{dom}(\mathfrak{s}_{\text{max}}^{-1}) \cdot \mathbf{v} \cdot \mathbf{w} / \text{img}(\mathfrak{s}_{\text{max}}^{-1}) \cdot \mathbf{a} \cdot \mathbf{e}],$$

where $\mathbf{e} \triangleq \langle e_1, e_2, \dots, e_{|\mathbf{w}|} \rangle$ and $\mathbf{a} \triangleq \langle a_1, a_2, \dots, a_{|\mathbf{v}|} \rangle$ denote disjoint sets of fresh variables.

The stack-projection is well-defined because $\text{dom}(\mathfrak{s}_{\text{max}}^{-1})$, \mathbf{w} and \mathbf{v} form a partitioning of $\text{locs}(\phi)$. Notice that the stack-projection is *unique* (w.r.t. the rewriting equivalence \equiv defined in Fig. 10): while the stack-projection involves picking an (arbitrary) order on the trees t_1, \dots, t_k and a choice of the fresh variables \mathbf{e} and \mathbf{a} , this does not matter because of the commutativity and associativity of \star and the possibility to rename quantified variables, which is allowed for by the rules of the rewriting equivalence \equiv .

EXAMPLE 7.24 (STACK-PROJECTION). *We consider three examples of stack-projections:*

- (1) *Let t be the Φ -tree from Example 7.2. Then, for $\mathfrak{f} = \{t\}$ and $\mathfrak{s} = \{x_1 \mapsto l_1, x_2 \mapsto l_2\}$, we have*

$$\text{heap}(\mathfrak{f}) = \{l_1 \mapsto b, b \mapsto l_2\} \quad \text{and} \quad \text{project}^{\text{Loc}}(\mathfrak{f}) = \forall a_1. \text{even}(l_2, a_1) \star \text{odd}(l_1, a_1).$$

As all locations in this formula are in the image of the stack, we have

$$\text{project}(\mathfrak{s}, \mathfrak{f}) = \text{project}^{\text{Loc}}(\mathfrak{f})[\text{dom}(\mathfrak{s}_{\text{max}}^{-1})/\text{img}(\mathfrak{s}_{\text{max}}^{-1})] = \forall a_1. \text{even}(x_2, a_1) \star \text{odd}(x_1, a_1).$$

(2) Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be the model and let $\mathfrak{t}_1, \mathfrak{t}_2, \mathfrak{t}_3$ be the Φ -trees from Example 6.7. Then,

$$\begin{aligned} \text{project}(\mathfrak{s}, \{\mathfrak{t}_1, \mathfrak{t}_3\}) &= \exists! a. (\text{lseg}(y, a) \star \text{cyclic}(x, y, z)) \star \text{lseg}(z, a), \text{ and} \\ \text{project}(\mathfrak{s}, \{\mathfrak{t}_2\}) &= \forall a'. \text{lseg}(z, a') \star \text{lseg}(y, a'). \end{aligned}$$

(3) Let $\mathfrak{t}_1, \mathfrak{t}_2, \mathfrak{t}_3, \mathfrak{t}_4$ be the Φ -trees from Example 6.3 for the state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ of Example 6.1. Then,

$$\begin{aligned} \text{project}(\mathfrak{s}, \{\mathfrak{t}_1, \mathfrak{t}_3, \mathfrak{t}_4\}) &= \exists! r. (\text{tll}(a, b, c) \star \text{tll}(x, y, z)) \star (b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star (c \mapsto \langle \text{nil}, \text{nil}, r \rangle), \text{ and} \\ \text{project}(\mathfrak{s}, \{\mathfrak{t}_2\}) &= \forall r'. ((b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star (c \mapsto \langle \text{nil}, \text{nil}, r' \rangle)) \star \text{tll}(a, b, c). \end{aligned}$$

In each of the above examples, we observe that $\langle \mathfrak{s}, \text{heap}(\mathfrak{f}) \rangle \models_{\Phi} \text{project}(\mathfrak{s}, \mathfrak{f})$. This is not a coincidence as eliminating locations preserves the soundness of forest projections:

LEMMA 7.25 (SOUNDNESS OF STACK-PROJECTION). *Let $\langle \mathfrak{s}, \mathfrak{h} \rangle \in \text{GStates}$. Moreover, let \mathfrak{f} be a Φ -forest with $\text{heap}(\mathfrak{f}) = \mathfrak{h}$. Then, we have $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{project}(\mathfrak{s}, \mathfrak{f})$.*

PROOF. See Appendix A.13. □

EXAMPLE 7.26 (WHY WE NEED GUARDED QUANTIFIERS). *We now have the machinery available to discuss why guarded quantifiers are needed. To this end, let us revisit the motivating example in Section 6.3.3: We consider the state $\langle \mathfrak{s}, \mathfrak{h} \rangle$, given by the stack $\mathfrak{s} = \{x \mapsto 1\}$ and the heap $\mathfrak{h} = \{1 \mapsto 2, 2 \mapsto \text{nil}\}$, and the SID Φ given by the following predicates:*

$$p(x, a, b) \Leftarrow \exists y. (x \mapsto y) \star q(y, a) \star x \neq a \star a \neq b \quad q(y, a) \Leftarrow (y \mapsto \text{nil}) \star y \neq a$$

We further consider the unfolding tree \mathfrak{t} consisting of the rule instance $p(1, 4, 5) \Leftarrow (1 \mapsto 2) \star q(2, 4) \star 1 \neq 4 \star 4 \neq 5$ at the root with a single child for the rule instance $q(2, 4) \Leftarrow (2 \mapsto \text{nil}) \star 2 \neq 4$. Hence, we have $\text{rootpred}(\mathfrak{t}) = p(1, 4, 5)$, $\text{allholepreds}(\mathfrak{t}) = \emptyset$, and $\text{locs}(p(1, 4, 5)) \setminus (\text{dom}(\mathfrak{h}) \cup \text{img}(\mathfrak{s})) = \{4, 5\}$. By Definition 7.16, we then obtain the tree projection

$$\begin{aligned} \text{project}(\mathfrak{s}, \mathfrak{t}) &= \forall a, b. ((\star \text{allholepreds}(p(1, a, b))) \star \text{rootpred}(\mathfrak{t}))[\langle 1, 4, 5 \rangle / \langle x, a, b \rangle] \\ &= \forall a, b. \text{emp} \star p(x, 4, 5)[\langle 4, 5 \rangle / \langle a, b \rangle] = \forall a, b. p(x, a, b). \end{aligned}$$

By Lemma 7.19, we have $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \forall a, b. p(x, a, b)$. In particular, the semantics of \forall guarantees that a and b refer to distinct locations that are neither allocated and nor the value of any stack variable. This is crucial to ensure soundness of the stack-projection: if we would use a standard universal quantifier instead of a guarded one, $\langle \mathfrak{s}, \mathfrak{h} \rangle$ would not be a model of $\text{project}^{\text{Loc}}(\mathfrak{t})$ as neither $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} p(x, 1, 5)$, $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} p(x, 2, 5)$ nor $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} p(x, 5, 5)$ hold.

7.3 Composing Projections

7.3.1 Motivation. Recall from Section 6.1.3 that our goal is the definition of a *composition* operator for the projections of forests. This operation should collect exactly those projections of forests $\mathfrak{f} \in \mathfrak{f}_1 \bullet_{\mathbf{F}} \mathfrak{f}_2$ (see Definition 7.15) that can be derived from $\text{project}(\mathfrak{s}, \mathfrak{f}_1)$ and $\text{project}(\mathfrak{s}, \mathfrak{f}_2)$, i.e.,

$$\text{project}(\mathfrak{s}, \mathfrak{f}_1) \bullet_{\mathbf{P}} \text{project}(\mathfrak{s}, \mathfrak{f}_2) \stackrel{?}{=} \{\text{project}(\mathfrak{s}, \mathfrak{f}) \mid \mathfrak{f} \in \mathfrak{f}_1 \bullet_{\mathbf{F}} \mathfrak{f}_2\}.$$

Put differently, we are looking for an operation $\bullet_{\mathbf{P}}$ such that $\text{project}(\mathfrak{s}, \cdot)$ is a homomorphism from the set of Φ -forests and $\bullet_{\mathbf{F}}$ to the set of projections and $\bullet_{\mathbf{P}}$.

How can we define such an operation $\bullet_{\mathbf{P}}$? Intuitively, we need to conjoin the projections via \star in order to simulate the operation $\mathfrak{f}_1 \uplus \mathfrak{f}_2$, and apply the generalized modus ponens rule (see Lemma 7.18) in order to simulate

the operation \blacktriangleright on trees. There is, however, one complication: our forest projections contain quantifiers. In particular, $\text{project}(s, \bar{f}_1) \star \text{project}(s, \bar{f}_2)$ is of the form $(\exists e_1. \forall a_1. \phi_1) \star (\exists e_2. \forall a_2. \phi_2)$, whereas $\text{project}(s, \bar{f}_1 \uplus \bar{f}_2)$ is of the form $\exists e. \forall a. \phi$, where ϕ_1 , ϕ_2 , and ϕ do not contain guarded quantifiers. In other words, \bullet_P has to push the guarded quantifiers to the front before the modus ponens rule can be applied.

7.3.2 Definition of the Composition Operation. We will define \bullet_P in terms of two operations: (1) An operator $\bar{\star}$ that captures all sound ways to move the guarded quantifiers to the front of the formula $\text{project}(s, \bar{f}_1) \star \text{project}(s, \bar{f}_2)$ (i.e., “re-scopes the guarded quantifiers”). (2) A derivation operator \triangleright that rewrites formulas based on the generalized modus ponens rule (Lemma 7.18).

DEFINITION 7.27 (RE-SCOPING). We say χ is a re-scoping of $\exists e_1. \forall a_1. \phi_1$ and $\exists e_2. \forall a_2. \phi_2$, in signs $\chi \in (\exists e_1. \forall a_1. \phi_1) \bar{\star} (\exists e_2. \forall a_2. \phi_2)$, if there are repetition-free sequences of variables \mathbf{a} , \mathbf{d}_i and $\mathbf{u}_i \subseteq \mathbf{a} \cup \mathbf{d}_{3-i}$, for $i = 1, 2$, such that (1) \mathbf{a} , \mathbf{d}_1 and \mathbf{d}_2 are pairwise disjoint, and

$$(2) \chi \equiv \exists \mathbf{d}_1 \cdot \mathbf{d}_2. \forall \mathbf{a}. \phi_1[\mathbf{e}_1 \cdot \mathbf{a}_1 / \mathbf{d}_1 \cdot \mathbf{u}_1] \star \phi_2[\mathbf{e}_2 \cdot \mathbf{a}_2 / \mathbf{d}_2 \cdot \mathbf{u}_2].$$

The re-scoping operation is sound with regard to the semantics of separation logic:

LEMMA 7.28 (SOUNDNESS OF RE-SCOPING). Let $\exists e_1. \forall a_1. \phi_1$ and $\exists e_2. \forall a_2. \phi_2$ be some formulas whose predicates are defined by some SID Φ . Then,

$$\chi \in (\exists e_1. \forall a_1. \phi_1) \bar{\star} (\exists e_2. \forall a_2. \phi_2) \text{ implies } (\exists e_1. \forall a_1. \phi_1) \star (\exists e_2. \forall a_2. \phi_2) \models_{\Phi} \chi.$$

PROOF. Follows directly from the semantics of the guarded quantifiers \exists and \forall . \square

DEFINITION 7.29 (DERIVABILITY). We say χ can be derived from $\exists e. \forall a. \phi$, in signs $\exists e. \forall a. \phi \triangleright \chi$, if χ can be obtained from $\exists e. \forall a. \phi$ by applying Lemma 7.18 and the rewriting equivalence \equiv (see Fig. 10), formally, if there are predicates $\text{pred}_1(x_1)$, $\text{pred}_2(x_2)$, and formulas ψ, ψ', ζ such that

- (1) $\phi \equiv (\text{pred}_2(x_2) \star \psi) \star \text{pred}_1(x_1) \star (\psi' \star \text{pred}_2(x_2)) \star \zeta$, and
- (2) $\chi \equiv \exists e. \forall a. (\psi \star \psi') \star \text{pred}_1(x_1) \star \zeta$.

The derivability relation is sound with regard to the semantics of separation logic:

LEMMA 7.30 (SOUNDNESS OF DERIVABILITY). Let $\exists e. \forall a. \phi$ be some formula whose predicates are defined by some SID Φ . Then, $\exists e. \forall a. \phi \triangleright \chi$ implies $\exists e. \forall a. \phi \models_{\Phi} \chi$.

PROOF. Follows directly from the soundness of the generalized modus ponens rule (see Lemma 7.18) and the soundness of the rewriting equivalence \equiv . \square

We now define composition based on the re-scoping and derivation operations:

DEFINITION 7.31 (COMPOSITION OPERATION). We define the composition of ϕ_1 and ϕ_2 by

$$\phi_1 \bullet_P \phi_2 \triangleq \{\phi \mid \zeta \triangleright^* \phi \text{ for some } \zeta \in \phi_1 \bar{\star} \phi_2\}.$$

COROLLARY 7.32 (SOUNDNESS OF \bullet_P). $\phi \in \phi_1 \bullet_P \phi_2$ implies $\phi_1 \star \phi_2 \models_{\Phi} \phi$.

PROOF. Follows immediately from Lemmas 7.28 and 7.30. \square

EXAMPLE 7.33. • For $\phi_1 = \text{ls}(x_2, x_3) \star \text{ls}(x_1, x_3)$ and $\phi_2 = \text{emp} \star \text{ls}(x_2, x_3)$, it holds that $\phi_1 \star \phi_2 \triangleright \text{emp} \star \text{ls}(x_1, x_3)$. Hence, $(\text{emp} \star \text{ls}(x_1, x_3)) \in \phi_1 \bullet_P \phi_2$.

- For $\phi_1 = \forall a. \text{ls}(x_2, a) \star \text{ls}(x_1, a)$ and $\phi_2 = \forall b. \text{ls}(x_3, b) \star \text{ls}(x_2, b)$, we have $\forall c. (\text{ls}(x_2, c) \star \text{ls}(x_1, c)) \star (\text{ls}(x_3, c) \star \text{ls}(x_2, c)) \in \phi_1 \bar{\star} \phi_2$. With $\forall c. (\text{ls}(x_2, c) \star \text{ls}(x_1, c)) \star (\text{ls}(x_3, c) \star \text{ls}(x_2, c)) \triangleright \forall c. (\text{ls}(x_3, c) \star \text{ls}(x_1, c))$, we have $\forall c. (\text{ls}(x_3, c) \star \text{ls}(x_1, c)) \in \phi_1 \bullet_P \phi_2$.

Let us also revisit our informal exposition in Example 6.8 and make it precise:

EXAMPLE 7.34 (COMPOSITION OPERATION ON PROJECTIONS).

- Let t_1, t_2, t_3 be the Φ -trees from Example 6.7. We set $\mathfrak{f}_1 = \{t_1, t_3\}$ and $\mathfrak{f}_2 = \{t_2\}$. We then have:

$$\text{project}(\mathfrak{s}, \mathfrak{f}_1) = \exists a. (\text{lseg}(y, a) \rightarrow \text{cyclic}(x, y, z)) \star \text{lseg}(z, a), \text{ and}$$

$$\text{project}(\mathfrak{s}, \mathfrak{f}_2) = \forall a'. \text{lseg}(z, a') \rightarrow \text{lseg}(y, a'). \text{ Then,}$$

$$\exists a. (\text{lseg}(y, a) \rightarrow \text{cyclic}(x, y, z)) \star \text{lseg}(z, a) \star (\text{lseg}(z, a) \rightarrow \text{lseg}(y, a))$$

$$\in \text{project}(\mathfrak{s}, \mathfrak{f}_1) \star \text{project}(\mathfrak{s}, \mathfrak{f}_2) \text{ and further}$$

$$\exists a. (\text{lseg}(y, a) \rightarrow \text{cyclic}(x, y, z)) \star \text{lseg}(z, a) \star (\text{lseg}(z, a) \rightarrow \text{lseg}(y, a)) \triangleright^* \text{cyclic}(x, y, z).$$

Hence, we have $\text{cyclic}(x, y, z) \in \text{project}(\mathfrak{s}, \mathfrak{f}_1) \bullet_P \text{project}(\mathfrak{s}, \mathfrak{f}_2)$.

- Let t_1, t_2, t_3, t_4 be the Φ -trees from Example 6.3. We set $\mathfrak{f}_1 = \{t_1, t_3, t_4\}$ and $\mathfrak{f}_2 = \{t_2\}$. We have

$$\text{project}(\mathfrak{s}, \mathfrak{f}_1) = \exists r. (\text{tll}(a, b, c) \rightarrow \text{tll}(x, y, z)) \star (b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star (c \mapsto \langle \text{nil}, \text{nil}, r \rangle), \text{ and}$$

$$\text{project}(\mathfrak{s}, \mathfrak{f}_2) = \forall r'. ((b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star (c \mapsto \langle \text{nil}, \text{nil}, r' \rangle)) \rightarrow \text{tll}(a, b, c). \text{ Then,}$$

$$\exists r. (\text{tll}(a, b, c) \rightarrow \text{tll}(x, y, z)) \star (b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star (c \mapsto \langle \text{nil}, \text{nil}, r \rangle) \star$$

$$(((b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star (c \mapsto \langle \text{nil}, \text{nil}, r \rangle)) \rightarrow \text{tll}(a, b, c)) \in \text{project}(\mathfrak{s}, \mathfrak{f}_1) \star \text{project}(\mathfrak{s}, \mathfrak{f}_2). \text{ Further,}$$

$$\exists r. (\text{tll}(a, b, c) \rightarrow \text{tll}(x, y, z)) \star (b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star (c \mapsto \langle \text{nil}, \text{nil}, r \rangle) \star$$

$$(((b \mapsto \langle \text{nil}, \text{nil}, c \rangle) \star (c \mapsto \langle \text{nil}, \text{nil}, r \rangle)) \rightarrow \text{tll}(a, b, c)) \triangleright^* \text{tll}(x, y, z).$$

Hence, we have $\text{tll}(x, y, z) \in \text{project}(\mathfrak{s}, \mathfrak{f}_1) \bullet_P \text{project}(\mathfrak{s}, \mathfrak{f}_2)$.

7.3.3 *Relating the Composition of Forests and of Projections.* Recall from Section 7.3.1 our design goal that the projection function $\text{project}(\mathfrak{s}, \cdot)$ should be a homomorphism from forests and forest composition \bullet_F (Definition 7.15) to projections and projection composition \bullet_P (Definition 7.31), i.e.,

$$\text{project}(\mathfrak{s}, \mathfrak{f}_1) \bullet_P \text{project}(\mathfrak{s}, \mathfrak{f}_2) \stackrel{?}{=} \{\text{project}(\mathfrak{s}, \mathfrak{f}) \mid \mathfrak{f} \in \mathfrak{f}_1 \bullet_F \mathfrak{f}_2\}.$$

Indeed, in one direction our composition operation achieves this:

LEMMA 7.35. *Let \mathfrak{s} be a stack and let $\mathfrak{f}_1, \mathfrak{f}_2$ be Φ -forests such that $\mathfrak{f}_1 \uplus \mathfrak{f}_2 \neq \perp$. Then,*

$$\mathfrak{f} \in \mathfrak{f}_1 \bullet_F \mathfrak{f}_2 \text{ implies } \text{project}(\mathfrak{s}, \mathfrak{f}) \in \text{project}(\mathfrak{s}, \mathfrak{f}_1) \bullet_P \text{project}(\mathfrak{s}, \mathfrak{f}_2).$$

PROOF. See Appendix A.14. □

Unfortunately, as demonstrated below, the homomorphism breaks in the other direction:

EXAMPLE 7.36 (PROJECTION IS NOT HOMOMORPHIC). *Consider the Φ -forests $\mathfrak{f}_1 = \{t_1\}$ and $\mathfrak{f}_2 = \{t_2\}$ and the stack $\mathfrak{s} \triangleq \{x_1 \mapsto l_1, x_2 \mapsto l_2, x_3 \mapsto l_3\}$, where*

$$t_1 = \{l_1 \mapsto \langle \emptyset, (\text{odd}(l_1, m_1) \leftarrow (l_1 \mapsto l_2) \star \text{even}(l_2, m_1)) \rangle\}, \text{ and}$$

$$t_2 = \{l_2 \mapsto \langle \emptyset, (\text{even}(l_2, m_2) \leftarrow (l_2 \mapsto l_3) \star \text{odd}(l_3, m_2)) \rangle\}.$$

The corresponding projections are

$$\text{project}(\mathfrak{s}, \mathfrak{f}_1) = \forall a. \text{even}(x_2, a) \rightarrow \text{odd}(x_1, a) \text{ and } \text{project}(\mathfrak{s}, \mathfrak{f}_2) = \forall a. \text{odd}(x_3, a) \rightarrow \text{even}(x_2, a).$$

Moreover, we have $\forall a. \text{odd}(x_3, a) \rightarrow \text{odd}(x_1, a) \in \text{project}(\mathfrak{s}, \mathfrak{f}_1) \bullet_P \text{project}(\mathfrak{s}, \mathfrak{f}_2)$.

However, since different locations, namely m_1 and m_2 , are unused in the two forests, there is only one forest in $\mathfrak{f}_1 \bullet_F \mathfrak{f}_2: \{t_1, t_2\}$. It is not possible to merge the trees, because the hole predicate of the first tree, $\text{even}(l_2, m_1)$, is different from the root of the second tree, $\text{even}(l_2, m_2)$. In particular, there does not exist a forest \mathfrak{f} with $\mathfrak{f} \in \mathfrak{f}_1 \bullet_F \mathfrak{f}_2$ and $\text{project}(\mathfrak{s}, \mathfrak{f}) \equiv \forall a. \text{odd}(x_3, a) \rightarrow \text{odd}(x_1, a)$.

The essence of Example 7.36 is that while \bullet_P allows renaming quantified universals, \bullet_F does not allow renaming locations, breaking the homomorphism. To get a correspondence between the two notions of composition, we therefore allow renaming all locations that do *not* occur as the value of any stack variable. We capture this in the notion of \mathfrak{s} -equivalence:

DEFINITION 7.37 (\mathfrak{s} -EQUIVALENCE). *Two Φ -forests $\mathfrak{f}_1, \mathfrak{f}_2$ are \mathfrak{s} -equivalent, denoted $\mathfrak{f}_1 \equiv_{\mathfrak{s}} \mathfrak{f}_2$, iff there is a bijective function $\sigma: \text{Loc} \rightarrow \text{Loc}$ such that $\sigma(l) = l$ for all $l \in \text{img}(\mathfrak{s})$, and $\sigma(\mathfrak{f}_1) = \mathfrak{f}_2$, where*

- $\sigma(\{t_1, \dots, t_k\}) \triangleq \{\sigma(t_1), \dots, \sigma(t_k)\}$,
- $\sigma(t) \triangleq \{\sigma(l) \mapsto \langle \sigma(\text{succ}_{t_1}(l)), \text{rule}_{t_1}(l)[\text{dom}(\sigma)/\text{img}(\sigma)] \mid l \in \text{dom}(t) \rangle\}$, and
- $(\text{pred}(l) \Leftarrow \phi)[v/w] \triangleq \text{pred}(l[v/w]) \Leftarrow \phi[v/w]$ for sequences of locations v and w .

Note that $\mathfrak{f}_1 \equiv_{\mathfrak{s}} \mathfrak{f}_2$ implies that $\langle \mathfrak{s}, \text{heap}(\mathfrak{f}_1) \rangle$ and $\langle \mathfrak{s}, \text{heap}(\mathfrak{f}_2) \rangle$ are isomorphic. In fact,

LEMMA 7.38. *If \mathfrak{f}_1 and \mathfrak{f}_2 are Φ -forests with $\mathfrak{f}_1 \equiv_{\mathfrak{s}} \mathfrak{f}_2$, then $\text{project}(\mathfrak{s}, \mathfrak{f}_1) \equiv \text{project}(\mathfrak{s}, \mathfrak{f}_2)$.*

PROOF. Direct from the definition of \mathfrak{s} -equivalence and the stack-projection. □

With the definition of \mathfrak{s} -equivalence in place, we indeed obtain the desired composition:

THEOREM 7.39. *If $\mathfrak{f}_1, \mathfrak{f}_2$ be Φ -forests with $\mathfrak{f}_1 \uplus \mathfrak{f}_2 \neq \perp$. Then,*

$$\text{project}(\mathfrak{s}, \mathfrak{f}_1) \bullet_P \text{project}(\mathfrak{s}, \mathfrak{f}_2) = \{\text{project}(\mathfrak{s}, \mathfrak{f}) \mid \mathfrak{f} \in \bar{\mathfrak{f}}_1 \bullet_F \bar{\mathfrak{f}}_2, \bar{\mathfrak{f}}_1 \equiv_{\mathfrak{s}} \mathfrak{f}_1, \bar{\mathfrak{f}}_2 \equiv_{\mathfrak{s}} \mathfrak{f}_2\}.$$

PROOF. See Appendix A.15. □

8 THE TYPE ABSTRACTION

We now formally introduce the abstraction on which our decision procedure for GSL will be built. As motivated in Section 6.4, we abstract every (guarded) state to a Φ -type, which is a set of stack-forest projections. In order to ensure the *finiteness* of the abstraction we need to restrict Φ -types to certain kinds of stack-forest projections. Let us denote by $\text{forests}_{\Phi}(\mathfrak{h}) \triangleq \{\mathfrak{f} \mid \text{heap}(\mathfrak{f}) = \mathfrak{h}\}$ the set of all Φ -forests whose induced heap is \mathfrak{h} . We will then abstract a state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ to a *subset* of the stack-forest projections whose induced heap is \mathfrak{h} , i.e., $\{\text{project}(\mathfrak{s}, \mathfrak{f}) \mid \mathfrak{f} \in \text{forests}_{\Phi}(\mathfrak{h})\}$.

We call the formulas $\text{project}(\mathfrak{s}, \mathfrak{f})$ *unfolded symbolic heaps* (USHs) with respect to SID Φ because any such stack-forest projection can be obtained by “partially unfolding” a symbolic heap (which might require adding appropriate (guarded) quantifiers). Intuitively, the USHs satisfied by a state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ capture all ways in which $\langle \mathfrak{s}, \mathfrak{h} \rangle$ relates to the predicates in SID Φ . While the entire set of USHs is finite for every fixed state $\langle \mathfrak{s}, \mathfrak{h} \rangle$, the set of all USHs w.r.t. an SID Φ is infinite in general:

EXAMPLE 8.1. *Assume the SID Φ defines the list-segment predicate lseg (see Example 3.3). Moreover, let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state with $|\mathfrak{h}| > n \in \mathbb{N}$ such that $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{lseg}(x, \text{nil})$. Then there exists a forest \mathfrak{f} with $\text{heap}(\mathfrak{f}) = \mathfrak{h}$ whose projection consists of n components, i.e.,*

$$\begin{aligned} \text{project}(\mathfrak{s}, \mathfrak{f}) = & \exists y_1, \dots, y_n. \text{lseg}(y_n, \text{nil}) \star (\text{lseg}(y_n, \text{nil}) \rightarrow \text{lseg}(y_{n-1}, \text{nil})) \\ & \star \dots \star (\text{lseg}(y_2, \text{nil}) \rightarrow \text{lseg}(y_1, \text{nil})) \star (\text{lseg}(y_1, \text{nil}) \rightarrow \text{lseg}(x, \text{nil})). \end{aligned}$$

As there exist such states $\langle \mathfrak{s}, \mathfrak{h} \rangle$ for arbitrary natural numbers n , there are infinitely many USHs w.r.t. Φ .

To obtain a *finite* abstraction, we restrict ourselves to *delimited USHs* (DUSHs), in which (1) all root parameters of predicate calls are free variables and (2) every variable occurs at most once as a root parameter on the left-hand side of a magic wand:

DEFINITION 8.2. *An unfolded symbolic heap ϕ is delimited iff*

- (1) for all $\text{pred}(z) \in \phi$, $\text{predroot}(\text{pred}(z)) \in \text{fvvars}(\phi)$, and

- (2) for every variable x there exists at most one predicate call $\text{pred}(z) \in \phi$ such that $\text{pred}(z)$ occurs on the left-hand side of a magic wand and $x = \text{predroot}(\text{pred}(z))$.

The notion of delimited unfolded symbolic heaps is motivated as follows: (1) For every guarded state, the targets of dangling pointers are in the image of the stack. (2) That every variable occurs at most once as the root of a predicate on the left-hand side of a magic wand is a prerequisite for “eliminating” the magic wand through the generalized modus ponens rule.

EXAMPLE 8.3. Recall from Example 3.3 the SID Φ_{tree} that defines binary trees. We consider the stack $\mathfrak{s} = \{x \mapsto l_1, y \mapsto l_2, z \mapsto l_3\}$.

- (1) We consider the following trees and corresponding forest-projections:

$$\begin{aligned} \mathfrak{t}_1 &\triangleq \{l_1 \mapsto \langle \emptyset, \text{tree}(l_1) \Leftarrow (l_1 \mapsto \langle l_2, l_3 \rangle) \star \text{tree}(l_2) \star \text{tree}(l_3) \rangle\} \\ \mathfrak{t}_2 &\triangleq \{l_2 \mapsto \langle \emptyset, \text{tree}(l_2) \Leftarrow l_2 \mapsto \langle \text{nil}, \text{nil} \rangle \rangle\}, \quad \mathfrak{t}_3 \triangleq \{l_3 \mapsto \langle \emptyset, \text{tree}(l_3) \Leftarrow l_3 \mapsto \langle \text{nil}, \text{nil} \rangle \rangle\}, \\ \text{project}(\mathfrak{s}, \{\mathfrak{t}_1, \mathfrak{t}_2, \mathfrak{t}_3\}) &= ((\text{tree}(y) \star \text{tree}(z)) \dashv \star \text{tree}(x)) \star \text{tree}(y) \star \text{tree}(z)) \\ \bar{\mathfrak{t}} &\triangleq \{l_1 \mapsto \langle \langle l_2, l_3 \rangle, \mathfrak{t}_1(l_1) \rangle\} \cup \mathfrak{t}_2 \cup \mathfrak{t}_3 \quad \text{project}(\mathfrak{s}, \{\bar{\mathfrak{t}}\}) = \text{tree}(x) \end{aligned}$$

We observe that $\text{project}(\mathfrak{s}, \{\mathfrak{t}_1, \mathfrak{t}_2, \mathfrak{t}_3\})$ and $\text{project}(\mathfrak{s}, \{\bar{\mathfrak{t}}\})$ are delimited. $\text{project}(\mathfrak{s}, \{\bar{\mathfrak{t}}\})$ can be obtained from $\text{project}(\mathfrak{s}, \{\mathfrak{t}_1, \mathfrak{t}_2, \mathfrak{t}_3\})$ by two applications of modus ponens.

- (2) We consider the following trees and the projection of the corresponding forest:

$$\begin{aligned} \mathfrak{t}_1 &\triangleq \{l_1 \mapsto \langle \emptyset, \text{tree}(l_1) \Leftarrow (l_1 \mapsto \langle l_2, l_2 \rangle) \star \text{tree}(l_2) \star \text{tree}(l_2) \rangle\} \\ \mathfrak{t}_2 &\triangleq \{l_2 \mapsto \langle \emptyset, \text{tree}(l_2) \Leftarrow l_2 \mapsto \langle \text{nil}, \text{nil} \rangle \rangle\} \\ \text{project}(\mathfrak{s}, \{\mathfrak{t}_1, \mathfrak{t}_2\}) &= ((\text{tree}(y) \star \text{tree}(y)) \dashv \star \text{tree}(x)) \star \text{tree}(y) \end{aligned}$$

We note that $\text{project}(\mathfrak{s}, \{\mathfrak{t}_1, \mathfrak{t}_2\})$ is not delimited because the variable y appears twice on the LHS of a magic wand; at most one occurrence of y can be eliminated using modus ponens.

We collect the set of all delimited unfolded symbolic heaps (DUSH) over the SID Φ in

$$\text{DUSH}_\Phi \triangleq \{\text{project}(\mathfrak{s}, \mathfrak{f}) \mid \mathfrak{s} \in \text{Stacks}, \mathfrak{f} \text{ is a } \Phi\text{-forest, project}(\mathfrak{s}, \mathfrak{f}) \text{ is delimited}\}.$$

We are now ready to introduce the type abstraction. Given a state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ and an SID Φ , we call the set of all projections of Φ -forests capturing the heap \mathfrak{h} in the DUSH fragment the Φ -type of $\langle \mathfrak{s}, \mathfrak{h} \rangle$:

DEFINITION 8.4 (Φ -TYPE). The Φ -type (type for short) of a state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ and an SID Φ is given by

$$\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}) \triangleq \{\text{project}(\mathfrak{s}, \mathfrak{f}) \mid \mathfrak{f} \in \text{forests}_\Phi(\mathfrak{h})\} \cap \text{DUSH}_\Phi.$$

In the remainder of this section, we discuss the main results and building blocks required for turning the type abstraction into a decision procedure for guarded separation logic (GSL).

8.1 Understanding Satisfiability as Computing Types

The main idea underlying our decision procedure is to reduce the satisfiability problem for GSL—“given a GSL formula ϕ , does ϕ have a model $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_\Phi \phi$?”—to the question of whether some type \mathcal{T} can be computed from a model of ϕ , i.e., $\mathcal{T} = \text{type}_\Phi(\mathfrak{s}, \mathfrak{h})$ should hold for some $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_\Phi \phi$; the set of all such types will be formally defined further below.

8.1.1 Aliasing Constraints. To conveniently reason about sets of types, we require that types in the same set have the same free variables and the same aliases, i.e., we will group types by *aliasing constraint*—an equivalence relation $\mathbf{ac} \subseteq \mathbf{Var} \times \mathbf{Var}$ representing all aliases under consideration. More formally, every stack \mathfrak{s} induces an aliasing constraint $\text{aliasing}(\mathfrak{s})$ given by

$$\text{aliasing}(\mathfrak{s}) \triangleq \{(x, y) \mid x, y \in \text{dom}(\mathfrak{s}) \text{ and } \mathfrak{s}(x) = \mathfrak{s}(y)\}.$$

We denote the *domain* of an aliasing constraint \mathbf{ac} by $\text{dom}(\mathbf{ac}) \triangleq \{x \mid (x, x) \in \mathbf{ac}\}$. Furthermore, we write $\mathbf{ac}(x)$ for the set of aliases of x given by the aliasing constraint \mathbf{ac} , i.e., the equivalence class $\mathbf{ac}(x) \triangleq \{y \mid (x, y) \in \mathbf{ac}\}$ of \mathbf{ac} that contains x . To obtain a canonical formalization, we frequently represent the equivalence class of x by its largest⁸ member; formally, $[x]_{=}^{\mathbf{ac}} \triangleq \max \mathbf{ac}(x)$.

8.1.2 From GSL satisfiability to types. As outlined at the beginning of Section 8.1, our decision procedure will be based on computing sets of types of the following form:

DEFINITION 8.5 (ac-TYPES). *Let \mathbf{ac} be an aliasing constraint (cf. Section 8.1.1). Then the set $\mathbf{Types}_{\Phi}^{\mathbf{ac}}(\phi)$ of \mathbf{ac} -types of GSL formula ϕ is defined as*

$$\mathbf{Types}_{\Phi}^{\mathbf{ac}}(\phi) \triangleq \{\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) \mid \mathfrak{s} \in \mathbf{Stacks}, \mathfrak{h} \in \mathbf{Heaps}, \text{aliasing}(\mathfrak{s}) = \mathbf{ac}, \langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi\}.$$

By the above definition, a GSL formula ϕ with at least one non-empty set of \mathbf{ac} -types is satisfiable: some type coincides with $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$, where $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$. Conversely, if ϕ is satisfiable, then there exists a model $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$ and the set $\mathbf{Types}_{\Phi}^{\text{aliasing}(\mathfrak{s})}(\phi)$ is non-empty. In summary:

$$\phi \text{ is satisfiable} \quad \text{iff} \quad \exists \mathbf{ac}. \mathbf{Types}_{\Phi}^{\mathbf{ac}}(\phi) \neq \emptyset.$$

On a first glance, finding a suitable aliasing constraint \mathbf{ac} and proving non-emptiness of $\mathbf{Types}_{\Phi}^{\mathbf{ac}}(\phi)$ might appear as difficult as finding a state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ such that $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$ holds due to three concerns:

- (1) There are, in general, both infinitely many aliasing constraints and infinitely many Φ -types, because the size of stacks—and thus the number of free variables to consider—is unbounded.
- (2) Even if the set $\mathbf{Types}_{\Phi}^{\mathbf{ac}}(\phi)$ is finite, effectively computing it is non-trivial.
- (3) Deciding whether a type \mathcal{T} belongs to $\mathbf{Types}_{\Phi}^{\mathbf{ac}}(\phi)$ is non-trivial: assume that $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$, $\langle \mathfrak{s}', \mathfrak{h}' \rangle \not\models_{\Phi} \phi$, and both states yield the same type, i.e., $\mathcal{T} = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) = \text{type}_{\Phi}(\mathfrak{s}', \mathfrak{h}')$. Determining that $\mathcal{T} \in \mathbf{Types}_{\Phi}^{\mathbf{ac}}(\phi)$ would then require us to know that \mathcal{T} can be computed from a specific state, namely $\langle \mathfrak{s}, \mathfrak{h} \rangle$.

As informally motivated in Section 6, our type abstraction can deal with each of the above concerns; we provide the formal details addressing each concern in the remainder of this section:

Regarding (1), we discuss in Section 8.2 how both aliasing constraints and types can safely be restricted to finite subsets. Determining whether $\exists \mathbf{ac}. \mathbf{Types}_{\Phi}^{\mathbf{ac}}(\phi) \neq \emptyset$ holds thus amounts to computing finitely many finite sets. This corresponds to achieving *finiteness* in Section 6.

Regarding (2), we introduce operations for effectively computing Φ -types from existing ones in Sections 8.3 and 8.4; they will be the building blocks of our decision procedure. This corresponds to achieving *compositionality* in Section 6.

Regarding (3), we show in Section 8.5 that one can decide whether a type \mathcal{T} belongs to $\mathbf{Types}_{\Phi}^{\mathbf{ac}}(\phi)$ *without reverting to any state underlying \mathcal{T}* . In particular, we will show that states yielding the same Φ -type satisfy the same GSL formulas. This corresponds to achieving *refinement* in Section 6.

⁸w.r.t. the linear ordering over variables we assume throughout this article; notice that the maximum is well-defined as long as the set of aliases of a variable is finite.

8.2 Finiteness

To ensure finiteness of the type abstraction, we only consider stacks with variables taken from some arbitrary, but fixed, finite set \mathbf{x} of variables. In particular, we denote by $\text{DUSH}_{\Phi}^{\mathbf{x}}$ the restriction of delimited unfolded symbolic heaps (DUSH_{Φ}) to formulas ϕ with free variables in \mathbf{x} , i.e., $\text{fvars}(\phi) \subseteq \mathbf{x}$. With this restriction in place, we are immediately able to argue the finiteness of the DUSH Fragment based on the following observation: Every variable can appear at most twice (once as the projection of a hole and once as the projection of a tree).

LEMMA 8.6. *Let $n \triangleq |\Phi| + |\mathbf{x}|$, where \mathbf{x} is a finite set of variables. Then $|\text{DUSH}_{\Phi}^{\mathbf{x}}| \in 2^{\mathcal{O}(n^2 \log(n))}$.*

PROOF. See Appendix A.16. \square

Analogously to $\text{DUSH}_{\Phi}^{\mathbf{x}}$, we only consider aliasing constraints and types over the finite set \mathbf{x} , i.e., we consider the finite set of aliasing constraints $\text{AC}^{\mathbf{x}} \triangleq \{\text{aliasing}(\mathfrak{s}) \mid \mathfrak{s} \in \text{Stacks}, \text{dom}(\mathfrak{s}) = \mathbf{x}\}$. We note that the number of aliasing constraints in $\text{AC}^{\mathbf{x}}$ equals the $|\mathbf{x}|$ -th Bell number, bounded by $n^n \in \mathcal{O}(2^{n \log(n)})$, where $n = |\mathbf{x}|$. Furthermore, we collect in $\text{Types}_{\Phi}^{\mathbf{x}}$ all ac-types over Φ and \mathbf{x} , i.e.,

$$\text{Types}_{\Phi}^{\mathbf{x}} \triangleq \bigcup_{\text{ac} \in \text{AC}^{\mathbf{x}}} \bigcup_{\phi \in \text{GSL}} \text{Types}_{\Phi}^{\text{ac}}(\phi).$$

The above restriction of types to variables in \mathbf{x} indeed ensures finiteness:

THEOREM 8.7. *Let $\mathbf{x} \subseteq \text{Var}$ be finite and $n \triangleq |\Phi| + |\mathbf{x}|$. Then $|\text{Types}_{\Phi}^{\mathbf{x}}| \in 2^{2^{\mathcal{O}(n^2 \log(n))}}$.*

PROOF. Recall from Lemma 8.6 that the set $\text{DUSH}_{\Phi}^{\mathbf{x}}$ of DUSHs over Φ with free variables taken from \mathbf{x} is of size $2^{\mathcal{O}(n^2 \log(n))}$. We show below that every type $\mathcal{T} \in \text{Types}_{\Phi}^{\mathbf{x}}$ is a subset of $\text{DUSH}_{\Phi}^{\mathbf{x}}$. Hence, the size of $\text{Types}_{\Phi}^{\mathbf{x}}$ is bounded by the number of subsets of $\text{DUSH}_{\Phi}^{\mathbf{x}}$, i.e., $|\text{Types}_{\Phi}^{\mathbf{x}}| \in 2^{2^{\mathcal{O}(n^2 \log(n))}}$.

It remains to show that, for every $\mathcal{T} \in \text{Types}_{\Phi}^{\mathbf{x}}$, we have $\mathcal{T} \subseteq \text{DUSH}_{\Phi}^{\mathbf{x}}$: By definition of $\text{Types}_{\Phi}^{\mathbf{x}}$, there exists an aliasing constraint $\text{ac} \in \text{AC}^{\mathbf{x}}$ and a GSL formula ϕ such that $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\phi)$. By Definition 8.5, there exists a state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ such that $\text{dom}(\mathfrak{s}) = \text{dom}(\text{ac}) \subseteq \mathbf{x}$, $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$, and $\mathcal{T} = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$. By Definition 8.4, $\mathcal{T} = \{\text{project}(\mathfrak{s}, \mathfrak{f}) \mid \mathfrak{f} \in \text{forests}_{\Phi}(\mathfrak{h})\} \cap \text{DUSH}_{\Phi} \subseteq \text{DUSH}_{\Phi}^{\mathbf{x}}$. \square

8.3 \mathfrak{s} -Delimited Forests

To introduce the forests that correspond to DUSHs we make use of the notions of an *interface* of a Φ -forest, which is the set of locations that appear in some tree either as the root or as a hole:

DEFINITION 8.8 (INTERFACE). *The interface of a Φ -forest $\mathfrak{f} = \{t_1, \dots, t_k\}$ is given by*

$$\text{interface}(\mathfrak{f}) \triangleq \bigcup_{1 \leq i \leq k} (\{\text{root}(t_i)\} \cup \text{allholes}(t_i)).$$

EXAMPLE 8.9 (INTERFACE). *Recall the forest \mathfrak{f} from Example 7.5. We have $\text{interface}(\mathfrak{f}) = \{l_1, l_2, l_3\}$: the locations l_1, l_2, l_3 all occur as the roots of a tree; l_2 and l_3 additionally occur as holes (of t_3 and t_1 , respectively); l_4 occurs neither as root nor as hole of a tree and is thus not part of the interface.*

An \mathfrak{s} -delimited forest is a Φ -forest whose interface consists only of locations covered by some stack variable and which does not have any duplicate holes:

DEFINITION 8.10 (\mathfrak{s} -DELIMITED Φ -FOREST). *A Φ -forest \mathfrak{f} is \mathfrak{s} -delimited iff (1) $\text{interface}(\mathfrak{f}) \subseteq \text{img}(\mathfrak{s})$, and (2) for every $l \in \text{allholes}(\mathfrak{f})$ in some tree $t \in \mathfrak{f}$, there is exactly one rule instance (for $l' \in \text{dom}(t)$)*

$$\text{rule}_t(l') = \text{pred}(\mathfrak{v}) \Leftarrow (a \mapsto b) \star \text{pred}_1(\mathfrak{v}_1) \star \dots \star \text{pred}_m(\mathfrak{v}_m) \star \Pi$$

and exactly one index $i \in [1, m]$ such that $\text{predroot}(\text{pred}_i(\mathfrak{v}_i)) = l$.

EXAMPLE 8.11. We consider the forests from Example 8.3: We note that $\{t_1, t_2, t_3\}$ from (1) is \mathfrak{s} -delimited, while $\{t_1, t_2\}$ from (2) is not.

A Φ -forest is \mathfrak{s} -delimited precisely when its projection is delimited (see Appendix A.17 for a proof):

LEMMA 8.12. Let \mathfrak{f} be a forest and let \mathfrak{s} be a stack. Then \mathfrak{f} is \mathfrak{s} -delimited iff $\text{project}(\mathfrak{s}, \mathfrak{f})$ is delimited.

We now state that the \mathfrak{s} -delimitedness of forests is preserved under decomposition; this result will allow us to lift the composition of DUSH formulas (resp. \mathfrak{s} -delimited forests) to types.

THEOREM 8.13. Let $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle, \langle \mathfrak{s}, \mathfrak{h}_2 \rangle \in \mathbf{GStates}$ be guarded states, and let \mathfrak{f} be an \mathfrak{s} -delimited forest with $\mathfrak{f} \in \text{forests}_{\Phi}(\mathfrak{h}_1 \uplus \mathfrak{h}_2)$. Then, there exist \mathfrak{s} -delimited forests $\mathfrak{f}_1, \mathfrak{f}_2$ with $\text{heap}(\mathfrak{f}_i) = \mathfrak{h}_i$ and $\mathfrak{f} \in \mathfrak{f}_1 \bullet_{\mathbf{F}} \mathfrak{f}_2$.

PROOF. See Appendix A.18. □

8.4 Operations on Types

type composition, renaming of variables, forgetting variables, and type extension. These operations will be the building blocks of our decision procedure for GSL.

8.4.1 *Type Composition.* We define an operation \bullet on the level of Φ -types such that $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2) = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1) \bullet \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_2)$, i.e., $\text{type}_{\Phi}(\mathfrak{s}, \cdot)$ is a homomorphism w.r.t. to the operation \uplus on heaps and the operation \bullet on types. As justified below, we can define \bullet by applying our composition operation for forest projections, $\bullet_{\mathbf{P}}$ (cf. Definition 7.31), to all elements of the involved types.

THEOREM 8.14 (COMPOSITIONALITY OF Φ -TYPES). For all guarded states $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle$ and $\langle \mathfrak{s}, \mathfrak{h}_2 \rangle$ with $\mathfrak{h}_1 \uplus \mathfrak{h}_2 \neq \perp$, $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2)$ can be computed from $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1)$ and $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_2)$ as follows:

$$\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2) = \{\phi \in \text{DUSH}_{\Phi} \mid \text{ex. } \psi_1 \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1), \psi_2 \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_2) \text{ such that } \phi \in \psi_1 \bullet_{\mathbf{P}} \psi_2\}.$$

PROOF. See Appendix A.19. □

Our second consideration for defining the composition operation \bullet on Φ -types is that the operation \uplus is only defined on disjoint heaps. In order to be able to express a corresponding condition on the level of types, we will make use of the following notion:

DEFINITION 8.15 (ALLOCATED VARIABLES OF A TYPE). The set of allocated variables of Φ -type \mathcal{T} is

$$\text{alloted}(\mathcal{T}) \triangleq \{x \mid \text{there ex. } \phi \in \mathcal{T} \text{ and } (\psi \star \text{pred}(z)) \text{ in } \phi \text{ s.t. } x = \text{predroot}(\text{pred}(z))\}.$$

The above notion is motivated by the fact that, for each non-empty type, the allocated variables of the type agree with the allocated variables of every state having that type.

LEMMA 8.16. Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state with $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) \neq \emptyset$. Then, $\text{alloted}(\mathfrak{s}, \mathfrak{h}) = \text{alloted}(\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}))$.

PROOF. See Appendix A.20. □

We note that, for every model $\langle \mathfrak{s}, \mathfrak{h} \rangle$ of some predicate call $\text{pred}(z_1, \dots, z_k)$, there is at least one tree \mathfrak{t} with $\text{heap}(\{\mathfrak{t}\}) = \mathfrak{h}$ (see Lemma 7.3); hence, $\text{project}(\mathfrak{s}, \{\mathfrak{t}\}) \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$ and the non-emptiness requirement of Lemma 8.16 is fulfilled—a fact that generalizes to all models of guarded formulas:

LEMMA 8.17. Let $\phi \in \text{GSL}$ and let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state with $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$. Then, $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) \neq \emptyset$.

PROOF. See Appendix A.21. □

We are now ready to state our composition operation \bullet on Φ -types:

DEFINITION 8.18 (TYPE COMPOSITION). The composition $\mathcal{T}_1 \bullet \mathcal{T}_2$ of Φ -types \mathcal{T}_1 and \mathcal{T}_2 is given by

$$\mathcal{T}_1 \bullet \mathcal{T}_2 \triangleq \begin{cases} \perp, & \text{if } \text{allocated}(\mathcal{T}_1) \cap \text{allocated}(\mathcal{T}_2) \neq \emptyset, \\ \left(\bigcup_{\phi_1 \in \mathcal{T}_1, \phi_2 \in \mathcal{T}_2} \phi_1 \bullet_P \phi_2 \right) \cap \text{DUSH}_\Phi, & \text{otherwise.} \end{cases}$$

We now state two results that \bullet has the desired properties, i.e., that $\text{type}_\Phi(\mathfrak{s}, \cdot)$ is a homomorphism w.r.t. to the operation \uplus on heaps and the operation \bullet on types (cf. Appendices A.22 and A.23):

COROLLARY 8.19 (COMPOSITIONALITY OF TYPE ABSTRACTION). For guarded states $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle$ and $\langle \mathfrak{s}, \mathfrak{h}_2 \rangle$ with $\mathfrak{h}_1 \uplus \mathfrak{h}_2 \neq \perp$, we have $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2) = \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_1) \bullet \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_2)$.

LEMMA 8.20. For $i \in \{1, 2\}$, let $\langle \mathfrak{s}, \mathfrak{h}_i \rangle$ be states with $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_i) = \mathcal{T}_i \neq \emptyset$ and $\mathcal{T}_1 \bullet \mathcal{T}_2 \neq \perp$. Then, there are states $\langle \mathfrak{s}, \mathfrak{h}'_i \rangle$ such that $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}'_i) = \mathcal{T}_i$ and $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}'_1 \uplus \mathfrak{h}'_2) = \mathcal{T}_1 \bullet \mathcal{T}_2$.

8.4.2 *Renaming Variables.* To compute the types of predicate calls $\text{pred}(y)$ compositionally, we need a mechanism to rename variables in Φ -types: Once we know the types of a predicate call $\text{pred}(x)$ over the formal arguments $x = \text{fvars}(\text{pred})$, we can compute the types of $\text{pred}(y)$ by renaming x to y . Such a renaming amounts to a simple variable substitution:

DEFINITION 8.21 (VARIABLE RENAMING). Let x be a sequence of pairwise distinct variables, let y be an arbitrary sequence of variables with $|y| = |x|$, and let ac be an aliasing constraint with $y \subseteq \text{dom}(\text{ac})$. Moreover, let y' be the sequence obtained by replacing every variable in $y \in y$ by the maximal variable in its equivalence class, i.e., by $[y]_{\text{ac}}^{\text{ac}}$. Then, the $[x/y]$ -renaming of type \mathcal{T} w.r.t. aliasing constraint ac is given by $\mathcal{T}[\text{ac} : x/y] \triangleq \{\phi[x/y'] \mid \phi \in \mathcal{T}\}$.

Variable renaming is compositional as it corresponds to first renaming variables at the level of stacks and then computing the type of the resulting state. More formally, assume a state $\langle \mathfrak{s}, \mathfrak{h} \rangle$, where we already renamed x to y in stack \mathfrak{s} ; in particular, $x \cap \text{dom}(\mathfrak{s}) = \emptyset$. Computing $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h})$ then coincides with the $[x/y]$ -renaming of $\text{type}_\Phi(\mathfrak{s}', \mathfrak{h})$, where $\mathfrak{s}' = \mathfrak{s}[x/y] \triangleq \mathfrak{s}[x/\mathfrak{s}(y)]$ is the stack \mathfrak{s} in which the variables x have not been renamed to y yet (cf. Appendix A.24).⁹

LEMMA 8.22. For x, y as above and a stack \mathfrak{s} with $y \subseteq \text{dom}(\mathfrak{s})$ and $x \cap \text{dom}(\mathfrak{s}) = \emptyset$, we have

$$\text{type}_\Phi(\mathfrak{s}[x/y], \mathfrak{h})[\text{aliasing}(\mathfrak{s}) : x/y] = \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}).$$

8.4.3 *Forgetting Variables.* Our third operation on types removes a free variable x from a type \mathcal{T} . Intuitively, for every formula $\phi \in \mathcal{T}$, there are two cases:

- (1) If x aliases with some free variable, then we replace x by its largest alias.
- (2) If x does not alias with any free variable, then we remove it from the set of free variables by introducing a (guarded) existential quantifier.

Formally, we fix an aliasing constraint ac (cf. Section 8.1.1) characterizing which free variables are aliases. Forgetting a variable x in a formula ϕ with respect to ac is then defined as follows:

$$\text{forget}_{\text{ac}, x}(\phi) \triangleq \begin{cases} \phi[x/\max(\text{ac}(x) \setminus \{x\})], & \text{if } x \in \text{fvars}(\phi) \text{ and } \text{ac}(x) \neq \{x\}, \\ \exists x. \phi, & \text{if } x \in \text{fvars}(\phi) \text{ and } \text{ac}(x) = \{x\}, \\ \phi, & \text{if } x \notin \text{fvars}(\phi). \end{cases}$$

Forgetting a variable in a type \mathcal{T} corresponds to applying the above operation to all $\phi \in \mathcal{T}$. However, $\text{forget}_{\text{ac}, x}(\phi)$ does—in general—not belong to the fragment DUSH_Φ because we might existentially quantify over a root variable of ϕ . Hence, we additionally intersect the result with DUSH_Φ :

⁹Recall that $\mathfrak{s}[u/v]$ denotes a stack *update* in which variables in u are added to the domain of stack \mathfrak{s} if necessary.

DEFINITION 8.23 (FORGETTING A VARIABLE). *The Φ -type obtained from forgetting variable x in Φ -type \mathcal{T} w.r.t. aliasing constraint ac is defined by $\text{forget}_{\text{ac},x}(\mathcal{T}) \triangleq \{\text{forget}_{\text{ac},x}(\phi) \mid \phi \in \mathcal{T}\} \cap \text{DUSH}_{\Phi}$.*

The above operation is compositional as forgetting an allocated variable in the type of a guarded state coincides with first removing the variable from the state and then computing its type:

LEMMA 8.24. *Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a guarded state such that $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h})$ holds for some variable x . Then,*

$$\text{forget}_{\text{aliasing}(\mathfrak{s}),x}(\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})) = \text{type}_{\Phi}(\mathfrak{s}[x/\perp], \mathfrak{h}).$$

PROOF. See Appendix A.25. □

8.4.4 *Type Extension.* Our fourth and final operation is concerned with extending types to stacks over larger domains. To this end, we instantiate universally quantified variables with free variables that do not appear in the type so far. Formally, let $\phi = \exists \mathbf{e}. \forall (\mathbf{a} \cdot \mathbf{u} \cdot \mathbf{b}). \psi$ be a formula and let x be a fresh variable, i.e., $x \notin \text{fvars}(\phi)$. We call the formula $\exists \mathbf{e}. \forall (\mathbf{a} \cdot \mathbf{b}). \psi[u/x]$ an x -instantiation of ϕ . Extending a type by variable x then corresponds to adding all x -instantiations of its members:

DEFINITION 8.25 (x -EXTENSION OF A TYP). *The x -extension of a Φ -type \mathcal{T} by a fresh variable x is*

$$\text{extend}_x(\mathcal{T}) \triangleq \mathcal{T} \cup \{\phi' \mid \phi' \text{ is an } x\text{-instantiation of } \phi \mid \phi \in \mathcal{T}\}.$$

As for the other operations, the x -extension of a type is compositional in the sense that it coincides with computing the type of a state with an already extended stack (cf. Appendix A.26):

LEMMA 8.26. *For every state $\langle \mathfrak{s}, \mathfrak{h} \rangle$, variable x with $\mathfrak{s}(x) \notin \text{locs}(\mathfrak{h})$ and $\text{aliasing}(\mathfrak{s})(x) = \{x\}$,*

$$\text{extend}_x(\text{type}_{\Phi}(\mathfrak{s}[x/\perp], \mathfrak{h})) = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}).$$

Rather than extending a type by a single variable, it will be convenient to extend it by all variables in an *aliasing constraint* that are not aliases of an existing variable.

DEFINITION 8.27 (EXTENSION OF A TYPE WITH REGARD TO AN ALIASING CONSTRAINT). *Let $\text{ac} \subseteq \text{ac}'$ be aliasing constraints. Let \mathbf{y} be a repetition-free sequence of all maximal variables in $\text{dom}(\text{ac})$, and let \mathbf{y}' be the sequence obtained by replacing every variable in $\mathbf{y} \in \mathbf{y}$ by the corresponding maximal variable in ac' , i.e., by $[y]_{\perp}^{\text{ac}'}$.¹⁰ Moreover, let $\mathbf{z} = \langle z_1, \dots, z_n \rangle$, $n \geq 0$, be a repetition-free sequence of all maximal variables in $\text{dom}(\text{ac}')$ that are not aliases of variables in $\text{dom}(\text{ac})$.¹¹ Then the ac' -extension of a Φ -type \mathcal{T} w.r.t. aliasing constraint ac is defined as $\text{extend}_{\text{ac}'}(\mathcal{T}) \triangleq \mathcal{T}_n$, where*

$$\mathcal{T}_k = \begin{cases} \{\phi[\mathbf{y}/\mathbf{y}'] \mid \phi \in \mathcal{T}\}, & \text{if } k = 0 \\ \text{extend}_{z_k}(\mathcal{T}_{k-1}), & \text{if } 0 < k \leq n. \end{cases}$$

The above operation preserves compositionality as it boils down to multiple type extensions:

LEMMA 8.28. *Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state and ac be an aliasing constraint with $\text{ac} \subseteq \text{aliasing}(\mathfrak{s})$. Let \mathfrak{s}' be the restriction of \mathfrak{s} to the domain $\text{dom}(\text{ac})$. If $\mathfrak{s}(x) \notin \text{locs}(\mathfrak{h})$ for every variable $x \in \text{dom}(\mathfrak{s})$ that is not an alias of a variable in $\text{dom}(\text{ac})$, then $\text{extend}_{\text{aliasing}(\mathfrak{s})}(\text{type}_{\Phi}(\mathfrak{s}', \mathfrak{h})) = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$.*

PROOF. Let $k \geq 0$ be the number of variables in $\text{dom}(\mathfrak{s})$ that are no aliases of variables in $\text{dom}(\text{ac})$. By Definition 8.27, we have $\text{extend}_{\text{aliasing}(\mathfrak{s})}(\text{type}_{\Phi}(\mathfrak{s}', \mathfrak{h})) = \mathcal{T}_k$, i.e., we need to apply k type extensions. The claim then follows from Lemma 8.26 by induction on k . □

¹⁰We recall that we need maximal variables for maintaining canonic projections, i.e., type representations.

¹¹I.e., $z \in \mathbf{z}$ iff $z \in \text{dom}(\text{ac}')$, $z = [z]_{\perp}^{\text{ac}'}$ and $z \notin \text{ac}'(y)$ for all $y \in \text{dom}(\text{ac})$.

8.5 Type Refinement

The main insight required for effectively deciding whether a type \mathcal{T} belongs to $\mathbf{Types}_{\Phi}^{\text{ac}}(\phi)$ is that states with identical Φ -types satisfy the same GSL formulas—a statement we formalize below. This property is perhaps surprising, as types only contain formulas from the DUSH fragment, which is largely orthogonal to GSL. For example, GSL formulas allow guarded negation and guarded septraction, but neither quantifiers nor unguarded magic wands, whereas DUSHs allow limited use of guarded quantifiers and unguarded magic wands, but neither Boolean structure nor septraction.

THEOREM 8.29 (REFINEMENT THEOREM). *For all stacks \mathfrak{s} , heaps $\mathfrak{h}_1, \mathfrak{h}_2$, and GSL formulas ϕ ,*

$$\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1) = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_2) \quad \text{implies} \quad \langle \mathfrak{s}, \mathfrak{h}_1 \rangle \models_{\Phi} \phi \quad \text{iff} \quad \langle \mathfrak{s}, \mathfrak{h}_2 \rangle \models_{\Phi} \phi .$$

PROOF. See Appendix A.27. □

Theorem 8.29 immediately implies that, if the type of a state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ is equal to some already-known type of some other state $\langle \mathfrak{s}', \mathfrak{h}' \rangle$ satisfying formula ϕ , then $\langle \mathfrak{s}, \mathfrak{h} \rangle$ satisfies ϕ .

COROLLARY 8.30. *If there is a type $\mathcal{T} \in \mathbf{Types}_{\Phi}^{\text{aliasing}(\mathfrak{s})}(\phi)$ with $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) = \mathcal{T}$, then $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$.*

Moreover, recall that GSL formulas are quantifier-free (although quantifiers may appear in predicate definitions). As demonstrated below, this limitation is crucial for upholding Theorem 8.29.

EXAMPLE 8.31. *Recall Φ_{ls} from Example 3.3. Moreover, let $\langle \mathfrak{s}, \mathfrak{h}_k \rangle$, $k \in \mathbb{N}$, be a list of length k from x_1 to x_2 . It then holds for all $i, j \geq 2$ that $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_i) = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_j)$. However,*

$$\begin{aligned} \langle \mathfrak{s}, \mathfrak{h}_2 \rangle &\not\models_{\Phi} \exists \langle y_1, y_2 \rangle . \text{lseg}(x_1, y_1) \star \text{lseg}(y_1, y_2) \star \text{lseg}(y_2, x_2), \quad \text{whereas, for all } j \geq 3, \\ \langle \mathfrak{s}, \mathfrak{h}_j \rangle &\models_{\Phi} \exists \langle y_1, y_2 \rangle . \text{lseg}(x_1, y_1) \star \text{lseg}(y_1, y_2) \star \text{lseg}(y_2, x_2). \end{aligned}$$

Hence, the refinement theorem does not hold if we admit quantifiers in GSL formulas.

9 ALGORITHMS FOR COMPUTING TYPES

As discussed in Section 8.1, deciding whether a GSL formula ϕ is decidable boils down to computing finite sets of types $\mathbf{Types}_{\Phi}^{\text{ac}}(\phi)$ for suitable aliasing constraints ac . We now present two algorithms for effectively computing $\mathbf{Types}_{\Phi}^{\text{ac}}(\phi)$: Section 9.1 deals with computing types of predicate calls defined by SIDs and Section 9.2 shows how to compute types of GSL formulas, respectively.¹²

9.1 Computing the Types of Predicate Calls

We first compute, for every predicate $\text{pred} \in \mathbf{Preds}(\Phi)$, the set of all ac -types of pred . Specifically, for every aliasing constraint $\text{ac} \in \mathbf{AC}^{\text{x} \cup \text{fvars}(\text{pred})}$, where $\text{x} \subseteq \mathbf{Var}$ finite, we will compute

$$\mathbf{Types}_{\Phi}^{\text{ac}}(\text{pred}) \triangleq \mathbf{Types}_{\Phi}^{\text{ac}}(\text{pred}(\text{fvars}(\text{pred}))).$$

Once we have a way to compute these types, we can also compute types for any GSL formula with free variables x , as we will see in Section 9.2.

9.1.1 Assumptions. Throughout this section, we fix a pointer-closed SID $\Phi \in \mathbf{ID}_{\text{btw}}$ and a finite set of variables x ; we assume w.l.o.g. that $\text{x} \cap \text{fvars}(\text{pred}) = \emptyset$ for all predicates $\text{pred} \in \mathbf{Preds}(\Phi)$.

$$\begin{aligned}
\text{ptypes}_p^x(x \approx y, \mathbf{ac}) &\triangleq \text{if } \langle x, y \rangle \in \mathbf{ac} \text{ then } \{\{\mathbf{emp}\}\} \text{ else } \emptyset \\
\text{ptypes}_p^x(x \not\approx y, \mathbf{ac}) &\triangleq \text{if } \langle x, y \rangle \notin \mathbf{ac} \text{ then } \{\{\mathbf{emp}\}\} \text{ else } \emptyset \\
\text{ptypes}_p^x(a \mapsto \mathbf{b}, \mathbf{ac}) &\triangleq \{\text{type}_\Phi(\text{ptrmodel}_{\mathbf{ac}}(a \mapsto \mathbf{b}))\} \\
\text{ptypes}_p^x(\text{pred}(y), \mathbf{ac}) &\triangleq \text{let } z \triangleq \text{fvars}(\text{pred}) \text{ in} \\
&\quad \text{extend}_{\mathbf{ac}[z/y]^{-1}}(p(\text{pred}, \mathbf{ac}[z/y]^{-1}|_{x \cup z})[\mathbf{ac} : z/y]) \\
\text{ptypes}_p^x(\phi_1 \star \phi_2, \mathbf{ac}) &\triangleq \text{ptypes}_p^x(\phi_1, \mathbf{ac}) \bullet \text{ptypes}_p^x(\phi_2, \mathbf{ac}) \\
\text{ptypes}_p^x(\exists y. \phi, \mathbf{ac}) &\triangleq \bigcup_{\mathbf{ac}' \in \mathbf{AC}^{\text{dom}(\mathbf{ac}) \cup \{y\}} \text{ with } \mathbf{ac}'|_{\text{dom}(\mathbf{ac})} = \mathbf{ac}} \text{forget}_{\mathbf{ac}', y}(\text{ptypes}_p^x(\phi, \mathbf{ac}'))
\end{aligned}$$

Fig. 11. Computing (a subset of) the Φ -types of existentially-quantified symbolic heap $\phi \in \mathbf{SH}^\exists$ for stacks with aliasing constraint \mathbf{ac} under the assumption that p maps every predicate symbol pred and every aliasing constraint to (a subset of) the types $\text{Types}_\Phi^{\mathbf{ac}}(\text{pred})$. Here, $\mathbf{ac}[u/v]^{-1}$ denotes the addition of the variables u into the aliasing constraint \mathbf{ac} such that the variables u are aliases of the variables v respectively; see Definition 9.1. We denote by $\mathbf{ac}|_y$ the restriction of \mathbf{ac} to the variables in y , i.e., $\mathbf{ac}|_y \triangleq \mathbf{ac} \cap (y \times y)$.

9.1.2 A Fixed-Point Algorithm for Computing the Types of Predicates. We compute $\text{Types}_\Phi^{\mathbf{ac}}(\text{pred})$ for all choices of \mathbf{ac} and pred by a simultaneous fixed-point computation. Specifically, our goal is to compute a (partial) function $p: \text{Preds}(\Phi) \times \mathbf{AC} \rightarrow 2^{\text{Types}_\Phi}$ that maps every predicate pred and every aliasing constraint $\mathbf{ac} \in \mathbf{AC}^{x \cup \text{fvars}(\text{pred})}$ to the set of types $\text{Types}_\Phi^{\mathbf{ac}}(\text{pred})$. We start off the fixed-point computation with $p(\text{pred}, \mathbf{ac}) = \emptyset$ for all pred and \mathbf{ac} ; each iteration adds to p some more types such that $p(\text{pred}, \mathbf{ac}) \subseteq \text{Types}_\Phi^{\mathbf{ac}}(\text{pred})$; and when we reach the fixed point, $p(\text{pred}, \mathbf{ac}) = \text{Types}_\Phi^{\mathbf{ac}}(\text{pred})$ will hold for all pred and \mathbf{ac} . Each iteration amounts to applying the function $\text{ptypes}_p^x(\phi, \mathbf{ac})$ defined in Fig. 11 to all rule bodies $\phi \in \mathbf{SH}^\exists$ of the SID Φ and all aliasing constraints \mathbf{ac} . Here, p is the pre-fixed point from the previous iteration. The function ptypes operates on sets of types. Hence, we need to lift $\bullet, \cdot[\cdot : \cdot/\cdot]$, forget and extend from types to sets of types in a point-wise manner, i.e.,

$$\begin{aligned}
\{\mathcal{T}_1, \dots, \mathcal{T}_m\} \bullet \{\mathcal{T}'_1, \dots, \mathcal{T}'_n\} &\triangleq \{\mathcal{T}_i \bullet \mathcal{T}'_j \mid 1 \leq i \leq m, 1 \leq j \leq n, \mathcal{T}_i \bullet \mathcal{T}'_j \neq \perp\}, \\
\{\mathcal{T}_1, \dots, \mathcal{T}_m\}[\mathbf{ac} : \mathbf{x}/\mathbf{y}] &\triangleq \{\mathcal{T}_1[\mathbf{ac} : \mathbf{x}/\mathbf{y}], \dots, \mathcal{T}_m[\mathbf{ac} : \mathbf{x}/\mathbf{y}]\}, \\
\text{forget}_{\mathbf{ac}, y}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\}) &\triangleq \{\text{forget}_{\mathbf{ac}, y}(\mathcal{T}_1), \dots, \text{forget}_{\mathbf{ac}, y}(\mathcal{T}_m)\}, \text{ and} \\
\text{extend}_{\mathbf{ac}}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\}) &\triangleq \{\text{extend}_{\mathbf{ac}}(\mathcal{T}_1), \dots, \text{extend}_{\mathbf{ac}}(\mathcal{T}_m)\}.
\end{aligned}$$

Further, ptypes uses the following operation on aliasing constraints:

DEFINITION 9.1 (REVERSE RENAMING OF ALIASING CONSTRAINTS). *Let \mathbf{x} be a sequence of pairwise distinct variables and let \mathbf{y} be a sequence of (not necessarily pairwise distinct) variables with $|\mathbf{y}| = |\mathbf{x}|$. Moreover, let \mathbf{ac} be an aliasing constraint with $\mathbf{x} \cap \text{dom}(\mathbf{ac}) = \emptyset$ and $\mathbf{y} \subseteq \text{dom}(\mathbf{ac})$. Then, the reverse renaming \mathbf{x} to \mathbf{y} in \mathbf{ac} is given by the aliasing constraint $\mathbf{ac}[\mathbf{x}/\mathbf{y}]^{-1} \in \mathbf{AC}^{\text{dom}(\mathbf{ac}) \cup \mathbf{x}}$ defined by*

$$\mathbf{ac}[\mathbf{x}/\mathbf{y}]^{-1} \triangleq \{(a_1, a_2) \mid \text{there is } (b_1, b_2) \in \mathbf{ac} \text{ with } b_1 = a_1[\mathbf{x}/\mathbf{y}] \text{ and } b_2 = a_2[\mathbf{x}/\mathbf{y}]\}.$$

Informally, the function $\text{ptypes}_p^x(\phi, \mathbf{ac})$ works as follows:

- If $\phi = x \approx y$ or $\phi = x \not\approx y$, we use the aliasing constraint \mathbf{ac} to check whether the (dis)equality ϕ holds and then return either the type of the empty model or no type. This is justified because our semantics enforces that (dis)equalities only hold in the empty heap.
- If $\phi = a \mapsto \mathbf{b}$, there is—up to isomorphism—only one state with aliasing constraint \mathbf{ac} that satisfies ϕ . We denote this state by $\text{ptrmodel}_{\mathbf{ac}}(a \mapsto \mathbf{b})$ and return its type.

¹²Recall that the formulas in SIDs are symbolic heaps and *not* GSL formulas; for example, they may contain quantifiers.

- If $\phi = \text{pred}(y)$, we look up the types of $\text{pred}(\text{fvvars}(\text{pred}))$ in the pre-fixed point p and then appropriately rename the formal parameters $\text{fvvars}(\text{pred})$ to the actual arguments y :
 - For the look-up we use the aliasing constraint $\text{ac}[z/y]^{-1}$, which is obtained from the aliasing constraint ac by adding the formal parameters $z \triangleq \text{fvvars}(\text{pred})$ to ac such that the z are aliases of the variables y respectively; see Definition 9.1 for details.
 - Crucially, we restrict $\text{ac}[z/y]^{-1}$ to the variables $x \cup z$ before we look up the types of $\text{pred}(\text{fvvars}(\text{pred}))$. This restriction guarantees that the computation of $\text{ptypes}_p^x(\phi, \text{ac})$ does not diverge by considering larger and larger aliasing constraints in recursive calls. (An illustration of the problem as well as an argument why our solution does not lead to divergence can be found in Appendix A.29.2).
 - After the loop-up we extend the types over aliasing constraint $\text{ac}[z/y]^{-1}|_{x \cup z}$ to types over aliasing constraint $\text{ac}[z/y]^{-1}$, undoing the earlier restriction.
 - Finally, we rename the formal parameters z of the recursive call with the actual parameters y and obtain types over aliasing constraint ac .
- If $\phi = \phi_1 \star \phi_2$, we apply the type composition operator developed in previous sections.
- If $\phi = \exists y. \phi'$, we consider all ways to extend the aliasing constraint ac with y and recurse. Our treatment of predicate calls outlined above guarantees that this does not lead to divergence.

Fixed Point Computation. We use the following wrapper for ptypes :

$$\begin{aligned} \text{unfold}_x &: (\mathbf{Preds}(\Phi) \times \mathbf{AC} \rightarrow 2^{\text{Types}_\Phi}) \rightarrow (\mathbf{Preds}(\Phi) \times \mathbf{AC} \rightarrow 2^{\text{Types}_\Phi}), \\ \text{unfold}_x(p) &= \lambda(\text{pred}, \text{ac}). p(\text{pred}, \text{ac}) \cup \bigcup_{(\text{pred}(y) \Leftarrow \phi) \in \Phi} \text{ptypes}_p^x(\text{ac}, \phi). \end{aligned}$$

We observe that unfold_x is a monotone function defined over a finite complete lattice:

- (1) The considered order \sqsubseteq of $\mathbf{Preds}(\Phi) \times \mathbf{AC} \rightarrow 2^{\text{Types}_\Phi}$ is the point-wise comparison of functions:

$$f \sqsubseteq g \triangleq \forall \text{pred} \forall \text{ac}. f(\text{pred}, \text{ac}) \subseteq g(\text{pred}, \text{ac}).$$

- (2) $\mathbf{Preds}(\Phi) \times \mathbf{AC} \rightarrow 2^{\text{Types}_\Phi}$ is a finite lattice because the image 2^{Types_Φ} and the domain $\{\langle \text{pred}, \text{ac} \rangle \mid \text{pred} \in \mathbf{Preds}(\Phi), \text{ac} \in \mathbf{AC}^{x \cup \text{fvvars}(\text{pred})}\}$ of the considered functions are finite.
- (3) $\mathbf{Preds}(\Phi) \times \mathbf{AC} \rightarrow 2^{\text{Types}_\Phi}$ is complete because the image 2^{Types_Φ} of the considered functions is a complete lattice (the subset lattice over Types_Φ).

By Tarski's and Knaster's fixed point theorem the least fixed point of unfold_x exists. This fixed point can be obtained in finitely many steps by Kleene iteration¹³:

$$\text{lfp}(\text{unfold}_x) \triangleq \lim_{n \in \mathbb{N}} \text{unfold}_x^n(\lambda(\text{pred}', \text{ac}'). \emptyset)$$

Moreover, since the lattice is finite, finitely many iterations suffice to reach the least fixed point.

Correctness and Complexity. We analyze the correctness of our construction, i.e.,

$$\text{for all } \text{pred} \in \mathbf{Preds}(\Phi) \text{ and } \text{ac} \in \mathbf{AC}^{x \cup \text{fvvars}(\text{pred})}. \text{lfp}(\text{unfold}_x)(\text{pred}, \text{ac}) = \text{Types}_\Phi^{\text{ac}}(\text{pred}),$$

as well as its complexity in three steps, which can be found in A.29:

- (1) We show $\text{lfp}(\text{unfold}_x)(\text{pred}, \text{ac}) \subseteq \text{Types}_\Phi^{\text{ac}}(\text{pred})$.
- (2) We show $\text{lfp}(\text{unfold}_x)(\text{pred}, \text{ac}) \supseteq \text{Types}_\Phi^{\text{ac}}(\text{pred})$.
- (3) We show that $\text{lfp}(\text{unfold}_x)$ is computable in $2^{2^{O(n^2 \log(n))}}$, where $n \triangleq |\Phi| + |x|$.

¹³To be precise we invoke a constructive version of Tarski's and Knaster's fixed point theorem [Cousot and Cousot 1979], which supports the computation of the least fixed point by transfinite induction; the finiteness of the lattice, however, ensures that the fixed point is already reached after finitely many steps.

$$\begin{aligned}
\text{types}(\mathbf{emp}, \mathbf{ac}) &\triangleq \{\{\mathbf{emp}\}\} \\
\text{types}(x \approx y, \mathbf{ac}) &\triangleq \text{if } \langle x, y \rangle \in \mathbf{ac} \text{ then } \{\{\mathbf{emp}\}\} \text{ else } \emptyset \\
\text{types}(x \not\approx y, \mathbf{ac}) &\triangleq \text{if } \langle x, y \rangle \notin \mathbf{ac} \text{ then } \{\{\mathbf{emp}\}\} \text{ else } \emptyset \\
\text{types}(a \mapsto \mathbf{b}, \mathbf{ac}) &\triangleq \{\text{type}_{\Phi}(\text{ptrmodel}_{\mathbf{ac}}(a \mapsto \mathbf{b}))\} \\
\text{types}(\text{pred}(y), \mathbf{ac}) &\triangleq \text{lfp}(\text{unfold}_{\text{dom}(\mathbf{ac})})(\text{pred}, \mathbf{ac}[\text{fvvars}(\text{pred})/y]^{-1})[\mathbf{ac} : \text{fvvars}(\text{pred})/y] \\
\text{types}(\phi_1 \star \phi_2, \mathbf{ac}) &\triangleq \text{types}(\phi_1, \mathbf{ac}) \bullet \text{types}(\phi_2, \mathbf{ac}) \\
\text{types}(\phi_1 \wedge \phi_2, \mathbf{ac}) &\triangleq \text{types}(\phi_1, \mathbf{ac}) \cap \text{types}(\phi_2, \mathbf{ac}) \\
\text{types}(\phi_1 \vee \phi_2, \mathbf{ac}) &\triangleq \text{types}(\phi_1, \mathbf{ac}) \cup \text{types}(\phi_2, \mathbf{ac}) \\
\text{types}(\phi_1 \wedge \neg \phi_2, \mathbf{ac}) &\triangleq \text{types}(\phi_1, \mathbf{ac}) \setminus \text{types}(\phi_2, \mathbf{ac}) \\
\text{types}(\phi_0 \wedge (\phi_1 \boxplus \phi_2), \mathbf{ac}) &\triangleq \{\mathcal{T} \in \text{types}(\phi_0, \mathbf{ac}) \mid \exists \mathcal{T}' \in \text{types}(\phi_1, \mathbf{ac}). \mathcal{T} \bullet \mathcal{T}' \in \text{types}(\phi_2, \mathbf{ac})\} \\
\text{types}(\phi_0 \wedge (\phi_1 \blackrightarrow \phi_2), \mathbf{ac}) &\triangleq \{\mathcal{T} \in \text{types}(\phi_0, \mathbf{ac}) \mid \forall \mathcal{T}' \in \text{types}(\phi_1, \mathbf{ac}). \mathcal{T} \bullet \mathcal{T}' \in \text{types}(\phi_2, \mathbf{ac})\}
\end{aligned}$$

Fig. 12. Computation of Φ -types for quantifier-free GSL formula ϕ and stacks with aliasing constraint \mathbf{ac} .

9.2 Computing the Types of Guarded Formulas

After we have established how to compute the types of predicate calls, we are now ready to define a function $\text{types}(\phi, \mathbf{ac})$ that computes the types of arbitrary GSL formulas ϕ —i.e., quantifier-free guarded formulas—for some fixed stack-aliasing constraint \mathbf{ac} ; the function is defined in Fig. 12.

THEOREM 9.2 (CORRECTNESS AND COMPLEXITY OF THE TYPE COMPUTATION). *Let $\phi \in \text{GSL}$ with $\text{fvvars}(\phi) = \mathbf{x}$ and $\text{locs}(\phi) = \emptyset$. Further, let $\mathbf{ac} \in \text{AC}^{\mathbf{x}}$. Then, $\text{Types}_{\Phi}^{\mathbf{ac}}(\phi) = \text{types}(\phi, \mathbf{ac})$. Moreover, $\text{types}(\phi, \mathbf{ac})$ can be computed in $2^{2^{O(n^2 \log(n))}}$, where $n \triangleq |\Phi| + |\phi|$.*

We now state the main result of this article:

THEOREM 9.3 (DECIDABILITY OF GSL). *Let $\phi \in \text{GSL}$ and $n \triangleq |\Phi| + |\phi|$. It is decidable in time $2^{2^{O(n^2 \log(n))}}$ whether ϕ is satisfiable.*

PROOF. Let $\mathbf{x} \triangleq \text{fvvars}(\phi)$. Note that $|\mathbf{x}| \leq n$. The formula ϕ is satisfiable iff there exists a state $\langle \mathbf{s}, \mathbf{h} \rangle$ with $\langle \mathbf{s}, \mathbf{h} \rangle \models_{\Phi} \phi$. By Lemma 8.17, $\text{type}_{\Phi}(\mathbf{s}, \mathbf{h}) \neq \emptyset$. Hence, it is sufficient to compute $\text{Types}_{\Phi}^{\mathbf{ac}}(\phi)$ for all aliasing constraints \mathbf{ac} with $\text{dom}(\mathbf{ac}) = \mathbf{x}$ and check whether $\text{Types}_{\Phi}^{\mathbf{ac}}(\phi) \neq \emptyset$.

By Theorem 9.2 we can compute $\text{Types}_{\Phi}^{\mathbf{ac}}(\phi) = \text{types}(\phi, \mathbf{ac})$ in $2^{2^{O(n^2 \log(n))}}$ for a fixed aliasing constraints \mathbf{ac} . Since there are at most $n^n \in O(2^{n \log(n)})$ stack-aliasing constraints, we can conclude that we can perform the satisfiability check in time $O(2^{n \log(n)}) \cdot 2^{2^{O(n^2 \log(n))}} = 2^{2^{O(n^2 \log(n))}}$. \square

Since the entailment query $\phi \models_{\Phi} \psi$ is equivalent to checking the unsatisfiability of $\phi \wedge \neg \psi$, and the negation in $\phi \wedge \neg \psi$ is guarded, we obtain an entailment checker with the same complexity:

COROLLARY 9.4 (DECIDABILITY OF ENTAILMENT FOR GSL). *Let $\phi, \psi \in \text{GSL}$ and $n \triangleq |\Phi| + |\phi| + |\psi|$. The entailment problem $\phi \models_{\Phi} \psi$ is decidable in time $2^{2^{O(n^2 \log(n))}}$.*

PROOF. If $\phi, \psi \in \text{GSL}$, then $\phi \wedge \neg \psi \in \text{GSL}$. The entailment $\phi \models_{\Phi} \psi$ is valid iff $\phi \wedge \neg \psi$ is unsatisfiable. Since 2ExpTime is closed under complement, the claim follows from Theorem 9.3. \square

EXAMPLE 9.5. *The entailments in Example 6.8 can be proven using our decision procedure.*

Finally, our decision procedure is also applicable to (quantifier-free) symbolic heaps over inductive predicate definitions of bounded treewidth, because these formulas are always guarded.¹⁴ Hence, we also obtain a tighter complexity bound for the original decidability result of Iosif et al. [2013]:

COROLLARY 9.6 (DECIDABILITY OF ENTAILMENT FOR SL_{btw}). *Let $\phi, \psi \in \text{SL}_{\text{btw}}$ be quantifier-free and $n \triangleq |\Phi| + |\phi| + |\psi|$. The entailment problem $\phi \models_{\Phi} \psi$ is decidable in time $2^{2^{O(n^2 \log(n))}}$.*

PROOF. Follows from Corollary 9.4, since every quantifier-free SL_{btw} formula is in **GSL**. \square

COROLLARY 9.7 (DECIDABILITY OVER VALUES WITH NULL-POINTER). *Let $\phi, \psi \in \text{GSL}$ with $\text{locs}(\phi) \cup \text{locs}(\psi) \subseteq \{\text{nil}\}$. Then, the satisfiability of ϕ resp. the entailment $\phi \models_{\Phi} \psi$ over $\text{Val} \triangleq \text{Loc} \cup \{\text{nil}\}$ is decidable in $2^{2^{O(n^2 \log(n))}}$ for $n \triangleq |\Phi| + |\phi|$ resp. $n \triangleq |\Phi| + |\phi| + |\psi|$.*

10 CONCLUSION

We have given a unified and revised presentation of the decision procedures developed in [Katelaan et al. 2019; Katelaan and Zuleger 2020] covering (1) the satisfiability of quantifier-free guarded separation logic and (2) the entailment problem of (quantifier-free) symbolic heaps over SIDs of bounded treewidth. In particular, we have established a 2EXPTIME upper bound for both problems. A corresponding lower bound has been proven recently [Echenim et al. 2020b]. Hence, we can conclude that our decision procedures have optimal computational complexity.

To the best of our knowledge, our decision procedure for **GSL** is the first decision procedure to support an **SL** fragment combining user-supplied inductive definitions, Boolean structure, and magic wands. We obtained an almost tight delineation between decidability and undecidability: We showed that any extension of **GSL** in which one of the guards is dropped, leads to undecidability.

In this article, we considered the quantifier-free fragment of **GSL**; quantifiers can only appear in SID rules. An interesting question for future research is to what extent quantifiers can also be admitted in **GSL** formulas without sacrificing decidability. This question is also of practical interest as quantifiers naturally appear in entailments obtained from verification condition generators.

We mention that recent follow-up work by Echenim et al. [2021] generalizes the decidability of the entailment problem for SL_{btw} by weakening the establishment requirement. The result employs an abstraction inspired by the type abstraction presented in this article. It is an interesting question whether this result can be lifted to guarded separation logic as well. Further, our undecidability results require an unbounded number of dangling pointers. While [Echenim et al. 2021] supports classes of structures with unbounded treewidth, the entailment needs only to be checked for so-called normal structures of bounded treewidth. It would be interesting to interpret this result in terms of the number of dangling pointers that need to be considered in order to understand whether a bounded number of dangling pointers is fundamental for decidability.

Finally, apart from implementing our decision procedure, it would be interesting whether one can extract a proof certificate from our type abstraction that can also be checked independently by other proof systems based on separation logic.

REFERENCES

- Timos Antonopoulos, Nikos Gorogiannis, Christoph Haase, Max I. Kanovich, and Joël Ouaknine. 2014. Foundations for Decision Problems in Separation Logic with General Inductive Predicates. In *FOSSACS*. 411–425.
- Andrew W. Appel. 2014. *Program Logics - for Certified Compilers*. Cambridge University Press.

¹⁴Notice that only the formulas in the entailment query need to be quantifier-free; quantifiers are permitted in inductive definitions. Since we can use an arbitrary number of free variables at the top-level, this is a mild restriction.

- Yehoshua Bar-Hillel, Micha Perles, and Eli Shamir. 1961. On formal properties of simple phrase structure grammars. *Sprachtypologie und Universalienforschung* 14 (1961), 143–172.
- Kevin Batz, Benjamin Lucien Kaminski, Joost-Pieter Katoen, Christoph Matheja, and Thomas Noll. 2019. Quantitative separation logic: a logic for reasoning about probabilistic pointer programs. *PACMPL* 3, POPL (2019), 34:1–34:29.
- Josh Berdine, Cristiano Calcagno, Byron Cook, Dino Distefano, Peter W. O’Hearn, Thomas Wies, and Hongseok Yang. 2007. Shape Analysis for Composite Data Structures. In *CAV*. 178–192.
- Josh Berdine, Cristiano Calcagno, and Peter W. O’Hearn. 2004. A Decidable Fragment of Separation Logic. In *FSTTCS*. 97–109.
- Josh Berdine, Cristiano Calcagno, and Peter W. O’Hearn. 2005a. Smallfoot: Modular Automatic Assertion Checking with Separation Logic. In *FMCO*. 115–137.
- Josh Berdine, Cristiano Calcagno, and Peter W. O’Hearn. 2005b. Symbolic Execution with Separation Logic. In *APLAS*. 52–68.
- Josh Berdine, Byron Cook, and Samin Ishtiaq. 2011. SLayer: Memory Safety for Systems-Level Code. In *CAV*. 178–183.
- Stefan Blom and Marieke Huisman. 2015. Witnessing the elimination of magic wands. *Int. J. Softw. Tools Technol. Transf.* 17, 6 (2015), 757–781.
- Rémi Brochenin, Stéphane Demri, and Étienne Lozes. 2012. On the almighty wand. *Inf. Comput.* 211 (2012), 106–137.
- James Brotherston. 2007. Formalised Inductive Reasoning in the Logic of Bunched Implications. In *SAS (LNCS, Vol. 4634)*, Hanne Riis Nielson and Gilberto Filé (Eds.). Springer, 87–103.
- James Brotherston, Dino Distefano, and Rasmus Lerche dahl Petersen. 2011. Automated Cyclic Entailment Proofs in Separation Logic. In *CADE-23*. 131–146.
- James Brotherston, Carsten Fuhs, Juan Antonio Navarro Pérez, and Nikos Gorogiannis. 2014. A decision procedure for satisfiability in separation logic with inductive predicates. In *CSL-LICS*. 25:1–25:10.
- Cristiano Calcagno and Dino Distefano. 2011. Infer: An Automatic Program Verifier for Memory Safety of C Programs. In *NFM*. 459–465.
- Cristiano Calcagno, Dino Distefano, Jérémy Dubreil, Dominik Gabi, Pieter Hooimeijer, Martino Luca, Peter W. O’Hearn, Irene Papakonstantinou, Jim Purbrick, and Dul ma Rodriguez. 2015. Moving Fast with Software Verification. In *NFM*. 3–11.
- Cristiano Calcagno, Dino Distefano, Peter W. O’Hearn, and Hongseok Yang. 2006. Beyond Reachability: Shape Abstraction in the Presence of Pointer Arithmetic. In *SAS*. 182–203.
- Cristiano Calcagno, Dino Distefano, Peter W. O’Hearn, and Hongseok Yang. 2011. Compositional Shape Analysis by Means of Bi-Abduction. *J. ACM* 58, 6 (2011), 26:1–26:66.
- Cristiano Calcagno, Peter W. O’Hearn, and Hongseok Yang. 2007. Local Action and Abstract Separation Logic. In *LICS*. 366–378.
- Cristiano Calcagno, Hongseok Yang, and Peter W. O’Hearn. 2001. Computability and Complexity Results for a Spatial Assertion Language for Data Structures. In *APLAS*. 289–300.
- Wei-Ngan Chin, Cristina David, Huu Hai Nguyen, and Shengchao Qin. 2012. Automated verification of shape, size and bag properties via user-defined predicates in separation logic. *Sci. Comput. Program.* 77, 9 (2012), 1006–1036.
- Byron Cook, Christoph Haase, Joël Ouaknine, Matthew J. Parkinson, and James Worrell. 2011. Tractable Reasoning in a Fragment of Separation Logic. In *CONCUR*. 235–249.
- Bruno Courcelle and Joost Engelfriet. 2012. *Graph Structure and Monadic Second-Order Logic - A Language-Theoretic Approach*. Encyclopedia of mathematics and its applications, Vol. 138. Cambridge University Press.
- Patrick Cousot and Radhia Cousot. 1979. Constructive versions of Tarski’s fixed point theorems. *Pacific journal of Mathematics* 82, 1 (1979), 43–57.
- Reinhard Diestel. 2016. *Graph Theory, 5th Edition*. Graduate texts in mathematics, Vol. 173. Springer.
- Mnacho Echenim, Radu Iosif, and Nicolas Peltier. 2020a. The Bernays-Schönfinkel-Ramsey Class of Separation Logic with Uninterpreted Predicates. *ACM Trans. Comput. Log.* 21, 3 (2020), 19:1–19:46.
- Mnacho Echenim, Radu Iosif, and Nicolas Peltier. 2020b. Entailment Checking in Separation Logic with Inductive Definitions is 2-EXPTIME hard. *73* (2020), 191–211.
- Mnacho Echenim, Radu Iosif, and Nicolas Peltier. 2021. Decidable Entailments in Separation Logic with Inductive Definitions: Beyond Establishment. In *CSL (LIPICS)*.
- Constantin Enea, Ondrej Lengál, Mihaela Sighireanu, and Tomáš Vojnar. 2017. SPEN: A Solver for Separation Logic. In *NFM*. 302–309.
- Mihaela Sighireanu et al. 2019. SL-COMP: Competition of Solvers for Separation Logic. In *TACAS*. 116–132.
- Alexey Gotsman, Josh Berdine, Byron Cook, and Mooly Sagiv. 2007. Thread-modular shape analysis. In *PLDI*. 266–277.
- John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. 2007. *Introduction to automata theory, languages, and computation, 3rd Edition*. Addison-Wesley.
- Radu Iosif, Adam Rogalewicz, and Jiri Simáček. 2013. The Tree Width of Separation Logic with Recursive Definitions. In *CADE-24*. 21–38.
- Radu Iosif, Adam Rogalewicz, and Tomáš Vojnar. 2014. Deciding Entailments in Inductive Separation Logic with Tree Automata. In *ATVA*. 201–218.
- Samin S. Ishtiaq and Peter W. O’Hearn. 2001. BI as an Assertion Language for Mutable Data Structures. In *POPL*. 14–26.
- Bart Jacobs, Jan Smans, Pieter Philippaerts, Frédéric Vogels, Willem Penninckx, and Frank Piessens. 2011. VeriFast: A Powerful, Sound, Predictable, Fast Verifier for C and Java. In *NFM*. 41–55.

- Christina Jansen, Jens Katelaan, Christoph Matheja, Thomas Noll, and Florian Zuleger. 2017. Unified Reasoning About Robustness Properties of Symbolic-Heap Separation Logic. In *ESOP*. 611–638.
- Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.* 28 (2018), e20.
- Jens Katelaan, Christoph Matheja, and Florian Zuleger. 2019. Effective Entailment Checking for Separation Logic with Inductive Definitions. In *TACAS*. 319–336.
- Jens Katelaan and Florian Zuleger. 2020. Beyond Symbolic Heaps: Deciding Separation Logic With Inductive Definitions. In *LPAR (EPIc Series in Computing, Vol. 73)*. EasyChair, 390–408.
- Quang Loc Le, Makoto Tatsuta, Jun Sun, and Wei-Ngan Chin. 2017. A Decidable Fragment in Separation Logic with Inductive Predicates and Arithmetic. In *CAV*. 495–517.
- Christoph Matheja. 2020. *Automated reasoning and randomization in separation logic*. Dissertation. RWTH Aachen University.
- Peter Müller, Malte Schwerhoff, and Alexander J. Summers. 2017. Viper: A Verification Infrastructure for Permission-Based Reasoning. In *Dependable Software Systems Engineering*. 104–125.
- Jens Pagel, Christoph Matheja, and Florian Zuleger. 2020. Complete Entailment Checking for Separation Logic with Inductive Definitions. *CoRR* abs/2002.01202 (2020). arXiv:2002.01202
- Ruzica Piskac, Thomas Wies, and Damien Zufferey. 2013. Automating Separation Logic Using SMT. In *CAV*. 773–789.
- Ruzica Piskac, Thomas Wies, and Damien Zufferey. 2014a. Automating Separation Logic with Trees and Data. In *CAV*. 711–728.
- Ruzica Piskac, Thomas Wies, and Damien Zufferey. 2014b. GRASShopper - Complete Heap Verification with Mixed Specifications. In *TACAS*. 124–139.
- John C. Reynolds. 2002. Separation Logic: A Logic for Shared Mutable Data Structures. In *LICS*. 55–74.
- Malte Schwerhoff and Alexander J. Summers. 2015. Lightweight Support for Magic Wands in an Automatic Verifier. In *ECOOP*. 614–638.
- Quang-Trung Ta, Ton Chanh Le, Siau-Cheng Khoo, and Wei-Ngan Chin. 2016. Automated Mutual Explicit Induction Proof in Separation Logic. In *FM (LNCS, Vol. 9995)*. 659–676.
- Quang-Trung Ta, Ton Chanh Le, Siau-Cheng Khoo, and Wei-Ngan Chin. 2018. Automated lemma synthesis in symbolic-heap separation logic. *PACMPL* 2, POPL (2018), 9:1–9:29.
- Aditya V. Thakur, Jason Breck, and Thomas W. Reps. 2014. Satisfiability modulo abstraction for separation logic with linked lists. In *SPIN*. 58–67.
- Hongseok Yang. 2001. *Local Reasoning for Stateful Programs*. Ph. D. Dissertation. University of Illinois at Urbana-Champaign, Champaign, IL, USA. Advisor(s) Reddy, Uday S. AAI3023240.

A ELECTRONIC APPENDIX

A.1 Formal definition of substitution

For $\mathbf{y} = \langle y_1, \dots, y_k \rangle$ and $\mathbf{z} = \langle z_1, \dots, z_k \rangle$, the substitution $\phi[\mathbf{y}/\mathbf{z}]$ is defined by the table below. Since quantified variables can be renamed before performing a substitution, we assume w.l.o.g. that \mathbf{y} contains no variables that are bound by a quantifier in ϕ .

ϕ	$\phi[\mathbf{y}/\mathbf{z}]$	
y_i	z_i	$1 \leq i \leq k$
u	u	$u \notin \mathbf{y}$
$\langle v_1, \dots, v_n \rangle$	$\langle v_1[\mathbf{y}/\mathbf{z}], \dots, v_n[\mathbf{y}/\mathbf{z}] \rangle$	
emp	emp	
$u \approx v$	$u[\mathbf{y}/\mathbf{z}] \approx v[\mathbf{y}/\mathbf{z}]$	
$u \not\approx v$	$u[\mathbf{y}/\mathbf{z}] \not\approx v[\mathbf{y}/\mathbf{z}]$	
$u \mapsto \mathbf{v}$	$u[\mathbf{y}/\mathbf{z}] \mapsto \mathbf{v}[\mathbf{y}/\mathbf{z}]$	
$\text{pred}(\mathbf{x})$	$\text{pred}(\mathbf{x}[\mathbf{y}/\mathbf{z}])$	
$\neg\psi$	$\neg(\psi[\mathbf{y}/\mathbf{z}])$	
$\psi \oplus \theta$	$\psi[\mathbf{y}/\mathbf{z}] \oplus \theta[\mathbf{y}/\mathbf{z}]$	$\oplus \in \{\wedge, \vee, \star, \neg\star\}$
$\exists x. \psi$	$\exists x. \psi[\mathbf{y}/\mathbf{z}]$	$x \notin \mathbf{y}$
$\forall x. \psi$	$\forall x. \psi[\mathbf{y}/\mathbf{z}]$	$x \notin \mathbf{y}$

A.2 Proof of Lemma 4.4

Claim. For all predicates $\text{pred} \in \mathbf{Preds}(\Phi)$ and all states $\langle \mathfrak{s}, \mathfrak{h} \rangle$, we have

$$\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{pred}(\mathbf{x}) \quad \text{implies} \quad \langle \mathfrak{s}, \mathfrak{h} \rangle \in \mathbf{GStates}.$$

PROOF. Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state such that $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{pred}(\mathbf{x})$ for some predicate $\text{pred} \in \mathbf{Preds}(\Phi)$. The proof proceeds by strong mathematical induction on the number of rule applications needed to establish $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{pred}(\mathbf{x})$:

According to the semantics, there is a rule $(\text{pred}(\mathbf{x}) \Leftarrow \phi) \in \Phi$, for some $\phi = \exists e. \phi'$ with $\phi' = (y \mapsto z) \star \text{pred}_1(z_1) \star \dots \star \text{pred}_k(z_k) \star \Pi$, Π pure (note that for $k = 0$ there are no recursive rule applications). Because of $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{pred}(\mathbf{x})$, there exists a stack $\mathfrak{s}' = \mathfrak{s}[e/\mathbf{v}]$ for some locations \mathbf{v} such that $\langle \mathfrak{s}', \mathfrak{h} \rangle \models \phi'$. Thus, there are heaps $\mathfrak{h}_0, \mathfrak{h}_1, \dots, \mathfrak{h}_k$ with $\mathfrak{h} = \mathfrak{h}_0 \uplus \mathfrak{h}_1 \uplus \dots \uplus \mathfrak{h}_k$, $\mathfrak{h}_0 \models (y \mapsto z)$ and $\langle \mathfrak{s}', \mathfrak{h}_i \rangle \models \text{pred}_i(z_i)$ for all $1 \leq i \leq k$. By I.H., we have $\langle \mathfrak{s}', \mathfrak{h}_i \rangle \in \mathbf{GStates}$. Hence, $\text{dangling}(\mathfrak{h}_0) \subseteq \text{img}(\mathfrak{s}')$ for all $0 \leq i \leq k$. By the definition of establishment (see Section 3.4.3), we have $\mathfrak{s}'(e) \subseteq \text{dom}(\mathfrak{h}) \cup \mathfrak{s}(\mathbf{x}) \cup \{0\}$. Because of $\mathfrak{h}_i \subseteq \mathfrak{h}$ and $\text{dangling}(\mathfrak{h}) = \text{locs}(\mathfrak{h}) \setminus \text{dom}(\mathfrak{h})$ we get that $\text{dangling}(\mathfrak{h}_i) \subseteq \mathfrak{s}(\mathbf{x})$. Hence, $\langle \mathfrak{s}, \mathfrak{h} \rangle \in \mathbf{GStates}$ because

$$\begin{aligned} \text{dangling}(\mathfrak{h}) &= \text{dangling}(\mathfrak{h}_0 \uplus \mathfrak{h}_1 \uplus \dots \uplus \mathfrak{h}_k) \\ &\subseteq \text{dangling}(\mathfrak{h}_0) \cup \text{dangling}(\mathfrak{h}_1) \cup \dots \cup \text{dangling}(\mathfrak{h}_k) \\ &\subseteq \mathfrak{s}(\mathbf{x}) = \text{img}(\mathfrak{s}). \end{aligned} \quad \square$$

A.3 Proof of Corollary 4.6

Claim. For all $\phi \in \mathbf{GSL}$ and all states $\langle \mathfrak{s}, \mathfrak{h} \rangle$, we have

$$\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi \quad \text{implies} \quad \langle \mathfrak{s}, \mathfrak{h} \rangle \in \mathbf{GStates}.$$

PROOF. Let $\phi \in \mathbf{GSL}$ be a guarded formula and let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state with $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$. The proof proceeds by structural induction on ϕ :

Case $\phi = \mathbf{emp}$, $\phi = x \approx y$, $x \not\approx y$: Clearly, there are no dangling pointers in the empty heap.

Case $\phi = x \mapsto y$. Immediate because of $\text{dangling}(\mathfrak{h}) \subseteq \mathfrak{s}(y) \subseteq \text{img}(\mathfrak{s})$.

Case $\phi = \text{pred}(\mathbf{x})$. By Lemma 4.4.

Case $\phi = \phi_1 \star \phi_2$. Since $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$, there exist heaps $\mathfrak{h}_1, \mathfrak{h}_2$ with $\mathfrak{h} = \mathfrak{h}_1 \uplus \mathfrak{h}_2$ and $\langle \mathfrak{s}, \mathfrak{h}_i \rangle \models_{\Phi} \phi_i$. By I.H., we have $\langle \mathfrak{s}, \mathfrak{h}_i \rangle \in \mathbf{GStates}$, i.e., $\text{dangling}(\mathfrak{h}_i) \subseteq \text{img}(\mathfrak{s})$. Thus, $\text{dangling}(\mathfrak{h}) = \text{dangling}(\mathfrak{h}_1 \uplus \mathfrak{h}_2) \subseteq \text{dangling}(\mathfrak{h}_1) \cup \text{dangling}(\mathfrak{h}_2) \subseteq \text{img}(\mathfrak{s})$. Hence, $\langle \mathfrak{s}, \mathfrak{h} \rangle \in \mathbf{GStates}$.

Case $\phi = \phi_1 \wedge \phi_2$. By the semantics of \wedge , this in particular means $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi_1$. By I.H., it then follows that $\langle \mathfrak{s}, \mathfrak{h} \rangle \in \mathbf{GStates}$. Notice that this case covers all standard conjunctions including the guarded negation, the guarded magic wand, and the guarded septraction.

Case $\phi = \phi_1 \vee \phi_2$. Assume w.l.o.g. that $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi_1$. By I.H., we have $\langle \mathfrak{s}, \mathfrak{h} \rangle \in \mathbf{GStates}$. \square

A.4 Proof of Lemma 4.5

Claim. Let $\phi \in \mathbf{GSL}$ be a guarded formula with $\text{fvars}(\phi) = \mathbf{x}$. Then, for every state $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$, there are predicates $\text{pred}_i \in \mathbf{Preds}(\Phi)$ and variables $\mathbf{z}_i \subseteq \mathbf{x}$ such that $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \star_{1 \leq i \leq k} \text{pred}_i(\mathbf{z}_i)$.

PROOF. By structural induction on ϕ .

Case $\phi = \mathbf{emp}$. This case is immediate because, for $k = 0$, \mathbf{emp} coincides with $\star_{1 \leq i \leq k} \dots$

Case $\phi = x \approx y$. Since, in our semantics, the equality $x \approx y$ entails \mathbf{emp} , we immediately obtain $\langle \mathfrak{s}, \mathfrak{h} \rangle \in \mathbf{GStates}$. The case for disequalities $x \not\approx y$ is analogous.

Case $\phi = x \mapsto y$. Since Φ is pointer-closed, there exists a predicate $\text{pred} \in \Phi$ such that $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{pred}(x \cdot y)$.

Case $\phi = \text{pred}(\mathbf{x})$. Clearly, the claim holds.

Case $\phi = \phi_1 \star \phi_2$. Since $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$, there exist domain-disjoint heaps $\mathfrak{h}_1, \mathfrak{h}_2$ with $\mathfrak{h} = \mathfrak{h}_1 \cup \mathfrak{h}_2$ such that $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle \models_{\Phi} \phi_1$ and $\langle \mathfrak{s}, \mathfrak{h}_2 \rangle \models_{\Phi} \phi_2$. By the induction hypothesis, there exist predicate calls $\text{pred}_{1,1}(\mathbf{x}_1), \dots, \text{pred}_{1,m}(\mathbf{x}_m)$ and $\text{pred}_{2,1}(\mathbf{y}_1), \dots, \text{pred}_{2,n}(\mathbf{y}_n)$ such that

$$\langle \mathfrak{s}, \mathfrak{h}_1 \rangle \models_{\Phi} \star_{1 \leq i \leq m} \text{pred}_{1,i}(\mathbf{x}_i) \quad \text{and} \quad \langle \mathfrak{s}, \mathfrak{h}_2 \rangle \models_{\Phi} \star_{1 \leq j \leq n} \text{pred}_{2,j}(\mathbf{y}_j).$$

The semantics of the separating conjunction \star and the fact $\mathfrak{h} = \mathfrak{h}_1 \cup \mathfrak{h}_2$ then yield

$$\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \star_{1 \leq i \leq m} \text{pred}_{1,i}(\mathbf{x}_i) \star \star_{1 \leq j \leq n} \text{pred}_{2,j}(\mathbf{y}_j).$$

Case $\phi = \phi_1 \wedge \phi_2$. By the semantics of \wedge , this in particular means $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi_1$. By the induction hypothesis, the claim then holds for $\langle \mathfrak{s}, \mathfrak{h} \rangle$ and ϕ_1 . Notice that this case covers all standard conjunctions including the guarded negation, the guarded magic wand, and the guarded septraction.

Case $\phi = \phi_1 \vee \phi_2$. Assume w.l.o.g. that $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi_1$. By the induction hypothesis, the claim then holds for $\langle \mathfrak{s}, \mathfrak{h} \rangle$ and ϕ_1 . \square

A.5 Proof of Lemma 5.8

Claim (Completeness of the encoding). Let $\mathbf{G} = \langle \mathbf{N}, \mathbf{T}, \mathbf{R}, \mathbf{S} \rangle$ and let Φ be the corresponding SID encoding. Let $1 \leq i \leq 2$, $x_1, x_2, x_3 \in \mathbf{Var}$, and let $w \in \mathcal{L}(\mathbf{G})$. Then there exists a model $\langle \mathfrak{s}, \mathfrak{h} \rangle$ of $\mathbf{S}(x_1, x_2, x_3)$ with $\text{wordof}_{\mathbf{S}}(\mathfrak{s}, \mathfrak{h}, x_2, x_3) = w$.

PROOF. We show the stronger claim that, for all $x_1, x_2, x_3 \in \mathbf{Var}$, $w \in \mathcal{L}(\mathbf{G})$, and $N \in \mathbf{N}$, if $N \Rightarrow^+ w$, then there exists a model $\langle \mathfrak{s}, \mathfrak{h} \rangle$ of $N(x_1, x_2, x_3)$ with $\text{wordof}_N(\mathfrak{s}, \mathfrak{h}, x_2, x_3) = w$. We proceed by mathematical induction on the number m of \Rightarrow steps in a (minimal-length) derivation $N \Rightarrow^+ w$.

If $m = 1$, $w = a_i$ for some $1 \leq i \leq n$ and there exists a rule $N \rightarrow a_i$. Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a model of $\exists a. (x_1 \mapsto \langle x_3, a \rangle) \star \text{letter}_i(a) \star x_1 \approx x_2$. Note that this is a rule of the predicate N , so it holds that $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} N(x_1, x_2, x_3)$. Moreover, $\text{wordof}_N(\mathfrak{s}, \mathfrak{h}, x_2, x_3) = a_i$.

If $m > 1$, there exists a rule $N \rightarrow AB$ such that $N \Rightarrow AB \Rightarrow^+ w$. Then, there exist words w_A, w_B with $w = w_A \cdot w_B$, $A \Rightarrow^+ w_A$, and $B \Rightarrow^+ w_B$.

Observe that both of the above derivations consist of strictly fewer than m steps. Now, fix some variables $l, m, r \in \text{Var}$. By I.H., there exist states (s_1, h_1) and (s_2, h_2) such that

- $\langle s_1, h_1 \rangle \models_{\Phi} A(l, x_2, m)$ and $\text{wordof}_A(s_1, h_1, x_2, m) = w_A$ as well as
- $\langle s_2, h_2 \rangle \models_{\Phi} B(r, m, x_3)$ and $\text{wordof}_B(s_2, h_2, m, x_3) = w_B$.

Assume w.l.o.g. that (1) $\text{dom}(s_1) \cap \text{dom}(s_2) = m$, (2) $s_1(m) = s_2(m)$, (3) and $h_1 \uplus h_2 \neq \perp$; if this is not the case, simply replace $\langle s_1, h_1 \rangle$ and $\langle s_2, h_2 \rangle$ with appropriate isomorphic models. We choose some location $k \in \text{Loc}$ such that $k \notin \text{locs}(h_1 \uplus h_2)$.

Let $s \triangleq s_1 \cup s_2 \cup \{x_1 \mapsto k\}$ and $h \triangleq h_1 \cup h_2 \cup \{k \mapsto \langle s(l), s(r) \rangle\}$. We obtain $\langle s, h \rangle \models_{\Phi} (x_1 \mapsto \langle l, r \rangle) \star A(l, x_2, m) \star B(r, m, x_3)$ and thus also $\langle s, h \rangle \models_{\Phi} \exists \langle l, r, m \rangle . (x_1 \mapsto \langle l, r \rangle) \star A(l, x_2, m) \star B(r, m, x_3)$. By definition of Φ , we conclude $\langle s, h \rangle \models_{\Phi} N(x_1, x_2, x_3)$. Furthermore, observe that

$$\begin{aligned} \text{wordof}_N(s, h, x_2, x_3) &= \text{wordof}_A(s, h_1, x_2, m) \cdot \text{wordof}_B(s, h_2, m, x_3) \\ &= w_A \cdot w_B = w. \end{aligned} \quad \square$$

A.6 Proof of Lemma 5.9

Claim (Soundness of the encoding). Let $G = \langle N, T, R, S \rangle$ and let Φ be the corresponding SID encoding. Let $x_1, x_2, x_3 \in \text{Var}$ and let $\langle s, h \rangle \models_{\Phi} S(x_1, x_2, x_3)$. Then $\text{wordof}_S(s, h, x_2, x_3) \in \mathcal{L}(G)$.

PROOF. We show the stronger claim that for all $x_1, x_2, x_3 \in \text{Var}$, all models $\langle s, h \rangle$, and all $N \in N$, if $\langle s, h \rangle \models_{\Phi} N(x_1, x_2, x_3)$ then $N \Rightarrow^+ \text{wordof}_N(s, h, x_2, x_3)$. Observe that h is a tree overlaid with a linked list. We proceed by mathematical induction on the height h of the tree in h .

If $h = 0$, then Φ contains a rule

$$N(x_1, x_2, x_3) \Leftarrow \exists a. (x_1 \mapsto \langle x_3, a \rangle) \star \text{letter}_k(a) \star x_1 \approx x_2$$

whose right-hand side is satisfied by $\langle s, h \rangle$. Then $\text{wordof}_N(s, h, x_2, x_3) = a_k$. By definition of Φ , this implies $N \rightarrow a_k \in R_1 \cup R_2$ and, consequently, $N \Rightarrow a_k$. Hence, $N \Rightarrow^+ a_k = \text{wordof}_N(s, h, x_2, x_3)$.

If $h > 0$, there exists a rule $(N(x_1, x_2, x_3) \Leftarrow \psi) \in \Phi$ such that $\langle s, h \rangle \models_{\Phi} \psi$, where ψ is of the form

$$\exists l, r, m. (x_1 \mapsto \langle l, r \rangle) \star A(l, x_2, m) \star B(r, m, x_3).$$

Recall that by definition of Φ , we have $N \rightarrow AB \in R_1 \cup R_2 \quad (\dagger)$.

By the semantics of \exists and \star there then are a stack s' with $\text{dom}(s') = \text{dom}(s) \cup \{l, r, m\}$ and heaps h_0, h_A, h_B such that $h = h_0 \uplus h_A \uplus h_B$, $\langle s', h_0 \rangle \models_{\Phi} (x_1 \mapsto \langle l, r \rangle)$, $\langle s', h_A \rangle \models_{\Phi} A(l, x_2, m)$, and $\langle s', h_B \rangle \models_{\Phi} B(r, m, x_3)$. Note that the height of the trees in h_A and h_B is at most $h - 1$, so we can apply the induction hypotheses for these models to obtain

$$A \Rightarrow^+ \text{wordof}_A(s', h_1, x_2, m) \quad \text{and} \quad B \Rightarrow^+ \text{wordof}_B(s', h_2, m, x_3).$$

Together with (\dagger) , we derive

$$\begin{aligned} N &\Rightarrow AB \\ &\Rightarrow^+ \text{wordof}_A(s', h_1, x_2, m) \cdot \text{wordof}_B(s', h_2, m, x_3) \\ &= \text{wordof}_N(s', h, x_2, x_3) \\ &= \text{wordof}_N(s, h, x_2, x_3). \end{aligned} \quad \square$$

A.7 Proof of Theorem 5.10

Claim. The satisfiability problem for the fragment $\text{SL}_{\text{btw}}(\wedge, \star, \text{t})$ is undecidable.

PROOF. Let Φ be the encoding of the CFGs $G_1 = \langle N_1, T, R_1, S_1 \rangle$ and $G_2 = \langle N_2, T, R_2, S_2 \rangle$ as described in Section 5.1. Moreover, consider the $SL_{btw}(\wedge, \star, t)$ formula

$$\phi \triangleq (S_1(a, x, y) \star t) \wedge (S_2(b, x, y) \star t).$$

We claim that ϕ is satisfiable iff $\mathcal{L}(G_1) \cap \mathcal{L}(G_2) \neq \emptyset$; both implications are proven separately:

If ϕ is satisfiable, there exists a state $\langle s, h \rangle$ with $\langle s, h \rangle \models_{\Phi} \phi$. By Lemma 5.6, there exist heaps $h_{w_1}, h_{w_2} \subseteq h$ such that $\text{wordof}_{S_i}(s, h, x, y) = \text{letters}(s, h_{w_i}, x, y)$ for $i \in \{1, 2\}$.

Observe that both $\langle s, h_{w_1} \rangle \models_{\Phi} \text{word}(x, y)$ and $\langle s, h_{w_2} \rangle \models_{\Phi} \text{word}(x, y)$. Hence, $h_{w_1} = h_{w_2}$ and thus

$$w \triangleq \text{wordof}_{S_2}(s, h, x, y) = \text{wordof}_{S_1}(s, h, x, y).$$

By Lemma 5.9, we have $w \in \mathcal{L}(G_1)$ and $w \in \mathcal{L}(G_2)$, i.e., $w \in \mathcal{L}(G_1) \cap \mathcal{L}(G_2)$.

Conversely, assume $\mathcal{L}(G_1) \cap \mathcal{L}(G_2) \neq \emptyset$. Then there exists a word $w \in \mathcal{L}(G_1) \cap \mathcal{L}(G_2)$. By Lemma 5.8, there exist states $\langle s, h_1 \rangle, \langle s, h_2 \rangle$ with $\langle s, h_1 \rangle \models_{\Phi} S_1(a, x, y)$ and $\langle s, h_2 \rangle \models_{\Phi} S_2(b, x, y)$. Let $h_{w_1} \subseteq h_1, h_{w_2} \subseteq h_2$ be the unique heaps with $\text{wordof}_{S_1}(s, h_1, x, y) = \text{letters}(s, h_{w_1}, x, y) = w = \text{letters}(s, h_{w_2}, x, y) = \text{wordof}_{S_2}(s, h_2, x, y)$.

Observe that $\langle s, h_{w_1} \rangle \cong \langle s, h_{w_2} \rangle$ (see Definition 3.4). Consequently, we can reason about an isomorphic state in which we replace h_2 by a heap that contains h_{w_1} (as opposed to h_{w_2}) as sub-heap and is otherwise disjoint from h_1 . That is, there exists a heap h'_2 such that $\langle s, h_2 \rangle \cong \langle s, h'_2 \rangle$, $\text{locs}(h'_2) \cap \text{locs}(h_1) = \text{locs}(h_{w_1})$, and $\text{wordof}_{S_2}(s, h'_2, x, y) = \text{letters}(s, h_{w_1}, x, y) = w$. In particular, $\langle s, h'_2 \rangle \models_{\Phi} S_2(b, x, y)$, because isomorphic states satisfy the same formulas (Lemma 3.5). Now let $h \triangleq h_1 \cup h'_2$ be the (non-disjoint) union of h_1 and h'_2 . Since $h_1 \subseteq h$ and $\langle s, h_1 \rangle \models_{\Phi} S_1(a, x, y)$, we have $\langle s, h \rangle \models_{\Phi} S_1(a, x, y) \star t$; and similarly, since $h'_2 \subseteq h$ and $\langle s, h'_2 \rangle \models_{\Phi} S_2(a, x, y)$, we have that $\langle s, h \rangle \models_{\Phi} S_2(b, x, y)$. Consequently, $\langle s, h \rangle \models_{\Phi} \phi$. \square

A.8 Proof of Lemma 5.13

Claim. Let $G_2 = \langle N_2, T, R_2, S_2 \rangle$ be the CFG fixed in Section 5.1. Moreover, let Φ be the corresponding SID encoding, $\text{word}_2(x, y) \triangleq (\text{word}(x, y) \oplus S_2(a, x, y)) \oplus S_2(a, x, y)$, and let $\langle s, h \rangle$ be a state. Then $\langle s, h \rangle \models_{\Phi} \text{word}_2(x, y)$ iff $\langle s, h \rangle \models_{\Phi} \text{word}(x, y)$ and $\text{letters}(s, h, x, y) \in \mathcal{L}(G_2)$.

PROOF. Assume $\langle s, h \rangle \models_{\Phi} \text{word}_2(x, y)$. By the semantics of \oplus , there exists a heap h_1 with $\langle s, h_1 \rangle \models_{\Phi} \text{word}(x, y) \oplus S_2(a, x, y)$ such that $\langle s, h \uplus h_1 \rangle \models_{\Phi} S_2(a, x, y)$. Observe that h_1 contains precisely the inner nodes of $\langle s, h \uplus h_1 \rangle$, i.e., everything *except* the part of the state that induces the word. Consequently, h is the part of the state that induces the word, i.e., $\text{wordof}_{S_2}(s, h \uplus h_1, x, y) = \text{letters}(s, h, x, y)$ and $\langle s, h \rangle \models_{\Phi} \text{word}(x, y)$. Lemma 5.9 then yields $\text{letters}(s, h, x, y) \in \mathcal{L}(G_2)$.

Conversely, assume a state $\langle s, h \rangle$ be such that $w \triangleq \text{letters}(s, h, x, y) \in \mathcal{L}(G_2)$. As a consequence of Lemma 5.8, there exists a heap h_1 with $\langle s, h \uplus h_1 \rangle \models_{\Phi} S_2(a, x, y)$. Because $\langle s, h \rangle \models_{\Phi} \text{word}(x, y)$ by assumption, the semantics of \oplus yields that $\langle s, h_1 \rangle \models_{\Phi} \text{word}(x, y) \oplus S_2(a, x, y)$. Because $\langle s, h \uplus h_1 \rangle \models_{\Phi} S_2(a, x, y)$, we obtain by the semantics of \oplus that $\langle s, h \rangle \models_{\Phi} (\text{word}(x, y) \oplus S_2(a, x, y)) \oplus S_2(a, x, y)$. \square

A.9 Proof of Theorem 5.14

Claim. The satisfiability problem of $SL_{btw}(\oplus)$ is undecidable.

PROOF. We claim that $\psi \triangleq \text{word}_2(x, y) \oplus S_1(a, x, y)$ is satisfiable iff $\mathcal{L}(G_1) \cap \mathcal{L}(G_2) \neq \emptyset$.

Assume ψ is satisfiable, i.e., there exists a state $\langle s, h \rangle \models_{\Phi} \psi$. By the semantics of \oplus , there exists a heap $h_0 \subseteq h$ with $h_0 \models_{\Phi} \text{word}_2(x, y)$ and $\langle s, h \uplus h_0 \rangle \models_{\Phi} S_1(a, x, y)$. As $\text{letters}(s, h_0, x, y) \in \mathcal{L}(G_2)$, by Lemma 5.13, we have that $\langle s, h_0 \rangle \models_{\Phi} \text{word}(x, y)$. It follows that h_0 is the unique sub-heap of $h \uplus h_0$ with $\text{wordof}_{S_1}(s, h \uplus h_0) = \text{letters}(s, h_0, x, y)$. By Lemma 5.9, $\text{letters}(s, h_0, x, y) \in \mathcal{L}(G_1)$. Together with Lemma 5.13, we thus have that $\text{letters}(s, h_0, x, y) \in \mathcal{L}(G_1) \cap \mathcal{L}(G_2)$.

Conversely, assume there exists a word $w \in \mathcal{L}(G_1) \cap \mathcal{L}(G_2)$. As shown in the proof of Theorem 5.10, there exist states $\langle s, h_1 \rangle, \langle s, h_2 \rangle, \langle s, h \rangle$ with $\langle s, h_1 \rangle \models_{\Phi} S_1(a, x, y)$, $\langle s, h_2 \rangle \models_{\Phi} S_2(b, x, y)$ and $\text{locs}(h_1) \cap \text{locs}(h_2) = \text{locs}(h)$ such that

$$\text{wordof}_{S_1}(s, h_1, x, y) = \text{wordof}_{S_2}(s, h_2, x, y) = \text{letters}(s, h, x, y) = w.$$

Let $h_0 \subseteq h_1$ be the sub-heap of h_1 with $h \uplus h_0 = h_1$. By Lemma 5.13, $\langle s, h \rangle \models_{\Phi} \text{word}_2(x, y)$. Consequently, $\langle s, h_0 \rangle \models_{\Phi} \psi$, i.e., ψ is satisfiable. \square

A.10 Proof of Lemma 7.10

Claim. For all $I \subseteq \text{Loc}$, every Φ -forest has a unique I -split $\text{split}(\mathfrak{f}, I)$.

PROOF. Let \mathfrak{f} be a Φ -forest with $\text{graph}(\mathfrak{f}) = \langle V_{\mathcal{G}}, E_{\mathcal{G}} \rangle$. Moreover, consider the graph

$$\mathcal{G} \triangleq \langle V_{\mathcal{G}}, E_{\mathcal{G}} \setminus \{ \langle a, b \rangle \mid a \in \text{Loc}, b \in I \} \rangle.$$

Since $\text{graph}(\mathfrak{f})$ is a forest and $\mathcal{G} \subseteq \text{graph}(\mathfrak{f})$, \mathcal{G} is a forest, i.e., all connected components C_1, \dots, C_k of \mathcal{G} are trees. Formally, let $\text{locs}(C_i)$ be all locations in C_i and let $\text{succ}_{C_i}(a)$ be the largest set of locations such that every edge in $\{a\} \times \text{succ}_{C_i}(a)$ appears in C_i . We then define:

$$\begin{aligned} t_i &\triangleq \{ a \mapsto \langle \text{succ}_{C_i}(a), \text{rule}_{\mathfrak{f}}(a) \rangle \mid a \in \text{locs}(C_i) \}, && \text{(tree induced by component } C_i) \\ \bar{\mathfrak{f}} &\triangleq \{ t_1, \dots, t_n \}. && \text{(\Phi-forest induced by the connected components)} \end{aligned}$$

By construction, the forest $\bar{\mathfrak{f}}$ is an I -split of \mathfrak{f} . Moreover, since every I -split must have the same domain and the same rule instances as \mathfrak{f} and because every connected component gives rise to a single Φ -tree, the I -split $\text{split}(\mathfrak{f}, I) = \bar{\mathfrak{f}}$ is unique. \square

A.11 Proof of Lemma 7.20

Let $\langle s, h \rangle \in \text{GStates}$ and ϕ be a quantifier free SL formula with $\langle s, h \rangle \models_{\Phi} \phi$. Moreover, let $v \in (\text{Loc} \setminus (\text{dom}(h) \cup \text{img}(s)))^*$ be a repetition-free sequence of locations. Then, for every set $\mathbf{a} \triangleq \{ a_1, \dots, a_{|v|} \}$ of fresh variables (i.e., $\mathbf{a} \cap \text{dom}(s) = \emptyset$), we have $\langle s, h \rangle \models_{\Phi} \forall \mathbf{a}. \phi[v/\mathbf{a}]$.

PROOF SKETCH. Let $w \in (\text{Loc} \setminus (\text{dom}(h) \cup \text{img}(s)))^*$ be a repetition-free sequence of locations with $|v| = |w|$. We note that $\text{dangling}(h) \subseteq \text{img}(s)$ because of $\langle s, h \rangle \in \text{GStates}$. Hence, neither v nor w intersect with $\text{locs}(h)$ or $\text{img}(s)$. Thus, it follows that $\langle s, h \rangle \models_{\Phi} \phi[v/w]$. Since w was arbitrary, $\langle s, h \rangle \models_{\Phi} \forall \mathbf{a}. \phi[v/\mathbf{a}]$ by the semantics of \forall . \square

A.12 Proof of Lemma 7.19

Claim. Let t be a Φ -tree. Then, $\langle _ , \text{heap}(t) \rangle \models_{\Phi} \text{project}^{\text{Loc}}(t)$ (where $_$ denotes an arbitrary stack).

PROOF. We prove the claim by mathematical induction on the height of t .

By construction, t has a root $r = \text{root}(t)$ with $m \geq 0$ successors that are the root of subtrees t_1, \dots, t_m . Hence, there is a rule instance $\text{rule}_t(r)$ (up to applying commutativity of \star) of the form¹⁵

$$\text{rootpred}(t) \Leftarrow (a \mapsto b) \star (\star_{1 \leq i \leq m} \text{rootpred}(t_i)) \star \star \text{holepred}_t(r),$$

By the semantics of \star and \rightarrow_{\star} , we have

$$\langle _ , \{ a \mapsto b \} \rangle \models_{\Phi} ((\star_{1 \leq i \leq m} \text{rootpred}(t_i)) \star \star \text{holepred}_t(r)) \rightarrow_{\star} \text{rootpred}(t)$$

¹⁵For $\text{height}(t) = 0$, there are no successors, i.e., $\star_{1 \leq i \leq m} \text{rootpred}(t_i)$ is equivalent to emp .

We apply the I.H. for each tree t_i and obtain

$$\langle _ , \text{heap}(t_i) \rangle \models_{\Phi} (\star \text{allholepreds}(t_i)) \rightarrow \star \text{rootpred}(t_i).$$

On the level of heaps, we have $\text{heap}(t) = \{a \mapsto b\} \uplus \text{heap}(t_1) \uplus \dots \uplus \text{heap}(t_m)$. Applying the semantics of \star and the definition of ψ_i then yields

$$\begin{aligned} \langle _ , \text{heap}(t) \rangle \models_{\Phi} & \left(((\star_{1 \leq i \leq m} \text{rootpred}(t_i)) \star \star \text{holepreds}_t(r)) \rightarrow \star \text{rootpred}(t) \right) \\ & \star \star_{1 \leq i \leq m} ((\star \text{allholepreds}(t_i)) \rightarrow \star \text{rootpred}(t_i)). \end{aligned}$$

Applying Lemma 7.18 m times, we then obtain

$$\begin{aligned} \langle _ , \text{heap}(t) \rangle \models_{\Phi} & ((\star_{1 \leq i \leq m} (\star \text{allholepreds}(t_i))) \star \star \text{holepreds}_t(r)) \rightarrow \star \text{rootpred}(t) \\ & = (\star \text{allholepreds}(t)) \rightarrow \star \text{rootpred}(t). \end{aligned} \quad (\text{Def. of allholepreds}(t))$$

□

A.13 Proof of Lemma 7.25

Claim (Soundness of stack-projection). Let $\langle s, h \rangle \in \text{GStates}$. Moreover, let \mathfrak{f} be a Φ -forest with $\text{heap}(\mathfrak{f}) = h$. Then, we have $\langle s, h \rangle \models_{\Phi} \text{project}(s, \mathfrak{f})$.

PROOF. Let $\mathfrak{f} = \{t_1, \dots, t_k\}$ and $\phi = \star_{1 \leq i \leq k} \text{project}^{\text{Loc}}(t_i)$. By Lemma 7.19, we know for each i that $\langle s, \text{heap}(t_i) \rangle \models_{\Phi} \text{project}^{\text{Loc}}(t_i)$. By definition, $\text{heap}(\mathfrak{f}) = \text{heap}(t_1) \uplus \dots \uplus \text{heap}(t_k)$. Applying the semantics of \star then yields $\langle s, h \rangle \models_{\Phi} \phi$ (†).

Let $w = \text{locs}(\phi) \cap (\text{dom}(\mathfrak{f}) \setminus \text{img}(s))$ be the locations that occur in ϕ and are allocated in $\text{heap}(\mathfrak{f})$ but are not the value of any stack variable, and let $v = \text{locs}(\phi) \setminus (\text{img}(s) \cup \text{dom}(\mathfrak{f}))$ be the locations that occur in the formula ϕ and are neither allocated nor the value of any stack variable.

Then, we have

$$\text{project}(s, \mathfrak{f}) = \exists e. \forall a. \phi[\text{dom}(s_{\text{max}}^{-1}) \cdot v \cdot w / \text{img}(s_{\text{max}}^{-1}) \cdot a \cdot e],$$

where $e \triangleq \langle e_1, e_2, \dots, e_{|w|} \rangle$ and $a \triangleq \langle a_1, a_2, \dots, a_{|v|} \rangle$ denote some disjoint sets of fresh variables.

The claim then follows by the implications below:

$$\begin{aligned} \langle s, h \rangle \models_{\Phi} \phi & \quad (\text{by } (\dagger)) \\ \implies \langle s, h \rangle \models_{\Phi} \phi[\text{dom}(s_{\text{max}}^{-1}) / \text{img}(s_{\text{max}}^{-1})] & \quad (\text{stack-heap semantics}) \\ \implies \langle s, h \rangle \models_{\Phi} \phi[\text{dom}(s_{\text{max}}^{-1}) / \text{img}(s_{\text{max}}^{-1})][v \cdot w / a \cdot e][a \cdot e / v \cdot w] & \quad (a \text{ and } e \text{ are disjoint sets of fresh variables}) \\ \implies \langle s, h \rangle \models_{\Phi} \phi[\text{dom}(s_{\text{max}}^{-1}) \cdot v \cdot w / \text{img}(s_{\text{max}}^{-1}) \cdot a \cdot e][a \cdot e / v \cdot w] & \quad (v \cap \text{img}(s) = \emptyset \text{ and } w \cap \text{img}(s) = \emptyset) \\ \implies \langle s, h \rangle \models_{\Phi} \forall a. \phi[\text{dom}(s_{\text{max}}^{-1}) \cdot v \cdot w / \text{img}(s_{\text{max}}^{-1}) \cdot a \cdot e][e / w] & \quad (\text{by Lemma 7.20}) \\ \implies \langle s, h \rangle \models_{\Phi} \exists e. \forall a. \phi[\text{dom}(s_{\text{max}}^{-1}) \cdot v \cdot w / \text{img}(s_{\text{max}}^{-1}) \cdot a \cdot e] & \quad (\text{semantics of } \exists) \\ \implies \langle s, h \rangle \models_{\Phi} \text{project}(s, \mathfrak{f}) & \quad \square \end{aligned}$$

A.14 Proof of Lemma 7.35

Before we prove Lemma 7.35, we need two auxiliary results.

LEMMA A.1. *Let s be a stack and let $\mathfrak{f}_1, \mathfrak{f}_2$ be Φ -forests with $\mathfrak{f}_1 \uplus \mathfrak{f}_2 \neq \perp$. Then, $\text{project}(s, \mathfrak{f}_1 \uplus \mathfrak{f}_2) \in \text{project}(s, \mathfrak{f}_1) \bullet_{\mathcal{P}} \text{project}(s, \mathfrak{f}_2)$.*

PROOF. We set $\mathfrak{f}_0 \triangleq \mathfrak{f}_1 \uplus \mathfrak{f}_2$. For $i \in \{0, 1, 2\}$, let $\phi_i = \star_{t \in \mathfrak{f}_i} \text{project}^{\text{Loc}}(t)$, let $\mathbf{w}_i = \text{locs}(\phi_i) \cap (\text{dom}(\mathfrak{f}_i) \setminus \text{img}(\mathfrak{s}))$ be the locations that occur in the formula ϕ and are allocated in $\text{heap}(\mathfrak{f}_i)$ but are not the value of any stack variable, and let $\mathbf{v}_i = \text{locs}(\phi_i) \setminus (\text{img}(\mathfrak{s}) \cup \text{dom}(\mathfrak{f}_i))$ be the locations that occur in the formula ϕ_i and are neither allocated nor the value of any stack variable. Then, we have

$$\text{project}(\mathfrak{s}, \mathfrak{f}_i) = \exists \mathbf{e}_i. \forall \mathbf{a}_i. \phi_i[\text{dom}(\mathfrak{s}_{\max}^{-1}) \cdot \mathbf{v}_i \cdot \mathbf{w}_i / \text{img}(\mathfrak{s}_{\max}^{-1}) \cdot \mathbf{a}_i \cdot \mathbf{e}_i],$$

where $\mathbf{e}_i \triangleq \langle e_1, e_2, \dots, e_{|\mathbf{w}_i|} \rangle$ and $\mathbf{a}_i \triangleq \langle a_1, a_2, \dots, a_{|\mathbf{v}_i|} \rangle$ denote some disjoint sets of fresh variables. Because of $\mathfrak{f}_0 = \mathfrak{f}_1 \uplus \mathfrak{f}_2$ we have $\mathbf{w}_0 = \mathbf{w}_1 \cdot \mathbf{w}_2$ and hence can choose \mathbf{e}_0 such that $\mathbf{e}_0 = \mathbf{e}_1 \cdot \mathbf{e}_2$.

We now argue that we can find sequences of variables $\mathbf{u}_i \subseteq \mathbf{a}_0 \cup \mathbf{e}_{3-i}$, for $i = 1, 2$, such that

$$\begin{aligned} \phi_0[\text{dom}(\mathfrak{s}_{\max}^{-1}) \cdot \mathbf{v}_0 \cdot \mathbf{w}_0 / \text{img}(\mathfrak{s}_{\max}^{-1}) \cdot \mathbf{a}_0 \cdot \mathbf{e}_0] \equiv \\ \phi_1[\text{dom}(\mathfrak{s}_{\max}^{-1}) \cdot \mathbf{v}_1 \cdot \mathbf{w}_2 / \text{img}(\mathfrak{s}_{\max}^{-1}) \cdot \mathbf{a}_1 \cdot \mathbf{e}_1][\mathbf{a}_1 / \mathbf{u}_1] \star \\ \phi_2[\text{dom}(\mathfrak{s}_{\max}^{-1}) \cdot \mathbf{v}_2 \cdot \mathbf{w}_2 / \text{img}(\mathfrak{s}_{\max}^{-1}) \cdot \mathbf{a}_2 \cdot \mathbf{e}_2][\mathbf{a}_2 / \mathbf{u}_2] (*) \end{aligned}$$

We consider a location $l \in \mathbf{v}_i$ for $i \in \{1, 2\}$. If $l \in \text{dom}(\mathfrak{f}_{3-i})$, then there is a variable $e \in \mathbf{e}_{3-i}$ which replaces l in the projection $\text{project}(\mathfrak{s}, \mathfrak{f}_0)$. If $l \notin \text{dom}(\text{heap}(\mathfrak{f}_{3-i}))$, then there is a variable $a \in \mathbf{a}_0$ which replaces l in the projection $\text{project}(\mathfrak{s}, \mathfrak{f}_0)$. Hence, we can choose sequences of variables $\mathbf{u}_i \subseteq \mathbf{a}_0 \cup \mathbf{e}_{3-i}$, for $i \in \{1, 2\}$, such that the following holds for all $l \in \mathbf{v}_i$ and $i \in \{1, 2\}$:

$$l[\mathbf{v}_0 \cdot \mathbf{w}_0 / \mathbf{a}_0 \cdot \mathbf{e}_0] = l[\mathbf{v}_i / \mathbf{a}_i][\mathbf{a}_i / \mathbf{u}_i]$$

The above then implies (*). □

LEMMA A.2. *Let \mathfrak{s} be a stack and let $\mathfrak{f}_1, \mathfrak{f}_2$ be Φ -forests with $\mathfrak{f}_1 \blacktriangleright \mathfrak{f}_2$. Then, $\text{project}(\mathfrak{s}, \mathfrak{f}_1) \triangleright \text{project}(\mathfrak{s}, \mathfrak{f}_2)$.*

PROOF. Since $\mathfrak{f}_1 \blacktriangleright \mathfrak{f}_2$, there exists a forest \mathfrak{f} and trees t_1, t_2, t such that

- (1) $\mathfrak{f}_1 = \mathfrak{f} \cup \{t_1, t_2\}$,
- (2) $\mathfrak{f}_2 = \mathfrak{f} \cup \{t\}$,
- (3) $\text{rootpred}(t_1) \in \text{allholepreds}(t_2)$,
- (4) $\text{rootpred}(t) = \text{rootpred}(t_2)$, and
- (5) $\text{allholepreds}(t) = \text{allholepreds}(t_1) \cup (\text{allholepreds}(t_2) \setminus \{\text{rootpred}(t_1)\})$.

Intuitively, this implies that the projections of t_1 and t_2 can be merged into the projection of t via the generalized modus ponens (see Lemma 7.18). In the following we make this claim formal.

For $i \in \{1, 2\}$, let $\phi_i = \star_{t \in \mathfrak{f}_i} \text{project}^{\text{Loc}}(t)$, let $\mathbf{w}_i = \text{locs}(\phi_i) \cap \text{dom}(\mathfrak{f}_i) \setminus \text{img}(\mathfrak{s})$ be the locations that occur in the formula ϕ and are allocated in $\text{heap}(\mathfrak{f}_i)$ but are not the value of any stack variable, and let $\mathbf{v}_i = \text{locs}(\phi_i) \setminus (\text{img}(\mathfrak{s}) \cup \text{dom}(\mathfrak{f}_i))$ be the locations that occur in the formula ϕ_i and are neither allocated nor the value of any stack variable. Then, we have

$$\text{project}(\mathfrak{s}, \mathfrak{f}_i) = \exists \mathbf{e}_i. \forall \mathbf{a}_i. \phi_i[\text{dom}(\mathfrak{s}_{\max}^{-1}) \cdot \mathbf{v}_i \cdot \mathbf{w}_i / \text{img}(\mathfrak{s}_{\max}^{-1}) \cdot \mathbf{a}_i \cdot \mathbf{e}_i],$$

where $\mathbf{e}_i \triangleq \langle e_1, e_2, \dots, e_{|\mathbf{w}_i|} \rangle$ and $\mathbf{a}_i \triangleq \langle a_1, a_2, \dots, a_{|\mathbf{v}_i|} \rangle$ denote some disjoint sets of fresh variables. We note that

$$\phi_1 \equiv \star_{t' \in \mathfrak{f}} \text{project}^{\text{Loc}}(t') \star (\star \text{allholepreds}(t_1)) \star \text{rootpred}(t_1) \star (\star \text{allholepreds}(t_2)) \star \text{rootpred}(t_2)$$

and

$$\phi_2 \equiv \star_{t' \in \mathfrak{f}} \text{project}^{\text{Loc}}(t') \star (\star \text{allholepreds}(t_1)) \star \star (\text{allholepreds}(t_2) \setminus \{\text{rootpred}(t_1)\}) \star \text{rootpred}(t_2).$$

In particular, we have $\text{locs}(\phi_2) \subseteq \text{locs}(\phi_1)$. Hence, we can assume without loss of generality that $\mathbf{e}_2 \subseteq \mathbf{e}_1$ and $\mathbf{a}_2 \subseteq \mathbf{a}_1$. By (*), we have

$$\begin{aligned} \text{project}(\mathfrak{s}, \mathfrak{f}_1) \equiv \exists \mathbf{e}_1. \forall \mathbf{a}_1. \star_{t' \in \mathfrak{f}} \text{project}^{\text{Loc}}(t') \star (\star \text{allholepreds}(t_1)) \star \text{rootpred}(t_1) \star \\ (\star \text{allholepreds}(t_2)) \star \text{rootpred}(t_2) [\text{dom}(\mathfrak{s}_{\text{max}}^{-1}) \cdot \mathbf{v}_i \cdot \mathbf{w}_i / \text{img}(\mathfrak{s}_{\text{max}}^{-1}) \cdot \mathbf{a}_i \cdot \mathbf{e}_i], \end{aligned}$$

and

$$\begin{aligned} \text{project}(\mathfrak{s}, \mathfrak{f}_2) \equiv \exists \mathbf{e}_1. \forall \mathbf{a}_1. \star_{t' \in \mathfrak{f}} \text{project}^{\text{Loc}}(t') \star (\star \text{allholepreds}(t_1)) \star \\ \star (\text{allholepreds}(t_2) \setminus \{\text{rootpred}(t_1)\}) \star \text{rootpred}(t_2) [\text{dom}(\mathfrak{s}_{\text{max}}^{-1}) \cdot \mathbf{v}_i \cdot \mathbf{w}_i / \text{img}(\mathfrak{s}_{\text{max}}^{-1}) \cdot \mathbf{a}_i \cdot \mathbf{e}_i], \end{aligned}$$

We now recognize that $\text{project}(\mathfrak{s}, \mathfrak{f}_2)$ can be obtained from $\text{project}(\mathfrak{s}, \mathfrak{f}_1)$ by applying the generalized modus ponens rule and dropping the quantified variables $\mathbf{e}_1 \setminus \mathbf{e}_2$ and $\mathbf{a}_1 \setminus \mathbf{a}_2$, which is supported by our rewriting rules (see Fig. 10) because these variables do not appear in ϕ_2 . \square

Claim (Lemma 7.35). Let \mathfrak{s} be a stack and let $\mathfrak{f}_1, \mathfrak{f}_2$ be Φ -forests such that $\mathfrak{f}_1 \uplus \mathfrak{f}_2 \neq \perp$. Then,

$$\mathfrak{f} \in \mathfrak{f}_1 \bullet_{\mathbf{F}} \mathfrak{f}_2 \quad \text{implies} \quad \text{project}(\mathfrak{s}, \mathfrak{f}) \in \text{project}(\mathfrak{s}, \mathfrak{f}_1) \bullet_{\mathbf{P}} \text{project}(\mathfrak{s}, \mathfrak{f}_2).$$

PROOF. The claim is an immediate consequence of Lemmas A.1 and A.2. \square

A.15 Proof of Theorem 7.39

Before we prove Theorem 7.39, we need two auxiliary results.

LEMMA A.3. *Let \mathfrak{s} be a stack, let $\mathfrak{f}_1, \mathfrak{f}_2$ be Φ -forests with $\text{dom}(\mathfrak{f}_1) \cap \text{dom}(\mathfrak{f}_2) \cap \text{img}(\mathfrak{s}) = \emptyset$, and let $\chi \in \text{project}(\mathfrak{s}, \mathfrak{f}_1) \bullet_{\mathbf{P}} \text{project}(\mathfrak{s}, \mathfrak{f}_2)$. Then, there exist forests $\mathfrak{f}'_1, \mathfrak{f}'_2$ with $\mathfrak{f}_1 \equiv_{\mathfrak{s}} \mathfrak{f}'_1$, $\mathfrak{f}_2 \equiv_{\mathfrak{s}} \mathfrak{f}'_2$ and $\text{project}(\mathfrak{s}, \mathfrak{f}'_1 \uplus \mathfrak{f}'_2) \equiv \chi$.*

PROOF. For $i \in \{1, 2\}$, let $\phi_i = \star_{t \in \mathfrak{f}_i} \text{project}^{\text{Loc}}(t)$, let $\mathbf{w}_i = \text{locs}(\phi_i) \cap \text{dom}(\mathfrak{f}_i) \setminus \text{img}(\mathfrak{s})$ be the locations that occur in the formula ϕ and are allocated in $\text{heap}(\mathfrak{f}_i)$ but are not the value of any stack variable, and let $\mathbf{v}_i = \text{locs}(\phi_i) \setminus (\text{img}(\mathfrak{s}) \cup \text{dom}(\mathfrak{f}_i))$ be the locations that occur in the formula ϕ_i and are neither allocated nor the value of any stack variable. Then, we have

$$\text{project}(\mathfrak{s}, \mathfrak{f}_i) = \exists \mathbf{e}_i. \forall \mathbf{a}_i. \phi_i [\text{dom}(\mathfrak{s}_{\text{max}}^{-1}) \cdot \mathbf{v}_i \cdot \mathbf{w}_i / \text{img}(\mathfrak{s}_{\text{max}}^{-1}) \cdot \mathbf{a}_i \cdot \mathbf{e}_i],$$

where $\mathbf{e}_i \triangleq \langle e_1, e_2, \dots, e_{|\mathbf{w}_i|} \rangle$ and $\mathbf{a}_i \triangleq \langle a_1, a_2, \dots, a_{|\mathbf{v}_i|} \rangle$ denote some disjoint sets of fresh variables.

By the definition of the re-scoping operation, we have $\chi = \exists \mathbf{e}. \forall \mathbf{a}. \phi$, where

- (1) $\mathbf{e} = \mathbf{e}_1 \cdot \mathbf{e}_2$, and
- (2) $\phi = \phi_1 [\text{dom}(\mathfrak{s}_{\text{max}}^{-1}) \cdot \mathbf{v}_1 \cdot \mathbf{w}_1 / \text{img}(\mathfrak{s}_{\text{max}}^{-1}) \cdot \mathbf{a}_1 \cdot \mathbf{e}_1] [\mathbf{a}_1 / \mathbf{u}_1] \star \phi_2 [\text{dom}(\mathfrak{s}_{\text{max}}^{-1}) \cdot \mathbf{v}_2 \cdot \mathbf{w}_2 / \text{img}(\mathfrak{s}_{\text{max}}^{-1}) \cdot \mathbf{a}_2 \cdot \mathbf{e}_2] [\mathbf{a}_2 / \mathbf{u}_2]$ for some sequences $\mathbf{u}_i \subseteq \mathbf{a} \cup \mathbf{e}_{3-i}$.

We can now choose bijective functions $\sigma_1: \text{Loc} \rightarrow \text{Loc}$ and $\sigma_2: \text{Loc} \rightarrow \text{Loc}$ such that

- $\sigma_1(l) = \sigma_2(l) = l$ for all $l \in \text{img}(\mathfrak{s})$,
- $\sigma_1(l) = \sigma_2(k)$ if and only if $l[\mathbf{v}_1 \cdot \mathbf{w}_1 / \mathbf{a}_1 \cdot \mathbf{e}_1] [\mathbf{a}_1 / \mathbf{u}_1] = k[\mathbf{v}_2 \cdot \mathbf{w}_2 / \mathbf{a}_2 \cdot \mathbf{e}_2] [\mathbf{a}_2 / \mathbf{u}_2]$ for all $l \in \mathbf{v}_1 \cdot \mathbf{w}_1, k \in \mathbf{v}_2 \cdot \mathbf{w}_2$, and
- $\text{dom}(\sigma_1(\mathfrak{f}_1)) \cap \text{dom}(\sigma_1(\mathfrak{f}_2)) = \emptyset$.

We set $\mathfrak{f}'_1 \triangleq \sigma(\mathfrak{f}_1)$ and $\mathfrak{f}'_2 \triangleq \sigma(\mathfrak{f}_2)$. By the above we have $\mathfrak{f}_1 \equiv_{\mathfrak{s}} \mathfrak{f}'_1$, $\mathfrak{f}_2 \equiv_{\mathfrak{s}} \mathfrak{f}'_2$ and $\mathfrak{f}'_1 \uplus \mathfrak{f}'_2 \neq \perp$. Further, we get that $\text{project}^{\text{Loc}}(\mathfrak{f}'_1 \uplus \mathfrak{f}'_2) \equiv \text{project}^{\text{Loc}}(\mathfrak{f}_1) [\text{dom}(\sigma_1) / \text{img}(\sigma_1)] \star \text{project}^{\text{Loc}}(\mathfrak{f}_2) [\text{dom}(\sigma_2) / \text{img}(\sigma_2)]$. Finally, we get that $\text{project}(\mathfrak{s}, \mathfrak{f}'_1 \uplus \mathfrak{f}'_2) \equiv \exists \mathbf{e}. \forall \mathbf{a}. \phi$ because we can appropriately rename the quantified variables by rewrite equivalence \equiv . \square

We note that the below lemma does not require the notion of \mathfrak{s} -equivalence:

LEMMA A.4. Let \bar{f}_1 be a Φ -forest and let χ be a formula such that $\text{project}(\bar{s}, \bar{f}_1) \triangleright \chi$. Then, there exist a forest \bar{f}_2 with $\bar{f}_1 \blacktriangleright \bar{f}_2$ and $\text{project}(\bar{s}, \bar{f}_2) \equiv \chi$.

PROOF. Let $\phi_1 = \star_{t \in \bar{f}_1} \text{project}^{\text{Loc}}(t)$, let $\mathbf{w}_1 = \text{locs}(\phi_1) \cap (\text{dom}(\bar{f}_1) \setminus \text{img}(\bar{s}))$ be the locations that occur in the formula ϕ_1 and are allocated in $\text{heap}(\bar{f}_1)$ but are not the value of any stack variable, and let $\mathbf{v}_1 = \text{locs}(\phi_1) \setminus (\text{img}(\bar{s}) \cup \text{dom}(\bar{f}_1))$ be the locations that occur in the formula ϕ_1 and are neither allocated nor the value of any stack variable. Then, we have

$$\text{project}(\bar{s}, \bar{f}_1) = \exists \mathbf{e}_1. \forall \mathbf{a}_1. \phi_1[\text{dom}(\bar{s}_{\max}^{-1}) \cdot \mathbf{v}_1 \cdot \mathbf{w}_1 / \text{img}(\bar{s}_{\max}^{-1}) \cdot \mathbf{a}_1 \cdot \mathbf{e}_1],$$

where $\mathbf{e}_1 \triangleq \langle e_1, e_2, \dots, e_{|\mathbf{w}_1|} \rangle$ and $\mathbf{a}_1 \triangleq \langle a_1, a_2, \dots, a_{|\mathbf{v}_1|} \rangle$ denote some disjoint sets of fresh variables. By definition of the projection of forests, we have

$$\phi_1 \equiv \star_{t \in \bar{f}_1} ((\star \text{allholepreds}(t)) \rightarrow \text{rootpred}(t)).$$

By the definition of \triangleright we have that there are predicates $\text{pred}_1(\mathbf{x}_1)$, $\text{pred}_2(\mathbf{x}_2)$, and formulae ψ, ψ', ζ such that

- (1) $\phi_1 \equiv (\text{pred}_2(\mathbf{x}_2) \star \psi) \rightarrow \text{pred}_1(\mathbf{x}_1) \star (\psi' \rightarrow \text{pred}_2(\mathbf{x}_2)) \star \zeta$, and
- (2) $\chi \equiv \exists \mathbf{e}_1. \forall \mathbf{a}_1. (\psi \star \psi') \rightarrow \text{pred}_1(\mathbf{x}_1) \star \zeta$.

Hence, there must be a forest \bar{f} and trees t_1, t_2 such that

- (1) $\bar{f}_1 = \bar{f} \cup \{t_1, t_2\}$,
- (2) $\text{rootpred}(t_2) \in \text{allholepreds}(t_1)$, and
- (3) $\text{rootpred}(t_2)[\text{dom}(\bar{s}_{\max}^{-1}) \cdot \mathbf{v}_1 \cdot \mathbf{w}_1 / \text{img}(\bar{s}_{\max}^{-1}) \cdot \mathbf{a}_1 \cdot \mathbf{e}_1] = \text{pred}_2(\mathbf{x}_2)$.

Let $l = \text{root}(t_2)$. Then, there is a tree t with $\{t_1, t_2\} = \text{split}(\{t\}, \{l\})$. We set $\bar{f}_2 = \bar{f} \cup \{t\}$. We note that $\bar{f}_1 \blacktriangleright \bar{f}_2$. It remains to argue that $\text{project}(\bar{s}, \bar{f}_2) \equiv \chi$.

Let $\phi_2 = \star_{t \in \bar{f}_2} \text{project}^{\text{Loc}}(t)$, let $\mathbf{w}_2 = \text{locs}(\phi_2) \cap (\text{dom}(\bar{f}_2) \setminus \text{img}(\bar{s}))$ be the locations that occur in the formula ϕ_2 and are allocated in $\text{heap}(\bar{f}_2)$ but are not the value of any stack variable, and let $\mathbf{v}_2 = \text{locs}(\phi_2) \setminus (\text{img}(\bar{s}) \cup \text{dom}(\bar{f}_2))$ be the locations that occur in the formula ϕ_2 and are neither allocated nor the value of any stack variable. Then, we have

$$\text{project}(\bar{s}, \bar{f}_2) = \exists \mathbf{e}_2. \forall \mathbf{a}_2. \phi_2[\text{dom}(\bar{s}_{\max}^{-1}) \cdot \mathbf{v}_2 \cdot \mathbf{w}_2 / \text{img}(\bar{s}_{\max}^{-1}) \cdot \mathbf{a}_2 \cdot \mathbf{e}_2],$$

where $\mathbf{e}_2 \triangleq \langle e_1, e_2, \dots, e_{|\mathbf{w}_2|} \rangle$ and $\mathbf{a}_2 \triangleq \langle a_1, a_2, \dots, a_{|\mathbf{v}_2|} \rangle$ denote some disjoint sets of fresh variables. We now note that $\text{locs}(\phi_2) \cup \{l\} = \text{locs}(\phi_1)$. Hence, we can assume without loss of generality that $\mathbf{e}_2 \subseteq \mathbf{e}_1$ and $\mathbf{a}_2 \subseteq \mathbf{a}_1$. Thus, we have $(\psi \star \psi') \rightarrow \text{pred}_1(\mathbf{x}_1) \star \zeta \equiv \phi_2[\text{dom}(\bar{s}_{\max}^{-1}) \cdot \mathbf{v}_2 \cdot \mathbf{w}_2 / \text{img}(\bar{s}_{\max}^{-1}) \cdot \mathbf{a}_2 \cdot \mathbf{e}_2]$. Finally, we note that

$$\begin{aligned} \chi &\equiv \exists \mathbf{e}_1. \forall \mathbf{a}_1. \phi_2[\text{dom}(\bar{s}_{\max}^{-1}) \cdot \mathbf{v}_2 \cdot \mathbf{w}_2 / \text{img}(\bar{s}_{\max}^{-1}) \cdot \mathbf{a}_2 \cdot \mathbf{e}_2] \equiv \\ &\quad \exists \mathbf{e}_2. \forall \mathbf{a}_2. \phi_2[\text{dom}(\bar{s}_{\max}^{-1}) \cdot \mathbf{v}_2 \cdot \mathbf{w}_2 / \text{img}(\bar{s}_{\max}^{-1}) \cdot \mathbf{a}_2 \cdot \mathbf{e}_2], \end{aligned}$$

because we can drop the quantified variables $\mathbf{e}_1 \setminus \mathbf{e}_2$ and $\mathbf{a}_1 \setminus \mathbf{a}_2$, which is supported by our rewriting rules (see Fig. 10) because these variables do not appear in ϕ_2 . \square

Claim (Theorem 7.39). If \bar{f}_1, \bar{f}_2 be Φ -forests with $\bar{f}_1 \equiv_s \bar{f}_2$, then

$$\text{project}(\bar{s}, \bar{f}_1) \bullet_P \text{project}(\bar{s}, \bar{f}_2) = \{\text{project}(\bar{s}, \bar{f}) \mid \bar{f} \in \bar{f}_1 \bullet_F \bar{f}_2, \bar{f}_1 \equiv_s \bar{f}_1, \bar{f}_2 \equiv_s \bar{f}_2\}.$$

PROOF. The claim is an immediate consequence of Lemmas 7.35, A.3 and A.4. \square

A.16 Proof of Lemma 8.6

Claim. Let $n \triangleq |\Phi| + |\mathbf{x}|$, where \mathbf{x} is a finite set of variables. Then $|\mathbf{DUSH}_{\Phi}^{\mathbf{x}}| \in 2^{O(n^2 \log(n))}$.

PROOF. We first show the following claim (\dagger): every element of $\mathbf{DUSH}_{\Phi}^{\mathbf{x}}$ can be encoded as a string of length $O(n^2)$ over the alphabet $Z \triangleq \mathbf{Preds}(\Phi) \cup \mathbf{x} \cup \{e_1, \dots, e_{n^2}\} \cup \{a_1, \dots, a_{n^2}\} \cup \{\mathbf{emp}, \star, \star, (,)\}$ where e_1, \dots, e_{n^2} and a_1, \dots, a_{n^2} are fresh variables.

By definition, every DUSH $\phi \in \mathbf{DUSH}_{\Phi}^{\mathbf{x}}$ is of the form

$$\begin{aligned} \phi &= \exists \mathbf{e}. \forall \mathbf{a}. \psi_1 \star \dots \star \psi_m, \\ \psi_i &= \zeta_i \star \text{pred}_i(\mathbf{z}_i) \text{ for } 1 \leq i \leq m. \end{aligned}$$

Since ϕ is delimited, $\text{predroot}(\text{pred}_i(\mathbf{z}_i)) \in \mathbf{x}$. Moreover, ϕ is the projection of a Φ -forest \mathfrak{f} . Hence, every variable $x \in \mathbf{x}$ can appear as a root parameter in at most one subformula ψ_i —otherwise, the value corresponding to x would be in the domain of two trees in \mathfrak{f} , which contradicts the fact that \mathfrak{f} is a Φ -forest. Consequently, the number m of subformulas ψ_i is bounded by $|\mathbf{x}| \leq n$.

Next, consider the subformulas ζ_i appearing on the left-hand side of magic wands. For every predicate call $\text{pred}'(\mathbf{z}')$ in ζ_i , $\text{predroot}(\text{pred}'(\mathbf{z}'))$ is a hole. Since the forest \mathfrak{f} is delimited, it follows that $\text{predroot}(\text{pred}'(\mathbf{z}')) \in \mathbf{x}$. Since no hole may occur more than once in a delimited USH, the *total* number of predicate calls across all ζ_i is also bounded by $|\mathbf{x}| \leq n$.

Overall, ϕ thus contains at most $2n \in O(n)$ predicate calls. Each predicate call takes at most $|\Phi| \leq n$ parameters. Since there are no superfluous quantified variables, this means that ϕ contains at most $n^2 - |\mathbf{x}| \leq n^2$ different variables. We can thus assume w.l.o.g. that all existentially-quantified variables in ϕ are among the variables e_1, \dots, e_{n^2} and all universally-quantified variables are among a_1, \dots, a_{n^2} . There then is no need to include the quantifiers explicitly in the string encoding. After dropping the quantifiers, we obtain a formula ϕ' that consists exclusively of letters from the alphabet Z . Moreover, this formula consists of at most $O(n^2)$ letters. This concludes the proof of (\dagger).

Now observe that $|Z| \in O(n^2)$. Consequently, every letter of Z can be encoded by $O(\log(n^2)) = O(\log(n))$ bits. Therefore, every $\phi \in \mathbf{DUSH}_{\Phi}^{\mathbf{x}}$ can be encoded by a bit string of length $O(n^2 \log(n))$. Since there are $2^{O(n^2 \log(n))}$ such strings, the claim follows. \square

A.17 Proof of Lemma 8.12

Claim. Let \mathfrak{f} be a forest and let \mathfrak{s} be a stack. Then \mathfrak{f} is \mathfrak{s} -delimited iff $\text{project}(\mathfrak{s}, \mathfrak{f})$ is delimited.

PROOF. Recall that the projection contains predicate calls corresponding to the roots and holes of the forest. It thus holds for all forests that

$$\text{interface}(\mathfrak{f}) = \{\text{predroot}(\text{pred}(\mathbf{z})) \mid \text{pred}(\mathbf{z}) \in \text{project}(\mathfrak{f})\}. \quad (\dagger)$$

We show that if \mathfrak{f} is \mathfrak{s} -delimited then $\text{project}(\mathfrak{s}, \mathfrak{f})$ is delimited. The proof of the other direction is completely analogous.

If \mathfrak{f} is \mathfrak{s} -delimited then $\text{interface}(\mathfrak{f}) \subseteq \text{img}(\mathfrak{s})$ and thus, by (\dagger),

$$\{\text{predroot}(\text{pred}(\mathbf{z})) \mid \text{pred}(\mathbf{z}) \in \text{project}(\mathfrak{f})\} \subseteq \text{img}(\mathfrak{s}).$$

Trivially, the set of root locations in the projection is a subset of the set of *all* locations in the projection.

$$\{\text{predroot}(\text{pred}(\mathbf{z})) \mid \text{pred}(\mathbf{z}) \in \text{project}(\mathfrak{f})\} \subseteq \text{locs}(\text{project}(\mathfrak{f})).$$

Combining the above two observations, we conclude

$$\begin{aligned} & \{\text{predroot}(\text{pred}(z)) \mid \text{pred}(z) \in \text{project}(\mathfrak{f})\} \\ & \subseteq \text{img}(\mathfrak{s}) \cap \text{locs}(\text{project}(\mathfrak{f})). \end{aligned}$$

We apply \mathfrak{s}_{\max}^{-1} on both sides to obtain that

$$\begin{aligned} & \mathfrak{s}_{\max}^{-1}(\{\text{predroot}(\text{pred}(z)) \mid \text{pred}(z) \in \text{project}(\mathfrak{f})\}) \\ & \subseteq \text{dom}(\mathfrak{s}) \cap \underbrace{(\text{fvars}(\text{project}(\mathfrak{s}, \mathfrak{f})))}_{\{\mathfrak{s}_{\max}^{-1}(l) \mid l \in \text{img}(\mathfrak{s}) \cap \text{locs}(\text{project}(\mathfrak{f}))\}} \subseteq \text{fvars}(\text{project}(\mathfrak{s}, \mathfrak{f})). \end{aligned}$$

Moreover, since there are no duplicate holes in \mathfrak{f} , and the holes of \mathfrak{f} are mapped to the predicate calls on the left-hand side of magic wands in $\text{project}(\mathfrak{s}, \mathfrak{f})$, no variable can occur twice as root parameter on the left-hand side of magic wands in $\text{project}(\mathfrak{s}, \mathfrak{f})$.

Consequently, $\text{project}(\mathfrak{s}, \mathfrak{f})$ is delimited. \square

A.18 Proof of Theorem 8.13

We first give some auxiliary definitions and results. Recall that we described how Φ -forests are merged in terms of splitting them at suitable locations (cf. Definition 7.9). Every split adds these locations to the interface of the resulting forest—provided they did not appear in the forest to begin with.

LEMMA A.5. *Let \mathfrak{f} be a forest and $\mathbf{l} \subseteq \text{Loc}$. Then, $\text{interface}(\text{split}(\mathfrak{f}, \mathbf{l})) = \text{interface}(\mathfrak{f}) \cup (\mathbf{l} \cap \text{dom}(\mathfrak{f}))$.*

PROOF. In the following, let $\text{locs}(\text{graph}(\mathfrak{f}))$ denote all those locations that occur in the relation $\text{graph}(\mathfrak{f})$.

$$\begin{aligned} & \text{roots}(\text{split}(\mathfrak{f}, \mathbf{l})) \\ & = \text{roots}(\mathfrak{f}) \cup \{b \in \mathbf{l} \mid \exists a. (a, b) \in \text{graph}(\mathfrak{f})\} \\ & = \text{roots}(\mathfrak{f}) \cup \{b \in \mathbf{l} \cap \text{dom}(\mathfrak{f}) \mid \exists a. (a, b) \in \text{graph}(\mathfrak{f})\} && (\text{locs}(\text{graph}(\mathfrak{f})) \subseteq \text{dom}(\mathfrak{f})) \\ & = \{b \in \text{dom}(\mathfrak{f}) \mid \forall a. (a, b) \notin \text{graph}(\mathfrak{f})\} \\ & \quad \cup \{b \in \mathbf{l} \cap \text{dom}(\mathfrak{f}) \mid \exists a. (a, b) \in \text{graph}(\mathfrak{f})\} && (\text{all and only roots have no predecessor}) \\ & = \{b \in \text{dom}(\mathfrak{f}) \mid \forall a. (a, b) \notin \text{graph}(\mathfrak{f})\} \\ & \quad \cup \{b \in \mathbf{l} \cap \text{dom}(\mathfrak{f}) \mid \forall a. (a, b) \notin \text{graph}(\mathfrak{f})\} \\ & \quad \cup \{b \in \mathbf{l} \cap \text{dom}(\mathfrak{f}) \mid \exists a. (a, b) \in \text{graph}(\mathfrak{f})\} && (\text{second set subset of first set}) \\ & = \text{roots}(\mathfrak{f}) \cup \{b \in \mathbf{l} \cap \text{dom}(\mathfrak{f}) \mid \forall a. (a, b) \notin \text{graph}(\mathfrak{f})\} \\ & \quad \cup \{b \in \mathbf{l} \cap \text{dom}(\mathfrak{f}) \mid \exists a. (a, b) \in \text{graph}(\mathfrak{f})\} \\ & = \text{roots}(\mathfrak{f}) \cup \{b \in \mathbf{l} \cap \text{dom}(\mathfrak{f})\} \end{aligned}$$

Similarly,

$$\begin{aligned} \text{allholes}(\text{split}(\mathfrak{f}, \mathbf{l})) & = \text{allholes}(\mathfrak{f}) \cup \{b \in \mathbf{l} \mid \exists a. (a, b) \in \text{graph}(\mathfrak{f})\} \\ & = \text{allholes}(\mathfrak{f}) \cup \{b \in \mathbf{l} \cap \text{dom}(\mathfrak{f})\}. \end{aligned}$$

By definition of interfaces, we thus obtain

$$\begin{aligned} \text{interface}(\text{split}(\bar{f}, \mathbf{I})) &= \text{roots}(\text{split}(\bar{f}, \mathbf{I})) \cup \text{allholes}(\text{split}(\bar{f}, \mathbf{I})) \\ &= \text{roots}(\bar{f}) \cup \{t \in \mathbf{I} \cap \text{dom}(\bar{f})\} \\ &\quad \cup \text{allholes}(\bar{f}) \cup \{b \in \mathbf{I} \cap \text{dom}(\bar{f})\} \\ &= \text{interface}(\bar{f}) \cup \{b \in \mathbf{I} \cap \text{dom}(\bar{f})\}. \end{aligned}$$

□

DEFINITION A.6. Let \bar{f} be an \mathfrak{s} -delimited forest. We call $\text{split}(\bar{f}, \text{img}(\mathfrak{s}))$ the \mathfrak{s} -decomposition of \bar{f} .

LEMMA A.7. The \mathfrak{s} -decomposition of an \mathfrak{s} -delimited forest is \mathfrak{s} -delimited.

PROOF. Let \bar{f} be the \mathfrak{s} -decomposition of an \mathfrak{s} -delimited forest f . By definition, $\bar{f} = \text{split}(f, \text{img}(\mathfrak{s}))$. By Lemma A.5, we have $\text{interface}(\bar{f}) \subseteq \text{interface}(f) \cup \text{img}(\mathfrak{s})$. Since f is \mathfrak{s} -delimited, $\text{interface}(f) \subseteq \text{img}(\mathfrak{s})$. Overall, we thus obtain $\text{interface}(\bar{f}) \subseteq \text{img}(\mathfrak{s})$, i.e., \bar{f} is \mathfrak{s} -delimited. □

We observe that, since the \mathfrak{s} -decomposition of a forest is obtained by splitting the trees of the forest at all locations in $\text{img}(\mathfrak{s})$, only the roots of the trees in an \mathfrak{s} -decomposition of forest \bar{f} can be locations in $\text{img}(\mathfrak{s})$:

LEMMA A.8. For every Φ -tree \bar{t} in an \mathfrak{s} -decomposition, we have $\text{img}(\mathfrak{s}) \cap \text{dom}(\bar{t}) = \{\text{root}(\bar{t})\}$.

PROOF. Let \bar{f} be an \mathfrak{s} -decomposition and let $\bar{t} \in \bar{f}$. Since \bar{f} is \mathfrak{s} -delimited by Lemma A.7, we have $\{\text{root}(\bar{t})\} \subseteq \text{img}(\mathfrak{s})$. Since $\text{root}(\bar{t}) \in \text{dom}(\bar{t})$, $\{\text{root}(\bar{t})\} \subseteq \text{img}(\mathfrak{s}) \cap \text{dom}(\bar{t})$.

Conversely, since $\bar{f} = \text{split}(f, \text{img}(\mathfrak{s}))$, we have $\text{roots}(\bar{f}) = \text{roots}(f) \cup (\text{img}(\mathfrak{s}) \cap \text{dom}(f))$, i.e., every location in $\text{img}(\mathfrak{s}) \cap \text{dom}(\bar{f})$ is a root of \bar{f} . Consequently, $\text{img}(\mathfrak{s}) \cap \text{dom}(\bar{t}) \subseteq \{\text{root}(\bar{t})\}$. □

LEMMA A.9. Let $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle, \langle \mathfrak{s}, \mathfrak{h}_2 \rangle \in \mathbf{GStates}$ be guarded states, and let f be a \mathfrak{s} -delimited forest with $f \in \text{forests}_{\Phi}(\mathfrak{h}_1 \uplus \mathfrak{h}_2)$. Then, there exist forests f_1, f_2 with $f_1 \uplus f_2 = f$ and $\text{heap}(f_i) = \mathfrak{h}_i$, where \bar{f} is the \mathfrak{s} -decomposition of f .

PROOF. We let $f_i \triangleq \{\bar{t} \in \bar{f} \mid \text{root}(\bar{t}) \in \text{dom}(\mathfrak{h}_i)\}$. Since $f_1 \uplus f_2 = f$ and thus $\text{heap}(f_1) \uplus \text{heap}(f_2) = \text{heap}(f)$ by Lemma 7.7, it suffices to show that for every tree \bar{t} in f_i that $\text{heap}(\bar{t}) \subseteq \mathfrak{h}_i$.

To this end, let $\bar{t} \in f_i$. Assume towards a contradiction that $\text{dom}(\bar{t}) \cap \text{dom}(\mathfrak{h}_{3-i}) \neq \emptyset$. Then there exist locations $l_1 \in \text{dom}(\bar{t}) \cap \text{dom}(\mathfrak{h}_i)$ and $l_2 \in \text{dom}(\bar{t}) \cap \text{dom}(\mathfrak{h}_{3-i})$ with $l_2 \in \text{succ}_i(l_1)$. In particular, $l_2 \in \text{img}(\mathfrak{h}_i)$ and $l_2 \in \text{dom}(\mathfrak{h}_{3-i})$, implying that $l_2 \in \text{dangling}(\mathfrak{h}_i)$. However, since $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle, \langle \mathfrak{s}, \mathfrak{h}_2 \rangle \in \mathbf{GStates}$, we have that $l_2 \in \text{img}(\mathfrak{s})$. Since $l_2 \neq \text{root}(\bar{t})$, this contradicts Lemma A.8. □

We restate the claim of Theorem 8.13: Let $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle, \langle \mathfrak{s}, \mathfrak{h}_2 \rangle \in \mathbf{GStates}$ be guarded states, and let f be a \mathfrak{s} -delimited forest with $f \in \text{forests}_{\Phi}(\mathfrak{h}_1 \uplus \mathfrak{h}_2)$. Then there exist \mathfrak{s} -delimited forests f_1, f_2 with $\text{heap}(f_i) = \mathfrak{h}_i$ and $f \in f_1 \bullet_{\mathbf{F}} f_2$.

PROOF. Let \bar{f} be the \mathfrak{s} -decomposition of f . In particular, we then have $\bar{f} \blacktriangleright^* f$ by definition of \blacktriangleright^* . Let f_1, f_2 be such that $f_1 \uplus f_2 = \bar{f}$ and $\text{heap}(f_i) = \mathfrak{h}_i$. Such forests exist by Lemma A.9. Then $f_1 \uplus f_2 = \bar{f} \blacktriangleright^* f$, i.e., $f \in f_1 \bullet_{\mathbf{F}} f_2$. Since \bar{f} is \mathfrak{s} -delimited (by Lemma A.7), so are f_1 and f_2 . □

A.19 Proof of Theorem 8.14

Claim. For all guarded states $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle$ and $\langle \mathfrak{s}, \mathfrak{h}_2 \rangle$ with $\mathfrak{h}_1 \uplus \mathfrak{h}_2 \neq \perp$, $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2)$ can be computed from $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1)$ and $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_2)$ as follows:

$$\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2) = \{\phi \in \mathbf{DUSH}_{\Phi} \mid \text{ex. } \psi_1 \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1), \psi_2 \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_2) \text{ such that } \phi \in \psi_1 \bullet_{\mathbf{P}} \psi_2\}.$$

PROOF. \subseteq Let $\phi \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2)$. By Definition 8.4, we know that (1) $\phi = \text{project}(\mathfrak{s}, f)$ for some forest $f \in \text{forests}_{\Phi}(\mathfrak{h}_1 \uplus \mathfrak{h}_2)$ and (2) ϕ is delimited. By (2) and Lemma 8.12, f is delimited as well. Moreover, by Theorem 8.13, there exist \mathfrak{s} -delimited forests f_1 and f_2 with $f \in f_1 \bullet_{\mathbf{F}} f_2$ and, for $i \in \{1, 2\}$, $\text{heap}(f_i) = \mathfrak{h}_i$. By

Lemma 8.12, both $\psi_1 \triangleq \text{project}(\mathfrak{s}, \bar{f}_1)$ and $\psi_2 \triangleq \text{project}(\mathfrak{s}, \bar{f}_2)$ are delimited—hence, $\psi_1 \in \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_1)$ and $\psi_2 \in \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_2)$. Furthermore, by Theorem 7.39, we have $\phi = \text{project}(\mathfrak{s}, \bar{f}) \in \psi_1 \bullet_P \psi_2$.

\supseteq Assume there exist formulas $\phi \in \text{DUSH}_\Phi$, $\psi_1 \in \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_1)$, and $\psi_2 \in \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_2)$ such that $\phi \in \psi_1 \bullet_P \psi_2$. By Definition 8.4, there exist forests $\bar{f}, \bar{f}_1, \bar{f}_2$ such that, for $i \in \{1, 2\}$, we have $\psi_i = \text{project}(\mathfrak{s}, \bar{f}_i)$ and $\text{heap}(\bar{f}_i) = \mathfrak{h}_i$. Then, by Theorem 7.39, there exist forests \bar{f}_1, \bar{f}_2 such that $\bar{f}_1 \equiv_s \bar{f}_1, \bar{f}_2 \equiv_s \bar{f}_2, \bar{f} \in \bar{f}_1 \bullet_F \bar{f}_2$, and $\text{project}(\mathfrak{s}, \bar{f}) = \phi$. By Definition 7.37, we have, for $i \in \{1, 2\}$, $\text{heap}(\bar{f}_i) = \text{heap}(\bar{f}_i) = \mathfrak{h}_i$. Moreover, by Lemma 7.7, we have $\mathfrak{h}_1 \uplus \mathfrak{h}_2 = \text{heap}(\bar{f}_1 \uplus \bar{f}_2)$. Since $\bar{f} \in \bar{f}_1 \bullet_F \bar{f}_2$, Lemma 7.14 and Definition 7.15 yield that $\mathfrak{h}_1 \uplus \mathfrak{h}_2 = \text{heap}(\bar{f})$. Hence, $\bar{f} \in \text{forests}_\Phi(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2)$ and thus also $\phi \in \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2)$. \square

A.20 Proof of Lemma 8.16

We first show an auxiliary result, namely that the stack-allocated variables of a state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ correspond precisely to the roots of the \mathfrak{s} -decomposed forests of \mathfrak{h} :

LEMMA A.10. *Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state and let $\bar{f} \in \text{forests}_\Phi(\mathfrak{h})$ an \mathfrak{s} -delimited forest. Then, we have $\text{alloted}(\mathfrak{s}, \mathfrak{h}) = \{x \mid \mathfrak{s}(x) \in \text{roots}(\bar{f})\}$, where \bar{f} is the \mathfrak{s} -decomposition of \bar{f} .*

PROOF. By Lemma 7.14, $\text{heap}(\bar{f}) = \mathfrak{h}$ and thus, in particular, $\text{dom}(\bar{f}) = \text{dom}(\mathfrak{h})$. Consequently,

$$\mathfrak{s}(\text{alloted}(\mathfrak{s}, \mathfrak{h})) = \text{img}(\mathfrak{s}) \cap \text{dom}(\bar{f}).$$

By Lemma A.8, we have $\text{img}(\mathfrak{s}) \cap \text{dom}(\bar{t}) = \{\text{root}(\bar{t})\}$ for all $\bar{t} \in \bar{f}$. Hence,

$$\text{img}(\mathfrak{s}) \cap \text{dom}(\bar{f}) = \text{roots}(\bar{f}).$$

Overall, we thus have $\mathfrak{s}(\text{alloted}(\mathfrak{s}, \mathfrak{h})) = \text{roots}(\bar{f})$. By taking the inverse \mathfrak{s}^{-1} on both sides of the equation, we obtain that $\text{alloted}(\mathfrak{s}, \mathfrak{h}) = \{x \mid \mathfrak{s}(x) \in \text{roots}(\bar{f})\}$. \square

We restate the claim of Lemma 8.16: Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state with $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}) \neq \emptyset$. Then, $\text{alloted}(\mathfrak{s}, \mathfrak{h}) = \text{alloted}(\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}))$.

PROOF. By definition of DUSHs, all root parameters of all DUSHs in $\text{alloted}(\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}))$ are in $\text{img}(\mathfrak{s})$. Consequently, $\text{alloted}(\mathfrak{s}, \mathfrak{h}) \supseteq \text{alloted}(\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}))$.

For the other implication, let \bar{f} be a forest with $\text{heap}(\bar{f}) = \mathfrak{h}$ and $\text{project}(\mathfrak{s}, \bar{f}) \in \text{type}_\Phi(\mathfrak{s}, \mathfrak{h})$. Such a forest must exist, as $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}) \neq \emptyset$ by assumption. Let \bar{f} be the \mathfrak{s} -decomposition of \bar{f} . By Lemma A.7, \bar{f} is delimited and by Lemma 7.14, $\text{heap}(\bar{f}) = \mathfrak{h}$, implying $\text{project}(\mathfrak{s}, \bar{f}) \in \text{type}_\Phi(\mathfrak{s}, \mathfrak{h})$ and we can apply Lemma A.10 to obtain that

$$\text{alloted}(\mathfrak{s}, \mathfrak{h}) = \{x \mid \mathfrak{s}(x) \in \text{roots}(\bar{f})\}.$$

Consequently, all variables in $\text{alloted}(\mathfrak{s}, \mathfrak{h})$ occur as root parameters on the right-hand side of magic wands in $\text{project}(\mathfrak{s}, \bar{f})$. Therefore, $\text{alloted}(\mathfrak{s}, \mathfrak{h}) \subseteq \text{alloted}(\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}))$. \square

A.21 Proof of Lemma 8.17

Claim. Let $\phi \in \text{GSL}$ and let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state with $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_\Phi \phi$. Then, $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}) \neq \emptyset$.

PROOF. By Corollary 4.6, $\langle \mathfrak{s}, \mathfrak{h} \rangle$ is a guarded state. Lemma 4.5 then yields that there exist $k \geq 1$ predicate calls such that $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_\Phi \star_{1 \leq i \leq k} \text{pred}_i(\mathbf{x}_i)$.

We split the heap \mathfrak{h} into disjoint heaps $\mathfrak{h}_1 \uplus \dots \uplus \mathfrak{h}_k$ such that $\langle \mathfrak{s}, \mathfrak{h}_i \rangle \models_\Phi \text{pred}_i(\mathbf{x}_i)$ for each $i \in [1, k]$. Next, consider the forest $\bar{f} \triangleq \{t_1, \dots, t_k\}$, where each t_i is a Φ -tree with $\text{heap}(t_i) = \mathfrak{h}_i$; such trees exist by Lemma 7.3. Observe further that each of these trees is delimited, because they do not have holes and their root is in $\mathfrak{s}(\mathbf{x}_i)$. By Lemma 7.7, we have $\text{heap}(\bar{f}) = \mathfrak{h}$. Finally, Lemma 7.25 yields $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_\Phi \text{project}(\mathfrak{s}, \bar{f})$ and thus $\text{project}(\mathfrak{s}, \bar{f}) \in \text{type}_\Phi(\mathfrak{s}, \mathfrak{h})$. Hence, $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}) \neq \emptyset$. \square

A.22 Proof of Corollary 8.19

Claim. For guarded states $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle$ and $\langle \mathfrak{s}, \mathfrak{h}_2 \rangle$ with $\mathfrak{h}_1 \uplus \mathfrak{h}_2 \neq \perp$, we have $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2) = \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_1) \bullet \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_2)$.

PROOF. We need to show that $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_1) \bullet \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_2) \neq \perp$. Assume that $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_i) = \emptyset$ for $i = 1$ or $i = 2$. Then, $\text{allocated}(\mathcal{T}_i) = \emptyset$ and we get that $\text{allocated}(\mathcal{T}_1) \cap \text{allocated}(\mathcal{T}_2) = \emptyset$. Otherwise, we have $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_i) \neq \emptyset$ for $i = 1, 2$. Then, $\text{allocated}(\mathfrak{s}, \mathfrak{h}_i) = \text{allocated}(\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_i))$. $\mathfrak{h}_1 \uplus \mathfrak{h}_2 \neq \perp$ then implies that $\text{allocated}(\mathcal{T}_1) \cap \text{allocated}(\mathcal{T}_2) = \emptyset$. The claim then follows from Theorem 8.14. \square

A.23 Proof of Lemma 8.20

Claim. For $i \in \{1, 2\}$, let $\langle \mathfrak{s}, \mathfrak{h}_i \rangle$ be states with $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_i) = \mathcal{T}_i \neq \emptyset$ and $\mathcal{T}_1 \bullet \mathcal{T}_2 \neq \perp$. Then, there are states $\langle \mathfrak{s}, \mathfrak{h}'_i \rangle$ such that $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}'_i) = \mathcal{T}_i$ and $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}'_1 \uplus \mathfrak{h}'_2) = \mathcal{T}_1 \bullet \mathcal{T}_2$.

PROOF. We choose some states $\langle \mathfrak{s}, \mathfrak{h}'_i \rangle$ that are isomorphic to $\langle \mathfrak{s}, \mathfrak{h}_i \rangle$ such that $\text{locs}(\mathfrak{h}'_1) \cap \text{locs}(\mathfrak{h}'_2) \subseteq \text{img}(\mathfrak{s})$. We have that $\langle \mathfrak{s}, \mathfrak{h}'_i \rangle = \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}_i) = \mathcal{T}_i$ because isomorphic states have the same types (observe that the stack-projection replaces location that are not in the image of the stack by quantified variables). By Lemma 8.16, we have $\text{allocated}(\mathcal{T}_i) = \text{allocated}(\langle \mathfrak{s}, \mathfrak{h}'_i \rangle) = \text{allocated}(\mathfrak{s}, \mathfrak{h}'_i)$. Thus, we get $\mathfrak{h}'_1 \uplus \mathfrak{h}'_2 \neq \perp$ from $\mathcal{T}_1 \bullet \mathcal{T}_2 \neq \perp$ and $\text{locs}(\mathfrak{h}'_1) \cap \text{locs}(\mathfrak{h}'_2) \subseteq \text{img}(\mathfrak{s})$. Then, Theorem 8.14 yields that $\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}'_1 \uplus \mathfrak{h}'_2) = \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}'_1) \bullet \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}'_2)$. \square

A.24 Proof of Lemma 8.22

Claim. For $\mathfrak{x}, \mathfrak{y}$ as above and a stack \mathfrak{s} with $\mathfrak{y} \subseteq \text{dom}(\mathfrak{s})$ and $\mathfrak{x} \cap \text{dom}(\mathfrak{s}) = \emptyset$, we have

$$\text{type}_\Phi(\mathfrak{s}[\mathfrak{x}/\mathfrak{y}], \mathfrak{h})[\text{aliasing}(\mathfrak{s}) : \mathfrak{x}/\mathfrak{y}] = \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}).$$

PROOF. Let \mathfrak{y}' be the sequence obtained by replacing every variable in $\mathfrak{y} \in \mathfrak{y}$ by the maximal variable y' with $\mathfrak{s}(y') = \mathfrak{s}(y)$. We consider some $\phi \in \text{type}_\Phi(\mathfrak{s}[\mathfrak{x}/\mathfrak{y}], \mathfrak{h})$. Then, there exists a forest $\mathfrak{f} \in \text{forests}_\Phi(\mathfrak{h})$ such that $\phi = \text{project}(\mathfrak{s}[\mathfrak{x}/\mathfrak{y}], \mathfrak{f})$. By construction of stack-projections (cf. Definition 7.23), we obtain that

$$\phi[\mathfrak{x}/\mathfrak{y}'] = \text{project}(\mathfrak{s}[\mathfrak{x}/\mathfrak{y}], \mathfrak{f})[\mathfrak{x}/\mathfrak{y}'] = \text{project}(\mathfrak{s}, \mathfrak{f}) \in \text{type}_\Phi(\mathfrak{s}, \mathfrak{h}).$$

The converse direction is analogous. \square

A.25 Proof of Lemma 8.24

Claim. Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a guarded state such that $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h})$ holds for some variable x . Then,

$$\text{forget}_{\text{aliasing}(\mathfrak{s}), x}(\text{type}_\Phi(\mathfrak{s}, \mathfrak{h})) = \text{type}_\Phi(\mathfrak{s}[x/\perp], \mathfrak{h}).$$

PROOF. We will use the following fact (\dagger) based on the construction of projections (cf. Definition 7.23): If $\mathfrak{f} \in \text{forests}_\Phi(\mathfrak{h})$, then $\mathfrak{s}(x)$ is replaced in $\text{project}(\mathfrak{s}[x/\perp], \mathfrak{f})$ by an existentially-quantified variable iff x is replaced in $\text{forget}_{\text{aliasing}(\mathfrak{s}), x}(\text{project}(\mathfrak{s}, \mathfrak{f}))$ by an existentially-quantified variable.

Now, consider some $\phi \in \text{type}_\Phi(\mathfrak{s}[x/\perp], \mathfrak{h})$. Then, there is some $\mathfrak{f} \in \text{forests}_\Phi(\mathfrak{h})$ such that $\text{project}(\mathfrak{s}[x/\perp], \mathfrak{f}) = \phi \in \text{DUSH}_\Phi$. Clearly, we have $\text{project}(\mathfrak{s}, \mathfrak{f}) \in \text{DUSH}_\Phi$ and hence $\text{project}(\mathfrak{s}, \mathfrak{f}) \in \text{type}_\Phi(\mathfrak{s}, \mathfrak{h})$. Applying (\dagger), we conclude that

$$\phi = \text{forget}_{\text{aliasing}(\mathfrak{s}), x}(\text{project}(\mathfrak{s}, \mathfrak{f})) \in \text{forget}_{\text{aliasing}(\mathfrak{s}), x}(\text{type}_\Phi(\mathfrak{s}, \mathfrak{h})).$$

Conversely, let $\phi \in \text{forget}_{\text{aliasing}(\mathfrak{s}), x}(\text{type}_\Phi(\mathfrak{s}, \mathfrak{h}))$. By construction, $\phi \in \text{DUSH}_\Phi$ and there is some $\mathfrak{f} \in \text{forests}_\Phi(\mathfrak{h})$ such that $\text{forget}_{\text{aliasing}(\mathfrak{s}), x}(\text{project}(\mathfrak{s}, \mathfrak{f})) = \phi$. Applying (\dagger), we can conclude that $\phi = \text{project}(\mathfrak{s}[x/\perp], \mathfrak{f}) \in \text{type}_\Phi(\mathfrak{s}[x/\perp], \mathfrak{h})$. \square

A.26 Proof of Lemma 8.26

Claim. For every state $\langle s, h \rangle$, variable x with $s(x) \notin \text{locs}(h)$ and aliasing $(s)(x) = \{x\}$,

$$\text{extend}_x(\text{type}_\Phi(s[x/\perp], h)) = \text{type}_\Phi(s, h).$$

PROOF. Assume ϕ is the x -extension of some $\phi' \in \text{type}_\Phi(s[x/\perp], h)$. Then there exists some $\bar{f} \in \text{forests}_\Phi(h)$ such that $\phi' = \text{project}(s[x/\perp], \bar{f}) \in \text{DUSH}_\Phi$. We can now choose a forest \bar{f}' with $\bar{f} \equiv_s \bar{f}'$ such that $\phi = \text{project}(s, \bar{f}')$. Because of $\bar{f} \equiv_s \bar{f}'$ and $\text{project}(s[x/\perp], \bar{f}) \in \text{DUSH}_\Phi$ we get that $\text{project}(s, \bar{f}') \in \text{DUSH}_\Phi$. Hence, $\phi \in \text{type}_\Phi(s, h)$.

Conversely, let $\phi \in \text{type}_\Phi(s, h)$. Then there exists some $\bar{f} \in \text{forests}_\Phi(h)$ such that $\phi = \text{project}(s, \bar{f}) \in \text{DUSH}_\Phi$. We note that $s(x) \notin \text{interface}(\bar{f})$ since $s(x) \notin \text{locs}(h)$. Hence, $\text{project}(s[x/\perp], \bar{f}) \in \text{type}_\Phi(s[x/\perp], h)$. We distinguish two cases: First, if $x \notin \text{fvvars}(\phi)$, then $\phi = \text{project}(s, \bar{f}) = \text{project}(s[x/\perp], \bar{f}) \in \text{type}_\Phi(s[x/\perp], h)$. Second, if $x \in \text{fvvars}(\phi)$, then $s(x)$ corresponds to a universally quantified variable in $\text{project}(s[x/\perp], \bar{f})$. Hence, ϕ is an x -instantiation of $\text{project}(s[x/\perp], \bar{f})$. By Definition 8.25, this means $\phi \in \text{extend}_x(\text{type}_\Phi(s[x/\perp], h))$. \square

A.27 Proof of Theorem 8.29

For a concise formalization, we assume—in addition to our global assumptions stated in Section 3.4.5—that all formulas ϕ under consideration are **GSL** formulas without constant locations.

We will prove Theorem 8.29 by structural induction on the syntax of **GSL** formulas. For most base cases—those that involve the heap—we rely on the fact that a state satisfies all formulas in its type.

LEMMA A.11. *If $\phi \in \text{type}_\Phi(s, h)$ for some state $\langle s, h \rangle$, then $\langle s, h \rangle \models_\Phi \phi$.*

PROOF. Since $\phi \in \text{type}_\Phi(s, h)$, there exists a Φ -forest \bar{f} with $\text{heap}(\bar{f}) = h$ and $\text{project}(s, \bar{f}) = \phi$. By Lemma 7.25, we have $\langle s, \text{heap}(\bar{f}) \rangle \models_\Phi \text{project}(s, \bar{f})$ and thus also $\langle s, h \rangle \models_\Phi \phi$. \square

Finally, to deal with the separating conjunction, we need another auxiliary result. In Corollary 8.19, we showed how two types can be composed into a single one, i.e.,

$$\text{type}_\Phi(s, h_1) \bullet \text{type}_\Phi(s, h_2) = \text{type}_\Phi(s, h_1 \uplus h_2).$$

To prove Theorem 8.29, we need the reverse: Given a composed type, say $\text{type}_\Phi(s, h) = \mathcal{T}_1 \bullet \mathcal{T}_2$, we need to *decompose* h into two heaps whose types (in conjunction with stack s) are \mathcal{T}_1 and \mathcal{T}_2 .

LEMMA A.12 (TYPE DECOMPOSABILITY). *Let $\langle s, h \rangle$ be a state with $\emptyset \neq \text{type}_\Phi(s, h) = \mathcal{T}_1 \bullet \mathcal{T}_2$. Then, there exist h_1, h_2 such that $h = h_1 \uplus h_2$, $\mathcal{T}_1 = \text{type}_\Phi(s, h_1)$, and $\mathcal{T}_2 = \text{type}_\Phi(s, h_2)$.*

Since proving type decomposability involves a bit more technical machinery, we refer the interested reader to Appendix A.28 for a detailed proof.

With the above three lemmas at hand, we can now prove the refinement theorem.

Claim (Refinement theorem). For all stacks s , heaps h_1, h_2 , and **GSL** formulas ϕ ,

$$\text{type}_\Phi(s, h_1) = \text{type}_\Phi(s, h_2) \quad \text{implies} \quad \langle s, h_1 \rangle \models_\Phi \phi \quad \text{iff} \quad \langle s, h_2 \rangle \models_\Phi \phi.$$

PROOF. We only show that if $\langle s, h_1 \rangle \models_\Phi \phi$ then $\langle s, h_2 \rangle \models_\Phi \phi$; the converse direction is symmetric. We proceed by induction on the structure of the **GSL** formula ϕ . Let us assume that $\langle s, h_1 \rangle \models_\Phi \phi$.

Case $\phi = \text{emp}$. By the semantics of **emp**, we have $h_1 = \emptyset$. Let \bar{f} be the empty forest. Then $\bar{f} \in \text{forests}_\Phi(h_1)$ and thus $\text{emp} = \text{project}(s, \bar{f}) \in \text{type}_\Phi(s, h_1) = \text{type}_\Phi(s, h_2)$. Hence, by Lemma A.11, we have $\langle s, h_2 \rangle \models_\Phi \text{emp}$.

Cases $\phi = x \approx y$, $\phi = x \not\approx y$. We observe that the states $\langle s, h_1 \rangle$ and $\langle s, h_2 \rangle$ have the same stack. Then we proceed as in the case for $\phi = \text{emp}$.

Case $\phi = x \mapsto \langle y_1, \dots, y_k \rangle$. By assumption, Φ is pointer-closed (see Definition 3.7), i.e., $\langle s, h_1 \rangle \models_{\Phi} \text{ptr}_k(x, y_1, \dots, y_k)$. We define a Φ -forest $\mathfrak{f} = \{t\}$, where t is

$$t = \{s(x) \mapsto \langle \emptyset, \text{ptr}_k(s(x), s(y_1), \dots, s(y_k)) \Leftarrow s(x) \mapsto \langle s(y_1), \dots, s(y_k) \rangle \rangle\}.$$

Observe that $\mathfrak{f} \in \text{forests}_{\Phi}(h_1)$ and

$$\text{ptr}_k(x, y_1, \dots, y_k) = \text{project}(s, \mathfrak{f}) \in \text{type}_{\Phi}(s, h_1) = \text{type}_{\Phi}(s, h_2).$$

Hence, by Lemma A.11, we have $\langle s, h_2 \rangle \models_{\Phi} \text{ptr}_k(x, y_1, \dots, y_k)$. By definition of the predicate ptr_k , we conclude that $\langle s, h_2 \rangle \models_{\Phi} x \mapsto \langle y_1, \dots, y_k \rangle$.

Case $\phi = \text{pred}(z_1, \dots, z_k)$. By Lemma 7.3, there exists a Φ -tree t such that $\text{rootpred}(t) = \text{pred}(s(z_1), \dots, s(z_k))$, $\text{allholepreds}(t) = \emptyset$, and $\text{heap}(\{t\}) = h_1$. Let

$$\psi \triangleq \text{pred}(s(z_1), \dots, s(z_k))[\text{dom}(s_{\max}^{-1})/\text{img}(s_{\max}^{-1})] = \text{project}(s, \{t\}).$$

Then, $\psi \in \text{type}_{\Phi}(s, h_1) = \text{type}_{\Phi}(s, h_2)$ and, by Lemma A.11, $\langle s, h_2 \rangle \models_{\Phi} \psi$. Observe that while $\psi \neq \text{pred}(z)$ is possible, we have by definition of s_{\max}^{-1} that the parameters of the predicate call in ψ evaluate to the same locations as the parameters z . Hence, $\langle s, h_2 \rangle \models_{\Phi} \text{pred}(z_1, \dots, z_k)$.

Case $\phi = \phi_1 \wedge \phi_2$. We then have $\langle s, h_1 \rangle \models_{\Phi} \phi_1$ and $\langle s, h_1 \rangle \models_{\Phi} \phi_2$. By I.H., $\langle s, h_2 \rangle \models_{\Phi} \phi_1$ and $\langle s, h_2 \rangle \models_{\Phi} \phi_2$. Hence, $\langle s, h_2 \rangle \models_{\Phi} \phi_1 \wedge \phi_2$.

Cases $\phi = \phi_1 \vee \phi_2$, $\phi = \phi_1 \wedge \neg \phi_2$. Analogous to previous case.

Case $\phi = \phi_1 \star \phi_2$. By the semantics of \star , there exist heaps $h_{1,1}$ and $h_{1,2}$ such that $\langle s, h_{1,i} \rangle \models_{\Phi} \phi_i$ for $1 \leq i \leq 2$.

Let $\mathcal{T}_i \triangleq \text{type}_{\Phi}(s, h_{1,i})$. By Corollary 4.6, we have that $\langle s, h_{1,i} \rangle \in \text{GStates}$ for $1 \leq i \leq 2$. By Corollary 8.19 we have that $\text{type}_{\Phi}(s, h_1 \uplus h_2) = \mathcal{T}_1 \bullet \mathcal{T}_2$. By Lemma 8.17, we have that $\text{type}_{\Phi}(s, h_1) \neq \emptyset$. We can then apply Lemma A.12 to $\langle s, h_2 \rangle$, \mathcal{T}_1 and \mathcal{T}_2 in order to obtain states $\langle s, h_{2,1} \rangle$ and $\langle s, h_{2,2} \rangle$ with $h_2 = h_{2,1} \uplus h_{2,2}$, $\text{type}_{\Phi}(s, h_{2,1}) = \mathcal{T}_1$, and $\text{type}_{\Phi}(s, h_{2,2}) = \mathcal{T}_2$.

We can thus apply the I.H. to both $h_{1,1}$, $h_{1,2}$, ϕ_1 and $h_{2,1}$, $h_{2,2}$, ϕ_2 to conclude that $\langle s, h_{2,1} \rangle \models_{\Phi} \phi_1$ and $\langle s, h_{2,2} \rangle \models_{\Phi} \phi_2$. Since $h_{2,1} \uplus h_{2,2} = h_2$, the semantics of \star then yields $\langle s, h_2 \rangle \models_{\Phi} \phi$.

Case $\phi = \phi_0 \wedge (\phi_1 \oplus \phi_2)$. Then there exists a heap h_0 with $\langle s, h_0 \rangle \models_{\Phi} \phi_1$ and $\langle s, h_1 \uplus h_0 \rangle \models_{\Phi} \phi_2$.

Since $\langle s, h_1 \rangle$ and $\langle s, h_2 \rangle$ have the same type, we have $\text{allocated}(s, h_1) = \text{allocated}(s, h_2)$. We can therefore assume w.l.o.g. that $h_2 \uplus h_0$ is defined—if this is not the case, simply replace h_0 with a heap h'_0 such that $\langle s, h_0 \rangle \cong \langle s, h'_0 \rangle$ and both $h_1 \uplus h'_0$ and $h_2 \uplus h'_0$ are defined. Then, by Lemma 3.5, we can conclude that $\langle s, h_1 \uplus h'_0 \rangle \models_{\Phi} \phi$.

By Corollary 4.6 we have $\langle s, h_0 \rangle, \langle s, h_1 \rangle \in \text{GStates}$. Corollary 8.19 then yields that

$$\begin{aligned} \text{type}_{\Phi}(s, h_1 \uplus h_0) &= \text{type}_{\Phi}(s, h_1) \bullet \text{type}_{\Phi}(s, h_0) \\ &= \text{type}_{\Phi}(s, h_2) \bullet \text{type}_{\Phi}(s, h_0) \\ &= \text{type}_{\Phi}(s, h_2 \uplus h_0). \end{aligned}$$

Now, we apply the I.H. for ϕ_0 , h_1 and h_2 to conclude that $\langle s, h_2 \rangle \models_{\Phi} \phi_0$, as well as for ϕ_2 , $\langle s, h_1 \uplus h_0 \rangle$ and $\langle s, h_2 \uplus h_0 \rangle$ to conclude that $\langle s, h_2 \uplus h_0 \rangle \models_{\Phi} \phi_2$. Hence, by the semantics of \oplus and \wedge , we have $\langle s, h_2 \rangle \models_{\Phi} \phi_0 \wedge (\phi_1 \oplus \phi_2)$.

Case $\phi = \phi_0 \wedge (\phi_1 \star \phi_2)$. Analogous to the previous case for guarded septraction, except that we must consider *arbitrary* models h_0 with $\langle s, h_0 \rangle \models_{\Phi} \phi_1$ and $\langle s, h_1 \uplus h_0 \rangle \models_{\Phi} \phi_2$. \square

A.28 Proof of Lemma A.12 (type decomposability)

We need a couple of auxiliary definitions and lemmata before we can show this result in Lemma A.12 at the end of this section.

DEFINITION A.13. We call \bar{f} \mathfrak{s} -decomposed iff $\bar{f} = \text{split}(\bar{f}, \text{img}(\mathfrak{s}))$.

LEMMA A.14. Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state with $\emptyset \neq \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$ and let $\mathcal{T}_1, \mathcal{T}_2 \in \mathbf{Types}_{\Phi}^{\text{aliasing}(\mathfrak{s})}$ be types with $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) = \mathcal{T}_1 \bullet \mathcal{T}_2$. Then there exist \mathfrak{s} -decomposed, \mathfrak{s} -delimited Φ -forests $\bar{f}, \bar{f}_1, \bar{f}_2$ such that

- (1) $\text{project}(\mathfrak{s}, \bar{f}_i) \in \mathcal{T}_i, 1 \leq i \leq 2$,
- (2) $\bar{f}_1 \uplus \bar{f}_2 = \bar{f}$, and
- (3) $\text{heap}(\bar{f}) = \mathfrak{h}$.

PROOF. By assumption we have $\emptyset \neq \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$. Take an arbitrary forest \bar{f} with $\text{project}(\mathfrak{s}, \bar{f}) \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$. By definition, \bar{f} is \mathfrak{s} -delimited. Let $\bar{f} \triangleq \text{split}(\bar{f}, \text{img}(\mathfrak{s}))$ be the \mathfrak{s} -decomposition of \bar{f} . By Lemma A.7, \bar{f} is \mathfrak{s} -delimited. By Lemma 7.13, $\bar{f} \blacktriangleright^* \bar{f}$. Lemma 7.14 thus gives us that $\text{heap}(\bar{f}) = \mathfrak{h}$. Hence, $\text{project}(\mathfrak{s}, \bar{f}) \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$.

By definition of \bullet , there exist formulas $\psi_1 \in \mathcal{T}_1, \psi_2 \in \mathcal{T}_2$ with $\text{project}(\mathfrak{s}, \bar{f}) \in \psi_1 \bullet_{\mathbf{P}} \psi_2$. By definition, there exist Φ -forests $\mathfrak{g}_1, \mathfrak{g}_2$ with $\text{project}(\mathfrak{s}, \mathfrak{g}_i) = \psi_i$. Because $\mathcal{T}_1 \bullet \mathcal{T}_2$ is defined, we have $\text{allocated}(\mathcal{T}_1) \cap \text{allocated}(\mathcal{T}_2) = \emptyset$, allowing us to assume w.l.o.g. that $\mathfrak{g}_1 \uplus \mathfrak{g}_2 \neq \perp$. By Theorem 7.39, there then exist forests \bar{f}_1, \bar{f}_2 with $\text{project}(\mathfrak{s}, \bar{f}_i) = \psi_i$ and $\bar{f} \in \bar{f}_1 \bullet_{\mathbf{F}} \bar{f}_2$. Because \bar{f} is \mathfrak{s} -decomposed, this implies that zero \blacktriangleright -steps were taken by $\bullet_{\mathbf{F}}$, i.e., $\bar{f} = \bar{f}_1 \uplus \bar{f}_2$. Moreover, because \bar{f} is \mathfrak{s} -decomposed and \mathfrak{s} -delimited, we get that \bar{f}_1 and \bar{f}_2 are \mathfrak{s} -decomposed and \mathfrak{s} -delimited as well. \square

DEFINITION A.15 (ROOTS OF A DUSH). Let $\phi = \exists e. \forall a. \star_{1 \leq i \leq k} (\zeta_i \star \text{pred}_i(z_i))$ be a DUSH. The roots of ψ are the set

$$\text{dushroots}_{\mathfrak{s}}(\phi) \triangleq \bigcup_{1 \leq i \leq k} \text{aliasing}(\mathfrak{s})(\text{predroot}(\text{pred}_i(z_i))).$$

Clearly, the roots of a forest are connected to the roots of a DUSH via the stack:

LEMMA A.16. Let \bar{f} be a Φ -forest. Then, $\text{dushroots}_{\mathfrak{s}}(\text{project}(\mathfrak{s}, \bar{f})) = \{x \mid \mathfrak{s}(x) \in \text{roots}(\bar{f})\}$.

PROOF. Let $\phi \triangleq \text{project}(\mathfrak{s}, \bar{f})$. By definition of DUSHs, $\text{roots}(\bar{f}) \subseteq \text{img}(\mathfrak{s})$. Every root $l \in \text{roots}(\bar{f})$ is therefore replaced by a variable in $\mathfrak{s}_{\max}^{-1}(l)$ by stack-forest projection. Since $\text{dushroots}_{\mathfrak{s}}(\phi)$ closes the set of roots under all aliasing(\mathfrak{s})(\cdot), we obtain that $\text{dushroots}_{\mathfrak{s}}(\phi)$ contains all variables x with $\mathfrak{s}(x) \in \text{roots}(\bar{f})$. \square

LEMMA A.17. Let \mathcal{T} be a Φ -type and $\psi \in \mathcal{T}$. There exists a formula $\psi' \in \mathcal{T}$ such that $\psi' \triangleright^* \psi$ and $\text{allocated}(\mathcal{T}) = \text{dushroots}_{\mathfrak{s}}(\psi')$.

PROOF. Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be such that $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) = \mathcal{T}$. By Lemma 8.16 we then have that $\text{allocated}(\mathfrak{s}, \mathfrak{h}) = \text{allocated}(\mathcal{T})$ (\dagger). By definition of Φ -types, there then exists a Φ -forest \bar{f} with $\text{heap}(\bar{f}) = \mathfrak{h}$ and $\text{project}(\mathfrak{s}, \bar{f}) = \psi$. Let $\bar{f} \triangleq \text{split}(\bar{f}, \text{img}(\mathfrak{s}))$ be the \mathfrak{s} -decomposition of \bar{f} and write $\psi' \triangleq \text{project}(\mathfrak{s}, \bar{f})$. We show that ψ' has the desired properties.

First, $\bar{f} \blacktriangleright^* \bar{f}$ by Lemma 7.13. Observe that $\bar{f} \in \bar{f} \bullet_{\mathbf{F}} \emptyset$ (where \emptyset is the empty forest). Lemma 7.35 therefore guarantees that $\psi \in \psi' \bullet_{\mathbf{P}} \text{emp}$ and thus $\psi' \triangleright^* \psi$.

Second, by Lemma 7.14, $\text{heap}(\bar{f}) = \text{heap}(\bar{f})$ and thus $\psi' \in \mathcal{T}$. Moreover, by Lemma A.10, $\text{allocated}(\mathfrak{s}, \mathfrak{h}) = \{x \mid \mathfrak{s}(x) \in \text{roots}(\bar{f})\}$. We combine the above with (\dagger) to derive $\text{allocated}(\mathcal{T}) = \{x \mid \mathfrak{s}(x) \in \text{roots}(\bar{f})\}$. Lemma A.16 then yields that $\text{allocated}(\mathcal{T}) = \text{dushroots}_{\mathfrak{s}}(\psi')$. \square

Claim (Lemma A.12). Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state with $\emptyset \neq \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) = \mathcal{T}_1 \bullet \mathcal{T}_2$. Then, there exist $\mathfrak{h}_1, \mathfrak{h}_2$ such that $\mathfrak{h} = \mathfrak{h}_1 \uplus \mathfrak{h}_2$, $\mathcal{T}_1 = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1)$, and $\mathcal{T}_2 = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_2)$.

PROOF. Let $\mathcal{T} \triangleq \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$. By Lemma A.14, there exist \mathfrak{s} -decomposed, \mathfrak{s} -delimited forests $\bar{f}, \bar{f}_1, \bar{f}_2$ with

- (1) $\text{project}(\mathfrak{s}, \bar{f}_i) \in \mathcal{T}_i, 1 \leq i \leq 2$,
- (2) $\bar{f}_1 \uplus \bar{f}_2 = \bar{f}$, and
- (3) $\text{heap}(\bar{f}) = \mathfrak{h}$.

Define $\mathfrak{h}_1 \triangleq \text{heap}(\mathfrak{f}_1)$, $\mathfrak{h}_2 \triangleq \text{heap}(\mathfrak{f}_2)$. Then, $\mathfrak{h}_1 \uplus \mathfrak{h}_2 = \mathfrak{h}$ because of $\text{heap}(\mathfrak{f}_1 \uplus \mathfrak{f}_2) = \mathfrak{h}$. Because the \mathfrak{f}_i are \mathfrak{s} -delimited, we have $\text{dangling}(\mathfrak{h}_i) \subseteq \text{img}(\mathfrak{s})$. Hence, $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle, \langle \mathfrak{s}, \mathfrak{h}_2 \rangle \in \mathbf{GStates}$. By Lemma A.10 and Lemma 8.16 we have

$$\text{allocated}(\mathfrak{s}, \mathfrak{h}) = \{x \mid \mathfrak{s}(x) \in \text{roots}(\mathfrak{f})\} = \text{allocated}(\mathcal{T}). \quad (*)$$

Further, we have

$$\text{allocated}(\mathfrak{s}, \mathfrak{h}_i) = \{x \mid \mathfrak{s}(x) \in \text{roots}(\mathfrak{f}_i)\} = \text{allocated}(\mathcal{T}_i), \quad (\dagger)$$

where the first equality follows from Lemma A.10, and the second equality holds by (*) and because the \mathfrak{f}_i are \mathfrak{s} -decomposed.

We will show that $\mathcal{T}_1 = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1)$; the argument for $\mathcal{T}_2 = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_2)$ is symmetrical. We prove the inclusions $\mathcal{T}_1 \subseteq \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1)$ and $\mathcal{T}_1 \supseteq \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1)$ separately.

“ $\mathcal{T}_1 \subseteq \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1)$.” Let $\psi_1 \in \mathcal{T}_1$. By Lemma A.17, there exists a formula $\psi'_1 \in \mathcal{T}_1$ such that $\psi'_1 \triangleright^* \psi_1$, and $\text{allocated}(\mathcal{T}_1) = \text{dushoots}_{\mathfrak{s}}(\psi'_1)$ (#). Let $\psi'_2 \triangleq \text{project}(\mathfrak{s}, \mathfrak{f}_2) \in \mathcal{T}_2$. By the definition of projections, we have $\psi'_i = \exists \mathbf{e}_i. \forall \mathbf{a}_i. \phi_i$ for some $\mathbf{e}_i, \mathbf{a}_i, \phi_i$, where we can assume w.l.o.g. that $\mathbf{e}_1, \mathbf{e}_2, \mathbf{a}_1, \mathbf{a}_2$ are pairwise disjoint. By definition of type composition, \bullet , it follows that $\psi \triangleq \exists \mathbf{e}_1, \mathbf{e}_2. \forall \mathbf{a}_1, \mathbf{a}_2. \phi_1 \star \phi_2 \in \mathcal{T}$. Then, there is an Φ -forest $\mathfrak{g} \in \text{forests}_{\Phi}(\mathfrak{s}, \mathfrak{h})$ such that $\text{project}(\mathfrak{s}, \mathfrak{g}) = \psi$. From (#), (*) and (\dagger) we obtain that $\text{allocated}(\mathfrak{s}, \mathfrak{h}) = \text{dushoots}_{\mathfrak{s}}(\psi)$. By Theorem 8.13 there exist Φ -forests $\mathfrak{g}_1, \mathfrak{g}_2$ with $\mathfrak{g} = \mathfrak{g}_1 \uplus \mathfrak{g}_2$, $\text{heap}(\mathfrak{g}_i) = \mathfrak{h}_i$ and $\mathfrak{g} \in \mathfrak{g}_1 \bullet_{\mathfrak{F}} \mathfrak{g}_2$. With $\text{project}(\mathfrak{s}, \mathfrak{g}_1 \uplus \mathfrak{g}_2) = \exists \mathbf{e}_1, \mathbf{e}_2. \forall \mathbf{a}_1, \mathbf{a}_2. \phi_1 \star \phi_2$ we then must have that $\text{project}(\mathfrak{s}, \mathfrak{g}_i) = \psi'_i$. Therefore, $\psi'_1 \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1)$. Since $\psi'_1 \triangleright^* \psi_1$, Lemma A.4 then gives us a forest \mathfrak{g}'_1 s.t. $\mathfrak{g}_1 \blacktriangleright^* \mathfrak{g}'_1$ and $\text{project}(\mathfrak{s}, \mathfrak{g}'_1) = \psi_1$. Because also $\text{heap}(\mathfrak{g}'_1) = \text{heap}(\mathfrak{g}_1)$ by Lemma 7.14, $\mathfrak{g}'_1 \in \text{forests}_{\Phi}(\mathfrak{h}_1)$ and $\text{project}(\mathfrak{s}, \mathfrak{g}'_1) \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1)$. Thus, $\psi_1 \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1)$.

“ $\mathcal{T}_1 \supseteq \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1)$.” Let $\psi_1 \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1)$. By Lemma A.17, there exists a formula $\psi'_1 \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1)$ such that $\psi'_1 \triangleright^* \psi_1$, and $\text{allocated}(\mathcal{T}_1) = \text{dushoots}_{\mathfrak{s}}(\psi'_1)$ (#). Let $\psi'_2 \triangleq \text{project}(\mathfrak{s}, \mathfrak{f}_2) \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_2)$. By the definition of projections, we have $\psi'_i = \exists \mathbf{e}_i. \forall \mathbf{a}_i. \phi_i$ for some $\mathbf{e}_i, \mathbf{a}_i, \phi_i$, where we can assume w.l.o.g. that $\mathbf{e}_1, \mathbf{e}_2, \mathbf{a}_1, \mathbf{a}_2$ are pairwise disjoint. By Corollary 8.19 we then have that $\psi \triangleq \exists \mathbf{e}_1, \mathbf{e}_2. \forall \mathbf{a}_1, \mathbf{a}_2. \phi_1 \star \phi_2 \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1) \bullet \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_2) = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1 \uplus \mathfrak{h}_2) = \mathcal{T}$. From (#), (*) and (\dagger) we obtain that $\text{allocated}(\mathfrak{s}, \mathfrak{h}) = \text{dushoots}_{\mathfrak{s}}(\psi)$. Because of $\mathcal{T} = \mathcal{T}_1 \bullet \mathcal{T}_2$ there are some $\psi''_i \in \mathcal{T}_i$ such that $\psi \in \psi''_1 \bullet_{\mathfrak{P}} \psi''_2$. Because of $\text{allocated}(\mathfrak{s}, \mathfrak{h}_i) = \text{allocated}(\mathcal{T}_i)$ we must have that $\psi''_i = \psi'_i$. Therefore, $\psi'_1 \in \mathcal{T}_1$. We now consider some forest \mathfrak{g}_1 with $\text{type}_{\Phi}(\mathfrak{s}, \text{heap}(\mathfrak{g}_1)) = \mathcal{T}_1$ and $\text{project}(\mathfrak{s}, \mathfrak{g}'_1) = \psi'_1$. Since $\psi'_1 \triangleright^* \psi_1$, Lemma A.4 then gives us a forest \mathfrak{g}'_1 s.t. $\mathfrak{g}_1 \blacktriangleright^* \mathfrak{g}'_1$ and $\text{project}(\mathfrak{s}, \mathfrak{g}'_1) = \psi_1$. Because also $\text{heap}(\mathfrak{g}'_1) = \text{heap}(\mathfrak{g}_1)$ by Lemma 7.14, we get that $\text{project}(\mathfrak{s}, \mathfrak{g}_1) \in \text{type}_{\Phi}(\mathfrak{s}, \text{heap}(\mathfrak{g}_1)) = \mathcal{T}_1$. Hence, $\psi_1 \in \mathcal{T}_1$. \square

A.29 Correctness of the Fixed Point Algorithm For Computing Types of Predicate Calls

A.29.1 Soundness of the Type Computation. We organize the soundness proof into a sequence of simple lemmata about the base cases of the fixed point algorithm and about the operations $\bullet, \cdot[\cdot \cdot \cdot]$, forget and extend. The soundness of the overall algorithm is then be a direct consequence of these lemmata. We first characterize the types of atomic formulas.

LEMMA A.18. *For all aliasing constraints \mathbf{ac} , $\text{Types}_{\Phi}^{\mathbf{ac}}(\mathbf{emp}) = \{\{\mathbf{emp}\}\}$.*

PROOF. Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state with $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) \in \text{Types}_{\Phi}^{\mathbf{ac}}(\mathbf{emp})$ and aliasing(\mathfrak{s}) = \mathbf{ac} . By definition, $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \mathbf{emp}$ and thus $\mathfrak{h} = \emptyset$. We now argue that $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) = \{\mathbf{emp}\}$.

We note that $\emptyset \in \text{forests}_{\Phi}(\mathfrak{h})$ (\emptyset is the forest that does not contain any trees) and $\mathbf{emp} = \text{project}(\mathfrak{s}, \{\emptyset\})$. Hence, $\mathbf{emp} \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$.

Conversely, let $\psi \in \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$. By definition, there is a Φ -forest $\mathfrak{f} = \{t_1, \dots, t_k\}$ with $\text{project}(\mathfrak{s}, \mathfrak{f}) = \psi$ and $\text{heap}(\mathfrak{f}) = \mathfrak{h}$. Since $\mathfrak{h} = \emptyset$, we have $\text{heap}(\mathfrak{f}) = \emptyset$ and thus $\text{heap}(t_i) = \emptyset$ for all $i \in [1, k]$. Hence, $k = 0$. By Definition 7.23 (stack projections), we then have $\psi = \mathbf{emp}$. \square

LEMMA A.19. *Let \mathbf{ac} be a aliasing constraint and $x, y \in \text{dom}(\mathbf{ac})$.*

- *If $\langle x, y \rangle \in \mathbf{ac}$, then $\text{Types}_{\Phi}^{\mathbf{ac}}(x \approx y) = \{\{\mathbf{emp}\}\}$ and $\text{Types}_{\Phi}^{\mathbf{ac}}(x \neq y) = \emptyset$.*

- Otherwise, $\text{Types}_{\Phi}^{\text{ac}}(x \neq y) = \{\{\text{emp}\}\}$ and $\text{Types}_{\Phi}^{\text{ac}}(x \approx y) = \emptyset$.

PROOF. We only consider the case $x \approx y$ as the argument for $x \neq y$ is completely analogous. If $\langle x, y \rangle \in \text{ac}$, our semantics of equalities enforces $\text{Types}_{\Phi}^{\text{ac}}(\text{emp}) = \text{Types}_{\Phi}^{\text{ac}}(x \approx y)$. The claim then follows from Lemma A.18. If $\langle x, y \rangle \notin \text{ac}$, it holds for all \mathfrak{s} with $\text{aliasing}(\mathfrak{s}) = \text{ac}$ that $\mathfrak{s}(x) \neq \mathfrak{s}(y)$. The semantics of $x \approx y$ then yields, for all heaps \mathfrak{h} , that $\langle \mathfrak{s}, \mathfrak{h} \rangle \not\models_{\Phi} x \approx y$. Hence, $\text{Types}_{\Phi}^{\text{ac}}(x \approx y) = \emptyset$. \square

LEMMA A.20. *Let ac be an aliasing constraint, let $a \in \text{dom}(\text{ac})$, and let $\mathbf{b} \in \text{Var}^*$ with $\mathbf{b} \subseteq \text{dom}(\text{ac})$. Then, $\text{Types}_{\Phi}^{\text{ac}}(a \mapsto \mathbf{b}) = \{\text{type}_{\Phi}(\text{ptrmodel}_{\text{ac}}(a \mapsto \mathbf{b}))\}$.*

PROOF. “ \supseteq ” Let $\langle \mathfrak{s}, \mathfrak{h} \rangle \triangleq \text{ptrmodel}_{\text{ac}}(a \mapsto \mathbf{b})$. By definition, $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} a \mapsto \mathbf{b}$. Hence, $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) \in \text{Types}_{\Phi}^{\text{ac}}(a \mapsto \mathbf{b})$.

“ \subseteq ” Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state such that $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) \in \text{Types}_{\Phi}^{\text{ac}}(a \mapsto \mathbf{b})$ and $\text{aliasing}(\mathfrak{s}) = \text{ac}$. By definition, $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} a \mapsto \mathbf{b}$ and thus, by the semantics of points-to assertions, $\mathfrak{h} = \{\mathfrak{s}(a) \mapsto \mathfrak{s}(\mathbf{b})\}$. Consequently, $\langle \mathfrak{s}, \mathfrak{h} \rangle \cong \text{ptrmodel}_{\text{ac}}(a \mapsto \mathbf{b})$ and $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) = \text{type}_{\Phi}(\text{ptrmodel}_{\text{ac}}(a \mapsto \mathbf{b}))$. Since $\langle \mathfrak{s}, \mathfrak{h} \rangle$ was an arbitrary model of ϕ with $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) \in \text{Types}_{\Phi}^{\text{ac}}(a \mapsto \mathbf{b})$ and $\text{aliasing}(\mathfrak{s}) = \text{ac}$, we have $\text{Types}_{\Phi}^{\text{ac}}(a \mapsto \mathbf{b}) \subseteq \{\text{type}_{\Phi}(\text{ptrmodel}_{\text{ac}}(a \mapsto \mathbf{b}))\}$. \square

We next consider the operations \bullet (type composition), $\cdot[\cdot : \cdot/\cdot]$ (variable renaming), forget and extend (lifted to sets of types). The lemmata for these operations below are more general than what is needed for the soundness of our fixed point algorithm because we will also use them for proving the correctness of our algorithm dealing with guarded formulas, see Section 9.2.

LEMMA A.21 (TYPE COMPOSITION). *For $\phi_1, \phi_2 \in \text{GSL}$, $\text{Types}_{\Phi}^{\text{ac}}(\phi_1 \star \phi_2) = \text{Types}_{\Phi}^{\text{ac}}(\phi_1) \bullet \text{Types}_{\Phi}^{\text{ac}}(\phi_2)$.*

PROOF. We show each inclusion separately.

- Let $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\phi_1 \star \phi_2)$. Moreover, fix a state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be such that $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi_1 \star \phi_2$ and $\mathcal{T} = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$. Then, there exist heaps $\mathfrak{h}_1, \mathfrak{h}_2$ such that $\langle \mathfrak{s}, \mathfrak{h}_i \rangle \models_{\Phi} \phi_i$ and $\mathfrak{h} = \mathfrak{h}_1 \uplus \mathfrak{h}_2$. By Corollary 4.6, we have $\langle \mathfrak{s}, \mathfrak{h}_1 \rangle, \langle \mathfrak{s}, \mathfrak{h}_2 \rangle \in \text{GStates}$. By Corollary 8.19, $\mathcal{T} = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_1) \bullet \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_2)$. Since $\langle \mathfrak{s}, \mathfrak{h}_i \rangle \models_{\Phi} \phi_i$, we have $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_i) \in \text{Types}_{\Phi}^{\text{ac}}(\phi_i)$. Hence, $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\phi_1) \bullet \text{Types}_{\Phi}^{\text{ac}}(\phi_2)$.
- Let $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\phi_1) \bullet \text{Types}_{\Phi}^{\text{ac}}(\phi_2)$. Then, there are $\mathcal{T}_1 \in \text{Types}_{\Phi}^{\text{ac}}(\phi_1)$ and $\mathcal{T}_2 \in \text{Types}_{\Phi}^{\text{ac}}(\phi_2)$ such that $\mathcal{T} = \mathcal{T}_1 \bullet \mathcal{T}_2$. Moreover, there are states $\langle \mathfrak{s}, \mathfrak{h}_i \rangle$ such that $\text{aliasing}(\mathfrak{s}) = \text{ac}$, $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}_i) = \mathcal{T}_i$ and $\langle \mathfrak{s}, \mathfrak{h}_i \rangle \models_{\Phi} \phi_i$. By Lemma 4.5 we have $\mathcal{T}_i \neq \emptyset$. By Lemma 8.20 there are states $\langle \mathfrak{s}, \mathfrak{h}'_i \rangle$ such that $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}'_i) = \mathcal{T}_i$ and $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}'_1 \uplus \mathfrak{h}'_2) = \mathcal{T}_1 \bullet \mathcal{T}_2$. By Corollary 8.30 we have $\langle \mathfrak{s}, \mathfrak{h}'_i \rangle \models_{\Phi} \phi_i$. By the semantics of \star , $\langle \mathfrak{s}, \mathfrak{h}'_1 \uplus \mathfrak{h}'_2 \rangle \models_{\Phi} \phi_1 \star \phi_2$. Hence, $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\phi_1 \star \phi_2)$. \square

LEMMA A.22 (RENAMING OF TYPE SETS). *Let \mathbf{x} and \mathbf{y} be sequences of variables as in Definition 9.1 from above. Then, for every GSL formula ϕ , we have*

$$\text{Types}_{\Phi}^{\text{ac}[\mathbf{x}/\mathbf{y}]^{-1}}(\phi)[\text{ac} : \mathbf{x}/\mathbf{y}] = \text{Types}_{\Phi}^{\text{ac}}(\phi[\mathbf{x}/\mathbf{y}]).$$

PROOF. Let $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) \in \text{Types}_{\Phi}^{\text{ac}}(\phi[\mathbf{x}/\mathbf{y}])$. By definition, this means $\text{aliasing}(\mathfrak{s}) = \text{ac}$. Moreover, we note that $\text{ac}[\mathbf{x}/\mathbf{y}]^{-1} = \text{aliasing}(\mathfrak{s}[\mathbf{x}/\mathbf{y}])$. By Lemma 8.22, we have $\text{type}_{\Phi}(\mathfrak{s}[\mathbf{x}/\mathbf{y}], \mathfrak{h})[\text{ac} : \mathbf{x}/\mathbf{y}] = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$, i.e., $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) \in \text{Types}_{\Phi}^{\text{ac}[\mathbf{x}/\mathbf{y}]^{-1}}(\phi)[\text{ac} : \mathbf{x}/\mathbf{y}]$. The converse direction is analogous. \square

LEMMA A.23 (FORGETTING A VARIABLE IN TYPE SETS). *Let $\phi \in \text{GSL}$ be a formula with free variables $\mathbf{x} \cup \{y\}$ such that $y \notin \mathbf{x}$. Moreover, assume that, for every state $\langle \mathfrak{s}, \mathfrak{h} \rangle$, $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$ implies $\mathfrak{s}(y) \in \text{dom}(\mathfrak{h})$. Then, for every aliasing constraint ac with $\text{dom}(\text{ac}) = \mathbf{x}$, we have*

$$\text{Types}_{\Phi}^{\text{ac}}(\exists y. \phi) = \bigcup_{\text{ac}' \in \text{AC}^{\mathbf{x} \cup \{y\}} \text{ with } \text{ac}'|_{\mathbf{x}} = \text{ac}} \text{forget}_{\text{ac}', y}(\text{Types}_{\Phi}^{\text{ac}'}(\phi)).$$

PROOF. Let $\mathcal{T} \in \text{forget}_{\text{ac}',y}(\text{Types}_{\Phi}^{\text{ac}'}(\phi))$, where $\text{ac}' \in \text{AC}^{\text{x} \cup \{y\}}$ is an aliasing constraint satisfying $\text{ac}'|_{\text{x}} = \text{ac}$. Then there exists a state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ such that $\mathcal{T} = \text{forget}_{\text{ac}',y}(\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}))$, $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \phi$, $\text{dom}(\mathfrak{s}) = \text{x}$, and $\text{aliasing}(\mathfrak{s})|_{\text{dom}(\text{ac})} = \text{ac}$. By assumption, $\mathfrak{s}(y) \in \text{dom}(\mathfrak{h})$. Lemma 8.24 then yields

$$\text{type}_{\Phi}(\mathfrak{s}[y/\perp], \mathfrak{h}) = \text{forget}_{\text{ac}',y}(\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})) = \mathcal{T}.$$

By the semantics of existential quantifiers, we have $\langle \mathfrak{s}[y/\perp], \mathfrak{h} \rangle \models_{\Phi} \exists y. \phi$. Hence, $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\exists y. \phi)$.

Conversely, let $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\exists y. \phi)$. Then, there is a state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ such that $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \exists y. \phi$, $\mathcal{T} = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$, $\text{dom}(\mathfrak{s}) = \text{x} \setminus \{y\}$ and $\text{aliasing}(\mathfrak{s}) = \text{ac}$. By the semantics of the existential quantifier, there is a location v such that $\langle \mathfrak{s}[y/v], \mathfrak{h} \rangle \models_{\Phi} \phi$. By assumption we have $\mathfrak{s}(y) \in \text{dom}(\mathfrak{h})$. Then, for $\text{ac}' = \text{aliasing}(\mathfrak{s}[y/v]) \in \text{AC}^{\text{x} \cup \{y\}}$, Lemma 8.24 yields

$$\mathcal{T} = \text{forget}_{\text{ac}',y}(\text{type}_{\Phi}(\mathfrak{s}[y/v], \mathfrak{h})) = \text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) \in \text{forget}_{\text{ac}',y}(\text{Types}_{\Phi}^{\text{ac}'}(\phi)). \quad \square$$

LEMMA A.24 (EXTENDING TYPE SETS). *Let $\phi \in \text{GSL}$ be a formula with free variables x . Moreover, let $\text{ac} \subseteq \text{ac}'$ be alias constraints with $\text{dom}(\text{ac}) = \text{x}$. Then, $\text{Types}_{\Phi}^{\text{ac}'}(\phi) \subseteq \text{extend}_{\text{ac}'}(\text{Types}_{\Phi}^{\text{ac}}(\phi))$.*

PROOF. We consider some $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}'}(\phi)$. Then, there is a state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ with $\text{dom}(\mathfrak{s}) = \text{x} = \text{dom}(\text{ac})$ and $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) = \mathcal{T}$. We now choose some extension \mathfrak{s}' of \mathfrak{s} to $\text{dom}(\text{ac}')$ such that $\mathfrak{s}'(x) \notin \text{locs}(\mathfrak{h})$ for every variable $x \in \text{dom}(\mathfrak{s})$ that is not an alias of a variable in $\text{dom}(\text{ac})$. By Lemma 8.28 we then have $\text{type}_{\Phi}(\mathfrak{s}', \mathfrak{h}) = \text{extend}_{\text{ac}'}(\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}))$. Hence, $\text{extend}_{\text{ac}'}(\mathcal{T}) \in \text{extend}_{\text{ac}'}(\text{Types}_{\Phi}^{\text{ac}}(\phi))$. \square

We are now ready to prove the soundness of the fixed point computation, i.e.,

$$\text{lfp}(\text{unfold}_{\text{x}})(\text{pred}, \text{ac}) \subseteq \text{Types}_{\Phi}^{\text{ac}}(\text{pred}).$$

We first need to establish that $\text{ptypes}_{\Phi}^{\text{x}}(\text{ac}, \phi)$ is sound when ϕ is a rule of predicate pred , i.e., that $\text{ptypes}_{\Phi}^{\text{x}}(\text{ac}, \phi) \subseteq \text{Types}_{\Phi}^{\text{ac}}(\phi)$ holds, under the assumption that p maps every pair of predicate identifier and aliasing constraint to a subset of the corresponding types. As SID rules are guaranteed to be existentially-quantified symbolic heaps, it suffices to prove this result for arbitrary $\phi \in \text{SH}^{\exists}$:

LEMMA A.25. *Let $\phi \in \text{SH}^{\exists}$ and $\text{ac} \in \text{AC}$ with $\text{dom}(\text{ac}) \supseteq \text{fv}(\phi)$. Moreover, let*

$$p: \text{Preds}(\Phi) \times \text{AC} \rightarrow 2^{\text{Types}_{\Phi}}$$

be such that for all $\text{pred} \in \text{Preds}(\Phi)$ and all $\text{ac}' \in \text{AC}^{\text{x} \cup \text{fv}(\text{pred})}$, it holds that $p(\text{pred}, \text{ac}') \subseteq \text{Types}_{\Phi}^{\text{ac}'}(\text{pred})$. Then $\text{ptypes}_{\Phi}^{\text{x}}(\phi, \text{ac}) \subseteq \text{Types}_{\Phi}^{\text{ac}}(\phi)$.

PROOF. We proceed by induction on the structure of ϕ and apply the lemmata from above.

Cases $\phi = x \approx y$, $\phi = x \neq y$. The claim follows from Lemma A.19.

Case $\phi = a \mapsto b$. The claim follows from Lemma A.20.

Case $\phi = \text{pred}(y)$. Let $\text{z} = \text{fv}(\text{pred})$. By I.H., we have

$$p(\text{pred}, \text{ac}[z/y]^{-1}|_{\text{z} \cup \text{x}}) \subseteq \text{Types}_{\Phi}^{\text{ac}[z/y]^{-1}|_{\text{z} \cup \text{x}}}(\text{pred}).$$

By Lemma A.24, we have

$$\text{extend}_{\text{ac}[z/y]^{-1}}(\text{Types}_{\Phi}^{\text{ac}[z/y]^{-1}|_{\text{z} \cup \text{x}}}(\text{pred})) \subseteq \text{Types}_{\Phi}^{\text{ac}[z/y]^{-1}}(\text{pred}).$$

Moreover, by Lemma A.22 we have

$$\text{Types}_{\Phi}^{\text{ac}[z/y]^{-1}}(\text{pred})[\text{ac} : z/y] = \text{Types}_{\Phi}^{\text{ac}}(\text{pred}[z/y]).$$

Hence, we get

$$\text{extend}_{\text{ac}}(p(\text{pred}, \text{ac}[z/y]^{-1}|_{\text{z} \cup \text{x}}))[\text{ac} : z/y] \subseteq \text{Types}_{\Phi}^{\text{ac}}(\text{pred}[z/y]).$$

Case $\phi = \phi_1 \star \phi_2$. The claim follows from Lemma A.21 and the induction hypothesis.

Case $\phi = \exists y. \phi$. The claim follows from Lemma A.23 and the induction hypothesis. \square

LEMMA A.26 (SOUNDNESS OF TYPE COMPUTATION). $\text{lfp}(\text{unfold}_x)(\text{pred}, \text{ac}) \subseteq \text{Types}_{\Phi}^{\text{ac}}(\text{pred})$.

PROOF. A straightforward induction on top of Lemma A.25. \square

A.29.2 Completeness of the Type Computation. We now establish the completeness of the fixed point computation, i.e., $\text{lfp}(\text{unfold}_x)(\text{pred}, \text{ac}) \supseteq \text{Types}_{\Phi}^{\text{ac}}(\text{pred})$. The main challenge is our treatment of predicate calls $\text{ptypes}_p^x(\text{pred}(y), \text{ac})$, for which the recursive look-up $p(\text{pred}, \text{ac}[z/y]^{-1}|_{x \cup z})$ restricts the stack-aliasing constraint $\text{ac}[z/y]^{-1}$ to $x \cup z$, where $z \triangleq \text{fvars}(\text{pred})$. This restriction of the variables is necessary in order to avoid having to consider larger and larger sequences of variables, which would lead to divergence as we illustrate below. Hence, our goal will be to establish that ptypes discovers all types even though we restrict the variables in the recursive look-up.

We now illustrate the need for restricting the variables in the recursive look-up: We assume a stack \mathfrak{s} with $\text{dom}(\mathfrak{s}) = x \cup y$ and pick a rule

$$\text{pred}(\text{fvars}(\text{pred})) \Leftarrow \exists e. (a \mapsto b) \star \text{pred}_1(z_1) \star \cdots \star \text{pred}_k(z_k).$$

We extend \mathfrak{s} to a stack \mathfrak{s}' with $\text{dom}(\mathfrak{s}') = \text{dom}(\mathfrak{s}) \cup e$ and are left with computing the types of

$$((a \mapsto b) \star \text{pred}_1(z_1) \star \cdots \star \text{pred}_k(z_k))[\text{fvars}(\text{pred})/y].$$

At a first glance, this implies recursively computing the types of the calls, $\text{pred}_i(z_i)$, w.r.t. the variables $x' \triangleq x \cup y \cup e$, i.e., we additionally have to consider the existentially quantified variables e . We attempt to do so by picking a rule for each predicate, say we first pick a rule for predicate pred_i . Then, we need to consider an extension \mathfrak{s}'' of \mathfrak{s}' with $\text{dom}(\mathfrak{s}'') = \text{dom}(\mathfrak{s}') \cup e_i$ for the existentially quantified variables e_i on the right hand side of the picked rule. Continuing in this fashion, the computation diverges as we have to extend the set of considered variables x again and again.

However, a more careful analysis reveals that restricting the aliasing constraints to $\text{ac}[z/y]^{-1}|_{x \cup z}$ for the recursive look-up (followed by extending and renaming the obtained set of types) is sufficient. This is a consequence of the *establishment* property, which we require for all SIDs in ID_{btw} . We formalize this insight in the notion of a *tree closure*, which restricts the locations a subtree can share with the variables appearing in the rule instance of its parent node:

DEFINITION A.27 (TREE CLOSURE). Let u be a set of locations and let t be a Φ -tree. Moreover, let t_{sub} be a proper subtree of t and let $\text{pred}(w) = \text{rootpred}(t_{\text{sub}})$. Let $\ell \in \text{dom}(t)$ be the parent location of the root of t_{sub} and let $\text{rule}_t(\ell) = \text{pred}(v) \Leftarrow \phi[\text{fvars}(\text{pred}) \cdot e/v \cdot m]$ be the rule instance at location ℓ . We say t is u -closed for t_{sub} , if $\text{ptrlocs}(t_{\text{sub}}) \cap (v \cup m) \subseteq w \cup u$.

Furthermore, we say t is u -closed, if t is u -closed for all proper subtrees of t .

EXAMPLE A.28. The tree from Fig. 6b is $\langle \rangle$ -closed, where $\langle \rangle$ is the empty sequence. If we replace location 8 everywhere in the tree with location 7, then the resulting tree is not $\langle \rangle$ -closed anymore (consider the subtree rooted at location 2), but 7-closed.

LEMMA A.29. Let t be some Φ -tree with $\text{allholepreds}(t) = \emptyset$ and $\text{pred}(u) = \text{rootpred}(t)$. Let $\ell \in \text{dom}(t)$ be some location and let $\text{rule}_t(\ell) = \text{pred}(v) \Leftarrow \phi[\text{fvars}(\text{pred}) \cdot e/v \cdot m]$ be the rule instance at location ℓ of t . Then, we have $v \cup m \subseteq \text{dom}(t) \cup u$.

PROOF. A direct consequence of establishment. \square

Recall that, by Lemma 7.3, we have $\langle s, h \rangle \models_{\Phi} \text{pred}(\text{fvars}(\text{pred}))$ if and only if there exists a Φ -tree t with $\text{rootpred}(t) = \text{pred}(s(\text{fvars}(\text{pred})))$, $\text{allholes}(t) = \emptyset$ and $\text{heap}(\{t\}) = h$. The completeness of our fixed point algorithm relies on the observation that such trees t are $s(\text{fvars}(\text{pred}))$ -closed:

LEMMA A.30. *Every Φ -tree t with $\text{allholepreds}(t) = \emptyset$ and $\text{pred}(u) = \text{rootpred}(t)$ is u -closed.*

PROOF. Let t_{sub} be a proper subtree of t and let $\text{pred}(w) = \text{rootpred}(t_{\text{sub}})$. Let $\ell \in \text{dom}(t)$ be the parent location of the root of t_{sub} and let $\text{rule}_t(\ell) = \text{pred}(v) \Leftarrow \phi[\text{fvars}(\text{pred}) \cdot e/v \cdot m]$ be the rule instance at location ℓ . We consider some $k \in \text{ptrlocs}(t_{\text{sub}}) \cap (v \cup m)$. Then k is either dangling or an allocated location in $\text{heap}(t_{\text{sub}})$.

Assume k is dangling in $\text{heap}(t_{\text{sub}})$: We define the stack $s : \text{fvars}(\text{pred}) \rightarrow \text{Loc}$ by setting $s(\text{fvars}(\text{pred}_{\text{sub}})) = w$. By Lemma 7.3, we have $\langle s, \text{heap}(t_{\text{sub}}) \rangle \models_{\Phi} \text{pred}(\text{fvars}(\text{pred}))$. By Lemma 4.4 we have $\langle s, \text{heap}(t_{\text{sub}}) \rangle \in \text{GStates}$. Hence, $k \in \text{img}(s) = w$. In other words, dangling locations do not invalidate that t is u -closed.

Assume k is allocated in $\text{heap}(t_{\text{sub}})$, i.e., $k \in \text{dom}(t_{\text{sub}})$: To prove that t is u -closed, it is sufficient that $k \notin w$ implies that $k \in u$. Hence, let us assume that $k \notin w$. Let $t_{\text{rem}} = t \setminus t_{\text{sub}}$ be the remainder of t after splitting off t_{sub} , i.e., $\text{split}(\{t\}, \{\text{root}(t_{\text{sub}})\}) = \{t_{\text{sub}}, t_{\text{rem}}\}$. We note that $k \notin \text{dom}(t_{\text{rem}})$ because a location cannot be allocated in two subtrees. We choose a fresh location $k' \in \text{Loc} \setminus \text{ptrlocs}(t)$ and then create a tree t'_{sub} as a copy of t_{sub} except that we replace every occurrence of location k in t_{sub} with k' . We note that $\text{rootpred}(t'_{\text{sub}}) = \text{pred}_{\text{sub}}(w) = \text{rootpred}(t_{\text{sub}})$ because of $k \notin w$. Hence, there is a tree t' such that $\text{split}(\{t'\}, \{\ell\}) = \{t'_{\text{sub}}, t_{\text{rem}}\}$. We note that, by construction of t' , we have that (1) $k' \notin \text{dom}(t')$, (2) $\text{rootpred}(t') = \text{rootpred}(t) = \text{pred}(u)$, (3) $\text{allholepreds}(t') = \emptyset$, (4) ℓ is the parent of t'_{sub} in t' , and (5) $\text{rule}_{t'}(\ell) = \text{pred}(v) \Leftarrow \phi[\text{fvars}(\text{pred}) \cdot e/v \cdot m]$. Lemma A.29 then yields $v \cup m \subseteq \text{dom}(t') \cup u$. Since k is allocated, this means $k \in \text{dom}(t') \cup u$. With (1) we then obtain $k \in u$. \square

We are now ready to prove the completeness of our fixed-point algorithm for computing types. We will show that the fixed-point algorithm discovers, for all predicates pred and aliasing constraints $\text{ac} \in \text{AC}^{\text{x} \cup \text{fvars}(\text{pred})}$, all types in the following set:

$$\begin{aligned} \{ \text{type}_{\Phi}(s, \text{heap}(t)) \mid & \text{aliasing}(s) = \text{ac}, \\ & \text{rootpred}(t) = \text{pred}(s(\text{fvars}(\text{pred}))), \\ & \text{allholes}(t) = \emptyset \}, \end{aligned}$$

where—as shown above—we can rely on the assumption that the considered trees t are $s(x)$ -closed.

LEMMA A.31. *Let s be a stack with $\text{dom}(s) = \text{x} \cup \text{fvars}(\text{pred})$ and $s(\text{fvars}(\text{pred})) = v$. Let t be an $s(x)$ -closed Φ -tree with $\text{rootpred}(t) = \text{pred}(v)$ and $\text{allholepreds}(t) = \emptyset$. Then,*

$$\text{type}_{\Phi}(s, \text{heap}(t)) \in \text{unfold}_{\text{x}}^{\text{height}(t)+1}(\lambda(\text{pred}', \text{ac}') . \emptyset)(\text{pred}, \text{aliasing}(s)).$$

PROOF. We prove the claim by strong mathematical induction on $\text{height}(t)$:

Let $h \triangleq \text{heap}(t)$, $r \triangleq \text{root}(t)$, and $\text{ac} \triangleq \text{aliasing}(s)$. Since t is an Φ -tree with $\text{rootpred}(t) = \text{pred}(v)$, there is a rule $(\text{pred}(x) \Leftarrow \phi) \in \Phi$ with $\phi = \exists e. \phi'$, $\phi' = (y \mapsto z) \star \text{pred}_1(z_1) \star \dots \star \text{pred}_k(z_k) \star \Pi$, Π pure¹⁶, such that $\text{rule}_t(r) = \text{pred}(v) \Leftarrow \phi'[\text{fvars}(\text{pred}) \cdot e/v \cdot m]$ for some $m \in \text{Loc}^*$ (i.e., the root r of t is labeled with an instance of the rule).

Let $s' = s[\text{fvars}(\text{pred})/v][e/m]$. Moreover, for $1 \leq i \leq k$, let t_i be the subtree of t such that $\text{rootpred}(t_i) = \text{pred}_i(z_i)[\text{fvars}(\text{pred}) \cdot e/v \cdot m]$; let $h_i = \text{heap}(t_i)$. By Lemma 7.3, we have $\langle s', h_i \rangle \models_{\Phi} \text{pred}_i(z_i)$ and, by Lemma 4.4, we have $\langle s', h_i \rangle \in \text{GStates}$. Finally, we denote by h_0 the unique heap such that $\langle s', h_0 \rangle \models_{\Phi} y \mapsto z$. Clearly, $\langle s', h_0 \rangle \in \text{GStates}$ and $h = h_0 \uplus \dots \uplus h_k$.

¹⁶In case of $\text{height}(t) = 0$, the rule $(\text{pred}(x) \Leftarrow \phi)$ is non-recursive, and we have $k = 0$ and there are no existentially quantified variables, i.e., $e = \epsilon$.

We use the following abbreviations:

$$\begin{aligned}\mathbf{ac}' &\triangleq \text{aliasing}(\mathfrak{s}') \\ \mathcal{T}_0 &\triangleq \text{type}_\Phi(\text{ptrmodel}_{\mathbf{ac}'}(y \mapsto z)) \\ \mathcal{T}_i &\triangleq \text{type}_\Phi(\mathfrak{s}', \mathfrak{h}_i), i \geq 1\end{aligned}$$

Since t is $\mathfrak{s}(\mathbf{x})$ -closed we have that $\text{ptrlocs}(t_i) \cap (\mathbf{v} \cup \mathbf{m}) \subseteq \mathfrak{s}'(z_i) \cup \mathfrak{s}(\mathbf{x}) = \mathfrak{s}'(z_i) \cup \mathfrak{s}'(\mathbf{x})$. Furthermore, due to $\text{img}(\mathfrak{s}') = \mathbf{v} \cup \mathbf{m}$, we have $\mathfrak{s}'(x) \notin \text{ptrlocs}(t_i) \supseteq \text{locs}(\mathfrak{h}_i)$ for all $x \in \text{dom}(\mathbf{ac}')$ for which there is no $y \in z_i \cup \mathbf{x}$ with $\mathfrak{s}'(x) = \mathfrak{s}'(y)$. Let \mathfrak{s}_i be the restriction of \mathfrak{s}' to $\mathbf{x} \cup z_i$. Lemma 8.28 then yields

$$\text{type}_\Phi(\mathfrak{s}', \mathfrak{h}_i) = \text{extend}_{\mathbf{ac}'}(\text{type}_\Phi(\mathfrak{s}_i, \mathfrak{h}_i)). \quad (\dagger)$$

We introduce some more abbreviations:

$$\begin{aligned}\mathfrak{s}'_i &\triangleq \mathfrak{s} \cup \{\text{fvars}(\text{pred}_i) \mapsto \mathfrak{s}'(z_i)\}, \\ \mathcal{T}'_i &\triangleq \text{type}_\Phi(\mathfrak{s}'_i, \mathfrak{h}_i) \\ \mathbf{ac}_i &\triangleq \text{aliasing}(\mathfrak{s}'_i)\end{aligned}$$

Observe that $\mathfrak{s}_i = \mathfrak{s}'_i[\text{fvars}(\text{pred}_i)/z_i]$. By Lemma 8.22, $\mathcal{T}'_i[\text{fvars}(\text{pred}_i)/z_i] = \text{type}_\Phi(\mathfrak{s}_i, \mathfrak{h}_i)$. We then apply (\dagger) to obtain

$$\mathcal{T}_i = \text{type}_\Phi(\mathfrak{s}', \mathfrak{h}_i) = \text{extend}_{\mathbf{ac}'}(\text{type}_\Phi(\mathfrak{s}_i, \mathfrak{h}_i)) = \text{extend}_{\mathbf{ac}'}(\mathcal{T}'_i[\text{fvars}(\text{pred}_i)/z_i]). \quad (\ddagger)$$

Finally, we note that t_i is $\mathfrak{s}'_i(\mathbf{x})$ -closed, $\text{rootpred}(t_i) = \text{pred}(\mathfrak{s}'_i(\text{fvars}(\text{pred}_i)))$ and $\text{allholepreds}(t_i) = \emptyset$. Since $\text{height}(t_i) < \text{height}(t)$, we can then apply the induction hypothesis and conclude that

$$\begin{aligned}\mathcal{T}'_i &= \text{type}_\Phi(\mathfrak{s}'_i, \mathfrak{h}_i) \\ &\in \text{unfold}_x^{\text{height}(t_i)+1}(\lambda(\text{pred}', \mathbf{ac}'). \emptyset)(\text{pred}_i, \text{aliasing}(\mathfrak{s}'_i)) \\ &\sqsubseteq \text{unfold}_x^{\text{height}(t)}(\lambda(\text{pred}', \mathbf{ac}'). \emptyset)(\text{pred}_i, \mathbf{ac}_i).\end{aligned} \quad (\clubsuit)$$

To finish the proof, we set $p \triangleq \text{unfold}_x^{\text{height}(t)}(\lambda(\text{pred}', \mathbf{ac}'). \emptyset)$ and proceed as follows:

$$\begin{aligned}&\text{unfold}_x^{\text{height}(t)+1}(\lambda(\text{pred}', \mathbf{ac}'). \emptyset)(\text{pred}, \mathbf{ac}) \\ &= \bigcup_{(\text{pred}(\text{fvars}(\text{pred})) \Leftarrow \phi) \in \Phi} \text{ptypes}_p^x(\mathbf{ac}, \phi) && \text{(by definition)} \\ &\supseteq \text{ptypes}_p^x(\mathbf{ac}, \exists e. \phi') && (\phi = \exists e. \phi') \\ &\supseteq \text{forget}_{\mathbf{ac}', e}(\text{ptypes}_p^x(\phi', \mathbf{ac}')) && \text{(Lemma A.23)} \\ &= \text{forget}_{\mathbf{ac}', e}(\text{ptypes}_p^x(y \mapsto z, \mathbf{ac}') \bullet \text{ptypes}_p^x(\text{pred}_1(z_1), \mathbf{ac}')) && \text{(Def. of } \phi', \text{ Lemma A.21)} \\ &\quad \bullet \dots \bullet \text{ptypes}_p^x(\text{pred}_k(z_k), \mathbf{ac}') \\ &= \text{forget}_{\mathbf{ac}', e}(\{\mathcal{T}_0\} \bullet \text{extend}_{\mathbf{ac}'}(p(\text{pred}_1, \mathbf{ac}_1)[\text{fvars}(\text{pred}_1) : z_1/]) \\ &\quad \bullet \dots \bullet \text{extend}_{\mathbf{ac}'}(p(\text{pred}_k, \mathbf{ac}_k)[\text{fvars}(\text{pred}_k) : z_k/])) && \text{(Def. of } \mathcal{T}_0, \text{ Lemma A.24)} \\ &\supseteq \text{forget}_{\mathbf{ac}', e}(\{\mathcal{T}_0\} \bullet \text{extend}_{\mathbf{ac}'}(\{\mathcal{T}'_1\}[\text{fvars}(\text{pred}_1) : z_1/]) \\ &\quad \bullet \dots \bullet \text{extend}_{\mathbf{ac}'}(\{\mathcal{T}'_k\}[\text{fvars}(\text{pred}_k) : z_k/])) && \text{(by } (\clubsuit)) \\ &= \text{forget}_{\mathbf{ac}', e}(\{\mathcal{T}_0 \bullet \mathcal{T}_1 \dots \bullet \mathcal{T}_k\}) && \text{(by } (\ddagger)) \\ &= \{\text{forget}_{\mathbf{ac}', e}(\text{type}_\Phi(\mathfrak{s}', \mathfrak{h}))\} && \text{(by Corollary 8.19)}\end{aligned}$$

$$= \{\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})\}. \quad (\text{by Lemma 8.24, as } \mathfrak{s}(\mathbf{e}) \subseteq \text{dom}(\mathfrak{h}))$$

Read from bottom to top, we have $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) \in \text{unfold}_{\mathbf{x}}^{\text{height}(\mathfrak{t})+1}(\lambda(\text{pred}', \text{ac}'). \emptyset)(\text{pred}, \text{ac})$. \square

Completeness then follows by exploiting the one-to-one correspondence between $\mathfrak{s}(\mathbf{x})$ -closed Φ -trees without holes and the models of a predicate:

LEMMA A.32 (COMPLETENESS OF TYPE COMPUTATION). *Let $\text{pred} \in \mathbf{Preds}(\Phi)$ such that $\text{fvars}(\text{pred}) = \mathbf{z} = \langle z_1, \dots, z_k \rangle \subseteq \mathbf{x}$. Moreover, let $\langle \mathfrak{s}, \mathfrak{h} \rangle \models_{\Phi} \text{pred}(\mathbf{z})$ for some state $\langle \mathfrak{s}, \mathfrak{h} \rangle$ with $\mathbf{x} = \text{dom}(\mathfrak{s})$. Then,*

$$\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) \in \text{lfp}(\text{unfold}_{\mathbf{x}})(\text{pred}, \text{aliasing}(\mathfrak{s})).$$

PROOF. By Lemma 7.3, there exists a Φ -tree \mathfrak{t} such that $\text{rootpred}(\mathfrak{t}) = \text{pred}(\mathfrak{s}(\mathbf{z}))$, $\text{allholes}(\mathfrak{t}) = \emptyset$, and $\text{heap}(\{\mathfrak{t}\}) = \mathfrak{h}$. By Lemma A.30, the tree \mathfrak{t} is $\mathfrak{s}(\mathbf{z})$ -closed. Since $\mathbf{z} \subseteq \mathbf{x}$, \mathfrak{t} is also $\mathfrak{s}(\mathbf{x})$ -closed. By Lemma A.31, we know that

$$\text{type}_{\Phi}(\mathfrak{s}, \text{heap}(\mathfrak{t})) \in \text{unfold}_{\mathbf{x}}^{\text{height}(\mathfrak{t})+1}(\lambda(\text{pred}', \text{ac}'). \emptyset)(\text{pred}, \text{aliasing}(\mathfrak{s})).$$

Recalling that $\text{lfp}(\text{unfold}_{\mathbf{x}}) = \lim_{n \in \mathbb{N}} \text{unfold}_{\mathbf{x}}^n(\lambda(\text{pred}', \text{ac}'). \emptyset)$ and $\mathfrak{h} = \text{heap}(\mathfrak{t})$, we conclude that

$$\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) \in \text{lfp}(\text{unfold}_{\mathbf{x}})(\text{pred}, \text{aliasing}(\mathfrak{s})). \quad \square$$

A.29.3 *Complexity of the Fixed-Point Computation.* We now establish that the types of predicates can be computed in *doubly-exponential time*. As a first step, we consider a special case: the complexity of computing the types of single points-to assertions $a \mapsto \mathbf{b}$. Intuitively, single points-to assertions correspond to Φ -trees of size one. To compute their types, we systematically enumerate all such trees and check for each tree whether the points-to assertion in the tree node coincides with $\mathfrak{s}(a) \mapsto \mathfrak{s}(\mathbf{b})$.

LEMMA A.33. *Let ac be an aliasing constraint, let $a \in \text{dom}(\text{ac})$, and let $\mathbf{b} \in \mathbf{Var}^*$ with $\mathbf{b} \subseteq \text{dom}(\text{ac})$. Let $n \triangleq \max\{|\Phi|, |\text{dom}(\text{ac})|\}$. Then, $\text{type}_{\Phi}(\text{ptrmodel}_{\text{ac}}(a \mapsto \mathbf{b}))$ is computable in $2^{O(n \log(n))}$.*

PROOF. Let $\langle \mathfrak{s}, \mathfrak{h} \rangle = \text{ptrmodel}_{\text{ac}}(a \mapsto \mathbf{b})$. W.l.o.g. we can assume that $\text{img}(\mathfrak{s}) \subseteq \{1, \dots, n\}$ (otherwise we can select an isomorphic model with this property). We observe that a single location is allocated in \mathfrak{h} . Hence, in order to compute the type of $\langle \mathfrak{s}, \mathfrak{h} \rangle$ we need to consider exactly those forests that consist of a single tree with a single rule instance whose points-to assertion agrees with $a \mapsto \mathbf{b}$. We collect those rule instances in the set \mathbf{R} :

$$\begin{aligned} \mathbf{R} \triangleq \{ & \text{pred}(\mathbf{l}) \Leftarrow ((v \mapsto \mathbf{w}) \star \text{pred}_1(z_1) \star \dots \star \text{pred}_k(z_k) \star \Pi)[\text{fvars}(\text{pred}) \cdot \mathbf{e}/\mathbf{l} \cdot \mathbf{m}] \in \mathbf{RuleInst}(\Phi) \mid \\ & \mathbf{l} \cdot \mathbf{m} \in \mathcal{L}^*, v[\text{fvars}(\text{pred}) \cdot \mathbf{e}/\mathbf{l} \cdot \mathbf{m}] = \mathfrak{s}(a), \mathbf{w}[\text{fvars}(\text{pred}) \cdot \mathbf{e}/\mathbf{l} \cdot \mathbf{m}] = \mathfrak{s}(\mathbf{b}) \} \end{aligned}$$

We note that $|\mathbf{l} \cdot \mathbf{m}| \leq n$ for all rule instances. Then, $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$ is given by the projections of the forests that consist of a single tree with a rule instance from \mathbf{R} :

$$\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h}) = \{\text{project}(\mathfrak{s}, \{\{a \mapsto \langle \emptyset, \mathcal{R} \rangle\}) \mid \mathcal{R} \in \mathbf{R}\} \cap \mathbf{DUSH}_{\Phi}\}.$$

For computing $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$, we only need those rules instances in \mathbf{R} such that $\mathbf{l} \cdot \mathbf{m} \subseteq \{1, \dots, n\}$: We have $\text{img}(\mathfrak{s}) \subseteq \{1, \dots, n\}$ and we can rename locations not in $\{1, \dots, n\}$ to obtain a \mathfrak{s} -equivalent forest with the desired property; such forests have the same projections due to Lemma 7.38. Thus, we can compute $\text{type}_{\Phi}(\mathfrak{s}, \mathfrak{h})$ by considering $n \cdot n^n \in 2^{O(n^2 \log(n))}$ rule instances. \square

THEOREM A.34 (COMPLEXITY OF TYPE COMPUTATION). *Let $n \triangleq |\Phi| + |\mathbf{x}|$. Then, one can compute the set $\text{lfp}(\text{unfold}_{\mathbf{x}})$ assigning sets of types to predicates in $2^{2^{O(n^2 \log(n))}}$.*

PROOF. Theorem 8.7 gives us a bound on the the size of all types over aliasing constraints in AC^x : $|\text{Types}_\Phi^x| \in 2^{2^{O(n^2 \log(n))}}$. Moreover, the number of predicates of Φ is bounded by n . Consequently, the number of functions with signature $\text{Preds}(\Phi) \times \text{AC} \rightarrow 2^{\text{Types}_\Phi}$ is bounded by

$$n \cdot 2^{2^{O(n^2 \log(n))}} = 2^{2^{O(n^2 \log(n))+1}} = 2^{2^{O(n^2 \log(n))}}.$$

Since every iteration of the fixed-point computation discovers at least one new type, the computation terminates after at most $2^{2^{O(n^2 \log(n))}}$ many iterations. We will show that each iteration takes at most $2^{2^{O(n^2 \log(n))}}$ steps. This is sufficient to establish the claim because

$$\underbrace{2^{2^{O(n^2 \log(n))}}}_{\text{number of iterations}} \cdot \underbrace{2^{2^{O(n^2 \log(n))}}}_{\text{cost per iteration}} = 2^{2^{O(n^2 \log(n))}}.$$

We now study the time spent in each iteration: Given some predicate $\text{pred} \in \text{Preds}(\Phi)$ and aliasing constraint $\text{ac} \in \text{AC}^x$, we need to compute

- the function $\text{ptypes}_p^x(\phi, \text{ac})$ for each rule $\text{pred}(\text{fvars}(\text{pred})) \Leftarrow \phi \in \Phi$, where p is the pre-fixed point from the previous iteration, and
- the union of the results of these function calls (note that we need to compute at most one union operation per rule $\phi \in \Phi$).

We argue below that each call $\text{ptypes}_p^x(\phi, \text{ac})$ can be done in at most $2^{2^{O(n^2 \log(n))}}$ many steps. We further observe that the union over a set of types is linear in the number of types, i.e., linear in $2^{2^{O(n^2 \log(n))}}$. Hence, each iteration takes at most $2^{2^{O(n^2 \log(n))}}$ many steps because

$$\underbrace{n}_{\text{number of rules}} \cdot \underbrace{O(2^{n \log(n)})}_{\text{number of aliasing constraints}} \cdot \underbrace{2^{2^{O(n^2 \log(n))}}}_{\text{cost for a fixed rule and aliasing constraint}} = 2^{2^{O(n^2 \log(n))}}.$$

To conclude the proof, we consider the cost of evaluating $\text{ptypes}_p^x(\phi, \text{ac})$ for a fixed rule body ϕ and aliasing constraint ac . Since the type for each right-hand side of a rule is computed at most once, we note that the recursive calls of ptypes lead to at most $|\phi| \leq n$ evaluations of base cases, i.e., (dis-)equalities and points-to assertions, and operations \bullet , $\cdot[\cdot : \cdot/\cdot]$, forget and extend. It remains to establish the cost of these operations:

- (1) Evaluating a (dis-)equality takes constant time.
- (2) The evaluation of a points-to assertions can be done in time $O(2^{n \log(n)})$ by Lemma A.33 (observing that $|\text{dom}(\text{ac})| \leq |\Phi| + |\mathbf{x}| = n$).
- (3) The evaluation of the operations \bullet , $\cdot[\cdot : \cdot/\cdot]$, forget and extend each takes time polynomial in the size of the types, i.e., $2^{O(n^2 \log(n))}$ (see Lemma 8.6). For $\cdot[\cdot : \cdot/\cdot]$, and forget this is trivial. For the composition operation, \bullet , the polynomial bound follows because (1) the number of formulas that can be obtained by re-scoping is bounded by the number of types, and (2) the number of formulas that can be obtained by \triangleright steps is also bounded by the number of types. Similarly, the number of formulas that can be obtained by extend is bounded by the number of types. As the number of types to which each function is applied is bounded by $2^{2^{O(n^2 \log(n))}}$ we obtain the following cost for each \bullet , $\cdot[\cdot : \cdot/\cdot]$, forget and extend:

$$\underbrace{\text{poly}(2^{O(n^2 \log(n))})}_{\text{cost of operation for a single type}} \cdot \underbrace{2^{2^{O(n^2 \log(n))}}}_{\text{number of types}} = 2^{2^{O(n^2 \log(n))}}.$$

Hence, the cost of evaluating $\text{ptypes}_p^x(\phi, \text{ac})$ for a fixed rule body ϕ and aliasing constraint ac is

$$\underbrace{n}_{\text{size of the rule } \phi} \cdot \underbrace{2^{2^{O(n^2 \log(n))}}}_{\text{cost of each of the } n \text{ operations}} = 2^{2^{O(n^2 \log(n))}}. \quad \square$$

A.30 Correctness of the Algorithm For Computing the Types of Guarded Formulas

The correctness of types is almost immediate from our previous results established for computing types of predicates; we only need two additional lemmata which we state below:

LEMMA A.35. *Let $\phi_1, \phi_2 \in \text{GSL}$ be two formulas and let ac be a stack-aliasing constraint. Then, $\text{Types}_{\Phi}^{\text{ac}}(\phi_1 \wedge \phi_2) = \text{Types}_{\Phi}^{\text{ac}}(\phi_1) \cap \text{Types}_{\Phi}^{\text{ac}}(\phi_2)$, $\text{Types}_{\Phi}^{\text{ac}}(\phi_1 \vee \phi_2) = \text{Types}_{\Phi}^{\text{ac}}(\phi_1) \cup \text{Types}_{\Phi}^{\text{ac}}(\phi_2)$ and $\text{Types}_{\Phi}^{\text{ac}}(\phi_1 \wedge \neg \phi_2) = \text{Types}_{\Phi}^{\text{ac}}(\phi_1) \setminus \text{Types}_{\Phi}^{\text{ac}}(\phi_2)$*

PROOF. We only show the first claim, the other two claims are shown analogously.

By definition of types, the inclusion $\text{Types}_{\Phi}^{\text{ac}}(\phi_1 \wedge \phi_2) \subseteq \text{Types}_{\Phi}^{\text{ac}}(\phi_1) \cap \text{Types}_{\Phi}^{\text{ac}}(\phi_2)$ is straightforward. For the converse direction, we consider some $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\phi_1) \cap \text{Types}_{\Phi}^{\text{ac}}(\phi_2)$. Because of $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\phi_1)$ there is a state $\langle s, h \rangle$ with $\mathcal{T} = \text{type}_{\Phi}(\langle s, h \rangle)$ and $\langle s, h \rangle \models_{\Phi} \phi_1$. By Corollary 4.6, we have $\langle s, h \rangle \in \text{GStates}$. Thus, Corollary 8.30 yields $\langle s, h \rangle \models_{\Phi} \phi_2$. Hence, $\langle s, h \rangle \models_{\Phi} \phi_1 \wedge \phi_2$ and we obtain that $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\phi_1 \wedge \phi_2)$. \square

LEMMA A.36. *Let $\phi_0, \phi_1, \phi_2 \in \text{GSL}$ be three formulas and let ac be a stack-aliasing constraint. Then,*

$$\text{Types}_{\Phi}^{\text{ac}}(\phi_0 \wedge (\phi_1 \oplus \phi_2)) = \{\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\phi_0) \mid \exists \mathcal{T}' \in \text{Types}_{\Phi}^{\text{ac}}(\phi_1). \mathcal{T} \bullet \mathcal{T}' \in \text{Types}_{\Phi}^{\text{ac}}(\phi_2)\},$$

$$\text{Types}_{\Phi}^{\text{ac}}(\phi_0 \wedge (\phi_1 \star \phi_2)) = \{\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\phi_0) \mid \forall \mathcal{T}' \in \text{Types}_{\Phi}^{\text{ac}}(\phi_1). \mathcal{T} \bullet \mathcal{T}' \in \text{Types}_{\Phi}^{\text{ac}}(\phi_2)\}.$$

PROOF. We only show the first claim, the second claim is shown analogously.

Let $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\phi_0 \wedge (\phi_1 \oplus \phi_2))$. Then, there is a state $\langle s, h \rangle$ with $\mathcal{T} = \text{type}_{\Phi}(\langle s, h \rangle)$, $\langle s, h \rangle \models_{\Phi} \phi_0$, and $\langle s, h \rangle \models_{\Phi} \phi_1 \oplus \phi_2$. By the semantics of \oplus , there exists a heap h_1 with $\langle s, h_1 \rangle \models_{\Phi} \phi_1$ and $\langle s, h \uplus h_1 \rangle \models_{\Phi} \phi_2$. Let $\mathcal{T}_1 \triangleq \text{type}_{\Phi}(\langle s, h_1 \rangle)$ and $\mathcal{T}_2 \triangleq \text{type}_{\Phi}(\langle s, h \uplus h_1 \rangle)$. By Corollary 8.19, $\mathcal{T}_2 = \mathcal{T} \bullet \mathcal{T}_1$. Hence, $\mathcal{T} \in \{\text{Types}_{\Phi}^{\text{ac}}(\phi_0) \mid \exists \mathcal{T}' \in \text{Types}_{\Phi}^{\text{ac}}(\phi_1). \mathcal{T} \bullet \mathcal{T}' \in \text{Types}_{\Phi}^{\text{ac}}(\phi_2)\}$.

Conversely, let $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\phi_0)$ such that there is an $\mathcal{T}' \in \text{Types}_{\Phi}^{\text{ac}}(\phi_1)$ with $\mathcal{T} \bullet \mathcal{T}' \in \text{Types}_{\Phi}^{\text{ac}}(\phi_2)$. Then, there is a state $\langle s, h \rangle$ with $\mathcal{T} = \text{type}_{\Phi}(\langle s, h \rangle)$ and $\langle s, h \rangle \models_{\Phi} \phi_0$. Further, there is a state $\langle s, h_1 \rangle$ with $\text{type}_{\Phi}(\langle s, h_1 \rangle) = \mathcal{T}'$ and $\langle s, h_1 \rangle \models_{\Phi} \phi_1$. We can assume w.l.o.g. that $h \uplus h_1 \neq \perp$ —otherwise, replace h_1 with an isomorphic heap that has this property. Corollary 8.19 yields $\text{type}_{\Phi}(\langle s, h \uplus h_1 \rangle) = \mathcal{T} \bullet \mathcal{T}' \in \text{Types}_{\Phi}^{\text{ac}}(\phi_2)$. Since $\phi_0, \phi_1 \in \text{GSL}$, we have $\langle s, h \rangle \in \text{GStates}$ and $\langle s, h_1 \rangle \in \text{GStates}$ by Corollary 4.6. Thus, also $\langle s, h \uplus h_1 \rangle \in \text{GStates}$. Corollary 8.30 then gives us that $\langle s, h \uplus h_1 \rangle \models_{\Phi} \phi_2$. Therefore, $\langle s, h \rangle \models_{\Phi} \phi_1 \oplus \phi_2$, which implies that $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\phi_0 \wedge (\phi_1 \oplus \phi_2))$. Hence, $\mathcal{T} \in \text{Types}_{\Phi}^{\text{ac}}(\phi_0 \wedge (\phi_1 \oplus \phi_2))$. \square

We restate the claim of Theorem 9.2: Let $\phi \in \text{GSL}$ with $\text{fvars}(\phi) = \mathbf{x}$ and $\text{locs}(\phi) = \emptyset$. Further, let $\text{ac} \in \text{AC}^x$. Then, $\text{Types}_{\Phi}^{\text{ac}}(\phi) = \text{types}(\phi, \text{ac})$. Moreover, $\text{types}(\phi, \text{ac})$ can be computed in $2^{2^{O(n^2 \log(n))}}$, where $n \triangleq |\Phi| + |\phi|$.

PROOF. We first prove that $\text{Types}_{\Phi}^{\text{ac}}(\phi) = \text{types}(\phi, \text{ac})$. The proof proceeds by induction on ϕ :

Case $\phi = \text{emp}$. By Lemma A.18.

Case $\phi = x \approx y, \phi = x \neq y$. By Lemma A.19.

Case $\phi = a \mapsto \mathbf{b}$. By Lemma A.20.

Case $\phi = \text{pred}(y)$. By Lemma A.22, Lemma A.26 and Lemma A.32.

Case $\phi = \phi_1 \star \phi_2$. By Lemma A.21 and the I.H..

Case $\phi = \phi_1 \wedge \phi_2, \phi = \phi_1 \vee \phi_2, \phi = \phi_1 \wedge \neg \phi_2$. By Lemma A.35 and the I.H.

Case $\phi = \phi_0 \wedge (\phi_1 \oplus \phi_2), \phi = \phi_0 \wedge (\phi_1 \star \phi_2)$. By Lemma A.36 and the I.H.

We now turn to the complexity claim:

We recall that the number of types in $\mathbf{Types}_{\Phi}^{\text{ac}}(\phi)$ is bounded by $2^{2^{O(n^2 \log(n))}}$ (see Theorem 8.7). The evaluation of $\text{types}(\phi, \text{ac})$ consists of at most $|\phi| \leq n$ invocations of the form $\text{types}(\cdot, \text{ac})$. We will show that each of these invocations can be evaluated in time at most $2^{2^{O(n^2 \log(n))}}$; this is sufficient to establish the claim because of $n \cdot 2^{2^{O(n^2 \log(n))}} = 2^{2^{O(n^2 \log(n))}}$:

- For **emp** and (dis-)equalities, the evaluation time is constant.
- For points-to assertions, this follows from Lemma A.33.
- For predicate calls, this follows from Theorem A.34.
- For \wedge , \vee , and \neg , the bound follows because each of these operations can be implemented in linear time in terms of the number of types.
- For \star , this follows because (1) \bullet is applied to at most $2^{2^{O(n^2 \log(n))}} \cdot 2^{2^{O(n^2 \log(n))}} = 2^{2^{O(n^2 \log(n))}}$ many types and (2) the composition $\mathcal{T}_1 \bullet \mathcal{T}_2$ takes time at most $\text{poly}(2^{2^{O(n^2 \log(n))}})$, as argued in the proof of Theorem A.34. Hence, the cost of \bullet is $\text{poly}(2^{2^{O(n^2 \log(n))}}) \cdot 2^{2^{O(n^2 \log(n))}} = 2^{2^{O(n^2 \log(n))}}$.
- For septraction and the magic wand, this is analogously to the cases for \wedge resp. \vee and \star . □

A.31 Correctness of the Reduction over Values with the Null-Pointer to Values without the Null-Pointer (Cor. 9.7)

We only show the claim about satisfiability. The claim about entailment follows from the first claim as in the proof of Cor. 9.4.

The proof of the first claim proceeds by a reduction to the satisfiability of a formula over the set of values $\mathbf{Val} \triangleq \mathbf{Loc} = \mathbb{N}_{>0}$: Let x be a fresh variable that does not appear in ϕ and Φ . Let ϕ' be the formula $x \mapsto x \star (\phi^\circ \star x \mapsto x)$, where ϕ° is obtained from ϕ by replacing every occurrence of nil by x . Further, let Φ' be the SID obtained from Φ by replacing every occurrence of nil by x and adding x as an additional parameter to every predicate. We now claim that ϕ is satisfiable over $\mathbf{Val} \triangleq \mathbf{Loc} \cup \{\text{nil}\}$ wrt SID Φ iff ϕ' is satisfiable over $\mathbf{Val} \triangleq \mathbf{Loc}$ wrt SID Φ' : Let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state with $\langle \mathfrak{s}, \mathfrak{h} \rangle \models \phi$. Let $\mathfrak{s}' = \mathfrak{s}[x/\ell]$ be the stack \mathfrak{s} extended by the mapping of x to some fresh location $\ell \in \mathbf{Loc} \setminus (\text{img}(\mathfrak{s}) \cup \text{locs}(\mathfrak{h}))$ (in particular, ℓ does not appear as a constant in ϕ) and let \mathfrak{h}' be the heap obtained from \mathfrak{h} by mapping all locations that map to nil to ℓ . Then it is easy to verify that $\langle \mathfrak{s}', \mathfrak{h}' \rangle \models \phi^\circ$. Moreover we have that $\langle \mathfrak{s}', \mathfrak{h}' \rangle \models x \mapsto x \star (\phi^\circ \star x \mapsto x)$ because of $\ell \notin \text{dom}(\mathfrak{h})$ (since ℓ has been chosen as a fresh location). For the other direction, let $\langle \mathfrak{s}, \mathfrak{h} \rangle$ be a state with $\langle \mathfrak{s}, \mathfrak{h} \rangle \models \phi'$. We observe that $\langle \mathfrak{s}, \mathfrak{h} \rangle \models \phi^\circ$ and $\mathfrak{s}(x) \notin \text{dom}(\mathfrak{h})$ because of $\langle \mathfrak{s}, \mathfrak{h} \rangle \models x \mapsto x \star (\phi^\circ \star x \mapsto x)$. Let \mathfrak{s}' be the stack obtained from \mathfrak{s} by removing the mapping for x and let \mathfrak{h}' be the heap obtained from \mathfrak{h} by mapping all locations that map to $\mathfrak{s}(x)$ to nil . Then it is easy to verify that $\langle \mathfrak{s}', \mathfrak{h}' \rangle \models \phi$.

The claim then follows from Theorem 9.3 and the observation that $|\Phi'| + |\phi'| = O(|\Phi| + |\phi|)$.