

Foundations for Decision Problems in Separation Logic with General Inductive Predicates

Timos Antonopoulos¹, Nikos Gorogiannis², Christoph Haase^{3*},
Max Kanovich⁴, and Joël Ouaknine¹

¹ Department of Computer Science, University of Oxford, UK

² Department of Computer Science, Middlesex University London, UK

³ LSV, CNRS & École Normale Supérieure (ENS) de Cachan, France

⁴ Department of Computer Science, Queen Mary University of London, UK

Abstract. We establish foundational results on the computational complexity of deciding entailment in Separation Logic with general inductive predicates whose underlying base language allows for pure formulas, pointers and existentially quantified variables. We show that entailment is in general undecidable, and EXPTIME-hard in a fragment recently shown to be decidable by Iosif *et al.* Moreover, entailment in the base language is Π_2^P -complete, the upper bound even holds in the presence of list predicates. We additionally show that entailment in essentially any fragment of Separation Logic allowing for general inductive predicates is intractable even when strong syntactic restrictions are imposed.

1 Introduction

Separation Logic (SL) is an extension of Hoare logic for reasoning about programs which manipulate heap data structures. Introduced in the early 2000s by O’Hearn, Ishtiaq, Reynolds and Yang [18, 20], it has been the starting point of a line of research that has led to a large body of theoretical and practical work.

In the early days, the potential of Separation Logic was recognised by proving the (partial) correctness of the Schorr-Waite graph marking algorithm [22] and Cheney’s copying garbage collector [6]. Those proofs were essentially carried out in a pen-and-paper fashion and demonstrated the strength of the paradigm underlying Separation Logic: *local reasoning*. The latter means that correctness proofs for heap-manipulating code should only depend on the portions of the heap accessed by the code and not the entire memory. Motivated by these positive results, research has been conducted on automating such proofs. On the one hand, support for Separation Logic has been integrated into proof assistants such as Coq, enabling semi-automated verification of program code, see *e.g.* [2]. On the other hand, a number of fully-automatic tools such as SMALLFOOT, SLAYER, SPACE INVADER or SLAD have been developed and successfully used to show absence of memory errors in low-level real-world code, see *e.g.* [11, 7, 5].

* Supported by the French ANR, REACHARD (grant ANR-11-BS02-001).

A crucial requirement for any of these tools is the ability to check applications of the consequence rule in Hoare logic, as it is this rule that underpins most methods of proof based on Separation Logic. The consequence rule, in turn, requires the ability to check entailment between Separation Logic formulas. However, entailment checking in full Separation Logic is undecidable [12, 10], thus these tools have to work with restricted, decidable fragments. Decidability, though, comes at the cost of reduced expressive power and, often, reduced generality. For instance, the fragment used by SMALLFOOT allows for reasoning about memory shapes built upon the hard-coded primitives of pointers and linked lists, essentially limiting its applicability to programs only involving those data structures; some efforts have been made in order to allow for reasoning about more generic list data structures, see *e.g.* [3]. The limitations of hard-coded inductive predicates have been realised by the community, and recent research has been conducted to enable automated reasoning about generic user-defined inductive predicates, inside the framework of Separation Logic, see *e.g.* [13, 9], or in related frameworks such as forest automata [16]. Notable recent progress has been made by Iosif *et al.* who showed decidability of satisfiability and entailment for a syntactic fragment of Separation Logic with general recursively defined predicates by establishing a reduction to Monadic Second Order Logic on graphs with bounded tree width [17]. Finally, Brotherston *et al.* have developed an EXPTIME-complete decision procedure for satisfiability of Separation Logic with general inductively defined predicates [8]. In the same paper it is shown that the problem becomes NP-complete if the arity of all predicates is bounded by a constant.

The goal of this paper is to contribute to this line of research and to establish foundational results on the inherent computational complexity of reasoning problems in Separation Logic with general inductively defined predicates. In order to obtain meaningful lower bounds, we restrict our analysis to the most basic syntactic fragment of Separation Logic comprising (positive) Boolean combinations of judgments on stack variables, both fixed and existentially quantified, and pointers. This fragment also forms the basis of the decidable fragments of Separation Logic from [17, 8]. Standard inductive data types are inductively expressible in this fragment, for instance singly-linked lists as used by SMALLFOOT [4] can be defined as follows:

$$\text{ls}(a, b) := \text{emp} \wedge a = b \mid \exists c. \text{pt}(a, c) * \text{ls}(c, b) \wedge a \neq b \quad (1)$$

Informally speaking, supposing that $\text{ls}(x, y)$ holds in a memory model, this definition states that there is a singly-linked list segment from the memory cell labeled with the stack variable x to the memory cell labeled with y if either the heap is empty and x is equal to y , or x is not equal to y and the heap can be split into two disjoint parts, indicated by the $*$ -conjunction, such that on the first part x is allocated and points to some cell a and heap cell a is the starting point of a singly-linked list segment ending in y in the other part.

The main results of our paper are as follows. In the first part, we consider entailment in Separation Logic with general inductive predicates. Given two assertions α, α' and a finite set of inductive predicates P referred to by α and α' ,

entailment is to decide whether the set of memory models of α is contained in the set of memory models of α' with respect to P . We show that this problem is undecidable in general and EXPTIME-hard when restricted to the decidable syntactic fragment defined by Iosif *et al.* In the second part, we take a closer look at entailment in the basic fragment of Separation Logic in the absence of inductive predicates, *i.e.*, Separation Logic with positive pure formulas, existentially quantified variables and pointers. We show that this problem is complete for Π_2^P , the second level of the polynomial hierarchy. The upper bound also holds when allowing for the above list predicate hard-coded in the syntax. Subsequently, we analyse the Π_2^P lower bound and define a natural syntactic fragment for which entailment is decidable in polynomial time, yet NP-hard in the presence of a list predicate, *i.e.*, one of the simplest possible inductive predicates. We discuss the results obtained in the conclusion at the end of the paper.

Some proofs have been omitted due to lack of space, but are included in a longer, online version of the paper, obtainable from the authors' webpages.

2 Preliminaries

Let X and Y be sets, and let $R \subseteq X \times Y$. We say that R is *functional* if for every $x \in X$ there is at most one $y \in Y$ with $(x, y) \in R$. Let Y be a countable, possibly infinite, set. We write $X \subseteq_{\text{fin}} Y$ if X is a finite subset of Y . Moreover, given countable sets X, Y , we write $f : X \rightarrow_{\text{fin}} Y$ if f is a function whose domain is a finite subset of X and its co-domain is Y . Given $f : X \rightarrow Y$, $x \in X$, $y \in Y$, we write $f[x \mapsto y]$ to denote the function f' such that $f'(z) \stackrel{\text{def}}{=} y$ if $z = x$, and $f'(z) \stackrel{\text{def}}{=} f(z)$ otherwise. Finally, given $i \leq j \in \mathbb{N}$, we write $[i, j]$ to denote the set $\{i, \dots, j\} \subseteq \mathbb{N}$ and $[i]$ as an abbreviation for $[1, i]$.

Graphs. Let L be a countable set of *labels*. We define *directed labeled graphs* (just *graphs* in the following) as tuples $G = (V, E, \ell)$, where V denotes the set of *nodes* or *vertices*, E is a subset of $V \times V$, and $\ell : L \rightarrow_{\text{fin}} V$ is a *labeling function*. If L is empty we omit ℓ and just write $G = (V, E)$. A graph $G = (V, E, \ell)$ is *undirected* if E is symmetric. If G is a graph, we also denote its set of nodes by $V(G)$ and its set of edges by $E(G)$. The *size of G* is defined as $|G| \stackrel{\text{def}}{=} |V(G)|$.

For interpretations below, we require a slightly more general class of graphs which we call selector graphs, inspired by [17]. A *selector graph* is a tuple $G = (V, E, \ell, s)$, where V and ℓ are as above, $s : V \rightarrow \mathbb{N}$ assigns an *arity* to each vertex, and $E : V \times \mathbb{N} \rightarrow V$ is a partial function such that E is defined precisely for every pair (v, i) with $v \in V$ and $i \in [s(v)]$. If $s(v) \in \{0, 1\}$ for all $v \in V$, we obtain partial functional graphs.

Formulas of Separation Logic. The subsequent definitions are partly adapted or inspired from [17]. Let **Vars** be a totally ordered countably infinite set of *variable names*, which is partitioned into disjoint infinite sets **EVars** and **FVars** representing sets of *existential variables* and *fixed stack variables*, respectively. We will usually use **a, b, c** for elements from **EVars**, and **x, y, z** will usually be elements from **Vars**. Variables in **FVars** will be used to represent fixed stack

variables. The purpose of the distinction between EVars and FVars is to help the reader to easily identify in which context a variable occurs. Let PNames be a finite set of *predicate names*, where each predicate has an associated arity $k \in \mathbb{N}$, and is written as $\text{pred}(\mathbf{a}_1, \dots, \mathbf{a}_k)$ with $\mathbf{a}_i \in \text{EVars}$ for $i \in [k]$. The syntax of *SL-assertions or SL-formulas over PNames* is given by the following grammar, where $\mathbf{x}, \mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}, \mathbf{y}_1, \dots, \mathbf{y}_k \in \text{Vars}$, $\mathbf{a}_1, \dots, \mathbf{a}_n \in \text{EVars}$, $m \geq 1$ and $n \geq 0$:

$$\begin{aligned} \varphi &::= \top \mid \perp \mid \mathbf{x} = \mathbf{y} \mid \neg\varphi \mid \varphi \wedge \varphi && \text{(pure formulas)} \\ \sigma &::= \text{emp} \mid \text{tt} \mid \text{pt}(\mathbf{x}, (\mathbf{x}_1, \dots, \mathbf{x}_m)) \mid \text{pred}(\mathbf{y}_1, \dots, \mathbf{y}_k) \mid \sigma * \sigma \mid \alpha && \text{(spatial formulas)} \\ \alpha &::= \exists \mathbf{a}_1, \dots, \mathbf{a}_n. \sigma \wedge \varphi && \text{(SL-assertions)} \end{aligned}$$

Here, $\text{pt}(\mathbf{x}, \mathbf{y})$ is the *points-to* or *pointer predicate* of an arbitrary arity m , and the $*$ -conjunction is commutative, *i.e.*, SL-assertions are considered equivalent up to permutations of $*$ -connected subformulas. We say that the SL-assertion $\alpha = \exists \mathbf{a}_1, \dots, \mathbf{a}_n. \sigma \wedge \varphi$ is *flat* if σ contains no SL-assertion α' as a subformula. Given an SL-assertion $\alpha = \exists \mathbf{a}_1, \dots, \mathbf{a}_m. (\sigma * \exists \mathbf{b}_1, \dots, \mathbf{b}_n. (\sigma' \wedge \varphi')) \wedge \varphi$, and supposing without loss of generality that $\{\mathbf{a}_i : i \in [m]\} \cap \{\mathbf{b}_j : j \in [n]\} = \emptyset$, we can exhaustively rewrite the formula α as $\exists \mathbf{a}_1, \dots, \mathbf{a}_m, \mathbf{b}_1, \dots, \mathbf{b}_n. (\sigma * \sigma') \wedge \varphi \wedge \varphi'$. Thus we may assume with no loss of generality that an SL-assertion is flat.

Remark 1. We have imposed flatness and $*$ -commutativity as syntactic properties. This is merely for presentational convenience in order to save space, as these properties follow from the semantics definition below.

We call φ *positive* if φ is a conjunction of literals $\mathbf{x} = \mathbf{y}$ and $\mathbf{x} \neq \mathbf{y}$. Moreover, we say that α is positive if φ is positive, and that α is *reduced* if no $\text{pred}(\mathbf{y}_1, \dots, \mathbf{y}_k)$ occurs in σ . By $\text{vars}(\alpha) \subseteq \text{Vars}$ we denote the *set of all variables occurring in α* . The *size of α* , denoted by $|\alpha|$, is defined to be the number of symbols in α .

A *set P of inductive predicates over PNames* is a finite set of *definitions*

$$\text{pred}(\mathbf{a}_1, \dots, \mathbf{a}_k) := \alpha_1 \mid \dots \mid \alpha_m$$

such that each $\mathbf{a}_i \in \text{EVars}$, α_i is a flat SL-assertion $\alpha_i = \exists \mathbf{b}_1, \dots, \mathbf{b}_n. \sigma \wedge \varphi$ over PNames such that $\{\mathbf{a}_i : i \in [m]\} \cap \{\mathbf{b}_j : j \in [n]\} = \emptyset$, and each predicate name $\text{pred}(\mathbf{a}_1, \dots, \mathbf{a}_k)$ occurs exactly once on the left-hand side of a definition in P . Moreover, we require that each α_i has no unbounded existential variables, *i.e.*, for each α_m as above, $\text{vars}(\alpha_m) \subseteq \text{FVars} \cup \{\mathbf{a}_i, \mathbf{b}_j : i \in [k], j \in [n]\}$.⁵ Given $\mathbf{x}_1, \dots, \mathbf{x}_k \in \text{Vars}$, define $\text{pred}[\mathbf{x}_1/\mathbf{a}_1, \dots, \mathbf{x}_k/\mathbf{a}_k] \stackrel{\text{def}}{=} \{\alpha_i[\mathbf{x}_1/\mathbf{a}_1, \dots, \mathbf{x}_k/\mathbf{a}_k] : i \in [k]\}$, where $\alpha_i[\mathbf{x}_1/\mathbf{a}_1, \dots, \mathbf{x}_k/\mathbf{a}_k]$ is obtained from α_i by replacing each occurrence of \mathbf{a}_j with \mathbf{x}_j for $j \in [k]$. Given a flat assertion $\alpha = \exists \mathbf{c}_1, \dots, \mathbf{c}_p. \sigma \wedge \varphi$, an *unraveling of α with respect to P* is obtained by replacing each $\text{pred}(\mathbf{x}_1, \dots, \mathbf{x}_k)$ occurring as a subformula in σ with some $\alpha' \in \text{pred}[\mathbf{x}_1/\mathbf{a}_1, \dots, \mathbf{x}_k/\mathbf{a}_k]$. We write $\alpha \rightarrow_P \beta$ if β is an unraveling of α with respect to P , and denote the reflexive transitive closure of \rightarrow_P by \rightarrow_P^* . An unraveling $\alpha \rightarrow_P^* \beta$ is *complete* if no $\text{pred}(\mathbf{x}_1, \dots, \mathbf{x}_k)$ occurs in β .

⁵ In a slight departure from convention, for presentational convenience we allow free variables to appear in the body of definitions; such predicates can always be transformed to equivalent ones where the previously free variables are parameters, w.l.o.g.

Example 2. Taking P to be the singleton set consisting of the definition of $\text{ls}(\mathbf{a}, \mathbf{b})$ from Equation (1), we have

$$\text{ls}(\mathbf{x}, \mathbf{y}) \rightarrow_P^* \exists c_1, c_2. \text{pt}(\mathbf{x}, c_1) * \text{pt}(c_1, c_2) * \text{ls}(c_2, \mathbf{y}) \wedge \mathbf{x} \neq \mathbf{y} \wedge c_1 \neq \mathbf{y}$$

with respect to P , which is *not* a complete unraveling. A complete unraveling is $\text{ls}(\mathbf{x}, \mathbf{y}) \rightarrow_P^* \exists c. \text{pt}(\mathbf{x}, c) * \text{emp} \wedge \mathbf{x} \neq \mathbf{y} \wedge c = \mathbf{y}$.

For convenience, we sometimes use a generalised $*$ -conjunction and, given spatial formulas $\sigma_1, \dots, \sigma_n$, write $*_{1 \leq i \leq n} \sigma_i$ for $\sigma_1 * \dots * \sigma_n$. Likewise, we write $\text{pred}(\mathbf{a}_1, \dots, \mathbf{a}_k) := \prod_{i \in [k]} \alpha_i$ for $\text{pred}(\mathbf{a}_1, \dots, \mathbf{a}_k) := \alpha_1 \mid \dots \mid \alpha_n$. Moreover, $\exists \mathbf{a}_1, \dots, \mathbf{a}_n. \sigma$ abbreviates $\exists \mathbf{a}_1, \dots, \mathbf{a}_n. \sigma \wedge \top$; we may also write *e.g.* $\text{pt}(\mathbf{x}, (-, \mathbf{y}))$ as a shorthand for $\exists \mathbf{a}. \text{pt}(\mathbf{x}, (\mathbf{a}, \mathbf{y}))$, where $\mathbf{a} \in \text{FVars}$ is a fresh existential variable. Furthermore, $\exists_{i \in [n]} \mathbf{a}_i. \sigma \wedge \varphi$ abbreviates $\exists \mathbf{a}_1, \dots, \mathbf{a}_n. \sigma \wedge \varphi$. If PNames is clear from the context we will omit stating it explicitly.

Interpretations. As stated above, interpretations are given in terms of selector graphs. This diversion from the more commonly found “heap-and-stack model” found in the literature is for technical convenience only, and it is easy to translate between the two interpretation domains.

An *SL-interpretation*, or simply *interpretation*, \mathcal{I} is a selector graph $\mathcal{I} = (V^{\mathcal{I}}, E^{\mathcal{I}}, \ell^{\mathcal{I}}, s^{\mathcal{I}})$ such that $\ell^{\mathcal{I}} : \text{FVars} \rightarrow_{\text{fin}} V$. For $\mathbf{x}_1, \dots, \mathbf{x}_n \in \text{FVars}$ and for $v_1, \dots, v_n \in V^{\mathcal{I}}$, we denote by $\mathcal{I}[\mathbf{x}_1 \mapsto v_1; \dots; \mathbf{x}_n \mapsto v_n]$ the SL-interpretation $\hat{\mathcal{I}} = (V^{\mathcal{I}}, E^{\mathcal{I}}, \hat{\ell}^{\mathcal{I}}, s^{\mathcal{I}})$, where $\hat{\ell}^{\mathcal{I}} \stackrel{\text{def}}{=} \ell^{\mathcal{I}}[\mathbf{x}_1 \mapsto v_1; \dots; \mathbf{x}_n \mapsto v_n]$, and we call such an interpretation an *extension* of \mathcal{I} .

In our interpretations, nodes with arity greater than zero are the equivalent to allocated heap cells in the “heap-and-stack model”, while record fields are represented by the different selectors. We define the $*$ -decomposition of \mathcal{I} as follows: $\mathcal{I} = \mathcal{I}_1 * \mathcal{I}_2$ iff $\mathcal{I}_1 = (V^{\mathcal{I}_1}, E^{\mathcal{I}_1}, \ell^{\mathcal{I}_1}, s^{\mathcal{I}_1})$ and $\mathcal{I}_2 = (V^{\mathcal{I}_2}, E^{\mathcal{I}_2}, \ell^{\mathcal{I}_2}, s^{\mathcal{I}_2})$ such that $V^{\mathcal{I}} = V^{\mathcal{I}_1} = V^{\mathcal{I}_2}$; $\ell^{\mathcal{I}}, \ell^{\mathcal{I}_1}$ and $\ell^{\mathcal{I}_2}$ coincide; for $i \in \{1, 2\}$, either $s^{\mathcal{I}_i}(v) = 0$ or $s^{\mathcal{I}_i}(v) = s^{\mathcal{I}}(v)$, and $s^{\mathcal{I}}(v) = s^{\mathcal{I}_1}(v) + s^{\mathcal{I}_2}(v)$ for all $v \in V^{\mathcal{I}}$; for $i \in \{1, 2\}$, if $s^{\mathcal{I}_i}(v) > 0$ then $E^{\mathcal{I}_i}(v, j) = E^{\mathcal{I}}(v, j)$ for all $v \in V^{\mathcal{I}}$ and $j \in [s^{\mathcal{I}_i}(v)]$.

Semantics of SL-assertions. The semantics of flat reduced SL-assertions is defined by structural induction. Let $\mathcal{I} = (V^{\mathcal{I}}, E^{\mathcal{I}}, \ell^{\mathcal{I}}, s^{\mathcal{I}})$ be an SL-interpretation and φ a pure formula only over variable names from FVars, the *satisfaction relation* $\mathcal{I} \models \varphi$ is defined such that $\mathcal{I} \models \top$ holds always, $\mathcal{I} \models \perp$ never holds, $\mathcal{I} \models \mathbf{x} = \mathbf{y}$ iff $\ell^{\mathcal{I}}(\mathbf{x}) = \ell^{\mathcal{I}}(\mathbf{y})$, $\mathcal{I} \models \neg \varphi$ iff $\mathcal{I} \not\models \varphi$, and $\mathcal{I} \models \varphi_1 \wedge \varphi_2$ iff $\mathcal{I} \models \varphi_1$ and $\mathcal{I} \models \varphi_2$.

For reduced flat spatial formulas σ such that $\text{vars}(\sigma) \subseteq \text{FVars}$, we define $\mathcal{I} \models \text{emp}$ iff $s^{\mathcal{I}}(v) = 0$ for all $v \in V^{\mathcal{I}}$ and $\mathcal{I} \models \text{tt}$ holds always. Moreover, $\mathcal{I} \models \sigma_1 * \sigma_2$ iff $\mathcal{I} = \mathcal{I}_1 * \mathcal{I}_2$ such that $\mathcal{I}_1 \models \sigma_1$ and $\mathcal{I}_2 \models \sigma_2$, and finally, $\mathcal{I} \models \text{pt}(\mathbf{x}, (\mathbf{x}_1, \dots, \mathbf{x}_m))$ iff

- $v = \ell^{\mathcal{I}}(\mathbf{x})$, $s^{\mathcal{I}}(v) = m$, $s^{\mathcal{I}}(v') = 0$ for all $v' \in V^{\mathcal{I}} \setminus v$; and
- $E^{\mathcal{I}}(v, i) = \ell^{\mathcal{I}}(\mathbf{x}_i)$ for all $i \in [m]$.

For a flat reduced SL-assertion $\alpha = \exists \mathbf{a}_1, \dots, \mathbf{a}_n. \sigma \wedge \varphi$, we define $\mathcal{I} \models \alpha$ iff there is an extension $\hat{\mathcal{I}} = \mathcal{I}[\mathbf{x}_1 \mapsto v_1, \dots, \mathbf{x}_n \mapsto v_n]$ for fresh variables $\mathbf{x}_1, \dots, \mathbf{x}_n \in \text{FVars}$

such that $\hat{\mathcal{I}} \models \sigma[x_1/a_1, \dots, x_n/a_n]$ and $\hat{\mathcal{I}} \models \varphi[x_1/a_1, \dots, x_n/a_n]$. We call \mathcal{I} a *model* of α if $\mathcal{I} \models \alpha$. Given α and a set of inductive predicates P , we write $\mathcal{I} \models_P \alpha$ if $\mathcal{I} \models \alpha'$ for some α' obtained from a complete unraveling $\alpha \rightarrow_P^* \alpha'$. Given α and α' over a set of inductive predicates P , we write $\alpha \models_P \alpha'$ iff whenever $\mathcal{I} \models_P \alpha$ then $\mathcal{I} \models_P \alpha'$. Given α over a set of inductive predicates P , *satisfiability* is to decide whether there is an interpretation \mathcal{I} such that $\mathcal{I} \models_P \alpha$. Given \mathcal{I} , *model checking* is to decide $\mathcal{I} \models_P \alpha$. The main decision problem of interest in this paper is entailment, defined as follows.

ENTAILMENT

INPUT: SL assertions α, α' with respect to a set P of inductive predicates.
QUESTION: Does $\alpha \models_P \alpha'$?

3 Entailment in the Presence of Inductive Predicates.

In this section, we show that entailment with general inductive predicates is undecidable when no restrictions are imposed. Subsequently, we give an EXPTIME lower bound for the fragment introduced by Iosif *et al.* [17].

General undecidability. We show undecidability via a reduction from the undecidable Post Correspondence Problem [19].

POST CORRESPONDENCE PROBLEM (PCP)

INPUT: A finite set of tiles $(v_1, w_1), \dots, (v_k, w_k)$, $v_i, w_i \in \{0, 1\}^*$.
QUESTION: Does there exist a sequence $s_1 s_2 \dots s_\ell \in \{1, \dots, k\}^\ell$, $\ell > 0$ such that $v_{s_1} v_{s_2} \dots v_{s_\ell} = w_{s_1} w_{s_2} \dots w_{s_\ell}$?

For any $u \in \{0, 1\}^*$, denote by $|u|$ the length of each tile, and by $u(i)$ the i -th symbol of u , for $1 \leq i \leq |u|$. For example, if $u = 01101$, we have $|u| = 5$ and $u(3) = 1$. Let $(v_1, w_1), \dots, (v_k, w_k)$ be an instance of PCP. The set of predicates P in Figure 1 establishes a reduction such that this instance has a solution iff $\text{PCP}(x, y) \wedge x_0 \neq x_1 \not\models_P \overline{\text{PCP}}(x, y)$. The intuition behind these definitions is as follows. For $x, y \in \text{FVars}$, $\text{PCP}(x, y)$ generates the set of all possible tilings for a given instance: in any model the v_i -tilings begin at x and the w_i -tilings at y . The fixed stack variables $x_0, x_1 \in \text{FVars}$ are used to represent the corresponding symbols 0 and 1. Likewise, $\overline{\text{PCP}}(x, y)$ generates all tilings which are incorrect. This is the case if the model is empty, there are two symbols at the same position which are different (*cf.* $\text{NEqualPair}(x, y)$), or the length of the strings encoded in a model is different (*cf.* $\text{NEqualLen}(x, y)$).

Theorem 3. *Entailment in Separation Logic with general inductive predicates is undecidable.*

Remark 4. An anonymous referee remarked that our reduction can also be applied for showing that satisfiability in the presence of conjunction over spatial formulas and general inductive predicates is undecidable. The models of the subsequent predicate encode all pairs of equal strings:

$$\text{EqPairs}(a, b) = \text{emp} \mid \parallel_{i \in \{0, 1\}} \exists p, r. \text{pt}(x, (x_i, p)) * \text{pt}(y, (x_i, r)) * \text{EqPairs}(p, r).$$

$$\begin{aligned}
\text{PCP}(\mathbf{a}, \mathbf{b}) &:= \text{emp} \mid \text{Tile}_1(\mathbf{a}, \mathbf{b}) \mid \cdots \mid \text{Tile}_k(\mathbf{a}, \mathbf{b}) \\
\text{Tile}_i(\mathbf{a}, \mathbf{b}), i \in [k] &:= \exists \mathbf{p}_0, \dots, \mathbf{p}_{|v_i|}, \mathbf{r}_0, \dots, \mathbf{r}_{|w_i|}. * \text{pt}(\mathbf{p}_j, (\mathbf{x}_{v_i(j+1)}, \mathbf{p}_{j+1})) * \\
&\quad * \text{pt}(\mathbf{r}_j, (\mathbf{x}_{w_i(j+1)}, \mathbf{r}_{j+1})) * \text{PCP}(\mathbf{p}_{|v_i|}, \mathbf{r}_{|w_i|}) \wedge \mathbf{a} = \mathbf{p}_0 \wedge \mathbf{b} = \mathbf{r}_0 \\
\text{NEqualPair}(\mathbf{a}, \mathbf{b}) &:= \exists \mathbf{p}, \mathbf{r}. \text{pt}(\mathbf{a}, (-, \mathbf{p})) * \text{pt}(\mathbf{b}, (-, \mathbf{r})) * \text{NEqualPair}(\mathbf{p}, \mathbf{r}) \\
&\quad \mid \exists \mathbf{c}, \mathbf{d}. \text{pt}(\mathbf{a}, (\mathbf{c}, -)) * \text{pt}(\mathbf{b}, (\mathbf{d}, -)) * \text{tt} \wedge \mathbf{c} \neq \mathbf{d} \\
\text{Tail}(\mathbf{a}) &:= \text{emp} \mid \exists \mathbf{b}. \text{pt}(\mathbf{a}, (-, \mathbf{b})) * \text{Tail}(\mathbf{b}) \\
\text{NEqualLen}(\mathbf{a}, \mathbf{b}) &:= \exists \mathbf{x}, \mathbf{p}, \mathbf{r}. \text{pt}(\mathbf{a}, (\mathbf{x}, \mathbf{p})) * \text{pt}(\mathbf{b}, (\mathbf{x}, \mathbf{r})) * \text{NEqualLen}(\mathbf{p}, \mathbf{r}) \\
&\quad \mid \exists \mathbf{p}. \text{pt}(\mathbf{a}, (-, \mathbf{p})) * \text{Tail}(\mathbf{p}) \mid \exists \mathbf{r}. \text{pt}(\mathbf{b}, (-, \mathbf{r})) * \text{Tail}(\mathbf{r}) \\
\overline{\text{PCP}}(\mathbf{a}, \mathbf{b}) &:= \text{emp} \mid \text{NEqualPair}(\mathbf{a}, \mathbf{b}) \mid \text{NEqualLen}(\mathbf{a}, \mathbf{b})
\end{aligned}$$

Fig. 1: The set P of inductive predicates for the reduction from PCP.

It is then easy to conjoin $\text{EqPairs}(\mathbf{x}, \mathbf{y})$ with $\text{PCP}(\mathbf{x}, \mathbf{y})$ such that a model exists if, and only if, the given PCP instance has a solution.

Inductive Predicates with Bounded Tree Width. Iosif *et al.* define in [17] a fragment of Separation Logic by syntactically restricting the definitions of inductive predicates such that all models have bounded tree width. In particular, their fragment requires that there is *exactly one* points-to predicate in any definition, which is clearly not the case in the reduction from PCP. Moreover, briefly speaking, further restrictions require that in each predicate definition, if a predicate name occurs in the body of a predicate definition then a points-to predicate occurs in the definition as well, that every existentially quantified variable is eventually allocated, and some further subtle technical conditions. We omit further details for space reason and show that entailment is EXPTIME-hard in this fragment. The reader can easily verify that our reduction fulfils the requirements defined in [17].

Our reduction is from the language inclusion problem for non-deterministic top-down binary finite tree automata. A *prefix closed set of strings* over $\{0, 1\}$ is a set of strings S such that for each $s \in S$ and any prefix s_p of s , s_p is also in S . A *binary ordered tree* t over a finite *alphabet* Σ is a tuple (N, Σ, ℓ) , where N is a prefix closed set of strings over $\{0, 1\}$ denoting the *nodes of the tree*, where for each $s \in N$, $s \cdot 1 \in N$, if and only if $s \cdot 0 \in N$, and $\ell : N \rightarrow \Sigma$ is a function assigning *labels* to nodes of the tree. The root of a tree is the empty string ϵ , and for any two nodes s and $s \cdot i$, for $i \in \{0, 1\}$, $s \cdot i$ is a child node of s . We say that a node $s \in N$ is a *leaf node* if it has no child nodes, and a node is *internal* otherwise.

Recall that a *finite non-deterministic top-down tree automaton (NFTA)* A is a tuple (Σ, Q, δ, I) , where Σ is a finite *alphabet*, Q is a finite set of *states* with a designated state q_{leaf} , $I \subseteq Q$ is the set of *initial or accepting states*, and $\delta : Q \times \Sigma \rightarrow 2^{Q \times Q}$ is the *transition function* such that for all $\sigma \in \Sigma$,

$\delta(q_{leaf}, \sigma) = \emptyset$. A *run of A* on a tree $t = (N, \Sigma, \ell)$ is a function $\rho : N \rightarrow Q$ assigning states to the nodes of t such that for each internal node $s \in N$, $(\rho(s \cdot 0), \rho(s \cdot 1)) \in \delta(\rho(s), \ell(s))$ and for each leaf node $s \in N$, $\rho(s) = q_{leaf}$. A run is *accepting* if $\rho(\epsilon) \in I$, and the *language $\mathcal{L}(A)$ accepted by a NFTA A* is the set of trees t for which there is an accepting run of A on t .

NFTA LANGUAGE INCLUSION PROBLEM

INPUT: Two NFTA A_1 and A_2 .

QUESTION: Does $\mathcal{L}(A_1) \subseteq \mathcal{L}(A_2)$ hold?

A classical result states that the language inclusion problem for non-deterministic tree automata is complete for EXPTIME [21]. Let $A = (\Sigma, Q, \delta, I)$ be an NFTA. We define the subsequent set P of inductive predicates, where $\text{Tree}_{q,\sigma}(\mathbf{a})$ is defined for every $\sigma \in \Sigma$ and $q \in Q$ for which $\delta(q, \sigma)$ is non-empty:

$$\begin{aligned} \text{Tree}_{q,\sigma}(\mathbf{a}) &:= \parallel_{\substack{(q_1, q_2) \in \delta(q, \sigma) \\ \sigma_1, \sigma_2 \in \Sigma}} \exists l, r. \text{pt}(\mathbf{a}, (\mathbf{s}_\sigma, l, r)) * \text{Tree}_{q_1, \sigma_1}(l) * \text{Tree}_{q_2, \sigma_2}(r) \\ &\quad \delta(q_1, \sigma_1) \neq \emptyset \vee q_1 = q_{leaf} \\ &\quad \delta(q_2, \sigma_2) \neq \emptyset \vee q_2 = q_{leaf} \\ \text{Tree}_{q_{leaf}, \sigma}(\mathbf{a}) &:= \text{pt}(\mathbf{a}, \mathbf{s}_\sigma) \\ \text{Trees}_A(\mathbf{a}) &:= \parallel_{\substack{q \in I \\ \sigma \in \Sigma}} \exists \mathbf{b}. \text{pt}(\mathbf{a}, \mathbf{b}) * \text{Tree}_{q, \sigma}(\mathbf{b}) \end{aligned}$$

In any model, the predicate $\text{Trees}_A(\mathbf{x})$ encodes all trees in $\mathcal{L}(A)$: apart from the node labeled with \mathbf{x} , each allocated vertex represents a node of the tree, the first selector represents the label of the node, and the subsequent selectors represent respectively the left and right descendants, if the node is an internal node. The additional pointer at \mathbf{x} is for technical reasons in order to comply with the restrictions defined in [17]. It is now easily checked that given two NFTA A_1 and A_2 over some alphabet $\Sigma = \{\sigma_1, \dots, \sigma_n\}$,

$$\text{Trees}_{A_1}(\mathbf{x}) \wedge \bigwedge_{1 \leq i \neq j \leq n} \mathbf{s}_{\sigma_i} \neq \mathbf{s}_{\sigma_j} \models_P \text{Trees}_{A_2}(\mathbf{x})$$

is a valid entailment if, and only if, $\mathcal{L}(A_1) \subseteq \mathcal{L}(A_2)$.

Theorem 5. *Deciding entailment in Separation Logic with inductive predicates with bounded tree width as defined in [17] is EXPTIME-hard.*

It is worth emphasizing that hardness already holds if the arity of the pointer predicates is fixed to three. Also note that the EXPTIME-hardness proof for satisfiability provided in [8] does not trivially establish that entailment in the fragment defined in [17] is EXPTIME-hard since the definitions given in [8] are not in the fragment of [17].

Also, note that in [21] it is shown that for two NFTA that accept finite languages, the language inclusion problem is PSPACE-complete, and therefore the proof of Theorem 5 can be adapted to show PSPACE-hardness of entailment with inductive predicates not involving cyclic definitions.

4 Entailment for Fixed Fragments

The primary goal of this section is to first study the complexity of entailment in the base language defined in Section 2, and subsequently in the base language equipped with a fixed list predicate as defined in (1), which is a fragment commonly found in the program verifiers discussed in the introduction.

We first show that entailment in the base language is Π_2^P -complete. Moreover, we additionally outline that the upper bound even holds in the presence of the aforementioned list predicate. This result complements the previous section in that it indicates that for specific and fixed natural decidable fragments involving cyclic definitions of small arity the EXPTIME lower bound can be avoided.

In the second part we analyse the lower bound from the first part and consider natural syntactic fragments defined in terms of structural properties of graphs representing SL-assertions. It has been shown that such restrictions can lead to polynomial-time decision procedures for entailment when dropping existentially quantified variables [14] and also decidability results for more expressive extensions of our base assertion language [7]. We show that basically there is no hope of achieving polynomial-time decision procedures in the presence of list predicates and existentially quantified variables, even when strong syntactic restrictions on the assertions are imposed.

Entailment in the General Case. We begin with the lower bound and show hardness for the base language, *i.e.* the language having only points-to predicates, via a reduction from a generalisation of *graph three-colorability* that has been defined in [1]. Recall that given an undirected graph $G = (V, E)$, graph three-colorability is to decide whether there is a *three-coloring* $f : V \rightarrow \{1, 2, 3\}$ such that $f(v) \neq f(w)$ for all $\{v, w\} \in E$. A *leaf coloring* of G is a function $f : V_l \rightarrow \{1, 2, 3\}$, where V_l is the set of vertices of G with degree one, *i.e.*, those nodes that have exactly one incident edge. The generalisation of graph three-colorability is given as follows.

2-ROUND 3-COLORABILITY

INPUT: Undirected graph $G = (V, E)$.

QUESTION: For every fixed leaf coloring f of G , can f be extended to a three-coloring of G ?

It has been shown in [1] that 2-ROUND 3-COLORABILITY is Π_2^P -complete. We now show hardness of entailment for SL-formulas via a reduction from 2-ROUND 3-COLORABILITY. To this end, we construct flat reduced SL-assertions α, α' such that the graph $G = (V, E)$ is a valid instance of 2-ROUND 3-COLORABILITY iff $\alpha \models \alpha'$. We partition V into disjoint sets $V' = \{v_1, \dots, v_n\}$ of nodes with degree greater than one and $V'' = \{v_{n+1}, \dots, v_m\}$ of nodes with degree equal to one, and define α and α' such that

$$\alpha \stackrel{\text{def}}{=} \bigstar_{\substack{i \in [3] \\ j \in [n]}} \text{pt}(x_{i,j}, y_i) * \bigstar_{n < j \leq m} \text{pt}(x_j, z_j) \wedge \bigwedge_{n < i \leq m} \bigvee_{j \in [3]} z_i = y_j \wedge \bigwedge_{1 \leq i \neq j \leq 3} y_i \neq y_j$$

$$\alpha' \stackrel{\text{def}}{=} \exists_{i \in [n]} \mathbf{a}_i \cdot \exists_{j \in [m]} \mathbf{b}_j \cdot \bigstar_{i \in [n]} \text{pt}(\mathbf{a}_i, \mathbf{b}_i) * \bigstar_{n < j \leq m} \text{pt}(x_j, \mathbf{b}_j) * \text{tt} \wedge \bigwedge_{(v_i, v_j) \in E} \mathbf{b}_i \neq \mathbf{b}_j.$$

Intuitively speaking, the pointers $\text{pt}(x_j, z_j)$ in α can choose in any model \mathcal{I} of α a coloring of the leaves of G , represented by the variables y_i , $i \in [3]$. The $x_{i,j}$ are allocated in order to enable α' to choose a coloring of the nodes which are not leaves. Consequently in a model \mathcal{I} of α , an extension of \mathcal{I} determining the existentially quantified variables b_i of α' determines a three-coloring of G . Conversely, if $\alpha \not\models \alpha'$ then the counter-model \mathcal{I} encodes a coloring of the leaves which cannot be extended to a three-coloring.

It is not difficult to see that Π_2^P -hardness of entailment can already be established for SL-assertions without disjunction, by only requiring $y_i \neq y_j$ for all $1 \leq i \neq j \leq 3$ in the pure part of α : if $\alpha \models \alpha'$ then, in particular, any leaf coloring can be extended to a three-coloring since a subset of the models of α will encode all possible leaf colorings. The converse direction then follows as above. In addition, by introducing additional existentially quantified slack variables, the hardness proof can also be tightened in a way such that no “tt” is required in spatial formulas, *i.e.*, hardness holds in non-intuitionistic fragments. Finally, this reduction can be reused in order to show NP-hardness of the model checking problem via a reduction from 3-COLORABILITY.

Since the size of all relevant models is *a priori* fixed by the size of the formulas under consideration, a Π_2^P upper bound follows trivially.

Theorem 6. *Entailment for flat reduced SL-assertions is Π_2^P -complete.*

For the remainder of this section, we turn towards entailment in the base language equipped with an additional fixed list predicate as defined in (1) and restrict our attention to pointer predicates of arity one, and consequently to interpretations which are functional graphs.

First, we can also establish a Π_2^P upper bound for entailment in this fragment by showing a small-model property. The main idea is that for a sufficiently large \mathcal{I} such that $\mathcal{I} \models \alpha$ and $\mathcal{I} \not\models \alpha'$, we can find an instantiation of an $\text{ls}(x, y)$ predicate in \mathcal{I} such that the path between x and y is long enough that we can obtain a new \mathcal{I}' by removing a vertex occurring on this path while ensuring that the newly obtained \mathcal{I}' is still a counter-model witnessing $\alpha \not\models \alpha'$.

Lemma 7. *Let α, α' be SL-assertions, let $n = |\text{vars}(\alpha)| + |\text{vars}(\alpha')|$ and suppose that $\alpha \not\models \alpha'$. Then there exists an \mathcal{I} witnessing $\alpha \not\models \alpha'$ with $|V^{\mathcal{I}}| \in O(n^2)$.*

This lemma immediately yields a Π_2^P upper bound: in order to show $\alpha \not\models \alpha'$, we can guess a small model \mathcal{I} of α and then verify with an NP oracle that $\mathcal{I} \not\models \alpha'$.

Theorem 8. *Entailment for flat reduced SL-assertions with a fixed list predicate is Π_2^P -complete.*

Entailment under Structural Restrictions. The goal of this section is to argue that entailment in essentially any useful fragment is intractable in the presence of existential quantification and list predicates, even under severe syntactic restrictions. In the following, we will only consider positive SL-assertions, since otherwise non-entailment is trivially coNP-hard.

In order to identify syntactic fragments with potentially polynomial-time entailment problems, we look at properties of graphs representing SL-formulas. Let $G = (V, E)$ be a directed graph and $v \in V$ a vertex of G . Then, define functions $pred(v) \stackrel{\text{def}}{=} \{v' \in V : (v', v) \in E\}$ and $succ(v) \stackrel{\text{def}}{=} \{v' \in V : (v, v') \in E\}$. A node $v \in V$ is a *source node* if $pred(v) = \emptyset$, and v is a *sink node* if $succ(v) = \emptyset$. Let $\alpha = \sigma \wedge \varphi$ be an SL-assertion, and $x \in vars(\alpha)$ be a variable of α . Then define the set $\text{Eq}(\varphi, x) = \{y \in vars(\alpha) : \text{for all } \mathcal{I}, \text{ if } \mathcal{I} \models \varphi \text{ then } \mathcal{I} \models x = y\}$. Next, we define the *graph* $G(\alpha)$ *corresponding to* α as $G(\alpha) = (V_\alpha, E_\alpha, \ell_\alpha)$, where the set of vertices is defined as $V_\alpha \stackrel{\text{def}}{=} \{\text{Eq}(\varphi, x) : x \in vars(\alpha)\}$, and the set of edges as $E_\alpha \stackrel{\text{def}}{=} \{(\text{Eq}(\varphi, x), \text{Eq}(\varphi, y)) : \text{pt}(x, y) \text{ or } \text{ls}(x, y) \text{ occurs in } \sigma\}$. Finally, ℓ_α is such that $\ell_\alpha(x) = \text{Eq}(\varphi, x)$ for all $x \in vars(\alpha)$. A node $v \in V_\alpha$ is *fixed* if there is $x \in \text{FVars}$ such that $\ell_\alpha(x) = v$.

Inspecting the Π_2^P -hardness proof above, we see that one fundamental source of complexity comes from having pointers $\text{pt}(\mathbf{a}, \mathbf{b})$ with $\mathbf{a}, \mathbf{b} \in \text{EVars}$, *i.e.*, the graph corresponding to α' above has source and sink nodes which are not fixed. On the one hand, when not allowing for list predicates, if all source nodes of an SL-assertion were to be fixed, entailment between formulas in such a suitably defined fragment would trivially become polynomial-time decidable. The main reason for this is that in any model \mathcal{I} of $\exists \mathbf{a}.\text{pt}(x, \mathbf{a})$, \mathbf{a} would have to be instantiated with the *unique* successor of the vertex labeled with x . However, such a fragment would only allow for reasoning about models whose size is *a priori* fixed by the size of the antecedent, which limits its usefulness. On the other hand, when allowing for list predicates, the proof of NP-hardness of abduction (given in [15]) can be easily adapted to show that entailment becomes intractable even if source nodes are required to be fixed.

This leaves us with an interesting case, which we introduce by considering an example of an instance of entailment:

$$\text{ls}(x, y) \wedge x \neq y \models \exists \mathbf{a}.\text{pt}(x, \mathbf{a}) * \text{ls}(\mathbf{a}, y).$$

The validity of this entailment rests on the fact that x has a successor in any model containing a non-empty list from x to y . In this example, the consequent is a formula of the fragment of the assertion language which we are going to consider in the remainder of this section: an SL-assertion α is in the *fixed endpoints fragment* if all source and sink nodes of the graph $G(\alpha)$ corresponding to α are fixed. We show CONP-hardness of entailment in this fragment via a reduction from an NP-complete variant of the Hamiltonian path problem.

FIXED VERTEX HAMILTONIAN PATH (FVHP)

INPUT: A directed graph $G = (V, E)$ and $v \in V$.

QUESTION: Does there exist a Hamiltonian path in G ending in v ?

Given a graph $G = (V, E)$, a vertex $v \in V$ and an instance of FVHP such that $V = \{v_1, \dots, v_{N+1}\}$ and $v = v_{N+1}$, we show how to compute in polynomial time SL-formulas α, α' in the fixed endpoints fragment such that $\alpha \not\models \alpha'$ if and only if G is a valid instance of FVHP. For $i \in [N + 1]$ and $j \in [N]$, for the spatial

parts of α and α' we define:

$$\begin{aligned} \text{node}_j &\stackrel{\text{def}}{=} \text{ls}(e_j^s, a_j^0) * \bigstar_{k \in [0, N-1]} \text{ls}(a_j^k, b_j^{k+1}) * \bigstar_{k \in [N-1]} \text{ls}(b_j^k, a_j^k) * \text{ls}(b_j^N, e_j^f) \\ \text{order}_i^0 &\stackrel{\text{def}}{=} \text{pt}(s_i^0, f_i^0) \\ \text{order}_i^j &\stackrel{\text{def}}{=} \text{ls}(s_i^j, d_i^j) * \text{ls}(d_i^j, f_i^j) \\ \sigma &\stackrel{\text{def}}{=} \bigstar_{j \in [N]} \text{node}_j * \bigstar_{\substack{i \in [N+1] \\ j \in [0, N]}} \text{order}_i^j \end{aligned}$$

A graphical illustration of the formulas node_j , order_i^j and order_i^0 is given in Figure 2, where list predicates are represented as dashed arrows and pointer predicates as full arrows. Consider the submodels of each of the formulas above. The intuition behind these formulas, in conjunction with the inequalities introduced below, is that there will be a model comprising a long concatenation, loosely speaking, of such submodels, if and only if there is a hamiltonian path in the given graph. The inequalities introduced below, will additionally ensure that such a long concatenation can happen only in the models of α and not of α' , and thus entailment will not hold in such a case. The variables a_j^k and b_j^k are essentially used in a way to count how long the concatenation is.

We now turn towards the pure parts of α and α' . For notational convenience, given $x \in \text{vars}(\alpha)$ and $S \subseteq \text{vars}(\alpha)$, subsequently $x \approx S$ abbreviates $\bigwedge_{y \in \text{vars}(\alpha) \setminus (S \cup \{x\})} x \neq y$. In other words, in any model \mathcal{I} of $x \approx S$, if $\ell^{\mathcal{I}}(x) = \ell^{\mathcal{I}}(y)$ for some $y \in \text{vars}(\alpha)$ then $y \in S$ or $x \equiv y$. We define Dvars to be the set $\{d_k^\ell : k \in [N+1], \ell \in [0, N]\}$, and we define φ to be the conjunction of the subsequent pure formulas:

$$\begin{aligned} s_i^0 &\approx \emptyset \wedge f_i^N \approx \text{Dvars} \wedge d_i^0 = f_i^0, & i \in [N+1] \\ s_i^j &\approx \bigcup_{\substack{v_p \in \text{pred}(v_i), \\ N-j < k \leq N-1}} \{a_p^k, b_p^k, b_p^N, e_p^f\} \cup \text{Dvars}, & i \in [N+1], j \in [N] \\ f_i^j &\approx \{e_i^s, a_i^0, b_i^{N-j}\} \cup \bigcup_{k \in [N-j-1]} \{a_i^k, b_i^k\} \cup \text{Dvars}, & i \in [N], j \in [0, N-1] \\ f_{N+1}^j &\approx \text{Dvars}, & j \in [0, N-1] \\ e_i^f &\neq e_j^f, & 1 \leq i \neq j \leq N \\ d_i^j &\neq b_i^{N-j}, & i \in [N], j \in [0, N-1] \\ \bigwedge_{k \in [N] \setminus \{i\}} d_i^j &\neq b_k^{N-j+1} & i \in [N+1], j \in [N] \end{aligned}$$

Finally, we define α and α' using the set of variables \mathbf{V} shown below:

$$\begin{aligned} \mathbf{V} &\stackrel{\text{def}}{=} \{a_i^0, a_i^j, b_i^j, b_i^N : i \in [N], j \in [N-1]\} \cup \{d_i^j : i \in [N+1], j \in [0, N]\} \\ \alpha &\stackrel{\text{def}}{=} \exists_{x \in \mathbf{V}} x. \sigma \wedge \varphi \quad \text{and} \quad \alpha' \stackrel{\text{def}}{=} \exists_{x \in \mathbf{V}} x. \sigma \wedge \varphi \wedge d_{N+1}^N \neq f_{N+1}^N \end{aligned}$$




Fig. 2: Graphical representation of the formulas node_i and order_i^j .

Note that φ includes $f_i^N \approx \text{Dvars}$ for all $i \in [N + 1]$. Given the additional disequality in α' , we now have $f_{N+1}^N \approx \text{Dvars} \setminus \{d_{N+1}^N\}$ in the pure part of α' .

Also note that in order to simplify the presentation, we have defined α and α' such that we use the same existentially quantified variables both in α and α' . It is important to note that given a model \mathcal{I} of both α and α' , the extension $\hat{\mathcal{I}}_1$ of \mathcal{I} that witnesses the satisfaction of the formula α and the extension $\hat{\mathcal{I}}_2$ that witnesses the satisfaction of the formula α' do not in general agree on the mapping of those existentially quantified variables. The existentially quantified variables in α could also be seen as fixed variables with names different from the existentially quantified variables of α' . As these variables act in the same way in both formulas, in order to avoid writing the above definitions twice and to simplify our proof, we have decided to treat them as existentially quantified and define them such that they have the same name in both formulas.

Lemma 9. $G = (V, E)$ and $v \in V$ is a valid instance of FVHP iff $\alpha \not\models \alpha'$.

Proof (sketch). First, a crucial observation is that for any \mathcal{I} such that $\mathcal{I} \models \alpha$ and $\mathcal{I} \not\models \alpha'$, in any extension $\hat{\mathcal{I}}$ witnessing $\mathcal{I} \models \alpha$, we have that the instantiation of d_{N+1}^N coincides with f_{N+1}^N . Suppose this was not the case, then $\hat{\mathcal{I}}$ would also witness $\mathcal{I} \models \alpha'$, contradicting our assumption. In this case we say that d_{N+1}^N is forced on f_{N+1}^N . We can show that forcing d_{N+1}^N on f_{N+1}^N is only possible if b_p^1 is forced on s_{N+1}^N for some unique predecessor $v_p \in \text{pred}(v_N)$ of v_N . In order to force b_p^1 on s_{N+1}^N , it can then be shown that d_p^{N-1} and therefore f_p^{N-1} is forced on a_p^0 . In summary, we can establish the following chain of inductive reasoning:

$$\begin{array}{ll}
\text{if } d_{i_0}^N \text{ is forced on } f_{i_0}^N & \text{then } b_{i_1}^1 \text{ is forced on } s_{i_0}^N \text{ for some } v_{i_1} \in \text{pred}(v_{i_0}) \\
\text{if } b_{i_1}^1 \text{ is forced on } s_{i_0}^N & \text{then } d_{i_1}^{N-1} \text{ is forced on } f_{i_1}^{N-1} \\
\vdots & \vdots \\
\text{if } d_{i_{N-1}}^N \text{ is forced on } f_{i_{N-1}}^N & \text{then } b_{i_N}^N \text{ is forced on } s_{i_0}^1 \text{ for some } v_{i_N} \in \text{pred}(v_{i_{N-1}}) \\
\text{if } b_{i_N}^N \text{ is forced on } s_{i_{N-1}}^1 & \text{then } d_{i_N}^0 \text{ is forced on } f_{i_N}^0
\end{array}$$

Now considering the variable names b_i^j in the implication chain, we obtain a sequence $b_{i_1}^1, b_{i_2}^2, \dots, b_{i_{N-1}}^{N-1}, b_{i_N}^N$ such that for all $1 \leq j < N$, v_{i_j} is a successor of $v_{i_{j+1}}$ in G . Consequently, the sequence of nodes $\pi = v_{i_N} \cdots v_{i_2} v_{i_1} v_{N+1}$ is a path of length $N + 1$ in G ending in v_{N+1} . Using the definition of φ , it follows that any b_i^j can only be “used” once, hence π does not contain duplicate nodes and thus is a Hamiltonian path ending in v_{N+1} . \square

It is readily checked that source and sink nodes in the graph corresponding to the definition of α and α' are fixed. Hence, we have established the following.

Theorem 10. *Entailment in the fixed endpoints fragment is CONP-hard.*

5 Conclusion

The results in this paper can be interpreted as follows: when considering fragments of Separation Logic which allow for existential quantification and unbounded data structures, having tractable, polynomial-time decision procedures will require severe syntactic restriction, since entailment is CONP-hard even in the strongly constrained fixed endpoints fragment. In the presence of general inductive predicates, although satisfiability is decidable [8], we have shown that entailment becomes undecidable. This result complements the decidability result obtained by Iosif *et al.* [17] and shows that the syntactic restrictions defined in [17] are not only natural but also crucial for decidability. However, we have shown that entailment in this fragment is EXPTIME-hard. On the more positive side, we have shown that entailment is “only” Π_2^P -complete in the presence of existential quantification, pointers and linked lists. Since this is a fragment that has been shown to be useful in program verifiers, this result may be seen as an argument in favour of supporting the development of decision procedures for domain-specific fragments of Separation Logic with a fixed set of predicates.

A number of problems remain open which we intend to investigate in the future. For instance, an open issue is whether a restriction to a one-selector fragment leads to decidable entailment. Also, although we have shown EXPTIME-hardness for the bounded-tree width fragment, we currently do not know whether this is a tight bound. This is also true for the fixed endpoints fragment and its CONP-hardness. Finally, of great interest would be decision procedures for entailment in these fragments, since most tools use incomplete heuristics.

Acknowledgments. We would like to thank the referees for their helpful comments. In particular, we wish to thank one referee who suggested the reduction from the inclusion problem for tree automata for the proof of Theorem 5.

References

1. M. Ajtai, R. Fagin, and L. Stockmeyer. The closure of monadic NP. *Journal of Computer and System Sciences*, 60(3):660–716, 2000.
2. J. Bengtson, J. Braband Jensen, and L. Birkedal. Charge! In L. Beringer and A. Felty, editors, *Interactive Theorem Proving*, volume 7406 of *LNCS*, pages 315–331. Springer, 2012.
3. J. Berdine, C. Calcagno, B. Cook, D. Distefano, P. W. O’Hearn, T. Wies, and H. Yang. Shape analysis for composite data structures. In W. Damm and H. Hermanns, editors, *Computer Aided Verification*, volume 4590 of *LNCS*, pages 178–192. Springer, 2007.
4. J. Berdine, C. Calcagno, and P. O’Hearn. A decidable fragment of separation logic. In *Foundations of Software Technology and Theoretical Computer Science*, volume 3328 of *LNCS*, pages 110–117. Springer, 2005.

5. J. Berdine, B. Cook, and S. Ishtiaq. SLAyer: Memory safety for systems-level code. In G. Gopalakrishnan and S. Qadeer, editors, *Computer Aided Verification*, volume 6806 of *LNCS*, pages 178–183. Springer, 2011.
6. L. Birkedal, N. Torp-Smith, and J. C. Reynolds. Local reasoning about a copying garbage collector. In *Principles of Programming Languages*, pages 220–231, New York, NY, USA, 2004. ACM.
7. A. Bouajjani, C. Drăgoi, C. Enea, and M. Sighireanu. Accurate invariant checking for programs manipulating lists and arrays with infinite data. In *Automated Technology for Verification and Analysis*, LNCS, pages 167–182. Springer, 2012.
8. J. Brotherston, C. Fuhs, N. Gorogiannis, and J. Navarro Pérez. A decision procedure for satisfiability in separation logic with inductive predicates. Technical Report RN/13/15, University College London, 2013.
9. J. Brotherston, N. Gorogiannis, and R.L. Petersen. A generic cyclic theorem prover. In *Asian Symposium on Programming Languages and Systems*, pages 350–367. Springer, 2012.
10. J. Brotherston and M. Kanovich. Undecidability of propositional separation logic and its neighbours. In *Logic in Computer Science*, pages 137–146. IEEE Computer Society, 2010.
11. C. Calcagno, D. Distefano, P. W. O’Hearn, and H. Yang. Beyond reachability: Shape abstraction in the presence of pointer arithmetic. In K. Yi, editor, *Static Analysis*, volume 4134 of *LNCS*, pages 182–203. Springer, 2006.
12. C. Calcagno, H. Yang, and P.W. O’Hearn. Computability and complexity results for a spatial assertion language for data structures. In *Asian Symposium on Programming Languages and Systems*, pages 289–300, 2001.
13. W.-N. Chin, C. David, H. H. Nguyen, and S. Qin. Automated verification of shape, size and bag properties via user-defined predicates in separation logic. *Science of Computer Programming*, 77(9):1006 – 1036, 2012.
14. B. Cook, C. Haase, J. Ouaknine, M. Parkinson, and J. Worrell. Tractable reasoning in a fragment of separation logic. In *Concurrency Theory*, volume 6901 of *LNCS*, pages 235–249. Springer, 2011.
15. N. Gorogiannis, M. Kanovich, and P. O’Hearn. The complexity of abduction for separated heap abstractions. In *Static Analysis*, volume 6887 of *LNCS*, pages 25–42. Springer, 2011.
16. P. Habermehl, L. Holík, A. Rogalewicz, J. Šimáček, and T. Vojnar. Forest automata for verification of heap manipulation. In G. Gopalakrishnan and S. Qadeer, editors, *Computer Aided Verification*, volume 6806 of *LNCS*, pages 424–440. Springer, 2011.
17. R. Iosif, A. Rogalewicz, and J. Šimáček. The tree width of separation logic with recursive definitions. In M. P. Bonacina, editor, *Automated Deduction – CADE*, volume 7898 of *LNCS*, pages 21–38. Springer, 2013.
18. S. Ishtiaq and P. O’Hearn. BI as an assertion language for mutable data structures. In *Principles of Programming Languages*, pages 14–26. ACM, 2001.
19. E. L. Post. A variant of a recursively unsolvable problem. *Bulletin of the American Mathematical Society*, 52(4):264–268, 1946.
20. J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Logic in Computer Science*. IEEE Computer Society, 2002.
21. H. Seidl. Deciding equivalence of finite tree automata. *SIAM Journal on Computing*, 19(3):424–437, June 1990.
22. H. Yang. *Local Reasoning for Stateful Programs*. PhD thesis, University of Illinois at Urbana-Champaign, 2001. (Technical Report UIUCDCS-R-2001-2227).

A Missing Proof Details from Section 4

A.1 Entailment in the General Case

We first formally prove Π_2^P -hardness. Recall the definitions of α and α' from the main text. Let $G = (V, E)$ be graph $\alpha \models \alpha'$ whose vertex set V is decomposed into disjoint sets $V' = \{v_1, \dots, v_n\}$ of nodes with degree greater than one and $V'' = \{v_{n+1}, \dots, v_m\}$ of nodes with degree equal to one. We define

$$\alpha \stackrel{\text{def}}{=} \bigstar_{i \in [3]} \bigstar_{j \in [n]} \text{pt}(x_{i,j}, y_i) * \bigstar_{n < i \leq m} \text{pt}(x_i, z_i) \wedge \bigwedge_{n < i \leq m} \bigvee_{j \in [3]} z_i = y_j \wedge \bigwedge_{1 \leq i \neq j \leq 3} y_i \neq y_j$$

$$\alpha' \stackrel{\text{def}}{=} \exists_{i \in [n]} \mathbf{a}_i. \exists_{i \in [m]} \mathbf{b}_i. \bigstar_{i \in [n]} \text{pt}(\mathbf{a}_i, \mathbf{b}_i) * \bigstar_{n < j \leq m} \text{pt}(x_j, \mathbf{b}_j) * \text{tt} \wedge \bigwedge_{(v_i, v_j) \in E} \mathbf{b}_i \neq \mathbf{b}_j.$$

Subsequently, the spatial part of α' is called σ' and the pure part φ' .

Lemma 11. *G is a valid instance of 2-ROUND 3-COLORABILITY iff $\alpha \models \alpha'$.*

Proof. Suppose $G = (V, E)$ is a valid instance of 2-ROUND 3-COLORABILITY and let \mathcal{I} be such that $\mathcal{I} \models \alpha$. Define $f : V'' \rightarrow \{1, 2, 3\}$ to be a leaf coloring such that for all $n < i \leq m$, $f(i) \stackrel{\text{def}}{=} j$ iff $\ell^{\mathcal{I}}(z_i) = y_j$. By assumption, f can be extended to a three coloring \hat{f} , so in particular $\hat{f}(v_i) \neq \hat{f}(v_j)$ for all $\{v_i, v_j\} \in E$. Let $\hat{\mathcal{I}}$ be an extension of \mathcal{I} such that additionally $\ell^{\hat{\mathcal{I}}}(\mathbf{b}_i) = \ell^{\mathcal{I}}(y_{\hat{f}(v_i)})$ for all $i \in [m]$ and $\ell^{\hat{\mathcal{I}}}(\mathbf{a}_i) = \ell^{\mathcal{I}}(x_{\hat{f}(v_i), i})$ for all $i \in [n]$; here and in the next paragraph, for brevity we abuse notation and treat the \mathbf{a}_i and \mathbf{b}_i in interpretation as if they come from FVars. It is not difficult to verify that $\hat{\mathcal{I}} \models \varphi' \wedge \sigma'$ and hence $\mathcal{I} \models \alpha'$.

Conversely, suppose 2-ROUND 3-COLORABILITY does not hold for $G = (V, E)$ and let $f : V'' \rightarrow \{1, 2, 3\}$ be the coloring of the leaves such that f cannot be extended to a three-coloring. Let \mathcal{I} be such that $\mathcal{I} \models \alpha \wedge \bigwedge_{v_i \in V''} z_i = y_{f(i)}$ and hence $\mathcal{I} \models \alpha$. We claim $\mathcal{I} \not\models \alpha'$. To the contrary, assume there is an extension $\hat{\mathcal{I}}$ such that $\hat{\mathcal{I}} \models \sigma' \wedge \varphi'$. Since $\ell^{\hat{\mathcal{I}}}(\mathbf{b}_i) \neq \ell^{\hat{\mathcal{I}}}(\mathbf{b}_j)$ for all $\{v_i, v_j\} \in E$, defining \hat{f} such that $\hat{f}(v_i) = c$ iff $\ell^{\hat{\mathcal{I}}}(\mathbf{b}_i) = y_c$ yields an extension of f to a three coloring, contradicting our assumption. \square

We now prove the Π_2^P upper bound. Since we are only dealing with pointer predicates with arity one, we can view interpretations $\mathcal{I} = (V^{\mathcal{I}}, E^{\mathcal{I}}, \ell^{\mathcal{I}}, s^{\mathcal{I}})$ as $\mathcal{I} = (V^{\mathcal{I}}, E^{\mathcal{I}}, \ell^{\mathcal{I}})$, where $E^{\mathcal{I}} \subseteq V^{\mathcal{I}} \times V^{\mathcal{I}}$ is a functional relation. This allows us to give a direct semantics definition of the ls predicate defined in Equation (1) as follows: we have $\mathcal{I} \models \text{ls}(x, y)$ iff either

- $E^{\mathcal{I}} = \emptyset$ and $\ell^{\mathcal{I}}(x) = \ell^{\mathcal{I}}(y)$; or
- $V^{\mathcal{I}} \supseteq \{v_1, \dots, v_m\}$, $E^{\mathcal{I}} = \{(v_i, v_{i+1}) : 1 \leq i < m\}$, $\ell^{\mathcal{I}}(x) = v_1 \neq v_m = \ell^{\mathcal{I}}(y)$ and $v_i \neq v_j$ for all $1 \leq i \neq j \leq m$.

In other words, either \mathcal{I} has no edges (*i.e.* allocated variables) at all and x and y coincide, or there is a non-empty chain of pointers from x to y with no duplicates.

Before giving the proof of Lemma 7, we provide three further definitions. A path π in \mathcal{I} of length n is a sequence $\pi = v_1 \cdots v_{n+1}$ of nodes such that $(v_i, v_{i+1}) \in E$ for all $1 \leq i \leq n$. The set $\text{nodes}(\pi)$ of nodes along π is $\text{nodes}(\pi) \stackrel{\text{def}}{=} \{v_1, \dots, v_{n+1}\}$. Given an interpretation \mathcal{I} and $v \in V^{\mathcal{I}}$, we define $\text{vars}^{\mathcal{I}}(v) \stackrel{\text{def}}{=} \{x \in \text{Vars} : \ell^{\mathcal{I}}(x) = v\}$.

Lemma 12 (Lemma 7 in the main text). *Let α, α' be SL-assertions such that $\alpha \not\models \alpha'$, and let $n = |\text{vars}(\alpha)| + |\text{vars}(\alpha')|$. Then there exists an \mathcal{I} witnessing $\alpha \not\models \alpha'$ with $|V^{\mathcal{I}}| \in O(n^2)$.*

Proof. Let $\alpha = \sigma \wedge \varphi$ and $\alpha' = \exists a_1 \dots a_r. \varphi' \wedge \sigma'$. Furthermore, let \mathcal{J} be a counter model witnessing $\alpha \not\models \alpha'$, so in particular for any extension $\hat{\mathcal{J}}$ of \mathcal{J} , $\hat{\mathcal{J}} \not\models \sigma' \wedge \varphi'$, and furthermore, let \mathcal{J} be such an interpretation with the least number of nodes. Suppose for contradiction that $|V^{\mathcal{J}}| > n^2$, then in \mathcal{J} there exists a v - w path $\pi = vv_1 \dots v_m w$ such that $m \geq n$, $\text{vars}^{\mathcal{J}}(v_i) = \emptyset$ for all $1 \leq i \leq m$, $\text{vars}^{\mathcal{J}}(v) \neq \emptyset$, and $\text{vars}^{\mathcal{J}}(w) \neq \emptyset$. The latter fact implies that each v_i has exactly one incoming from v_{i-1} , where $v_0 = v$, since otherwise $\text{vars}^{\mathcal{I}}(v_i)$ would not be empty for some $i \in [m]$, c.f. the semantics of the ls predicate above.

Let \mathcal{I} be obtained from \mathcal{J} by removing v_m from V_a and the in- and outgoing edges connected to it, and by adding an edge from v_{m-1} to w . We claim that \mathcal{I} witnesses $\alpha \not\models \alpha'$. To the contrary, assume $\mathcal{I} \models \alpha'$ and let $\hat{\mathcal{I}}$ be the extension of \mathcal{I} such that $\hat{\mathcal{I}} \models \varphi' \wedge \sigma'$. We are going to show how to obtain from $\hat{\mathcal{I}}$ an extension $\hat{\mathcal{J}}$ of \mathcal{J} such that $\hat{\mathcal{J}} \models \varphi' \wedge \sigma'$, which contradicts our assumption. Let $\pi' = vv_1 \dots v_{m-1}w$ be the shortened v - w path in $E^{\hat{\mathcal{I}}}$ of length m . Since π' consists of $m+1 \geq n+1$ nodes, and since $\text{vars}(v) \neq \emptyset$ and $\text{vars}(w) \neq \emptyset$, we have $\text{vars}(v'_j) = \emptyset$ for some $1 \leq j \leq m-1$. In the following, let j be the smallest such number. Let $\hat{\mathcal{J}}$ be the extension of \mathcal{J} such that

- $\text{vars}^{\hat{\mathcal{J}}}(v) = \text{vars}^{\hat{\mathcal{I}}}(v)$ for all $v \notin \text{nodes}^{\mathcal{J}}(\pi)$;
- $\text{vars}^{\hat{\mathcal{J}}}(v) = \text{vars}^{\hat{\mathcal{I}}}(v)$ and $\text{vars}^{\hat{\mathcal{J}}}(w) = \text{vars}^{\hat{\mathcal{I}}}(w)$;
- $\text{vars}^{\hat{\mathcal{J}}}(v_i) = \text{vars}^{\hat{\mathcal{I}}}(v'_i)$ for all $1 \leq i < j$;
- $\text{vars}^{\hat{\mathcal{J}}}(v_j) = \emptyset$; and
- $\text{vars}^{\hat{\mathcal{J}}}(v_i) = \text{vars}^{\hat{\mathcal{I}}}(v_{i-1})$ for all $j < i < m$.

We have $\hat{\mathcal{J}} \models \varphi'$, since for all $x, y \in \text{vars}(\alpha')$, $\ell^{\hat{\mathcal{J}}}(x) = \ell^{\hat{\mathcal{J}}}(y)$ iff $\ell^{\hat{\mathcal{I}}}(x) = \ell^{\hat{\mathcal{I}}}(y)$. Moreover, it is easily checked that for any σ'' , $\hat{\mathcal{J}} \models \text{pt}(x, y) * \sigma''$ iff $\hat{\mathcal{I}} \models \text{pt}(x, y) * \sigma''$, and $\hat{\mathcal{J}} \models \text{ls}(x, y) * \sigma''$ iff $\hat{\mathcal{I}} \models \text{ls}(x, y) * \sigma''$. Consequently, $\hat{\mathcal{J}} \models \varphi' \wedge \sigma'$, which is a contradiction. Therefore, \mathcal{I} is an interpretation witnessing that $\alpha \not\models \alpha'$, which is a contradiction to \mathcal{J} being the smallest such interpretation. \square

A.2 Entailment in the Fixed Endpoints Fragment

We begin by introducing some terminology. Let P be an SL-interpretation, and recall that since we are only considering interpretations in which all nodes have arity at most one they can be viewed as partial functional graphs. We call P a




Fig. 3: The connected weave $(P_0, P_1, P_2, P_3, P_4)$.

path SL-interpretation with endpoints (s, t) if there is $n > 0$ such that $E(P) = \{(v_i, v_{i+1}) : i \in [n-1], s = v_1, t = v_n\}$. Let (P_0, \dots, P_n) be an ordered sequence of path SL-interpretations with corresponding endpoints

$$((s_0, f_0), \dots, (s_n, f_n)).$$

We say that the sequence is a *connected weave* if it is an interpretation $P_0 * \dots * P_n$ and there are nodes v_1, \dots, v_n in V^{P_1}, \dots, V^{P_n} respectively, such that for all $i \in [n]$, $f_{i-1} = v_i$. In other words, a connected weave is of the form shown in Figure 3.

Furthermore, if α_1 is a subformula of some formula α , and \mathcal{I} is a model of α , we denote with $\mathcal{I}(\alpha_1)$ the submodel of \mathcal{I} corresponding to α_1 . Let \mathcal{I} be a model of an SL-assertion α , and let $\mathbf{x} \in \text{Vars}(\alpha)$. We denote by $\text{mapped}^{\mathcal{I}}(\mathbf{x})$ the set of all $v \in V^{\mathcal{I}}$ for which there is an extension $\hat{\mathcal{I}}$ of \mathcal{I} , with $\hat{\mathcal{I}} \models \alpha$, such that $\hat{\ell}^{\hat{\mathcal{I}}}(\mathbf{x}) = v$. For a node $v \in V^{\mathcal{I}}$, we say that \mathbf{x} is *forced on v in \mathcal{I}* if $\text{mapped}^{\mathcal{I}}(\mathbf{x}) = \{v\}$. For what follows, we identify the instantiation of an existentially quantified variable by the same name.

Below, we reduce the FIXED VERTEX HAMILTONIAN PATH problem over directed graphs to the problem of deciding whether an SL formula α entails an SL formula α' .

Before we begin proving the construction from the main text correct, let us recall the problem we are reducing from.

FIXED VERTEX HAMILTONIAN PATH (FVHP)

INPUT: A directed graph $G = (V, E)$ and $v \in V$.

QUESTION: Does there exist a Hamiltonian path in G ending in v ?




Fig. 4: $(\mathcal{I}(\text{order}_{i_N}^0), \mathcal{I}(\text{node}_{i_N}), \mathcal{I}(\text{order}_{i_{N-1}}^1), \dots, \mathcal{I}(\text{order}_{i_1}^{N-1}), \mathcal{I}(\text{node}_{i_1}), \mathcal{I}(\text{order}_{i_0}^N))$.

Lemma 13. *Given $G = (V, E)$ and $v \in V$, FVHP is NP-complete. Hardness even holds if G has no self-loops.*

Proof. Membership in NP is trivial. For hardness we reduce from the well-known NP-complete HAMILTONIAN PATH PROBLEM. Let $G' = (V', E')$ be an instance of this problem such that $|V'| = n$. We define $G = (V, E)$ as follows:

- $V \stackrel{\text{def}}{=} V' \cup \{v^*\}$ for some $v^* \notin V'$, and
- $E \stackrel{\text{def}}{=} (E' \setminus \{(u, u) : u \in V'\}) \cup \{(u, v^*) : u \in V'\}$.

We claim that G' is a valid instance of the HAMILTONIAN PATH PROBLEM if, and only if, G and v^* is an instance of FHVP. Suppose there is a Hamiltonian path $\pi' = v_1 \cdots v_n$ in G' . There is no (v_i, v_{i+1}) occurring along π' such that $v_i = v_{i+1}$, hence π' is a path in G , and $\pi \stackrel{\text{def}}{=} v_1 \cdots v_n v^*$ is a Hamiltonian path in G ending in v^* . Conversely, if $\pi = v_1 \cdots v_n v^*$ is a Hamiltonian path in G ending in v^* , then $\pi' \stackrel{\text{def}}{=} v_1 \cdots v_n$ is a path in G' visiting every vertex once, *i.e.*, a Hamiltonian path. \square

The following lemma now establishes CONP-hardness of entailment in the fixed endnodes fragment of Separation Logic.

Lemma 14. $G = (V, E)$ and $v \in V$ is a valid instance of FVHP if, and only if, $\alpha \not\equiv \alpha'$.

Proof. Let $G = (V, E)$ be a directed graph with no self loops, and let $|V| = N+1$ and with no loss of generality $v = v_{N+1}$. Let then α and α' be defined as it is described in Section 4 given the graph G and the vertex $v = v_{N+1}$.

Suppose first that there is a model \mathcal{I} of α that is not a model of α' , which by definition implies that \mathbf{d}_{N+1}^N is forced on $\ell(\mathbf{f}_{N+1}^N)$. We proceed by showing that for any $j \in [0, N]$ and $i \in [N+1]$, if in some model \mathcal{I} the variable \mathbf{d}_i^j is forced on $\ell(\mathbf{f}_i^j)$ then \mathcal{I} there exists a sequence of path graphs

$$(\mathcal{I}(\text{order}_{\pi(1)}^0), \mathcal{I}(\text{node}_{\pi(1)}), \mathcal{I}(\text{order}_{\pi(2)}^1), \dots, \\ \mathcal{I}(\text{order}_{\pi(j)}^{j-1}), \mathcal{I}(\text{node}_{\pi(j)}), \mathcal{I}(\text{order}_{\pi(j+1)}^j))$$

which is a connected weave in the model \mathcal{I} , and where π is a mapping from $[N+1]$ to $[N+1]$, where

- $\pi(j+1) = i$,
- for all $1 \leq i < k \leq j$, $\pi(i) \neq \pi(k)$
- for all $\ell \in [j]$, $v_{\pi(\ell)} \in \text{pred}(v_{\pi(\ell+1)})$.

Such a weave is shown in Figure 4 for $j = N$, where for all $k \in [0, N]$, $\pi(k+1) = i_{N-k}$. We proceed by induction on $j \in [0, N]$. For the base case, let $j = 0$. Then, for all $i \in [N+1]$, in the subgraph $\mathcal{I}(\text{order}_i^0)$, the variable \mathbf{d}_i^0 is forced on $\ell(\mathbf{f}_i^0)$, and therefore, for any mapping π , the variable $\mathbf{d}_{\pi(1)}^0$ is forced on $\ell(\mathbf{f}_{\pi(1)}^0)$, as required. The function π that maps 1 to i satisfies the conditions stated.

For the induction step, suppose the statement holds for all $j < J$ for some $J \in [N]$, and consider the case for J . We will show that

- if $\mathbf{d}_{i_1}^J$ is forced on $\mathbf{f}_{i_1}^J$ for some i_1 , then there is i_2 such that $\mathbf{b}_{i_2}^{N-J+1}$ is forced on $\mathbf{s}_{i_2}^J$, where $v_{i_2} \in \text{pred}(v_{i_1})$, and
- if $\mathbf{b}_{i_2}^{N-J+1}$ is forced on $\mathbf{s}_{i_2}^J$, then $\mathbf{d}_{i_2}^{J-1}$ is forced on $\mathbf{f}_{i_2}^{J-1}$.

We can then use the inductive hypothesis to construct the mapping π .

Suppose then that \mathcal{I} is a model of α such that there is an $i_1 \in [N+1]$, such that in $\mathcal{I}(\text{order}_{i_1}^J)$ the variable $\mathbf{d}_{i_1}^J$ is forced on $\ell(\mathbf{f}_{i_1}^J)$.

Claim. There is i_2 such that $\mathbf{b}_{i_2}^{N-J+1}$ is forced on $\mathbf{s}_{i_2}^J$, where $v_{i_2} \in \text{pred}(v_{i_1})$

For every other node of $\mathcal{I}(\text{order}_{i_1}^J)$ other than $\ell(\mathbf{f}_{i_1}^J)$, and every extension $\hat{\mathcal{I}}$ of \mathcal{I} , either $\mathbf{b}_{i_1}^{N-J}$ or \mathbf{b}_r^{N-J+1} for $r \neq i_1$ is equal to that node on that extension. If there was a node v in $\mathcal{I}(\text{order}_{i_1}^J)$ other than $\ell(\mathbf{f}_{i_1}^J)$, and an extension $\hat{\mathcal{I}}$ where v is labelled with neither $\mathbf{b}_{i_1}^{N-J}$ nor \mathbf{b}_r^{N-J+1} , for $r \neq i_1$, then $v \in \text{mapped}^{\mathcal{I}}(\mathbf{d}_{i_1}^J)$, where $v \neq \ell(\mathbf{f}_{i_1}^J)$, and hence $|\text{mapped}^{\mathcal{I}}(\mathbf{d}_{i_1}^J)| > 1$, which is a contradiction. In particular, this is also the case for the node $\ell(\mathbf{s}_{i_1}^J)$ in every extension $\hat{\mathcal{I}}$ of \mathcal{I} with

$\hat{\mathcal{I}} \models \alpha$. First, we argue that since $\ell(\mathbf{s}_{i_1}^J)$ has an outgoing edge belonging to the submodel $\mathcal{I}(\text{order}_{i_1}^J)$, it holds that for any variable \mathbf{b}_p^k , if \mathbf{b}_p^k is equal to $\ell(\mathbf{s}_{i_1}^J)$ in some extension $\hat{\mathcal{I}}$ of \mathcal{I} , then \mathbf{e}_p^f will also have to be mapped to $\ell(\mathbf{s}_{i_1}^J)$ in that extension, as otherwise, $\ell(\mathbf{s}_{i_1}^J)$ would have two outgoing edges, one belonging to the submodel $\mathcal{I}(\text{order}_{i_1}^J)$, and one belonging to the submodel $\mathcal{I}(\text{node}_p)$.

But then, by definition of φ and σ , and the equalities and inequalities these formulas impose on $\mathbf{s}_{i_1}^J$, $\mathbf{s}_{i_1}^J$ is only allowed to be equal to a variable \mathbf{e}_r^f when $v_r \in \text{pred}(v_{i_1})$. Since G has no self loops, the variable $\mathbf{b}_{i_1}^{N-J}$ cannot be mapped to $\ell(\mathbf{s}_{i_1}^J)$. Therefore, only the variables \mathbf{b}_r^{N-J+1} , for $r \neq i_1$ where $v_r \in \text{pred}(v_{i_1})$ can be mapped to $\ell(\mathbf{s}_{i_1}^J)$. Furthermore, if there are extensions $\hat{\mathcal{I}}_1$ and $\hat{\mathcal{I}}_2$ of \mathcal{I} such that the variable $\mathbf{b}_{k_1}^{N-J+1}$ is mapped to $\ell(\mathbf{s}_{i_1}^J)$ in $\hat{\mathcal{I}}_1$ and the variable $\mathbf{b}_{k_2}^{N-J+1}$ is mapped to $\ell(\mathbf{s}_{i_1}^J)$ in $\hat{\mathcal{I}}_2$, with $k_1 \neq k_2$, then, since $\mathbf{e}_{k_1}^f$ and $\mathbf{e}_{k_2}^f$ are fixed variables, both of them have to be mapped to $\ell(\mathbf{s}_{i_1}^J)$, in \mathcal{I} , but this would contradict the set of inequalities $\mathbf{e}_{k_1}^f \neq \mathbf{e}_{k_2}^f$ for distinct k_1, k_2 . Therefore, there is a single $i_2 \neq i_1$, with $v_{i_2} \in \text{pred}(v_{i_1})$, such that in all extensions $\hat{\mathcal{I}}$ of \mathcal{I} , the variable $\mathbf{b}_{i_2}^{N-J+1}$ is equal to $\ell(\mathbf{s}_{i_1}^J)$, and therefore, $\mathbf{b}_{i_2}^{N-J+1}$ is forced on $\ell(\mathbf{s}_{i_1}^J)$.

Claim. The variable $\mathbf{d}_{i_2}^{J-1}$ is forced on $\ell(\mathbf{f}_{i_2}^{J-1})$.

Again, since $\mathbf{s}_{i_1}^J$ has an outgoing edge that belongs to the submodel $\mathcal{I}(\text{order}_{i_1}^J)$, it has to be the case that all variables $\mathbf{a}_{i_2}^k$, and all variables $\mathbf{b}_{i_2}^k$, for $k \geq N - J + 1$ are forced on $\ell(\mathbf{e}_{i_2}^f)$ in $\mathcal{I}(\text{node}_{i_2})$, and also forced on $\ell(\mathbf{s}_{i_1}^J)$. But by definition, $\mathbf{s}_{i_1}^J$, allows exactly those variables from the subformula node_{i_2} , and no others, to be equal to it, and furthermore, it allows such variables only for i_2 such that $v_{i_2} \in \text{pred}(v_{i_1})$. Therefore, for all $k < N - J + 1$ (and most importantly for $k = N - J$), the variable $\mathbf{a}_{i_2}^k$ is forced to be *not* equal to the node $\ell(\mathbf{e}_{i_2}^f) = \ell(\mathbf{s}_{i_1}^J)$.

Notice that the length of the submodel $\mathcal{I}(\text{node}_{i_2})$ is at least 1. For each extension $\hat{\mathcal{I}}$ of \mathcal{I} , let then \hat{v} be the node of the submodel $\mathcal{I}(\text{node}_{i_2})$, on which the variable $\mathbf{a}_{i_2}^{N-J}$ is mapped, and notice that from what is shown above, this node is different from the one on which the variable $\mathbf{e}_{i_2}^f$ is mapped. Therefore \hat{v} has an outgoing edge that belongs to the submodel $\mathcal{I}(\text{node}_{i_2})$. Notice that the only variables that cannot be equal to $\mathbf{b}_{i_2}^{N-J+1}$ are either some variables \mathbf{d}_u^p or some variables \mathbf{f}_u^p or some variables \mathbf{s}_u^p , for appropriate p and u . Since \hat{v} has an outgoing edge belonging to the submodel $\mathcal{I}(\text{node}_{i_2})$, no variable \mathbf{s}_u^p can be mapped to \hat{v} , as there would then be two outgoing edges from \hat{v} , one belonging to $\mathcal{I}(\text{node}_{i_2})$ and one to $\mathcal{I}(\text{order}_u^p)$. Furthermore, there is no variable \mathbf{f}_u^p that is allowed to be equal to $\mathbf{a}_{i_2}^{N-J}$ and not $\mathbf{b}_{i_2}^{N-J+1}$, and therefore a variable \mathbf{d}_u^p for some p and u is equal to \hat{v} in $\hat{\mathcal{I}}$.

Furthermore, notice that the variable \mathbf{d}_u^p is such that either $p = N - (N - J + 1) = J - 1$ and $u = i_2$, or $p = N - (N - J + 1) + 1 = J$ and $u \neq i_2$. Since in every extension $\hat{\mathcal{I}}$ of \mathcal{I} , the node \hat{v} has an outgoing edge that belongs to the submodel $\mathcal{I}(\text{node}_{i_2})$, it holds that for all k and q , if $\mathbf{d}_{i_2}^k$ is equal to $\hat{\ell}(\mathbf{a}_{i_2}^{N-J})$ then \mathbf{f}_q^k will also be equal to $\hat{\ell}(\mathbf{a}_{i_2}^{N-J})$. Otherwise, \hat{v} would have two outgoing edges, one belonging to the submodel $\mathcal{I}(\text{node}_{i_2})$ and one to the submodel $\mathcal{I}(\text{order}_q^k)$.

By definition φ and σ and the inequality constraints they impose on \mathbf{f}_q^k , for all k and q , all variables \mathbf{a}_t^p with $t \neq q$ are not allowed to be equal to \mathbf{f}_q^k , and hence $u = i_2$. But then, according to the above, $p = J - 1$, and therefore only one of the variables mentioned, namely the variable $\mathbf{d}_{i_2}^{J-1}$ can be equal to $\hat{\ell}(\mathbf{a}_{i_2}^{N-J}) = \hat{v}$ for any extension $\hat{\mathcal{I}}$. Hence for all extensions $\hat{\mathcal{I}}$ of \mathcal{I} , $\mathbf{d}_{i_2}^{J-1}$ is equal to $\ell(\mathbf{f}_{i_2}^{J-1})$ in $\mathcal{I}(\text{order}_{i_2}^{J-1})$, and therefore $\mathbf{d}_{i_2}^{J-1}$ is forced on $\ell(\mathbf{f}_{i_2}^{J-1})$.

By the inductive hypothesis then, there is a weave

$$(\mathcal{I}(\text{order}_{\pi(1)}^0), \mathcal{I}(\text{node}_{\pi(1)}), \mathcal{I}(\text{order}_{\pi(2)}^1), \dots, \\ \mathcal{I}(\text{order}_{\pi(J-1)}^{J-2}), \mathcal{I}(\text{node}_{\pi(J-1)}), \mathcal{I}(\text{order}_{\pi(J)}^{J-1})),$$

with $\pi(J) = i_2$ such that for all $k \in [J - 1]$, $v_{\pi(k)} \in \text{pred}(v_{\pi(k+1)})$. Hence, according to the above, there is a larger weave containing the two submodels $\mathcal{I}(\text{node}_{i_2})$ and $\mathcal{I}(\text{order}_{i_1}^J)$, namely the sequence

$$(\mathcal{I}(\text{order}_{\pi(1)}^0), \mathcal{I}(\text{node}_{\pi(1)}), \mathcal{I}(\text{order}_{\pi(2)}^1), \dots, \\ \mathcal{I}(\text{order}_{\pi(J)}^{J-1}), \mathcal{I}(\text{node}_{\pi(J)}), \mathcal{I}(\text{order}_{\pi(J+1)}^J)),$$

where $\pi(J+1) = i_1$, and where additionally, $v_{\pi(J)} \in \text{pred}(v_{\pi(J+1)})$, since $\pi(J) = i_2$, $\pi(J+1) = i_1$ and $v_{i_2} \in \text{pred}(v_{i_1})$.

Notice that all submodels $\mathcal{I}(\text{order}_i^j)$ in the weave are distinct, and furthermore, none of the $\mathcal{I}(\text{node}_i)$ can be used twice, as all of them have their sink attached to a non-sink node of a non-empty submodel $\mathcal{I}(\text{order}_i^j)$, distinct from all other ones. Therefore, $\pi(J) \neq \pi(k)$, for all $k \in [J - 1]$, and therefore π satisfies the conditions above.

Hence, if \mathbf{d}_{N+1}^N is forced on $\ell(\mathbf{f}_{N+1}^N)$, there is a connected weave

$$(\mathcal{I}(\text{order}_{\pi(1)}^0), \mathcal{I}(\text{node}_{\pi(1)}), \mathcal{I}(\text{order}_{\pi(2)}^1), \dots, \\ \mathcal{I}(\text{order}_{\pi(N)}^{N-1}), \mathcal{I}(\text{node}_{\pi(N)}), \mathcal{I}(\text{order}_{\pi(N+1)}^N)),$$

in the model \mathcal{I} , where for all $k \in [N]$, $v_{\pi(k)} \in \text{pred}(v_{\pi(k+1)})$, and $\pi(N+1) = N+1$. Since the formulas node_i are not defined for $i = N+1$, $\pi(N+1) \neq \pi(k)$, for all $k \in [N]$. From the conditions of π defined above, it follows that $(v_{\pi(1)}, \dots, v_{\pi(N+1)})$ is a hamiltonian path in G ending in v_{N+1} .

For the other direction, suppose there is a hamiltonian path $p_1 = (v_{i_1}, \dots, v_{i_N}, v_{i_{N+1}})$ in G with v_{N+1} the last vertex, and define π to be the mapping, where for all $k \in [N]$, $\pi(k) = v_{i_k}$, and $\pi(N+1) = N+1$, and let W be the weave

$$(\mathcal{I}(\text{order}_{\pi(1)}^0), \mathcal{I}(\text{node}_{\pi(1)}), \dots, \mathcal{I}(\text{order}_{\pi(N)}^{N-1}), \mathcal{I}(\text{node}_{\pi(N)}), \mathcal{I}(\text{order}_{\pi(N+1)}^N)),$$

where each submodel in the weave is a single edge. Let all other submodels, not in the weave stand disconnected from each other and the weave, and let this be the model \mathcal{I} . Then by the arguments presented for the previous direction $\alpha \models \mathcal{I}$

and $\alpha' \not\models \mathcal{I}$. It follows, that there is a model \mathcal{I} of α and not α' if and only if there is a fixed vertex hamiltonian path in G , and hence checking whether $\alpha \models \alpha'$ is coNP-hard, since finding a hamiltonian path is NP-hard for directed graphs with no self loops. \square