



HAL
open science

The Effects of Adding Reachability Predicates in Propositional Separation Logic

Stéphane Demri, Etienne Lozes, Alessio Mansutti

► **To cite this version:**

Stéphane Demri, Etienne Lozes, Alessio Mansutti. The Effects of Adding Reachability Predicates in Propositional Separation Logic. 22nd International Conference on Foundations of Software Science and Computation Structures FOSSACS, 2018, Thessaloniki, Greece. hal-01920563

HAL Id: hal-01920563

<https://hal.science/hal-01920563v1>

Submitted on 13 Nov 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Effects of Adding Reachability Predicates in Propositional Separation Logic

S. Demri¹, E. Lozes², and A. Mansutti¹

¹LSV, CNRS, ENS Paris-Saclay, Université Paris-Saclay, France

²I3S, Université Côte d’Azur, France

Abstract. The list segment predicate `ls` used in separation logic for verifying programs with pointers is well-suited to express properties on singly-linked lists. We study the effects of adding `ls` to the full propositional separation logic with the separating conjunction and implication, which is motivated by the recent design of new fragments in which all these ingredients are used indifferently and verification tools start to handle the magic wand connective. This is a very natural extension that has not been studied so far. We show that the restriction without the separating implication can be solved in polynomial space by using an appropriate abstraction for memory states whereas the full extension is shown undecidable by reduction from first-order separation logic. Many variants of the logic and fragments are also investigated from the computational point of view when `ls` is added, providing numerous results about adding reachability predicates to propositional separation logic.

1 Introduction

Separation logic [20,25,28] is a well-known assertion logic for reasoning about programs with dynamic data structures. Since the implementation of Smallfoot and the evidence that the method is scalable [3,33], many tools supporting separation logic as an assertion language have been developed [3,16,33,8,9,17]. Even though the first tools could handle relatively limited fragments of separation logic, like symbolic heaps, there is a growing interest and demand to consider extensions with richer expressive power. We can point out three particular extensions of symbolic heaps (without list predicates) that have been proved decidable.

- Symbolic heaps with generalised inductive predicates, adding a fixpoint combinator to the language, is a convenient logic for specifying data structures that are more advanced than lists or trees. The entailment problem is known to be decidable by means of tree automata techniques for the bounded tree-width fragment [19,1], whereas satisfiability is EXPTIME-complete [6]. Other related results can be found in [21].
- List-free symbolic heaps with all classical Boolean connectives \wedge and \neg (and with the separating conjunction $*$), called herein $SL(*)$, is a convenient extension when combinations of results of various analysis need to be expressed, or when the analysis requires a complementation. This extension already is PSPACE-complete [11].

- Propositional separation logic with separating implication, a.k.a. magic wand (\multimap), is a convenient fragment (called herein $\text{SL}(*, \multimap)$) in which can be solved two problems of frame inference and abduction, that play an important role in static analysers and provers built on top of separation logic. $\text{SL}(*, \multimap)$ can be decided in PSPACE thanks to a small model property [32].

A natural question is how to combine these extensions, and which separation logic fragment that allows Boolean connectives, magic wand and generalised recursive predicates can be decided with some adequate restrictions. As already advocated in [7,31,29,18,24], dealing with the separating implication \multimap is a desirable feature for program verification and several semi-automated or automated verification tools support it in some way, see e.g. [31,29,24,18].

Our contribution. In this paper, we address the question of combining magic wand and inductive predicates in the extremely limited case where the only inductive predicate is the gentle list segment predicate ls . So the starting point of this work is this puzzling question: what is the complexity/decidability status of propositional separation logic $\text{SL}(*, \multimap)$ enriched with the list segment predicate ls (herein called $\text{SL}(*, \multimap, \text{ls})$)? More precisely, we study the decidability/complexity status of extensions of propositional separation logic $\text{SL}(*, \multimap)$ by adding one of the reachability predicates among ls (precise predicate as usual in separation logic), reach (existence of a path, possibly empty) and reach^+ (existence of a non-empty path).

First, we establish that the satisfiability problem for the propositional separation logic $\text{SL}(*, \multimap, \text{ls})$ is undecidable. Our proof is by reduction from the undecidability of first-order separation logic [5,14], using an encoding of the variables as heap cells (see Theorem 1). As a consequence, we also establish that $\text{SL}(*, \multimap, \text{ls})$ is not finitely axiomatisable. Moreover, our reduction requires a rather limited expressive power of the list segment predicate, and we can strengthen our undecidability results to some fragments of $\text{SL}(*, \multimap, \text{ls})$. For instance, surprisingly, the extension of $\text{SL}(*, \multimap)$ with the atomic formulae of the form $\text{reach}(\mathbf{x}, \mathbf{y}) = 2$ and $\text{reach}(\mathbf{x}, \mathbf{y}) = 3$ (existence of a path between \mathbf{x} and \mathbf{y} of respective length 2 or 3) is already undecidable, whereas the satisfiability problem for $\text{SL}(*, \multimap, \text{reach}(\mathbf{x}, \mathbf{y}) = 2)$ is known to be in PSPACE [15].

Second, we show that the satisfiability problem for $\text{SL}(*, \text{reach}^+)$ is PSPACE-complete, extending the well-known result on $\text{SL}(*)$. The PSPACE upper bound relies on a small heap property based on the techniques of test formulae, see e.g. [23,22,4,15], and the PSPACE-hardness of $\text{SL}(*)$ is inherited from [11]. The PSPACE upper bound can be extended to the fragment of $\text{SL}(*, \multimap, \text{reach}^+)$ made of Boolean combinations of formulae from $\text{SL}(*, \text{reach}^+) \cup \text{SL}(*, \multimap)$ (see the developments in Section 4). Even better, we show that the fragment of $\text{SL}(*, \multimap, \text{reach}^+)$ in which reach^+ is not in the scope of \multimap is decidable. As far as we know, this is the largest fragment including full Boolean expressivity, \multimap and ls for which decidability is established.

2 Preliminaries

Let $\text{PVAR} = \{\mathbf{x}, \mathbf{y}, \dots\}$ be a countably infinite set of *program variables* and $\text{LOC} = \{\ell_0, \ell_1, \ell_2, \dots\}$ be a countable infinite set of *locations*. A *memory state* is a pair (s, h) such that $s : \text{PVAR} \rightarrow \text{LOC}$ is a variable valuation (known as the *store*) and $h : \text{LOC} \rightarrow_{\text{fin}} \text{LOC}$ is a partial function with finite domain, known as the *heap*. We write $\text{dom}(h)$ to denote its domain and $\text{ran}(h)$ to denote its range. Given a heap h with $\text{dom}(h) = \{\ell_1, \dots, \ell_n\}$, we also write $\{\ell_1 \mapsto h(\ell_1), \dots, \ell_n \mapsto h(\ell_n)\}$ to denote h . Each $\ell_i \mapsto h(\ell_i)$ is understood as a *memory cell* of h .

As usual, the heaps h_1 and h_2 are said to be *disjoint*, written $h_1 \perp h_2$, if $\text{dom}(h_1) \cap \text{dom}(h_2) = \emptyset$; when this holds, we write $h_1 + h_2$ to denote the heap corresponding to the disjoint union of the graphs of h_1 and h_2 , hence $\text{dom}(h_1 + h_2) = \text{dom}(h_1) \uplus \text{dom}(h_2)$. When the domains of h_1 and h_2 are not disjoint, the composition $h_1 + h_2$ is not defined. Moreover, we write $h' \sqsubseteq h$ to denote that $\text{dom}(h') \subseteq \text{dom}(h)$ and for all locations $\ell \in \text{dom}(h')$, we have $h'(\ell) = h(\ell)$. The formulae φ of the separation logic $\text{SL}(*, \neg*, \mathbf{ls})$ and its atomic formulae π are built from $\pi ::= \mathbf{x} = \mathbf{y} \mid \mathbf{x} \hookrightarrow \mathbf{y} \mid \mathbf{ls}(\mathbf{x}, \mathbf{y}) \mid \mathbf{emp} \mid \top$ and $\varphi ::= \pi \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi * \varphi \mid \varphi \neg* \varphi$, where $\mathbf{x}, \mathbf{y} \in \text{PVAR}$ ($\Rightarrow, \Leftrightarrow$ and \vee are defined as usually). Models of the logic $\text{SL}(*, \neg*, \mathbf{ls})$ are memory states and the satisfaction relation \models is defined as follows (omitting standard clauses for \neg, \wedge):

$$\begin{aligned}
(s, h) \models \mathbf{x} = \mathbf{y} &\iff s(\mathbf{x}) = s(\mathbf{y}) \\
(s, h) \models \mathbf{emp} &\iff \text{dom}(h) = \emptyset \\
(s, h) \models \mathbf{x} \hookrightarrow \mathbf{y} &\iff s(\mathbf{x}) \in \text{dom}(h) \text{ and } h(s(\mathbf{x})) = s(\mathbf{y}) \\
(s, h) \models \mathbf{ls}(\mathbf{x}, \mathbf{y}) &\iff \text{either } (\text{dom}(h) = \emptyset \text{ and } s(\mathbf{x}) = s(\mathbf{y})) \text{ or} \\
&\quad h = \{\ell_0 \mapsto \ell_1, \ell_1 \mapsto \ell_2, \dots, \ell_{n-1} \mapsto \ell_n\} \text{ with } n \geq 1, \\
&\quad \ell_0 = s(\mathbf{x}), \ell_n = s(\mathbf{y}) \text{ and for all } i \neq j \in [0, n], \ell_i \neq \ell_j \\
(s, h) \models \varphi_1 * \varphi_2 &\iff \text{there are } h_1 \text{ and } h_2 \text{ such that } (h_1 \perp h_2, (h_1 + h_2) = h, \\
&\quad (s, h_1) \models \varphi_1 \text{ and } (s, h_2) \models \varphi_2) \\
(s, h) \models \varphi_1 \neg* \varphi_2 &\iff \forall h_1 \text{ if } (h_1 \perp h \text{ and } (s, h_1) \models \varphi_1) \text{ then } (s, h + h_1) \models \varphi_2.
\end{aligned}$$

Note that the semantics for $*$, $\neg*$, \hookrightarrow , \mathbf{ls} and for all other ingredients is the usual one in separation logic and \mathbf{ls} is the *precise* list segment predicate. In the sequel, we use the following abbreviations: $\mathbf{size} \geq 0 \stackrel{\text{def}}{=} \top$ and for all $\beta \geq 0$, $\mathbf{size} \geq \beta + 1 \stackrel{\text{def}}{=} (\mathbf{size} \geq \beta) * \neg \mathbf{emp}$, $\mathbf{size} \leq \beta \stackrel{\text{def}}{=} \neg(\mathbf{size} \geq \beta + 1)$ and $\mathbf{size} = \beta \stackrel{\text{def}}{=} (\mathbf{size} \leq \beta) \wedge (\mathbf{size} \geq \beta)$. Moreover, $\varphi_1 \oplus \varphi_2 \stackrel{\text{def}}{=} \neg(\varphi_1 \neg* \neg\varphi_2)$ (*septraction connective*), $\mathbf{alloc}(\mathbf{x}) \stackrel{\text{def}}{=} (\mathbf{x} \hookrightarrow \mathbf{x}) \neg* \perp$ and $\mathbf{x} \mapsto \mathbf{y} \stackrel{\text{def}}{=} (\mathbf{x} \hookrightarrow \mathbf{y}) \wedge \mathbf{size} = 1$. W.l.o.g., we can assume that $\text{LOC} = \mathbb{N}$ since none of the developments depend on the elements of LOC as the only predicate involving locations is the equality. We write $\text{SL}(*, \neg*)$ to denote the restriction of $\text{SL}(*, \neg*, \mathbf{ls})$ without \mathbf{ls} . Similarly, we write $\text{SL}(*)$ to denote the restriction of $\text{SL}(*, \neg*)$ without $\neg*$. Given two formulae φ, φ' (possibly from different logical languages), we write $\varphi \equiv \varphi'$ whenever for all (s, h) , we have $(s, h) \models \varphi$ iff $(s, h) \models \varphi'$. When $\varphi \equiv \varphi'$, the formulae φ and φ' are said to be *equivalent*.

Variants with other reachability predicates. We use two additional reachability predicates $\mathbf{reach}(\mathbf{x}, \mathbf{y})$ and $\mathbf{reach}^+(\mathbf{x}, \mathbf{y})$ and we write $\text{SL}(*, \neg*, \mathbf{reach})$ (resp.

$\text{SL}(*, \neg, \text{reach}^+)$) to denote the variant of $\text{SL}(*, \neg, \text{ls})$ in which ls is replaced by reach (resp. by reach^+). The relation \models is extended as follows: $(s, h) \models \text{reach}(x, y)$ holds when there is $i \geq 0$ such that $h^i(s(x)) = s(y)$ (i functional composition(s) of h is denoted by h^i) and $(s, h) \models \text{reach}^+(x, y)$ holds when there is $i \geq 1$ such that $h^i(s(x)) = s(y)$. As $\text{ls}(x, y) \equiv \text{reach}(x, y) \wedge \neg(\neg \text{emp} * \text{reach}(x, y))$ and $\text{reach}(x, y) \equiv \top * \text{ls}(x, y)$, the logics $\text{SL}(*, \neg, \text{reach})$ and $\text{SL}(*, \neg, \text{ls})$ have identical decidability status. As far as computational complexity is concerned, a similar analysis can be done as soon as $*$, \neg , \wedge and emp are parts of the fragments (the details are omitted here). Similarly, we have the equivalences: $\text{reach}(x, y) \equiv x = y \vee \text{reach}^+(x, y)$ and $\text{ls}(x, y) \equiv (x = y \wedge \text{emp}) \vee (\text{reach}^+(x, y) \wedge \neg(\neg \text{emp} * \text{reach}^+(x, y)))$. So clearly, $\text{SL}(*, \text{reach})$ and $\text{SL}(*, \text{ls})$ can be viewed as fragments of $\text{SL}(*, \text{reach}^+)$ and, $\text{SL}(*, \neg, \text{ls})$ as a fragment of $\text{SL}(*, \neg, \text{reach}^+)$. It is therefore stronger to establish decidability or complexity upper bounds with reach^+ and to show undecidability or complexity lower bounds with ls or reach . Herein, we provide the optimal results.

Decision problems. Let \mathcal{L} be a logic defined above. As usual, the *satisfiability problem* for \mathcal{L} takes as input a formula φ from \mathcal{L} and asks whether there is (s, h) such that $(s, h) \models \varphi$. The *validity problem* is also defined as usual. The *model-checking problem* for \mathcal{L} takes as input a formula φ from \mathcal{L} , (s, h) and asks whether $(s, h) \models \varphi$ (s is restricted to the variables occurring in φ and h is encoded as a finite and functional graph). Unless otherwise specified, the *size* of a formula φ is understood as its tree size, i.e. approximately its number of symbols.

The main purpose of this paper is to study the decidability/complexity status of $\text{SL}(*, \neg, \text{ls})$ and its fragments.

3 Undecidability of $\text{SL}(*, \neg, \text{ls})$

In this section, we show that $\text{SL}(*, \neg, \text{ls})$ has an undecidable satisfiability problem even though it does not admit first-order quantification.

Let $\text{SL}(\forall, \neg)$ be the first-order extension of $\text{SL}(\neg)$ obtained by adding the universal quantifier \forall . The formulae φ of $\text{SL}(\forall, \neg)$ are built from $\pi ::= x = y \mid x \leftrightarrow y$ and $\varphi ::= \pi \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi * \varphi \mid \forall x \varphi$, where $x, y \in \text{PVAR}$. Note that emp can be easily defined by $\forall x, x' \neg(x \leftrightarrow x')$. Models of the logic $\text{SL}(\forall, \neg)$ are memory states and the satisfaction relation \models is defined as for $\text{SL}(\neg)$ with the additional clause:

$$(s, h) \models \forall x \varphi \iff \text{for all } \ell \in \text{LOC}, \text{ we have } (s[x \leftarrow \ell], h) \models \varphi.$$

Without any loss of generality, we can assume that the satisfiability [resp. validity] problem for $\text{SL}(\forall, \neg)$ is defined by taking as inputs closed formulae (i.e. without free occurrences of the variables).

Proposition 1. [5,14] *The satisfiability problem for $\text{SL}(\forall, \neg)$ is undecidable and the set of valid formulae for $\text{SL}(\forall, \neg)$ is not recursively enumerable.*

In a nutshell, we establish the undecidability of $\text{SL}(*, -*, \mathbf{1s})$ by reduction from the satisfiability problem for $\text{SL}(\forall, -*)$. The reduction is nicely decomposed in two intermediate steps: (1) the undecidability of $\text{SL}(*, -*)$ extended with a few atomic predicates, to be defined soon, and (2) a *tour de force* resulting in the encoding of these atomic predicates in $\text{SL}(*, -*, \mathbf{1s})$.

3.1 Encoding quantified variables as cells in the heap

In this section, we assume for a moment that we can express three atomic predicates $\text{alloc}^{-1}(\mathbf{x})$, $n(\mathbf{x}) = n(\mathbf{y})$ and $n(\mathbf{x}) \leftrightarrow n(\mathbf{y})$, that will be used in the translation and have the following semantics:

- $(s, h) \models \text{alloc}^{-1}(\mathbf{x})$ holds whenever $s(\mathbf{x}) \in \text{ran}(h)$,
- $(s, h) \models n(\mathbf{x}) = n(\mathbf{y})$ holds iff $\{s(\mathbf{x}), s(\mathbf{y})\} \subseteq \text{dom}(h)$ and $h(s(\mathbf{x})) = h(s(\mathbf{y}))$,
- $(s, h) \models n(\mathbf{x}) \leftrightarrow n(\mathbf{y})$ holds iff $\{s(\mathbf{x}), s(\mathbf{y})\} \subseteq \text{dom}(h)$ and $h^2(s(\mathbf{x})) = h(s(\mathbf{y}))$.

Let us first intuitively explain how the two last predicates will help encoding $\text{SL}(\forall, -*)$. By definition, the satisfaction of the quantified formula $\forall \mathbf{x} \psi$ from $\text{SL}(\forall, -*)$ requires the satisfaction of the formula ψ for all the values in LOC assigned to \mathbf{x} . The principle of the encoding is to use a set L of locations initially not in the domain or range of the heap to mimic the store by modifying how they are allocated. In this way, a variable will be interpreted by a location in the heap and, instead of checking whenever $\mathbf{x} \leftrightarrow \mathbf{y}$ (or $\mathbf{x} = \mathbf{y}$) holds, we will check if $n(\mathbf{x}) \leftrightarrow n(\mathbf{y})$ (or $n(\mathbf{x}) = n(\mathbf{y})$) holds, where \mathbf{x} and \mathbf{y} correspond, after the translation, to the locations in L that mimic the store for those variables. Let X be the set of variables needed for the translation. In order to properly encode the store, each location in L only mimics exactly one variable, i.e. there is a bijection between X and L , and cannot be reached by any location. As such, the formula $\forall \mathbf{x} \psi$ will be encoded by the formula $(\text{alloc}(\mathbf{x}) \wedge \text{size} = 1) -* (\text{OK}(X) \Rightarrow \text{T}(\psi))$, where $\text{OK}(X)$ (formally defined below) checks whenever the locations in L still satisfy the auxiliary conditions just described, whereas $\text{T}(\psi)$ is the translation of ψ .

Unfortunately, the formula $\psi_1 -* \psi_2$ cannot simply be translated into $\text{T}(\psi_1) -* (\text{OK}(X) \Rightarrow \text{T}(\psi_2))$ because the evaluation of $\text{T}(\psi_1)$ in a disjoint heap may need the values of free variables occurring in ψ_1 but our encoding of the variable valuations via the heap does not allow to preserve these values through disjoint heaps. In order to solve this problem, for each variable \mathbf{x} in the formula, X will contain an auxiliary variable $\bar{\mathbf{x}}$, or alternatively we define on X an involution $(\bar{\cdot})$. If the translated formula has q variables then the set X of variables needed for the translation will have cardinality $2q$. In the translation of a formula whose outermost connective is the magic wand, the locations corresponding to variables of the form $\bar{\mathbf{x}}$ will be allocated on the left side of the magic wand, and checked to be equal to their non-bar versions on the right side of the magic wand. As such, the left side of the magic wand will be translated into

$$((\bigwedge_{z \in Z} \text{alloc}(\bar{z})) \wedge (\bigwedge_{z \in X \setminus Z} \neg \text{alloc}(z)) \wedge \text{OK}(Z) \wedge \text{T}(\psi_1)[z \leftarrow \bar{z} \mid z \in X]),$$

where Z is the set of free variables in ψ_1 , whereas the right side will be

$$(((\bigwedge_{z \in Z} n(\mathbf{z}) = n(\bar{\mathbf{z}})) \wedge \text{OK}(X)) \Rightarrow ((\bigwedge_{z \in Z} \text{alloc}(\bar{\mathbf{z}}) \wedge \text{size} = \text{card}(Z)) * \text{T}(\psi_2))).$$

The use of the separating conjunction before the formula $\text{T}(\psi_2)$ separates the memory cells corresponding to $\bar{\mathbf{x}}$ from the rest of the heap. By doing this, we can reuse $\bar{\mathbf{x}}$ whenever a magic wand appears in $\text{T}(\psi_2)$.

For technical convenience, we consider a slight alternative for the semantics of the logics $\text{SL}(\forall, *)$ and $\text{SL}(*, *, \mathbf{1s})$, which does not modify the notion of satisfiability/validity and such that the set of formulae and the definition of the satisfaction relation \models remain unchanged. So far, the memory states are pairs of the form (s, h) with $s : \text{PVAR} \rightarrow \text{LOC}$ and $h : \text{LOC} \rightarrow_{\text{fin}} \text{LOC}$ for a *fixed* countably infinite set of locations LOC , say $\text{LOC} = \mathbb{N}$. Alternatively, the models for $\text{SL}(\forall, *)$ and $\text{SL}(*, *, \mathbf{1s})$ can be defined as triples (LOC_1, s_1, h_1) such that LOC_1 is a countable infinite set, $s_1 : \text{PVAR} \rightarrow \text{LOC}_1$ and $h_1 : \text{LOC}_1 \rightarrow_{\text{fin}} \text{LOC}_1$. As shown below, this does not change the notion of satisfiability and validity, but this generalisation will be handy in a few places. Most of the time, a generalised memory state (LOC_1, s_1, h_1) shall be written (s_1, h_1) when no confusion is possible.

Given a bijection $\mathfrak{f} : \text{LOC}_1 \rightarrow \text{LOC}_2$ and a heap $h_1 : \text{LOC}_1 \rightarrow_{\text{fin}} \text{LOC}_1$ equal to $\{\ell_1 \mapsto h_1(\ell_1), \dots, \ell_n \mapsto h_1(\ell_n)\}$, we write $\mathfrak{f}(h_1)$ to denote the heap $h_2 : \text{LOC}_2 \rightarrow_{\text{fin}} \text{LOC}_2$ with $h_2 = \{\mathfrak{f}(\ell_1) \mapsto \mathfrak{f}(h_1(\ell_1)), \dots, \mathfrak{f}(\ell_n) \mapsto \mathfrak{f}(h_1(\ell_n))\}$.

Definition 1. *Let (LOC_1, s_1, h_1) and (LOC_2, s_2, h_2) be generalised memory states and $X \subseteq \text{PVAR}$. A partial isomorphism with respect to X from (LOC_1, s_1, h_1) to (LOC_2, s_2, h_2) is a bijection $\mathfrak{f} : \text{LOC}_1 \rightarrow \text{LOC}_2$ such that $h_2 = \mathfrak{f}(h_1)$ and for all $\mathbf{x} \in X$, $\mathfrak{f}(s_1(\mathbf{x})) = s_2(\mathbf{x})$ (we write $(\text{LOC}_1, s_1, h_1) \approx_X (\text{LOC}_2, s_2, h_2)$).*

A folklore result states that isomorphic memory states satisfy the same formulae since the logics $\text{SL}(\forall, *)$, $\text{SL}(*, *, \mathbf{1s})$ can only perform equality tests.


Lemma 1. *Let (LOC_1, s_1, h_1) and (LOC_2, s_2, h_2) be two generalised memory states such that $(\text{LOC}_1, s_1, h_1) \approx_X (\text{LOC}_2, s_2, h_2)$, for some $X \subseteq \text{PVAR}$. (I) For all formulae φ in $\text{SL}(\forall, *)$ whose free variables are among X , we have $(\text{LOC}_1, s_1, h_1) \models \varphi$ iff $(\text{LOC}_2, s_2, h_2) \models \varphi$. (II) For all formulae φ in $\text{SL}(*, *, \mathbf{1s})$ built on variables among X , we have $(\text{LOC}_1, s_1, h_1) \models \varphi$ iff $(\text{LOC}_2, s_2, h_2) \models \varphi$.*

As a direct consequence, satisfiability in $\text{SL}(*, *, \mathbf{1s})$ as defined in Section 2, is equivalent to satisfiability with generalised memory states, the same holds for $\text{SL}(\forall, *)$. Next, we define the encoding of a generalised memory state. This can be seen as the semantical counterpart of the syntactical translation process and, as such, formalise the intuition of using part of a heap to mimic the store.

Definition 2. *Let $X = \{\mathbf{x}_1, \dots, \mathbf{x}_{2q}\}$, $Y \subseteq \{\mathbf{x}_1, \dots, \mathbf{x}_q\}$ and, (LOC_1, s_1, h_1) and (LOC_2, s_2, h_2) be two (generalised) memory states. We say that (LOC_1, s_1, h_1) is encoded by (LOC_2, s_2, h_2) w.r.t. X, Y , written $(\text{LOC}_1, s_1, h_1) \triangleright_q^Y (\text{LOC}_2, s_2, h_2)$, if the following conditions hold:*

- $\text{LOC}_1 = \text{LOC}_2 \setminus \{s_2(\mathbf{x}) \mid \mathbf{x} \in X\}$,
- for all $\mathbf{x} \neq \mathbf{y} \in X$, $s_2(\mathbf{x}) \neq s_2(\mathbf{y})$,
- $h_2 = h_1 + \{s_2(\mathbf{x}) \mapsto s_1(\mathbf{x}) \mid \mathbf{x} \in Y\}$.

Notice that h_2 is equal to h_1 plus the heap $\{s_2(\mathbf{x}) \mapsto s_1(\mathbf{x}) \mid \mathbf{x} \in Y\}$ that encodes the store s_1 . The picture below presents a memory state (left) and its encoding (right), where $Y = \{\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k\}$. From the encoding, we can retrieve the initial heap by removing the memory cells corresponding to \mathbf{x}_i , \mathbf{x}_j and \mathbf{x}_k . By way of example, the memory state on the left satisfies the formulae $\mathbf{x}_i = \mathbf{x}_j$, $\mathbf{x}_i \hookrightarrow \mathbf{x}_k$ and $\mathbf{x}_k \hookrightarrow \mathbf{x}_k$ whereas its encoding satisfies the formulae $n(\mathbf{x}_i) = n(\mathbf{x}_j)$, $n(\mathbf{x}_i) \hookrightarrow n(\mathbf{x}_k)$ and $n(\mathbf{x}_k) \hookrightarrow n(\mathbf{x}_k)$.



3.2 The translation

We are now ready to define the translation of a first-order formula in propositional separation logic extended with the three predicates introduced at the beginning of the section. Let φ be a closed formula of $\text{SL}(\forall, *)$ with quantified variables $\{\mathbf{x}_1, \dots, \mathbf{x}_q\}$. W.l.o.g., we can assume that distinct quantifications involve distinct variables. Moreover, let $X = \{\mathbf{x}_1, \dots, \mathbf{x}_{2q}\}$ and $(\bar{\cdot})$ be the involution on X such that for all $i \in [1, q]$ $\bar{\mathbf{x}}_i \stackrel{\text{def}}{=} \mathbf{x}_{i+q}$.

We write $\text{OK}(X)$ to denote the formula $(\bigwedge_{i \neq j} \mathbf{x}_i \neq \mathbf{x}_j) \wedge (\bigwedge_i \neg \text{alloc}^{-1}(\mathbf{x}_i))$. The translation function T has two arguments: the formula in $\text{SL}(\forall, *)$ to be recursively translated and the total set of variables potentially appearing in the target formula (useful to check that $\text{OK}(X)$ holds on every heap involved in the satisfaction of the translated formula). Let us come back to the definition of $\text{T}(\psi, X)$ (homomorphic for Boolean connectives) with the assumption that the variables in ψ are among $\mathbf{x}_1, \dots, \mathbf{x}_q$.

$$\text{T}(\mathbf{x}_i = \mathbf{x}_j, X) \stackrel{\text{def}}{=} n(\mathbf{x}_i) = n(\mathbf{x}_j)$$

$$\text{T}(\mathbf{x}_i \hookrightarrow \mathbf{x}_j, X) \stackrel{\text{def}}{=} n(\mathbf{x}_i) \hookrightarrow n(\mathbf{x}_j)$$

$$\text{T}(\forall \mathbf{x}_i \psi, X) \stackrel{\text{def}}{=} (\text{alloc}(\mathbf{x}_i) \wedge \text{size} = 1) * (\text{OK}(X) \Rightarrow \text{T}(\psi, X))$$

Lastly, the translation $\text{T}(\psi_1 * \psi_2, X)$ is defined as

$$((\bigwedge_{z \in Z} \text{alloc}(\bar{z})) \wedge (\bigwedge_{z \in X \setminus Z} \neg \text{alloc}(\bar{z})) \wedge \text{OK}(X) \wedge \text{T}(\psi_1, X)[\mathbf{x} \leftarrow \bar{\mathbf{x}}]) *$$

$$(((\bigwedge_{z \in Z} n(\mathbf{z}) = n(\bar{\mathbf{z}})) \wedge \text{OK}(X)) \Rightarrow ((\bigwedge_{z \in Z} \text{alloc}(\bar{z}) \wedge \text{size} = \text{card}(Z)) * \text{T}(\psi_2, X))),$$

where $Z \subseteq \{\mathbf{x}_1, \dots, \mathbf{x}_q\}$ is the set of free variables in ψ_1 .

Here is the main result of this section, which is essential for the correctness of $\mathcal{T}_{\text{SAT}}(\varphi)$, defined below.

Lemma 2. *Let $X = \{\mathbf{x}_1, \dots, \mathbf{x}_{2q}\}$, $Y \subseteq \{\mathbf{x}_1, \dots, \mathbf{x}_q\}$, ψ be a formula in $\text{SL}(\forall, \neg)$ with free variables among Y that does not contain any bound variable of ψ and $(\text{LOC}_1, s_1, h_1) \triangleright_q^Y (\text{LOC}_2, s_2, h_2)$. We have $(s_1, h_1) \models \psi$ iff $(s_2, h_2) \models \text{T}(\psi, X)$.*

We define the translation $\mathcal{T}_{\text{SAT}}(\varphi)$ in $\text{SL}(*, \neg, \text{ls})$ where $\text{T}(\varphi, X)$ is defined recursively.

$$\mathcal{T}_{\text{SAT}}(\varphi) \stackrel{\text{def}}{=} \left(\bigwedge_{i \in [1, 2q]} \neg \text{alloc}(\mathbf{x}_i) \right) \wedge \text{OK}(X) \wedge \text{T}(\varphi, X).$$

The first two conjuncts specify initial conditions, namely each variable \mathbf{y} in X is interpreted by a location that is unallocated, it is not in the heap range and it is distinct from the interpretation of all other variables; in other words, the value for \mathbf{y} is isolated. Similarly, let $\mathcal{T}_{\text{VAL}}(\varphi)$ be the formula in $\text{SL}(*, \neg, \text{ls})$ defined by $((\bigwedge_{i \in [1, 2q]} \neg \text{alloc}(\mathbf{x}_i)) \wedge \text{OK}(X)) \Rightarrow \text{T}(\varphi, X)$. As a consequence of Lemma 2, φ and $\mathcal{T}_{\text{SAT}}(\varphi)$ are shown equisatisfiable, whereas φ and $\mathcal{T}_{\text{VAL}}(\varphi)$ are shown equivalent.

Corollary 1. *Let φ be a closed formula in $\text{SL}(\forall, \neg)$ using quantified variables among $\{\mathbf{x}_1, \dots, \mathbf{x}_q\}$. (I) φ and $\mathcal{T}_{\text{SAT}}(\varphi)$ are equisatisfiable. (II) φ and $\mathcal{T}_{\text{VAL}}(\varphi)$ are equivalent.*

3.3 Expressing the auxiliary atomic predicates

To complete the reduction, we briefly explain how to express the formulae $\text{alloc}^{-1}(\mathbf{x})$, $n(\mathbf{x}) = n(\mathbf{y})$ and $n(\mathbf{x}) \hookrightarrow n(\mathbf{y})$ within $\text{SL}(*, \neg, \text{ls})$. Let us introduce a few macros that shall be helpful.

- Given φ in $\text{SL}(*, \neg, \text{reach}^+)$ and $\gamma \geq 0$, we write $[\varphi]_\gamma$ to denote the formula $(\text{size} = \gamma \wedge \varphi) * \top$. It is easy to show that for any memory state (s, h) , $(s, h) \models [\varphi]_\gamma$ iff there is $h' \sqsubseteq h$ such that $\text{card}(\text{dom}(h')) = \gamma$ and $(s, h') \models \varphi$.
- We write $\text{reach}(\mathbf{x}, \mathbf{y}) = \gamma$ to denote the formula $[\text{ls}(\mathbf{x}, \mathbf{y})]_\gamma$, which is satisfied in any memory state (s, h) where $h^\gamma(s(\mathbf{x})) = s(\mathbf{y})$. Lastly, we write $\text{reach}(\mathbf{x}, \mathbf{y}) \leq \gamma$ to denote the formula $\bigvee_{0 \leq \gamma' \leq \gamma} \text{reach}(\mathbf{x}, \mathbf{y}) = \gamma'$.

In order to define the existence of a predecessor (i.e. $\text{alloc}^{-1}(\mathbf{x})$) in $\text{SL}(*, \neg, \text{ls})$, we need to take advantage of an auxiliary variable \mathbf{y} whose value is different from the one for \mathbf{x} . Let $\text{alloc}_{\mathbf{y}}^{-1}(\mathbf{x})$ be the formula

$$\mathbf{x} \hookrightarrow \mathbf{x} \vee \mathbf{y} \hookrightarrow \mathbf{x} \vee [(\text{alloc}(\mathbf{y}) \wedge \neg(\mathbf{y} \hookrightarrow \mathbf{x}) \wedge \text{size} = 1) \oplus \text{reach}(\mathbf{y}, \mathbf{x}) = 2]_1$$

Lemma 3. *Let $\mathbf{x}, \mathbf{y} \in \text{PVAR}$. (I) For all memory states (s, h) such that $s(\mathbf{x}) \neq s(\mathbf{y})$, we have $(s, h) \models \text{alloc}_{\mathbf{y}}^{-1}(\mathbf{x})$ iff $s(\mathbf{x}) \in \text{ran}(h)$. (II) In the translation, $\text{alloc}^{-1}(\mathbf{x})$ can be replaced with $\text{alloc}_{\bar{\mathbf{x}}}^{-1}(\mathbf{x})$.*

As stated in Lemma 3(II), we can exploit the fact that in the translation of a formula with variables in $\{x_1, \dots, x_q\}$, we use $2q$ variables that correspond to $2q$ distinguished locations in the heap in order to retain the soundness of the translation while using $\text{alloc}_{\bar{x}}^{-1}(x)$ as $\text{alloc}^{-1}(x)$. Moreover, $\text{alloc}_{\bar{y}}^{-1}(x)$ allows to express in $\text{SL}(*, -*, \text{ls})$ whenever a location corresponding to a program variable reaches itself in exactly two steps (we use this property in the definition of $n(x) \leftrightarrow n(y)$). We write $x \leftrightarrow_y^2 x$ to denote the formula $\neg(x \leftrightarrow x) \wedge (x \leftrightarrow y \Leftrightarrow y \leftrightarrow x) \wedge [\text{alloc}(x) \wedge \text{alloc}_{\bar{y}}^{-1}(x) \wedge (\top \text{ } -* \text{ } \neg \text{reach}(x, y) = 2)]_2$. For any memory state (s, h) such that $s(x) \neq s(y)$, we have $(s, h) \models x \leftrightarrow_y^2 x$ if and only if $h^2(s(x)) = s(x)$ and $h(s(x)) \neq s(x)$.

The predicate $n(x) = n(y)$ can be defined in $\text{SL}(*, -*, \text{ls})$ as

$$\begin{aligned} & (x \neq y \Rightarrow [\text{alloc}(x) \wedge \text{alloc}(y) \wedge ((x \leftrightarrow y \wedge y \leftrightarrow x) \vee (y \leftrightarrow x \wedge x \leftrightarrow y) \\ & ((\bigwedge_{z, z' \in \{x, y\}} \neg(z \leftrightarrow z')) \wedge (\top \text{ } -* \text{ } \neg(\text{reach}(x, y) = 2 \wedge \text{reach}(y, x) = 2))))]_2) \wedge \text{alloc}(x) \end{aligned}$$

Lemma 4. *Let $x, y \in \text{PVAR}$. For all memory states (s, h) , we have $(s, h) \models n(x) = n(y)$ iff $h(s(x)) = h(s(y))$.*

Similarly to $\text{alloc}^{-1}(x)$, we can show that $n(x) \leftrightarrow n(y)$ is definable in $\text{SL}(*, -*, \text{ls})$ by using one additional variable z whose value is different from both x and y . Let $\varphi_{\leftrightarrow}(x, y, z)$ be $(n(x) = n(y) \wedge \varphi_{\bar{\leftrightarrow}}(x, y, z)) \vee (n(x) \neq n(y) \wedge \varphi_{\neq \leftrightarrow}(x, y, z))$ where $\varphi_{\bar{\leftrightarrow}}(x, y, z)$ is defined as

$$\begin{aligned} & (x \leftrightarrow x \wedge y \leftrightarrow x) \vee (y \leftrightarrow y \wedge x \leftrightarrow y) \vee (x \leftrightarrow z \wedge z \leftrightarrow z) \\ & \vee [\text{alloc}(x) \wedge \neg \text{alloc}_{\bar{z}}^{-1}(x) \wedge (\top \text{ } -* \text{ } \neg \text{reach}(x, z) \leq 3)]_2 \end{aligned}$$

whereas $\varphi_{\neq \leftrightarrow}(x, y)$ is defined as

$$\begin{aligned} & (x \leftrightarrow y \wedge \text{alloc}(y)) \vee (y \leftrightarrow y \wedge \text{reach}(x, y) = 2) \vee (y \leftrightarrow x \wedge x \leftrightarrow_y^2 x) \vee \\ & [\text{alloc}(x) \wedge \text{alloc}(y) \wedge (\bigwedge_{z, z' \in \{x, y\}} \neg z \leftrightarrow z') \wedge \neg \text{reach}(x, y) \leq 3 \\ & \wedge ((\text{size} = 1 \wedge \text{alloc}_x^{-1}(y)) \oplus (\text{reach}(x, y) = 3 \wedge y \leftrightarrow_x^2 y))]_3 \end{aligned}$$

Lemma 5. *Let $x, y, z \in \text{PVAR}$. (I) For all memory states (s, h) such that $s(x) \neq s(z)$ and $s(y) \neq s(z)$, we have $(s, h) \models \varphi_{\leftrightarrow}(x, y, z)$ iff $\{s(x), s(y)\} \subseteq \text{dom}(h)$ and $h(h(s(x))) = h(s(y))$; (II) In the translation, $n(x) \leftrightarrow n(y)$ can be replaced by $\varphi_{\leftrightarrow}(x, y, \bar{x})$.*

As for $\text{alloc}_{\bar{y}}^{-1}(x)$, the properties of the translation imply the equivalence between $n(x) \leftrightarrow n(y)$ and $\varphi_{\leftrightarrow}(x, y, \bar{x})$ (as stated in Lemma 5(II)). By looking at the formulae herein defined, the predicate reach only appears bounded, i.e. in the form of $\text{reach}(x, y) = 2$ and $\text{reach}(x, y) = 3$. The three new predicates can therefore be defined in $\text{SL}(*, -*)$ enriched with $\text{reach}(x, y) = 2$ and $\text{reach}(x, y) = 3$.

3.4 Undecidability results and non-finite axiomatization

It is time to collect the fruits of all our efforts and to conclude this part about undecidability. As a direct consequence of Corollary 1 and the undecidability of $\text{SL}(\forall, \neg)$, here is one of the main results of the paper.

Theorem 1. *The satisfiability problem for $\text{SL}(*, \neg, \text{ls})$ is undecidable.*

As a by-product, the set of valid formulae for $\text{SL}(*, \neg, \text{ls})$ is not recursively enumerable. Indeed, suppose that the set of valid formulae for $\text{SL}(*, \neg, \text{ls})$ were r.e., then one can enumerate the valid formulae of the form $\mathcal{T}_{\text{VAL}}(\varphi)$ as it is decidable in PTIME whether ψ in $\text{SL}(*, \neg, \text{ls})$ is syntactically equal to $\mathcal{T}_{\text{VAL}}(\varphi)$ for some $\text{SL}(\forall, \neg)$ formula φ . This leads to a contradiction since this would allow the enumeration of valid formulae in $\text{SL}(\forall, \neg)$.

The essential ingredients to establish the undecidability of $\text{SL}(*, \neg, \text{ls})$ are the fact that the following properties $n(\mathbf{x}) = n(\mathbf{y})$, $n(\mathbf{x}) \leftrightarrow n(\mathbf{y})$ and $\text{alloc}^{-1}(\mathbf{x})$ are expressible in the logic.

Corollary 2. *$\text{SL}(*, \neg)$ augmented with built-in formulae of the form $n(\mathbf{x}) = n(\mathbf{y})$, $n(\mathbf{x}) \leftrightarrow n(\mathbf{y})$ and $\text{alloc}^{-1}(\mathbf{x})$ (resp. of the form $\text{reach}(\mathbf{x}, \mathbf{y}) = 2$ and $\text{reach}(\mathbf{x}, \mathbf{y}) = 3$) admits an undecidable satisfiability problem.*

This is the addition of $\text{reach}(\mathbf{x}, \mathbf{y}) = 3$ that is crucial for undecidability since the satisfiability problem for $\text{SL}(*, \neg, \text{reach}(\mathbf{x}, \mathbf{y}) = 2)$ is in PSPACE [15]. Following a similar analysis, let $\text{SL1}(\forall, *, \neg)$ be the restriction of $\text{SL}(\forall, *, \neg)$ (i.e. $\text{SL}(\forall, *, \neg)$ plus $*$) to formulae of the form $\exists \mathbf{x}_1 \cdots \exists \mathbf{x}_q \varphi$, where $q \geq 1$, the variables in φ are among $\{\mathbf{x}_1, \dots, \mathbf{x}_{q+1}\}$ and the only quantified variable in φ is \mathbf{x}_{q+1} . The satisfiability problem for $\text{SL1}(\forall, *, \neg)$ is PSPACE-complete [15]. Note that $\text{SL1}(\forall, *, \neg)$ can easily express $n(\mathbf{x}) = n(\mathbf{y})$ and $\text{alloc}^{-1}(\mathbf{x})$. The distance between the decidability for $\text{SL1}(\forall, *, \neg)$ and the undecidability for $\text{SL}(*, \neg, \text{ls})$, is best witnessed by the corollary below, which solves an open problem [15, Section 6].

Corollary 3. *$\text{SL1}(\forall, *, \neg)$ augmented with $n(\mathbf{x}) \leftrightarrow n(\mathbf{y})$ (resp. $\text{SL1}(\forall, *, \neg)$ augmented with ls) admits an undecidable satisfiability problem.*

4 $\text{SL}(*, \text{reach}^+)$ and other PSPACE variants

As already seen in Section 2, $\text{SL}(*, \text{ls})$ can be understood as a fragment of $\text{SL}(*, \text{reach}^+)$. Below, we show that the satisfiability problem for $\text{SL}(*, \text{reach}^+)$ can be solved in polynomial space. Refining the arguments used in our proof, we also show the decidability of the fragment of $\text{SL}(*, \neg, \text{reach}^+)$ where reach^+ is constrained not to occur in the scope of \neg , i.e. φ belongs to that fragment iff for any subformula ψ of φ of the form $\psi_1 \neg \psi_2$, reach^+ does not occur in ψ_1 and in ψ_2 .

The proof relies on a small heap property: a formula φ is satisfiable if and only if it admits a model with a polynomial amount of memory cells. The PSPACE upper bound then follows by establishing that the model-checking problem for $\text{SL}(*, \text{reach}^+)$ is in PSPACE too. To establish the small heap property, an equivalence relation on memory states with finite index is designed, following the standard approach in [32,10] and using test formulae as in [23,22,4,15].

4.1 Introduction to test formulae

Before presenting the test formulae for $\text{SL}(*, \text{reach}^+)$, let us recall the standard result for $\text{SL}(*, -*)$ (that will be also used at some point later on).

Proposition 2. [32,22] *Any formula φ in $\text{SL}(*, -*)$ built over variables in $\mathbf{x}_1, \dots, \mathbf{x}_q$ is logically equivalent to a Boolean combination of formulae among $\mathbf{x}_i = \mathbf{x}_j$, $\text{alloc}(\mathbf{x}_i)$, $\mathbf{x}_i \hookrightarrow \mathbf{x}_j$ and $\text{size} \geq \beta$ ($i, j \in \{1, \dots, q\}$, $\beta \in \mathbb{N}$).*

By way of example, $(\neg \text{emp} * ((\mathbf{x}_1 \hookrightarrow \mathbf{x}_2) -* \perp))$ is equivalent to $\text{size} \geq 2 \wedge \text{alloc}(\mathbf{x}_1)$. As a corollary of the proof of Proposition 2, in $\text{size} \geq \beta$ we can enforce that $\beta \leq 2 \times |\varphi|$ (rough upper bound) where $|\varphi|$ is the size of φ . Similar results will be shown for $\text{SL}(*, \text{reach}^+)$ and for some of its extensions.

In order to define a set of test formulae that captures the expressive power of $\text{SL}(*, \text{reach}^+)$, we need to study which basic properties on memory states can be expressed by $\text{SL}(*, \text{reach}^+)$ formulae. For example, consider the memory states from Figure 1.




Fig. 1. Memory states $(s_1, h_1), \dots, (s_4, h_4)$ (from left to right)


The fragment memory states (s_1, h_1) and (s_2, h_2) can be distinguished by the formula $\top * (\text{reach}(\mathbf{x}_i, \mathbf{x}_j) \wedge \text{reach}(\mathbf{x}_j, \mathbf{x}_k) \wedge \neg \text{reach}(\mathbf{x}_k, \mathbf{x}_i))$. Indeed, (s_1, h_1) satisfies this formula by considering a subheap that does not contain a path from $s(\mathbf{x}_k)$ to $s(\mathbf{x}_i)$, whereas it is impossible to find a subheap for (s_2, h_2) that retains the path from $s(\mathbf{x}_i)$ to $s(\mathbf{x}_j)$, the one from $s(\mathbf{x}_j)$ to $s(\mathbf{x}_k)$ but where the path from $s(\mathbf{x}_k)$ to $s(\mathbf{x}_i)$ is lost. This suggests that $\text{SL}(*, \text{reach}^+)$ can express whether, for example, any path from $s(\mathbf{x}_i)$ to $s(\mathbf{x}_j)$ also contains $s(\mathbf{x}_k)$. We will introduce the test formula $\text{sees}_q(\mathbf{x}_i, \mathbf{x}_j) \geq \beta$ to capture this property.

Similarly, the memory states (s_3, h_3) and (s_4, h_4) can be distinguished by the formula $(\text{size} = 1) * (\text{reach}(\mathbf{x}_j, \mathbf{x}_k) \wedge \neg \text{reach}(\mathbf{x}_i, \mathbf{x}_k) \wedge \neg \text{reach}^+(\mathbf{x}_k, \mathbf{x}_k))$. The memory state (s_3, h_3) satisfies this formula by separating $\{\ell \mapsto \ell'\}$ from the rest of the heap, whereas the formula is not satisfied by (s_4, h_4) . Indeed, there is no way to break the loop from $s(\mathbf{x}_k)$ to itself by removing just one location from the heap while retaining the path from $s(\mathbf{x}_j)$ to $s(\mathbf{x}_k)$ and loosing the path from $s(\mathbf{x}_i)$ to $s(\mathbf{x}_k)$. This suggests that the two locations ℓ and ℓ' are particularly interesting since they are reachable from several locations corresponding to program variables. Therefore by separating them from the rest of the heap, several paths are lost. In order to capture this, we introduce the notion of *meet-points*.

Let Terms_q be the set $\{\mathbf{x}_1, \dots, \mathbf{x}_q\} \cup \{m_q(\mathbf{x}_i, \mathbf{x}_j) \mid i, j \in [1, q]\}$ understood as the set of *terms* that are either variables or expressions denoting a meet-point. We write $\llbracket \mathbf{x}_i \rrbracket_{s,h}^q$ to denote $s(\mathbf{x}_i)$ and $\llbracket m_q(\mathbf{x}_i, \mathbf{x}_j) \rrbracket_{s,h}^q$ to denote (if it exists) the first location reachable from $s(\mathbf{x}_i)$ that is also reachable from $s(\mathbf{x}_j)$. Moreover we require that this location can reach another location corresponding to a program variable. Formally, $\llbracket m_q(\mathbf{x}_i, \mathbf{x}_j) \rrbracket_{s,h}^q$ is defined as the unique location ℓ such that

- there are $L_1, L_2 \geq 0$ such that $h^{L_1}(s(\mathbf{x}_i)) = h^{L_2}(s(\mathbf{x}_j)) = \ell$, and
- for all $L'_1 < L_1$ and for all $L'_2 \geq 0$, $h^{L'_1}(s(\mathbf{x}_i)) \neq h^{L'_2}(s(\mathbf{x}_j))$, and
- there exist $k \in [1, q]$ and $L \geq 0$ such that $h^L(\ell) = s(\mathbf{x}_k)$.

These conditions hold for at most one location ℓ . One can easily show that the notion $\llbracket m_q(\mathbf{x}_i, \mathbf{x}_j) \rrbracket_{s,h}^q$ is well-defined. The picture below provides a taxonomy of meet-points, where arrows labelled by ‘+’ represent paths of non-zero length and zig-zag arrows any path (possibly of zero length). Symmetrical cases, obtained by swapping \mathbf{x}_i and \mathbf{x}_j , are omitted.



Notice how the asymmetrical definition of meet-points is captured in the two rightmost heaps. Consider the memory states from Figure 1, (s_3, h_3) and (s_4, h_4) can be seen as an instance of the third case of the taxonomy and, as such, it holds that $\llbracket m_q(\mathbf{x}_i, \mathbf{x}_j) \rrbracket_{s_3, h_3}^q = \ell$ and $\llbracket m_q(\mathbf{x}_j, \mathbf{x}_i) \rrbracket_{s_3, h_3}^q = \ell'$.

Given $q, \alpha \geq 1$, we write $\text{Test}(q, \alpha)$ to denote the following set of atomic formulae (also called *test formulae*):

$$v = v' \quad v \hookrightarrow v' \quad \text{alloc}(v) \quad \text{sees}_q(v, v') \geq \beta + 1 \quad \text{sizeR}_q \geq \beta,$$

where $v, v' \in \text{Terms}_q$ and $\beta \in [1, \alpha]$. It is worth noting that the $\text{alloc}(v)$'s are not needed for the logic $\text{SL}(*, \text{reach}^+)$ but it is required for extensions.

We identify as special locations the $s(\mathbf{x}_i)$'s and the meet-points of the form $\llbracket m_q(\mathbf{x}_i, \mathbf{x}_j) \rrbracket_{s,h}^q$ when it exists ($i, j \in [1, q]$). We call such locations, *labelled locations*, and the set of labelled locations is written $\text{Labels}_{s,h}^q$. The formal semantics of the test formulae is provided below:

$$\begin{aligned} (s, h) \models v = v' & \iff \llbracket v \rrbracket_{s,h}^q, \llbracket v' \rrbracket_{s,h}^q \text{ are defined, } \llbracket v \rrbracket_{s,h}^q = \llbracket v' \rrbracket_{s,h}^q \\ (s, h) \models \text{alloc}(v) & \iff \llbracket v \rrbracket_{s,h}^q \text{ is defined and belongs to } \text{dom}(h) \\ (s, h) \models v \hookrightarrow v' & \iff h(\llbracket v \rrbracket_{s,h}^q) = \llbracket v' \rrbracket_{s,h}^q \\ (s, h) \models \text{sees}_q(v, v') \geq \beta + 1 & \iff \exists L \geq \beta + 1, h^L(\llbracket v \rrbracket_{s,h}^q) = \llbracket v' \rrbracket_{s,h}^q \text{ and} \\ & \quad \forall 0 < L' < L, h^{L'}(\llbracket v \rrbracket_{s,h}^q) \notin \text{Labels}_{s,h}^q \\ (s, h) \models \text{sizeR}_q \geq \beta & \iff \text{card}(\text{Rem}_{s,h}^q) \geq \beta \end{aligned}$$

where $\text{Rem}_{s,h}^q$ is the set of programs that neither belong to a path between two locations interpreted by program variables nor are equal to program variable interpretations, i.e. $\text{Rem}_{s,h}^q \stackrel{\text{def}}{=} \{\ell \in \text{dom}(h) \mid \forall i \in [1, q], s(\mathbf{x}_i) \neq \ell \text{ and } \forall j \in [1, q] \nexists L, L' \geq 1, h^L(s(\mathbf{x}_i)) = \ell \text{ and } h^{L'}(\ell) = s(\mathbf{x}_j)\}$. There is no need for test formulae of the form $\text{sees}_q(v, v') \geq 1$ since they are equivalent to $v \hookrightarrow v' \vee \text{sees}_q(v, v') \geq 2$. One can check whether $\llbracket m_q(\mathbf{x}_i, \mathbf{x}_j) \rrbracket_{s,h}^q$ is defined thanks to the formula $m_q(\mathbf{x}_i, \mathbf{x}_j) = m_q(\mathbf{x}_i, \mathbf{x}_j)$. By contrast, $\text{sizeR}_q \geq \beta$ states that the cardinality of the set $\text{Rem}_{s,h}^q$ is at least β . Furthermore, $\text{sees}_q(v, v') \geq \beta + 1$ states that there is a minimal path between v and v' of length at least $\beta + 1$ and strictly between v and v' , there are no labelled locations. The satisfaction of $\text{sees}_q(v, v') \geq \beta + 1$ entails the exclusion of labelled locations in the witness path, which is reminiscent to $T \xrightarrow{h \setminus T''} T'$ in the logic GRASS [26]. So, the test formulae are quite expressive since they capture the atomic formulae from $\text{SL}(*, \text{reach}^+)$ and the test formulae for $\text{SL}(*, -*)$.

Lemma 6. *Given $\alpha, q \geq 1, i, j \in [1, q]$, for any atomic formula among $\text{ls}(\mathbf{x}_i, \mathbf{x}_j)$, $\text{reach}(\mathbf{x}_i, \mathbf{x}_j)$, $\text{reach}^+(\mathbf{x}_i, \mathbf{x}_j)$, emp and $\text{size} \geq \beta$ with $\beta \leq \alpha$, there is a Boolean combination of test formulae from $\text{Test}(q, \alpha)$ logically equivalent to it.*

4.2 Expressive power and small model property

The sets of test formulae $\text{Test}(q, \alpha)$ are sufficient to capture the expressive power of $\text{SL}(*, \text{reach}^+)$ (as shown below, Theorem 2) and deduce the small heap property of this logic (Theorem 3). We introduce an indistinguishability relation between memory states based on test formulae, see analogous relations in [22,13,15].

Definition 3. *Given $q, \alpha \geq 1$, we write $(s, h) \approx_\alpha^q (s', h') \stackrel{\text{def}}{\iff} \text{for all } \psi \in \text{Test}(q, \alpha), (s, h) \models \psi \text{ iff } (s', h') \models \psi$.*

Theorem 2(I) states that if $(s, h) \approx_\alpha^q (s', h')$, then the two memory states cannot be distinguished by formulae whose syntactic resources are bounded in some way by q and α (details will follow, see the definition for $\text{msize}(\varphi)$).

Below, we state the key intermediate result of the section that can be viewed as a distributivity lemma. The expressive power of the test formulae allows us to mimic the separation between two equivalent memory states with respect to the relation \approx_α^q , which is essential in the proof of Theorem 2(I).

Lemma 7. *Let $q, \alpha, \alpha_1, \alpha_2 \geq 1$ with $\alpha = \alpha_1 + \alpha_2$ and $(s, h), (s', h')$ be such that $(s, h) \approx_\alpha^q (s', h')$. For all heaps h_1, h_2 such that $h = h_1 + h_2$ there are heaps h'_1, h'_2 such that $h = h'_1 + h'_2$, $(s, h_1) \approx_{\alpha_1}^q (s', h'_1)$ and $(s, h_2) \approx_{\alpha_2}^q (s', h'_2)$.*

For each formula φ in $\text{SL}(*, \text{reach}^+)$, we define its *memory size* $\text{msize}(\varphi)$ following the clauses below (see also [32]).

$$\begin{aligned} \text{msize}(\pi) &\stackrel{\text{def}}{=} 1 \quad \text{for any atomic formula } \pi \\ \text{msize}(\psi * \psi') &\stackrel{\text{def}}{=} \text{msize}(\psi) + \text{msize}(\psi') \\ \text{msize}(\psi \wedge \psi') &\stackrel{\text{def}}{=} \max(\text{msize}(\psi), \text{msize}(\psi')) \\ \text{msize}(\neg\psi) &\stackrel{\text{def}}{=} \text{msize}(\psi). \end{aligned}$$

We have $1 \leq \text{msize}(\varphi) \leq |\varphi|$. Theorem 2 below establishes the properties that formulae in $\text{SL}(*, \text{reach}^+)$ can express.

Theorem 2. *Let φ be in $\text{SL}(*, \text{reach}^+)$ built over the variables in $\mathbf{x}_1, \dots, \mathbf{x}_q$. (I) For all $\alpha \geq 1$ such that $\text{msize}(\varphi) \leq \alpha$ and for all memory states $(s, h), (s', h')$ such that $(s, h) \approx_\alpha^q (s', h')$, we have $(s, h) \models \varphi$ iff $(s', h') \models \varphi$. (II) φ is logically equivalent to a Boolean combination of test formulae from $\text{Test}(q, \text{msize}(\varphi))$.*

The proof of Theorem 2(I) is by structural induction on φ . The basic cases for atomic formulae follow from Lemma 6 whereas the inductive cases for Boolean connectives are immediate. For the separating conjunction, suppose $(s, h) \models \varphi_1 * \varphi_2$ and $\text{msize}(\varphi_1 * \varphi_2) \leq \alpha$. There are heaps h_1 and h_2 such that $h = h_1 + h_2$, $(s, h_1) \models \psi_1$ and $(s, h_2) \models \psi_2$. As $\alpha \geq \text{msize}(\psi_1 * \psi_2) = \text{msize}(\psi_1) + \text{msize}(\psi_2)$, there exist α_1 and α_2 such that $\alpha = \alpha_1 + \alpha_2$, $\alpha_1 \geq \text{msize}(\psi_1)$ and $\alpha_2 \geq \text{msize}(\psi_2)$. By Lemma 7, there exist heaps h'_1 and h'_2 such that $h' = h'_1 + h'_2$, $(s, h_1) \approx_{\alpha_1}^q (s', h'_1)$ and $(s, h_2) \approx_{\alpha_2}^q (s', h'_2)$. By the induction hypothesis, we get $(s', h'_1) \models \psi_1$ and $(s', h'_2) \models \psi_2$. Consequently, we obtain $(s', h') \models \psi_1 * \psi_2$.

As an example, we can apply this result to the memory states from Figure 1. We have already shown how we can distinguish (s_1, h_1) from (s_2, h_2) using a formula with only one separating conjunction. Theorem 2 ensures that these two memory states do not satisfy the same set of test formulae for $\alpha \geq 2$. Indeed, only (s_1, h_1) satisfies $\text{sees}_q(\mathbf{x}_i, \mathbf{x}_j) \geq 2$. The same argument can be used with (s_3, h_3) and (s_4, h_4) : only (s_3, h_3) satisfies the test formula $m_q(\mathbf{x}_i, \mathbf{x}_j) \hookrightarrow m_q(\mathbf{x}_j, \mathbf{x}_i)$. Clearly, Theorem 2(II) relates separation logic with classical logic as advocated also in the works [23,10]. Now, it is possible to establish a small heap property.

Theorem 3. *Let φ be a satisfiable $\text{SL}(*, \text{reach}^+)$ formula built over $\mathbf{x}_1, \dots, \mathbf{x}_q$. There is (s, h) such that $(s, h) \models \varphi$ and $\text{card}(\text{dom}(h)) \leq (q^2 + q) \cdot (|\varphi| + 1) + |\varphi|$.*

The small heap property for $\text{SL}(*, \text{reach}^+)$ is inherited from the small heap property for the Boolean combinations of test formulae, which is analogous to the small model property for other theories of singly linked lists, see e.g. [27,13].

4.3 Complexity upper bounds

Let us draw some consequences of Theorem 3. First, for the logic $\text{SL}(*, \text{reach}^+)$, we get a PSPACE upper, which matches the lower bound for $\text{SL}(*)$ [11].

Theorem 4. *The satisfiability problem for $\text{SL}(*, \text{reach}^+)$ is PSPACE-complete.*

Besides, we may consider restricting the usage of Boolean connectives. We note $\text{Bool}(\text{SHF})$ for the Boolean combinations of formulae from the symbolic heap fragment [2]. A PTIME upper bound for the entailment/satisfiability problem for the symbolic heap fragment is successfully solved in [12,17], whereas the satisfiability problem for a slight variant of $\text{Bool}(\text{SHF})$ is shown in NP in [26, Theorem 4]. Theorem 3 allows us to conclude this NP upper bound result as a by-product (we conjecture that our quadratic upper bound on the number of cells could be improved to a linear one in that case).

Corollary 4. *The satisfiability problem for $\text{Bool}(\text{SHF})$ is NP-complete.*

It is possible to push further the PSPACE upper bound by allowing occurrences of -* in a controlled way. Let $\text{SL}(*, \text{reach}^+, \bigcup_{q,\alpha} \text{Test}(q, \alpha))$ be the extension of $\text{SL}(*, \text{reach}^+)$ augmented with the test formulae. The memory size function is also extended: $\text{msize}(v \hookrightarrow v') \stackrel{\text{def}}{=} 1$, $\text{msize}(\text{sees}_q(v, v') \geq \beta + 1) \stackrel{\text{def}}{=} \beta + 1$, $\text{msize}(\text{sizeR} \geq \beta) \stackrel{\text{def}}{=} \beta$ and $\text{msize}(\text{alloc}(v)) \stackrel{\text{def}}{=} 1$. When formulae are encoded as trees, we have $1 \leq \text{msize}(\varphi) \leq |\varphi|\alpha_\varphi$ where α_φ is the maximal constant in φ . Theorem 2(I) admits a counterpart for $\text{SL}(*, \text{reach}^+, \bigcup_{q,\alpha} \text{Test}(q, \alpha))$ and consequently, any formula built over $\mathbf{x}_1, \dots, \mathbf{x}_q$ can be shown equivalent to a Boolean combination of test formulae from $\text{Test}(q, |\varphi|\alpha_\varphi)$. By Theorem 3, any satisfiable formula has therefore a model with $\text{card}(\text{dom}(h)) \leq (q^2 + q) \cdot (|\varphi|\alpha_\varphi + 1) + |\varphi|\alpha_\varphi$. Hence, the satisfiability problem for $\text{SL}(*, \text{reach}^+, \bigcup_{q,\alpha} \text{Test}(q, \alpha))$ is in PSPACE when the constants are encoded in unary. Now, we can state the new PSPACE upper bound for Boolean combinations of formulae from $\text{SL}(*, \text{-*}) \cup \text{SL}(*, \text{reach}^+)$.

Theorem 5. *The satisfiability problem for Boolean combinations of formulae from $\text{SL}(*, \text{-*}) \cup \text{SL}(*, \text{reach}^+)$ is PSPACE-complete.*

To conclude, let us introduce the largest fragment including -* and 1s for which decidability can be established so far.

Theorem 6. *The satisfiability problem for the fragment of $\text{SL}(*, \text{-*}, \text{reach}^+)$ in which reach^+ is not in the scope of -* is decidable.*

5 Conclusion

We studied the effects of adding 1s to $\text{SL}(*, \text{-*})$ and variants. $\text{SL}(*, \text{-*}, \text{1s})$ is shown undecidable (Theorem 1) and non-finitely axiomatisable, which remains quite unexpected since there are no first-order quantifications. This result is strengthened to even weaker extensions of $\text{SL}(*, \text{-*})$ such as the one augmented with $n(\mathbf{x}) = n(\mathbf{y})$, $n(\mathbf{x}) \hookrightarrow n(\mathbf{y})$ and $\text{alloc}^{-1}(\mathbf{x})$, or the one augmented with $\text{reach}(\mathbf{x}, \mathbf{y}) = 2$ and $\text{reach}(\mathbf{x}, \mathbf{y}) = 3$. If the magic wand is discarded, we have established that the satisfiability problem for $\text{SL}(*, \text{1s})$ is PSPACE-complete by introducing a class of test formulae that captures the expressive power of $\text{SL}(*, \text{1s})$ and that leads to a small heap property. Such a logic contains the Boolean combinations of symbolic heaps and our proof technique allows us to get an NP upper bound for such formulae. Moreover, we show that the satisfiability problem for $\text{SL}(*, \text{-*}, \text{reach}^+)$ restricted to formulae in which reach^+ is not in the scope of -* is decidable, leading to the largest known decidable fragment for which -* and reach^+ (or 1s) cohabit. So, we have provided proof techniques to establish undecidability when $*$, -* and 1s are present and to establish decidability based on test formulae. This paves the way to investigate the decidability status of $\text{SL}(\text{-*}, \text{1s})$ as well as of the positive fragment of $\text{SL}(*, \text{-*}, \text{1s})$ from [30,31].

References

1. T. Antonopoulos, N. Gorogiannis, C. Haase, M. Kanovich, and J. Ouaknine. Foundations for decision problems in separation logic with general inductive predicates. In *FOSSACS'14*, volume 8412 of *LNCS*, pages 411–425. Springer, 2014.
2. J. Berdine, C. Calcagno, and P. O’Hearn. A decidable fragment of separation logic. In *FSTTCS'04*, volume 3328 of *LNCS*, pages 97–109. Springer, 2004.
3. J. Berdine, C. Calcagno, and P. O’Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In *FMCO'05*, volume 4111 of *LNCS*, pages 115–137. Springer, 2005.
4. R. Brochenin, S. Demri, and E. Lozes. Reasoning about sequences of memory states. *APAL*, 161(3):305–323, 2009.
5. R. Brochenin, S. Demri, and E. Lozes. On the almighty wand. *IC*, 211:106–137, 2012.
6. J. Brotherston, C. Fuhs, N. Gorogiannis, and J. Navarro Perez. A decision procedure for satisfiability in separation logic with inductive predicates. In *CSL-LICS'14*, 2014.
7. J. Brotherston and J. Villard. Parametric completeness for separation theories. In *POPL'14*, pages 453–464. ACM, 2014.
8. C. Calcagno and D. Distefano. Infer: An automatic program verifier for memory safety of C programs. In *NASA Formal Methods*, volume 6617 of *LNCS*, pages 459–465. Springer, 2011.
9. C. Calcagno, D. Distefano, P. O’Hearn, and H. Yang. Compositional shape analysis by means of bi-abduction. *JACM*, 58(6):26:1–26:66, 2011.
10. C. Calcagno, P. Gardner, and M. Hague. From separation logic to first-order logic. In *FOSSACS'05*, volume 3441 of *LNCS*, pages 395–409. Springer, 2005.
11. C. Calcagno, P. O’Hearn, and H. Yang. Computability and complexity results for a spatial assertion language for data structures. In *FSTTCS'01*, volume 2245 of *LNCS*, pages 108–119. Springer, 2001.
12. B. Cook, C. Haase, J. Ouaknine, M. Parkinson, and J. Worrell. Tractable reasoning in a fragment of separation logic. In *CONCUR'11*, volume 6901 of *LNCS*, pages 235–249. Springer, 2011.
13. C. David, D. Kroening, and M. Lewis. Propositional reasoning about safety and termination of heap-manipulating programs. In *ESOP'15*, volume 9032 of *LNCS*, pages 661–684. Springer, 2015.
14. S. Demri and M. Deters. Expressive completeness of separation logic with two variables and no separating conjunction. *ACM ToCL*, 17(2):12, 2016.
15. S. Demri, D. Galmiche, D. Larchey-Wendling, and D. Mery. Separation logic with one quantified variable. *Theory of Computing Systems*, 61:371–461, 2017.
16. D. Distefano, P. O’Hearn, and H. Yang. A local shape analysis based on separation logic. In *TACAS'06*, volume 3920 of *LNCS*, pages 287–302. Springer, 2006.
17. C. Haase, S. Ishtiaq, J. Ouaknine, and M. Parkinson. SeLogger: A tool for graph-based reasoning in separation logic. In *CAV'13*, volume 8044 of *LNCS*, pages 790–795. Springer, 2013.
18. Z. Hou, R. Goré, and A. Tiu. Automated theorem proving for assertions in separation logic with all connectives. In *CADE'15*, volume 9195 of *LNCS*, pages 501–516. Springer, 2015.
19. R. Iosif, A. Rogalewicz, and J. Simacek. The tree width of separation logic with recursive definitions. In *CADE'13*, volume 7898 of *LNCS*, pages 21–38. Springer, 2013.

20. S. Ishtiaq and P. O’Hearn. BI as an assertion language for mutable data structures. In *POPL’01*, pages 14–26. ACM, 2001.
21. Q. Le, M. Tatsuta, J. Sun, and W. Chin. A decidable fragment in separation logic with inductive predicates and arithmetic. In *CAV’17*, volume 10427 of *LNCS*, pages 495–517. Springer, 2017.
22. E. Lozes. *Expressivité des Logiques Spatiales*. Phd thesis, ENS Lyon, 2004.
23. E. Lozes. Separation logic preserves the expressive power of classical logic. In *SPACE’04*, 2004.
24. P. Müller, M. Schwerhoff, and A. Summers. Viper: A verification infrastructure for permission-based reasoning. In *VMCAI’16*, volume 9583 of *LNCS*, pages 41–62. Springer, 2016.
25. P. O’Hearn, J. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *CSL’01*, volume 2142 of *LNCS*, pages 1–19. Springer, 2001.
26. R. Piskac, T. Wies, and D. Zufferey. Automating separation logic using SMT. In *CAV’13*, volume 8044 of *LNCS*, pages 773–789. Springer, 2013.
27. S. Ranise and C. Zarba. A theory of singly-linked lists and its extensible decision procedure. In *SEFM’06*, pages 206–215. IEEE, 2006.
28. J. Reynolds. Separation logic: a logic for shared mutable data structures. In *LICS’02*, pages 55–74. IEEE, 2002.
29. M. Schwerhoff and A. Summers. Lightweight support for magic wands in an automatic verifier. In *ECOOP’15*, pages 999–1023. Leibniz-Zentrum für Informatik, LIPICs, 2015.
30. A. Thakur. *Symbolic Abstraction: Algorithms and Applications*. PhD thesis, University of Wisconsin-Madison, 2014.
31. A. Thakur, J. Breck, and T. Reps. Satisfiability modulo abstraction for separation logic with linked lists. In *SPIN’14*, pages 58–67. ACM, 2014.
32. H. Yang. *Local Reasoning for Stateful Programs*. PhD thesis, University of Illinois, Urbana-Champaign, 2001.
33. H. Yang, O. Lee, J. Berdine, C. Calcagno, B. Cook, D. Distefano, and P. O’Hearn. Scalable shape analysis for systems code. In *CAV’08*, volume 5123 of *LNCS*, pages 385–398. Springer, 2008.