EasyChair Preprint

№ 433

The Complexity of Prenex Separation Logic with One Selector

Mnacho Echenim, Radu Iosif and Nicolas Peltier

August 16, 2018

The Complexity of Prenex Separation Logic with One Selector

M. Echenim[1], R. Iosif[2] and N. Peltier[1]

[1] Univ. Grenoble Alpes, CNRS, LIG, F-38000 Grenoble France
[2] Univ. Grenoble Alpes, CNRS, VERIMAG, F-38000 Grenoble France

Abstract. We first show that infinite satisfiability can be reduced to finite satisfiability for all prenex formulas of Separation Logic with $k \geq 1$ selector fields (SL^k). Second, we show that this entails the decidability of the finite and infinite satisfiability problem for the class of prenex formulas of SL^1, by reduction to the first-order theory of one unary function symbol and unary predicate symbols. We also prove that the complexity is not elementary, by reduction from the first-order theory of one unary function symbol. Finally, we prove that the Bernays-Schönfinkel-Ramsey fragment of prenex SL^1 formulae with quantifier prefix in the language $\exists^*\forall^*$ is PSPACE-complete. The definition of a complete (hierarchical) classification of the complexity of prenex SL^1, according to the quantifier alternation depth is left as an open problem.

1 Introduction

Separation Logic [9,13] (SL) is a logical framework used in program verification to describe properties of the heap memory, such as the placement of pointer variables within the topology of complex data structures, such as lists or trees. The features that make SL attractive for program verification are the ability of defining (i) weakest pre- and post-condition calculi that capture the semantics of programs with pointers, and (ii) compositional verification methods, based on the principle of local reasoning, which consists of infering separate specifications of different parts of a program and combining these specifications a posteriori, in a global verification condition.

The search for automated push-button program verification methods motivates the need for understanding the decidability, complexity and expressive power of various dialects thereof, that are used as assertion languages in Hoare-style proofs [9], or logic-based abstract domains in static analysis [4].

Essentially, one can view SL as the first order theory of the heap using quantification over heap locations, to which two non-classical connectives are added: (i) the *separating conjunction* $\phi_1 * \phi_2$, that asserts a split of the heap into disjoint heaps satisfying ϕ_1 and ϕ_2 respectively, and (ii) the *separating implication* or *magic wand* $\phi_1 \mathrel{-\!\!*} \phi_2$, stating that each extension of the heap by a heap satisfying ϕ_1 must satisfy ϕ_2.

Let us consider the following Hoare triple defining the weakest precondition of a selector update in a program handling lists, such as the classical in-place list reversal example [13]:

$$\{\exists x \,.\, \mathsf{i} \mapsto x * (\mathsf{i} \mapsto \mathsf{j} \mathrel{-\!\!*} \phi)\}\ \mathsf{i.next} = \mathsf{j}\ \{\phi\}$$

A typical verification condition asks whether the weakest precondition formula is entailed by another precondition ψ, generated by a program verifier or supplied by the user. The entailment $\psi \rightarrow \exists x \,.\, \mathsf{i} \mapsto x * (\mathsf{i} \mapsto \mathsf{j} \mathrel{-\!\!*} \phi)$ is valid if and only if the formula $\theta = \psi \wedge \forall x \,.\, \neg(\mathsf{i} \mapsto x * (\mathsf{i} \mapsto \mathsf{j} \mathrel{-\!\!*} \phi))$ is unsatisfiable.

Assume now that ϕ and ψ are formulae of the form $Q_1 x_1 \ldots Q_n x_n \,.\, \varphi$, where $Q_1, \ldots, Q_n$ are the first order quantifiers $\exists$ and $\forall$ and φ is quantifier-free. These formulae are said to be in *prenex* form. Because the assertions $\mathsf{i} \mapsto x$ and $\mathsf{i} \mapsto \mathsf{j}$ define *precise* parts of the heap, the quantifiers of ϕ can be hoisted and the entire formula θ can be written in prenex form, following the result of [11, Lemma 3].

Deciding the satisfiability of prenex SL formulae is thus an important ingredient for push-button program verification. In general, unlike first order logic, SL formulae do not have a prenex form because e.g. $\phi * \forall x \,.\, \psi(x) \not\equiv \forall x \,.\, \phi * \psi(x)$ and $\phi \mathrel{-\!\!*} \exists x \,.\, \psi(x) \not\equiv \exists x \,.\, \phi \mathrel{-\!\!*} \psi(x)$. Moreover, it was proved that, for heaps with only one selector, SL is undecidable in the presence of $*$ and $\mathrel{-\!\!*}$ (in fact SL^1 is as expressive as second order logic), whereas the fragment of SL without $\mathrel{-\!\!*}$ is decidable but not elementary recursive [3].

In this paper we answer several open problems, by showing that:

1. the prenex fragment of SL^1 with $*$ and $\mathrel{-\!\!*}$ is decidable but not elementary recursive, and
2. the Bernays-Schönfinkel-Ramsey fragment of SL^1 with $*$ and $\mathrel{-\!\!*}$ is PSPACE-complete.

All results in this paper have been obtained using reductions to and from first order logic with one monadic function symbol, denoted as $[all, (\omega), (1)]_=$ in [2]. The decidability of this fragment is a consequence of the celebrated Rabin Tree Theorem [12], which established the decidability of monadic second order logic of the infinite binary tree (S2S). Furthermore, the $[all, (\omega), (1)]_=$ fragment is shown to be nonelementary, by a direct reduction from domino problems of size equal to a tower of exponentials and, finally, the $[\exists^*\forall^*, (\omega), (1)]_=$ fragment is proved to be Σ_2^P-complete [2].

Essential to our reductions to and from $[all, (\omega), (1)]_=$ is a result stating that each quantifier-free SL^k formula, for $k \geq 1$, is equivalent to a boolean combination of patterns, called *test formulae* [8]. Similar translations exist for quantifier-free SL^1 [10,3] and for SL^1 with one quantified variable [6]. In our previous work [8], we have considered both the finite and infinite satisfiability problems

separately. In this paper we also show that the infinite satisfiability reduces to the finite satisfiability for the prenex fragment of SL^k.

For space reasons, some proofs are given in the extended technical report [7].

2 Preliminaries

In this section, we briefly review some usual definitions and notations. We denote by $\mathbb{Z}$ the set of integers and by $\mathbb{N}$ the set of positive integers including zero. We define $\mathbb{Z}_\infty = \mathbb{Z} \cup \{\infty\}$ and $\mathbb{N}_\infty = \mathbb{N} \cup \{\infty\}$, where for each $n \in \mathbb{Z}$ we have $n + \infty = \infty$ and $n < \infty$. For a countable set S we denote by $\|S\| \in \mathbb{N}_\infty$ the cardinality of S. A decision problem is in $(\mathsf{N})\mathsf{SPACE}(n)$ if it can be decided by a (nondeterministic) Turing machine in space $O(n)$ and in PSPACE if it is in $\mathsf{SPACE}(n^c)$ for some input independent integer $c \geq 1$.

2.1 First Order Logic

Let Var be a countable set of variables, denoted as x, y, z and U be a sort. A *function symbol* f has $\#(f) \geq 0$ arguments of sort U and a sort $\sigma(f)$, which is either the boolean sort Bool or U. If $\#(f) = 0$, we call f a *constant*. We use $\bot$ and $\top$ for the boolean constants false and true, respectively. First-order (FO) terms t and formulae φ are defined by the following grammar:

$$t := x \mid f(t_1, \ldots, t_{\#(f)}) \quad \varphi := \bot \mid \top \mid \varphi_1 \wedge \varphi_2 \mid \neg\varphi_1 \mid \exists x \,.\, \varphi_1 \mid t_1 \approx t_2 \mid p(t_1, \ldots, t_{\#(p)})$$

where $x \in \mathsf{Var}$, f and p are function symbols, $\sigma(f) = U$ and $\sigma(p) = \mathsf{Bool}$. We write $\varphi_1 \vee \varphi_2$ for $\neg(\neg\varphi_1 \wedge \neg\varphi_2)$, $\varphi_1 \rightarrow \varphi_2$ for $\neg\varphi_1 \vee \varphi_2$, $\varphi_1 \leftrightarrow \varphi_2$ for $\varphi_1 \rightarrow \varphi_2 \wedge \varphi_2 \rightarrow \varphi_1$ and $\forall x \,.\, \varphi$ for $\neg\exists x \,.\, \neg\varphi$. The size of a formula φ, denoted as $\mathsf{size}(\varphi)$, is the number of occurrences of symbols needed to write it down. Let $\mathsf{var}(\varphi)$ be the set of variables that occur free in φ, i.e. not in the scope of a quantifier.

First-order formulae are interpreted over FO-structures (called structures, when no confusion arises) $\mathcal{S} = (\mathfrak{U}, \mathfrak{s}, \mathfrak{i})$, where $\mathfrak{U}$ is a countable set, called the *universe*, the elements of which are called *locations*, $\mathfrak{s} : \mathsf{Var} \rightharpoonup \mathfrak{U}$ is a mapping of variables to locations, called a *store* and $\mathfrak{i}$ interprets each function symbol f by a function $f^{\mathfrak{i}} : \mathfrak{U}^{\#(f)} \rightarrow \mathfrak{U}$, if $\sigma(f) = U$ and $f^{\mathfrak{i}} : \mathfrak{U}^{\#(f)} \rightarrow \{\bot^{\mathfrak{i}}, \top^{\mathfrak{i}}\}$ if $\sigma(f) = \mathsf{Bool}$, with $\bot^{\mathfrak{i}} \neq \top^{\mathfrak{i}}$. A structure $(\mathfrak{U}, \mathfrak{s}, \mathfrak{i})$ is *finite* when $\|\mathfrak{U}\| \in \mathbb{N}$ and *infinite* otherwise.

We write $\mathcal{S} \models \varphi$ iff φ is true when interpreted in $\mathcal{S}$. This relation is defined recursively on the structure of φ, as usual. When $\mathcal{S} \models \varphi$, we say that $\mathcal{S}$ is a *model* of φ. A formula is *satisfiable* when it has a model. We write $\varphi_1 \models \varphi_2$ when every model of φ_1 is also a model of φ_2 and by $\varphi_1 \equiv \varphi_2$ we mean $\varphi_1 \models \varphi_2$ and $\varphi_1 \models \varphi_2$.

The *(in)finite satisfiability problem* asks, given a formula φ, whether a (in)finite model exists for this formula.

The Bernays-Schönfinkel-Ramsey fragment of FO [$\mathsf{BSR}(\mathsf{FO})$] is the set of sentences $\exists x_1 \ldots \exists x_n \forall y_1 \ldots \forall y_m \,.\, \varphi$, where φ is a quantifier-free formula in which all function symbols f of arity $\#(f) > 0$ have sort $\sigma(f) = \mathsf{Bool}$.

2.2 Separation Logic

Let $k \in \mathbb{N}$ be a strictly positive integer. The logic SL^k is the set of formulae generated by the grammar:

$$\varphi := \bot \mid \top \mid \mathsf{emp} \mid x \approx y \mid x \mapsto (y_1, \ldots, y_k) \mid \varphi \wedge \varphi \mid \neg\varphi \mid \varphi * \varphi \mid \varphi \mathrel{-\!\!*} \varphi \mid \exists x \,.\, \varphi$$

where $x, y, y_1, \ldots, y_k \in \mathsf{Var}$. The connectives $*$ and $\mathrel{-\!\!*}$ are respectively called the *separating conjunction* and *separating implication* (*magic wand*). We denote by $\mathbf{y}$ the tuple $(y_1, \ldots, y_k) \in \mathsf{Var}^k$. The size of an SL^k formula φ, denoted $\mathsf{size}(\varphi)$, is the number of symbols needed to write it down.

SL^k formulae are interpreted over SL-*structures* (called structures when no confusion arises) $\mathcal{I} = (\mathfrak{U}, \mathfrak{s}, \mathfrak{h})$, where $\mathfrak{U}$ and $\mathfrak{s}$ are as before and $\mathfrak{h} : \mathfrak{U} \rightharpoonup_{fin} \mathfrak{U}^k$ is a finite partial mapping of locations to k-tuples of locations, called a *heap*. As before, a structure $(\mathfrak{U}, \mathfrak{s}, \mathfrak{h})$ is finite when $\|\mathfrak{U}\| \in \mathbb{N}$ and infinite otherwise.

Given a heap $\mathfrak{h}$, we denote by $\mathrm{dom}(\mathfrak{h})$ the domain of the heap, by $\mathrm{img}(\mathfrak{h}) \stackrel{\text{def}}{=} \{\ell_i \mid \exists \ell \in \mathrm{dom}(\mathfrak{h}), \mathfrak{h}(\ell) = (\ell_1, \ldots, \ell_k), i \in [1, k]\}$ and by $\mathrm{elems}(\mathfrak{h}) = \mathrm{dom}(\mathfrak{h}) \cup \mathrm{img}(\mathfrak{h})$ the set of elements either in the domain or the image of the heap. For a store $\mathfrak{s}$, we define $\mathrm{img}(\mathfrak{s}) \stackrel{\text{def}}{=} \{\ell \mid x \in \mathsf{Var},\ \mathfrak{s}(x) = \ell\}$. Two heaps $\mathfrak{h}_1$ and $\mathfrak{h}_2$ are *disjoint* iff $\mathrm{dom}(\mathfrak{h}_1) \cap \mathrm{dom}(\mathfrak{h}_2) = \emptyset$, in which case $\mathfrak{h}_1 \uplus \mathfrak{h}_2$ denotes their union, where $\uplus$ is undefined for non-disjoint heaps. The relation $(\mathfrak{U}, \mathfrak{s}, \mathfrak{h}) \models \varphi$ is defined inductively, as follows:

$$\begin{array}{ll}
(\mathfrak{U}, \mathfrak{s}, \mathfrak{h}) \models \mathsf{emp} & \Leftrightarrow \mathfrak{h} = \emptyset \\
(\mathfrak{U}, \mathfrak{s}, \mathfrak{h}) \models x \mapsto (y_1, \ldots, y_k) & \Leftrightarrow \mathfrak{h} = \{\langle \mathfrak{s}(x), (\mathfrak{s}(y_1), \ldots, \mathfrak{s}(y_k)) \rangle\} \\
(\mathfrak{U}, \mathfrak{s}, \mathfrak{h}) \models \varphi_1 * \varphi_2 & \Leftrightarrow \text{there exist disjoint heaps } \mathfrak{h}_1, \mathfrak{h}_2 \text{ such that } \mathfrak{h} = \mathfrak{h}_1 \uplus \mathfrak{h}_2 \\
 & \quad \text{and } (\mathfrak{U}, \mathfrak{s}, \mathfrak{h}_i) \models \varphi_i, \text{ for } i = 1, 2 \\
(\mathfrak{U}, \mathfrak{s}, \mathfrak{h}) \models \varphi_1 \mathrel{-\!\!*} \varphi_2 & \Leftrightarrow \text{for all heaps } \mathfrak{h}' \text{ disjoint from } \mathfrak{h} \text{ such that } (\mathfrak{U}, \mathfrak{s}, \mathfrak{h}') \models \varphi_1, \\
 & \quad \text{we have } (\mathfrak{U}, \mathfrak{s}, \mathfrak{h}' \uplus \mathfrak{h}) \models \varphi_2
\end{array}$$

The semantics of equality, boolean and first-order connectives is the usual one. Satisfiability, entailment and equivalence are defined for SL^k as for FO formulae. The (in)finite satisfiability problem for SL^k asks whether a (in)finite model exists for a given formula. We write $\phi \equiv^{fin} \psi$ [$\phi \equiv^{inf} \psi$] whenever $(\mathfrak{U}, \mathfrak{s}, \mathfrak{h}) \models \phi \Leftrightarrow (\mathfrak{U}, \mathfrak{s}, \mathfrak{h}) \models \psi$ for every finite [infinite] structure $(\mathfrak{U}, \mathfrak{s}, \mathfrak{h})$.

The prenex fragment of SL^k [$\mathsf{PRE}(\mathsf{SL}^k)$] is the set of sentences $Q_1 x_1 \ldots Q_n x_n \,.\, \phi$, where $Q_1, \ldots, Q_n \in \{\exists, \forall\}$ and ϕ is a quantifier-free SL^k formula. Unlike FO, where each formula is equivalent to a linear-size formula in prenex form, there are SL^k formulae that do not have a prenex form equivalent. For instance, $\phi * \forall x \,.\, \psi(x) \not\equiv \forall x \,.\, \phi * \psi(x)$ and dually, $\phi \mathrel{-\!\!*} \exists x \,.\, \psi(x) \not\equiv \exists x \,.\, \phi \mathrel{-\!\!*} \psi(x)$, where ϕ and ψ are arbitrary SL^k formulae.

The Bernays-Schönfinkel-Ramsey fragment of SL^k [$\mathsf{BSR}(\mathsf{SL}^k)$] is the set of sentences $\exists x_1 \ldots \exists x_n \forall y_1 \ldots \forall y_m \,.\, \phi$, where ϕ is a quantifier-free SL^k formula. Since there are no function symbols of arity greater than zero in SL^k, there are no restrictions, other than the form of the quantifier prefix, defining $\mathsf{BSR}(\mathsf{SL}^k)$.

2.3 Test Formulae for SL^k

This section contains a number of definitions and results from [8], needed for self-containment. For more details, the interested reader is pointed towards [8].

Definition 1. *The following patterns are called* test formulae*:*

$$x \hookrightarrow \mathbf{y} \stackrel{\text{def}}{=} x \mapsto \mathbf{y} * \top \qquad |U| \geq n \stackrel{\text{def}}{=} \top \mathrel{-\!\!\circ} |h| \geq n,\ n \in \mathbb{N}$$

$$\mathsf{alloc}(x) \stackrel{\text{def}}{=} x \mapsto \underbrace{(x, \ldots, x)}_{k \text{ times}} \mathrel{-\!\!*} \bot \qquad |h| \geq |U| - n \stackrel{\text{def}}{=} |h| \geq n+1 \mathrel{-\!\!*} \bot, n \in \mathbb{N}$$

$$|h| \geq n \stackrel{\text{def}}{=} \begin{cases} |h| \geq n-1 * \neg\mathsf{emp}, & \text{if } n > 0 \\ \top, & \text{if } n = 0 \\ \bot, & \text{if } n = \infty \end{cases}$$

and $x \approx y$, where $x, y \in \mathsf{Var}$, $\mathbf{y} \in \mathsf{Var}^k$ and $n \in \mathbb{N}_\infty$ is a positive integer or ∞.

The test formulae of the form $|U| \geq n$ and $|h| \geq |U| - n$ are called *domain dependent* and the rest *domain independent*. A *literal* is a test formula or its negation.

The semantics of test formulae is intuitive: $x \hookrightarrow \mathbf{y}$ holds when x denotes a location and $\mathbf{y}$ is the image of that location in the heap, $\mathsf{alloc}(x)$ holds when x denotes a location in the domain of the heap (allocated), $|h| \geq n$, $|U| \geq n$ and $|h| \geq |U| - n$ are cardinality constraints involving the size of the heap, denoted $|h|$ and that of the universe, denoted $|U|$. We recall that $|h|$ ranges over $\mathbb{N}$, whereas $|U|$ is always interpreted as a number larger than $|h|$ and possibly infinite.

Observe that not all atoms of SL^k are test formulae, for instance $x \mapsto \mathbf{y}$ and emp are not test formulae. However, we have the equivalences $x \mapsto \mathbf{y} \equiv x \hookrightarrow \mathbf{y} \wedge \neg|h| \geq 2$ and $\mathsf{emp} \equiv \neg|h| \geq 1$. Moreover, for any $n \in \mathbb{N}$, the test formulae $|U| \geq n$ and $|h| \geq |U| - n$ become trivially true and false, respectively, if we consider the universe to be infinite.

The following result establishes a translation of quantifier-free SL^k formulae into boolean combinations of test formulae. This translation relies on the notion of a minterm.

Definition 2. *A* minterm *M is a set (conjunction) of literals containing:*

- *exactly one literal $|h| \geq \mathrm{hmin}_M$ and one literal $|h| < \mathrm{hmax}_M$, where $\mathrm{hmin}_M \in \mathbb{N} \cup \{|U| - n \mid n \in \mathbb{N}\}$ and $\mathrm{hmax}_M \in \mathbb{N}_\infty \cup \{|U| - n \mid n \in \mathbb{N}\}$, and*
- *exactly one literal of the form $|U| \geq n$ and at most one literal of the form $|U| < n$.*

One of the results in [8] is that, for each quantifier-free SL^k formula ϕ, it is possible to define a disjunction on minterms that preserves the finite models of ϕ. We denote the set of minterms in the disjunction as $\mu^{fin}(\phi)$, where $\mu^{fin}(.)$ is an effectively computable function, defined recursively on the structure of ϕ.

Lemma 1. *Given a quantifier-free SL^k formula ϕ, $\mu^{fin}(\phi)$ is a finite set of minterms and we have $\phi \equiv^{fin} \bigvee_{M \in \mu^{fin}(\phi)} M$.*

Proof. See [8, Lemma 5]. □

Given a quantifier-free SL^k formula ϕ, the number of minterms occurring in $\mu^{fin}(\phi)$ is exponential in the size of ϕ, in the worst case. Therefore, an optimal decision procedure cannot generate and store these sets explicitly, but rather must enumerate minterms lazily. The next lemma shows that it is possible to check whether $M \in \mu^{fin}(\phi)$ using space bounded by a polynomial in $\mathsf{size}(\phi)$. For a boolean combination of test formulae ϕ, we denote by $\mathcal{N}(\phi)$ the maximum $n \in \mathbb{N}$ that occurs in an atom of the form $|h| \geq n$ or $|U| \geq n$ in ϕ.

Lemma 2. *For every SL^k formula ϕ, the size of every minterm $\mu^{fin}(\phi)$ is polynomial w.r.t. $\mathsf{size}(\phi)$. In particular, $\max_{M \in \mu^{fin}(\phi)} \mathcal{N}(M)$ is polynomial w.r.t. $\mathsf{size}(\phi)$. Furthermore, given a minterm M, the problem of checking whether $M \in \mu^{fin}(\phi)$ is in PSPACE.*

Proof. See [8, Lemma 8 and Corollary 1]. □

3 The $\mathsf{PRE}(\mathsf{SL}^1)$ Fragment is Decidable

The first result of this paper is the decidability of the prenex fragment of SL^1. In particular, this shows that $\mathsf{PRE}(\mathsf{SL}^k)$ is strictly less expressive than SL^k, because SL^1 has been shown to be at least as expressive as Second Order Logic, thus having an undecidable satisfiability problem [3, Theorem 6.11].

3.1 From Infinite to Finite Satisfiability

We begin by showing that the infinite satisfiability problem can be reduced to the finite satisfiability problem for prenex SL-formulae. The intuition is that two SL-structures defined on the same heap and store can be considered equivalent if both have enough locations outside of the heap.

Definition 3. *Let X be a set of variables and let $n \in \mathbb{N}$. Two* SL*-structures $\mathcal{I} = (\mathfrak{U}, \mathfrak{s}, \mathfrak{h})$ and $\mathcal{I}' = (\mathfrak{U}', \mathfrak{s}', \mathfrak{h}')$ are* (X,n)-similar *(written $\mathcal{I} \sim^n_X \mathcal{I}'$) iff the following conditions hold:*

1. $\mathfrak{h} = \mathfrak{h}'$.
2. *For all $x, y \in X$, $\mathfrak{s}(x) = \mathfrak{s}(y) \Leftrightarrow \mathfrak{s}'(x) = \mathfrak{s}'(y)$.*
3. *For every $x \in X$, if $\mathfrak{s}(x) \in \mathrm{elems}(\mathfrak{h})$ or $\mathfrak{s}'(x) \in \mathrm{elems}(\mathfrak{h})$ then $\mathfrak{s}(x) = \mathfrak{s}'(x)$.*
4. *$\|\mathfrak{U} \setminus \mathrm{elems}(\mathfrak{h})\| \geq n + \|X\|$ and $\|\mathfrak{U}' \setminus \mathrm{elems}(\mathfrak{h})\| \geq n + \|X\|$.*

Note that Condition 1 entails that $\mathrm{elems}(\mathfrak{h}) \subseteq \mathfrak{U} \cap \mathfrak{U}'$. Next, we prove that any two SL-structures that are $(\mathsf{var}(\phi), m)$-similar are also indistinguishable by any formula ϕ prefixed by m quantifiers.

Proposition 1. *Let $\phi = \mathrm{Q}_1 x_1 \ldots \mathrm{Q}_m x_m \,.\, \psi$ be a prenex SL^k formula, with $\mathrm{Q}_i \in \{\forall, \exists\}$ for $i = 1, \ldots, m$. Assume that ψ is a quantifier-free boolean combination of domain independent test formulae. If $\mathcal{I} \sim^m_{fv(\phi)} \mathcal{I}'$ and $\mathcal{I} \models \phi$ then $\mathcal{I}' \models \phi$.*

The formulas $x \in h$ and $\mathsf{distinct}(x_1, \ldots, x_n)$ are shorthands for the formulas $\exists y_0, y_1, \ldots y_k \,.\, (y_0 \hookrightarrow (y_1, \ldots, y_k) \wedge \bigvee_{i=0}^{k} x \approx y_i)$ and $\bigwedge_{i=1}^{n} \bigwedge_{j=1}^{i-1} \neg(x_i \approx x_j)$, respectively. We define the formula:

$$\lambda_p \stackrel{\text{def}}{=} \exists x_1, \ldots, x_p \,.\, (\mathsf{distinct}(x_1, \ldots, x_p) \wedge \bigwedge_{i=1}^{p} \neg x_i \in h)$$

It is clear that $(\mathfrak{U}, \mathfrak{s}, \mathfrak{h}) \models \lambda_p$ iff $\|\mathfrak{U} \setminus \mathrm{elems}(\mathfrak{h})\| \geq p$. In particular, λ_p is always true on infinite domains. Observe, moreover, that λ_p belongs to the $\mathsf{PRE}(\mathsf{SL}^k)$ fragment, for any $p \geq 2$ and any $k \geq 1$.

The following lemma reduces the infinite satisfiability problem to the finite version of it. This is done by adding an axiom ensuring that there are enough locations outside of the heap. Note that there is no need to consider test formulae of the form $|U| \geq n$ and $|h| \geq |U| - n$ because they alway evaluate to true and, respectively, false, on infinite SL-structures.

Lemma 3. *Let $\phi = \mathrm{Q}_1 x_1 \ldots \mathrm{Q}_m x_m \,.\, \psi$ be a prenex SL^k formula, where $\mathrm{Q}_i \in \{\forall, \exists\}$ for $i = 1, \ldots, m$ and $fv(\phi) = \emptyset$. Assume that ψ is a boolean combination of test formulas of the form $x \approx y$ or $x \hookrightarrow (y_1, \ldots, y_k)$ or $\mathsf{alloc}(x)$ or $|h| \geq n$. The two following assertions are equivalent.*

1. *ϕ admits an infinite model.*
2. *$\phi \wedge \lambda_m$ admits a finite model.*

Proof. (1) $\Rightarrow$ (2): Assume that ϕ admits an infinite model $(\mathfrak{U}, \mathfrak{s}, \mathfrak{h})$. Let $\mathfrak{U}'$ be a finite subset of $\mathfrak{U}$ including $\mathrm{elems}(\mathfrak{h})$ plus m additional elements. It is clear that $(\mathfrak{U}, \mathfrak{s}, \mathfrak{h}) \sim^m_{\emptyset} (\mathfrak{U}', \mathfrak{s}, \mathfrak{h})$. Indeed, Condition 1 holds since the two structures share the

same heap, Conditions 2 and 3 trivially hold since the considered set of variables is empty, and Condition 4 holds since $\mathfrak{U}$ is infinite and the additional elements in $\mathfrak{U}'$ do not occur in elems($\mathfrak{h}$). Thus $(\mathfrak{U}', \mathfrak{s}, \mathfrak{h}) \models \phi$ by Proposition 1. Furthermore, $(\mathfrak{U}', \mathfrak{s}, \mathfrak{h}) \models \lambda_m$, by definition of $\mathfrak{U}'$.

(2) $\Rightarrow$ (1): Assume that $\phi \wedge \lambda_m$ has a finite model $(\mathfrak{U}, \mathfrak{s}, \mathfrak{h})$. Let $\mathfrak{U}'$ be any infinite set containing $\mathfrak{U}$. Again, we have $(\mathfrak{U}, \mathfrak{s}, \mathfrak{h}) \sim^m_{\emptyset} (\mathfrak{U}', \mathfrak{s}, \mathfrak{h})$. As in the previous case, Conditions 1, 2 and 3 trivially hold, and Condition 4 holds since $\mathfrak{U}'$ is infinite and $(\mathfrak{U}, \mathfrak{s}, \mathfrak{h}) \models \lambda_m$. By Proposition 1, we deduce that $(\mathfrak{U}', \mathfrak{s}, \mathfrak{h}) \models \phi$. □

3.2 Translating PRE(SL^1) into First-Order Logic

After reduction of the infinite to the finite satisfiability problem, the decidability of the latter for PRE(SL^1) is established by reduction to the finite satisfiability of the $[all, (\omega), (1)]_=$ fragment of FO, with an arbitrary number of monadic boolean function symbols and one function symbol f of sort $\sigma(f) = U$. The decidability of this fragment is a consequence of the celebrated Rabin's Tree Theorem, which established the decidability of the monadic theory of the infinite binary tree [12].

In the following, we define an equivalence-preserving (on finite structures) translation of SL^k into FO. Let $\mathfrak{d}$ be a unary predicate symbol and let $\mathfrak{f}_i$ (for $i = 1, \ldots, k$) be unary function symbols. We define the following transformation from quantified boolean combinations of test formulae into first order formulae:

$$
\begin{aligned}
\Theta(x \approx y) &\stackrel{\text{def}}{=} x \approx y \\
\Theta(x \hookrightarrow (y_1, \ldots, y_k)) &\stackrel{\text{def}}{=} \mathfrak{d}(x) \wedge \bigwedge_{i=1}^{k} y_i \approx \mathfrak{f}_i(x) \\
\Theta(\mathsf{alloc}(x)) &\stackrel{\text{def}}{=} \mathfrak{d}(x) \\
\Theta(\neg\phi) &\stackrel{\text{def}}{=} \neg\Theta(\phi) \\
\Theta(\phi_1 \bullet \phi_2) &\stackrel{\text{def}}{=} \Theta(\phi_1) \bullet \Theta(\phi_2) \qquad \text{if } \bullet \in \{\wedge, \vee, \rightarrow, \leftrightarrow\} \\
\Theta(\mathrm{Q}x \,.\, \phi) &\stackrel{\text{def}}{=} \mathrm{Q}x \,.\, \Theta(\phi) \qquad \text{if } \mathrm{Q} \in \{\exists, \forall\} \\
\Theta(|U| \geq n) &\stackrel{\text{def}}{=} \exists x_1, \ldots, x_n \,.\, \mathsf{distinct}(x_1, \ldots, x_n) \\
\Theta(|h| \geq n) &\stackrel{\text{def}}{=} \exists x_1, \ldots, x_n \,.\, \mathsf{distinct}(x_1, \ldots, x_n) \wedge \bigwedge_{i=1}^{n} \mathfrak{d}(x_i) \\
\Theta(|h| \geq |U| - n) &\stackrel{\text{def}}{=} \exists x_1, \ldots, x_n \forall y \,.\, \bigwedge_{i=1}^{n} y \not\approx x_i \rightarrow \mathfrak{d}(y)
\end{aligned}
$$

Proposition 2. *Let ϕ be a quantified boolean combination of test formulae. The formula ϕ has a finite* SL *model iff $\Theta(\phi)$ has a finite* FO *model.*

Proof. A FO-structure $\mathcal{I} = (\mathfrak{U}, \mathfrak{s}, \mathfrak{i})$ on the signature $\mathfrak{d}$, $\mathfrak{f}_1, \ldots, \mathfrak{f}_k$ *corresponds* to an SL-structure $\mathcal{I}' = (\mathfrak{U}', \mathfrak{s}', \mathfrak{h})$ iff $\mathfrak{U} = \mathfrak{U}'$, $\mathfrak{s} = \mathfrak{s}'$, $\mathfrak{d}^{\mathfrak{i}} = \mathrm{dom}(\mathfrak{h})$ and for every $j \in [\![1 \,..\, k]\!]$, $\mathfrak{f}^{\mathfrak{i}}_j(x) = y_j$ if $\mathfrak{h}(x) = (y_1, \ldots, y_k)$. It is clear that for every finite first-order structure $\mathcal{I}$ there exists a finite SL-structure $\mathcal{I}'$ such that $\mathcal{I}$ corresponds to $\mathcal{I}'$

and vice-versa. Furthermore, if $\mathcal{I}$ corresponds to $\mathcal{I}'$ then it is straightforward to check that $\mathcal{I}' \models \phi \Leftrightarrow \mathcal{I} \models \Theta(\phi)$. □

Given a formula $\psi = Q_1 x_1 \ldots Q_n x_n \,.\, \phi$ of $\mathsf{PRE}(\mathsf{SL}^1)$, where ϕ is a quantifier-free SL^1 formula, consider the expansion of ϕ as a disjunction of minterms $\mu \stackrel{\text{def}}{=} \bigvee_{M \in \mu^{fin}(\phi)} M$. By Lemma 1, we have $\phi \equiv^{fin} \mu$, thus $\psi \equiv^{fin} Q_1 x_1 \ldots Q_n x_n \,.\, \mu$. By Proposition 2, ψ has a finite SL model iff $\Theta(Q_n x_n \,.\, \mu)$ has a finite FO model. Moreover, it is easy to see that $\Theta(Q_n x_n \,.\, \mu)$ belongs to the $[all, (\omega), (1)]_=$ fragment of FO, whose finite satisfiability problem is decidable [2, Corollary 7.2.12]. The following theorem summarizes the result:

Theorem 1. *The finite and infinite satisfiability problems are decidable for* $\mathsf{PRE}(\mathsf{SL}^1)$.

4 The $\mathsf{PRE}(\mathsf{SL}^1)$ Fragment is not Elementary Recursive

This section is concerned with the computational complexity of the (in)finite satisfiability problem(s) for the $\mathsf{PRE}(\mathsf{SL}^1)$ fragment. We use the fact that the $[all, (\omega), (1)]_=$ fragment of FO is nonelementary and obtain a similar lower bound by an opposite reduction, from the satisfiabilty of $[all, (\omega), (1)]_=$ to that of $\mathsf{PRE}(\mathsf{SL}^1)$. This reduction, in the finite and infinite case, respectivelly, is carried out by the following propositions:

Proposition 3. *There is a polynomial reduction of the finite satisfiability problem for* FO *formulae with one monadic function symbol to the finite satisfiability problem for* $\mathsf{PRE}(\mathsf{SL}^1)$ *formulae.*

Proof. The reduction is immediate: it suffices to add the axiom: $\forall x \,.\, \mathsf{alloc}(x)$ (i.e., the heap is total) and replace all equations of the form $f(x) \approx y$ by $x \hookrightarrow y$ (by flattening we may assume that all the equations occurring in the formula are of the form $f(x) \approx y$ or $x \approx y$, where x, y are variables). It is straightforward to check that satisfiability is preserved. □

Proposition 4. *There is a polynomial reduction of the finite satisfiability problem for* FO *formulae with one monadic function symbol to the infinite satisfiability problem for* $\mathsf{PRE}(\mathsf{SL}^1)$ *formulae.*

Proof. We may apply the same transformation as above on equations $f(x) \approx y$, but this time the axiom $\forall x \,.\, \mathsf{alloc}(x)$ cannot be added as it would make the resulting formula unsatisfiable. Instead, we add the axiom $\neg\mathsf{emp} \wedge \forall x, y \,.\, x \hookrightarrow y \rightarrow \mathsf{alloc}(y)$, and we replace every quantification $\forall x \,.\, \phi$ (resp. $\exists x \,.\, \phi$) by a quantification over the domain of the heap: $\forall x \,.\, \mathsf{alloc}(x) \rightarrow \phi$ (resp. $\exists x \,.\, \mathsf{alloc}(x) \wedge \phi$). Again, it is straightforward to check that satisfiability is preserved. Note that

infinite satisfiability is equivalent to finite satisfiability here since the quantifications range over elements occurring in the heap. The domain of the (finite) first-order interpretation is encoded as the domain of the heap. □

The main difficulty here is the lack of a direct result stating that the *finite* satisfiability problem for $[all,(\omega),(1)]_{=}$ is nonelementary. Instead the result of [2, Theorem 7.2.15] considers arbitrary FO structures, in which the cardinality of the universe is not necessarily finite. In the following we show that this result can be strenghtened to considering finite structures only. Observe that this is not automatically the case for FO formulae with one monadic function symbol, for instance, the formula $\exists x \forall y \,.\, x \not\approx f(y) \wedge \forall y, z, \,.\, f(y) \approx f(z) \rightarrow y \approx z$ is satisfiable only on infinite FO structures. However, this is the case for the formula obtained in [2, Theorem 7.2.15] by reduction from domino the problem of nonelementary size, defined below:

Definition 4. *A* domino system *is a tuple $\mathcal{D} = (D, H, V)$, where D is a finite set of tiles and $H, V \subseteq D \times D$. For some $t \geq 2$, let $Z_t \times Z_t$ be a* torus*, where $Z_t = ([0, t-1], succ)$ and $succ(n) \stackrel{\text{def}}{=} (n+1) \mod t$, for all $n \in [0, t-1]$. We say that $\mathcal{D}$* tiles *$Z_t \times Z_t$ with initial condition $d_0 \ldots d_{m-1} \in D^m$ iff there exists a mapping $\tau : Z_t \times Z_t \rightarrow D$ such that, for all $(x, y) \in Z_t \times Z_t$, we have $H(\tau(x, y), \tau(succ(x), y))$ and $V(\tau(x, y), \tau(x, succ(y)))$, and moreover $\tau(i, 0) = d_i$, for all $i \in [0, m-1]$.*

Given a tower of exponentials $\mathrm{T}(n) = \underbrace{2^{2^{\cdots^2}}}_{n}$, the existence of a tiling of $Z_{\mathrm{T}(n)} \times Z_{\mathrm{T}(n)}$ with a given initial condition is a nonelementary recursive problem [2, Theorem 6.1.2]. For the sake of self-containment, we describe the main ingredients of the reduction from this problem to the satisfiability of $[all,(\omega),(1)]_{=}$ on arbitrary FO-structures.

Suppose that $D = \{d_1, \ldots, d_r\}$. First, we express the tiling conditions (Definition 4) by a formula θ, using $r+1$ binary boolean functions $P_0, \ldots, P_r$, where:

1. $P_0(x, y)$ encodes the successor relation $succ(x) = y$,
2. $P_i(x, y)$ holds iff $\tau(x, y) = d_i$, for all $i \in [1, r]$,
3. the horizontal and vertical adjacency conditions H and V are respected, and
4. there is an element x_0 such that the points $(x_0, succ^i(x_0))$ are labeled with w_i, for all $i \in [0, m-1]$.

Next, we assume that the FO-structures encoding the tiling are models of the formula $\alpha \stackrel{\text{def}}{=} \exists x \forall y \,.\, f(x) \approx x \wedge f^{n+1}(y) \approx x$, which states that the domain can be viewed as a tree of height at most $n+1$, where the (necessarily unique) element assigned to the variable x is the root of the tree, and where f maps every other node to its parent.

Intuitively, the domain $[0, \mathrm{T}(n) - 1]$ will be represented by the direct sons of the root. The main problem is ensuring that the universe $Z_{\mathrm{T}(n)}$ has size (at most) $\mathrm{T}(n)$. To this end we define inductively the equivalence relations $E_0, \ldots, E_n$ as:

1. all nodes are E_0-equivalent, and
2. for $m \geq 1$, two nodes are E_m-equivalent if for every E_{m-1}-equivalence class K, either both nodes have no child in K or both nodes have a child in K.

Then, in each model of α, there are at most $\mathrm{T}(m)$ E_m-equivalence classes, for each $m \geq 0$. If we denote by N_m the index of E_m, then we have $N_m \leq 2^{N_{m-1}}$, because each equivalence class E_m is uniquely identified by a binary sequence associating either 0 (no children in K) or 1 (at least one child) to each equivalence class of E_{m-1}. Clearly, there are at most $2^{N_{m-1}}$ such sequences. Moreover, we have $E_{m-1} \subseteq E_m$, for all $m \geq 1$, therefore $E_n = \cap_{i=0}^{n} E_i$.

We consider formulae $\beta_m(x, y)$ stating that x and y have height at most m and are E_m-equivalent and a formula $\delta(x)$, stating that x is a child of the root (asserted by α) with at most one child in each E_{n-1} equivalence class. Then let $\gamma \stackrel{\mathsf{def}}{=} \forall x, y \,.\, \delta(x) \wedge \delta(y) \wedge \beta_n(x, y) \rightarrow x \approx y$. In any model $(\mathfrak{U}, \mathfrak{s}, \mathfrak{i}) \models \alpha \wedge \gamma$ there are at most $\mathrm{T}(n)$ elements a such that $(\mathfrak{U}, \mathfrak{s}[x \leftarrow a], \mathfrak{i}) \models \delta(x)$, because there is at most one element in each E_n-equivalence class and there are at most $\mathrm{T}(n)$ such classes.

It remains to encode the fact that an element $(x, y) \in Z_{\mathrm{T}}(n) \times Z_{\mathrm{T}}(n)$ is labeled by the tile d_i, i.e. that $P_i(x, y)$ holds in any model of θ. Since we assumed that $\delta(x) \wedge \delta(y)$ holds, x and y have at most one child in each E_{n-1} equivalence class, thus each element can be distinguished by the tuple $(n_1, \ldots, n_s)$ of numbers of children in each E_{n-1} equivalence class $K_1, \ldots, K_s$. We encode $P_i(x, y)$ by assuming the existence of a node z with $g_i(j, k) = 2 + 4i + 2j + k$ children in each class $K_1, \ldots, K_s$. This is encoded by a formula $\pi_i(x, y)$.

Finally, the $[all, (\omega), (1)]_{=}$ formula that states the existence of a tiling of $Z_{\mathrm{T}}(n) \times Z_{\mathrm{T}}(n)$ is obtained from θ by replacing each quantifier $\exists x \,.\, \phi$ by $\exists x \,.\, \delta(x) \wedge \phi$ and $\forall x \,.\, \phi$ by $\forall x \,.\, \delta(x) \rightarrow \phi$ and each occurrence of a predicate symbol $P_i(x, y)$ by $\pi_i(x, y)$.

Lemma 4. *The finite satisfiability problem is not elementary recursive for first order formulae built on a signature containing only one function symbol of arity* 1 *and the equality predicate.*

Proof. Let φ be the formula encoding the existence of a tiling of $Z_{\mathrm{T}}(n) \times Z_{\mathrm{T}}(n)$ by a tiling system $\mathcal{D} = (D, H, V)$ and $\mathcal{I} = (\mathfrak{U}, \mathfrak{s}, \mathfrak{i})$ be a model of φ, with $\mathfrak{f} = f^{\mathfrak{i}}$. We denote by $\mathfrak{r}$ the root of the tree, i.e., the unique element of $\mathfrak{U}$ with $(\mathfrak{U}, \mathfrak{s}[x \mapsto \mathfrak{r}], \mathfrak{i}) \models \forall y \,.\, f(x) \approx x \wedge f^{n+1}(y) \approx x$. Given $i \in [0, r]$ and $a, b \in \mathfrak{U}$, if $(\mathfrak{U}, \mathfrak{s}[x \mapsto a, y \mapsto b], \mathfrak{i}) \models \pi_i(x, y)$, then we denote by $\mu(i, a, b)$ a set containing an arbitrarily chosen element z satisfying $\mathcal{P}(i, a, b)$ in the definition of $\pi_i(x, y)$ along with all

the children of z, otherwise $\mu(i,a,b)$ is empty. Observe that $\mu(i,a,b)$ is always finite because the number of children of z in each equivalence class is bounded by $g_i(1,1) = 2+4\times i+2+1 \leq 2+4\times r+2+1$, moreover the number of E_n-equivalence classes is finite.

We show that φ admits a finite model $\mathcal{I}'$. The set B of elements b such that $(\mathfrak{U}, \mathfrak{s}[x \mapsto b], \mathfrak{i}) \models \delta(x)$ is finite. Let $\Pi \stackrel{\text{def}}{=} \bigcup\{\mu(i,a,b) \mid a,b \in B, i \in [0,r]\}$. Since B is finite and every set $\mu(i,a,b)$ is finite, Π is also finite. With each element $a \in \mathfrak{U}$ and each E_n-equivalence class K, we associate a set $\nu(a,K)$ containing exactly one child of a in K if such a child exists, otherwise $\nu(a,K)$ is empty. We now consider the subset $\mathfrak{U}'$ of $\mathfrak{U}$ defined as the set of elements a such that for every $m \in \mathbb{N}$, $\mathfrak{f}^m(a)$ occurs either in $\{\mathfrak{r}\} \cup B \cup \Pi$ or in a set $\nu(b,K)$, where $b \in \mathfrak{U}$ and K is an E-equivalence class. Note that $\mathfrak{r} \in \mathfrak{U}'$ and that if $a \in \mathfrak{U}'$ then necessarily $\mathfrak{f}(a) \in \mathfrak{U}'$. Furthermore, if $\mathfrak{f}(b) \in \mathfrak{U}'$ and $b \in \nu(\mathfrak{f}(b),K)$ then $b \in \mathfrak{U}'$.

It is easy to check that $\mathfrak{U}'$ is finite. Indeed, since $(\mathfrak{U}, \mathfrak{s}, \mathfrak{i}) \models \alpha$ and no new node or edge is added, all nodes are of height less or equal to $n+1$. Furthermore, all nodes have at most $\|B\| + \|\Pi\| + \#K$ children in $\mathfrak{U}'$, where $\#K$ denotes the number of E_n-equivalence classes.

We denote by $\mathcal{I}' = (\mathfrak{U}', \mathfrak{s}, \mathfrak{i}')$ the restriction of $\mathcal{I}$ to the elements of $\mathfrak{U}'$ (we may assume that $\mathfrak{s}$ is a store on $\mathfrak{U}'$ since φ is closed). We prove that $\mathcal{I}' \models \varphi$.

- Since $\mathfrak{U}'$ contains the root, and $\mathcal{I} \models \alpha$, we must have $\mathcal{I}' \models \alpha$.
- Observe that $\mathfrak{U}'$ necessarily contains $\nu(b,K)$, for every $b \in \mathfrak{U}'$, since by definition the parent of the (unique) element of $\nu(b,K)$ is b. Thus at least one child of b is kept in each equivalence class. Thus the relations E_m on elements of $\mathfrak{U}'$ are preserved in the transformation: for every $a,b \in \mathfrak{U}'$, a,b are E_m-equivalent in the structure $\mathcal{I}$ iff they are equivalent in the structure $\mathcal{I}'$. Further, the height of the nodes cannot change. Therefore, for every $a,a' \in U'$:

  $$(\mathfrak{U}', \mathfrak{s}[x \mapsto a, y \mapsto a'], \mathfrak{i}') \models \beta_n(x,y) \text{ iff } (\mathfrak{U}, \mathfrak{s}[x \mapsto a, y \mapsto a'], \mathfrak{i}) \models \beta_n(x,y)$$

  By definition, for every $a \in B$ and $m \in \mathbb{N}$, $\mathfrak{f}^m(a) \in \{a, \mathfrak{r}\}$, thus $B \subseteq \mathfrak{U}'$. Because no new edges are added, we deduce:

  $$(\mathfrak{U}', \mathfrak{s}[x \mapsto a], \mathfrak{i}') \models \delta(x) \Leftrightarrow (\mathfrak{U}, \mathfrak{s}[x \mapsto a], \mathfrak{i}) \models \delta(x) \Leftrightarrow a \in B$$

  Consequently, since $\mathcal{I} \models \gamma$, we have $\mathcal{I}' \models \gamma$.
- All elements in $\mu(i,a,a')$ with $a,a' \in B$ occur in $\mathfrak{U}'$ (because if $b \in \mu(i,a,a')$ and $m \in \mathbb{N}$ then $\mathfrak{f}^m(b) \in \{\mathfrak{r}\} \cup B \cup \mu(i,a,a')$), thus, for all $a,a' \in B$:

  $$(\mathfrak{U}', \mathfrak{s}[x \mapsto a, y \mapsto a'], \mathfrak{i}') \models \pi_i(x,y) \Leftrightarrow (\mathfrak{U}, \mathfrak{s}[x \mapsto a, y \mapsto a'], \mathfrak{i}) \models \pi_i(x,y)$$

  Since all quantifications in η' range over elements in B, we deduce, by a straightforward induction on the formula, that $\mathcal{I}$ and $\mathcal{I}'$ necessarily agree

on the formula $\eta'[D(x)/\delta(x), P_i(x,y)/\pi_i(x,y)]$. Consequently, we must have $\mathcal{I}' \models \eta'[D(x)/\delta(x), P_i(x,y)/\pi_i(x,y)]$. □

Theorem 2. *The finite and infinite satisfiability problems are not elementary recursive for prenex formulae of* SL^1.

5 The $\mathsf{BSR}(\mathsf{SL}^1)$ Fragment is PSPACE-complete

The last result concerns the tight complexity of the $\mathsf{BSR}(\mathsf{SL}^1)$ fragment. For $k \geq 2$, we showed that $\mathsf{BSR}(\mathsf{SL}^k)$ is undecidable, in general, and PSPACE-complete if the positive occurrences of the magic wand are forbidden[3]. Here we answer the problem concerning the exact complexity of $\mathsf{BSR}(\mathsf{SL}^1)$, by showing its PSPACE-completeness.

Let $\mathcal{I} = (\mathfrak{U}, \mathfrak{s}, \mathfrak{h})$ be a structure, X a non-empty set of variables and $L \subseteq \mathrm{dom}(\mathfrak{h})$ be a set of locations. We define:

$$
\begin{array}{l}
V_{X,L} \stackrel{\text{def}}{=} L \cup \mathfrak{s}(X) \\
\overline{V}_{X,L} \stackrel{\text{def}}{=} \left\{ \ell \in \mathfrak{U} \mid \exists i \geq 0 \; \exists \ell' \in V_{X,L} \,.\, \mathfrak{h}^i(\ell') = \ell \right\} \\
W_{X,L} \stackrel{\text{def}}{=} V_{X,L} \cup \left\{ \ell \in \overline{V}_{X,L} \mid \exists \ell', \ell'' \in \overline{V}_{X,L} \,.\, \ell' \neq \ell'' \wedge \mathfrak{h}(\ell') = \mathfrak{h}(\ell'') = \ell \right\}
\end{array}
$$

Intuitively, $\overline{V}_{X,L}$ contains all locations reachable via the heap from a location either in L or labelled with a variable from X and $W_{X,L}$ contains all locations from $V_{X,L}$ and those from $\overline{V}_{X,L}$ that have two or more predecessors via the heap.

Given a location $\ell_0 \in \mathrm{dom}(\mathfrak{h})$, the *segment* $S(\ell_0) = \langle \ell_0, \ell_1, \ldots, \ell_n \rangle$, for some $n \geq 0$, is the unique sequence of locations such that $\ell_1, \ldots, \ell_n \in \mathrm{dom}(\mathfrak{h}) \setminus W_{X,L}$, $\mathfrak{h}(\ell_i) = \ell_{i+1}$ for all $i \in [0, n-1]$ and either $\mathfrak{h}(\ell_n) \in W_{X,L}$ or $\mathfrak{h}^2(\ell_n) = \bot$. Note that because the domain of $\mathfrak{h}$ is necessarily finite, such a sequence is well defined. We denote by $|S(\ell_0)| = n+1$ the number of locations in the segment. For an integer $N \geq 0$, we denote by $S^N(\ell_0)$ the restriction of $S(\ell_0)$ to its first $\min(|S(\ell_0)| - 1, N) + 1$ elements. We sometimes blur the distinction between a segment and the set of its elements and write $\ell \in S(\ell_0)$ iff ℓ is one of the elements of $S(\ell_0)$.

Given a structure $\mathcal{I} = (\mathfrak{U}, \mathfrak{s}, \mathfrak{h})$, the (N, X, L)-*contraction of* $\mathcal{I}$ is the structure $C^N_{X,L}(\mathcal{I}) = (\mathfrak{U}', \mathfrak{s}, \mathfrak{h}')$ defined as follows:

- $\mathfrak{U}' \stackrel{\text{def}}{=} (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup \bigcup_{\ell_0 \in W_{X,L}} S^N(\ell_0)$,
- for each $\ell \in (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$, $\mathfrak{h}'(\ell) \stackrel{\text{def}}{=} \mathfrak{h}(\ell)$,
- for each $\ell_0 \in W_{X,L}$ such that $S^N(\ell_0) = \langle \ell_0, \ldots, \ell_M \rangle$ and $M = \min(|S(\ell_0)| - 1, N)$, we define:

[3] For infinite satisfiability, it is enough to forbid positive occurrences of the magic wand containing universally quantified variables only.

- $\mathfrak{h}'(\ell_i) \stackrel{\text{def}}{=} \mathfrak{h}(\ell_i)$ $[= \ell_{i+1}]$ for all $i \in [1, M-1]$, and
- $\mathfrak{h}'(\ell_M) \stackrel{\text{def}}{=} \mathfrak{h}^i(\ell_M)$, where $i > 0$ is the smallest integer such that either $\mathfrak{h}^i(\ell_M) \in W_{X,L}$ or $\mathfrak{h}^{i+1}(\ell_M) = \bot$. Such an integer necessarily exists by definition of $S(\ell_0)$.

Proposition 5. *Given a structure $\mathcal{I} = (\mathfrak{U}, \mathfrak{s}, \mathfrak{h})$, for any (N, X, L)-contraction $C^N_{X,L}(\mathcal{I}) = (\mathfrak{U}', \mathfrak{s}, \mathfrak{h}')$, we have $\|\mathfrak{U}'\| - \|(\mathfrak{U} \setminus \overline{V}_{X,L})\| \leq 2N(\|\mathfrak{s}(X)\| + \|L\|)$.*

Proof. By induction on $\|V_{X,L}\| \geq 1$, one shows that $\|W_{X,L} \setminus V_{X,L}\| \leq \|V_{X,L}\|$, which implies $\|W_{X,L} \setminus V_{X,L}\| \leq \|\mathfrak{s}(X)\| + \|L\|$. If $\|V_{X,L}\| = 1$ then there exists at most one location $\ell \in W_{X,L}$ such that $\ell = \mathfrak{h}^i(\ell_0) = \mathfrak{h}^j(\ell)$, for some $\ell_0 \in V_{X,L}$ and some $i, j > 0$. Thus $\|W_{X,L} \setminus V_{X,L}\| \leq 1$. Let $\ell_0 \in V_{X,L}$ be a location, $V^0_{X,L} = V_{X,L} \setminus \{\ell_0\}$ and $\overline{V^0}_{X,L}$, $W^0_{X,L}$ be the sets defined using $V^0_{X,L}$ instead of $V_{X,L}$. We distinguish the following cases:

- If all locations reachable from ℓ_0 are outside $\overline{V^0}_{X,L}$, then there exists at most one location ℓ such that $\ell = \mathfrak{h}^i(\ell_0) = \mathfrak{h}^j(\ell)$, for some $i, j > 0$, thus either $W_{X,L} = W^0_{X,L}$ or $W_{X,L} = W^0_{X,L} \cup \{\ell\}$.
- Otherwise, there exists a location $\ell \in \overline{V^0}_{X,L}$ such that $\ell = \mathfrak{h}^i(\ell_0)$, for some $i > 0$ and let i be the minimal such number. Then we have $W_{X,L} = W^0_{X,L} \cup \{\ell\}$.

In both cases we have $W_{X,L} \subseteq W^0_{X,L} \cup \{\ell\}$, for some location ℓ. We compute:

$$\begin{aligned} W_{X,L} \setminus V_{X,L} &\subseteq W_{X,L} \setminus V^0_{X,L} \\ &\subseteq (W^0_{X,L} \cup \{\ell\}) \setminus V^0_{X,L} \\ &= (W^0_{X,L} \setminus V^0_{X,L}) \cup (\{\ell\} \setminus V^0_{X,L}) \end{aligned}$$

Then we obtain:

$$\begin{aligned} \|W_{X,L} \setminus V_{X,L}\| &\leq \|W^0_{X,L} \setminus V^0_{X,L}\| + \|\{\ell\} \setminus V^0_{X,L}\| \\ &\leq \|W^0_{X,L} \setminus V^0_{X,L}\| + 1 \\ \text{(induction hypothesis)} \quad &\leq \|V^0_{X,L}\| + 1 \leq \|V_{X,L}\| \end{aligned}$$

Since every segment in $C_{N,X,L}$ has length at most N, we obtain that $\mathfrak{U}'$ contains at most $\|(\mathfrak{U} \setminus \overline{V}_{X,L})\| + 2N(\|\mathfrak{s}(X)\| + \|L\|)$ locations. □

Lemma 5. *Let $\psi = \exists y_1 \ldots \exists y_m \,.\, \phi(x_1, \ldots, x_n, y_1, \ldots, y_m)$ be a formula, where such that $n, m \geq 1$ and ϕ is a quantifier-free boolean combination of test formulae. Let $X = \{x_1, \ldots, x_n\}$ and consider a structure $\mathcal{I} = (\mathfrak{U}, \mathfrak{s}, \mathfrak{h})$ such that there exists a set of locations $L \subseteq \mathfrak{U}$ with $\|L \cap \mathrm{dom}(\mathfrak{h})\| \geq \mathcal{N}(\phi)$. If $C^m_{X,L}(\mathcal{I}) \models \psi$ then $\mathcal{I} \models \psi$.*

Proof. Let $C^m_{X,L}(\mathcal{I}) = (\mathfrak{U}', \mathfrak{s}, \mathfrak{h}')$. If $(\mathfrak{U}', \mathfrak{s}, \mathfrak{h}') \models \psi$ then there exists a sequence of locations $\ell'_1, \ldots, \ell'_m \in \mathfrak{U}'$ such that $(\mathfrak{U}', \mathfrak{s}[y_1 \leftarrow \ell'_1, \ldots, y_m \leftarrow \ell'_m], \mathfrak{h}') \models \phi$. We

shall build a sequence $\ell_1, \ldots, \ell_m \in \mathfrak{U}$ such that $(\mathfrak{U}, \mathfrak{s}[y_1 \leftarrow \ell_1, \ldots, y_m \leftarrow \ell_m], \mathfrak{h}) \models \phi$. Initially, for each $\ell'_i \in (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$, let $\ell_i \stackrel{\text{def}}{=} \ell'_i$ and mark the index i as visited. Then repeat the following steps, until there are no more unmarked indices in $[1,m]$:

1. For each unmarked index i such that $\ell'_i = \ell'_j$ for some marked index j, let $\ell_i \stackrel{\text{def}}{=} \ell_j$ and mark i.
2. Choose an unmarked index i. Since i is unmarked, necessarily $\ell'_i \notin (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$ hence $\ell'_i \in S^m(\ell'_0)$, for some $\ell'_0 \in W_{X,L}$. Let $i_1 < \ldots < i_q$ be the set of unmarked indices such that $\ell'_{i_1}, \ldots, \ell'_{i_q} \in S^m(\ell'_0)$, and consider the numbers $r_1, \ldots, r_q$ such that:

$$\mathfrak{h}'^{r_1}(\ell'_0) = \ell'_{i_1}, \ldots, \mathfrak{h}'^{r_{j+1}}(\ell'_{i_j}) = \ell'_{i_{j+1}}, \ldots, \mathfrak{h}'^{r_q}(\ell'_{i_q}) = \mathfrak{h}^t(\ell'_0) \tag{1}$$

   where $t > 0$ is the smallest number such that either $\mathfrak{h}^t(\ell'_0) \in W_{X,L}$ or $\mathfrak{h}^{t+1}(\ell'_0) = \bot$. Note that in particular $\sum_{i=1}^{q} r_i = t$. If $t \leq m$ then let $\ell_{i_j} \stackrel{\text{def}}{=} \ell'_{i_j}$ for all $j \in [1,q]$. Otherwise, since $r_1 + \ldots + r_q > m$ and $q \leq m$, there exists $h \in [1,q]$ such that $r_h \geq 2$. Let h be the maximal such number. Then let $\ell_{i_j} \stackrel{\text{def}}{=} \ell'_{i_j}$ if $j \in [1,h]$ and $\ell_{i_j} \stackrel{\text{def}}{=} \mathfrak{h}^{t-\sum_{s=j}^{q} r_s}(\ell'_0)$ if $j \in [h+1,q]$. Finally mark $i_1, \ldots, i_q$ as visited.

Now we show that, for any literal λ, if $(\mathfrak{U}', \mathfrak{s}[y_1 \leftarrow \ell'_1, \ldots, y_m \leftarrow \ell'_m], \mathfrak{h}') \models \lambda$ then $(\mathfrak{U}, \mathfrak{s}[y_1 \leftarrow \ell_1, \ldots, y_m \leftarrow \ell_m], \mathfrak{h}) \models \lambda$, by a case split on the form of λ:

– $x \approx y$, $\neg x \approx y$
  - If $x, y \in X$ then $\mathfrak{s}[y_1 \leftarrow \ell'_1, \ldots, y_m \leftarrow \ell'_m]$ and $\mathfrak{s}[y_1 \leftarrow \ell_1, \ldots, y_m \leftarrow \ell_m]$ agree on the values assigned to x and y.
  - If $x \in X$ and $y = y_i$ for some $i \in [1,m]$ then $\mathfrak{s}(x) \in V_{X,L}$. If we also have $\ell'_i \in V_{X,L}$ then $\ell_i \stackrel{\text{def}}{=} \ell'_i$ and $\mathfrak{s}[y_1 \leftarrow \ell'_1, \ldots, y_m \leftarrow \ell'_m]$ and $\mathfrak{s}[y_1 \leftarrow \ell_1, \ldots, y_m \leftarrow \ell_m]$ agree on the values assigned to x and y because both values are in $W_{X,L}$. Otherwise, $\ell'_i \notin V_{X,L}$ and suppose, by contradiction, that $\ell_i \in V_{X,L}$. We distinguish the following cases:
    * if ℓ_i is assigned initially, then we have $\ell_i = \ell'_i \notin V_{X,L}$, contradiction.
    * else, if ℓ_i is assigned at step 2, it is necessarily assigned to some location not in $V_{X,L}$, contradiction.
    * otherwise, if ℓ_i is assigned to some ℓ_j (step 1) because $\ell'_i = \ell'_j$ then we obtain $\ell_j \in V_{X,L}$, $\ell'_j \notin V_{X,L}$ and the argument is repeated inductively, until a contradiction is reached.

    Then the values assigned to x and y are different for both $\mathfrak{s}[y_1 \leftarrow \ell'_1, \ldots, y_m \leftarrow \ell'_m]$ and $\mathfrak{s}[y_1 \leftarrow \ell_1, \ldots, y_m \leftarrow \ell_m]$.
  - Otherwise, $x = y_i$ and $y = y_j$ for some $i, j \in [1,m]$. Then $\ell_i = \ell_j$ iff $\ell'_i = \ell'_j$, by definition (step 1).

– $\mathsf{alloc}(x)$:

- If $x \in X$, then since $\mathfrak{s}(x) \in \mathrm{dom}(\mathfrak{h}')$ we must have $\mathfrak{s}(x) \in \mathrm{dom}(\mathfrak{h})$, because $\mathrm{dom}(\mathfrak{h}') \subseteq \mathrm{dom}(\mathfrak{h})$.
- Otherwise $x = y_i$ for some $i \in [1,m]$ and $\ell'_i \in \mathrm{dom}(\mathfrak{h}')$. We distinguish the following cases, based on the definition of ℓ_i:
  * if ℓ_i is assigned initially, we have $\ell_i = \ell'_i \in \mathrm{dom}(\mathfrak{h}') \subseteq \mathrm{dom}(\mathfrak{h})$,
  * else, if ℓ_i is assigned at step 2 then necessarily $\ell_i \in \mathrm{dom}(\mathfrak{h})$,
  * otherwise, if ℓ_i is assigned to some ℓ_j (step 1) because $\ell'_i = \ell'_j$ then we are left with proving $\ell_j \in \mathrm{dom}(\mathfrak{h})$, repeating the argument inductively.

– $\neg\mathsf{alloc}(x)$:
- If $x \in X$ then $\mathfrak{s}(x) \in V_{X,L} \subseteq W_{X,L}$. By construction, $\mathrm{dom}(\mathfrak{h}') \cap W_{X,L} = \mathrm{dom}(\mathfrak{h}) \cap W_{X,L}$, thus $\mathfrak{s}(x) \notin \mathrm{dom}(\mathfrak{h}')$ implies $\mathfrak{s}(x) \notin \mathrm{dom}(\mathfrak{h})$.
- Otherwise $x = y_i$ for some $i \in [1,m]$ and $\ell'_i \notin \mathrm{dom}(\mathfrak{h}')$. Then either $\ell'_i \in (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$, in which case $\ell_i = \ell'_i$ by definition and $\mathrm{dom}(\mathfrak{h}') \cap [(\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}] = \mathrm{dom}(\mathfrak{h}) \cap [(\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}]$, or $\ell'_i \in S^m(\ell_0)$ for some $\ell_0 \in W_{X,L}$. The latter case, however, contradicts the fact that $\ell'_i \notin \mathrm{dom}(\mathfrak{h}')$.

– $x \hookrightarrow y$:
- If $x, y \in X$, then since $\mathfrak{h}'(\mathfrak{s}(x)) = \mathfrak{s}(y)$ and $\mathfrak{s}(x) \in W_{X,L}$, we have $\mathfrak{h}(\mathfrak{s}(x)) = \mathfrak{s}(y)$ because $\mathfrak{h}'$ agrees with $\mathfrak{h}$ on $W_{X,L}$.
- If $x \in X$ and $y = y_i$ for some $i \in [1,m]$, we have $\mathfrak{h}'(\mathfrak{s}(x)) = \mathfrak{h}(\mathfrak{s}(x)) = \ell'_i$, because $\mathfrak{s}(x) \in W_{X,L}$ and $\mathfrak{h}'$ agrees with $\mathfrak{h}$ on $W_{X,L}$. There remains to show that $\ell'_i = \ell_i$ in this case. If $\ell'_i \in (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$ then this is the case by definition. Otherwise $\ell'_i \in S^m(\mathfrak{s}(x))$. Thus, $r_1 = 1$, where $r_1, \ldots, r_q$ is the sequence of numbers in step 2 of the construction above. If $t \leq m$, then $l_i = l'_i$ by construction. Otherwise, since $r_1 = 1$, the maximal number h such that $r_h \geq 2$ is strictly greater than 1 and once again, $l_i = l'_i$.
- If $x = y_i$ for some $i \in [1,m]$ and $y \in X$, we have $\mathfrak{h}'(\ell'_i) = \mathfrak{s}(y) \in W_{X,L}$. We distinguish the following cases:
  * If $\ell'_i \in (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$ then $\ell_i = \ell'_i$ by definition and moreover $\mathfrak{h}'$ agrees with $\mathfrak{h}$ on $(\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$.
  * Otherwise $\ell'_i \in S^m(\ell'_0)$ for some $\ell'_0 \in W_{X,L}$. Since $\mathfrak{s}(y) \in W_{X,L}$ it must be that ℓ'_i is the last location in $S^m(\ell'_0)$, hence $r_q = 1$, where $r_1, \ldots, r_q$ (1) is the sequence of numbers from the definition of $\ell'_1, \ldots, \ell'_m$ (step 2). Then either $r_j = 1$ for all $j \in [1,q]$, in which case $\ell'_i = \ell_i$, or $r_h \geq 2$ for some $h \in [1,q]$. However, since $h \neq q$, we also have that $\ell'_i = \ell_i$ in this case.
- If $x = y_i$ and $y = y_j$, for some $i, j \in [1,m]$, we have $\mathfrak{h}'(\ell'_i) = \ell'_j$ and we prove that $\mathfrak{h}(\ell_i) = \ell_j$ as well. We distinguish the following cases:
  * If $\ell'_i, \ell'_j \in (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$ then since $\ell'_i = \ell_i$, $\ell'_j = \ell_j$ and $\mathfrak{h}'$, $\mathfrak{h}$ agree on $W_{X,L}$, we have the result.

* Otherwise, if $\ell'_i \in S^m(\ell_0)$, for some $\ell_0 \in W_{X,L}$, let $r_p = 1$ be the number such that $\mathfrak{h}'^{r_p}(\ell'_i) = \ell'_j$ in (1), where $r_1, \ldots, r_q$ (1) is the sequence of numbers from the definition of $\ell'_1, \ldots, \ell'_m$ (step 2). Then either $r_j = 1$ for all $j \in [1,q]$, in which case $\ell'_i = \ell_i$ and $\ell'_j = \ell_j$, or $r_h \geq 2$ for some $h \in [1,q]$. However, since $h \neq p$, we also have that $\ell'_i = \ell_i$ and $\ell'_j = \ell_j$, in this case.

– $\neg x \hookrightarrow y$: If $\mathfrak{s}(x) \notin \mathrm{dom}(\mathfrak{h}')$ we show that $\mathfrak{s}(x) \notin \mathrm{dom}(\mathfrak{h})$, as in the $\neg\mathsf{alloc}(x)$ case above. Otherwise, $\mathfrak{s}(x) \in \mathrm{dom}(\mathfrak{h}')$ and $\mathfrak{h}'(\mathfrak{s}(x)) \neq \mathfrak{s}(y)$. We distinguish the following cases:
  - $x, y \in X$ is similar to the case $x \hookrightarrow y$ for $x, y \in X$, above.
  - If $x \in X$ and $y = y_i$, for some $i \in [1,m]$, we have $\mathfrak{h}'(\mathfrak{s}(x)) = \mathfrak{h}(\mathfrak{s}(x)) \neq \ell'_i$, because $\mathfrak{s}(x) \in W_{X,L}$ and $\mathfrak{h}'$, $\mathfrak{h}$ agree on $W_{X,L}$. Suppose, by contradiction, that $\mathfrak{h}(\mathfrak{s}(x)) = \ell_i$. Then $\ell_i \in S(\mathfrak{s}(x)) = \langle \mathfrak{s}(x), \ell_i, \ldots \rangle$ and since $m \geq 1$, also $\ell_i \in S^m(\mathfrak{s}(x))$, which leads to $\ell_i = \ell'_i$, in contradiction with $\mathfrak{h}'(\mathfrak{s}(x)) \neq \ell'_i$.
  - If $x = y_i$ for some $i \in [1,m]$ and $y \in X$, then $\mathfrak{h}'(\ell'_i) \neq \mathfrak{s}(y)$ and suppose, by contradiction, that $\mathfrak{h}(\ell_i) = \mathfrak{s}(y)$. We distinguish the following cases:
    * If $\ell_i \in (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$ then $\ell_i = \ell'_i$ by definition and moreover $\mathfrak{h}'$ agrees with $\mathfrak{h}$ on $\ell'_i \in (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$, which contradicts with $\mathfrak{h}'(\ell'_i) \neq \mathfrak{s}(y)$.
    * Otherwise $\ell_i \in S(\ell_0)$ for some $\ell_0 \in W_{X,L}$ and since $\mathfrak{s}(y) \in W_{X,L}$, we have $r_q = 1$ and $\ell'_i = \ell_i$, by definition, where $r_1, \ldots, r_q$ (1) is the sequence of numbers from the definition of $\ell'_1, \ldots, \ell'_m$ (step 2). Moreover, $\mathfrak{h}'(\ell'_i) = \mathfrak{s}(y)$ by the definition of $\mathfrak{h}'$, which contradicts with $\mathfrak{h}'(\ell'_i) \neq \mathfrak{s}(y)$.
  - If $x = y_i$ and $y = y_j$, for some $i, j \in [1,m]$, such that $\mathfrak{h}'(\ell'_i) \neq \ell'_j$. Suppose, by contradiction, that $\mathfrak{h}(\ell_i) = \ell_j$. We distinguish the following cases:
    * if $\ell_i, \ell_j \in (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$ then $\ell'_i = \ell_i$, $\ell'_j = \ell_j$ and $\mathfrak{h}'$ and $\mathfrak{h}$ agree on $(\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$, then $\mathfrak{h}'(\ell'_i) = \ell'_j$, contradiction.
    * if $\ell_i \in (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$ and $\ell_j \notin (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$, then $\ell'_i = \ell_i$, $\mathfrak{h}'(\ell'_i) = \mathfrak{h}(\ell_i)$ and $S(\ell'_i) = \langle \ell'_i, \ell_j, \ldots \rangle$. But then $r_1 = 1$ (1) and $\ell'_j = \ell_j$ by definition, contradiction.
    * if $\ell_i \notin (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$ and $\ell_j \in (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$, then $\ell'_j = \ell_j$ and $\ell_i \in S(\ell_0)$ for some $\ell_0 \in W$. But then ℓ_i is the last location in the segment, thus $r_q = 1$ (1) and $\ell'_i = \ell_i$, $\mathfrak{h}'(\ell'_i) = \ell'_j$ follows, contradiction.
    * if $\ell_i \notin (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$ and $\ell_j \notin (\mathfrak{U} \setminus \overline{V}_{X,L}) \cup W_{X,L}$, then $\ell_i, \ell_j \in S(\ell_0)$ for some $\ell_0 \in W_{X,L}$ and, moreover, ℓ_i and ℓ_j are consequtive locations in $S(\ell_0)$. Then the only possibility is that $\ell'_i, \ell'_j \in S^m(\ell_0)$ and $\mathfrak{h}'(\ell'_i) = \ell'_j$, contradiction.

– $|h| \geq |U| - n$, $|h| < |U| - n$: Let $T = \bigcup_{\ell_0 \in W_{X,L}} S(\ell_0) \setminus S^m(\ell_0)$. It is not hard to show that (i) $T \subseteq \mathrm{dom}(\mathfrak{h})$, (ii) $\mathfrak{U}' = \mathfrak{U} \setminus T$ and (iii) $\mathrm{dom}(\mathfrak{h}') = \mathrm{dom}(\mathfrak{h}) \setminus T$. Then $\|\mathfrak{U}\| - \|\mathrm{dom}(\mathfrak{h})\| = \|\mathfrak{U}'\| - \|\mathrm{dom}(\mathfrak{h}')\|$ and the result follows.
– $|h| \geq n$: we have $\|\mathrm{dom}(\mathfrak{h})\| \geq \|\mathrm{dom}(\mathfrak{h}')\| \geq n$.
– $|h| < n$: since $\|L \cap \mathrm{dom}(\mathfrak{h})\| \geq \mathcal{N}(\varphi)$, we have $\|\mathrm{dom}(\mathfrak{h}')\| \geq n$, thus $\mathcal{I}' \not\models |h| < n$.
– $|U| \geq n$: we have $\|\mathfrak{U}\| \geq \|\mathfrak{U}'\| \geq n$.
– $|U| < n$: since $\|L \cap \mathrm{dom}(\mathfrak{h})\| \geq \mathcal{N}(\varphi)$, we have $\|\mathfrak{U}'\| \geq n$, thus $\mathcal{I}' \not\models |U| < n$.

□

Given a set $L \subseteq \mathfrak{U}$, the (X,L)-*restriction* $R_{X,L}(\mathcal{I}) = (\mathfrak{U}', \mathfrak{s}, \mathfrak{h}')$ is defined as $\mathfrak{U}' \stackrel{\text{def}}{=} \overline{V}_{X,L}$, and for each $\ell \in \mathfrak{U}'$, $\mathfrak{h}'(\ell) \stackrel{\text{def}}{=} \mathfrak{h}(\ell)$. Observe that, because $\overline{V}_{X,L}$ is closed under applications of $\mathfrak{h}$, we have $\mathrm{dom}(\mathfrak{h}') \cup \mathrm{img}(\mathfrak{h}') \subseteq \mathfrak{U}'$.

Lemma 6. *Let $\psi = \exists y_1 \ldots \exists y_m \,.\, \phi(x_1, \ldots, x_n, y_1, \ldots, y_m)$ be a formula, where ϕ is a quantifier-free boolean combination of test formulae with free variables $x_1, \ldots, x_n, y_1, \ldots, y_m$. Let $X = \{x_1, \ldots, x_n\}$ and consider a structure $\mathcal{I} = (\mathfrak{U}, \mathfrak{s}, \mathfrak{h})$ such that there exists a set of locations $L \subseteq \mathfrak{U}$ with $\|L \cap \mathrm{dom}(\mathfrak{h})\| \geq \mathcal{N}(\phi)$ and $\|(L \cup \mathfrak{s}(X)) \setminus \mathrm{dom}(\mathfrak{h})\| = \min(\|U \setminus \mathrm{dom}(\mathfrak{h})\|, \mathcal{N}(\phi) + 1)$. If $R_{X,L}(\mathcal{I}) \models \psi$ then $\mathcal{I} \models \psi$.*

Proof. If $R_{X,L}(\mathcal{I}) \models \psi$ then there exist $\ell_1, \ldots, \ell_m \in \mathfrak{U}'$ such that $(\mathfrak{U}', \mathfrak{s}[y_1 \leftarrow \ell_1, \ldots, y_m \leftarrow \ell_m], \mathfrak{h}') \models \phi$. We show that, for each literal λ, we have $(\mathfrak{U}', \mathfrak{s}[y_1 \leftarrow \ell_1, \ldots, y_m \leftarrow \ell_m], \mathfrak{h}') \models \lambda \Rightarrow (\mathfrak{U}, \mathfrak{s}[y_1 \leftarrow \ell_1, \ldots, y_m \leftarrow \ell_m], \mathfrak{h}) \models \lambda$, using a case split on the form of λ:

1. $x \approx y$, $\neg x \approx y$: trivial, because the store does not change between $\mathcal{I}'$ and $\mathcal{I}$.
2. $\mathsf{alloc}(x)$: $\mathfrak{s}(x) \in \mathrm{dom}(\mathfrak{h}') \subseteq \mathrm{dom}(\mathfrak{h})$.
3. $\neg\mathsf{alloc}(x)$: $\mathfrak{s}(x) \in \mathfrak{U}' \setminus \mathrm{dom}(\mathfrak{h}')$ and suppose that $\mathfrak{s}(x) \in \mathrm{dom}(\mathfrak{h})$. Since $\mathrm{dom}(\mathfrak{h}') = \mathrm{dom}(\mathfrak{h}) \cap \mathfrak{U}'$, it must be the case that $\mathfrak{s}(x) \notin \mathfrak{U}'$, contradiction.
4. $x \hookrightarrow y$: $\mathfrak{s}(x), \mathfrak{s}(y) \in \mathfrak{U}'$, $\mathfrak{s}(x) \in \mathrm{dom}(\mathfrak{h}')$ and $\mathfrak{h}'$ agrees with $\mathfrak{h}$ on $\mathfrak{U}'$.
5. $\neg x \hookrightarrow y$: if $\mathfrak{s}(x) \in \mathrm{dom}(\mathfrak{h}')$ then $\mathfrak{s}(x) \in \mathrm{dom}(\mathfrak{h})$ and $\mathfrak{h}(\mathfrak{s}(x)) = \mathfrak{h}'(\mathfrak{s}(x))$, otherwise $\mathfrak{s}(x) \notin \mathrm{dom}(\mathfrak{h}')$ and $\mathfrak{s}(x) \notin \mathrm{dom}(\mathfrak{h})$ follows, by the argument used in the $\neg\mathsf{alloc}(x)$ case.
6. $|h| \geq |U| - n$: $\|\mathfrak{U}' \setminus \mathrm{dom}(\mathfrak{h}')\| \leq n$ and, since $\mathfrak{U}' = \overline{V}_{X,L}$ and $\mathrm{dom}(\mathfrak{h}') = \mathrm{dom}(\mathfrak{h}) \cap \overline{V}_{X,L}$, we compute:

$$\begin{aligned} \mathfrak{U}' \setminus \mathrm{dom}(\mathfrak{h}') &= \overline{V}_{X,L} \setminus (\mathrm{dom}(\mathfrak{h}) \cap \overline{V}_{X,L}) \\ &= \overline{V}_{X,L} \setminus \mathrm{dom}(\mathfrak{h}) \\ &\supseteq (L \cup \mathfrak{s}(X)) \setminus \mathrm{dom}(\mathfrak{h}) \end{aligned}$$

   thus $\|(L \cup \mathfrak{s}(X)) \setminus \mathrm{dom}(h)\| \leq \|\mathfrak{U}' \setminus \mathrm{dom}(\mathfrak{h}')\| \leq n$, hence, since $n < \mathcal{N}(\phi) + 1$, we have $\|(L \cup \mathfrak{s}(X)) \setminus \mathrm{dom}(h)\| = \|\mathfrak{U} \setminus \mathrm{dom}(\mathfrak{h})\| \leq n$.
7. $|h| < |U| - n$: we have $\|\mathfrak{U}' \setminus \mathrm{dom}(\mathfrak{h}')\| > n$. Since $\mathfrak{U}' \subseteq \mathfrak{U}$ and $\mathrm{dom}(\mathfrak{h}') = \mathrm{dom}(\mathfrak{h}) \cap \mathfrak{U}'$ this entails that $\|\mathfrak{U} \setminus \mathrm{dom}(\mathfrak{h})\| > n$.

8. $|h| \geq n$, $|h| < n$, $|U| \geq n$ and $|U| < n$: using the same argument as in the proof of Lemma 5. □

Theorem 3. *The finite and infinite satisfiability problems for* BSR(SL^1) *are PSPACE-complete.*

Proof. PSPACE-hardness follows from the proof that satisfiability of the quantifier free fragment of SL^2 is PSPACE-complete [5, Proposition 5]. This proof does not depend on the universe being infinite or $k = 2$. It remains to show PSPACE-membership for both problems.

Let $\psi = \forall y_1 \ldots \forall y_m \,.\, \phi(x_1, \ldots, x_n, y_1, \ldots, y_m)$, where ϕ is a quantifier-free SL^1 formula with free variables $x_1, \ldots, x_n, y_1, \ldots, y_m$. By Lemma 3, ψ has an infinite model iff $\psi \wedge \lambda_{n+m}$ has a finite model, where the size of λ_{n+m} is quadratic in $n+m$. Moreover, since λ_{n+m} is a BSR(SL) formula, $\psi \wedge \lambda_{n+m}$ is a BSR(SL) formula. We may therefore focus on the finite satisfiability problem.

By Proposition 1, ψ has a finite model iff it has a model $\mathcal{I} = (\mathfrak{U}, \mathfrak{s}, \mathfrak{h})$ such that $\|\mathfrak{U} \setminus \mathrm{elems}(\mathfrak{h})\| \leq m+n$. Suppose that $\mathcal{I} \models \psi$ where $\mathcal{I} = (\mathfrak{U}, \mathfrak{s}, \mathfrak{h})$ and $\|\mathfrak{U} \setminus \mathrm{elems}(\mathfrak{h})\| \leq m+n$. We prove that ψ has a model $\mathcal{I}' = (\mathfrak{U}', \mathfrak{s}, \mathfrak{h}')$ such that $\|\mathfrak{h}'\| \leq \|\mathfrak{U}'\| \leq \mathcal{P}(|\varphi|)$, for some polynomial function $\mathcal{P}(x)$.

Let $\mu = \bigvee_{M \in \mu^{fin}(\phi)} M$ be the expansion of ϕ as a disjunction of minterms that preserves all its finite models. By Lemma 1, the formula ψ is equivalent on finite models to $\forall y_1, \ldots, y_m \,.\, \mu$. Let $X = \{x_1, \ldots, x_n\}$ and $N = \max_{M \in \mu^{fin}(\phi)} \mathcal{N}(M)$. If there is no set $L \subseteq \mathfrak{U} \setminus \mathfrak{s}(X)$ such that $\|L \cap \mathrm{dom}(\mathfrak{h})\| = N$ and $\|(L \cup \mathfrak{s}(X)) \setminus \mathrm{dom}(\mathfrak{h})\| = \min(\|U \setminus \mathrm{dom}(\mathfrak{h})\|, N+1)$, then $\|\mathrm{dom}(\mathfrak{h})\| < N+n$ must be the case, as we show next. Suppose, by contradiction, that $\|\mathrm{dom}(\mathfrak{h})\| \geq N+n$. Then there exists a set $L_1 \subseteq (\mathfrak{U} \setminus \mathfrak{s}(X)) \cap \mathrm{dom}(\mathfrak{h})$ such that $\|L_1\| = N$. Let $n' = \|\mathfrak{s}(X) \setminus \mathrm{dom}(\mathfrak{h})\|$. By definition, $(\mathfrak{U} \setminus \mathfrak{s}(X)) \setminus \mathrm{dom}(\mathfrak{h})$ contains $\|\mathfrak{U} \setminus \mathrm{dom}(\mathfrak{h})\| - n'$ elements. Hence there exists a set $L_2 \subseteq (\mathfrak{U} \setminus \mathfrak{s}(X)) \setminus \mathrm{dom}(\mathfrak{h})$ such that $\|L_2\| = \min(\|U \setminus \mathrm{dom}(\mathfrak{h})\|, \mathcal{N}(\phi) + 1) - n'$. Let $L \stackrel{\text{def}}{=} L_1 \cup L_2$. The sets L_1, L_2 and $\mathfrak{s}(X)$ are pairwise disjoint, and since $L_1 \subseteq \mathrm{dom}(\mathfrak{h})$, we have $(L \cup \mathfrak{s}(X)) \setminus \mathrm{dom}(\mathfrak{h}) = L_2 \cup (\mathfrak{s}(X) \setminus \mathrm{dom}(\mathfrak{h}))$. We deduce that $\|(L \cup \mathfrak{s}(X)) \setminus \mathrm{dom}(\mathfrak{h})\| = \|L_2\| + n' = \min(\|U \setminus \mathrm{dom}(\mathfrak{h})\|, \mathcal{N}(\phi) + 1)$ and $\|L \cap \mathrm{dom}(\mathfrak{h})\| = \|L_1\| = N$.

Hence $\|\mathrm{dom}(\mathfrak{h})\| < N+n$ and $\|\mathrm{elems}(\mathfrak{h})\| < 2(N+n)$, since each allocated location points to exactly one location, allocated or not. Therefore, $\|\mathfrak{U}\| < m + n + 2(N+n) = 2N + 3n + m$ and since N is polynomially bounded by $\mathsf{size}(\varphi)$, by [8, Lemma 7], we are done, since we may assume that $\mathcal{P}$ is such that $\mathcal{P}(|\varphi|) \geq 2N + 3n + m$.

Otherwise, let L be such a set. By definition $\|L\| \leq N + N + 1$. By Lemma 6, since $\mathcal{I} \models \forall y_1 \ldots \forall y_m \,.\, \mu$, we have $R_{X,L}(\mathcal{I}) \models \forall y_1 \ldots \forall y_m \,.\, \mu$ and by Lemma 5, we obtain $C^m_{X,L}(R_{X,L}(\mathcal{I})) \models \forall y_1 \ldots \forall y_m \,.\, \mu$. Let $\mathcal{I}' = C^m_{X,L}(R_{X,L}(\mathcal{I})) = (\mathfrak{U}', \mathfrak{s}, \mathfrak{h}')$ and $\mathcal{I}'' = R_{X,L}(\mathcal{I}) = (\mathfrak{U}'', \mathfrak{s}, \mathfrak{h}'')$. By definition of $R_{X,L}(\mathcal{I})$, $\mathfrak{U}'' = \overline{V}_{X,L}$. By Proposition

5, we have $\|\mathfrak{U}'\| - \|\mathfrak{U}'' \setminus \overline{V}_{X,L}\| \leq 2m(n + \|L\|)$, hence we deduce that $\|\mathfrak{U}'\| \leq 2m(n + 2N + 1)$. Again, the proof is completed, taking $\mathcal{P}(|\varphi|) = 2m(n + 2N + 1)$.

We are left with proving that the model checking problem $\mathcal{I} \models \forall y_1 \ldots \forall y_m \,.\, \mu$ is in PSPACE. We prove that the complement problem $\mathcal{I} \not\models \forall y_1 \ldots \forall y_m \,.\, \mu$ is in PSPACE and use the fact that PSPACE is closed under complement [1, Corollary 4.21]. Let $\mathcal{I} = (\mathfrak{U}, \mathfrak{s}, \mathfrak{h})$. To check that $\mathcal{I} \models \exists y_1 \ldots \exists y_m \,.\, \neg\mu$, we guess locations $\ell_1, \ldots, \ell_m \in \mathfrak{U}$ and a $\mathcal{M}$-bounded minterm M. Then we check that $M \in \mu^{fin}(\neg\psi)$ and that $(\mathfrak{U}, \mathfrak{s}[y_1 \leftarrow \ell_1, \ldots, y_m \leftarrow \ell_m], \mathfrak{h}) \models M$. The first check is in PSPACE, according to Lemma 2 and the second is in P. □

6 Conclusion

We show that the prenex fragment of Separation Logic over heaps with one selector, denoted as SL^1, is decidable in time not elementary recursive. Moreover, the Bernays-Schönfinkel-Ramsey $\mathsf{BSR}(\mathsf{SL}^1)$ is PSPACE-complete. These results answer an open question raised in [8], which established the undecidability of SL^k, over heaps with $k \geq 2$ selector fields.

References

1. S. Arora and B. Barak. *Computational Complexity - A Modern Approach*. Cambridge University Press, 2009.
2. E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer, 1997.
3. R. Brochenin, S. Demri, and E. Lozes. On the almighty wand. *Information and Computation*, 211:106 – 137, 2012.
4. C. Calcagno and D. Distefano. Infer: An automatic program verifier for memory safety of c programs. In *Proc. of NASA Formal Methods'11*, volume 6617 of *LNCS*. Springer, 2011.
5. C. Calcagno, H. Yang, and P. W. O'hearn. Computability and complexity results for a spatial assertion language for data structures. In *FST TCS 2001, Proceedings*, pages 108–119. Springer, 2001.
6. S. Demri, D. Galmiche, D. Larchey-Wendling, and D. Méry. Separation logic with one quantified variable. In *CSR'14*, volume 8476 of *LNCS*, pages 125–138. Springer, 2014.
7. M. Echenim, R. Iosif, and N. Peltier. The complexity of prenex separation logic with one selector. *CoRR*, arXiv:1804.03556, 2018.
8. M. Echenim, R. Iosif, and N. Peltier. On the expressive completeness of bernays-schönfinkel-ramsey separation logic. *CoRR*, arXiv:1802.00195, 2018.
9. S. S. Ishtiaq and P. W. O'Hearn. Bi as an assertion language for mutable data structures. In *ACM SIGPLAN Notices*, volume 36, pages 14–26, 2001.
10. É. Lozes. *Expressivité des logiques spatiales*. Thèse de doctorat, Laboratoire de l'Informatique du Parallélisme, ENS Lyon, France, Nov. 2004.
11. P. W. O'Hearn, H. Yang, and J. C. Reynolds. Separation and information hiding. *SIGPLAN Not.*, 39(1):268–280, 2004.
12. M. O. Rabin. Decidability of second-order theories and automata on infinite trees. *Transactions of the American Mathematical Society*, 141:1–35, 1969.
13. J. Reynolds. Separation Logic: A Logic for Shared Mutable Data Structures. In *Proc. of LICS'02*, 2002.