Reasoning in the Bernays-Schönfinkel-Ramsey Fragment of Separation Logic

Andrew Reynolds[1], Radu Iosif[2], and Cristina Serban[2]

[1] The University of Iowa
[2] Verimag/CNRS/Université de Grenoble Alpes

Abstract. Separation Logic (SL) is a well-known assertion language used in Hoare-style modular proof systems for programs with dynamically allocated data structures. In this paper we investigate the fragment of first-order SL restricted to the Bernays-Schönfinkel-Ramsey quantifier prefix $\exists^*\forall^*$, where the quantified variables range over the set of memory locations. When this set is uninterpreted (has no associated theory) the fragment is PSPACE-complete, which matches the complexity of the quantifier-free fragment [7]. However, SL becomes undecidable when the quantifier prefix belongs to $\exists^*\forall^*\exists^*$ instead, or when the memory locations are interpreted as integers with linear arithmetic constraints, thus setting a sharp boundary for decidability within SL. We have implemented a decision procedure for the decidable fragment of $\exists^*\forall^*$SL as a specialized solver inside a DPLL(T) architecture, within the CVC4 SMT solver. The evaluation of our implementation was carried out using two sets of verification conditions, produced by (i) unfolding inductive predicates, and (ii) a weakest precondition-based verification condition generator. Experimental data shows that automated quantifier instantiation has little overhead, compared to manual model-based instantiation.

1 Introduction

Separation Logic (SL) is a popular logical framework for program verification, used by a large number of methods, ranging from static analysis [10, 27, 6] to Hoare-style proofs [19] and property-guided abstraction refinement [1]. The salient features that make SL particularly attractive for program verification are the ability of defining (i) recursive data structures using small and natural inductive definitions, (ii) weakest pre- and post-condition calculi that capture the semantics of programs with pointers, and (iii) compositional verification methods, based on the principle of local reasoning (analyzing separately pieces of program working on disjoint heaps).

Consider, for instance, the following inductive definitions, describing an acyclic and a possibly cyclic list segment, respectively:

$$\widehat{\mathsf{ls}}(x,y) \equiv \mathsf{emp} \wedge x = y \ \vee \ x \neq y \wedge \exists z \,.\, x \mapsto z * \widehat{\mathsf{ls}}(z,y) \quad \text{acyclic list segment from } x \text{ to } y$$
$$\mathsf{ls}(x,y) \equiv \mathsf{emp} \wedge x = y \ \vee \ \exists u \,.\, x \mapsto u * \mathsf{ls}(u,y) \quad \text{list segment from } x \text{ to } y$$

Intuitively, an acyclic list segment is either empty, in which case the head and the tail coincide ($\mathsf{emp} \wedge x = y$), or it contains at least one element which is disjoint from the rest of the list segment. We denote by $x \mapsto z$ the fact that x is an allocated memory location,

which points to z, and by $x \mapsto z * \widehat{ls}(z,y)$ the fact that $x \mapsto z$ and $\widehat{ls}(z,y)$ hold over disjoint parts of the heap. The constraint $x \neq y$, in the inductive definition of $\widehat{ls}$, captures the fact that the tail of the list segment is distinct from every allocated cell in the list segment, which ensures the acyclicity condition. Since this constraint is omitted from the definition of the second (possibly cyclic) list segment $ls(x,y)$, its tail y is allowed to point inside the set of allocated cells.

Automated reasoning is the key enabler of push-button program verification. Any procedure that checks the validity of a logical entailment between inductive predicates requires checking the satisfiability of formulae from the base (non-inductive) assertion language, as shown by the example below. Consider a fragment of the inductive proof showing that any acyclic list segment is also a list segment, given below:

$$\dfrac{\dfrac{\widehat{ls}(z,y) \vdash ls(z,y)}{x \neq y \wedge x \mapsto z * \widehat{ls}(z,y) \vdash \exists u \,.\, x \mapsto u * ls(u,y)}}{\widehat{ls}(x,y) \vdash ls(x,y)} \quad \begin{array}{l} x \neq y \wedge x \mapsto z \models \exists u \,.\, x \mapsto u \\ \text{by instantiation } u \leftarrow z \end{array}$$

The first (bottom) derivation in the proof corresponds to one of the two cases produced by unfolding both the antecedent and consequent of the entailment (the second case $emp \wedge x = y \vdash emp \wedge x = y$ is trivial and omitted for clarity). The second derivation is a simplification of the sequent obtained by unfolding, to a sequent matching the initial one (by renaming z to x), and allows to conclude this branch of the proof by an inductive argument, based on the principle of infinite descent [5].

The simplification applied by the second derivation above relies on the validity of the entailment $x \neq y \wedge x \mapsto z \models \exists u \,.\, x \mapsto u$, which reduces to the (un)satisfiability of the formula $x \neq y \wedge x \mapsto z \wedge \forall u \,.\, \neg x \mapsto u$. The latter falls into the Bernays-Schönfinkel-Ramsey fragment, defined by the $\exists^*\forall^*$ quantifier prefix, and can be proved unsatisfiable using the instantiation of the universally quantified variable u with the existentially quantified variable z (or a corresponding Skolem constant). In other words, this formula is unsatisfiable because the universal quantified subformula asks that no memory location is pointed to by x, which is contradicted by $x \mapsto z$. The instantiation of u that violates the universal condition is $u \leftarrow z$, which is carried over in the rest of the proof.

The goal of this paper is mechanizing satisfiability of the Bernays-Schönfinkel-Ramsey fragment of SL, without inductively defined predicates[3]. This fragment is defined by the quantifier prefix of the formulae in prenex normal form. We consider formulae $\exists x_1 \ldots \exists x_m \forall y_1 \ldots \forall y_n \,.\, \phi(x_1, \ldots, x_m, y_1, \ldots, y_n)$, where ϕ is any quantifier-free formula of SL, consisting of pure formulae from given base theory T, and points-to atomic propositions relating terms of T, combined with unrestricted Boolean and separation connectives, and the quantified variables range essentially over the set of memory locations. In a nutshell, the contributions of the paper are two-fold:

1. We draw a sharp boundary between decidability and undecidability, proving essentially that the satisfiability problem for the Bernays-Schönfinkel-Ramsey fragment of SL is PSPACE-complete, if the domain of memory locations is an uninterpreted

[3] Strictly speaking, the Bernays-Schönfinkel-Ramsey class refers to the $\exists^*\forall^*$ fragment of first-order logic with equality and predicate symbols, but no function symbols [17].

set, whereas interpreting memory locations as integers with linear arithmetic constraints, leads to undecidability. Moreover, undecidability occurs even for uninterpreted memory locations, if we extend the quantifier prefix to $\exists^*\forall^*\exists^*$.

2. We have implemented an effective decision procedure for quantifier instantiation, based on counterexample-driven learning of conflict lemmas, integrated within the DPLL(T) architecture [12] of the CVC4 SMT solver [2]. Experimental evaluation of our implementation shows that the overhead of the push-button quantifier instantiation is negligible, compared to the time required to solve a quantifier-free instance of the problem, obtained manually, by model inspection.

Related Work The first theoretical results on the decidability and computational complexity of SL (without inductive definitions) were found by Calcagno, Yang and O'Hearn [7]. They show that the satisfiability problem for SL is undecidable, in the presence of quantifiers, assuming that each memory location can point to two other locations, i.e. using atomic propositions of the form $\mathsf{x} \mapsto (\mathsf{y}, \mathsf{z})$. Decidability can be recovered by considering the quantifier-free fragment, proved to be PSPACE-complete, by a small model argument [7]. Refinements of these results consider decidable fragments of SL with one record field (atomic points-to propositions $\mathsf{x} \mapsto \mathsf{y}$), and one or two quantified variables. In a nutshell, SL with one record field and separating conjunction only is decidable with non-elementary time complexity, whereas adding the magic wand adjoint leads to undecidability [4]. Decidability, in the presence of the magic wand operator, is recovered by restricting the number of quantifiers to one, in which case the logic becomes PSPACE-complete [9]. This bound is sharp, because allowing two quantified variables leads to undecidability, and decidability with non-elementary time complexity if the magic wand is removed [8].

SMT techniques were applied to deciding the satisfiability of SL in the work of Piskac, Wies and Zufferey [21, 22]. They considered quantifier-free fragments of SL with separating conjunction in positive form (not occurring under negation) and without magic wand, and allow for hardcoded inductive predicates (list and tree segments). In a similar spirit, we previously define a translation to multi-sorted second-order logic combined with counterexample-driven instantiation for set quantifiers to define a decision procedure for the quantifier-free fragment of SL [24]. In a different vein, a tableau-based semi-decision procedure is given by Méry and Galmiche [11]. Termination of this procedure is guaranteed for the (decidable) quantifier-free fragment of SL, yet no implementation is available for comparison.

A number of automated theorem provers have efficient and complete approaches for the Bernays-Schönfinkel-Ramsey fragment of first-order-logic, also known as effectively propositional logic (EPR) [3, 16]. A dedicated approach for EPR in the SMT solver Z3 was developed in [20]. A approach based on finite model finding is implemented in CVC4 [25], which is model-complete for EPR. Our approach is based on counterexample-guided quantifier instantiation, which has been used in the context of SMT solving in previous works [13, 23].

2 Preliminaries

We consider formulae in multi-sorted first-order logic, over a signature consisting of a countable set of sort symbols and a set of function symbols, and we write $\top$ and $\bot$ for the Boolean constants *true* and *false*. A *signature* Σ consists of a set Σ^{s} of sort symbols and a set Σ^{f} of (sorted) *function symbols* $f^{S_1 \cdots S_n S}$, where $n \geq 0$ and $S_1, \ldots, S_n, S \in \Sigma^{\mathrm{s}}$. If $n = 0$, we call f^S a *constant symbol*. In this paper, we consider signatures where for any finite sequence of sorts $S_1, \ldots, S_n \in \Sigma^{\mathrm{s}}$, the *tuple* sort $S_1 \times \ldots \times S_n$ also belongs to Σ^{s}. For each $k > 0$, let S^k denote the k-tuple sort $S \times \ldots \times S$.

Let Vars be a countable set of first-order variables, each $x^S \in \mathsf{Vars}$ having an associated sort S. First-order terms and formulae over the signature Σ (called Σ-terms and Σ-formulae) are defined as usual. For a Σ-formula φ, we denote by $\mathrm{Fvc}(\varphi)$ the set of free variables and constant symbols in φ, and by writing $\varphi(x)$ we mean that $x \in \mathrm{Fvc}(\phi)$. Whenever $\mathrm{Fvc}(\phi) \cap \mathsf{Vars} = \emptyset$, we say that ϕ is a *sentence*, i.e. ϕ has no free variables. A *Σ-interpretation* $\mathcal{I}$ maps: (1) each sort symbol $S \in \Sigma$ to a non-empty set $S^{\mathcal{I}}$, (2) each function symbol $f^{S_1,\ldots,S_n,S} \in \Sigma$ to a total function $f^{\mathcal{I}} : S_1^{\mathcal{I}} \times \ldots \times S_n^{\mathcal{I}} \to S^{\mathcal{I}}$ where $n > 0$, and to an element of $S^{\mathcal{I}}$ when $n = 0$, and (3) each variable $x^S \in \mathsf{Vars}$ to an element of $S^{\mathcal{I}}$. For an interpretation $\mathcal{I}$ a sort symbol σ and a variable x, we denote by $\mathcal{I}[\sigma \leftarrow S]$ and, respectively $\mathcal{I}[x \leftarrow v]$, the interpretation associating the set S to σ, respectively the value v to x, and which behaves like $\mathcal{I}$ in all other cases. For a Σ-term t, we write $t^{\mathcal{I}}$ to denote the interpretation of t in $\mathcal{I}$, defined inductively, as usual. A satisfiability relation between Σ-interpretations and Σ-formulas, written $\mathcal{I} \models \varphi$, is also defined inductively, as usual. We say that $\mathcal{I}$ is *a model of* φ if $\mathcal{I}$ satisfies φ.

A (multi-sorted first-order) *theory* is a pair $T = (\Sigma, \mathbf{I})$ where Σ is a signature and $\mathbf{I}$ is a non-empty set of Σ-interpretations, the *models* of T. We assume that Σ contains always the equality predicate, which we denote by $\approx$, as well as projection functions for each tuple sort. A Σ-formula φ is *T-satisfiable* if it is satisfied by some interpretation in $\mathbf{I}$. We write E to denote the empty theory (with equality), whose signature consists of a sort U with no additional function symbols, and LIA to denote the theory of linear integer arithmetic, whose signature consists of the sort Int, the binary predicate symbol $\geq$, function $+$ denoting addition, and the constants $0, 1$ of sort Int, interpreted as usual. By ELIA we denote the theory obtained by extending the signature of LIA with the sort U of E and equality over U.

Let $T = (\Sigma, \mathbf{I})$ be a theory and let Loc and Data be two sorts from Σ, with no restriction other than the fact that Loc is always interpreted as a countable set. Also, we consider that Σ has a designated constant symbol $\mathsf{nil}^{\mathsf{Loc}}$. The *Separation Logic* $\mathsf{SL}(T)_{\mathsf{Loc},\mathsf{Data}}$ is the set of formulae generated by the following syntax:

$$\varphi := \phi \mid \mathsf{emp} \mid \mathsf{t} \mapsto \mathsf{u} \mid \varphi_1 * \varphi_2 \mid \varphi_1 \mathrel{-\!\!*} \varphi_2 \mid \neg\varphi_1 \mid \varphi_1 \wedge \varphi_2 \mid \exists x^S \ . \ \varphi_1(x)$$

where ϕ is a Σ-formula, and t, u are Σ-terms of sorts Loc and Data, respectively. As usual, we write $\forall x^S \ . \ \varphi(x)$ for $\neg\exists x^S \ . \ \neg\varphi(x)$. We omit specifying the sorts of variables and constants when they are clear from the context.

Given an interpretation $\mathcal{I}$, a *heap* is a finite partial mapping $h : \mathsf{Loc}^{\mathcal{I}} \rightharpoonup_{\mathrm{fin}} \mathsf{Data}^{\mathcal{I}}$. For a heap h, we denote by $\mathrm{dom}(h)$ its domain. For two heaps h_1 and h_2, we write $h_1 \# h_2$ for $\mathrm{dom}(h_1) \cap \mathrm{dom}(h_2) = \emptyset$ and $h = h_1 \uplus h_2$ for $h_1 \# h_2$ and $h = h_1 \cup h_2$. We define the

satisfaction relation $\mathcal{I}, h \models_{\mathsf{SL}} \phi$ inductively, as follows:

$$
\begin{array}{ll}
\mathcal{I}, h \models_{\mathsf{SL}} \phi & \iff \mathcal{I} \models \phi \text{ if } \phi \text{ is a } \Sigma\text{-formula} \\
\mathcal{I}, h \models_{\mathsf{SL}} \mathsf{emp} & \iff h = \emptyset \\
\mathcal{I}, h \models_{\mathsf{SL}} \mathsf{t} \mapsto \mathsf{u} & \iff h = \{(\mathsf{t}^{\mathcal{I}}, \mathsf{u}^{\mathcal{I}})\} \text{ and } \mathsf{t}^{\mathcal{I}} \not\approx \mathsf{nil}^{\mathcal{I}} \\
\mathcal{I}, h \models_{\mathsf{SL}} \phi_1 * \phi_2 & \iff \text{there exist heaps } h_1, h_2 \text{ such that } h = h_1 \uplus h_2 \text{ and } \mathcal{I}, h_i \models_{\mathsf{SL}} \phi_i, i = 1, 2 \\
\mathcal{I}, h \models_{\mathsf{SL}} \phi_1 \mathrel{-\!\!*} \phi_2 & \iff \text{for all heaps } h' \text{ if } h' \# h \text{ and } \mathcal{I}, h' \models_{\mathsf{SL}} \phi_1 \text{ then } \mathcal{I}, h' \uplus h \models_{\mathsf{SL}} \phi_2 \\
\mathcal{I}, h \models_{\mathsf{SL}} \exists x^S . \varphi(x) & \iff \mathcal{I}[x \leftarrow s], h \models_{\mathsf{SL}} \varphi(x), \text{ for some } s \in S^{\mathcal{I}}
\end{array}
$$

The satisfaction relation for Σ-formulae, Boolean connectives $\wedge$, $\neg$, and linear arithmetic atoms, are the classical ones from first-order logic. Notice that the range of a quantified variable x^S is the interpretation of its associated sort $S^{\mathcal{I}}$.

A formula φ is said to be *satisfiable* if there exists an interpretation $\mathcal{I}$ and a heap h such that $\mathcal{I}, h \models_{\mathsf{SL}} \varphi$. The (SL, T)-*satisfiability problem* asks, given an SL formula φ, whether there exists an interpretation $\mathcal{I}$ of T and a heap h such that $\mathcal{I}, h \models_{\mathsf{SL}} \varphi$. We write $\varphi \models_{\mathsf{SL}} \psi$ if for every interpretation $\mathcal{I}$ and heap h, if $\mathcal{I}, h \models_{\mathsf{SL}} \varphi$ then $\mathcal{I}, h \models_{\mathsf{SL}} \psi$, and we say that φ *entails* ψ in this case.

The Bernays-Schönfinkel-Ramsey Fragment of SL In this paper we address the satisfiability problem for the class of sentences $\phi \equiv \exists x_1 \ldots \exists x_m \forall y_1 \ldots \forall y_n \,.\, \varphi(x_1, \ldots, x_m, y_1, \ldots, y_n)$, where φ is a quantifier-free formula of $\mathsf{SL}(T)_{\mathsf{Loc,Data}}$. We shall denote this fragment by $\exists^* \forall^* \mathsf{SL}(T)_{\mathsf{Loc,Data}}$. It is easy to see that any sentence ϕ, as above, is satisfiable if and only if the sentence $\forall y_1 \ldots \forall y_n \,.\, \varphi[c_1/x_1, \ldots, c_m/x_m]$ is satisfiable, where $c_1, \ldots, c_m$ are (Skolem) constant symbols. The latter is called the *functional form* of ϕ.

We address the satisfiability problem for $\exists^* \forall^* \mathsf{SL}(T)_{\mathsf{Loc,Data}}$ in the following cases:

1. Loc is interpreted as the sort U of E and Data as U^k, for some $k \geq 1$. The satisfiability problem for the fragment $\exists^* \forall^* \mathsf{SL}(\mathsf{E})_{U,U^k}$ is PSPACE-complete, and the proof follows a small model property argument.
2. as above, with the further constraint that U is interpreted as an infinite countable set, i.e. of cardinality $\aleph_0$. In this case, we prove a cut-off property stating that all locations not in the domain of the heap and not used in the interpretation of constants, are equivalent from the point of view of an SL formula. This satisfiability problem is reduced to unconstrained one above, and also found to be PSPACE-complete.
3. both Loc and Data are interpreted as Int, equipped with addition and total order, in which case $\exists^* \forall^* \mathsf{SL}(\mathsf{LIA})_{\mathsf{Int,Int}}$ is undecidable.
4. Loc is interpreted as the sort U of E, and Data as $U \times \mathsf{Int}$. Then $\exists^* \forall^* \mathsf{SL}(\mathsf{ELIA})_{U,U \times \mathsf{Int}}$ is undecidable.

Additionally, we prove that the fragment $\exists^* \forall^* \exists^* \mathsf{SL}(\mathsf{E})_{U,U^k}$, with two quantifier alternations, is undecidable, if $k \geq 2$. The question whether the fragment $\exists^* \forall^* \mathsf{SL}(\mathsf{ELIA})_{U,\mathsf{Int}}$ is decidable is currently open, and considered for future work.

3 Decidability and Complexity Results

This section defines the decidable cases of the Bernays-Schönfinkel-Ramsey fragment of SL, with matching undecidable extensions. The decidable fragment $\exists^* \forall^* \mathsf{SL}(\mathsf{E})_{U,U^k}$ relies on a small model property, given in section 3.1. Undecidability of $\exists^* \forall^* \mathsf{SL}(\mathsf{LIA})_{\mathsf{Int,Int}}$

is obtained by a refinement of the undecidability proof for Presburger arithmetic with one monadic predicate [14], in section 3.4.

3.1 Small Model Property

For reasons of self-containment, we recall a number of definitions and results from [29]. Some of them are slightly modified for our purposes, but these changes have no effect on the validity of their original proofs. In the rest of this section, we consider formulae of $\mathsf{SL}(\mathsf{E})_{U,U^k}$, meaning that (i) $\mathsf{Loc} = U$, and (ii) there exists an integer $k > 0$ such that $\mathsf{Data} = U^k$, where U is the (uninterpreted) sort of E. We fix k for the rest of this section.

Definition 1. *[29, Definition 90] Given a set of locations S, the equivalence relation $=_S$ between k-tuples of locations is defined as $\langle v_1, \ldots, v_k \rangle =_S \langle v'_1, \ldots, v'_k \rangle$ if and only if*
- *if $v_i \in S$ then $v_i = v'_i$, and*
- *if $v_i \notin S$ then $v'_i \notin S$,*

for all $i = 1, \ldots, k$.

Intuitively, $=_S$ restricts the equality to the elements in S. Observe that $=_S$ is an equivalence relation and that $S \subseteq T$ implies $=_T \subseteq =_S$. For a set S, we write $\|S\|$ for its cardinality, in the following.

Definition 2. *[29, Definition 91] Given an interpretation $\mathcal{I}$, an integer $n > 0$, a set of variables $X \subseteq \mathsf{Vars}$ and a set of locations $S \subseteq U^{\mathcal{I}}$, for any two heaps $h, h' : U^{\mathcal{I}} \rightharpoonup_{\text{fin}} (U^{\mathcal{I}})^k$, we define $h \sim^{\mathcal{I}}_{n,X,S} h'$ if and only if*
1. *$\mathcal{I}(X) \cap \text{dom}(h) = \mathcal{I}(X) \cap \text{dom}(h')$,*
2. *for all $\ell \in \mathcal{I}(X) \cap \text{dom}(h)$, we have $h(\ell) =_{\mathcal{I}(X) \cup S} h'(\ell)$,*
3. *if $\|\text{dom}(h) \setminus \mathcal{I}(X)\| < n$ then $\|\text{dom}(h) \setminus \mathcal{I}(X)\| = \|\text{dom}(h') \setminus \mathcal{I}(X)\|$,*
4. *if $\|\text{dom}(h) \setminus \mathcal{I}(X)\| \geq n$ then $\|\text{dom}(h') \setminus \mathcal{I}(X)\| \geq n$.*

Observe that, for any $n \leq m$ and $S \subseteq T$ we have $\sim^{\mathcal{I}}_{m,X,T} \subseteq \sim^{\mathcal{I}}_{n,X,S}$. In addition, for any integer $k > 0$, subset $S \subseteq U^{\mathcal{I}}$ and location $\ell \in U^{\mathcal{I}}$, we consider the function $prun^{\ell}_{k,S}(\ell_1, \ldots, \ell_k)$, which replaces each value $\ell_i \notin S$ in its argument list by ℓ.

Lemma 1. *[29, Lemma 94] Given an interpretation $\mathcal{I}$ and a heap $h : U^{\mathcal{I}} \rightharpoonup_{\text{fin}} (U^{\mathcal{I}})^k$, for each integer $n > 0$, each set of variables $X \subseteq \mathsf{Vars}$, each set of locations $L \subseteq U^{\mathcal{I}}$ such that $L \cap \mathcal{I}(X) = \emptyset$ and $\|L\| = n$, and each location $v \in U^{\mathcal{I}} \setminus (\mathcal{I}(X) \cup \{\mathsf{nil}^{\mathcal{I}}\} \cup L)$, there exists a heap $h' : U^{\mathcal{I}} \rightharpoonup_{\text{fin}} (U^{\mathcal{I}})^k$, with the following properties:*
1. *$h \sim^{\mathcal{I}}_{n,X,L} h'$,*
2. *$\text{dom}(h') \setminus \mathcal{I}(X) \subseteq L$,*
3. *for all $\ell \in \text{dom}(h')$, we have $h'(\ell) = prun^{v}_{k,\mathcal{I}(X) \cup L}(h(\ell))$.*

Next, we define the following measure on quantifier-free SL formulae:

$$|\phi * \psi| = |\phi| + |\psi| \quad |\phi \mathrel{-\!\!*} \psi| = |\psi| \quad |\phi \wedge \psi| = \max(|\phi|, |\psi|) \quad |\neg\phi| = |\phi|$$
$$|\mathsf{t} \mapsto \mathsf{u}| = 1 \quad |\mathsf{emp}| = 1 \quad |\phi| = 0 \text{ if } \phi \text{ is a } \Sigma\text{-formula}$$

Intuitively, $|\varphi|$ is the maximum number of invisible locations, that are not in $\mathcal{I}(\text{Fvc}(\varphi))$, and which can be distinguished by the quantifier-free $\mathsf{SL}(\mathsf{E})_{U,U^k}$ formula φ. The crux of

the PSPACE-completeness proof for quantifier-free $\mathsf{SL}(\mathsf{E})_{U,U^k}$ is that two heaps equivalent up to $|\varphi|$ invisible locations are also equivalent from the point of view of satisfiability of φ, which provides a small model property for this fragment [29, 7].

Lemma 2. *[29, Prop. 95] Given a quantifier-free $\mathsf{SL}(\mathsf{E})_{U,U^k}$ formula φ, an interpretation $\mathcal{I}$, and two heaps h and h', if $h \sim^{\mathcal{I}}_{|\varphi|,\mathrm{Fvc}(\varphi),\emptyset} h'$ and $\mathcal{I}, h \models_{\mathsf{SL}} \varphi$ then $\mathcal{I}, h' \models_{\mathsf{SL}} \varphi$.*

Our aim is to extend this result to $\exists^*\forall^*\mathsf{SL}(\mathsf{E})_{U,U^k}$, in the first place. This new small model property is given by the next lemma.

Lemma 3. *Let $\varphi(x_1^U, \ldots, x_n^U)$ be a quantifier-free $\mathsf{SL}(\mathsf{E})_{U,U^k}$-formula, and $\varphi^\forall \equiv \forall x_1^U \ldots \forall x_n^U \,.\, \varphi(x_1^U, \ldots, x_n^U)$ be its universal closure. Then $\varphi^\forall$ has a model if and only if there exists an interpretation $\mathcal{I}$ and a heap $h : U^\mathcal{I} \rightharpoonup_{\mathrm{fin}} (U^\mathcal{I})^k$ such that $\mathcal{I}, h \models_{\mathsf{SL}} \varphi^\forall$ and:*

1. $\|U^\mathcal{I}\| \leq |\varphi| + \|\mathrm{Fvc}(\varphi^\forall)\| + n$,
2. $\mathrm{dom}(h) \subseteq L \cup \mathcal{I}(\mathrm{Fvc}(\varphi^\forall))$,
3. *for all $\ell \in \mathrm{dom}(h)$, we have $h(\ell) \in (\mathcal{I}(\mathrm{Fvc}(\varphi^\forall)) \cup \{\mathsf{nil}^\mathcal{I}\} \cup L \cup \{v\})^k$,*

where $L \subseteq U^\mathcal{I} \setminus \mathcal{I}(\mathrm{Fvc}(\varphi^\forall))$ is a set of locations such that $\|L\| = |\varphi| + n$ and $v \in U^\mathcal{I} \setminus (\mathcal{I}(\mathrm{Fvc}(\varphi^\forall)) \cup \{\mathsf{nil}^\mathcal{I}\} \cup L)$ is an arbitrary location.

Proof. See Appendix A.1.

We are ready to prove two decidability results, based on the above small model property, concerning the cases where (i) Loc is interpreted as a countable set with equality, and (ii) Loc is interpreted as an infinite countable set with no other operators than equality.

3.2 Uninterpreted Locations without Cardinality Constraints

In this section, we consider the satisfiability problem for the fragment $\exists^*\forall^*\mathsf{SL}(\mathsf{E})_{U,U^k}$, where the location sort U can be interpreted by any (possibly finite) countable set, with no other operations than the equality, and the data sort consists of k-tuples of locations.

Theorem 1. *The satisfiability problem for $\exists^*\forall^*\mathsf{SL}(\mathsf{E})_{U,U^k}$, problem is* PSPACE*-complete.*

Proof. PSPACE-hardness follows from the fact that satisfiability is PSPACE-complete for quantifier-free $\mathsf{SL}(\mathsf{E})_{U,U^k}$ [7]. To prove membership in PSPACE, consider the formula $\phi \equiv \exists x_1 \ldots \exists x_m \forall y_1 \ldots \forall y_n \,.\, \varphi(\mathbf{x}, \mathbf{y})$, where φ is a quantifier-free $\mathsf{SL}(\mathsf{E})_{U,U^k}$ formula. Let $\mathbf{c} = \langle c_1, \ldots, c_m \rangle$ be a tuple of constant symbols, and $\widetilde{\phi} \equiv \forall y_1 \ldots \forall y_n \,.\, \varphi(\mathbf{c}, \mathbf{y})$ be the functional form of ϕ, obtained by replacing x_i with c_i, for all $i = 1, \ldots, m$. By Lemma 3, $\widetilde{\phi}$ has a model if and only if it has a model $\mathcal{I}, h$ such that:

- $\|U^\mathcal{I}\| \leq |\varphi| + n + m$,
- $\mathrm{dom}(h) \subseteq L \cup \mathbf{c}^\mathcal{I}$,
- $\forall \ell \in \mathrm{dom}(h) \,.\, h(\ell) \in (\mathcal{I}(\mathbf{c}) \cup \{\mathsf{nil}^\mathcal{I}\} \cup L \cup \{v\})^k$,

where $L \subseteq U^\mathcal{I} \setminus \mathcal{I}(\mathbf{c})$, $\|L\| = |\varphi| + m$ and $v \in U^\mathcal{I} \setminus (\mathcal{I}(\mathbf{c}) \cup \{\mathsf{nil}^\mathcal{I}\} \cup L)$. We describe below a nondeterministic polynomial space algorithm that decides satisfiability of $\widetilde{\phi}$. First, nondeterministically chose a model $\mathcal{I}, h$ that meets the above requirements. Then we check, for each tuple $\langle u_1, \ldots, u_n \rangle \in (U^\mathcal{I})^n$ that $\mathcal{I}[y_1 \leftarrow u_1] \ldots [y_n \leftarrow u_n], h \models_{\mathsf{SL}} \varphi$. In order

to enumerate all tuples from $(U^{\mathcal{I}})^n$ we need $n \cdot \lceil \log_2(|\varphi| + n + m) \rceil$ extra bits, and the check for each such tuple can be done in PSPACE, according to [7, §5]. □

This result is somewhat surprising, because the classical Bernays-Schönfinkel fragment of first-order formulae with predicate symbols (but no function symbols) and quantifier prefix $\exists^*\forall^*$ is known to be NEXPTIME-complete [17, §7]. The explanation lies in the fact that the interpretation of an arbitrary predicate symbol $P(\mathsf{x}_1, \ldots, \mathsf{x}_n)$ cannot be captured using only points-to atomic propositions, e.g. $\mathsf{x}_1 \mapsto (\mathsf{x}_2, \ldots, \mathsf{x}_n)$, between locations and tuples of locations, due to the interpretation of points-to's as heaps[4] (finite partial functions).

The following lemma sets a first decidability boundary for $\mathsf{SL}(\mathsf{E})_{U,U^k}$, by showing how extending the quantifier prefix to $\exists^*\forall^*\exists^*$ leads to undecidability.

Lemma 4. *The satisfiability problem for $\exists^*\forall^*\exists^*\mathsf{SL}(\mathsf{E})_{U,U^k}$ is undecidable, if $k \geq 2$.*

Proof. See Appendix A.2.

Observe that the result of Lemma 4 sets a fairly tight boundary between the decidable and undecidable fragments of SL. On one hand, simplifying the quantifier prefix to $\exists^*\forall^*$ yields a decidable fragment (Theorem 1), whereas $\mathsf{SL}(\mathsf{E})_{U,U}$ ($k = 1$) without the magic wand ($\mathrel{-\!\!*}$) is decidable with non-elementary time complexity, even when considering an unrestricted quantifier prefix [4].

3.3 Uninterpreted Locations with Cardinality $\aleph_0$

We consider the stronger version of the satisfiability problem for $\exists^*\forall^*\mathsf{SL}(\mathsf{E})_{U,U^k}$, where U is interpreted as an infinite countable set (of cardinality $\aleph_0$) with no function symbols, other than equality. Instances of this problem occur when, for instance, the location sort is taken to be Int, but no operations are used on integers, except for testing equality.

Observe that this restriction changes the satisfiability status of certain formulae. For instance, $\exists \mathsf{x} \forall \mathsf{y} \;.\; \mathsf{y} \mapsto \mathsf{x}$ is satisfiable if U is interpreted as a finite set, but becomes unsatisfiable when U is infinite. The reason is that this formula requires every location from $U^{\mathcal{I}}$ to be part of the domain of the heap, which is impossible due the fact that only finite heaps are considered by the semantics of SL.

In the following proof, we use the formula $\mathsf{alloc}(\mathsf{x}) \equiv \mathsf{x} \mapsto \mathsf{x} \mathrel{-\!\!*} \top$, expressing the fact that a location variable x is *allocated*, i.e. its interpretation is part of the heap's domain [4]. Intuitively, we reduce any instance of the $\exists^*\forall^*\mathsf{SL}(\mathsf{E})_{U,U^k}$ satisfiability problem, with U of cardinality $\aleph_0$, to an instance the same problem without this restriction, by the following cut-off argument: if a free variable is interpreted as a location which is neither part of the heap's domain, nor equal to the interpretation of some constant, then it is not important which particular location is chosen for that interpretation.

Theorem 2. *The satisfiability problem for $\exists^*\forall^*\mathsf{SL}(\mathsf{E})_{U,U^k}$ is* PSPACE*-complete if U is required to have cardinality $\aleph_0$.*

[4] If $\mathsf{x}_1 \mapsto (\mathsf{x}_2, \ldots, \mathsf{x}_n)$ and $\mathsf{x}_1 \mapsto (\mathsf{x}'_2, \ldots, \mathsf{x}'_n)$ hold, this forces $\mathsf{x}_i = \mathsf{x}'_i$, for all $i = 2, \ldots, n$.

Proof. PSPACE-hardness follows from the PSPACE-completeness of the satisfiability problem for quantifier-free SL, with uninterpreted locations [7, §5.2]. Since the reduction from [7, §5.2] involves no universally quantified variables, the $\aleph_0$ cardinality constraint has no impact on this result.

Let $\exists x_1 \ldots \exists x_m \forall y_1 \ldots \forall y_n \,.\, \varphi(\mathbf{x}, \mathbf{y})$ be a formula, and $\forall y_1 \ldots \forall y_n \,.\, \varphi(\mathbf{c}, \mathbf{y})$ be its functional form, obtained by replacing each x_i with c_i, for $i = 1, \ldots, m$. We consider the following formulae:

$$\begin{array}{rcl} \psi_0(y) & \equiv & \mathsf{alloc}(y) \\ \psi_1(y) & \equiv & \bigvee_{i=1}^{m} y = c_i \\ \psi_2(y) & \equiv & y = d_y \\ \mathsf{external} & \equiv & \bigwedge_{i=1}^{n} (\neg \mathsf{alloc}(d_{y_i}) \wedge \bigwedge_{j=1}^{m} d_{y_i} \neq c_j) \end{array}$$

where $\{d_y \mid y \in \mathbf{y}\}$ is a set of constant symbols not occurring in $\forall y_1 \ldots \forall y_n \,.\, \varphi(\mathbf{c}, \mathbf{y})$. Then we show the following fact:

Fact 1 *There exists an interpretation $\mathcal{I}$ and a heap h such that $\|U^{\mathcal{I}}\| = \aleph_0$ and $\mathcal{I}, h \models_{\mathsf{SL}} \forall y_1 \ldots \forall y_n \,.\, \varphi(\mathbf{c}, \mathbf{y})$ iff there exists an interpretation $\mathcal{I}'$, not constraining the cardinality of $U^{\mathcal{I}'}$, and a heap h' such that:*

$$\mathcal{I}', h' \models_{\mathsf{SL}} \mathsf{external} \wedge \forall y_1 \ldots \forall y_n \bigwedge_{\langle t_1, \ldots, t_n \rangle \in \{0,1,2\}^n} \underbrace{\bigwedge_{i=1}^{n} \Big(\psi_{t_i}(y_i) \Rightarrow \varphi(\mathbf{c}, \mathbf{y}) \Big)}_{\Psi_{\langle t_1, \ldots, t_n \rangle}}$$

Proof. See Appendix A.3.

To show membership in PSPACE, consider a nondeterministic algorithm that choses $\mathcal{I}'$ and h' and uses $2n$ extra bits to check that $\mathcal{I}', h' \models_{\mathsf{SL}} \mathsf{extern} \wedge \forall y_1 \ldots \forall y_n \,.\, \Psi_{\langle t_1, \ldots, t_n \rangle}$ separately, for each $\langle t_1, \ldots, t_n \rangle \in \{0, 1, 2\}^n$. By Lemma 3, the sizes of $\mathcal{I}'$ and h' are bounded by a polynomial in the size of $\Psi_{\langle t_1, \ldots, t_n \rangle}$, which is polynomial in the size of φ, and by Theorem 1, each of these checks can be done in polynomial space. □

3.4 Integer Locations with Linear Arithmetic

In the rest of this section we show that the Bernays-Schönfinkel-Ramsey fragment of SL becomes undecidable as soon as we use integers to represent the set of locations and combine SL with linear integer arithmetic (LIA). The proof relies on an undecidability argument for a fragment of Presburger arithmetic with one monadic predicate symbol, interpreted over finite sets. Formally, we denote by $(\exists^*\forall^* \cap \forall^*\exists^*) - \mathsf{LIA}$ the set of formulae consisting of a conjunction between two linear arithmetic formulae, one with quantifier prefix in the language $\exists^*\forall^*$, and another with quantifier prefix $\forall^*\exists^*$.

Theorem 3. *The satisfiability problem is undecidable for the fragment $(\exists^*\forall^* \cap \forall^*\exists^*) - \mathsf{LIA}$, with one monadic predicate symbol, interpreted over finite sets of integers.*

Proof. We reduce from the following variant of *Hilbert's 10th Problem*: given a multivariate Diophantine polynomial $R(x_1, \ldots, x_n)$, the problem "does $R(x_1, \ldots, x_n) = 0$ have a solution in $\mathbb{N}^n$?" is undecidable [18].

By introducing sufficiently many free variables, we encode the equation $R(x_1, \ldots, x_n) = 0$ as an equi-satisfiable Diophantine system of degree at most two, containing only equations of the form $x = yz$ (resp. $x = y^2$) and linear equations $\sum_{i=1}^{k} a_i x_i = b$, where $a_1, \ldots, a_k, b \in \mathbb{Z}$. Next, we replace each equation of the form $x = yz$, with y and z distinct variables, with the quadratic system $2x + t_y + t_z = t_{y+z} \wedge t_y = y^2 \wedge t_z = z^2 \wedge t_{y+z} = (y+z)^2$, where t_y, t_z and t_{y+z} are fresh (free) variables. In this way, we replace all multiplications between distinct variables by occurrences of the squaring function. Let $\Psi_{R(x_1,\ldots,x_n)=0}$ be the conjunction of the above equations. It is manifest that $R(x_1, \ldots, x_n) = 0$ has a solution in $\mathbb{N}^n$ iff $\Psi_{R(x_1,\ldots,x_n)=0}$ is satisfiable, with all free variables ranging over $\mathbb{N}$.

Now we introduce a monadic predicate symbol P, which is intended to denote a (possibly finite) set of consecutive perfect squares, starting with 0. To capture this definition, we require the following:

$$\begin{array}{l} P(0) \wedge P(1) \wedge \forall x \forall y \forall z \,.\, P(x) \wedge P(y) \wedge P(z) \wedge x < y < z \wedge \\ (\forall u \,.\, x < u < y \vee y < u < z \Rightarrow \neg P(u)) \Rightarrow z - y = y - x + 2 \end{array} \qquad \text{(sqr)}$$

Observe that this formula is a weakening of the definition of the infinite set of perfect squares given by Halpern [14], from which the conjunct $\forall x \exists y \,.\, y > x \wedge P(y)$, requiring that P is an infinite set of natural numbers, has been dropped. Moreover, notice that sqr has quantifier prefix $\forall^3\exists$, due to the fact that $\forall u$ occurs under negation, on the left-hand side of an implication. If P is interpreted as a finite set $P^{\mathcal{I}} = \{p_0, p_1, \ldots, p_N\}$ such that (w.l.o.g.) $p_0 < p_1 < \ldots < p_N$, it is easy to show, by induction on $N > 0$, that $p_i = i^2$, for all $i = 0, 1, \ldots, N$.

The next step is encoding the squaring function using the monadic predicate P. This is done by replacing each atomic proposition $x = y^2$ in $\Psi_{R(x_1,\ldots,x_n)=0}$ by the formula $\theta_{x=y^2} \equiv P(x) \wedge P(x + 2y + 1) \wedge \forall z \,.\, x < z < x + 2y + 1 \Rightarrow \neg P(z)$. Let us now prove the following fact:

Fact 2 *For each interpretation $\mathcal{I}$ mapping x and y into $\mathbb{N}$, $\mathcal{I} \models x = y^2$ iff $\mathcal{I}$ can be extended to an interpretation of P as a finite set of consecutive perfect squares such that $\mathcal{I} \models \theta_{x=y^2}$.*

Proof. See Appendix A.4.

Let $\Phi_{R(x_1,\ldots,x_n)=0}$ be the conjunction of sqr with the formula obtained by replacing each atomic proposition $x = y^2$ with $\theta_{x=y^2}$ in $\Psi_{R(x_1,\ldots,x_n)=0}$. Observe that each universally quantified variable in $\Phi_{R(x_1,\ldots,x_n)=0}$ occurs either in sqr or in some $\theta_{x=y^2}$, and moreover, each $\theta_{x=y^2}$ belongs to the $\exists^*\forall^*$ fragment of SLIA. $\Phi_{R(x_1,\ldots,x_n)=0}$ belongs thus to the $\exists^*\forall^* \cap \forall^*\exists^*$ fragment of SLIA, with P being the only monadic predicate symbol. Finally, we prove that $R(x_1, \ldots, x_n) = 0$ has a solution in $\mathbb{N}^n$ iff $\Phi_{R(x_1,\ldots,x_n)=0}$ is satisfiable.

"$\Rightarrow$" Let $\mathcal{I}$ be a valuation mapping $x_1, \ldots, x_n$ into $\mathbb{N}$, such that $\mathcal{I} \models R(x_1, \ldots, x_n) = 0$. Obviously, $\mathcal{I}$ can be extended to a model of $\Psi_{R(x_1,\ldots,x_n)=0}$ by assigning $t_x^{\mathcal{I}} = (x^{\mathcal{I}})^2$ for all auxiliary variables t_x occurring in $\Psi_{R(x_1,\ldots,x_n)=0}$. We extend $\mathcal{I}$ to a model of $\Phi_{R(x_1,\ldots,x_n)=0}$ by assigning $P^{\mathcal{I}} = \{n^2 \mid 0 \leq n \leq \sqrt{m}\}$, where $m = \max\{(x^{\mathcal{I}} + 1)^2 \mid x \in \mathrm{Fvc}(\Psi_{R(x_1,\ldots,x_n)=0})\}$. Clearly $P^{\mathcal{I}}$ meets the requirements of sqr. By Fact 2, we obtain that $\mathcal{I} \models \theta_{x=y^2}$ for each subformula $\theta_{x=y^2}$ of $\Phi_{R(x_1,\ldots,x_n)=0}$, thus $\mathcal{I} \models \Phi_{R(x_1,\ldots,x_n)=0}$.

"$\Leftarrow$" If $\mathcal{I} \models \Psi_{R(x_1,\ldots,x_n)=0}$ then, by sqr, $P^{\mathcal{I}}$ is a set of consecutive perfect squares, and, by Fact 2, $\mathcal{I} \models x = y^2$ for each subformula $\theta_{x=y^2}$ of $\Phi_{R(x_1,\ldots,x_n)=0}$. Then $\mathcal{I} \models \Psi_{R(x_1,\ldots,x_n)=0}$ and consequently $\mathcal{I} \models R(x_1, \ldots, x_n) = 0$. □

We consider now the satisfiability problem for the fragment $\exists^*\forall^*\mathsf{SL}(\mathsf{LIA})_{\mathsf{Int},\mathsf{Int}}$ where both Loc and Data are taken to be the Int sort, equipped with addition and total order. Observe that, in this case, the heap consists of a set of lists, possibly with aliases and circularities. Without losing generality, we consider that Int is interpreted as the set of positive integers[5].

The above theorem cannot be directly used for the undecidability of $\exists^*\forall^*\mathsf{SL}(\mathsf{LIA})_{\mathsf{Int},\mathsf{Int}}$, by interpreting the (unique) monadic predicate as the (finite) domain of the heap. The problem is with the sqr formula, that defines the interpretation of the monadic predicate as a set of consecutive perfect squares $0, 1, \ldots, n^2$, and whose quantifier prefix lies in the $\forall^*\exists^*$ fragment. We overcome this problem by replacing the sqr formula above with a definition of such sets in $\exists^*\forall^*\mathsf{SL}(\mathsf{LIA})_{\mathsf{Int},\mathsf{Int}}$. Let us first consider the following properties expressed in SL [4]:

$$\sharp\mathsf{x} \geq 1 \equiv \exists \mathsf{u} \,.\, \mathsf{u} \mapsto \mathsf{x} * \top$$
$$\sharp\mathsf{x} \leq 1 \equiv \forall \mathsf{u} \forall \mathsf{t} \,.\, \neg(\mathsf{u} \mapsto \mathsf{x} * \mathsf{t} \mapsto \mathsf{x} * \top)$$

Intuitively, $\sharp\mathsf{x} \geq 1$ states that x has at least one predecessor in the heap, whereas $\sharp\mathsf{x} \leq 1$ states that x has at most one predecessor. We use $\sharp\mathsf{x} = 0$ and $\sharp\mathsf{x} = 1$ as shorthands for $\neg(\sharp\mathsf{x} \geq 1)$ and $\sharp\mathsf{x} \geq 1 \wedge \sharp\mathsf{x} \leq 1$, respectively. The formula below states that the heap can be decomposed into a list segment starting with x and ending in y, and several disjoint cyclic lists:

$$\mathsf{x} \overset{\circlearrowleft}{\longrightarrow}^{+} \mathsf{y} \equiv \sharp\mathsf{x} = 0 \wedge \mathsf{alloc}(\mathsf{x}) \wedge \sharp\mathsf{y} = 1 \wedge \neg\mathsf{alloc}(\mathsf{y}) \wedge \forall \mathsf{z} \,.\, \mathsf{z} \not\approx \mathsf{y} \Rightarrow (\sharp\mathsf{z} = 1 \Rightarrow \mathsf{alloc}(\mathsf{z})) \wedge \forall \mathsf{z} \,.\, \sharp\mathsf{z} \leq 1$$

We forbid the existence of circular lists by adding the following arithmetic constraint:

$$\forall \mathsf{u} \forall \mathsf{t} \,.\, \mathsf{u} \mapsto \mathsf{t} * \top \Rightarrow \mathsf{u} < \mathsf{t} \qquad \text{(nocyc)}$$

We ask, moreover, that the elements of the list segment starting in x are consecutive perfect squares:

$$\mathsf{consqr}(\mathsf{x}) \equiv \mathsf{x} = 0 \wedge \mathsf{x} \mapsto 1 * \top \wedge \forall \mathsf{z} \forall \mathsf{u} \forall \mathsf{t} \,.\, \mathsf{z} \mapsto \mathsf{u} * \mathsf{u} \mapsto \mathsf{t} * \top \Rightarrow \mathsf{t} - \mathsf{u} = \mathsf{u} - \mathsf{z} + 2 \quad \text{(consqr)}$$

Observe that the formula $\exists \mathsf{x} \exists \mathsf{y} \,.\, \mathsf{x} \overset{\circlearrowleft}{\longrightarrow}^{+} \mathsf{y} \wedge \mathsf{nocyc} \wedge \mathsf{consqr}(\mathsf{x})$ belongs to $\exists^*\forall^*\mathsf{SL}(\mathsf{LIA})_{\mathsf{Int},\mathsf{Int}}$.

Theorem 4. *The satisfiability problem for* $\exists^*\forall^*\mathsf{SL}(\mathsf{LIA})_{\mathsf{Int},\mathsf{Int}}$ *is undecidable.*

Proof. We use the same reduction as in the proof of Theorem 3, with two differences:

- we replace sqr by $\exists \mathsf{x} \exists \mathsf{y} \,.\, \mathsf{x} \overset{\circlearrowleft}{\longrightarrow}^{+} \mathsf{y} \wedge \mathsf{nocyc} \wedge \mathsf{consqr}(\mathsf{x})$, and
- define $\theta_{x=y^2} \equiv \mathsf{alloc}(x) \wedge \mathsf{alloc}(x + 2y + 1) \wedge \forall z \,.\, x < z < x + 2y + 1 \Rightarrow \neg\mathsf{alloc}(z)$. □

It is tempting, at this point to ask whether interpreting locations as integers and considering subsets of LIA instead may help recover the decidability. For instance, it has been found that the Bernays-Schönfinkel-Ramsey class is decidable in presence of

[5] Extending the interpretation of Loc to include negative integers does not make any difference for the undecidability result.

integers with difference bounds arithmetic [28], and the same type of question can be asked about the fragment of $\exists^*\forall^*\mathsf{SL}(\mathsf{LIA})_{\mathsf{Int},\mathsf{Int}}$, with difference bounds constraints only.

Finally, we consider a variant of the previous undecidability result, in which locations are the (uninterpreted) sort U of E and the data consists of tuples of sort $U \times \mathsf{Int}$. This fragment of SL can be used to reason about lists with integer data. The undecidability of this fragment can be proved along the same lines as Theorem 4.

Theorem 5. *The satisfiability problem for $\exists^*\forall^*\mathsf{SL}(\mathsf{ELIA})_{U,U\times\mathsf{Int}}$ is undecidable.*

Proof. Along the same lines as the proof of Theorem 4, with the following shorthands:

$$\sharp \mathsf{x} \geq 1 \equiv \exists \mathsf{u}^U \exists d^{\mathsf{Int}} \,.\, \mathsf{u} \mapsto (d, \mathsf{x}) * \top$$
$$\sharp \mathsf{x} \leq 1 \equiv \forall \mathsf{u}^U \forall \mathsf{t}^U \forall d^{\mathsf{Int}} \,.\, \neg(\mathsf{u} \mapsto (d, \mathsf{x}) * \mathsf{t} \mapsto (d, \mathsf{x}) * \top)$$

$$\mathsf{nocyc} \equiv \forall \mathsf{u}^U \forall \mathsf{t}^U \forall \mathsf{v}^U \forall d^{\mathsf{Int}} \forall e^{\mathsf{Int}} \,.\, \mathsf{u} \mapsto (d, \mathsf{t}) * \mathsf{t} \mapsto (e, \mathsf{v}) * \top \Rightarrow d < e$$

$$\begin{aligned}\mathsf{consqr}(\mathsf{x}) \equiv\ & \exists \mathsf{y}^U \exists \mathsf{z}^U \,.\, \mathsf{x} \mapsto (0, \mathsf{y}) * \mathsf{y} \mapsto (1, \mathsf{z}) * \top \wedge \\ & \forall \mathsf{u}^U \forall \mathsf{t}^U \forall \mathsf{v}^U \forall \mathsf{w}^U \forall d^{\mathsf{Int}} \forall e^{\mathsf{Int}} \forall f^{\mathsf{Int}} \,.\, \mathsf{u} \mapsto (d, \mathsf{t}) * \mathsf{t} \mapsto (e, \mathsf{v}) * \mathsf{v} \mapsto (f, \mathsf{w}) * \top \Rightarrow \\ & f - e = e - d + 2\end{aligned}$$

□

4 A Procedure for $\exists^*\forall^*$ Separation Logic in an SMT Solver

In previous work [24], we developed a decision procedure for quantifier-free $\mathsf{SL}(T)_{\mathsf{Loc},\mathsf{Data}}$ inputs where the satisfiability problem for quantifier-free T-constraints is decidable. The procedure is implemented in the SMT solver CVC4 [2] [6]. This section presents a procedure for the satisfiability of $\exists^*\forall^*\mathsf{SL}(\mathsf{E})_{U,U^k}$ inputs which builds on this procedure. Like existing approaches for quantified formulas in SMT [13, 23], our approach is based on incremental quantifier instantiation based on a stream of candidate models returned by a solver for quantifier-free inputs. Our approach for this fragment exploits the small model property given in Lemma 3 to restrict the set of quantifier instantiations it considers to a finite set.

Figure 1 gives a counterexample-guided approach for establishing the satisfiability of input $\exists \mathbf{x} \forall \mathbf{y} \varphi(\mathbf{x}, \mathbf{y})$. We first introduce tuples of fresh constants $\mathbf{k}$ and $\mathbf{e}$ of the same type as $\mathbf{x}$ and $\mathbf{y}$ respectively. Our procedure will be based on finding a set of instantiations of $\forall \mathbf{y} \varphi(\mathbf{k}, \mathbf{y})$ that are either collectively unsatisfiable or are satisfiable and entail our input. Then, we construct a set L which is the union of constants $\mathbf{k}$ and a set L' of fresh constants whose cardinality is equal to the cardinality of $\varphi(\mathbf{x}, \mathbf{y})$ plus the number of universal variables n in our input. Conceptually, L is a finite set of terms from which the instantiations of $\mathbf{y}$ in $\forall \mathbf{y} \varphi(\mathbf{k}, \mathbf{y})$ can be built.

After constructing L, we call the recursive subprocedure solve_rec on Γ (initially empty) and L. This procedure incrementally adds instances of $\forall \mathbf{y} \varphi(\mathbf{k}, \mathbf{y})$ to Γ. In step 1, we first check if Γ is (SL, T)-unsatisfiable using the procedure from [24]. If so, our

[6] The procedure is incorporated into the master branch of CVC4 (`https://github.com/CVC4`), and can be enabled by command line parameter `--quant-epr`.

$\mathsf{solve}(\exists \mathbf{x} \forall \mathbf{y}\, \varphi(\mathbf{x}, \mathbf{y}))$ where $\mathbf{x} = (x_1, \ldots, x_m)$ and $\mathbf{y} = (y_1, \ldots, y_n)$:

Let $\mathbf{k} = (k_1, \ldots, k_m)$ and $\mathbf{e} = (e_1, \ldots, e_n)$ be fresh constants of the same type as $\mathbf{x}$ and $\mathbf{y}$.
Let $L = L' \cup \{k_1, \ldots, k_m\}$ where L' is a set of fresh constants s.t. $\|L'\| = |\varphi(\mathbf{x}, \mathbf{y})| + n$.
Return $\mathsf{solve\_rec}(\exists \mathbf{x} \forall \mathbf{y}\, \varphi(\mathbf{x}, \mathbf{y}), \emptyset, L)$.

$\mathsf{solve\_rec}(\exists \mathbf{x} \forall \mathbf{y}\, \varphi(\mathbf{x}, \mathbf{y}), \Gamma, L)$:

1. If Γ is $(\mathsf{SL}, \mathsf{E})$-unsat, return "unsat".
2. Assume $\exists \mathbf{x} \forall \mathbf{y}\, \varphi(\mathbf{x}, \mathbf{y})$ is equivalent to $\exists \mathbf{x} \forall \mathbf{y}\, \varphi_1(\mathbf{x}, \mathbf{y}) \wedge \ldots \wedge \forall \mathbf{y}\, \varphi_p(\mathbf{x}, \mathbf{y})$.
   If $\Gamma'_j = \Gamma \cup \{\neg \varphi_j(\mathbf{k}, \mathbf{e}) \wedge \bigwedge_{i=1}^{n} \bigvee_{t \in L} e_i \approx t\}$ is $(\mathsf{SL}, \mathsf{E})$-unsat for all $j = 1, \ldots, p$, return "sat".
3. Otherwise, let $\mathcal{I}, h \models_{\mathsf{SL}} \Gamma'_j$ for some $j \in \{1, \ldots, p\}$.
   Let $\mathbf{t} = (t_1, \ldots, t_n)$ be such that $e_i^{\mathcal{I}} = t_i^{\mathcal{I}}$ and $t_i \in L$ for each $i = 1, \ldots, n$.
   Return $\mathsf{solve\_rec}(\exists \mathbf{x} \forall \mathbf{y}\, \varphi(\mathbf{x}, \mathbf{y}), \Gamma \cup \{\varphi_j(\mathbf{k}, \mathbf{t})\}, L)$.

Fig. 1. A counterexample-guided procedure for $\exists^* \forall^* \mathsf{SL}(\mathsf{E})_{U,U^k}$ formulas $\exists \mathbf{x} \forall \mathbf{y}\, \varphi(\mathbf{x}, \mathbf{y})$, where U is an uninterpreted sort in the signature of E.

input is (SL, T)-unsatisfiable. Otherwise, in step 2 we consider the *miniscoped* form of our input $\exists \mathbf{x} \forall \mathbf{y}\, \varphi_1(\mathbf{x}, \mathbf{y}) \wedge \ldots \wedge \forall \mathbf{y}\, \varphi_p(\mathbf{x}, \mathbf{y})$. In the following, we may omit quantification on conjunctions φ_j that do not contain variables from $\mathbf{y}$. Given this formula, for each $j = 1, \ldots, p$, we check the (SL, T)-satisfiability of set Γ'_j containing Γ, the negation of $\forall \mathbf{y}\, \varphi_j(\mathbf{k}, \mathbf{y})$ where $\mathbf{y}$ is replaced by fresh contants $\mathbf{e}$, and a conjunction of constraints that says each e_i must be equal to at least one term in L for $i = 1, \ldots, n$. If Γ'_j is (SL, T)-unsatisfiable for each $j = 1, \ldots, p$, our input is (SL, T)-satisfiable. Otherwise in step 3, given an interpretation $\mathcal{I}$ and heap h satisfying Γ'_j, we construct a tuple of terms $\mathbf{t} = (t_1, \ldots, t_n)$ used for instantiating $\forall \mathbf{y}\, \varphi_j(\mathbf{k}, \mathbf{y})$. For each $i = 1, \ldots, n$, we choose t_i to be a term from L whose interpretation is the same as e_i. The existence of such a t_i is guaranteed by the fact that $\mathcal{I}$ satisfies the constraint from Γ'_j that tells us e_i is equal to at least one such term. This selection ensures that instantiations on each iteration are chosen from a finite set of possibilities and are unique. In practice, the procedure terminates, both for unsatisfiable and satisfiable inputs, before considering all $\mathbf{t}$ from L^n for each $\forall \mathbf{y}\, \varphi_j(\mathbf{x}, \mathbf{y})$.

Theorem 6. *Let U be an uninterpreted sort belonging to the signature of E. For all $\exists^* \forall^* \mathsf{SL}(\mathsf{E})_{U,U^k}$ formulae ψ of the form $\exists \mathbf{x} \forall \mathbf{y}\, \varphi(\mathbf{x}, \mathbf{y})$, $\mathsf{solve}(\psi)$:*

1. *Answers "unsat" only if ψ is $(\mathsf{SL}, \mathsf{E})$-unsatisfiable.*
2. *Answers "sat" only if ψ is $(\mathsf{SL}, \mathsf{E})$-satisfiable.*
3. *Terminates.*

Proof. To show (1), note that Γ contains only formulas of the form $\varphi_j(\mathbf{k}, \mathbf{t})$, which are consequences of our input. Thus, when Γ is $(\mathsf{SL}, \mathsf{E})$-unsatisfiable, our input is $(\mathsf{SL}, \mathsf{E})$-unsatisfiable as well.

To show (2), we have that Γ is $(\mathsf{SL}, \mathsf{E})$-satisfiable and $\Gamma' = \Gamma \cup \{\neg\varphi_j(\mathbf{k}, \mathbf{e}) \wedge A\}$ is $(\mathsf{SL}, \mathsf{E})$-unsatisfiable for each $j = 1, \ldots, p$, where:

$$A = \bigwedge_{i=1}^{n} \bigvee_{t \in L} e_i \approx t$$

where $L = L' \cup \{k_1, \ldots, k_m\}$ and L' is a set of fresh constants s.t. $\|L'\| = |\varphi(\mathbf{x}, \mathbf{y})| + n$. In other words, we have that all models of Γ satisfy $(\varphi_1(\mathbf{k}, \mathbf{e}) \wedge \ldots \wedge \varphi_p(\mathbf{k}, \mathbf{e})) \vee \neg A$, which is equivalent to $\varphi(\mathbf{k}, \mathbf{e}) \vee \neg A$. Since $\mathbf{e}$ is not contained in Γ, we have that all models of Γ satisfy $\forall \mathbf{y}\, (\varphi(\mathbf{k}, \mathbf{y}) \vee \neg A\{\mathbf{e} \mapsto \mathbf{y}\})$. Since Γ is $(\mathsf{SL}, \mathsf{E})$-satisfiable, we have that $\forall \mathbf{y}\, (\varphi(\mathbf{k}, \mathbf{y}) \vee \neg A\{\mathbf{e} \mapsto \mathbf{y}\})$ is $(\mathsf{SL}, \mathsf{E})$-satisfiable as well. Consider the formula $\forall \mathbf{y}\, \varphi(\mathbf{k}, \mathbf{y})$. By Lemma 3, $\forall \mathbf{y}\, \varphi(\mathbf{k}, \mathbf{y})$ has a model if and only if there exists an interpretation $\mathcal{I}$ and heap h such that $\mathcal{I}, h \models_{\mathsf{SL}} \forall \mathbf{y}\, \varphi(\mathbf{k}, \mathbf{y})$ and $\|U^{\mathcal{I}}\| \leq |\varphi| + \|\mathrm{Fvc}(\forall \mathbf{y}\, \varphi(\mathbf{x}, \mathbf{y}))\| + n$. Due to the construction of L, this implies that $\forall \mathbf{y}\, (\varphi(\mathbf{k}, \mathbf{y}) \vee \neg A\{\mathbf{e} \mapsto \mathbf{y}\})$ is $(\mathsf{SL}, \mathsf{E})$-satisfiable if and only if $\forall \mathbf{y}\, \varphi(\mathbf{k}, \mathbf{y})$ is $(\mathsf{SL}, \mathsf{E})$-satisfiable. Thus, $\exists \mathbf{x} \forall \mathbf{y}\, \varphi(\mathbf{x}, \mathbf{y})$ is $(\mathsf{SL}, \mathsf{E})$-satisfiable.

To show (3), clearly only a finite number of possible formulas can be added to Γ as a result of the procedure, since all terms $\mathbf{t}$ belong to the finite set L and p is finite. Furthermore, on every iteration, for any j, $\mathcal{I}$ satisfies Γ and $\neg\varphi_j(\mathbf{k}, \mathbf{e})$. Since $e_i^{\mathcal{I}} = t_i^{\mathcal{I}}$ for each $i = 1, \ldots, n$, we have that $\varphi_j(\mathbf{k}, \mathbf{t}) \notin \Gamma$, and thus a new formula is added to Γ on every call. Thus, only a finite number of recursive calls are made to solve_rec. Since the $(\mathsf{SL}, \mathsf{E})$-satisfiability of quantifier-free is decidable, all steps in the procedure are terminating, and thus solve terminates. □

We discuss a few important details regarding our implementation of the procedure.
Matching Heuristics When constructing the terms $\mathbf{t}$ for instantiation, it may be the case that $e_i^{\mathcal{I}} = u^{\mathcal{I}}$ for multiple $u \in L$ for some $i \in \{1, \ldots, n\}$. In such cases, the procedure will choose one such u for instantiation. To increase the likelihood of the instantiation being relevant to the satisfiability of our input, we use heuristics for selecting the best possible u among those whose interpretation is equal to e_i in $\mathcal{I}$. In particular, if $e_i^{\mathcal{I}} = u_1^{\mathcal{I}} = u_2^{\mathcal{I}}$, and Γ' contains predicates of the form $e_i \mapsto v$ and $u_1 \mapsto v_1$ for some v, v_1 where $v^{\mathcal{I}} = v_1^{\mathcal{I}}$ but no predicate of the form $u_2 \mapsto v_2$ for some v_2 where $v^{\mathcal{I}} = v_2^{\mathcal{I}}$, then we strictly prefer term u_1 over term u_2 when choosing term t_i for e_i.
Finding Minimal Models Previous work [25] developed efficient techniques for finding small models for uninterpreted sorts in CVC4. We have found these techniques to be beneficial to the performance of the procedure in Figure 1. In particular, we use these techniques to find Σ-interpretations $\mathcal{I}$ in solve_rec that interpret U as a finite set of minimal size. When combined with the aforementioned matching heuristics, these techniques lead to finding useful instantiations more quickly, since more terms are constrained to be equal to e_i for $i = 1, \ldots, n$ in interpretations $\mathcal{I}$.
Symmetry Breaking The procedure in Figure 1 introduces a set of fresh constants L, which in turn introduce the possibility of discovering Σ-interpretations $\mathcal{I}$ that are isomorphic, that is, identical up to renaming of constants in L'. Our procedure adds additional constraints to Γ that do not affect its satisfiability, but reduce the number of isomorphic models. In particular, we consider an ordering $\prec$ on the constants from L', and add constraints that ensure that all models $(\mathcal{I}, h)$ of Γ are such that if $\ell_1^{\mathcal{I}} \notin \mathrm{dom}(h)$, then $\ell_2^{\mathcal{I}} \notin \mathrm{dom}(h)$ for all ℓ_2 such that $\ell_1 \prec \ell_2$.

Example 1. Say we wish to show the validity of the entailment $\mathsf{x} \neq \mathsf{y} \wedge \mathsf{x} \mapsto \mathsf{z} \models_{\mathsf{SL}} \exists \mathsf{u} \exists \mathsf{x} \mapsto \mathsf{u}$, from the introductory example (section 1), where $\mathsf{x}, \mathsf{y}, \mathsf{z}, \mathsf{u}$ are of the (uninterpreted) sort U of E. This entailment is valid iff the $\exists^*\forall^*\mathsf{SL}(\mathsf{E})_{U,U^k}$ formula $\exists \mathsf{x} \exists \mathsf{y} \exists \mathsf{z} \forall \mathsf{u} \,.\, \mathsf{x} \not\approx \mathsf{y} \wedge \mathsf{x} \mapsto \mathsf{z} \wedge \neg \mathsf{x} \mapsto \mathsf{u}$ is $(\mathsf{SL}, \mathsf{E})$-unsatisfiable. A run of the procedure in Figure 1 on this input constructs tuples $\mathbf{k} = (k_x, k_y, k_z)$ and $\mathbf{e} = (e_u)$, and set $L = \{k_x, k_y, k_z, \ell_1, \ell_2\}$, noting that $|\mathsf{x} \not\approx \mathsf{y} \wedge \mathsf{x} \mapsto \mathsf{z} \wedge \neg \mathsf{x} \mapsto \mathsf{u}| = 1$. We then call solve_rec where Γ is initially empty. By miniscoping, our input is equivalent to $\exists \mathsf{x} \exists \mathsf{y} \exists \mathsf{z} \,.\, \mathsf{x} \not\approx \mathsf{y} \wedge \mathsf{x} \mapsto \mathsf{z} \wedge \forall \mathsf{u} \,.\, \neg \mathsf{x} \mapsto \mathsf{u}$. On the first two recursive calls to solve_rec, we may add $k_x \not\approx k_y$ and $k_x \mapsto k_z$ to Γ by trivial instantiation of the first two conjuncts. On the third recursive call, Γ is $(\mathsf{SL}, \mathsf{E})$-satisfiable, and we check the satisfiability of:

$$\Gamma' = \{k_x \neq k_y, k_x \mapsto k_z, k_x \mapsto e_u \wedge (e_u \approx k_x \vee e_u \approx k_y \vee e_u \approx k_z \vee e_u \approx \ell_1 \vee e_u \approx \ell_2)\}$$

Since $k_x \mapsto k_z$ and $k_x \mapsto e_u$ are in Γ', all Σ-interpretations $\mathcal{I}$ and heaps h such that $\mathcal{I}, h \models_{\mathsf{SL}} \Gamma'$ are such that $e_u^{\mathcal{I}} = k_z^{\mathcal{I}}$. Since $k_z \in L$, we may choose to add the instantiation $\neg k_x \mapsto k_z$ to Γ, after which Γ is $(\mathsf{SL}, \mathsf{E})$-unsatisfiable on the next recursive call to solve_rec. Thus, our input is $(\mathsf{SL}, \mathsf{E})$-unsatisfiable and the entailment is valid. ■

A modified version of the procedure in Figure 1 can be used for $\exists^*\forall^*\mathsf{SL}(T)_{\mathsf{Loc},\mathsf{Data}}$-satisfiability for theories T beyond equality, and where Loc and Data are not restricted to uninterpreted sorts. Notice that in such cases, we cannot restrict Σ-interpretations $\mathcal{I}$ in solve_rec to interpret each e_i as a member of finite set L, and hence we modify solve_rec to omit the constraint restricting variables in $\mathbf{e}$ to be equal to a term from L in the check in Step 2. This modification results in a procedure that is sound both for "unsat" and "sat", but is no longer terminating in general. Nevertheless, it may be used as a heuristic for determining $\exists^*\forall^*\mathsf{SL}(T)_{\mathsf{Loc},\mathsf{Data}}$-(un)satisfiability.

5 Experimental Evaluation

We implemented the solve procedure from Figure 1 within the CVC4 SMT solver[7] (version 1.5 prerelease). This implementation was tested on two kinds of benchmarks: (i) finite unfoldings of inductive predicates, mostly inspired by benchmarks used in the SL-COMP'14 solver competition [26], and (ii) verification conditions automatically generated by applying the weakest precondition calculus of [15] to the program loops in Figure 2. All experiments were run on a 2.80GHz Intel(R) Core(TM) i7 CPU machine with with 8MB of cache [8].

We compared our implementation with the results of applying the CVC4 decision procedure for the quantifier-free fragment of SL [24] to a variant of the benchmarks, obtained by manual quantifier instantiation, as follows. Consider checking the validity of the entailment $\exists \mathbf{x} \,.\, \phi(\mathbf{x}) \models_{\mathsf{SL}} \exists \mathbf{y} \,.\, \psi(\mathbf{y})$, which is equivalent to the unsatisfiability of the formula $\exists \mathbf{x} \forall \mathbf{y} \,.\, \phi(\mathbf{x}) \wedge \neg \psi(\mathbf{y})$. We first check the satisfiability of ϕ. If ϕ is not satisfiable, the entailment holds trivially, so let us assume that ϕ has a model. Second, we check the

[7] Available at `http://cvc4.cs.nyu.edu/web/`.

[8] The CVC4 binary and examples used in these experiments are available at `http://cs.uiowa.edu/~ajreynol/VMCAI2017-seplog-epr`.

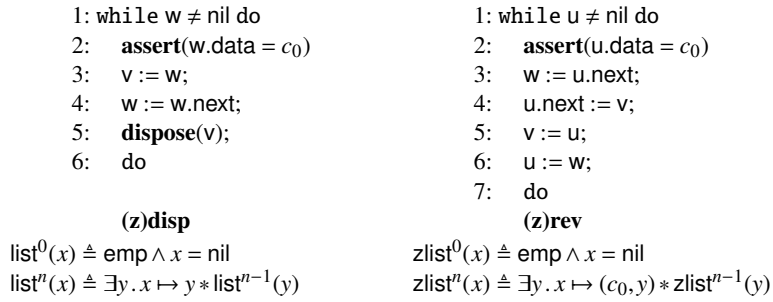

Fig. 2. Program Loops

satisfiability of $\phi \wedge \psi$. Again, if this is unsatisfiable, the entailment cannot hold, because there exists a model of ϕ which is not a model of ψ. Else, if $\phi \wedge \psi$ has a model, we add an equality $x = y$ for each pair of variables $(x, y) \in \mathbf{x} \times \mathbf{y}$ that are mapped to the same term in this model, the result being a conjunction $E(\mathbf{x}, \mathbf{y})$ of equalities. Finally, we check the satisfiability of the formula $\phi \wedge \neg\psi \wedge E$. If this formula is unsatisfiable, the entailment is valid, otherwise, the check is inconclusive. The times in Table 1 correspond to checking satisfiability of $\exists\mathbf{x}\forall\mathbf{y}\ .\ \phi(\mathbf{x}) \wedge \neg\psi(\mathbf{y})$ using the solve procedure (Figure 1), compared to checking satisfiability of $\phi \wedge \neg\psi \wedge E$, where E is manually generated.

In the first set of experiments (Table 1) we have considered inductive predicates commonly used as verification benchmarks [26]. Here we check the validity of the entailment between lhs and rhs, where both predicates are unfolded $n = 1, 2, 3, 4, 8$ times. The entailment between pos_2^1 and neg_4^1 is skipped because it is not valid (since the negated formula is satisfiable, we cannot generate the manual instantiation).

The second set of experiments considers the verification conditions of the forms $\varphi \Rightarrow \mathsf{wp}(\mathbf{l}, \phi)$ and $\varphi \Rightarrow \mathsf{wp}^n(\mathbf{l}, \phi)$, where $\mathsf{wp}(\mathbf{l}, \phi)$ denotes the weakest precondition of the SL formula ϕ with respect to the sequence of statements $\mathbf{l}$, and $\mathsf{wp}^n(\mathbf{l}, \phi) = \mathsf{wp}(\mathbf{l}, \ldots \mathsf{wp}(\mathbf{l}, \mathsf{wp}(\mathbf{l}, \phi))\ldots)$ denotes the iterative application of the weakest precondition n times in a row. We consider the loops depicted in Figure 2, where, for each loop $\mathbf{l}$, we consider the variant $\mathbf{zl}$ as well, which tests that the data values contained within the memory cells are equal to a constant c_0 of sort Loc, by the assertions on line 2. The postconditions are specified by finite unfoldings of the inductive predicates list and zlist.

We observed that the fully automated solver was less than 1.5 seconds slower than checking the manual instantiation, on 82% of the test cases. The automated solver experienced 3 timeouts, where the manual instantiation succeeds (for $\widehat{\mathsf{tree}}$ vs tree with $n = 8$, $\widehat{\mathsf{ts}}$ vs ts with $n = 3$, and $\mathsf{list}^n(u) * \mathsf{list}^0(v)$ vs $\mathsf{wp}^n(\mathbf{rev}, u = \mathsf{nil} \wedge \mathsf{list}^n(v))$ with $n = 8$). These timeouts are caused by the first call to the quantifier-free SL decision procedure, which fails to produce a model in less than 300 seconds (time not accounted for in the manually produced instance of the problem).

6 Conclusions and Future Work

We present theoretical and practical results for the existence of effective decision procedures for the fragment of Separation Logic obtained by restriction of formulae to

lhs	rhs		$n=1$	$n=2$	$n=3$	$n=4$	$n=8$
Unfoldings of inductive predicates							
$\widehat{\mathsf{ls}}(x,y)\triangleq\mathsf{emp}\wedge x=y\vee$	$\mathsf{ls}(x,y)\triangleq\mathsf{emp}\wedge x=y\vee$	solve	<0.01s	0.02s	0.03s	0.05s	0.21s
$\exists z.\, x\neq y\wedge x\mapsto z*\widehat{\mathsf{ls}}(z,y)$	$\exists z.\, x\mapsto z*\mathsf{ls}(z,y)$	manual	<0.01s	<0.01s	<0.01s	<0.01s	<0.01s
$\widehat{\mathsf{tree}}(x)\triangleq\mathsf{emp}\wedge x=\mathsf{nil}\vee$	$\mathsf{tree}(x)\triangleq\mathsf{emp}\wedge x=\mathsf{nil}\vee$	solve	<0.01s	0.04s	1.43s	23.42s	>300s
$\exists l\exists r.\, l\neq r\wedge x\mapsto(l,r)*\mathsf{tree}(l)*\mathsf{tree}(r)$	$\exists l\exists r.\, x\mapsto(l,r)*\mathsf{tree}(l)*\mathsf{tree}(r)$	manual	<0.01s	<0.01s	<0.01s	<0.01s	0.09s
$\widehat{\mathsf{ts}}(x,a)\triangleq\mathsf{emp}\wedge x=\mathsf{nil}\vee$	$\mathsf{ts}(x,a)\triangleq\mathsf{emp}\wedge x=\mathsf{nil}\vee$	solve	<0.01s	0.81s	>300s	>300s	>300s
$\exists l\exists r.\, x\neq y\wedge x\mapsto(l,r)*\widehat{\mathsf{ts}}(l,y)*\mathsf{tree}(r)\vee$	$\exists l\exists r.\, \wedge x\mapsto(l,r)*\mathsf{ts}(l,y)*\mathsf{tree}(r)\vee$	manual	<0.01s	0.03s	103.89s	>300s	>300s
$\exists l\exists r.\, x\neq y\wedge x\mapsto(l,r)*\mathsf{tree}(l)*\widehat{\mathsf{ts}}(r,y)$	$\exists l\exists r.\, \wedge x\mapsto(l,r)*\mathsf{tree}(l)*\mathsf{ts}(r,y)$						
$\mathsf{pos}_1(x,a)\triangleq x\mapsto a\vee\exists y\exists b.$	$\mathsf{neg}_1(x,a)\triangleq\neg x\mapsto a\vee\exists y\exists b.$	solve	0.34s	0.01s	0.31s	0.76s	21.19s
$x\mapsto a*\mathsf{pos}_1(y,b)$	$x\mapsto a*\mathsf{neg}_1(y,b)$	manual	0.04s	0.05s	0.08s	0.12s	0.53s
$\mathsf{pos}_1(x,a)\triangleq x\mapsto a\vee\exists y\exists b.$	$\mathsf{neg}_2(x,a)\triangleq x\mapsto a\vee\exists y\exists b.$	solve	0.03s	0.12s	0.23s	0.46s	3.60s
$x\mapsto a*\mathsf{pos}_1(y,b)$	$\neg x\mapsto a*\mathsf{neg}_2(y,b)$	manual	0.05s	0.08s	0.08s	0.12s	0.54s
$\mathsf{pos}_2(x,a)\triangleq x\mapsto a\vee\exists y.$	$\mathsf{neg}_3(x,a)\triangleq\neg x\mapsto a\vee\exists y.$	solve	0.04s	0.13s	0.28s	0.48s	4.20s
$x\mapsto a*\mathsf{pos}_2(a,y)$	$x\mapsto a*\mathsf{neg}_3(a,y)$	manual	0.01s	0.03s	0.05s	0.09s	0.45s
$\mathsf{pos}_2(x,a)\triangleq x\mapsto a\vee\exists y.$	$\mathsf{neg}_4(x,a)\triangleq x\mapsto a\vee\exists y.$	solve	—	0.08s	0.15s	0.26s	1.33s
$x\mapsto a*\mathsf{pos}_2(a,y)$	$\neg x\mapsto a*\mathsf{neg}_4(a,y)$	manual	—	0.03s	0.06s	0.09s	0.46s
Verification conditions							
$\mathsf{list}^n(w)$	$\mathsf{wp}(\mathbf{disp},\mathsf{list}^{n-1}(w))$	solve	0.01s	0.03s	0.08s	0.19s	1.47s
		manual	<0.01s	0.01s	0.02s	0.05s	0.26s
$\mathsf{list}^n(w)$	$\mathsf{wp}^n(\mathbf{disp},\mathsf{emp}\wedge w=\mathsf{nil})$	solve	0.01s	0.06s	0.17s	0.53s	7.08s
		manual	<0.01s	0.02s	0.08s	0.14s	2.26s
$\mathsf{zlist}^n(w)$	$\mathsf{wp}(\mathbf{zdisp},\mathsf{zlist}^{n-1}(w))$	solve	0.04s	0.05s	0.09s	0.19s	1.25s
		manual	<0.01s	0.01s	0.02s	0.04s	0.29s
$\mathsf{zlist}^n(w)$	$\mathsf{wp}^n(\mathbf{zdisp},\mathsf{emp}\wedge w=\mathsf{nil})$	solve	0.01s	0.10s	0.32s	0.87s	11.88s
		manual	0.01s	0.02s	0.07s	0.15s	2.20s
$\mathsf{list}^n(u)*\mathsf{list}^0(v)$	$\mathsf{wp}(\mathbf{rev},\mathsf{list}^{n-1}(u)*\mathsf{list}^1(v))$	solve	0.38s	0.06s	0.11s	0.16s	0.56s
		manual	0.07s	0.03s	0.07s	0.11s	0.43s
$\mathsf{list}^n(u)*\mathsf{list}^0(v)$	$\mathsf{wp}^n(\mathbf{rev},u=\mathsf{nil}\wedge\mathsf{list}^n(v))$	solve	0.38s	0.07s	0.30s	68.68s	>300s
		manual	0.08s	0.06s	0.11s	0.23s	1.79s
$\mathsf{zlist}^n(u)*\mathsf{zlist}^0(v)$	$\mathsf{wp}(\mathbf{zrev},\mathsf{zlist}^{n-1}(u)*\mathsf{zlist}^1(v))$	solve	0.22s	0.07s	0.15s	0.21s	0.75s
		manual	0.04s	0.02s	0.04s	0.06s	0.31s
$\mathsf{zlist}^n(u)*\mathsf{zlist}^0(v)$	$\mathsf{wp}^n(\mathbf{zrev},u=\mathsf{nil}\wedge\mathsf{zlist}^n(v))$	solve	0.23s	0.09s	0.17s	0.30s	2.06s
		manual	0.04s	0.02s	0.05s	0.09s	0.48s

Table 1. Experimental results

quantifier prefixes in the set $\exists^*\forall^*$. The theoretical results range from undecidability, when the set of memory locations is taken to be the set of integers and linear arithmetic constraints are allowed, to PSPACE-completeness, when locations and data in the cells belong to an uninterpreted sort, equipped with equality only. We have implemented a decision procedure for the latter case in the CVC4 SMT solver, using an effective counterexample-driven instantiation of the universal quantifiers. The procedure is shown to be sound, complete and termination is guaranteed when the input belongs to a decidable fragment of SL.

As future work, we aim at refining the decidability chart for $\exists^*\forall^*\mathsf{SL}(T)_{\mathsf{Loc,Data}}$, by considering the case where the locations are interpreted as integers, with weaker arithmetics, such as sets of difference bounds, or octagonal constraints. These results are likely to extend the application range of our tool, to e.g. solvers working on SL with inductive definitions and data constraints. The current implementation should also benefit from improvements of the underlying quantifier-free SL and set theory solvers.

References

1. Albargouthi, A., Berdine, J., Cook, B., Kincaid, Z.: Spatial Interpolants, pp. 634–660. Springer (2015)
2. Barrett, C., Conway, C., Deters, M., Hadarean, L., Jovanovic, D., King, T., Reynolds, A., Tinelli, C.: CVC4. In: Computer Aided Verification (CAV). Springer (2011)
3. Baumgartner, P., Fuchs, A., Tinelli, C.: Implementing the model evolution calculus. International Journal on Artificial Intelligence Tools 15(1), 21–52 (2006)
4. Brochenin, R., Demri, S., Lozes, E.: On the almighty wand. Information and Computation 211, 106 – 137 (2012)
5. Brotherston, J., Simpson, A.: Sequent calculi for induction and infinite descent. Journal of Logic and Computation 21(6), 1177–1216 (December 2011)
6. Calcagno, C., Distefano, D.: Infer: An automatic program verifier for memory safety of c programs. In: Proc. of NASA Formal Methods'11. LNCS, vol. 6617. Springer (2011)
7. Calcagno, C., Yang, H., Ohearn, P.W.: Computability and complexity results for a spatial assertion language for data structures. In: FST TCS 2001, Proceedings, pp. 108–119. Springer (2001)
8. Demri, S., Deters, M.: Two-variable separation logic and its inner circle. ACM Transactions on Computational Logic 16(2:15) (2015)
9. Demri, S., Galmiche, D., Larchey-Wendling, D., Méry, D.: Separation logic with one quantified variable. In: Proceedings of the 9th International Computer Science Symposium in Russia (CSR'14). Lecture Notes in Computer Science, vol. 8476, pp. 125–138. Springer (2014)
10. Dudka, K., Peringer, P., Vojnar, T.: Predator: A practical tool for checking manipulation of dynamic data structures using separation logic. In: Proc. of CAV'11. LNCS, vol. 6806. Springer (2011)
11. Galmiche, D., Méry, D.: Tableaux and resource graphs for separation logic. Journal of Logic and Computation 20(1), 189–231 (2010)
12. Ganzinger, H., Hagen, G., Nieuwenhuis, R., Oliveras, A., Tinelli, C.: Dpll (t): Fast decision procedures. In: CAV 2004, Proceedings, pp. 175–188 (2004)
13. Ge, Y., de Moura, L.: Complete instantiation for quantified formulas in satisfiability modulo theories. In: Proceedings of CAV'09. LNCS, vol. 5643 (2009)
14. Halpern, J.Y.: Presburger arithmetic with unary predicates is π_1^1 complete. The Journal of Symbolic Logic 56(2), 637–642 (1991)
15. Ishtiaq, S.S., O'Hearn, P.W.: Bi as an assertion language for mutable data structures. In: ACM SIGPLAN Notices. vol. 36, pp. 14–26 (2001)
16. Korovin, K.: iprover - an instantiation-based theorem prover for first-order logic (system description). In: Automated Reasoning, 4th International Joint Conference, IJCAR 2008, Sydney, Australia, August 12-15, 2008, Proceedings. pp. 292–298 (2008)
17. Lewis, H.R.: Complexity results for classes of quantificational formulas. Journal of Computer and System Sciences 21(3), 317 – 353 (1980)
18. Matiyasevich, Y.: Enumerable sets are diophantine. Journal of Sovietic Mathematics 11, 354 – 358 (1970)
19. Nguyen, H.H., Chin, W.N.: Enhancing program verification with lemmas. In: Proc of CAV'08. LNCS, vol. 5123. Springer (2008)
20. Piskac, R., de Moura, L.M., Bjørner, N.: Deciding effectively propositional logic using DPLL and substitution sets. J. Autom. Reasoning 44(4), 401–424 (2010)
21. Piskac, R., Wies, T., Zufferey, D.: Automating Separation Logic Using SMT, chap. CAV 2013, Proceedings, pp. 773–789 (2013)

22. Piskac, R., Wies, T., Zufferey, D.: Automating Separation Logic with Trees and Data, pp. 711–728 (2014)
23. Reynolds, A., Deters, M., Kuncak, V., Barrett, C.W., Tinelli, C.: Counterexample guided quantifier instantiation for synthesis in CVC4. In: CAV. Springer (2015)
24. Reynolds, A., Iosif, R., King, T., Serban, C.: A decision procedure for separation logic in SMT. CoRR abs/1603.06844 (2016)
25. Reynolds, A., Tinelli, C., Goel, A., Krstić, S.: Finite model finding in SMT. In: Sharygina, N., Veith, H. (eds.) Computer Aided Verification, Lecture Notes in Computer Science, vol. 8044, pp. 640–655. Springer Berlin Heidelberg (2013)
26. Sighireanu, M., Cok, D.: Report on sl-comp 2014. Journal on Satisfiability, Boolean Modeling and Computation 1 (2014)
27. Toubhans, A., Chang, B.Y.E., Rival, X.: An Abstract Domain Combinator for Separately Conjoining Memory Abstractions, pp. 285–301. Springer (2014)
28. Voigt, M., Weidenbach, C.: Bernays-schönfinkel-ramsey with simple bounds is nexptime-complete. CoRR abs/1501.07209 (2015)
29. Yang, H.: Local Reasoning for Stateful Programs. Ph.D. thesis, University of Illinois at Urbana-Champaign (2001)

A Additional Material

A.1 Proof of Lemma 3

"$\Rightarrow$" Suppose that $\varphi^\forall$ has a model, i.e. $\mathcal{I}', h' \models_{\mathsf{SL}} \forall x_1 \ldots \forall x_n . \varphi(x_1, \ldots, x_n)$ for some interpretation $\mathcal{I}'$ and some heap $h' : U^{\mathcal{I}'} \rightharpoonup_{\text{fin}} (U^{\mathcal{I}'})^k$. We consider the case in which $\|\text{dom}(h') \setminus \mathcal{I}'(\text{Fvc}(\varphi^\forall))\| > |\varphi| + n$ — the other case $\|\text{dom}(h') \setminus \mathcal{I}'(\text{Fvc}(\varphi^\forall))\| \leq |\varphi| + n$ is an easy check left to the reader. Because $U^{\mathcal{I}'}$ is countable, we can choose a subset consisting of $|\varphi| + n$ locations from $\text{dom}(h') \setminus \mathcal{I}'(\text{Fvc}(\varphi^\forall))$, and let $L = \{\ell_1, \ldots, \ell_{|\varphi|+n}\}$ be this set. By Lemma 1, there exists a heap $h : U^{\mathcal{I}'} \rightharpoonup_{\text{fin}} (U^{\mathcal{I}'})^k$ such that:

- $h \sim^{\mathcal{I}'}_{|\varphi|+n, \text{Fvc}(\varphi^\forall), L} h'$,
- $\text{dom}(h) \setminus \mathcal{I}'(\text{Fvc}(\varphi^\forall)) \subseteq L$ and
- $h(\ell) = prun^{\nu}_{\mathcal{I}'(\text{Fvc}(\varphi^\forall)) \cup L}(h'(\ell))$, for all $\ell \in \text{dom}(h)$.

We define $\mathcal{I} = \mathcal{I}'[U \leftarrow \text{dom}(h) \cup \mathcal{I}'(\text{Fvc}(\varphi^\forall))]$ and prove that $\mathcal{I}, h \models_{\mathsf{SL}} \varphi^\forall$. Clearly the pair $\mathcal{I}, h$ satisfies the requirements from the statement of the lemma. We have to prove that $\mathcal{I}[x_1 \leftarrow u_1] \ldots [x_n \leftarrow u_n], h \models_{\mathsf{SL}} \varphi(x_1, \ldots, x_n)$, for all $u_1, \ldots, u_n \in U^{\mathcal{I}} = \text{dom}(h) \cup \mathcal{I}'(\text{Fvc}(\varphi^\forall))$. By Lemma 2, it is sufficient to prove that:

$$h \sim^{\mathcal{I}[x_1 \leftarrow u_1] \ldots [x_n \leftarrow u_n]}_{|\varphi|, \text{Fvc}(\varphi), \emptyset} h', \text{ for all } u_1, \ldots, u_n \in \text{dom}(h) \cup \mathcal{I}'(\text{Fvc}(\varphi^\forall))$$

Since $\mathcal{I}$ and $\mathcal{I}'$ agree on all variables from Vars, and $\mathcal{I}'[x_1 \leftarrow u_1] \ldots [x_n \leftarrow u_n], h' \models_{\mathsf{SL}} \varphi(x_1, \ldots, x_n)$, for all $u_1, \ldots, u_n \in U^{\mathcal{I}'}$, by the hypothesis, we obtain $\mathcal{I}, h \models_{\mathsf{SL}} \forall x_1 \ldots \forall x_n \, . \, \varphi(x)$. The proof is by induction on $n > 0$.

The base case $n = 1$**.** Let us prove the requirements of Definition 2:

1. $\mathcal{I}[x_1 \leftarrow u_1](\text{Fvc}(\varphi)) \cap \text{dom}(h) = \mathcal{I}[x_1 \leftarrow u_1](\text{Fvc}(\varphi)) \cap \text{dom}(h')$: observe first that $\text{Fvc}(\varphi) = \text{Fvc}(\varphi^\forall) \cup \{x_1\}$, thus we have:

$$\mathcal{I}[x_1 \leftarrow u_1](\text{Fvc}(\varphi)) = \mathcal{I}(\text{Fvc}(\varphi^\forall)) \cup \{u_1\} = \mathcal{I}'(\text{Fvc}(\varphi^\forall)) \cup \{u_1\} \quad .$$

If $u_1 \in \mathrm{dom}(h) \setminus \mathcal{I}'(\mathrm{Fvc}(\varphi^\forall))$, then $u_1 \in L$, because $\mathrm{dom}(h) \subseteq L$, and implicitly $u_1 \in \mathrm{dom}(h')$, since $L \subseteq \mathrm{dom}(h') \setminus \mathcal{I}'(\mathrm{Fvc}(\varphi^\forall))$. In this case, we have:

$$\begin{aligned}
\mathcal{I}[x_1 \leftarrow u_1](\mathrm{Fvc}(\varphi)) \cap \mathrm{dom}(h) &= (\mathcal{I}'(\mathrm{Fvc}(\varphi^\forall)) \cup \{u_1\}) \cap \mathrm{dom}(h) \\
&= (\mathcal{I}'(\mathrm{Fvc}(\varphi^\forall)) \cap \mathrm{dom}(h)) \cup \{u_1\} \\
&= (\mathcal{I}'(\mathrm{Fvc}(\varphi^\forall)) \cap \mathrm{dom}(h')) \cup \{u_1\} \text{ since } h \sim^{\mathcal{I}'}_{|\varphi|+n,\mathrm{Fvc}(\varphi^\forall),L} h' \\
&= \mathcal{I}[x_1 \leftarrow u_1](\mathrm{Fvc}(\varphi)) \cap \mathrm{dom}(h') \quad .
\end{aligned}$$

On the other hand, if $u_1 \in \mathcal{I}'(\mathrm{Fvc}(\varphi^\forall))$, we have $\mathcal{I}[x_1 \leftarrow u_1](\mathrm{Fvc}(\varphi)) = \mathcal{I}'(\mathrm{Fvc}(\varphi^\forall))$ and the result follows immediately.

2. $h(\ell') =_{\mathcal{I}[x_1 \leftarrow u_1](\mathrm{Fvc}(\varphi))} h'(\ell')$, for all $\ell' \in \mathcal{I}[x_1 \leftarrow u_1](\mathrm{Fvc}(\varphi)) \cap \mathrm{dom}(h)$: by the definition of h (Lemma 1), we have $h(\ell') = prun^v_{\mathcal{I}'(\mathrm{Fvc}(\varphi^\forall))\cup L}(h'(\ell'))$, for all $\ell' \in \mathrm{dom}(h)$. Hence we have $h(\ell') =_{\mathcal{I}'(\mathrm{Fvc}(\varphi^\forall))\cup L} h'(\ell')$ and, consequently $h(\ell') =_{\mathcal{I}'(\mathrm{Fvc}(\varphi^\forall))\cup\{u_1\}} h'(\ell')$, for all $\ell' \in \mathrm{dom}(h)$, since $u_1 \in \mathcal{I}'(\mathrm{Fvc}(\varphi^\forall)) \cup L$.
3. $\|\mathrm{dom}(h') \setminus \mathcal{I}[x_1 \leftarrow u_1](\mathrm{Fvc}(\varphi))\| = \|\mathrm{dom}(h') \setminus (\mathcal{I}(\mathrm{Fvc}(\varphi^\forall)) \cup \{u_1\})\| \geq \|\mathrm{dom}(h') \setminus \mathcal{I}(\mathrm{Fvc}(\varphi^\forall))\| - 1 > |\varphi|$, by the previous assumption. Since $h' \sim^{\mathcal{I}}_{|\varphi|+1,\mathrm{Fvc}(\varphi^\forall),L} h$, we get $\|\mathrm{dom}(h) \setminus \mathcal{I}(\mathrm{Fvc}(\varphi^\forall))\| \geq |\varphi| + 1$, thus $\|\mathrm{dom}(h) \setminus \mathcal{I}[x_1 \leftarrow u_1](\mathrm{Fvc}(\varphi^\forall))\| = \|\mathrm{dom}(h) \setminus (\mathcal{I}(\mathrm{Fvc}(\varphi^\forall)) \cup \{u_1\})\| \geq |\varphi|$.

The induction step $n > 1$**.** We prove the points of Definition 2, similar to the base case:

1. $\mathcal{I}[x_1 \leftarrow u_1]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi)) \cap \mathrm{dom}(h) = I[x_1 \leftarrow u_1]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi)) \cap \mathrm{dom}(h')$: we distinguish the case (i) $u_1 \in \mathrm{dom}(h) \setminus \mathcal{I}[x_2 \leftarrow u_2]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi))$ from (ii) $u_1 \in \mathcal{I}[x_2 \leftarrow u_2]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi))$. In the first case, we have:

$$\begin{aligned}
\mathcal{I}[x_1 \leftarrow u_1]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi)) \cap \mathrm{dom}(h) &= (\mathcal{I}[x_2 \leftarrow u_2]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi)) \cup \{u_1\}) \cap \mathrm{dom}(h) \\
&= (\mathcal{I}[x_2 \leftarrow u_2]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi)) \cap \mathrm{dom}(h)) \cap \{u_1\} \\
\text{by the induction hypothesis} \quad &= (\mathcal{I}[x_2 \leftarrow u_2]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi)) \cap \mathrm{dom}(h')) \cap \{u_1\} \\
&= \mathcal{I}[x_1 \leftarrow u_1]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi)) \cap \mathrm{dom}(h') \quad .
\end{aligned}$$

If $u_1 \in \mathcal{I}[x_2 \leftarrow u_2]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi))$, we have $\mathcal{I}[x_1 \leftarrow u_1]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi)) = \mathcal{I}[x_2 \leftarrow u_2]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi))$ and an application of the induction hypothesis concludes the proof.

2. By the construction of h, we have $h(\ell) =_{\mathcal{I}(\mathrm{Fvc}(\varphi^\forall))\cup L} h'(\ell)$, for all $\ell \in \mathrm{dom}(h)$, thus $h(\ell) =_{\mathcal{I}[x_1 \leftarrow u_1]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi))} h'(\ell)$, for all $\ell \in dom(h) \cap \mathcal{I}[x_1 \leftarrow u_1]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi))$.
3. Similar to the base case, we have:

$$\begin{aligned}
\|\mathrm{dom}(h') \setminus \mathcal{I}[x_1 \leftarrow u_1]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi))\| &= \|\mathrm{dom}(h') \setminus (\mathcal{I}(\mathrm{Fvc}(\varphi^\forall)) \cup \{u_1, \ldots, u_n\})\| \\
&\geq \|\mathrm{dom}(h') \setminus \mathcal{I}(\mathrm{Fvc}(\varphi^\forall))\| - n \\
&> |\varphi|, \text{ and} \\
\|\mathrm{dom}(h) \setminus \mathcal{I}[x_1 \leftarrow u_1]\ldots[x_n \leftarrow u_n](\mathrm{Fvc}(\varphi))\| &= \|\mathrm{dom}(h') \setminus (\mathcal{I}(\mathrm{Fvc}(\varphi^\forall)) \cup \{u_1, \ldots, u_n\})\| \\
&\geq |\varphi| \quad .
\end{aligned}$$

Since the direction "⇐" is trivial, this concludes our proof. □

A.2 Proof of Lemma 4

By reduction from the undecidability of the following tiling problem. Let $\mathcal{T} = \{T_0, \ldots, T_s\}$ be a set of tile types and $H, V \subseteq \mathcal{T} \times \mathcal{T}$ be two relations between tile types. Given $n, m > 1$, a tiling of the $n \times m$ square is a function $\tau : [1,n] \times [1,m] \to \mathcal{T}$ such that:

- $\tau(1,1) = T_0$,
- $(\tau(i,j), \tau(i+1,j)) \in H$, for all $i \in [1, n-1]$ and $j \in [1,m]$,
- $(\tau(i,j), \tau(i,j+1)) \in V$, for all $i \in [1,n]$ and $j \in [1, m-1]$.

The existence of $n, m > 1$ and of a tiling of the $n \times m$ square is a well-known undecidable problem. We encode this as the satisfiability of a formula in $\exists^*\forall^*\exists^*\mathsf{SL}(\mathsf{E})_{U,U^k}$. Let $\ell_0, \ell_1, t_0, \ldots, t_s$ be constant symbols of sort U.

The basic idea is to represent each cell of the $n \times m$ grid by a heap cell, defined by the atomic proposition $\mathsf{x} \mapsto (b, r, \mathsf{y}, \mathsf{z}, T)$, where:

- $b, r \in \{\ell_0, \ell_1\}$ act as binary flags indicating whether the cell belongs to the bottom n row ($b = \ell_1$), and the rightmost m column ($r = \ell_1$), respectively,
- y and z are the horizontal (right) and vertical (below) successors of x, and
- $T \in \{t_0, \ldots, t_s\}$ is the type of the tile covering the cell pointed to by x.

The finite relations $H, V \subseteq \mathcal{T} \times \mathcal{T}$ are defined by the finite disjunctions $h(x,y), v(x,y)$ of equalities involving $t_0, \ldots, t_s$. We write $\mathsf{x} \hookrightarrow (b, r, \mathsf{y}, \mathsf{z}, T)$ for $\mathsf{x} \mapsto (b, r, \mathsf{y}, \mathsf{z}, T) * \top$, $\mathsf{x} \hookrightarrow (\_, r, \mathsf{y}, \mathsf{z}, T)$ for $\bigvee_{b \in \{\ell_0, \ell_1\}} \mathsf{x} \hookrightarrow (b, r, \mathsf{y}, \mathsf{z}, T)$, $\mathsf{x} \hookrightarrow (b, \_, \mathsf{y}, \mathsf{z}, T)$ for $\bigvee_{r \in \{\ell_0, \ell_1\}} \mathsf{x} \hookrightarrow (b, r, \mathsf{y}, \mathsf{z}, T)$, $\mathsf{x} \hookrightarrow (b, r, \mathsf{y}, \mathsf{z}, \_)$ for $\bigvee_{T \in \{t_0, \ldots, t_s\}} \mathsf{x} \hookrightarrow (b, r, \mathsf{y}, \mathsf{z}, T)$, and $\mathsf{x} \hookrightarrow (b, r, \_, \_, T)$ for $\exists \mathsf{y} \exists \mathsf{z} \,.\, \mathsf{x} \hookrightarrow (b, r, \mathsf{y}, \mathsf{z}, T)$. We define the formula Φ as the conjunction of the formulae below, with the pairwise disequality constraint $\bigwedge_{c, c' \in \{\ell_0, \ell_1, t_0, \ldots, t_s\}} c \not\approx c'$:

$$\exists \mathsf{u} \forall \mathsf{y} \forall \mathsf{z} \,.\, \mathsf{u} \hookrightarrow (0, 0, \_, \_, t_0) \wedge \neg \mathsf{y} \hookrightarrow (\_, \_, \mathsf{u}, \_, \_) \wedge \neg \mathsf{z} \hookrightarrow (\_, \_, \_, \mathsf{u}, \_) \quad (1)$$

$$\forall \mathsf{x} \forall \mathsf{y} \forall \mathsf{z} \,.\, \bigwedge_{b \in \{\ell_0, \ell_1\}} [\mathsf{x} \hookrightarrow (b, 0, \mathsf{y}, \mathsf{z}, \_) \Rightarrow \mathsf{y} \hookrightarrow (b, \_, \_, \_, \_) \wedge \mathsf{x} \not\approx \mathsf{y}] \quad (2)$$

$$\forall \mathsf{x} \forall \mathsf{y} \forall \mathsf{z} \,.\, \mathsf{x} \hookrightarrow (\_, 1, \mathsf{y}, \mathsf{z}, \_) \Rightarrow \mathsf{x} \approx \mathsf{y} \quad (3)$$

$$\forall \mathsf{x} \forall \mathsf{y} \forall \mathsf{z} \,.\, \bigwedge_{r \in \{\ell_0, \ell_1\}} [\mathsf{x} \hookrightarrow (0, r, \mathsf{y}, \mathsf{z}, \_) \Rightarrow \mathsf{y} \hookrightarrow (\_, r, \_, \_, \_) \wedge \mathsf{x} \not\approx \mathsf{z}] \quad (4)$$

$$\forall \mathsf{x} \forall \mathsf{y} \forall \mathsf{z} \,.\, \mathsf{x} \hookrightarrow (1, \_, \mathsf{y}, \mathsf{z}, \_) \Rightarrow \mathsf{x} \approx \mathsf{z} \quad (5)$$

$$\forall \mathsf{x} \forall \mathsf{y} \forall \mathsf{y}' \forall \mathsf{y}'' \forall \mathsf{z} \forall \mathsf{z}' \forall \mathsf{z}'' \,.\, \mathsf{x} \hookrightarrow (\_, \_, \mathsf{y}, \mathsf{z}, \_) * \mathsf{y} \mapsto (\_, \_, \mathsf{y}', \mathsf{z}', \_) * \mathsf{z} \mapsto (\_, \_, \mathsf{y}'', \mathsf{z}'', \_) \Rightarrow \mathsf{z}' \approx \mathsf{y}'' \quad (6)$$

$$\forall \mathsf{x} \forall \mathsf{y} \forall \mathsf{z} \,.\, \bigwedge_{T, T' \in \{t_0, \ldots, t_s\}} \mathsf{x} \hookrightarrow (\_, \_, \mathsf{y}, \mathsf{z}, T) * \mathsf{y} \mapsto (\_, \_, \_, \_, T') \Rightarrow h(T, T') \quad (7)$$

$$\forall \mathsf{x} \forall \mathsf{y} \forall \mathsf{z} \,.\, \bigwedge_{T, T' \in \{t_0, \ldots, t_s\}} \mathsf{x} \hookrightarrow (\_, \_, \mathsf{y}, \mathsf{z}, T) * \mathsf{z} \mapsto (\_, \_, \_, \_, T') \Rightarrow v(T, T') \quad (8)$$

The intuition of these formulae is as follows: (1) is the initial constraint, asking that the top-left corner is labeled with T_0, (2) requires that each cell not on the rightmost column has a distinct left successor, (3) is the dual constraint, asking that the left successor of each cell on the rightmost column is the cell itself, (4) and (5) are the similar constraints for the bottom successors, (6) is the grid constraint, and (7), (8) are the horizontal and vertical constraints on the types of tiles. Observe that the quantifier prefix of Φ belongs to the language of the regular expression $\exists^*\forall^*\exists^*$.

Observe that the formulae (1 - 8) use $k = 5$ record fields. We can reduce the value of k to 2, by using the following encoding of the atomic propositions in Φ:

$$\mathsf{x} \hookrightarrow (b, r, \mathsf{y}, \mathsf{z}, T) \equiv \exists \mathsf{x}_0 \exists \mathsf{x}_1 \exists \mathsf{x}_2 \exists \mathsf{x}_3 \ . \ \mathsf{x} \hookrightarrow (\mathsf{x}_0, \mathsf{x}_1) * \mathsf{x}_1 \mapsto (\mathsf{y}, \mathsf{z}) * \mathsf{x}_0 \mapsto (\mathsf{x}_2, b) * \mathsf{x}_2 \mapsto (\mathsf{x}_3, r) * \mathsf{x}_3 \mapsto (T, \mathsf{nil})$$

It is easy to show now that Φ has a model iff there exists $n, m > 1$ and a tiling of the $n \times m$ grid, which proves the undecidability of the satisfiability problem for $\exists^* \forall^* \exists^* \mathsf{SL}(\mathsf{E})_{U,U^k}$, when $k \geq 2$. □

A.3 Proof of Theorem 2

Fact 3 *There exists an interpretation $\mathcal{I}$ and a heap h such that $\|U^{\mathcal{I}}\| = \aleph_0$ and $\mathcal{I}, h \models_{\mathsf{SL}} \forall y_1 \ldots \forall y_n \ . \ \varphi(\mathbf{c}, \mathbf{y})$ iff there exists an interpretation $\mathcal{I}'$, not constraining the cardinality of $U^{\mathcal{I}'}$, and a heap h' such that:*

$$\mathcal{I}', h' \models_{\mathsf{SL}} \mathsf{external} \wedge \forall y_1 \ldots \forall y_n \bigwedge_{\langle t_1, \ldots, t_n \rangle \in \{0,1,2\}^n} \underbrace{\bigwedge_{i=1}^{n} \Big(\psi_{t_i}(y_i) \Rightarrow \varphi(\mathbf{c}, \mathbf{y}) \Big)}_{\Psi_{\langle t_1, \ldots, t_n \rangle}}$$

Proof. "⇒" This direction is immediate, because for each location $\ell \in U^{\mathcal{I}}$, we have $\mathcal{I}[y \leftarrow \ell], h \models_{\mathsf{SL}} \psi_0(y) \vee \psi_1(y) \vee \psi_2(y)$ and $\mathcal{I}[y_1 \leftarrow \ell_1] \ldots [y_n \leftarrow \ell_n], h \models_{\mathsf{SL}} \varphi(\mathbf{c}, \mathbf{y})$, for all $\langle \ell_1, \ldots, \ell_n \rangle \in (U^{\mathcal{I}})^n$. "⇐" Because the cardinality of $U^{\mathcal{I}'}$ is unconstrained, by Lemma 3, there exists an interpretation $\mathcal{I}''$ and a heap h such that $\mathcal{I}'' = \mathcal{I}'[U \leftarrow \mathrm{dom}(h) \cup \mathcal{I}'(\mathrm{Fvc}(\Psi))]$ and $\mathcal{I}'', h \models_{\mathsf{SL}} \mathsf{extern} \wedge \forall y_1 \ldots \forall y_n \bigwedge_{\langle t_1, \ldots, t_n \rangle \in \{0,1,2\}^n} \Psi_{\langle t_1, \ldots, t_n \rangle}$. We obtain that $\mathcal{I}''[y_1 \leftarrow \ell_1] \ldots [y_n \leftarrow \ell_n] \models_{\mathsf{SL}} \varphi(\mathbf{c}, \mathbf{y})$ for each tuple $\langle \ell_1, \ldots, \ell_n \rangle \in (U^{\mathcal{I}''})^n$ where either (i) $\ell_i \in \mathrm{dom}(h)$, (ii) $\ell_i = \mathcal{I}''(c_j)$ for some $j = 1, \ldots, m$, or (iii) $\ell_i = \mathcal{I}''(d_j)$ for some $j = 1, \ldots, n$ and neither of the previous hold. Because it is not important which location is used for the interpretation of d_j, $j = 1, \ldots, n$, we have $\mathcal{I}, h \models_{\mathsf{SL}} \forall y_1 \ldots \forall y_n \ . \ \varphi(\mathbf{c}, \mathbf{y})$ for every extension $\mathcal{I}$ of $\mathcal{I}''$ such that $\|U^{\mathcal{I}}\| = \aleph_0$. □

A.4 Proof of Theorem 3

Fact 4 *For each interpretation $\mathcal{I}$ mapping x and y into $\mathbb{N}$, $\mathcal{I} \models x = y^2$ iff $\mathcal{I}$ can be extended to an interpretation of P as a finite set of consecutive perfect squares such that $\mathcal{I} \models \theta_{x=y^2}$.*

Proof. "⇒" If $\mathcal{I} \models x = y^2$ then we have $(y^{\mathcal{I}} + 1)^2 = x^{\mathcal{I}} + 2y^{\mathcal{I}} + 1$. Let $P^{\mathcal{I}}$ be the set $\{0, 1, \ldots, (y^{\mathcal{I}} + 1)^2\}$. Clearly $x^{\mathcal{I}}, x^{\mathcal{I}} + 2y^{\mathcal{I}} + 1 \in P^{\mathcal{I}}$ and, since they are consecutive perfect squares, every number in between $x^{\mathcal{I}}$ and $x^{\mathcal{I}} + 2y^{\mathcal{I}} + 1$ does not belong to $P^{\mathcal{I}}$. Thus $\mathcal{I} \models \theta_{x=y^2}$. "⇐" If $\mathcal{I} \models \theta_{x=y^2}$ and $P^{\mathcal{I}}$ is a set of consecutive perfect squares, it follows that $x^{\mathcal{I}}$ and $x^{\mathcal{I}} + 2y^{\mathcal{I}} + 1$ are consecutive perfect squares, i.e. $x^{\mathcal{I}} = n^2$ and $x^{\mathcal{I}} + 2y^{\mathcal{I}} + 1 = (n+1)^2$ for some $n \in \mathbb{N}$. Then $y^{\mathcal{I}} = n$, thus $\mathcal{I} \models x = y^2$. □