

Deciding Boolean Separation Logic via Small Models *

(Technical Report)

Tomáš Dacík¹, Adam Rogalewicz¹, Tomáš Vojnar¹, and Florian Zuleger²

¹ Faculty of Information Technology, Brno University of Technology, Czech Republic

² Faculty of Informatics, Vienna University of Technology, Austria

Abstract. We present a novel decision procedure for a fragment of separation logic (SL) with arbitrary nesting of separating conjunctions with boolean conjunctions, disjunctions, and guarded negations together with a support for the most common variants of linked lists. Our method is based on a model-based translation to SMT for which we introduce several optimisations—the most important of them is based on bounding the size of predicate instantiations within models of larger formulae, which leads to a much more efficient translation of SL formulae to SMT. Through a series of experiments, we show that, on the frequently used symbolic heap fragment, our decision procedure is competitive with other existing approaches, and it can outperform them outside the symbolic heap fragment. Moreover, our decision procedure can also handle some formulae for which no decision procedure has been implemented so far.

1 Introduction

In the last decade, separation logic (SL) [14, 29] has become one of the most popular formalisms for reasoning about programs working with dynamically-allocated memory, including approaches based on deductive verification [31], abstract interpretation [33], symbolic execution [30], or bi-abductive analysis [6, 11, 17]. The key ingredients of SL used in these approaches include the separating conjunction $*$, which allows modular reasoning by stating that the program heap can be decomposed into disjoint parts satisfying operands of the separating conjunction, along with inductive predicates describing shapes of data structures, such as lists, trees, or their various combinations.

The high expressive power of SL comes with the price of high complexity and even undecidability when several of its features are combined together. The existing decision procedures are usually limited to the so-called *symbolic heap fragment* that disallows any boolean structure of spatial assertions.

In this paper, we present a novel decision procedure for a fragment of SL that we call *boolean separation logic* (BSL). The fragment allows arbitrary nesting of separating conjunctions and boolean connectives of conjunction, disjunction, and a limited form of negation of the form $\varphi \wedge \neg\psi$ called *guarded negation*. To the best of our knowledge, no existing, practically applicable decision procedure supports a fragment with such a rich boolean structure and at least basic inductive predicates. The decision procedure for SL in CVC5 [28] supports arbitrary nesting of boolean connectives (including even unguarded negation, which is considered very expensive in the context of SL) but no inductive predicates. A support for conjunctions and disjunctions under separating

* The work was supported by the Czech Science Foundation project GA23-06506S. Basic research funding of the Czech team was provided by the FIT BUT internal project FIT-S-23-8151 and the ERC.CZ project LL1908. Tomáš Dacík was supported by the Brno Ph.D. Talent Scholarship funded by the Brno City Municipality.

conjunctions is available in the backend solver of the GRASSHOPPER verifier [26, 27] though not described in the papers. In our experimental evaluation, we outperform both of these approaches on some benchmarks (and can decide some formulae beyond the capabilities of both of them). We further show that adding guarded negations to BSL makes its satisfiability problem PSPACE-hard.

To motivate the usefulness of the fragment we consider, we now give several examples when SL formulae with a rich boolean structure are useful. First, in symbolic execution of heap manipulating programs, one usually needs to consider functions that involve some non-determinism—typically, at least the `malloc()` statement has the non-deterministic contract $\{\text{emp}\} x = \text{malloc}() \{x \mapsto f \vee (x = \text{nil} \wedge \text{emp})\}$ (where f is a fresh variable) stating that when the statement is started in the empty heap, once it finishes, x is either allocated, or the allocation had failed and the heap is empty. Such contracts typically need a dedicated (and usually incomplete) treatment when no support of disjunctions is available.³ Further, the guarded negation semantically represents the set of counterexamples of the entailment $\varphi \models \psi$, and hence allows one to reduce entailment queries to UNSAT checking. Guarded negation can also be used when one needs to obtain several models of a formula φ by joining formulae representing the already obtained models to φ using guarded negations. One can also use the guarded negation to express interesting properties such as the fact that given a list $\text{sls}(x, y)$ and a pointer $y \mapsto z$, the pointer does not point back somewhere into the list closing a lasso. This can be expressed through the formula $(\text{sls}(x, y) \wedge \neg(\text{sls}(x, z) * \text{sls}(z, y))) * y \mapsto z$. Finally, boolean connectives can be introduced by translating quantitative separation logic into the classical SL [2].

In this work, we consider BSL with three fixed, built-in inductive predicates representing the most-common variants of lists: singly-linked (SLL), doubly-linked (DLL), and nested singly-linked (NLL). Our results can be easily extended for their variations such as nested doubly-linked lists of singly-linked lists and the like, but for the price of manually defining their semantics in the SMT encoding. We do, however, believe that our approach of bounding the sizes of models and instantiations of the individual predicates can be lifted to more complex inductive definitions and can serve as a starting point for allowing integration of SL with inductive definitions into SMT.

Contributions. Our approach to deciding BSL formulae is inspired by previous works on translation of SL to SMT. The early works [26] and [27] translate SL to intermediate theories first. Our approach is closer to the more recent approach of [15], which builds on small-model properties and axiomatizes reachability through pointer links directly. We extend the SL fragment considered in [15] by going beyond the so-called unique footprint property (under which it is much easier to obtain an efficient translation). Further, we define a more precise way to obtain global bounds on models of entire formulae, and, most importantly, we modify the translation of inductive predicates in a way that allows us to encode them succinctly by computing local bounds on their instantiations. According to our experiments, this makes the decision procedure efficient and competitive with the state-of-the-art approaches on the symbolic heap fragment (despite the increased decisive power).

³ Note that, while the post-condition with a single disjunction might seem simple, the formulae typically start growing in the further symbolic execution.

Related work. In [3], a proof system for deciding entailments of symbolic heaps with lists was proposed. This problem was later shown to be solvable in polynomial time in [8] via graph homomorphism checking. A superposition-based calculus for the fragment was presented in [22], and a model-based approach enhancing SMT solvers was proposed in [23]. In [23], a combination of SL with SMT theories is considered but still limited to the symbolic heap fragment. A more expressive boolean structure and integration with SMT theories was developed in [26] for lists and extended for trees in [27] but still without a support for guarded negations.

Other decision procedures are focusing on more general, *user-defined* inductive predicates (usually of some restricted form). They are based, e.g., on *cyclic proof systems* (CYCLIST [5], S2S [18, 19]); lemma synthesis (SONGBIRD [32]); or automata—tree automata are used in the tools SLIDE [12] and SPEN [10], and a specialised type of automata, called *heap automata*, is used in HARRSH [16]. These procedures do, however, not support nested use of boolean connectives and separating conjunctions.

There also exist works on deciding much more expressive fragments of SL such as [9, 13, 20, 25] but they do not lead to practically implementable decision procedures.

2 Preliminaries

Partial functions. We write $f : X \rightarrow Y$ to denote a *partial function* from X to Y . For a partial function f , $\text{dom}(f)$ and $\text{img}(f)$ denote its domain and image, respectively; $|f| = |\text{dom}(f)|$ denotes its size, and $f(x) = \perp$ denotes that f is undefined for x . A restriction $f|_A$ of f to $A \subseteq X$ is defined as $f(x)$ for $x \in A$ and undefined otherwise. To represent a finite partial function f , we often use the set notation $f = \{x_1 \mapsto y_1, \dots, x_n \mapsto y_n\}$ meaning that f maps each x_i to y_i , and is undefined for other values. We call partial functions f_1 and f_2 *disjoint* if $\text{dom}(f_1) \cap \text{dom}(f_2) = \emptyset$ and define their *disjoint union* $f_1 \uplus f_2$ as $f_1 \cup f_2$, which is otherwise undefined.

Graphs and paths. Let $G = (V, \rightarrow_1, \dots, \rightarrow_m)$ be a directed graph with vertices V and edges $\rightarrow = \rightarrow_1 \cup \dots \cup \rightarrow_m$. For $1 \leq f \leq m$, a sequence $\sigma = \langle v_0, v_1, \dots, v_n \rangle \in V^+$ is a path from v_0 to v_n via \rightarrow_f in G , denoted as $\sigma : v_0 \rightsquigarrow_f v_n$, if all elements of σ are distinct, and for all $0 \leq i < n$, it holds that $v_i \rightarrow_f v_{i+1}$. By the definition, paths cannot be cyclic. The *domain* of the path σ is the set $\text{dom}(\sigma) = \{v_0, v_1, \dots, v_{n-1}\}$, and the length of the path is defined as $|\sigma| = |\text{dom}(\sigma)| = n$.

Formulae. For a first-order formula φ , we denote by $\varphi[t/x]$ the formula obtained by simultaneously replacing all free occurrences of the variable x in φ with the term t . For a first-order model \mathcal{M} and a term t , we write $t^{\mathcal{M}}$ to denote the evaluation of t in \mathcal{M} defined as usual.

3 Separation Logic

Syntax. Let Vars be a countably infinite set of *sorted variables*. We denote by x^S a variable x of a sort $S \in \text{Sort} = \{\mathbb{S}, \mathbb{D}, \mathbb{N}\}$ representing a location in an SLL, DLL, or NLL, respectively. We omit the sorts when they are not relevant or clear from the context. We further assume that there exists a distinguished, unsorted variable nil . We write $\text{vars}(\varphi)$ to denote the set of all variables in φ plus nil (even when it does not appear in φ). Analogically, $\text{vars}_S(\varphi)$ stands for all variables of the sort S plus nil .

$(s, h) \models x \bowtie y$	iff $s(x) \bowtie s(y)$ and $\text{dom}(h) = \emptyset$ for $\bowtie \in \{=, \neq\}$
$(s, h) \models x \mapsto \langle f_i : f_i \rangle_{i \in I}$	iff $h = \{s(x) \mapsto \langle f_i : s(f_i) \rangle_{i \in I}\}$
$(s, h) \models \psi_1 \bowtie \psi_2$	iff $(s, h) \models \psi_1 \bowtie (s, h) \models \psi_2$ for $\bowtie \in \{\wedge, \vee, \neg\}$
$(s, h) \models \psi_1 * \psi_2$	iff $\exists h_1, h_2. h = h_1 \uplus h_2 \neq \perp$ and $(s, h_i) \models \psi_i$ for $i = 1, 2$
$(s, h) \models \exists x. \psi$	iff there exists ℓ such that $(s[x \mapsto \ell], h) \models \psi$
$(s, h) \models \text{sls}(x, y)$	iff $(s, h) \models x = y$, or $s(x) \neq s(y)$ and $(s, h) \models \exists n. x \mapsto n * \text{sls}(n, y)$
$(s, h) \models \text{dls}(x, y, x', y')$	iff $(s, h) \models x = y * x' = y'$, or $s(x) \neq s(y)$, $s(x') \neq s(y')$, and $(s, h) \models \exists n. x \mapsto \langle n : n, p : y' \rangle * \text{dls}(n, y, x', x)$
$(s, h) \models \text{nls}(x, y, z)$	iff $(s, h) \models x = y$, or $s(x) \neq s(y)$ and $(s, h) \models \exists n, t. x \mapsto \langle n : n, t : t \rangle * \text{sls}(n, z) * \text{nls}(t, y, z)$

Fig. 1: The semantics of the separation logic. The existential quantifier is used for the definition of the semantics of inductive predicates and it is not a part of our fragment.

The syntax of our fragment is given by the following grammar:

$p ::= x^{\mathbb{S}} \mapsto \langle n : n \rangle \mid x^{\mathbb{D}} \mapsto \langle n : n, p : p \rangle \mid x^{\mathbb{N}} \mapsto \langle n : n, t : t \rangle$	(points-to predicates)
$\pi ::= \text{sls}(x^{\mathbb{S}}, y^{\mathbb{S}}) \mid \text{dls}(x^{\mathbb{D}}, y^{\mathbb{D}}, x_b^{\mathbb{D}}, y_b^{\mathbb{D}}) \mid \text{nls}(x^{\mathbb{N}}, y^{\mathbb{N}}, z^{\mathbb{S}})$	(inductive predicates)
$\varphi_A ::= x = y \mid x \neq y \mid p \mid \pi$	(atomic formulae)
$\varphi ::= \varphi_A \mid \varphi * \varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \wedge \neg \varphi$	(formulae)

The *points-to* predicate $x \mapsto \langle f_1 : f_1, \dots, f_n : f_n \rangle$ denotes that x is a structure whose fields f_i point to values f_i . We often write $x \mapsto n$ instead of $x \mapsto \langle n : n \rangle$ and $x \mapsto _$ if the right-hand side is not relevant. We call x the *root* of the points-to predicate. If π is an inductive predicate $\text{sls}(x, y)$, $\text{dls}(x, y, x', y')$, or $\text{nls}(x, y, z)$, we again call x the root of π , y is the *sink* of π , and we write $\pi(x, y)$ to denote the root and the sink. We define the *sort* of the predicate π , denoted as S_π , as the sort of its root. Then, there is a one-to-one correspondence of predicates and sorts, which we often implicitly use.

Memory model. Let Loc be a countably infinite set of memory locations, and let $\text{Field} = \{n, p, t\}$ be the set of fields. A *stack* is a finite partial function $s : \text{Vars} \rightarrow \text{Loc}$. A *heap* is a finite partial function $h : \text{Loc} \rightarrow (\text{Field} \rightarrow \text{Loc})$. For succinctness, we write $h(\ell, f)$ instead of $h(\ell)(f)$. To represent heap elements in a readable way, we write functions $\text{Field} \rightarrow \text{Loc}$ as vectors with labels, i.e., $h(\ell) = \langle f : h(\ell, f) \mid f \in \text{Field} \wedge h(\ell, f) \neq \perp \rangle$ and we write $\text{img}(h)$ for $\{\ell \in \text{Loc} \mid \exists \ell', f. h(\ell', f) = \ell\}$. Moreover, we use $h(\ell) = n$ when $h(\ell) = \langle n : n \rangle$. A *stack-heap model* is a pair (s, h) where s is stack and h is a heap such that $s(\text{nil}) \neq \perp$ and $h(s(\text{nil})) = \perp$. We define the set of locations of the model (s, h) as $\text{locs}(s, h) = \text{img}(s) \cup \text{dom}(h) \cup \text{img}(h)$.

Semantics. The semantics of our SL over stack-heap models is given in Fig. 1. For pure formulae, we use the so-called *precise semantics*, which additionally requires that the heap must be empty⁴. The semantics of pointer assertions, boolean connectives, and

⁴ This is a common approach to avoid the atom true to be expressed as $\text{nil} = \text{nil}$. In our fragment, we forbid true in order not to introduce “unbounded” negations as $\neg \varphi \triangleq \text{true} \wedge \neg \varphi$. Due to this change, symbolic heaps are formulae of form $* \psi_i$ where each ψ_i is an atom.

separating conjunctions is as usual. The intuition behind the semantics of the inductive predicates is as follows. An SLL segment $\text{sls}(x, y)$ is either empty or represents an acyclic sequence of allocated locations starting from x and leading via the n field to y , which is not allocated. A DLL segment $\text{dls}(x, y, x', y')$ is either empty with $x = y$ and $x' = y'$, or it represents an acyclic sequence that is doubly-linked via the n and p fields and leads from the first allocated location x of the segment to its last allocated location x' (x and x' may coincide) with y/y' being the n/p -successors of x'/x , respectively. Both y and y' are not allocated. An NLL segment $\text{nls}(x, y, z)$ is a (possibly empty) acyclic sequence of locations starting from x and leading to y via the t (top) field in which successor of each locations starts a disjoint inner SLL to z via n .

Stack-heap graphs. We frequently identify stack-heap models with their graph representation. A stack-heap model (s, h) defines a graph $G[(s, h)] = (V, (\rightarrow_f)_{f \in \text{Field}})$ where $V = \text{locs}(s, h)$ and $u \rightarrow_f v$ iff $h(u, f) = v$. We frequently use the fact that if there exists a path $\sigma : x \rightsquigarrow_f y$ in a stack-heap graph, then it is uniquely determined because f -edges are given by a partial function.

4 Small-Model Property

Small-model properties, which state that each satisfiable formula has a model of bounded size, are frequently used for various fragments of SL to prove their decidability [7] or to design decision procedures [15, 25, 28]. The latter is also the case of our translation-based decision procedure which will heavily rely on enumeration over all locations, and, for its efficiency, it is therefore necessary to obtain location bounds that are as small as possible.

The way we obtain our small-model property is inspired by the approach of [15] and by insights from the so-called *strong-separation logic* [25]. The main idea is to define a satisfiability-preserving reduction $\downarrow^s h$ which takes a heap h (referenced from a stack s), decomposes it into basic sub-heaps (which we call *chunks*), and reduces it per the sub-heaps in such a way that its size can be easily bounded by a linear expression. To define the reduction, we first need to introduce some auxiliary notions related to stack-heap models.

We say that a model (s, h) is *positive* if there exists φ with $(s, h) \models \varphi$. A positive model (s, h) is *atomic* if it is non-empty, and for all positive models (s, h_1) and (s, h_2) , $h = h_1 \uplus h_2$ implies that $h_1 = \emptyset$ or $h_2 = \emptyset$. In other words, atomic models cannot be decomposed into two non-empty positive models. Several examples of atomic models are shown in Fig. 2. Observe that the models of dls (Figure 2b) and nls (Figure 2c) are indeed atomic as any of their decomposition, in particular the split at the location u , does not give two positive models.

A sub-heap $c \subseteq h$ is a *chunk* of a model (s, h) if c is a maximal sub-heap of h such that (s, c) is an atomic positive model. Notice that the way the definition of chunks is constructed excludes the possibility of using as a chunk a sub-heap of a heap that itself forms an atomic model. The reason is that otherwise the remaining part of the larger atomic model could not be described by the available predicates. For example, in nested lists as shown in Fig. 2c, one cannot take as a chunk a part of some inner list (e.g., the pointer $u \mapsto z$) as the heap shown in the figure itself forms an atomic model. Indeed, if $u \mapsto z$ was removed, one would need a more general version of the NLL predicate to cover the remaining heap by atomic models.

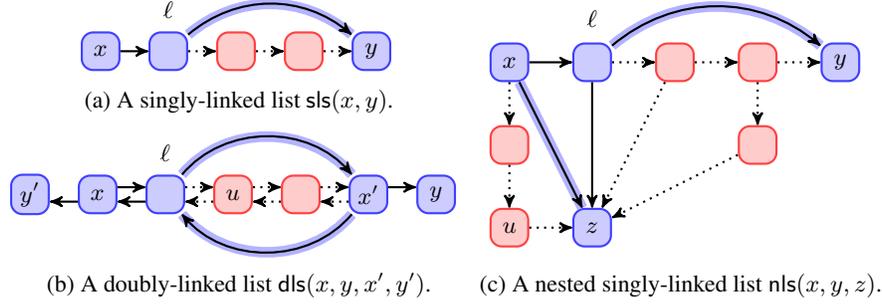


Fig. 2: An illustration of reductions of atomic models of inductive predicates. Removed heap locations are red, removed edges are dotted, and added edges are highlighted.

Lemma 1 (Chunk decomposition). *A positive model (s, h) can be uniquely decomposed into the set of its chunks, denoted $\text{chunks}(s, h)$, i.e., $h = \biguplus \text{chunks}(s, h)$.*

Minimal atomic models of inductive predicates. The key reason why the small-model property that we are going to state holds is that our fragment of SL cannot distinguish atomic models of the considered predicates beyond certain small sizes—namely, two for sls and nls , and three for dls . For further use, we will now state predicates describing exactly the sets of the indistinguishable lists of the different kinds.

We start with SLLs and use a disequality to exclude empty lists: $\text{sls}_{\geq 1}(x, y) \triangleq \text{sls}(x, y) * x \neq y$, and a guarded negation to exclude lists of length one consisting of a single pointer only: $\text{sls}_{\geq 2}(x, y) \triangleq \text{sls}_{\geq 1}(x, y) \wedge \neg(x \mapsto y)$. A similar predicate can be defined for NLLs too: $\text{nls}_{\geq 2}(x, y, z) \triangleq (\text{nls}(x, y, z) * x \neq y) \wedge \neg(x \mapsto \langle n : z, t : y \rangle)$.

For DLLs, we define $\text{dls}_{\geq 2}(x, y, x', y') \triangleq \text{dls}(x, y, x', y') * x \neq y * x \neq x'$ to exclude models that are either empty or consist of a single pointer; and $\text{dls}_{\geq 3}(x, y, x', y') \triangleq \text{dls}_{\geq 2}(x, y, x', y') \wedge \neg(x \mapsto \langle n : x', p : y' \rangle * x' \mapsto \langle n : y, p : x \rangle)$ to also exclude models consisting of exactly two pointers.

It holds that atomic models, and consequently also chunks, are precisely either models of single pointers or of the above predicates.

Lemma 2. *For atomic model (s, h) , exactly one of the following conditions holds.*

1. $(s, h) \models x \mapsto _$ for some x . (pointer-atom)
2. $(s, h) \models \text{sls}_{\geq 2}(x, y)$ for some x and y . (sls-atom)
3. $(s, h) \models \text{dls}_{\geq 3}(x, y, x', y')$ for some $x, y, x',$ and y' . (dls-atom)
4. $(s, h) \models \text{nls}_{\geq 2}(x, y, z)$ for some $x, y,$ and z . (nls-atom)

We can now define the reduction in the way we have already sketched.

Definition 1. *The heap of a positive model (s, h) reduces to $\downarrow^s h = \biguplus_{c \in \text{chunks}(s, h)} \downarrow^s c$ where the reduction of a chunk c with a root x as follows:*

- $\downarrow^s c = c$ if $(s, c) \models x \mapsto _$.
- $\downarrow^s c = \{s(x) \mapsto \ell, \ell \mapsto s(y)\}$ where $\ell = c(s(x), n)$ if $(s, c) \models \text{sls}_{\geq 2}(x, y)$ for some y .
- $\downarrow^s c = \{s(x) \mapsto \langle n : \ell, p : s(y') \rangle, \ell \mapsto \langle n : s(x'), p : s(x) \rangle, s(x') \mapsto \langle n : s(y), p : \ell \rangle\}$ where $\ell = c(s(x), n)$ if $(s, c) \models \text{dls}_{\geq 3}(x, y, x', y')$ for some x', y' and y .
- $\downarrow^s c = \{s(x) \mapsto \langle t : \ell, n : s(z) \rangle, \ell \mapsto \langle t : s(y), n : s(z) \rangle\}$ where $\ell = c(s(x), t)$ if $(s, c) \models \text{nls}_{\geq 2}(x, y, z)$ for some y and z .

We lift the reduction to stack-heap models as $\downarrow^X (s, h) = (s', \downarrow^{s'} h)$ where $s' = s|_X$ for some set of variables X and show that it preserves satisfiability when $X = \text{vars}(\varphi)$.

Theorem 1. *For a positive model (s, h) , it holds that $(s, h) \models \varphi$ iff $\downarrow^{\text{vars}(\varphi)} (s, h) \models \varphi$.*

The final step to show our small-model property is to find an upper bound on the size of the reduced models. We define the size of a variable x^S , $\|x^S\|$, which represents its contribution to the location bound, and is defined as 2 if $S \in \{\mathbb{S}, \mathbb{N}\}$ and 1.5 if $S = \mathbb{D}$ (this corresponds to the size of a reduced chunk of sort S divided by the number of variables which are allocated in it). We further define $\|\text{nil}\| = 0$. The location bound of φ is then given as $\text{bound}(\varphi) = 1 + \lfloor \sum_{x \in \text{vars}(\varphi)} \|x\| \rfloor$ (the additional location is for nil). Analogically, the location bound for a sort S is $\text{bound}_S(\varphi) = \lfloor \sum_{x \in \text{vars}_S(\varphi)} \|x\| \rfloor$.

Theorem 2 (Small-model property). *If a formula φ is satisfiable, then there exists a model $(s, h) \models \varphi$ such that $|\text{locs}(s, h)| \leq \text{bound}(\varphi)$.*

We conjecture that the bound can be further improved, e.g., by showing that each model can be transformed to an equivalent one (indistinguishable by BSL formulae) such that the number of its chunks is bounded by the number of roots of spatial predicates in φ . We demonstrate this on the formula $\text{sls}(x, y) * y \mapsto z$ and its model in which y points back into the middle of the list segment (thus splitting it into two chunks). Clearly, this model can be transformed by redirecting z outside of the list domain.

5 Translation-Based Decision Procedure

In this section, we present our translation of SL to SMT. We first present an SMT encoding of our memory model and a translation of basic predicates and boolean connectives. Then we discuss methods for efficient translation of separating conjunctions and inductive predicates with the focus on avoiding quantifiers by replacing them by small enumerations of their instantiations.

We fix an input formula φ and let $n_S = \text{bound}_S(\varphi)$ for each sort $S \in \text{Sort}$.

5.1 Encoding the Memory Model in SMT

To encode the heap, we use a classical approach which encodes its mapping and domain separately [15, 26, 28]. Namely, we use arrays to encode mappings and sets to encode domains. We also use the theory of datatypes to represent a finite sort of locations by a datatype $L \triangleq \text{loc}^{\text{nil}} \mid \text{loc}_1^{\mathbb{S}} \mid \dots \mid \text{loc}_{n_{\mathbb{S}}}^{\mathbb{S}} \mid \text{loc}_1^{\mathbb{D}} \mid \dots \mid \text{loc}_{n_{\mathbb{D}}}^{\mathbb{D}} \mid \text{loc}_1^{\mathbb{N}} \mid \dots \mid \text{loc}_{n_{\mathbb{N}}}^{\mathbb{N}}$.

Now, we define the signature of the translation's language over the sort L . For each $x \in \text{vars}(\varphi)$, we introduce a constant x of the same name—its interpretation represents the stack image $s(x)$. To represent the heap, we introduce a set symbol D representing the domain and an array symbol h_f for each field $f \in \text{Field}$ which represents the mapping of the partial function $\lambda \ell. h(\ell, f)$. To distinguish sorts of locations, we further introduce a set symbol D_S for each sort $S \in \text{Sort}$. We define meaning of these symbols by showing how a stack-heap model can be reconstructed from a first-order model.

Definition 2 (Inverse translation). *Let \mathcal{M} be a first-order model. We define its inverse translation $\top_{\varphi}^{-1}(\mathcal{M}) = (s, h)$ where $s(x) = x^{\mathcal{M}}$ if $x \in \text{vars}(\varphi)$ and*

$$h(\ell) = \begin{cases} \langle n: h_n[\ell]^{\mathcal{M}} \rangle & \text{if } \ell \in (D \cap D_{\mathbb{S}})^{\mathcal{M}} \\ \langle n: h_n[\ell]^{\mathcal{M}}, p: h_p[\ell]^{\mathcal{M}} \rangle & \text{if } \ell \in (D \cap D_{\mathbb{D}})^{\mathcal{M}} \\ \langle n: h_n[\ell]^{\mathcal{M}}, t: h_t[\ell]^{\mathcal{M}} \rangle & \text{if } \ell \in (D \cap D_{\mathbb{N}})^{\mathcal{M}}. \end{cases}$$

To ensure consistency of the translation with the memory model used, we define the following axioms that a result of translation needs to satisfy:

$$\mathcal{A}_\varphi \triangleq \text{nil} = \text{loc}^{\text{nil}} \wedge \text{nil} \notin D \wedge \bigwedge_{S \in \text{Sort}} (D_S = \{\text{loc}^{\text{nil}}, \text{loc}_1^S, \dots, \text{loc}_{n_S}^S\} \wedge \bigwedge_{x \in \text{vars}_S(\varphi)} x \in D_S).$$

The axioms ensure that nil is never allocated, that each variable is interpreted as a location of the corresponding sort and they fix the interpretation of the sets $D_S, D_{\mathbb{D}}, D_{\mathbb{N}}$, which we will later use in the translation to assign sorts to locations.

5.2 Translation of SL to SMT

We define the translation as a function $\mathbb{T}(\varphi) = \mathcal{A}_\varphi \wedge \mathbb{T}(\varphi, D)$ where \mathcal{A}_φ are the above defined axioms and $\mathbb{T}(\varphi, D)$ is a recursive translation function of the formula φ with the domain symbol D . The translation $\mathbb{T}(\cdot)$ together with the inverse translation of models $\mathbb{T}_\varphi^{-1}(\cdot)$ are linked by the following correctness theorem.

Theorem 3 (Translation correctness). *An SL formula φ is satisfiable iff its translation $\mathbb{T}(\varphi)$ is satisfiable. Moreover, if $\mathcal{M} \models \mathbb{T}(\varphi)$, then $\mathbb{T}_\varphi^{-1}(\mathcal{M}) \models \varphi$.*

The translation of non-inductive predicates and boolean connectives is defined as:

$$\begin{aligned} \mathbb{T}(x \bowtie y, F) &\triangleq x \bowtie y \wedge F = \emptyset && \text{for } \bowtie \in \{=, \neq\} \\ \mathbb{T}(\psi_1 \bowtie \psi_2, F) &\triangleq \mathbb{T}(\psi_1, F) \bowtie \mathbb{T}(\psi_2, F) && \text{for } \bowtie \in \{\wedge, \vee, \wedge \neg\} \\ \mathbb{T}(x \mapsto \langle f_i : f_i \rangle_{i \in I}, F) &\triangleq F = \{x\} \wedge \bigwedge_{i \in I} h_{f_i}[x] = f_i \end{aligned}$$

The translation of boolean connectives follows the boolean structure and propagates the domain symbol F to the operands. The translation of pointer assertions postulates content of memory cells represented by arrays and also requires the domain F to be $\{x\}$.

Translation of separating conjunctions. The semantics of separating conjunctions involves a quantification over sets (heap domains). The most direct way of translation is to use quantifiers over sets leading to decidable formulae due to the bounded location domain. This approach combined with a counterexample-guided quantifier instantiation is used in the decision procedure for a fragment of SL supported in CVC5 [28]. In some fragments, however, separating conjunctions can be translated in a way that completely avoids quantifiers. An example is the fragment of boolean combinations of symbolic heaps which has the so-called *unique footprint property* (UFP) [15, 26]—a formula ψ has a (unique) footprint in a model (s, h) with $(s, h) \models \psi * \text{true}$ ⁵, if there exists a (unique) set F such that $(s, h|_F) \models \psi$. The UFP-based approaches of [15, 26] axiomatize the footprints during translation and check operands of separating conjunctions just on the sub-heaps induced by their footprints.

However, UFP does not hold for BSL because of disjunctions. As an example, take the formula $\psi \triangleq x \mapsto y \vee \text{emp}$ and the heap $h = \{x \mapsto y\}$. Both $(s, h|_{\{s(x)\}}) \models \psi$ and $(s, h|_{\emptyset}) \models \psi$ hold. The sets $\{s(x)\}$ and \emptyset are, however, the only footprints of ψ in (s, h) , and this observation can be used to generalise the idea of footprints beyond the fragment in which they are unique.

⁵ Assuming the standard semantics of true which is not part of our logic.

Instead of axiomatizing the footprints, our translation builds a set of footprint terms for operands of separating conjunctions. This change can be also seen as a simplification of the former translations as it eliminates the need to deal with two kinds of formulae (the actual translation and footprint axioms), which must be treated differently during the translation. However, the precise computation of the set of all footprints of ψ in (s, h) , denoted as $\text{FP}_{(s,h)}(\psi)$, is as hard as satisfiability—when the set of footprints is non-empty, the formula ψ is satisfiable. Therefore, we compute just an over-approximation denoted as $\text{FP}^\#(\psi)$. This is justified by the following lemma which gives an equivalent semantics of the separating conjunction in terms of footprints.

Lemma 3. *Let $\varphi \triangleq \psi_1 * \psi_2$ and let (s, h) be a model. Let \mathcal{F}_1 and \mathcal{F}_2 be sets of locations such that $\text{FP}_{(s,h)}(\psi_i) \subseteq \mathcal{F}_i$. Then $(s, h) \models \psi_1 * \psi_2$ iff*

$$\bigvee_{F_1 \in \mathcal{F}_1} \bigvee_{F_2 \in \mathcal{F}_2} \bigwedge_{i=1,2} (s, h|_{F_i}) \models \psi_i \wedge F_1 \cap F_2 = \emptyset \wedge F_1 \cup F_2 = \text{dom}(h).$$

Intuitively, to check whether a separating conjunction holds in a model, it is not necessary to check all possible splits of the heap, but only the splits induced by (possibly over-approximated) footprints of its operands. The lemma is therefore a generalisation of UFP and leads to the following definition of the translation $\mathbb{T}(\psi_1 * \psi_2, F)$:

$$\exists F_1 \in \mathcal{F}_1. \exists F_2 \in \mathcal{F}_2. \mathbb{T}(\psi_1, F_1) \wedge \mathbb{T}(\psi_2, F_2) \wedge F_1 \cap F_2 = \emptyset \wedge F = F_1 \cup F_2.$$

Here, we use a quantifier expression of the form $\exists x \in X. \psi$ as a placeholder that helps us to define two methods which the translation can use for separating conjunctions:

- The method `SatEnum` computes sets of footprints \mathcal{F}_i as $\text{FP}^\#(\psi_i)$ (the computation is described below) and replaces expressions $\exists x \in X. \psi$ with $\bigvee_{x' \in X} \psi[x'/x]$ as in Lemma 3. This strategy is quite efficient in many practical cases when we can compute small sets of footprints \mathcal{F}_1 and \mathcal{F}_2 .
- The method `SatQuantif` does not compute sets \mathcal{F}_i at all and replaces $\exists x \in X. \psi$ simply with $\exists x. \psi$. This strategy is better when the existential quantifier can be later eliminated by Skolemization or when the set of footprints would be too large.

We now show how to compute the set of footprint terms $\text{FP}^\#(\psi)$. We again postpone inductive predicates to Section 5.3. We just note that their footprints are unique. The cases of pure formulae and pointer assertions follow directly from the definition of their semantics, which requires the heap to be empty and a single pointer, respectively.

$$\text{FP}^\#(x \bowtie y) = \{\emptyset\} \text{ for } \bowtie \in \{=, \neq\} \quad \text{FP}^\#(x \mapsto _) = \{\{x\}\}$$

For the boolean conjunction, we can select from footprints of its operand the one with the lesser cardinality. Since negations have many footprints (consider, e.g., $\neg \text{emp}$), we define the case of the guarded negation by taking footprints of its guard. The disjunction is the only case which brings non-uniqueness as we need to consider footprints of both of its operands.

$$\begin{aligned} \text{FP}^\#(\psi_1 \wedge \neg \psi_2) &= \text{FP}^\#(\psi_1) & \text{FP}^\#(\psi_1 \vee \psi_2) &= \text{FP}^\#(\psi_1) \cup \text{FP}^\#(\psi_2) \\ \text{FP}^\#(\psi_1 \wedge \psi_2) &= \text{if } |\text{FP}^\#(\psi_1)| \leq |\text{FP}^\#(\psi_2)| \text{ then } \text{FP}^\#(\psi_1) \text{ else } \text{FP}^\#(\psi_2) \end{aligned}$$

Finally, we define footprints of the separating conjunction by taking the union $F_1 \cup F_2$ for each pair (F_1, F_2) of footprints of its operands. Notice that here $F_1 \cup F_2$ represents an SMT term, therefore we cannot replace it with a disjoint union which is not available in the classical set theories in SMT. We can, however, use heuristics and filter out terms for which we can statically determine that interpretations of F_1 and F_2 are not disjoint.

$$\text{FP}^\#(\psi_1 * \psi_2) = \{F_1 \cup F_2 \mid F_1 \in \text{FP}^\#(\psi_1) \text{ and } F_2 \in \text{FP}^\#(\psi_2)\}$$

We state the correctness of the footprint computation in the following lemma.

Lemma 4. *Let \mathcal{M} be a first-order model with $\mathcal{M} \models \top(\varphi)$ and let $(s, h) = \top_\varphi^{-1}(\mathcal{M})$. Then we have $\text{FP}_{(s,h)}(\varphi) \subseteq \{F^\mathcal{M} \mid F \in \text{FP}^\#(\varphi)\}$.*

5.3 Translation of Inductive Predicates

To translate inductive predicates, we express them in terms of reachability and paths in the heaps. While unbounded reachability cannot be expressed in first-order logic, we can efficiently express bounded *linear* reachability in our encoding. The linearity means that each path uses only a single field (which is not the case, e.g., for paths in trees). All predicates in this section are parametrised with an interval $[m, n]$ which bounds the length of the considered paths. When we do not state the bounds explicitly, we assume conservative bounds $[0, \text{bound}_S(\varphi)]$ for a path starting from a root of a sort S . We show how to compute more precise bounds in Section 6. We start with the translation of reachability:

$$\text{reach}^{\equiv n}(h, x, y) \triangleq h^n[x] = y \quad \text{reach}^{[m,n]}(h, x, y) \triangleq \bigvee_{m \leq i \leq n} \text{reach}^{\equiv i}(h, x, y)$$

Here, the predicate $\text{reach}^{\equiv n}(h, x, y)$ expresses that x can reach y via a field represented by the array h in exactly n steps. Similarly, $\text{reach}^{[m,n]}$ expresses reachability in m to n steps. Besides reachability, we will need a macro $\text{path}_C(h, x, y)$ expressing the domain of a path from x to y , or the empty set if such a path does not exist:

$$\begin{aligned} \text{path}_C^{\equiv n}(h, x, y) &\triangleq \bigcup_{0 \leq i < n} C(h^i[x]) \\ \text{path}_C^{[m,n]}(h, x, y) &\triangleq \text{if } (\text{reach}^{\equiv m}(h, x, y)) \text{ then } (\text{path}_C^{\equiv m}(h, x, y)) \\ &\quad \dots \text{ else if } (\text{reach}^{\equiv n}(h, x, y)) \text{ then } (\text{path}_C^{\equiv n}(h, x, y)) \text{ else } (\emptyset) \end{aligned}$$

The additional parameter C is a function applied to each element of the path that can be used to define nested paths. We define a simple path $\text{path}_S^{[m,n]}(h, x, y) \triangleq \text{path}_C^{[m,n]}(h, x, y)$ with $C \triangleq \lambda \ell. \{\ell\}$ and a nested path as $\text{path}_N^{[m,n]}(h_1, h_2, x, y, z) \triangleq \text{path}_C^{[m,n]}(h_1, x, y)$ with $C \triangleq \lambda \ell. \text{path}_S(h_2, \ell, z)$. In the case of the nested path, the array h_1 represents the top-level path from x to y , and h_2 represents nested paths terminating in the common location z . Now we can define footprints of inductive predicates using path terms as follows:

$$\begin{aligned} \text{FP}^\#(\pi(x, y)) &= \{\text{path}_S(h_n, x, y)\} && \text{for } \pi \in \{\text{sls}, \text{dls}\} \\ \text{FP}^\#(\text{nls}(x, y, z)) &= \{\text{path}_N(h_t, h_n, x, y, z)\} \end{aligned}$$

The common part of the translation $T(\pi(x, y), F)$ postulates the existence of a top-level path from x to y and a domain F based on this path (formalised in the formula `main_path` below); and ensures that all locations have the correct sort (through the formula `typing`). For DLLs, we add an invariant which ensures that its locations are correctly doubly-linked (the `back_links` formula), and we further need a special treatment of the cases when the list is empty as well as a special treatment for its roots and sinks (cf. the formula `boundaries`). For NLLs, we add an invariant stating that an inner list starts from each location in its top-level path (the `inner_lists` formula) and that those inner paths are disjoint (the `disjoint` formula).

- $T(\text{sls}(x, y), F) \triangleq \text{main_path} \wedge \text{typing}$ where
 - $\text{main_path} \triangleq \text{reach}(h_n, x, y) \wedge F = \text{path}_S(h_n, x, y)$ and $\text{typing} \triangleq F \subseteq D_S$.
- $T(\text{dls}(x, y, x', y'), F) \triangleq \text{empty} \vee \text{nonempty}$ where
 - $\text{empty} \triangleq x = y \wedge x' = y' \wedge F = \emptyset$,
 - $\text{nonempty} \triangleq x \neq y \wedge x' \neq y' \wedge \text{main_path} \wedge \text{boundaries} \wedge \text{typing} \wedge \text{back_links}$,
 - $\text{main_path} \triangleq \text{reach}(h_n, x, y) \wedge F = \text{path}_S(h_n, x, y)$,
 - $\text{boundaries} \triangleq h_p[x] = y' \wedge h_n[x'] = y \wedge x' \in F \wedge y' \notin F$,
 - $\text{typing} \triangleq F \subseteq D_D$,
 - $\text{back_links} \triangleq \forall \ell. (\ell \in F \wedge \ell \neq x') \rightarrow h_p[h_n[\ell]] = \ell$.
- $T(\text{nls}(x, y, z), F) \triangleq \text{main_path} \wedge \text{typing} \wedge \text{inner_lists} \wedge \text{disjoint}$ where
 - $\text{main_path} \triangleq \text{reach}(h_t, x, y) \wedge F = \text{path}_N(h_t, h_n, x, y, z)$,
 - $\text{typing} \triangleq \text{path}_S(h_t, x, y) \subseteq D_N \wedge F \setminus \text{path}_S(h_t, x, y) \subseteq D_S$,
 - $\text{inner_lists} \triangleq \forall \ell. \ell \in F \cap D_N \rightarrow \text{reach}(h_n, h[\ell], z)$,
 - $\text{disjoint} \triangleq \forall \ell_1, \ell_2. (\{\ell_1, \ell_2\} \subseteq F \wedge \ell_1 \neq \ell_2 \wedge h_n[\ell_1] = h_n[\ell_2]) \rightarrow h_n[\ell_1] \notin F^6$.

Path quantifiers. Invariants of paths are naturally expressed using universal quantifiers. For quantifiers, however, we cannot directly take advantage of bounds on path lengths. Therefore, similarly as for separating conjunctions, we use the idea of replacing quantifiers by small enumerations of their instances, which is efficient when we can compute small enough bounds on the paths. For example, if we know that the length of an f -path with a root x is at most two, it is enough to instantiate its invariant for x , $h_f[x]$, and $h_f^2[x]$. This idea is formalised using expressions $\mathbb{P}_{(h,x)}^{\leq n} \ell. \psi$, which we call *path quantifiers* and which state that ψ holds for all locations of the path with the length n starting from x via the array h :

$$\mathbb{P}_{(h,x)}^{\leq n} \ell. \psi \triangleq \bigwedge_{0 \leq i \leq n} \psi[h^i[x]/\ell].$$

If we need to quantify over nested paths, we need to use two path quantifiers (one for the top-level path and one for the nested paths). The quantifiers in the last conjunct of the NLL translation can be rewritten as $\mathbb{P}_{(h_t,x)} \ell'_1. \mathbb{P}_{(h_t,x)} \ell'_2. \mathbb{P}_{(h_n,\ell'_1)} \ell_1. \mathbb{P}_{(h_n,\ell'_2)} \ell_2$. In this expression, ℓ'_1 and ℓ'_2 range over locations in the top-level list, and ℓ_1 and ℓ_2 range over locations in the nested paths starting from ℓ'_1 and ℓ'_2 , respectively.

⁶ We could also write $h_n[\ell_1] = z$ instead of $h_n[\ell_1] \notin F$, but the latter leads to better performance of SMT solvers.

5.4 Complexity

This section briefly discusses the complexity of the proposed decision procedure as well as the complexity lower bound for the satisfiability problem in the considered fragment of SL. We will use $\text{SAT}(\omega_1, \dots, \omega_n)$ to denote the satisfiability problem for a sub-fragment constructed of atomic formulae and the connectives ω_i and $\text{SAT}(\overline{\omega_1, \dots, \omega_n})$ to denote the fragment where none of the connectives ω_i appear.

Theorem 4. *The procedure `SatQuantif` produces formula of polynomial size, and, for $\text{SAT}(\overline{\wedge\neg})$, it runs in NP. The procedure `SatEnum` runs in NP for $\text{SAT}(\overline{\vee})$.*

Proof (sketch). When not considering the instantiation of quantifiers over footprints, both `SatQuantif` and `SatEnum` produce a formula $T(\varphi)$ of a polynomial size dominated by the translation of inductive predicates. For the variant of the translation of inductive predicates using universal quantifiers over locations, the size is $\mathcal{O}(n^3)$ for SLLs and DLLs (dominated by the $\mathcal{O}(n^3)$ size of the path_S term), and $\mathcal{O}(n^5)$ for NLLs (dominated by path_N). If the input formula does not contain guarded negations, then all quantifiers can be eliminated using Skolemization. The translated formulae are then in a theory decidable in NP (e.g., when sets are encoded as extended arrays [21]).

The procedure `SatEnum` can produce exponentially large formulae because of the footprint enumeration. This can be prevented if the input formula does not contain disjunctions, in which case the footprints of all sub-formulae are unique, i.e., singleton sets. The translated formulae are then again in a theory decidable in NP. \square

Theorem 5. *$\text{SAT}(\overline{\mapsto}, \overline{\wedge\neg}, \overline{\wedge}, \overline{\vee}, \overline{*})$ is PSPACE-complete.*

Proof (sketch). Membership in PSPACE was proved in [25] for a more expressive fragment. For the hardness part, we build on the reduction from QBF used in [7]. In this reduction, the boolean value of a variable is represented by the corresponding SL variable being allocated (always pointing to nil for simplicity). The fact that x is false is expressed using a negative points-to predicate stating that x is not allocated. The existential quantifier is expressed using the separating conjunction, and the universal quantifier is obtained using the (unguarded) negation. (For details, see [7].)

We show that this reduction can be done without the unguarded negation and the negative points-to assertion, using the guarded negation instead. The key observation is that, for a QBF formula with variables X , we can express that all variables in X can have arbitrary boolean values as $\text{arbitrary}[X] \triangleq *_{x \in X} (x \mapsto \text{nil} \vee \text{emp})$. In the context of variables X , we can then express negation as $\neg F \triangleq \text{arbitrary}[X] \wedge \neg F$ and the truth values of a variable x as $\neg x \triangleq \text{arbitrary}[X \setminus \{x\}]$ and $x \triangleq \text{arbitrary}[X] * x \mapsto \text{nil}$. The rest of the reduction then easily follows [7]. \square

6 Optimised Bound Computation

In many practical cases, the main source of complexity is the translation of inductive predicates, which heavily depends on the possible lengths of paths between locations. We now propose how to bound the length of these paths based on the so-called *SL-graphs* which are graph representations of constraints imposed by SL formulae. SL-graphs were originally used for representation and deciding of symbolic heaps with lists in [8]. Here, we use their generalised form which captures must-relations holding in all models of a given formula. Note that the nodes of the graphs are implicitly given by the domains of the involved relations, which themselves can be viewed as edges.

Definition 3. An SL-graph of φ is a tuple $G[\varphi] = (\ominus, \oplus, (\ominus_f, \ominus_f^*), \otimes_f)_{f \in \text{Field}}$ where:

- $\ominus \subseteq \text{vars}(\varphi) \times \text{vars}(\varphi)$ is an equivalence relation called *must-equality*,
- $\oplus \subseteq \text{vars}(\varphi) \times \text{vars}(\varphi)$ is a symmetric relation called *must-disequality*,
- $\ominus_f \subseteq \text{vars}(\varphi) \times \text{vars}(\varphi)$ is a *must-f-pointer relation*,
- $\ominus_f^* \subseteq \text{vars}(\varphi) \times \text{vars}(\varphi)$ is an *irreflexive must-f-path relation*,
- $\otimes_f \subseteq \text{vars}(\varphi)^2 \times \text{vars}(\varphi)^2$ is a symmetric relation called *must-f-path-disjointness*.

Except \otimes_f , the components of $G[\varphi]$ represent atomic formulae—equalities, disequalities, pointers, and paths (i.e., list segments)—holding within all models of φ . The fact that $(x_1, y_1) \otimes_f (x_2, y_2)$ states that, in all models of φ , the domains of f-paths from x_1 to y_1 and from x_2 to y_2 are disjoint.

To compute the SL-graph $G[\varphi]$, we define some auxiliary notation. We define G_\emptyset to be an SL-graph where all the relations are empty. We write $G \triangleleft \{x_i \bowtie_i y_i\}_{i \in I}$ to denote the SL-graph G' which is the same as G with the elements $x_i \bowtie_i y_i$ for $i \in I$ added to the corresponding relations. We use \sqcup and \sqcap as a component-wise union and intersection of SL-graphs, respectively. We define the disjoint union of SL-graphs as:

$$\begin{aligned} G_1 \boxplus G_2 &= (G_1 \sqcup G_2) \\ &\triangleleft \{x \oplus y \mid x \in \text{alloc}(G_1), y \in \text{alloc}(G_2), \text{ and } (x \text{ is not nil or } y \text{ is not nil})\} \\ &\triangleleft \{e_1 \otimes_f e_2 \mid f \in \text{Field}, e_1 \in \text{paths}_f(G_1), \text{ and } e_2 \in \text{paths}_f(G_2)\}. \end{aligned}$$

Here, $\text{paths}_f(G)$ is defined as $\ominus_f \cup \ominus_f^*$, and the set of must-allocated variables is $\text{alloc}(G) = \{x \mid \exists y, f. x \ominus_f y \text{ or } (x \ominus_f^* y \text{ and } x \oplus y)\} \cup \{\text{nil}\}$ (nil is added for technical reasons). We further assume that all operations on SL-graphs (\triangleleft , \sqcup , \sqcap , and \boxplus) preserve relational properties (symmetry, transitivity, etc.) of the components of SL-graphs by computing the corresponding closures after the operation is performed. We compute the SL-graph $G[\varphi]$ as follows.

$$\begin{aligned} G[x = y] &= G_\emptyset \triangleleft \{x \ominus y\} & G[x \mapsto \langle f_i : f_i \rangle_{i \in I}] &= G_\emptyset \triangleleft \{x \ominus_{f_i} f_i\}_{i \in I} \\ G[x \neq y] &= G_\emptyset \triangleleft \{x \oplus y\} & G[\text{sls}(x, y)] &= G_\emptyset \triangleleft \{x \ominus_n y\} \\ G[\psi_1 \wedge \neg \psi_2] &= G[\psi_1] & G[\text{dls}(x, y, x', y')] &= G_\emptyset \triangleleft \{x \ominus_n y, x' \ominus_p y'\} \\ G[\psi_1 \wedge \psi_2] &= G[\psi_1] \sqcup G[\psi_2] & G[\text{nls}(x, y, z)] &= G_\emptyset \triangleleft \{x \ominus_n z, x \ominus_t y\} \\ G[\psi_1 \vee \psi_2] &= G[\psi_1] \sqcap G[\psi_2] & G[\psi_1 * \psi_2] &= G[\psi_1] \boxplus G[\psi_2] \end{aligned}$$

Observe that we only approximate dls and nls. After the construction is finished, we apply the following rules for matching of pointers and for detection of inconsistencies.

$$\frac{x_1 \ominus_f y_1 \quad x_2 \ominus_f y_2 \quad x_1 \ominus x_2}{y_1 \ominus y_2} \text{ (}\mapsto\text{-match)} \qquad \frac{x \ominus y \quad x \oplus y}{\varphi \text{ is unsat}} \text{ (contradiction)}$$

Tighter location bounds. Using SL-graphs, we can slightly improve the location bound from Section 4 by considering equivalence classes of \ominus instead of individual variables (this can be also used to refine the later described path bound computation) and by defining $\|x\| = 1$ if x is a must-pointers, i.e., $x \ominus_f y$ for some f and y .

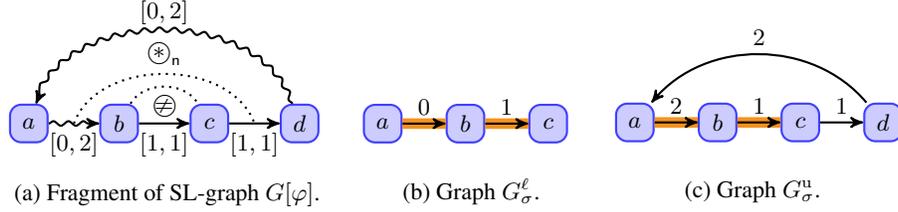


Fig. 3: An illustration of the bound computation for the path σ from a to c on a fragment of SL-graph of $\varphi \triangleq (\text{sls}(a, b) * b \mapsto c * c \mapsto d * \text{sls}(d, a)) \wedge \neg(\text{sls}(a, c) * \text{sls}(c, a))$. The highlighted edges denote the paths used to determine the bound $[1, 3]$.

Path bounds. We now fix an f-path σ from x^S to y and show how to compute an interval $[\ell, u]$ that gives bounds on its length. The computation of the path bounds runs in two steps. In the first step, we compute an initial bound $[\ell_e^0, u_e^0]$ for each edge $e \in \text{paths}_f(G)$. If e is a pointer edge, its bound is given as $[1, 1]$. For a path edge $e = (a, b)$, we define $\ell_e^0 = 1$ if $a \oplus b$ and 0 otherwise; while u_e^0 is defined as $\text{bound}_S(\varphi) - \sum_{v \in V} \|v\|$ where $V = \{v \in \text{vars}_S(\varphi) \mid v \text{ is not } x \text{ and } \exists u. (v, u) \oplus_f (x, y)\}$. This way, we exclude from the computation of the initial upper bound the source v of each path disjoint with σ and all locations possibly allocated in a chunk with the root v . Note that it can be the case that the actual size of this chunk has a lesser size than $\|v\|$, but this means that we were too conservative when computing the global location bound and can decrease the path bound by the same number anyway.

In the second phase, we compute the bounds of the path σ using initial bounds from the first step. The computation is based on two weighted directed graphs derived from the SL-graph G : G_σ^u for the upper bound and G_σ^ℓ for the lower bound (in both cases, the vertices are implicitly given as $\text{vars}(\varphi)$, and the edge weight of an edge e is given by u_e^0 and ℓ_e^0 computed in the previous step, respectively):

$$G_\sigma^u = \{a \rightarrow b \mid (a, b) \in \text{paths}_f(G)\},$$

$$G_\sigma^\ell = \{a \rightarrow b \mid (a \ominus_f b \text{ and } a \oplus_f y) \text{ or } (a \ominus_f b \text{ and } \exists w. \text{nonempty}(y, w) \text{ and } (y, w) \oplus_f (a, b))\}.$$

Here, the condition $\text{nonempty}(y, w)$ states that a directed SL-graph edge (y, w) is non-empty which holds if either $y \ominus_f w$, or when $y \ominus_f w$ and $y \oplus_f w$.

Intuitively, the upper bound u is computed as the length of the shortest path from x to y in G_σ^u . Since f-paths are uniquely determined, we know that no path can be longer than the shortest one, and thus u is indeed a correct upper bound. The lower bound ℓ is computed as the length of the longest path starting from x (ending anywhere) in G_σ^ℓ . By construction, G_σ^ℓ contains only those edges for which one can prove that they cannot contain y in their domains. A path from x of a length ℓ therefore implies that x cannot reach y in less than ℓ steps, and thus ℓ is indeed a correct lower bound.

Example. We demonstrate the path bound computation in Fig. 3, which shows a fragment of the SL-graph of a formula φ (it shows only those \oplus_n edges that are relevant in our example) and the graphs G_σ^ℓ and G_σ^u for the path σ from a to c . We have that $\|b\| = \|c\| = 1$ and $\|a\| = \|d\| = 2$. This gives us the location bound, which is 6. In the first phase, we compute the initial bound $[0, 2]$ for paths of the predicates $\text{sls}(a, b)$ and $\text{sls}(d, a)$ because both of them are disjoint with all the other paths in $G[\varphi]$. In the second phase, we get the bound for σ equal to $[1, 3]$ instead of the default bound $[0, 6]$.

7 Experimental Evaluation

We have implemented the proposed decision procedure in a new solver called `ASTRAL`⁷. `ASTRAL` is written in OCaml and can use multiple backend SMT solvers. With the encoding presented in Section 5, it can use either `CVC5` supporting set theory directly [1] or `Z3` supporting it by a reduction to the extended theory of arrays [21]. We have also developed an alternative encoding in which both locations and location sets are represented as bitvectors. The bitvector encoding differs only in expressing set operations on the level of bitvectors with additional axioms ensuring that all locations “can fit” into sets encoded by the bitvectors (for details, see Appendix D). With the bitvector encoding, a backend solver only needs to support theories of bitvectors and arrays, which are both standard and supported by many other SMT solvers. Another advantage is that the quantification on bitvectors seems to perform significantly better than on sets.

In our experiments, if we do not say explicitly which encoding and solver is used, we use the bitvector encoding and `BITWUZLA` [24] as the backend solver, which we found to be the best performing combination. We set a limit for the method `SatEnum` to 64 footprints. If this limit is exceeded, we dynamically switch to `SatQuantif`. We use path quantifiers when the path bound is at most half of the domain bound. These are design choices that can be revisited in the future.

All experiments were run on a machine with 2.5 GHz Intel Core i5-7300HQ CPU and 16 GiB RAM, running Ubuntu 18.04. The timeout was set to 60 s and the memory limit to 1 GB. Our experiments were conducted using `BENCHEXEC` [4], a framework for reliable benchmarking.

7.1 Entailments of Symbolic Heaps

In the first part of our evaluation, we focus on formulae from the symbolic heap fragment which is frequently used by verification tools and for which there exist many dedicated solvers. We therefore do not expect to outperform the best existing tools but rather to obtain a comparison with other translation-based decision procedures.

In Table 1a, we provide results for the category `QF_SHLID_ENTL` (entailments with SLLs). We divide the category into two subsets: verification conditions (which are simpler) and more complex artificially generated formulae “*bolognesa*” and “*clones*” from [22]. During the experiments, we found out that several “cloned” entailments contain root variables on the right-hand side of the entailment that do not appear on the left-hand side, making the entailment trivially invalid when its left-hand side is satisfiable. For a few hard clone instances, this makes a problem for `ASTRAL` as it cannot use the path bound computation as such roots do not appear in the `SL`-graph. We have therefore implemented a heuristic that detects entailments $\varphi \models \psi$ that can be reduced to satisfiability of φ . Since this is a benchmark-specific heuristic, we present also the version without this heuristic (`ASTRAL*`) in Table 1a. The optimised version of `ASTRAL` is able to solve all the formulae being faster than other translation-based solvers `GRASSHOPPER`⁸ and `SLOTH`. For illustration, the table further contains the second best solver in the latest edition of `SL-COMP`, `S2S`⁹.

⁷ <https://github.com/TDacik/Astral>

⁸ Since `GRASSHOPPER` is not a solver but a verification tool, we encode the entailment checking as a verification of an empty program.

⁹ We had technical issues running the winner `ASTERIX` [23]. The difference between those tools is, however, negligible.

Table 1: Experimental results for formulae from SL-COMP. The columns are: solved instances (OK), out of time/memory (RO), instances on which ASTRAL wins—ASTRAL can solve it and the other solver not or ASTRAL solves it faster (WIN), instances solved in the time limits of 0.1 s and 1 s, and the total time for solved instances in seconds.

(a) Results for the category QF_SHLS_ENTL.

Solver	Verification conditions (86)						bolognesa+clones (210)					
	OK	RO	WIN	<0.1	≤1	Total time	OK	RO	WIN	<0.1	≤1	Total time
ASTRAL	86	0	-	84	86	4.62	210	0	-	68	169	202.91
ASTRAL*	86	0	42	83	86	4.64	195	15	88	64	150	408.48
GRASSHOPPER	86	0	70	52	86	8.65	203	7	148	60	87	1229.35
S2S	86	0	5	86	86	2.08	210	0	3	203	210	8.18
SLOTH	64	3	86	0	28	235.28	70	140	210	0	50	149.42

(b) Results for a subset of the category QF_SHLID_ENTL.

Solver	Doubly-linked lists (17)						Nested singly-linked lists (19)					
	OK	RO	WIN	<0.1	≤1	Total time	OK	RO	WIN	<0.1	≤1	Total time
ASTRAL	17	0	-	11	17	2.72	19	0	-	3	9	86.93
GRASSHOPPER	17	0	16	3	15	7.53	-	-	-	-	-	-
HARRSH	17	0	17	0	0	95.18	14	5	18	0	0	183.01
S2S	17	0	0	17	17	0.15	19	0	0	19	19	0.43
SONGBIRD	11	5	14	5	9	13.39	11	5	8	4	11	1.38

In Table 1b, we provide results for a subset of the category QF_SHLID_ENTL (entailments with linear inductive definitions from which we selected DLLs and NLLs) for ASTRAL and three best-performing solvers competing in the latest edition of SL-COMP—S2S, SONGBIRD (in the version with automated lemma synthesis called SLS), and HARRSH. We also include GRASSHOPPER which supports DLLs only. Except S2S which solves almost all formulae virtually immediately, ASTRAL is the only one able to solve all the formulae in the given time limit.

7.2 Experiments on Formulae Outside of the Symbolic Heap Fragment

For formulae outside of the symbolic heap fragment and its top-level boolean closure, there are currently no existing benchmarks. For now, we therefore limit ourselves to randomly generated but extensive sets of formulae. In the future, we would like to develop a program analyser using symbolic execution over BSL and make more careful experiments on realistic formulae.

We first focus on the fragment with guarded negations but without inductive predicates, on which we can compare ASTRAL with CVC5. We have prepared a set of 1000 entailments of the form $\varphi \models \psi$ which are generated as random binary trees with depth 8 over 8 variables with the only atoms being pointer assertions. To reduce the number of trivial instances, we only generated formulae for which $\text{vars}(\psi) \subseteq \text{vars}(\varphi)$ and ASTRAL cannot deduce contradiction from their SL-graphs. To avoid any suspicion that the difference is caused by better performance of the backend solver rather than the design of our translation, we used ASTRAL with the CVC5 backend and direct set

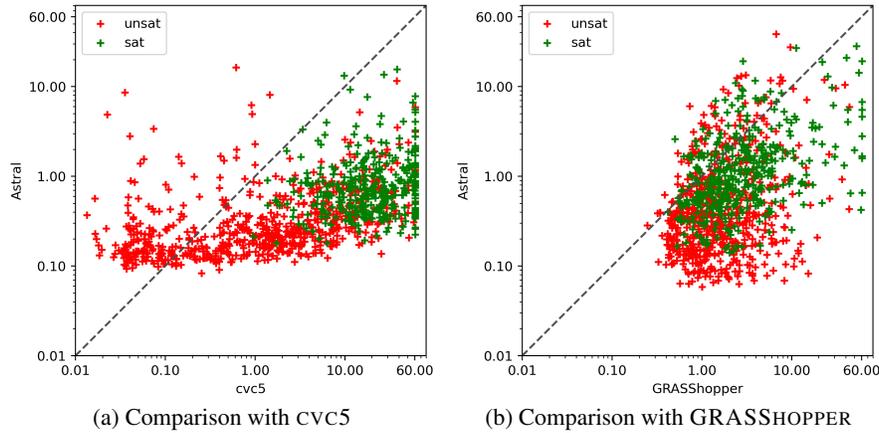


Fig. 4: A comparison of ASTRAL with CVC5 and GRASSHOPPER on randomly generated formulae. Times are in seconds, axes are logarithmic. The timeout was set to 60 s.

encoding (with BITWUZLA and bitvector encoding, our results would be even better). The results are given in Fig. 4a and suggest that our treatment of guarded negations really brings a better performance—ASTRAL can solve all the instances and almost all of them under 10 seconds. On the other hand, CVC5 timed out in 61 cases and is usually slower than ASTRAL, in particular on satisfiable formulae which represent invalid entailments.

In the second experiment, we compared our solver with GRASSHOPPER on the fragment which it supports, i.e., arbitrary nesting of conjunctions and disjunctions. We again generated 1000 entailments, this time with depth 6, 6 variables and with atoms being singly-linked lists (with 20 % probability) or pointer-assertions. The results are given in Fig. 4b. ASTRAL ran out of memory in 5 cases, and GRASSHOPPER timed out in 10 cases. In summary, ASTRAL is faster on more than 80 % of the formulae with an almost 3 times lesser running time.

Finally, to illustrate that ASTRAL can indeed handle formulae out of the fragments of all the other mentioned tools, we apply it on an entailment query that involves the formula mentioned at the end of the introduction: $((\text{sls}(x, y) \wedge \neg(\text{sls}(x, z) * \text{sls}(z, y))) * y \mapsto z) \models \text{sls}(x, z)$, converted to an unsatisfiability query. ASTRAL resolves the query in 0.12 s. Note that without the requirement $\neg(\text{sls}(x, z) * \text{sls}(z, y))$, the entailment does not hold as a cycle may be closed in the heap.

8 Conclusions and Future Work

We have presented a novel decision procedure based on a small-model property and translation to SMT. Our experiments have shown very promising results, especially for formulae with rich boolean structure for which our decision procedure outperforms other approaches (apart from being able to solve more formulae).

In the future, we would like to extend our approach with some class of user-defined inductive predicates, with more complex spatial connectives such as septractions and/or magic wands, consider a lazy and/or interactive translation instead of the current eager approach, and try ASTRAL within some SL-based program analyser.

References

1. Bansal, K., Barrett, C., Reynolds, A., Tinelli, C.: A New Decision Procedure for Finite Sets and Cardinality Constraints in SMT. In: IJCAR (2017)
2. Batz, K., Fesefeldt, I., Jansen, M., Katoen, J.P., Keßler, F., Matheja, C., Noll, T.: Foundations for Entailment Checking in Quantitative Separation Logic. In: ESOP (2022)
3. Berdine, J., Calcagno, C., O’Hearn, P.W.: A Decidable Fragment of Separation Logic. In: FSTTCS 2004. LNCS, vol. 3328 (2004)
4. Beyer, D., Löwe, S., Wendler, P.: Reliable Benchmarking: Requirements and Solutions. International Journal on Software Tools for Technology Transfer **21** (2017)
5. Brotherston, J., Gorogiannis, N., Petersen, R.L.: A Generic Cyclic Theorem Prover. In: APLAS. LNCS, vol. 7705 (2012)
6. Calcagno, C., Distefano, D., O’Hearn, P., Yang, H.: Compositional Shape Analysis by Means of Bi-Abduction. Journal of the ACM **58**(6) (2011)
7. Calcagno, C., Yang, H., O’Hearn, P.W.: Computability and Complexity Results for a Spatial Assertion Language for Data Structures. In: FST TCS (2001)
8. Cook, B., Haase, C., Ouaknine, J., Parkinson, M., Worrell, J.: Tractable Reasoning in a Fragment of Separation Logic. In: CONCUR. LNCS, vol. 3901 (2011)
9. Echenim, M., Iosif, R., Peltier, N.: The Bernays-Schönfinkel-Ramsey Class of Separation Logic with Uninterpreted Predicates. ACM Transactions on Computational Logic **21** (2019)
10. Enea, C., Lengál, O., Sighireanu, M., Vojnar, T.: Compositional Entailment Checking for a Fragment of Separation Logic. In: APLAS (2014)
11. Holík, L., Peringer, P., Rogalewicz, A., Šoková, V., Vojnar, T., Zuleger, F.: Low-level bi-abduction. In: ECOOP 2022. LIPIcs, vol. 222, pp. 19:1–19:30 (2022)
12. Iosif, R., Rogalewicz, A., Vojnar, T.: Deciding Entailments in Inductive Separation Logic with Tree Automata. In: ATVA (2014)
13. Iosif, R., Zuleger, F.: Expressiveness results for an inductive logic of separated relations. In: Pérez, G.A., Raskin, J. (eds.) CONCUR. LIPIcs, vol. 279, pp. 20:1–20:20 (2023). <https://doi.org/10.4230/LIPICS.CONCUR.2023.20>, <https://doi.org/10.4230/LIPICS.CONCUR.2023.20>
14. Ishtiaq, S., O’Hearn, P.: Separation and Information Hiding. In: Proc. of POPL’01. ACM (2001)
15. Katelaan, J., Jovanovic, D., Weissenbacher, G.: A Separation Logic with Data: Small Models and Automation. In: IJCAR (2018)
16. Katelaan, J., Matheja, C., Noll, T., Zuleger, F.: Harrsh: A Tool for Unied Reasoning about Symbolic-Heap Separation Logic. In: LPAR-22 Workshop and Short Paper Proceedings. vol. 9 (2018)
17. Le, Q.L., Gherghina, C., Qin, S., Chin, W.N.: Shape Analysis via Second-Order Bi-Abduction. In: Proc. of CAV’14. LNCS, vol. 8559. Springer (2014)
18. Le, Q.L.: Compositional Satisfiability Solving in Separation Logic. In: VMCAI. LNCS, vol. 12597 (2021)
19. Le, Q.L., Le, X.B.D.: An Efficient Cyclic Entailment Procedure in a Fragment of Separation Logic. In: FoSSaCS (2023)
20. Matheja, C., Pagel, J., Zuleger, F.: A Decision Procedure for Guarded Separation Logic Complete Entailment Checking for Separation Logic with Inductive Definitions. ACM Trans. Comput. Logic **24**(1) (2023)
21. de Moura, L., Bjørner, N.: Generalized, efficient array decision procedures. In: FMCAD (2009)
22. Navarro Pérez, J.A., Rybalchenko, A.: Separation Logic + Superposition Calculus = Heap Theorem Prover. In: PLDI (2011)

23. Navarro Pérez, J.A., Rybalchenko, A.: Separation Logic Modulo Theories. In: APLAS. LNCS, vol. 8301 (2013)
24. Niemetz, A., Preiner, M.: Bitwuzla. In: CAV. LNCS, vol. 13965 (2023)
25. Pagel, J., Zuleger, F.: Strong-separation logic. *ACM Trans. Program. Lang. Syst.* **44**(3), 16:1–16:40 (2022). <https://doi.org/10.1145/3498847>, <https://doi.org/10.1145/3498847>
26. Piskac, R., Wies, T., Zufferey, D.: Automating Separation Logic Using SMT. In: CAV (2013)
27. Piskac, R., Wies, T., Zufferey, D.: Automating Separation Logic with Trees and Data. In: CAV (2014)
28. Reynolds, A., Iosif, R., King, T.: A Decision Procedure for Separation Logic in SMT. In: ATVA (2016)
29. Reynolds, J.: Separation Logic: A Logic for Shared Mutable Data Structures. In: Proceedings 17th Annual IEEE Symposium on Logic in Computer Science (2002)
30. Santos, J., Maksimovic, P., Ayoun, S.E., Gardner, P.: Gillian, Part I: A Multi-Language Platform for Symbolic Execution. In: Proc. of PLDI'20. ACM (2020)
31. Summers, A.J., Müller, P.: Automating deductive verification for weak-memory programs (extended version). *Int. J. Softw. Tools Technol. Transf.* **22**(6), 709–728 (2020). <https://doi.org/10.1007/s10009-020-00559-y>
32. Ta, Q.T., Le, T.C., Khoo, S.C., Chin, W.N.: Automated Lemma Synthesis in Symbolic-Heap Separation Logic. In: POPL (2018)
33. Yang, H., Lee, O., Berdine, J., Calcagno, C., Cook, B., Distefano, D., O'Hearn, P.: Scalable Shape Analysis for Systems Code. In: Proc. of CAV'08. LNCS, vol. 5123. Springer (2008)

A Proof of Small-Model Property

This section presents proofs omitted in Section 4. We first introduce some additional notation. If a stack s is clear from the context, we call a heap h positive (atomic) if (s, h) is positive (atomic). We say that a model (s, h) is a canonical model of a formula φ if $(s, h) \models \varphi$ and $\downarrow^X(s, h) = (s, h)$ for $X = \text{vars}(\varphi)$.

Proof (Lemma 2). Let (s, h) be an atomic model with $(s, h) \models \varphi$. Consequently $h \neq \emptyset$ by the definition of atomicity. We proceed by structural induction on φ . Cases of pure formulae cannot happen since $h \neq \emptyset$. Cases of pointer assertions are trivial as they always meet precisely the condition (1). Cases of inductive predicates are as follows:

- If $(s, h) \models \text{sIs}(x, y)$, then h is either a single pointer $x \mapsto y$ that satisfies (1), or otherwise it satisfies (2).
- If $(s, h) \models \text{dIs}(x, y, x', y')$, then h is either a single pointer $x \mapsto \langle n : y, p : y' \rangle$ that satisfies (1), or otherwise it satisfies (3). The case when h consists of two pointers $x \mapsto \langle n : x', p : y' \rangle$ and $x' \mapsto \langle n : y, p : x \rangle$ cannot happen because it contradicts the assumption that (s, h) is atomic.
- If $(s, h) \models \text{nIs}(x, y, z)$, then h is either a single pointer $x \mapsto \langle n : z, t : y \rangle$ which satisfies (1), or otherwise it satisfies (4).

If $(s, h) \models \psi_1 * \psi_2$, then there exist disjoint heaps h_1 and h_2 with $h = h_1 \uplus h_2 \neq \perp$ satisfying ψ_1 and ψ_2 , respectively. Since (s, h) is atomic, exactly one of the heaps h_1 and h_2 needs to be empty. Assume w.l.o.g. that $h_1 = \emptyset$ and $h = h_2$. Hence $(s, h) \models \psi_2$ and we obtain the claim using IH. The cases of boolean connectives follow directly from IH. \square

To prove Lemma 1, we introduce a definition of *split point*. Intuitively, a split point is a location that can be used to split a model of an inductive predicate into two non-empty models of the same predicate type. If a model contains a split point, then it is not atomic and consequently also not a chunk.

Definition 4 (Split point). Let (s, h) be a model such that $h \neq \emptyset$ and $(s, h) \models \pi(x, y)$ for an inductive predicate $\pi(x, y)$. Let $\ell \in \text{dom}(h)$ be a location such that $\ell \in \text{img}(s)$. We say that ℓ is a split point of (s, h) if $s(x) \neq \ell$ and there exist x', y', z such that one of the following conditions holds:

- $\pi \triangleq \text{sIs}(x, y)$,
- $\pi \triangleq \text{nIs}(x, y, z)$ and $h(\ell, t) \neq \perp$,
- $\pi \triangleq \text{dIs}(x, y, x', y')$ and $h(\ell, p) \in \text{img}(s)$.

Lemma 5. Let (s, h) be a model such that $(s, h) \models \pi$ for some inductive predicate π . If (s, h) contains a split point, then it is not atomic.

Proof. Let ℓ be a split point of (s, h) and let p be some variable such that $s(p) = \ell$. We will show that (s, h) is not atomic as it can be split into two non-empty positive sub-heaps satisfying formulae ψ_1 and ψ_2 given as follows:

- If $\pi \triangleq \text{sIs}(x, y)$, then $\psi_1 \triangleq \text{sIs}(x, p)$ and $\psi_2 \triangleq \text{sIs}(p, y)$.
- If $\pi \triangleq \text{nIs}(x, y, z)$, then $\psi_1 \triangleq \text{nIs}(x, p, z)$ and $\psi_2 \triangleq \text{nIs}(p, y, z)$.

- If $\pi \triangleq \text{dls}(x, y, x', y')$, then $\psi_1 \triangleq \text{dls}(x, p, p', y')$ and $\psi_2 \triangleq \text{dls}(p, y, x', p')$, where p' is some variable such that $s(p') = h(\ell, p)$ whose existence is guaranteed by the definition of split points.

In all the cases, non-emptiness of the model of ψ_1 follows from the fact that $s(x) \neq s(p)$ by the definition of split points; and non-emptiness of the model of ψ_2 follows from the fact that $s(y) \neq s(p)$ because $s(p) \in \text{dom}(h)$, but $s(y) \notin \text{dom}(h)$. \square

We prove Lemma 1 by proving the following stronger claim.

Lemma 6. *Let (s, h) be a positive model. The the following claims hold:*

1. $h = \bigsqcup \text{chunks}(s, h)$
2. If h_1 and h_2 are non-empty positive sub-heaps such that $h = h_1 \uplus h_2$, then $\text{chunks}(s, h) = \text{chunks}(s, h_1) \uplus \text{chunks}(s, h_2)$.

Proof. By induction on $|h|$. If $|h| = 0$, then both claims hold as $\text{chunks}(s, h) = \emptyset$ and h cannot be split into non-empty sub-heaps. If $|h| = n + 1$ and h is atomic, the claims also hold because $\text{chunks}(s, h) = \{h\}$ and h cannot be further decomposed into positive sub-heaps. Otherwise, if h is not atomic, then there exist non-empty positive heaps h_1 and h_2 such that $h = h_1 \uplus h_2$. Since h_1 is non-empty, we have that $|h_2| \leq n$ and vice versa. Therefore we can apply IH to conclude that $h_i = \bigsqcup \text{chunks}(s, h_i)$ for $i = 1, 2$.

Since $h = h_1 \uplus h_2$, we directly have that $h = \bigsqcup (\text{chunks}(s, h_1) \uplus \text{chunks}(s, h_2))$. It remains to show that $\text{chunks}(s, h) = \text{chunks}(s, h_1) \uplus \text{chunks}(s, h_2)$ which directly implies both (1) and (2). Namely, we need to show that (a) there does not exist chunk c of (s, h) such that $c \notin \text{chunks}(s, h_i)$ for $i = 1, 2$ and (b) each element of $\text{chunks}(s, h_i)$ is also member of $\text{chunks}(s, h)$ for $i = 1, 2$.

- (a) By contradiction. Assume that there exists such a chunk c . It holds that c cannot be a pointer-chunk, because then it would be chunk in either (s, h_1) or (s, h_2) . Therefore c is a π -chunk for some inductive predicate $\pi(x, y)$. Since c is non-empty, we can w.l.o.g. assume that the location $s(x)$ is allocated in h_1 . As h_1 is fully decomposed into its chunks by IH, there exists a chunk $c' \in \text{chunks}(s, h_1)$ with root x' such that $s(x) \in \text{dom}(c')$. We proceed by case distinction on types of c and c' :
 - If c' is a-pointer chunk satisfying $x' \mapsto y'$ for some y' which is n , $\langle n : n, p : _ \rangle$, or $\langle n : _, t : t \rangle$, then n or t is a split point of the chunk c which is a contradiction.
 - If c' is $\pi(x', y')$ -chunk for some y' and $\pi \in \{\text{dls}, \text{nls}\}$, then c also needs to be π -chunk. If $s(x) = s(x')$, then either $c = c'$ which is a contradiction, or $c \subseteq c'$ which means that y is a split point of c' , or $c' \subseteq c$ which means that y' is a split point of c —both possibilities lead to a contradiction. If $s(x) \neq s(x')$, then $s(x)$ is a split point of c' which again leads to a contradiction.
 - If both c and c' are sls-chunks we can reach contradiction by the same reasoning as in the previous step.
 - The only remaining case is when c is sls-chunk and $(s, c) \models \text{nls}(x', y', z')$ for some y' and z' . It holds that $c \subset c'$ because otherwise $s(z')$ would be a split point of c or $c = c'$, both leading to a contradiction. From $c \subset c'$ and $c, c' \subseteq h$, we have that $c \notin \text{chunks}(s, h)$ as it is not maximal, which is a contradiction.
- (b) By contradiction. Assume w.l.o.g. that there exists c such that $c \in \text{chunks}(s, h_1)$ and $c \notin \text{chunks}(s, h)$. Consequently, there must exists $c' \subseteq h$ such that $c \subset c'$. It holds that $c' \notin \text{chunks}(s, h_1)$ because otherwise c would not be a chunk of (s, h_1) . As we have already shown in (a), an existence of such c' leads to a contradiction. \square

To prove Theorem 1, we need to show that for arbitrary stack s , \downarrow^s is homomorphism on the separation algebra of all positive heaps (w.r.t. the stack s).

Lemma 7 (\downarrow^s is homomorphism). *Let s be a stack and let h_1 and h_2 be positive heaps. If h_1 and h_2 are disjoint, then the following claims hold:*

1. $\downarrow^s h_1$ and $\downarrow^s h_2$ are disjoint,
2. $\downarrow^s h_1 \uplus \downarrow^s h_2 = \downarrow^s (h_1 \uplus h_2)$.

Proof.

1. Directly from the definition of \downarrow^s , we have that $\text{dom}(\downarrow^s h_i) \subseteq \text{dom}(h_i)$ for $i = 1, 2$. Hence, if $\text{dom}(h_1) \cap \text{dom}(h_2) = \emptyset$, then also $\text{dom}(\downarrow^s h_1) \cap \text{dom}(\downarrow^s h_2) = \emptyset$.
2. Let $\text{chunks}(s, h_1) = \{c_{1,1}, \dots, c_{1,m}\}$ and $\text{chunks}(s, h_2) = \{c_{2,1}, \dots, c_{2,n}\}$.

$$\begin{aligned}
& \downarrow^s h_1 \uplus \downarrow^s h_2 \\
&= \downarrow^s (c_{1,1} \uplus \dots \uplus c_{1,m}) \uplus \downarrow^s (c_{2,1} \uplus \dots \uplus c_{2,n}) \quad [\text{Lemma 6}] \\
&= \downarrow^s c_{1,1} \uplus \dots \uplus \downarrow^s c_{1,m} \uplus \downarrow^s c_{2,1} \uplus \dots \uplus \downarrow^s c_{2,n} \quad [\text{Definition of } \downarrow^s] \\
&= \downarrow^s (c_{1,1} \uplus \dots \uplus c_{1,m} \uplus c_{2,1} \uplus \dots \uplus c_{2,n}) \quad [\text{Definition of } \downarrow^s, \text{ Lemma 6}] \\
&= \downarrow^s (h_1 \uplus h_2) \quad [\text{Lemma 6}]
\end{aligned}$$

Notice that for the third equality to hold, we need Lemma 6 which guarantees that $\text{chunks}(s, h_1 \uplus h_2) = \text{chunks}(s, h_1) \uplus \text{chunks}(s, h_2)$. \square

Corollary 1. *For a positive model (s, h) , it holds that $(s, h) \models \varphi$ iff $(s|_{\text{vars}(\varphi)}, h) \models \varphi$.*

Proof. By straightforward structural induction on φ using the fact that whether it holds that $(s, h) \models \varphi$ does not depend on variables outside of $\text{vars}(\varphi)$.

Proof (Theorem 1). Let $s' = s|_{\text{vars}(\varphi)}$ and $h' = \downarrow^{s'} h$. From Corollary 1 we have that $(s, h) \models \varphi$ iff $(s', h) \models \varphi$. It remains to show that $(s', h) \models \varphi$ iff $(s', h') \models \varphi$. We proceed by structural induction on φ .

- *Base cases.* If φ is a pure formula or a pointer assertion, then h is empty or a singleton set, respectively, and the claim holds as in both cases $h' = h$. If φ is an inductive predicate $\pi(x, y)$, then both (s', h) and (s', h') consist of a single chunk because their only allocated variable is x and they therefore cannot be split into two non-empty positive models (one of splits would need to contain an allocated root different from x). The implication (\Rightarrow) then holds trivially from the definition of the reduction. The implication (\Leftarrow) is proved by rule inversion on the definition of the reduction—if h' is π -chunk, then h was also a π -chunk satisfying the same predicate π .
- *Induction steps.* The cases of boolean connectives follow directly from IH. The case of separating conjunction is more involved:
 - (\Rightarrow) By the assumption, there exist disjoint heaps h_1, h_2 such that $h = h_1 \uplus h_2$ and $(s', h_i) \models \psi_i$ for $i = 1, 2$. From IH, we have that $(s', \downarrow^{s'} h_i) \models \psi_i$ for $i = 1, 2$ and from Lemma 7, we have that $\downarrow^{s'} h_1$ and $\downarrow^{s'} h_2$ are also disjoint. Hence, $(s', \downarrow^{s'} h_1 \uplus \downarrow^{s'} h_2) \models \varphi$. From Lemma 7, we also have that $\downarrow^{s'} h_1 \uplus \downarrow^{s'} h_2 = \downarrow^{s'} (h_1 \uplus h_2)$. Thus, $(s', h') \models \varphi$.

(\Leftarrow) By the assumption, there exist disjoint heaps h_1, h_2 such that $\downarrow^{s'} h = h_1 \uplus h_2$ and $(s', h_i) \models \psi_i$ for $i = 1, 2$. Let $\mathcal{H}_i = \text{chunks}(s', h_i)$ for $i = 1, 2$ and let $\mathcal{H} = \text{chunks}(s', h)$. Since $\downarrow^{s'} h = h_1 \uplus h_2$, for each chunk $c \in \mathcal{H}_1 \cup \mathcal{H}_2$ there exists a unique chunk $c' \in \mathcal{H}$ such that $\downarrow^{s'} c' = c$. Let h'_1 and h'_2 be defined as $h'_i = \uplus\{c' \mid c \in \text{chunks}(s, h_i)\}$. In other words, we compose h'_i from chunks of h that were reduced to chunks of h_i . By construction, it holds that $\downarrow^{s'} h'_i = h_i$ and therefore by IH $(s', h'_i) \models \psi_i$. It also holds that $h'_1 \uplus h'_2 = h \neq \perp$. Thus $(s', h) \models \psi_1 * \psi_2$. \square

We split the proof of Theorem 2, into several lemmas giving bounds on components of canonical models.

Lemma 8 (Bound on the number of allocated locations). *Let φ be a formula and let (s, h) be its canonical model. Then, $|\text{dom}(h)| \leq \lfloor \sum_{x \in \text{vars}(\varphi)} \|x\| \rfloor$.*

Proof. From Lemma 1, we have that h is fully decomposed into its chunks. Since each chunk of a sort S must contain an allocated root variable, the number of S -chunks is bounded by the number of variables of the sort S . We will now discuss the upper bounds on sizes of reduced chunks. A pointer-chunk always has the size 1. The size of a chunk of a sort $S \in \{\mathbb{S}, \mathbb{N}\}$ is 2 by the definition of the reduction. The size of a reduced chunk of a sort \mathbb{D} is 3, but such a chunk needs to contain two allocated variables ($\text{dls}(x, y, x', y')$ -chunk allocates both x and x'). As the result, each variable of the sort \mathbb{D} adds 1.5 to the total bound.

In the worst case, almost all chunks are proper predicate chunks. The only exception is the case when the number of DLL variables is odd. In this case, the last “unpaired” variable can create just a pointer-chunk, not a proper dls-chunk. This is the reason why the sum is rounded down. \square

In the next step, we show that positive models do not have any unlabelled dangling location (a location ℓ is dangling if $\ell \in \text{img}(h)$, but $\ell \notin \text{dom}(h)$).

Lemma 9. *For a positive model (s, h) , it holds that $\text{img}(h) \setminus \text{dom}(h) \subseteq \text{img}(s)$.*

Proof. Since (s, h) is positive, there is a formula φ such that $(s, h) \models \varphi$. The claim is proved by structural induction on φ . The base cases are trivial as the dangling locations of $x \mapsto \langle \{f_i : y_i\}_{i \in I}, \text{sls}(x, y), \text{dls}(x, y, x', y'), \text{nls}(x, y, z) \rangle$ are precisely $\{y_i\}_{i \in I}$, $\{y\}$, $\{y, y'\}$ and $\{y, z\}$, respectively. The induction step follows directly from IH in all the cases, for the case of separating conjunction this is because a composition of two heaps cannot introduce a new dangling location. \square

Proof (Theorem 2). Let φ be a satisfiable formula. From Theorem 1, we know that there exists a canonical model $(s, h) \models \varphi$. Recall that the set of locations in (s, h) is defined as $\text{locs}(s, h) = \text{dom}(h) \cup \text{img}(h) \cup \text{img}(s)$, which can be simplified to $\text{locs}(s, h) = \text{dom}(h) \cup \text{img}(s)$ using Lemma 9.

From Lemma 8, we already have that $|\text{dom}(h)| \leq \lfloor \sum_{x \in \text{vars}(\varphi)} \|x\| \rfloor$. When computing this bound, we have already assigned at least one location to every variable and thus no additional locations are needed for $\text{img}(s)$, except one for nil which may not appear in φ , but we require $s(\text{nil})$ to be present in every stack-heap model. This gives us the bound $|\text{locs}(s, h)| \leq 1 + \lfloor \sum_{x \in \text{vars}(\varphi)} \|x\| \rfloor$. \square

B Translation Correctness

This section presents proofs omitted in Section 5. First, we develop the concept of correspondence of stack-heap and first-order models and its properties in Section B.1 and then we present correctness proofs in Section B.2. For this section, we fix a separation logic formula φ .

B.1 Model Correspondence

We define the translation's signature $\Sigma_\varphi = \{(x)_{x \in \text{vars}(\varphi)}, (h_f)_{f \in \text{Field}}, (D_S)_{S \in \text{Sort}}, D\}$. We call a first-order model over Σ_φ a *model of SMT encoding* (SMT model for short) if $\mathcal{M} \models \mathcal{A}_\varphi$ and *canonical* if for all $\ell \in \mathbb{L} \setminus D^\mathcal{M}$ and all $f \in \text{Field}$, it holds that $h_f[\ell]^\mathcal{M} = \text{loc}^{\text{nil}}$. We call models (s, h) and \mathcal{M} *corresponding* if $(s, h) = \mathbb{T}_\varphi^{-1}(\mathcal{M})$.

Corollary 2 (Unique corresponding models). *For a stack-heap model (s, h) , there is the unique canonical corresponding SMT model \mathcal{M} , and vice versa.*

Proof. For a first-order model \mathcal{M} , the stack-heap model $\mathbb{T}_\varphi^{-1}(\mathcal{M})$ is the unique corresponding model. For a stack-heap model (s, h) , the unique corresponding model is the canonical model \mathcal{M} with $\mathbb{T}_\varphi^{-1}(\mathcal{M})$. \square

We now define an operation of composition of SMT models which mimics the composition of stack-heap models with the same stack and disjoint heaps.

Definition 5 (Compatible models). *Let $F \subseteq \text{Loc}$ be a set of locations. SMT models \mathcal{M}_1 and \mathcal{M}_2 are F -compatible if the following conditions hold:*

- $D^{\mathcal{M}_1} \cap D^{\mathcal{M}_2} \subseteq F$,
- For all $x \in \text{vars}(\varphi)$, $x^{\mathcal{M}_1} = x^{\mathcal{M}_2}$,
- For all $\ell \in F$ and all $f \in \text{Field}$, $h_f[\ell]^{\mathcal{M}_1} = h_f[\ell]^{\mathcal{M}_2}$.

We say that models are compatible if they are \emptyset -compatible. Observe that compatible models correspond to stack-heap models with the same stack and disjoint heaps.

Definition 6 (Model composition). *The composition of \mathcal{M}_1 and \mathcal{M}_2 , $\mathcal{M}_1 \oplus \mathcal{M}_2$, is defined iff \mathcal{M}_1 and \mathcal{M}_2 are compatible as \mathcal{M}_1 except:*

- $D^{\mathcal{M}_1 \oplus \mathcal{M}_2} = D^{\mathcal{M}_1} \cup D^{\mathcal{M}_2}$
- $h_f[\ell]^{\mathcal{M}_1 \oplus \mathcal{M}_2} = \begin{cases} h_f[\ell]^{\mathcal{M}_1} & \text{if } \ell \in D^{\mathcal{M}_1}, \\ h_f[\ell]^{\mathcal{M}_2} & \text{if } \ell \in D^{\mathcal{M}_2}, \\ \text{loc}^{\text{nil}} & \text{otherwise.} \end{cases}$

Corollary 3. *Let (s, h_i) and \mathcal{M}_i be corresponding models. Then $(s, h_1 \uplus h_2)$ corresponds to $\mathcal{M}_1 \oplus \mathcal{M}_2$.*

Proof. First, observe that $\mathcal{M}_1 \oplus \mathcal{M}_2$ is SMT model iff \mathcal{M}_1 and \mathcal{M}_2 are SMT models. The rest of the claim follows directly from the definitions of model correspondence and model composition. \square

The translation $\mathsf{T}(\psi, F)$ has the invariant that it cannot distinguish F -compatible models.

Lemma 10. *For F -compatible SMT models \mathcal{M}_1 and \mathcal{M}_2 , it holds that:*

$$\mathcal{M}_1 \models \mathsf{T}(\psi, F) \iff \mathcal{M}_2 \models \mathsf{T}(\psi, F).$$

Proof. By the definition, F -compatible models \mathcal{M}_1 and \mathcal{M}_2 can differ only in their interpretation of the symbol D and of array images $h_f[\ell]$ such that $\ell \notin F$. It can be easily checked that D does not occur in $\mathsf{T}(\psi, F)$ which therefore does not restrict its interpretation at all. It remains to show that $\mathsf{T}(\psi, F)$ also does not restrict the interpretation of $h_f[\ell]$ such that $\ell \notin F$ which can be done by a straightforward structural induction on φ . \square

The following lemma captures the correctness of translation of the separating conjunction.

Lemma 11. *For compatible SMT models \mathcal{M}_1 and \mathcal{M}_2 , the following are equivalent:*

- $\mathcal{M}_1 \models \mathsf{T}(\psi_1, F_1) \wedge D = F_1$ and $\mathcal{M}_2 \models \mathsf{T}(\psi_2, F_2) \wedge D = F_2$,
- $\mathcal{M}_1 \oplus \mathcal{M}_2 \models \mathsf{T}(\psi_1, F_1) \wedge \mathsf{T}(\psi_2, F_2) \wedge D = F_1 \cup F_2$.

Proof. Since \mathcal{M}_1 and \mathcal{M}_2 are compatible, we know that $\mathcal{M}_1 \oplus \mathcal{M}_2$ is defined. Further, it holds that $\mathcal{M}_1 \oplus \mathcal{M}_2$ is F_i -compatible with \mathcal{M}_i . The rest of the claim follows from the fact that $\mathcal{M}_1 \oplus \mathcal{M}_2 \models \mathsf{T}(\psi_i, F_i)$ iff $\mathcal{M}_i \models \mathsf{T}(\psi_i, F_i)$ by Lemma 10. \square

B.2 Proof of Translation Correctness

Proof (Lemma 3).

- (\Rightarrow) Assume that $(s, h) \models \psi_1 * \psi_2$, then there exist disjoint heaps h_1 and h_2 with $h = h_1 \uplus h_2$ and $(s, h_i) \models \psi_i$. Then $\text{dom}(h_i) \in \mathcal{F}_i$ for $i = 1, 2$ because it is a footprint of ψ_i in (s, h) . The right-hand side then holds for $F_i := \text{dom}(h_i)$.
- (\Leftarrow) Assume that $F_1 \in \mathcal{F}_1$, $F_2 \in \mathcal{F}_2$ are disjoint sets satisfying $(s, h|_{F_i}) \models \psi_i$ and $F_1 \cup F_2 = \text{dom}(h)$, then $(s, h) \models \psi_1 * \psi_2$. \square

We will now define the semantics of inductive predicates in terms of paths in stack-heap graphs. For this we will use a property that such paths are uniquely determine.

Corollary 4 (Path uniqueness). *Let x and y be locations in a model (s, h) . If there is a path $\sigma : x \rightsquigarrow_f y$, then it is uniquely determined as $\sigma = \langle x, h(x, f), \dots, h^{|\sigma|}(x, f) \rangle$.*

Proof. The claim is a direct consequence of the fact that f -successor of a location ℓ in $G[(s, h)]$ is uniquely determined as $h(\ell, f)$; and of paths being defined as acyclic. \square

For reasoning about inductive predicates, we formally define the sorts of locations as follows:

$$\begin{aligned} \text{Loc}_{\text{sls}} &= \{\ell \in \text{Loc} \mid h(\ell, \text{n}) \neq \perp, h(\ell, \text{p}) = h(\ell, \text{t}) = \perp\} \\ \text{Loc}_{\text{dls}} &= \{\ell \in \text{Loc} \mid h(\ell, \text{n}) \neq \perp, h(\ell, \text{p}) \neq \perp, h(\ell, \text{t}) = \perp\} \\ \text{Loc}_{\text{nls}} &= \{\ell \in \text{Loc} \mid h(\ell, \text{n}) \neq \perp, h(\ell, \text{t}) \neq \perp, h(\ell, \text{p}) = \perp\} \end{aligned}$$

For a path σ , a field f and a location z , we further define the *nested domain* of σ in $G[(s, h)]$ as $\text{dom}_N(\sigma, f, z) = \text{dom}(\sigma) \cup \bigcup \{\sigma' \mid \exists \ell \in \text{dom}(\sigma). \sigma' : \ell \rightsquigarrow_f s(z)\}$.

Lemma 12 (Path semantics of inductive predicates). *Let (s, h) be a model and let π be an inductive predicate. The following claims hold:*

- $(s, h) \models \text{sls}(x, y)$ iff there is a path $\sigma : x \rightsquigarrow_n y$ s.t. $\text{dom}(h) = \text{dom}(\sigma) \subseteq \text{Loc}_{\text{sls}}$.
- $(s, h) \models \text{dls}(x, y, x', y')$ iff either $s(x) = s(y)$, $s(x') = s(y')$ and $\text{dom}(h) = \emptyset$, or $s(x) \neq s(y)$, $s(x') \neq s(y')$ and there exists a path $\sigma : x \rightsquigarrow_n y$ such that:
 1. $\text{dom}(h) = \text{dom}(\sigma) \subseteq \text{Loc}_{\text{dls}}$,
 2. $\ell \in \text{dom}(h)$ and $\ell \neq s(x')$ implies that $h(\text{p}, h(\text{n}, \ell)) = \ell$,
 3. $h(s(x'), \text{n}) = s(y)$ and $s(x') \in \text{dom}(h)$,
 4. $h(s(x), \text{p}) = s(y')$ and $s(y') \notin \text{dom}(h)$.
- $(s, h) \models \text{nls}(x, y, z)$ iff there exists a path $\sigma : x \rightsquigarrow_t y$ such that:
 1. $\ell \in \text{dom}(h)$ implies that there exists a path $\sigma' : \ell \rightsquigarrow_n z$,
 2. For all distinct ℓ_1 and ℓ_2 , if $\{\ell_1, \ell_2\} \subseteq \text{dom}_N(\sigma, \text{n}, z)$ and $h(\ell_1, \text{n}) = h(\ell_2, \text{n})$, then $h(\ell_1, \text{n}) \notin \text{dom}_N(\sigma, \text{n}, z)$.
 3. $\text{dom}(h) = \text{dom}_N(\sigma, \text{n}, z)$,
 4. $\text{dom}(\sigma) \subseteq \text{Loc}_{\text{nls}}$ and $\text{dom}(h) \setminus \text{dom}(\sigma) \subseteq \text{Loc}_{\text{sls}}$.

Proof. We only prove the more complex implication (\Leftarrow). For the other direction, it can be intuitively checked that models of inductive predicates satisfy the given conditions and its proof is analogical. We proceed by induction on $|\sigma|$ in all the cases.

SLL. If $|\sigma| = 0$, then the claim holds as $s(x) = s(y)$ and $\text{dom}(h) = \text{dom}(\sigma) = \emptyset$. Assume that $|\sigma| = n + 1$. Since there exists a path $\sigma : s(x) \rightsquigarrow_n s(y)$ of length $n + 1$ with $s(x) \in \text{dom}(\sigma) \subseteq \text{Loc}_{\text{sls}}$, there also exists a path $\sigma' : h(s(x), \text{n}) \rightsquigarrow_n s(y)$. Then, by IH, $(s, h) \models \exists u. x^{\text{S}} \mapsto u * \text{sls}(u, y)$. Moreover, $s(x) \neq s(y)$ because σ is non-empty. Thus $(s, h) \models \text{sls}(x, y)$.

DLL. If $|\sigma| = 0$, then the claim trivially holds. If $|\sigma| = 1$, then $h(s(x), \text{n}) = s(y)$ and $\text{dom}(\sigma) = \{s(x)\} \subseteq \text{Loc}_{\text{dls}}$. Hence, by the conditions (3) and (4), we have that $(s, h) \models x^{\text{D}} \mapsto \langle \text{n} : y, \text{p} : y' \rangle * x \neq y * x \neq y'$ where the first disequality follows from non-emptiness of the path σ and the second follows from the fact that $s(x') \in \text{dom}(\sigma)$ while $s(y') \notin \text{dom}(\sigma)$ by (4). Thus, $(s, h) \models \text{dls}(x, y, x', y')$.

If $|\sigma| = n + 2$, then $\{s(x), h(s(x), \text{n})\} \subseteq \text{dom}(\sigma) \subseteq \text{Loc}_{\text{dls}}$ and there exists a path $\sigma' : h(s(x), \text{n}) \rightsquigarrow_n s(y)$. We will show that for σ' , the conditions (1) – (4) hold. This is trivial for (1), (2) and (4) because $\text{dom}(\sigma') = \text{dom}(\sigma) \setminus \{s(x)\}$. The condition (3) follows from the fact that $s(x) \neq s(x')$ because $h(s(x), \text{n}) \neq h(s(x'), \text{n})$ (the first location is in $\text{dom}(\sigma)$ by assumption while the second is not as it is equal to $s(y)$ by (3)). Hence, by IH, $(s, h) \models \exists u. x^{\text{D}} \mapsto \langle \text{n} : u, \text{p} : y' \rangle * \text{dls}(u, y, x', x) * x \neq y * x' \neq y'$. Thus $(s, h) \models \text{dls}(x, y, x', y')$.

NLL. The base case is analogical to the case of SLL. If $|\sigma| = n + 1$, then $s(x) \in \text{dom}(\sigma)$ and there is a path $\sigma_t : h(s(x), \text{t}) \rightsquigarrow_t s(y)$ with $|\sigma_t| < n$. As $\text{dom}(\sigma_t) \subseteq \text{dom}(\sigma)$ we have that for a sub-heap h_t with $\text{dom}(h_t) = \text{dom}_N(\sigma', \text{n}, z)$, $(s, h_t) \models \exists t. \text{nls}(t, y, z)$ by IH. By (1), we further have that there is a path $\sigma_n : h(s(x), \text{n}) \rightsquigarrow_n s(z)$ for which $\text{dom}(\sigma_n) \subseteq \text{Loc}_{\text{sls}}$ by (4). Hence $(s, h_n) \models \exists n. \text{sls}(n, z)$. By (2), we have that h_t and h_n are disjoint. Combining together all the claims derived so far, we obtain that $(s, h) \models \exists t, n. x^{\text{N}} \mapsto \langle \text{n} : n, \text{t} : t \rangle * \text{sls}(n, z) * \text{nls}(t, y, z) * x \neq y$ where the last disequality follows from non-emptiness of the path σ . Thus, $(s, h) \models \text{nls}(x, y, z)$.

Lemma 13 (Translation of path semantics). *Let (s, h) and \mathcal{M} be corresponding models. The following claim holds:*

1. $\mathcal{M} \models \text{reach}^{[m,n]}(h_f, x, y)$ iff there is a path $\sigma : s(x) \rightsquigarrow_f s(y)$ with $m \leq |\sigma| \leq n$.

Further, let ψ be a formula with one free variable v . Assume that there exists a path $\sigma : s(x) \rightsquigarrow_f s(y)$ with $m \leq |\sigma| \leq n$, then the following claims hold:

2. $\text{path}_S^{[m,n]}(h_f, x, y)^{\mathcal{M}} = \text{dom}(\sigma)$,
3. $\mathcal{M} \models \mathbb{P}_{(h_f, x)}^{\leq n} v. v \in \text{path}_S^{[m,n]}(h_f, x, y) \rightarrow \psi$ iff $\mathcal{M} \models \psi[\ell/v]$ for all $\ell \in \text{dom}(\sigma)$.

Further, assume that for each $\ell \in \text{dom}(\sigma)$ there is $\sigma' : \ell \rightsquigarrow_{f'} s(z)$ s. t. $|\sigma'| \leq k$, then:

4. $\text{path}_N^{[m,n]}(h_f, h_{f'}, x, y, z)^{\mathcal{M}} = \text{dom}_N(\sigma, f', z)$,
5. $\mathcal{M} \models \mathbb{P}_{(h_f, x)}^{\leq n} v'. \mathbb{P}_{(h_{f'}, v')}^{\leq k} v. v \in \text{path}_N^{[m,n]}(h_f, h_{f'}, x, y, z) \rightarrow \psi$ iff $\mathcal{M} \models \psi[\ell/v]$ for all $\ell \in \text{dom}_N(\sigma, f', z)$.

Proof. The claims (1), (2) and (4) directly follow from Corollary 4. The claim (3) follows from the fact that the size of the path $\text{dom}(\sigma)$ is at most n . The proof of (5) is analogical to (3) using the fact that size of each inner path is at most k .

Proof (Lemma 4). We write $\text{FP}_{\mathcal{M}}^{\#}(\varphi)$ instead of $\{F^{\mathcal{M}} \mid F \in \text{FP}^{\#}(\varphi)\}$. We will show that $\text{FP}_{(s, h)}(\varphi) \subseteq \text{FP}_{\mathcal{M}}^{\#}(\varphi)$ writing just $\text{FP}(\varphi) \subseteq \text{FP}^{\#}(\varphi)$ for short.

If $(s, h) \not\models \varphi * \text{true}$, then there does not exist $F \subseteq \text{dom}(h)$ such that $(s, h|_F) \models \varphi$ and the claim trivially holds since $\text{FP}(\varphi) = \emptyset$. Assuming that $(s, h) \models \varphi * \text{true}$, we proceed by structural induction on φ .

- *Base cases.* For all the base cases, it holds that $\text{FP}(\varphi) = \text{FP}^{\#}(\varphi) = \{F\}$ where F is defined as follows. If φ is a pure atom or a pointer assertion $x \mapsto _$, then $F = \{\emptyset\}$ or $F = \{s(x)\}$, respectively. If φ is an inductive predicate, its unique footprint is given by Lemma 12 and its encoding in SMT is correct by Lemma 13.
- *Induction steps.* Assume that φ is a binary connective with operands ψ_1 and ψ_2 . By induction hypothesis, $\text{FP}(\psi_i) \subseteq \text{FP}^{\#}(\psi_i)$ for $i = 1, 2$.
 - If $\varphi \triangleq \psi_1 \wedge \psi_2$, let i denote the operand for which the set $\text{FP}^{\#}(\psi_i)$ has smaller cardinality. Then:

$$\text{FP}(\varphi) \subseteq \text{FP}(\psi_1) \cap \text{FP}(\psi_2) \subseteq \text{FP}(\psi_i) \subseteq \text{FP}^{\#}(\psi_i) = \text{FP}^{\#}(\varphi),$$

where the first inclusion follows from the definition of footprints and the last follows from IH.

- If $\varphi \triangleq \psi_1 \wedge \neg \psi_2$, then the proof is analogical to the previous case with $i = 1$.
- If $\varphi \triangleq \psi_1 \vee \psi_2$, then we have:

$$\text{FP}(\varphi) \subseteq \text{FP}(\psi_1) \cup \text{FP}(\psi_2) \subseteq \text{FP}^{\#}(\psi_1) \cup \text{FP}^{\#}(\psi_2) = \text{FP}^{\#}(\varphi),$$

where the first inclusion follows from the definition of footprints and the second from IH.

- If $\varphi \triangleq \psi_1 * \psi_2$, then we have:

$$\begin{aligned} \text{FP}(\varphi) &\subseteq \{F_1 \cup F_2 \mid F_1 \cap F_2 = \emptyset \text{ and } F_1 \in \text{FP}(\psi_1), F_2 \in \text{FP}(\psi_2)\} \\ &\subseteq \{F_1 \cup F_2 \mid F_1 \in \text{FP}(\psi_1), F_2 \in \text{FP}(\psi_2)\} \\ &\subseteq \{F_1 \cup F_2 \mid F_1 \in \text{FP}^{\#}(\psi_1), F_2 \in \text{FP}^{\#}(\psi_2)\} \\ &= \text{FP}^{\#}(\varphi). \end{aligned}$$

Proof (Theorem 3). We reformulate the original claim to obtain a stronger inductive hypothesis. We prove that for arbitrary SL formula ψ and corresponding models (s, h) and \mathcal{M} , it holds that:

$$\mathcal{M} \models \mathsf{T}(\psi, F) \wedge D = F \iff (s, h) \models \psi.$$

We proceed by structural induction on φ :

- For a pure formula $\psi \triangleq x \bowtie y$ where $\bowtie \in \{=, \neq\}$, we have:

$$\begin{aligned} (s, h) \models \psi &\iff s(x) \bowtie s(y) \wedge \text{dom}(h) = \emptyset && \text{(SL semantics)} \\ &\iff \mathcal{M} \models x \bowtie y \wedge D = \emptyset && \text{(Model correspondence)} \\ &\iff \mathcal{M} \models \mathsf{T}(\psi, F) \wedge D = F && \text{(Formula translation)} \end{aligned}$$

- For a pointer assertion $\psi \triangleq x^S \mapsto \langle f_i : f_i \rangle_{i \in I}$, we have that $(s, h) \models \psi$

$$\begin{aligned} &\iff \text{dom}(h) = \{s(x)\} \wedge h(s(x)) = \langle f_i : f_i \rangle_{i \in I} && \text{(SL semantics)} \\ &\iff \mathcal{M} \models D = \{x\} \wedge x \in D_S \wedge \bigwedge_{i \in I} h_{f_i}[x] = f_i && \text{(Model correspondence)} \\ &\iff \mathcal{M} \models \mathsf{T}(\psi, F) \wedge D = F && \text{(Formula translation)} \end{aligned}$$

Notice that the condition $x \in D_S$ is needed to correctly reconstruct $h(x)$ and in our translation, it is ensured by axioms \mathcal{A}_φ .

- The cases of inductive predicates follow from Lemma 12 and Lemma 13.
- For a boolean connective $\psi \triangleq \psi_1 \bowtie \psi_2$ where $\bowtie \in \{\wedge, \vee, \wedge \neg\}$, we have:

$$\begin{aligned} (s, h) \models \psi &\iff (s, h) \models \psi_1 \bowtie (s, h) \models \psi_2 && \text{(SL semantics)} \\ &\iff \mathcal{M} \models \mathsf{T}(\psi_1, F) \bowtie \mathsf{T}(\psi_2, F) \wedge D = F && \text{(Induction hypothesis)} \\ &\iff \mathcal{M} \models \mathsf{T}(\psi, F) \wedge D = F && \text{(Formula translation)} \end{aligned}$$

- For a separating conjunction $\psi_1 * \psi_2$, we proceed as follows. Let $\mathcal{F}_i = \text{FP}_{\mathcal{M}}^\#(\psi_i)$. From Lemma 3 and Lemma 4, we have that $(s, h) \models \varphi$ iff

$$\bigvee_{F_1 \in \mathcal{F}_1} \bigvee_{F_2 \in \mathcal{F}_2} \bigwedge_{i=1,2} (s, h|_{F_i}) \models \psi_i \wedge F_1 \cap F_2 = \emptyset \wedge F_1 \cup F_2 = \text{dom}(h).$$

Let \mathcal{M}_i be the corresponding model to $(s, h|_{F_i})$ for $i = 1, 2$. By IH:

$$\bigvee_{F_1 \in \mathcal{F}_1} \bigvee_{F_2 \in \mathcal{F}_2} \bigwedge_{i=1,2} (\mathcal{M}_i \models \mathsf{T}(\psi_i, F_i) \wedge D = F_i) \wedge F_1 \cap F_2 = \emptyset \wedge F_1 \cup F_2 = D^{\mathcal{M}}.$$

Using Corollary 3, we have that $\mathcal{M} = \mathcal{M}_1 \oplus \mathcal{M}_2$. Applying Lemma 11, we have an equivalent claim:

$$\bigvee_{F_1 \in \mathcal{F}_1} \bigvee_{F_2 \in \mathcal{F}_2} \mathcal{M} \models \mathsf{T}(\psi_1, F_1) \wedge \mathsf{T}(\psi_2, F_2) \wedge F_1 \cap F_2 = \emptyset \wedge F_1 \cup F_2 = D^{\mathcal{M}}.$$

Finally, this is equivalent to the following claim:

$$\mathcal{M} \models \bigvee_{F_1 \in \mathcal{F}_1} \bigvee_{F_2 \in \mathcal{F}_2} \bigwedge_{i=1,2} \mathsf{T}(\psi_i, F_i) \wedge F_1 \cap F_2 = \emptyset \wedge F_1 \cup F_2 = D,$$

which is the case iff $\mathcal{M} \models \mathsf{T}(\psi, F) \wedge D = F$. \square

C Complexity of BSL Satisfiability

In this section, we give details for the complexity results from Section 5.4. We use $n = \text{bound}(\varphi)$ to denote the location bound of the formula φ and note that n is linear w.r.t. the size of φ .

Proof (Theorem 4). We prove that `SatQuantif` produces formulae of polynomial size, in particular $\mathcal{O}(n^3)$ for SLLs and DLLs, and $\mathcal{O}(n^5)$ for NLLs. We proceed by analysis of asymptotic sizes of auxiliary predicates used for the translation of inductive predicates. Let x be a term of the size $\mathcal{O}(n)$ and let h, h', y, z be terms of the size $\mathcal{O}(1)$. This assumption is justified by the fact that parameters of all predicates in the translation except the roots of considered paths (denoted by x) are variables. The root x can be a term of the form $h_n^k[v]$ or $h_n^k[h_t^\ell[v]]$ (where v is a variable) and is therefore always of linear size because the numbers k and ℓ are never greater than the location bound n . The size of auxiliary predicates is then:

- $|\text{reach}^{\text{=k}}(h, x, y)| \in \mathcal{O}(n)$ when $k \leq n$ because its size is dominated by x .
- $|\text{reach}^{[0,n]}(h, x, y)| \in \mathcal{O}(n^2)$ as it contains n -times $\text{reach}^{\text{=k}}(h, x, y)$ for $k \leq n$.
- $|\text{path}_S^{\text{=k}}(h, x, y)| \in \mathcal{O}(n^2)$ when $k \leq n$ because it is a union of k terms of the form $\{h^i[x]\}$ for $i \leq k \leq n$ and thus of the size $\mathcal{O}(n)$.
- $|\text{path}_S^{[0,n]}(h, x, y)| \in \mathcal{O}(n^3)$ because it consists of $\mathcal{O}(n)$ branches, each of them dominated by the size of $\text{path}_S^{\text{=k}}(h, x, y)$ for $k \leq n$ which is $\mathcal{O}(n^2)$.
- $|\text{path}_N^{\text{=k}}(h, h', x, y, z)| \in \mathcal{O}(n^4)$ when $k \leq n$ because it is a union of k terms of the form $\text{path}_S^{[0,n]}(h', h^i[x], y)$ for $i \leq k \leq n$ and thus of the size $\mathcal{O}(n^3)$.
- $|\text{path}_N^{[0,n]}(h, h', x, y, z)| \in \mathcal{O}(n^5)$ because it consists of $\mathcal{O}(n)$ branches, each of them dominated by the size of $\text{path}_N^{\text{=k}}(h, h', x, y, z)$ for $k \leq n$ which is $\mathcal{O}(n^4)$.

Translation of SLLs and DLLs is dominated by the size of the path_S term and translation of NLLs is dominated by the size the path_N term. Note, that this holds for the translation version using universal quantifiers over locations. When the optimised version with path quantifiers is used for NLLs, we may obtain asymptotically larger formulae, but still polynomial.

For all the other atomic formulae, the size of translated formula is constant. The translation of all binary connective is linear in the size of their operands (also in the case of separating conjunction when the method `SatQuantif` is used). The size of axioms is also linear. Thus, the final size of the translated formula is dominated by the translation of inductive predicates. \square

Proof (Theorem 4). We give the full definition of the reduction from QBF based on [7]. Let $F \triangleq Q_1 x_1 \dots Q_m x_m. F'$ be a ground formula with variables $X = \{x_1, \dots, x_m\}$ (w.l.o.g., they are all distinct), quantifiers $Q_i \in \{\forall, \exists\}$, and propositional body F' . Recall that for a set of variables Y , we define $\text{arbitrary}[Y] \triangleq *_{y \in Y} (y \mapsto \text{nil} \vee \text{emp})$. The reduction of F is defined as $R(F) \wedge *_{x \in X} x \mapsto \text{nil}$ where $R(\cdot)$ is defined as:

$$\begin{aligned}
 R(x) &\triangleq \text{arbitrary}[X] * x \mapsto \text{nil} & R(F \wedge G) &\triangleq R(F) \wedge R(G) \\
 R(\neg x) &\triangleq \text{arbitrary}[X \setminus \{x\}] & R(F \vee G) &\triangleq R(F) \vee R(G) \\
 R(\exists x. F) &\triangleq (x \mapsto \text{nil} \vee \text{emp}) * R(F) & R(\neg F) &\triangleq \text{arbitrary}[X] \wedge \neg R(F) \\
 R(\forall x. F) &\triangleq \text{arbitrary}[X] \wedge \neg((x \mapsto \text{nil} \vee \text{emp}) * R(\neg F))
 \end{aligned}$$

D Bitvector Encoding

In this section, we sketch the bitvector encoding that we have used for the experimental evaluation in Section 7. We assume a subset of the theory of *fixed sized bitvectors*¹⁰ with the signature containing the following symbols: $|$ (bit-wise or), $\&$ (bit-wise and), $\overline{}$ (bit-wise negation), \ll (left shift), and $b[i]$ (extraction of the i -th least significant bit of the bitvector b , interpreted as a boolean value). We further write k_w to denote the unsigned representation of integer k as a bitvector of width w .

Set encoding. The main idea of the bitvector encoding is to represent location sets as bitvectors where the boolean value of i -th least significant bit represents the membership of i -th location in the set. We therefore define the width of considered bitvectors $w = n$, where n is the location bound of the input formula. To reduce the state space, one could consider computing a lower width by excluding locations that cannot be allocated (such as nil), but we have not done this so far. We encode the set operations used in our encoding as follows.

$$\begin{aligned} \emptyset &\triangleq 0_w & \{x_1, \dots, x_m\} &\triangleq 1_w \ll x_1 | \dots | 1_w \ll x_m \\ X_1 \cup X_2 &\triangleq X_1 | X_2 & x \in X &\triangleq X[x] \\ X_1 \cap X_2 &\triangleq X_1 \& X_2 & X_1 \subseteq X_2 \triangleq \overline{X_1} | X_2 = \overline{0_w} \end{aligned}$$

Memory model axioms. We represent locations also as bitvectors of the same width w . For consistency, we need to ensure that each location used in a model is interpreted as a bitvector lesser than n (greater bitvectors would not fit into sets represented as bitvectors which have size $w = n$). We first define $L = \{0_w, 1_w, \dots, (n-1)_w\}$ and axioms defining sets representing location sorts from Section 5.1 as:

$$\begin{aligned} \mathcal{A}_{D_{\mathbb{S}}} &\triangleq D_{\mathbb{S}} = \{0_w\} \cup \{1_w \dots, (n_{\mathbb{S}})_w\}, \\ \mathcal{A}_{D_{\mathbb{D}}} &\triangleq D_{\mathbb{D}} = \{0_w\} \cup \{(n_{\mathbb{S}} + 1)_w, \dots, (n_{\mathbb{S}} + n_{\mathbb{D}})_w\}, \\ \mathcal{A}_{D_{\mathbb{N}}} &\triangleq D_{\mathbb{N}} = \{0_w\} \cup \{(n_{\mathbb{S}} + n_{\mathbb{D}} + 1)_w, \dots, (n_{\mathbb{S}} + n_{\mathbb{D}} + n_{\mathbb{N}})_w\}. \end{aligned}$$

Notice that those axioms will ensure that all locations referenced by some variable are interpreted as bitvectors lesser than n (since $n = n_{\mathbb{S}} + n_{\mathbb{D}} + n_{\mathbb{N}} + 1$). To ensure the same for anonymous locations, we add the following axiom to restrict all heap images to be lesser than n :

$$\mathcal{A}_{\text{heap}} \triangleq \bigwedge_{f \in \text{Field}} \bigwedge_{\ell \in L} h_f[\ell] < n_w.$$

Putting it all together, the axioms in the bitvector encoding are defined as:

$$\mathcal{A}_{\varphi} \triangleq \text{nil} = 0_w \wedge \text{nil} \notin D \wedge \mathcal{A}_{\text{heap}} \wedge \bigwedge_{S \in \text{Sort}} (\mathcal{A}_{D_S} \wedge \bigwedge_{x \in \text{vars}_S(\varphi)} x \in D_S).$$

¹⁰ <https://smtlib.cs.uiowa.edu/theories-FixedSizeBitVectors.shtml>