

Satisfiability Modulo Heap-Based Programs

Quang Loc Le ¹

Jun Sun ¹

Wei-Ngan Chin ²

(1) Singapore University of Technology and Design (2) National University of Singapore

Abstract. In this work, we present a semi-decision procedure for a fragment of separation logic with user-defined predicates and Presburger arithmetic. To check the satisfiability of a formula, our procedure iteratively unfolds the formula and examines the derived disjuncts. In each iteration, it searches for a proof of either satisfiability or unsatisfiability. Our procedure is further enhanced with automatically inferred invariants as well as detection of cyclic proof. We also identify a syntactically restricted fragment of the logic for which our procedure is terminating and thus complete. This decidable fragment is relatively expressive as it can capture a range of sophisticated data structures with non-trivial pure properties, such as size, sortedness and near-balanced. We have implemented the proposed solver and a new system for verifying heap-based programs. We have evaluated our system on benchmark programs from a software verification competition.

Keywords: Decision Procedures · Satisfiability · Separation Logic · Inductive Predicates · Cyclic Proofs

1 Introduction

Satisfiability solvers, particularly those based on Satisfiability Modulo Theory (SMT) technology [3,19], have made tremendous practical advances in the last decade to the point where they are now widely used in tools for applications as diverse as bug finding [27], program analyses [4] to automated verification [2]. However, current SMT solvers are based primarily on first-order logic, and do not yet cater to the needs of resource-oriented logics, such as separation logic [26,40]. Separation logic has recently established a solid reputation for reasoning about programs that manipulate heap-based data structures. One of its strengths is the ability to concisely and locally describe program states that hold in separate regions of heap memory. In particular, a spatial conjunction (i.e., $\kappa_1 * \kappa_2$) asserts that a given heap can be decomposed into two disjoint regions and the formulas, κ_1 and κ_2 , hold respectively and separately in the two memory regions. In this work, we investigate the problem of verifying heap-manipulating programs in the framework of SMT. We reduce this problem to solving verification conditions representing precise program semantics [44,10,9,22].

Developing an SMT solver supporting separation logic with inductive predicates and Presburger arithmetic is challenging as the satisfiability problem for this fragment is undecidable [31,30]. We focus on an expressive fragment which consists of spatial predicates expressing empty heap assertion (`emp`), points-to assertion ($x \mapsto c(\bar{v})$), and inductive predicate assertions ($P(\bar{v})$). Moreover, it may include pure constraints on data values and capture desired properties of structural heaps (such as size, height, sortedness and even near-balanced tree properties). We thus face the challenge of handling

recursive predicates with pure properties, that are inherently infinite. Furthermore, we would like to support both satisfiability (SAT) and unsatisfiability (UNSAT) checks.

There have been a number of preliminary attempts in this direction. For instance, early proposals *fixed* the set of shape predicates that may be used, for example, to linked lists (in SeLogger [17,23], and SLLB [34]) or trees (GRIT [36]). There are few approaches supporting user-defined predicates [14,39,25]. Brotherston *et. al.* recently made an important contribution by introducing SLSAT, a decision procedure for a fragment of separation logic with arbitrary *shape-only* inductive predicates [12]. However, SLSAT is limited to the shape domain, whereas shape predicates extended with pure properties are often required for automated verification of functional correctness.

In this paper, we start by proposing a new procedure, called S2SAT, which combines under-approximation and over-approximation for simultaneously checking SAT and UNSAT properties for a sound and complete theory augmented with inductive predicates. S2SAT takes a set of user-defined predicates and a logic formula as inputs. It iteratively constructs *an unfolding tree* by unfolding the formula in a breadth-first, flow- and context-sensitive manner until either a symbolic model, or a proof of unsatisfiability or a fixpoint (e.g., a cyclic proof) is identified. In each iteration, it searches over the leaves of the tree (the disjunction of which is equivalent to the input formula) to check whether there is a satisfiable leaf (which proves satisfiability) or whether all leaves are unsatisfiable. In particular, to prove SAT, it considers *base disjuncts* which are derived from base-case branches of the predicates. These disjuncts are under-approximations of the input formula and critical for satisfiability. Disjuncts which have no inductive predicates are precisely decided. To prove UNSAT, S2SAT over-approximates the leaves prior to prove UNSAT. Our procedure systematically enumerates all disjuncts derived from a given inductive formula, so it is terminating for SAT. However, it may not be terminating for UNSAT with those undecidable augmented logic. To facilitate termination, we propose an approach for fixpoint computation. This fixpoint computation is useful for domains with finite model semantics i.e., collecting semantics for a given formula of such domains is finite. In other words, the input formula is unsatisfiable when the unfolding goes on forever without uncovering any models. We have implemented one instantiation of the fixpoint detection for inductive proving based on cyclic proof [13] s.t. the soundness of the cyclic proof guarantees the well-foundedness of all reasoning.

To explicitly handle heap-manipulating programs, we propose a separation logic instantiation of S2SAT, called S2SAT_{SL}. Our base theory is a combination of the aforementioned separation logic predicates except inductive predicates. We show that our decision procedure for this base theory is sound and complete. S2SAT_{SL} over-approximates formulas with soundly inferred predicate invariants. In addition, we describe some syntax restrictions such that S2SAT_{SL} is always able to construct a cyclic proof for a restricted formula so that our procedure is terminating and complete.

To summarize, we make the following technical contributions in this work.

- We introduce cyclic proof into a satisfiability procedure for a base theory augmented with inductive predicates (refer to Sec. 3).
- We propose a satisfiability procedure for separation logic with user-defined predicates and Presburger arithmetic (Sec. 4).

```

    struct node {int val; node next;}
1 int main(int n){
2   if(n<0) return 0;
3   node x=ll(n);
4   int r=test(x);
5   if(!r) ERROR();
6   return 1;
7 }
8 node ll(int i){
9   if(i==0) return null;
10  else return new node(i,ll(i-1)); }
11 int test(node p){
12  if(p==null) return 1;
13  else {if(p->val<0) return 0;
14        else return test(p->next);}}

```

Fig. 1. Motivating Example.

- We prove that $S2SAT_{SL}$ is: (i) sound for SAT and UNSAT; (ii) and terminating (i.e., proposing a new decision procedure) for restricted fragments (Sec. 5).
- We present a mechanism to automatically derive sound (over-approximated) invariants for user-defined predicates (Sec. 6).
- We have implemented the satisfiability solver $S2SAT_{SL}$ and the new verification system, called $S2_{td}$. We evaluated $S2SAT_{SL}$ and $S2_{td}$ with benchmarks from recent competitions. The experimental results show that our system is expressive, robust and efficient (Sec. 7).

Proofs of Lemmas and Theorems presented in this paper are available in the companion technical report [30].

2 Illustrative Example

We illustrate how our approach works with the example shown in Fig. 1. Our verification system proves that this program is memory safe and function `ERROR()` (line 5) is never called. Our system uses symbolic execution in [6,14] and large-block encoding [8] to provide a semantic encoding of verification conditions. For safety, one of the generated verification conditions is: $\Delta_0 \equiv ll(n,x)_0^0 * test(x,r_1)_0^1 \wedge n \geq 0 \wedge r_1 = 0$. If Δ_0 is unsatisfiable, function `ERROR()` is never called. In Δ_0 , `ll` and `test` are Interprocedural Control Flow Graph (ICFG) of the functions `ll` and `test`. Our system eludes these ICFGs as inductive predicates. For each predicate, a parameter `res` is appended at the end to model the return value of the function; for instance, the variables x (in `ll`) and r_1 (in `test`) of Δ_0 are the actual parameters corresponding to `res`. Each inductive predicate instance is also labeled with a subscript for the unfolding number and a superscript for the sequence number, which are used to control the unfolding in a breadth-first and flow-sensitive manner.

To discharge Δ_0 , $S2SAT_{SL}$ iteratively derives a series of unfolding trees \mathcal{T}_i . An unfolding tree is a tree such that each node is labeled with an unfolded disjunct, corresponding to a path condition in the program. We say that a leaf of \mathcal{T}_i is closed if it is unsatisfiable; otherwise it is open. During each iteration, $S2SAT_{SL}$ either proves SAT by identifying a satisfiable leaf of \mathcal{T}_i which contains no user-defined predicate instances or proves UNSAT by showing that an over-approximation of all leaves is unsatisfiable. Initially, \mathcal{T}_0 contains only one node Δ_0 . As Δ_0 contains inductive predicates, it is not considered for proving SAT. $S2SAT_{SL}$ then over-approximates Δ_0 to a first-order logic

formula by substituting each predicate instance with its corresponding sound invariants in order to prove UNSAT. We assume that `ll` (resp. `test`) is annotated with invariant $i \geq 0$ (resp. $0 \leq \text{res} \leq 1$). Hence, the over-approximation of Δ_0 is computed as: $\pi_0 \equiv n \geq 0 \wedge 0 \leq r_1 \leq 1 \wedge n \geq 0 \wedge r_1 = 0 \wedge n = 0 \wedge x = \text{null}$. Formula π_0 is then passed to an SMT solver, such as Z3 [19], for unsatisfiability checking. As expected, π_0 is not unsatisfiable.

Next, S2SAT_{SL} selects an open leaf for unfolding to derive \mathcal{T}_1 . A leaf is selected in a breadth-first manner; furthermore a predicate instance of the selected leaf is selected for unfolding if its sequence number is the smallest. With Δ_0 , the `ll` instance is selected. As so, \mathcal{T}_1 has two open leaves corresponding to two derived disjuncts:

$$\begin{aligned} \Delta_{11} &\equiv \text{test}(x, r_1)_0^1 \wedge n \geq 0 \wedge r_1 = 0 \wedge n = 0 \wedge x = \text{null} \\ \Delta_{12} &\equiv x \mapsto \text{node}(n, r_2) * ll(n_1, r_2)_1^0 * \text{test}(x, r_1)_0^1 \wedge n \geq 0 \wedge r_1 = 0 \wedge n \neq 0 \wedge n_1 = n - 1 \end{aligned}$$

Since Δ_{11} and Δ_{12} include predicate instances, they are not considered for SAT. To prove UNSAT, S2SAT_{SL} computes their over-approximated invariants:

$$\begin{aligned} \pi_{11} &\equiv 0 \leq r_1 \leq 1 \wedge n \geq 0 \wedge r_1 = 0 \wedge n = 0 \wedge x = \text{null} \\ \pi_{12} &\equiv x \neq \text{null} \wedge n_1 \geq 0 \wedge 0 \leq r_1 \leq 1 \wedge n \geq 0 \wedge r_1 = 0 \wedge n \neq 0 \wedge n_1 = n - 1 \end{aligned}$$

As neither π_{11} nor π_{12} is unsatisfiable, S2SAT_{SL} selects `test` of Δ_{11} for unfolding to construct \mathcal{T}_2 . For efficiency, unfolding is performed in a context-sensitive manner. A branch is infeasible (and pruned in advance) if its invariant is inconsistent with the (over-approximated) context. For instance, the invariant of the `then` branch at line 12 of `test` is $\text{inv}_{\text{test}_o} \equiv p = \text{null} \wedge \text{res} = 1$. As $\text{inv}_{\text{test}_o}$ (after proper renaming) is inconsistent with π_{11} , this branch is infeasible. Similarly, both `else` branches of `test` are infeasible. For \mathcal{T}_3 , the remaining leaf Δ_{12} is selected for unfolding. As the `test`'s unfolding number is smaller than `ll`'s, `test` is selected. After the `then` branch is identified as infeasible and pruned, \mathcal{T}_3 is left with two open leaves as shown in Fig. 2, where infeasible leaves are dotted-lined. Δ_{32} and Δ_{33} are as below.




Fig. 2. Unfolding Tree \mathcal{T}_3 .

$$\begin{aligned} \Delta_{32} &\equiv x \mapsto \text{node}(n, r_2) * ll(n_1, r_2)_1^0 \wedge n \geq 0 \wedge r_1 = 0 \wedge n \neq 0 \wedge n_1 = n - 1 \wedge x \neq \text{null} \wedge n < 0 \wedge r_1 = 0 \\ \Delta_{33} &\equiv x \mapsto \text{node}(n, r_2) * ll(n_1, r_2)_1^0 * \text{test}(r_2, r_1)_1^1 \wedge n \geq 0 \wedge r_1 = 0 \wedge n \neq 0 \wedge n_1 = n - 1 \\ &\quad \wedge x \neq \text{null} \wedge n \geq 0 \end{aligned}$$

As Δ_{32} and Δ_{33} include inductive predicate instances, SAT checking is not applicable. For UNSAT checking, S2SAT_{SL} proves that Δ_{32} is unsatisfiable (its unsatisfiable cores are underlined as above); and shows that Δ_{33} can be linked back to Δ_0 (i.e., subsumed by Δ_0). The latter is shown based on some weakening and substitution principles (see Sec. 4.2). In particular: (i) Substituting Δ_{33} with $\theta = [n_2/n, x_1/x, n/n_1, x/r_2]$ such that predicate instances in the substituted formula, i.e., Δ_{33_a} , and Δ_0 are identical; as such, Δ_{33_a} is computed as below.

$$\begin{aligned} \Delta_{33_a} &\equiv x_1 \mapsto \text{node}(n_2, x) * ll(n, x)_1^0 * \text{test}(x, r_1)_1^1 \wedge n_2 \geq 0 \wedge r_1 = 0 \wedge n_2 \neq 0 \wedge n = n_2 - 1 \\ &\quad \wedge x_1 \neq \text{null} \wedge n_2 \geq 0 \end{aligned}$$

Algorithm 1: S2SAT Procedure.

```
input :  $\lambda^{ind}$ 
output: SAT or UNSAT
1  $i \leftarrow 0$ ;  $\mathcal{T}_0 \leftarrow \{\lambda^{ind}\}$ ; /* initialize */
2 while true do
3    $(is\_sat, \mathcal{T}_i) \leftarrow UA\_test(\mathcal{T}_i)$ ; /* check SAT */
4   if is_sat then return SAT; /* SAT */
5   else
6      $\mathcal{T}_i \leftarrow OA\_test(\mathcal{T}_i)$ ; /* prune UNSAT */
7      $\mathcal{T}_i \leftarrow link\_back(\mathcal{T}_i)$ ; /* detect fixpoint */
8     if is_closed( $\mathcal{T}_i$ ) then return UNSAT; /* UNSAT */
9     else
10       $\lambda^{ind}_i \leftarrow choose\_bfs(\mathcal{T}_i)$ ; /* choose an open leaf */
11       $i \leftarrow i + 1$ ;
12       $\mathcal{T}_i \leftarrow unfold(\lambda^{ind}_i)$ ;
13    end
14  end
15 end
```

(ii) subtracting identical inductive predicates between Δ_{33_a} and Δ_0 ; (iii) weakening the remainder of Δ_{33_a} (i.e., $x_1 \mapsto node(n_2, x)$ is eliminated); (iv) checking validity of the implication between pure of the remainder of Δ_{33_a} with the pure part of the remainder of Δ_0 , i.e., $n_2 \geq 0 \wedge r_1 = 0 \wedge n_2 \neq 0 \wedge n = n_2 - 1 \wedge x_1 \neq null \wedge n_2 \geq 0 \implies n \geq 0 \wedge r_1 = 0$. The back-link between Δ_{33} and Δ_0 establishes a cyclic proof which then proves Δ_0 is unsatisfiable.

3 S2SAT Algorithm

In this section, we present S2SAT, a procedure for checking satisfiability of formula with inductive predicates. We start by defining our target formulas. Let \mathcal{L} be a *base theory* (logic) with the following properties: (i) \mathcal{L} is closed under propositional combination and supports boolean variables; (ii) there exists a complete decision procedure for \mathcal{L} . Let \mathcal{L}^{ind} be the extension of \mathcal{L} with inductive predicate instances defined in a system with a set of predicates $\mathcal{P} = \{P_1, \dots, P_k\}$. Each predicate may be annotated with a *sound* invariant. We use λ to denote a formula in \mathcal{L} and λ^{ind} to denote a formula in the extended theory. Semantically, $\lambda^{ind} \equiv \bigvee_{i=0}^n \lambda_i$, $n \geq 0$.

S2SAT is presented in Algorithm 1. S2SAT takes a formula λ^{ind} as input, systematically enumerates disjuncts λ_i and can produce two possible outcomes if it terminates: SAT with a satisfiable formula λ_i or UNSAT with a proof. We remark that non-termination is classified as UNKNOWN.

S2SAT maintains a set of open leaves of the unfolding tree \mathcal{T}_i that is derived from λ^{ind} . In each iteration, S2SAT selects and unfolds an open leaf so as either to include more reachable base formulas (with the hope to produce a SAT answer), or to refine inductive formulas (with the hope to produce an UNSAT answer). Specially, in each

iteration, S2SAT checks whether the formula is SAT at line 3; whether it is UNSAT at line 6; whether a fixpoint can be established at line 7. Function `UA_test` searches for a satisfiable *base* disjunct (i.e., `is_sat` is set to true). Simultaneously, it marks all unsatisfiable base disjuncts *closed*. Next, function `OA_test` uses predicate invariants to over-approximate open leaves of \mathcal{T}_i , and marks those with an *unsatisfiable* over-approximation closed. After that, function `link_back` attempts to link remaining open leaves back to interior nodes so as to form a fixpoint (i.e., a (partial) pre-proof for induction proving). The leaves which have been linked back are also marked as closed. Whenever all leaves are closed, S2SAT decides λ^{ind} as UNSAT (line 8). Otherwise, the `choose_bfs` (line 10) chooses an *open* leaf in breadth-first manner for unfolding.

Procedure `link_back` takes the unfolding tree \mathcal{T}_i as input and checks whether each open leaf $\lambda^{ind\text{bud}} \in \mathcal{T}_i$ matches with one interior node $\lambda^{ind\text{comp}}$ in \mathcal{T}_i via a *matching function* \mathbf{f}_{fix} . \mathbf{f}_{fix} is based on weakening and substitution principles [13]. Intuitively, \mathbf{f}_{fix} detects the case of (i) the unfolding goes forever if we keep unfolding $\lambda^{ind\text{bud}}$; and (ii) $\lambda^{ind\text{bud}}$ has no model when $\lambda^{ind\text{comp}}$ has no model. If $\mathbf{f}_{\text{fix}}(\lambda^{ind\text{bud}}, \Delta^{comp}) = \text{true}$, Δ^{bud} is marked closed.

Our procedure systematically enumerates all disjuncts derived from a given inductive formula, so it is terminating for SAT. However, it may not be terminating for UNSAT with those undecidable augmented logic. In the next paragraph, we discuss the soundness of the algorithm.

Soundness When S2SAT terminates, there are the following three cases.

- (case A) S2SAT produces SAT with a base satisfiable λ^{ind}_i ;
- (case B) S2SAT produces UNSAT with a proof that all leaves of \mathcal{T}_i are unsatisfiable;
- (case C) S2SAT produces UNSAT with a fixpoint: a proof that some leaves of \mathcal{T}_i are unsatisfiable and the remaining leaves are linked back.

Under the assumption that \mathcal{L} is both sound and complete, case A can be shown to be sound straightforwardly. Soundness of case B immediately follows the soundness of `OA_test`. In the following, we describe the cyclic proof instantiation of `link_back` for fixpoint detection and prove the soundness of case C.

We use CYCLIC to denote the cyclic proof for entailment procedure adapted from [13]. The following definitions are adapted from their analogues of CYCLIC.

Definition 1 (Pre-proof) A pre-proof derived for a formula λ^{ind} is a pair $(\mathcal{T}_i, \mathcal{L})$ where \mathcal{T}_i is an unfolding tree whose root labelled by λ^{ind} and \mathcal{L} is a back-link function assigning every open leaf λ^{ind}_l of \mathcal{T}_i to an interior node $\lambda^{ind}_c = \mathcal{L}(\lambda^{ind}_l)$ such that there exists some substitution θ i.e., $\lambda^{ind}_c = \lambda^{ind}_l[\theta]$. λ^{ind}_l is referred as a bud and λ^{ind}_c is referred as its companion.

A path in a pre-proof is a sequence of nodes $(\lambda^{ind}_i)_{i \geq 0}$.

Definition 2 (Trace) Let $(\lambda^{ind}_i)_{i \geq 0}$ be a path in a pre-proof \mathcal{PP} . A trace following $(\lambda^{ind}_i)_{i \geq 0}$ is a sequence $(\alpha_i)_{i \geq 0}$ such that, for all $i \geq 0$, α_i is a predicate instance $\mathcal{P}(\bar{t})$ in the formula λ^{ind}_i , and either:

1. α_{i+1} is the subformula according to $\mathcal{P}(\bar{t})$ occurrence in λ^{ind}_{i+1} , or

2. $\lambda^{ind}_i[\bar{t}/\bar{v}]$ where λ^{ind}_i is branches of inductive predicate $P(\bar{v})$. i is a progressing point of the trace.

To ensure that pre-proofs correspond to sound proofs, a global *soundness condition* must be imposed on such pre-proofs as follows.

Definition 3 (Cyclic proof) *A pre-proof is a cyclic proof if, for every infinite path $(\lambda^{ind}_i)_{i \geq 0}$, there is a tail of the path $p = (\lambda^{ind}_i)_{i \geq n}$ such that there is an infinitely progressing trace following p .*

Theorem 1 (Soundness). *If there is a cyclic proof of λ^{ind}_0 , λ^{ind}_0 is UNSAT.*

Proof We reduce our cyclic proof problem for satisfiability to the cyclic proof problem for entailment check, i.e., $\lambda^{ind}_0 \vdash \text{false}$ of CYCLIC. Assume there is a cyclic proof \mathcal{PP} of λ^{ind}_0 . From \mathcal{PP} we construct the pre-proof \mathcal{PP}_\vdash for the sequent $\lambda^{ind}_0 \vdash \text{false}$ as follows. For each node $(\lambda^{ind}_i)_{i \geq 0}$ in \mathcal{PP} , we replace the formula λ^{ind}_i by the sequent $\lambda^{ind}_i \vdash \text{false}$. Since \mathcal{PP} is a cyclic proof, it follows that for every infinite path $(\lambda^{ind}_i)_{i \geq 0}$, there is a tail of the path, $p = (\lambda^{ind}_i)_{i \geq n}$, such that there is an infinitely progressing trace following p (Definition 3). Since formulas in [13] are only traced through the LHS of the sequent and not its RHS, it is implied that for every infinite path $(\lambda^{ind}_i \vdash \text{false})_{i \geq 0}$, there is a tail of the path, $p = (\lambda^{ind}_i \vdash \text{false})_{i \geq n}$, such that there is an infinitely progressing trace following p . Thus, \mathcal{PP}_\vdash is a cyclic proof (Definition 3 of [13]). As such $\lambda^{ind}_0 \models \text{false}$ (Theorem 6 of [13]). In other words, λ^{ind}_0 is UNSAT. \square

To sum up, to implement a sound cyclic proof system besides the matching function, a global *soundness condition* must be established on pre-proofs to guarantee well-foundedness of all reasoning.

4 Separation Logic Instantiation of S2SAT

In this section, to explicitly handle heap-manipulating programs, we propose a separation logic instantiation of S2SAT, called S2SAT_{SL}. We start by presenting SLPA, a fragment of separation logic with inductive predicates and arithmetic.

4.1 A Fragment of Separation Logic

Syntax The syntax of SLPA formulas is presented in Fig. 3. We use \bar{x} to denote a sequence (e.g., \bar{v} for sequence of variables), and x_i to denote the i^{th} element. Whenever possible, we discard f_i of the points-to predicate and use its short form as $x \mapsto c(v_i)$. Note that $v_1 \neq v_2$ and $v \neq \text{null}$ are short forms for $\neg(v_1 = v_2)$ and $\neg(v = \text{null})$, respectively. All free variables are implicitly universally quantified at the outermost level. To express different scenarios for shape predicates, the fragment supports disjunction Φ over formulas. Each predicate instance is of the form $P(\bar{v})_u^o$ where o and u are labels used for context- and flow- sensitive unfolding. In particular, o captures the sequence number and u is the number of unfolding. For simplicity, we occasionally omit these two numbers if there is no ambiguity. A formula Δ is a *base formula* if it does not have any user-defined predicate instances. Otherwise, Δ is an *inductive formula*.

| | | |
|-------------------|--|---|
| Formula | $\Phi ::= \Delta \mid \Phi_1 \vee \Phi_2$ | $\Delta ::= \exists \bar{v}. (\kappa \wedge \pi)$ |
| Spatial formula | $\kappa ::= \text{emp} \mid x \mapsto c(f_i: v_i) \mid \mathbf{P}(\bar{v})_u^o \mid \kappa_1 * \kappa_2$ | |
| Pure formula | $\pi ::= \pi_1 \wedge \pi_2 \mid \mathbf{b} \mid \alpha \mid \phi$ | |
| Ptr (Dis)Equality | $\alpha ::= v_1 = v_2 \mid v = \text{null} \mid v_1 \neq v_2 \mid v \neq \text{null} \mid \alpha_1 \wedge \alpha_2$ | |
| Presburger arith. | $\phi ::= i \mid \exists v. \phi \mid \neg \phi \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2$ | |
| Boolean formula | $b ::= \text{true} \mid \text{false} \mid v \mid b_1 = b_2$ | |
| Linear arithmetic | $i ::= a_1 = a_2 \mid a_1 \leq a_2$ | |
| | $a ::= k^{\text{int}} \mid v \mid k^{\text{int}} \times a \mid a_1 + a_2 \mid -a \mid \max(a_1, a_2) \mid \min(a_1, a_2)$ | |
| | $\mathbf{P} \in \mathcal{P} \quad c \in \text{Node} \quad f_i \in \text{Fields} \quad v, v_i, x, y, \text{res}, \text{res}' \in \text{Var} \quad \bar{v} \equiv v_1, \dots, v_n$ | |

Fig. 3. Syntax.

User-Defined Predicate A user-defined predicate \mathbf{P} is of the following general form

$$\text{pred } \mathbf{P}(\bar{t}) \equiv \bigvee_{i=1}^n (\exists \bar{w}_i. \Delta_i \mid \pi_i^b) \quad \overline{inv}: \pi;$$

where \mathbf{P} is predicate name; \bar{t} is a set of formal parameters; and $\exists \bar{w}_i. \Delta_i$ ($i \in 1 \dots n$) is a branch. Each branch is optionally annotated with a sound invariant π_i^b which is a pure formula that over-approximates the branch. π is an optionally sound *predicate invariant*. It must be a superset of all possible models of the predicate \mathbf{P} via a pure constraint on stack. The default invariant of each inductive predicate is `true`. For efficiency, we infer more precise invariants automatically (See Sec.6). Inductive branches may be recursive. We assume that the recursion is direct, i.e., a recursive branch of predicate \mathbf{P} includes at least one predicate instance \mathbf{P} . In each branch, we require that variables which are not formal parameters must be existentially quantified i.e., $\forall i \in 1 \dots n. FV(\Delta_i) = \bar{t}$ and $\bar{w}_i \cap \bar{t} = \emptyset$ where $FV(\Delta)$ are all free variables in the formula Δ .

In the following, we apply SLPA to model two data structures: sorted lists (`sortll`) without an annotated invariant and AVL trees (`avl`) with annotated-invariant.

$$\begin{aligned} \text{pred } \text{sortll}(\text{root}, n, m) &\equiv \text{root} \mapsto \text{node}(m, \text{null}) \wedge n = 1 \\ &\vee \exists q, n_1, m_1. \text{root} \mapsto \text{node}(m, q) * \text{sortll}(q, n_1, m_1) \wedge n = n_1 + 1 \wedge m \leq m_1 \end{aligned}$$

`struct c2 { c2 left; c2 right; } // data structure declaration`

$$\begin{aligned} \text{pred } \text{avl}(\text{root}, n, h) &\equiv \text{emp} \wedge \text{root} = \text{null} \wedge n = 0 \wedge h = 0 \mid \text{root} = \text{null} \wedge n = 0 \wedge h = 0 \\ &\vee \exists l, r, n_1, n_2, h_1, h_2. \text{root} \mapsto c_2(l, r) * \text{avl}(l, n_1, h_1) * \text{avl}(r, n_2, h_2) \wedge \\ &\quad n = n_1 + n_2 + 1 \wedge h = \max(h_1, h_2) + 1 \wedge -1 \leq h_1 - h_2 \leq 1 \mid \text{root} \neq \text{null} \wedge n > 0 \wedge h > 0 \\ &\overline{inv}: n \geq 0 \wedge h \geq 0 \end{aligned}$$

Semantics In the following, we discuss the semantics of SLPA. Concrete heap models assume a fixed finite collection *Node*, a fixed finite collection *Fields*, a disjoint set *Loc* of locations (heap addresses), a set of non-address values *Val*, such that `null` \in *Val* and $\text{Val} \cap \text{Loc} = \emptyset$. Further, we define:

$$\begin{aligned} \text{Heaps} &\stackrel{\text{def}}{=} \text{Loc} \rightarrow_{\text{fin}} (\text{Node} \rightarrow \text{Fields} \rightarrow \text{Val} \cup \text{Loc}) \\ \text{Stacks} &\stackrel{\text{def}}{=} \text{Var} \rightarrow \text{Val} \cup \text{Loc} \end{aligned}$$

| | |
|---|---|
| $s, h \models \text{emp}$ | iff $h = \emptyset$ |
| $s, h \models v \mapsto c(f_i : v_i)$ | iff $l = s(v)$, $\text{dom}(h) = \{l \rightarrow r\}$ and $r(c, f_i) = s(v_i)$ |
| $s, h \models \mathbf{p}(\bar{v})$ | iff $(s(\bar{v}), h) \in \llbracket \mathbf{p}(\bar{v}) \rrbracket$ |
| $s, h \models \kappa_1 * \kappa_2$ | iff $\exists h_1, h_2. h_1 \# h_2$ and $h = h_1 \cdot h_2$ and $s, h_1 \models \kappa_1$ and $s, h_2 \models \kappa_2$ |
| $s, h \models \text{true}$ | iff always |
| $s, h \models \text{false}$ | iff never |
| $s, h \models \exists v_1, \dots, v_n. (\kappa \wedge \pi)$ | iff $\exists \alpha_1 \dots \alpha_n. s(v_1 \mapsto \alpha_1 * \dots * v_n \mapsto \alpha_n), h \models \kappa$ and $s(v_1 \mapsto \alpha_1 * \dots * v_n \mapsto \alpha_n) \models \pi$ |
| $s, h \models \Phi_1 \vee \Phi_2$ | iff $s, h \models \Phi_1$ or $s, h \models \Phi_2$ |

Fig. 4. Semantics.

The semantics is given by a forcing relation: $s, h \models \Phi$ that forces the stack s and heap h to satisfy the constraint Φ where $h \in \text{Heaps}$, $s \in \text{Stacks}$, and Φ is a formula.

The semantics is presented in Fig. 4. $\text{dom}(f)$ is the domain of function f ; $h_1 \# h_2$ denotes that heaps h_1 and h_2 are disjoint, i.e., $\text{dom}(h_1) \cap \text{dom}(h_2) = \emptyset$; and $h_1 \cdot h_2$ denotes the union of two disjoint heaps. Inductive predicates are interpreted using the least model semantics [42]. Semantics of pure formulas depend on stack valuations; it is straightforward and omitted in Fig. 4, for simplicity.

4.2 Implementation of Separation Logic Instantiation

In the following, we describe how S2SAT_{SL} is realized. In particular, we show how the functions `UA.test`, `OA.test`, `unfold`, and `link.back` are implemented.

Deciding Separation Logic Formula Given an SLPA formula, the functions `UA.test` and `OA.test` in S2SAT_{SL} work similarly, by reducing the formula to a first-order formula systematically and deciding the first-order formula. In the following, we define a function called `eXPure`, which transforms a separation logic formula into a first-order formula. `eXPure` is defined over the symbolic heap as follows:

$$\begin{aligned} \text{eXPure}(\exists \bar{w}. x_1 \mapsto c_1(\bar{v}_1) * \dots * x_n \mapsto c_n(\bar{v}_n) * P_1(\bar{t}_1) * \dots * P_m(\bar{t}_m) \wedge \pi) \equiv \\ \exists \bar{w}. \bigwedge \{x_i \neq \text{null} \mid i \in 1 \dots n\} \wedge \bigwedge \{x_i \neq x_j \mid i, j \in 1 \dots n \text{ and } i \neq j\} \wedge \\ \bigwedge \{\text{inv}(\mathcal{P}, P_j, \bar{t}_j) \mid j \in 1 \dots m\} \wedge \\ \pi \end{aligned}$$

where the reduction at the first line (after \equiv) is for points-to predicates, and the second line is for user-defined predicates. The auxiliary function $\text{inv}(\mathcal{P}, P, \bar{v})$ returns the invariant of the predicate P with a proper renaming.

Next, the auxiliary procedure $\text{sat}_p(\exists \bar{w}. \pi)$ takes a quantified first-order formula as input. It preprocesses the formula and then invokes an SMT solver to solve it. The preprocessing consists of two steps. First, the existential quantifiers \bar{w} are eliminated through a projection $\Pi(\pi, \bar{w})$. Second, remaining existential quantifiers are skolemized and `null` is substituted by special number (i.e., zero). The preprocessed formulas are

of the form of linear arithmetic with free function symbols. These formulas may contain existential (\exists) and universal (\forall) quantifiers but no $\exists\forall$ alternation. Hence, they are naively supported by SMT solvers.

Deriving Unfolding Tree Next, we describe how function `unfold` works in `S2SATSL`. Given a formula, `unfold` selects one predicate instance for unfolding as follows.

$$\frac{\pi_c \equiv \text{eXPure}(\kappa * P(\bar{v}) \wedge \pi) \quad \Gamma_i = \text{unfoldP}(P(\bar{v})_u^o, \pi_c)}{\text{unfold}(\exists \bar{w}_0. \kappa * P(\bar{v})_u^o \wedge \pi) \rightsquigarrow \{\exists \bar{w}_0. \kappa * \Delta_i \wedge \pi \mid \Delta_i \in \Gamma_i\}}$$

Predicate instances in κ are sorted by a pair of unfolding number and ordering number where the former has higher priority. The instance $P(\bar{v})_u^o$ is selected if u is the smallest number of unfoldings and o is the smallest number among instances which have the same unfolding number u . The procedure `unfold` outputs a set of disjuncts which are combined from branches of the predicate P with the remainder $\kappa \wedge \pi$. At the middle, the predicate instance is unfolded by the procedure `unfoldP`. This auxiliary procedure `unfoldP`($P(\bar{t})_u^o, \pi_c$) unfolds the user-defined predicate P with actual parameter \bar{t} under the context π_c . It outputs branches of the predicate P that are not inconsistent with the context. It is formalized as follows.

$$\frac{\pi_c^P \equiv \Pi(\pi_c, \bar{v}) \quad (\bigvee_{i=1}^m (\exists \bar{w}_i. \kappa_i \wedge \pi_i \mid \pi_i^b), \bar{t}) = \text{lookup}(\mathcal{P}, P) \quad \bar{w}'_i = \text{fresh}(\bar{w}_i)}{\frac{(\bar{v}', \pi_{eq}) = \text{freshEQ}(\bar{v}) \quad \rho_p = [\bar{v}'/\bar{t}] \quad \rho_i^{\exists} = [\bar{w}'_i/\bar{w}_i] \quad \rho_i = \rho_p \circ \rho_i^{\exists}}{\text{unfoldP}(P(\bar{v})_u^o, \pi_c) \rightsquigarrow \{\exists \bar{w}'_i. [\rho_i] \kappa_i \wedge [\rho_i] \pi_i \wedge \pi_{eq} \mid \text{sat}_P(\pi_c^P \wedge [\rho_i] \pi_i^b \wedge \pi_{eq}) \neq \text{unsat}, i \in 1 \dots m\}}}$$

In the first line, the procedure looks up the definition of P and refreshes the existential quantifiers (using the function `fresh(...)`). In the second line, formal parameters are substituted by the corresponding actual arguments. Finally, the substituted definition is combined and pruned as shown in the RHS of \rightsquigarrow . Function `freshEQ`(\bar{v}) above refreshes the sequence of variables \bar{v} and produces the equality constraints π_{eq} between the old and new ones, i.e. $\pi_{eq} \equiv \bigwedge v_i = v'_i$. Let $Q(\bar{t})_i^{o_i}$ denote a predicate instance of the derived κ_i , its unfolding number is set to $u+1$ if its corresponding branch Δ_i is recursive. Otherwise, it is u . Its sequence number is set to $o_i + o$.

The branch invariant is used as a *necessary condition* to unfold a branch. The formalism underlying the pruning process is as follows: given a context Δ_c with its over-approximation π_c and a branch Δ_i with its over-approximation π_i^b , if $\pi_c \wedge \pi_i^b$ is unsatisfiable, so is $\Delta_c * \Delta_i$. Similar to the specialization calculus [15], our unfolding mechanism also prunes infeasible disjuncts while unfolding user-defined predicates. However, the specialization calculus performs exhaustive pruning with multiple unfolding that may be highly costly and redundant compared with our one-step unfolding.

Detecting Cyclic Proof In the following, we implement the *matching function* $\mathbf{f}_{\text{cyclic}}$, an instantiation of \mathbf{f}_{fix} , to form a cyclic proof for fixpoint detection. $\mathbf{f}_{\text{cyclic}}$ checks whether there exists a *well-founded* ordering relation R between Δ^{comp} and Δ^{bud} so as to form an *infinite* path following the path between these two nodes. If Δ^{bud} matches with Δ^{comp} , Δ^{bud} is marked as closed. For global infinitary soundness, $\mathbf{f}_{\text{cyclic}}$ only considers those Δ^{bud} and Δ^{comp} of the restricted form as: $\Delta^{\text{comp}} \equiv \Delta_{b_1} * P_1(\bar{t}_1)_m^0 * \dots * P_i(\bar{t}_i)_m^i$, and $\Delta^{\text{bud}} \equiv \Delta_{b_2} * P_1(\bar{t}'_1)_n^0 * \dots * P_k(\bar{t}'_k)_n^k$, where $k \geq i$, $n > m$, Δ_{b_1} and Δ_{b_2} are base formulas.

Like [13], $\mathbf{f}_{\text{cyclic}}$ is implemented using the weakening and substitution principle. In particular, it looks for a substitution θ s.t. $\Delta^{\text{bud}} \theta \implies \Delta^{\text{comp}}$. $\mathbf{f}_{\text{cyclic}}(\Delta^{\text{bud}}, \Delta^{\text{comp}})$

$$\begin{array}{c}
\frac{\frac{\text{[EX-L]}}{\bar{w}' = \text{fresh } \bar{w} \quad \Delta_1[\bar{w}'/\bar{w}] \vdash_{lb} \Delta_2}}{\exists \bar{w} \cdot \Delta_1 \vdash_{lb} \Delta_2}}{\quad} \quad \frac{\frac{\text{[EX-R]}}{\bar{w}' = \text{fresh } \bar{w} \quad \Delta_1 \vdash_{lb} \Delta_2[\bar{w}'/\bar{w}]}}{\Delta_1 \vdash_{lb} \exists \bar{w} \cdot \Delta_2}}{\quad} \quad \frac{\text{[PURE]}}{\pi_1 \implies \pi_2 \quad \pi_1 \vdash_{lb} \pi_2} \\
\text{[SUBST]} \\
\frac{s \in \bar{v} \quad t \in \bar{w} \quad \exists R \cdot \cdot R(s, t) \quad t' = \text{fresh } t \quad (\kappa_1 * \mathbf{P}(\bar{v}) \wedge \pi_1)[t'/t; t/s] \vdash_{lb} \kappa_2 * \mathbf{P}(\bar{w}) \wedge \pi_2}{\kappa_1 * \mathbf{P}(\bar{v}) \wedge \pi_1 \vdash_{lb} \kappa_2 * \mathbf{P}(\bar{w}) \wedge \pi_2} \\
\frac{\frac{\text{[PRED-MATCH]}}{(\kappa_1 \wedge \pi_1)[\bar{v}/\bar{w}] \vdash_{lb} \kappa_2 \wedge \pi_2}}{\kappa_1 * \mathbf{P}(\bar{v}) \wedge \pi_1 \vdash_{lb} \kappa_2 * \mathbf{P}(\bar{w}) \wedge \pi_2}}{\quad} \quad \frac{\text{[PRED-WEAKEN]}}{\mathbf{P}(\bar{w}) \notin \kappa_2 \quad \bar{v} \cap FV(\kappa_1 \wedge \pi_1) = \emptyset \quad \kappa_1 \wedge \pi_1 \vdash_{lb} \kappa_2 \wedge \pi_2}}{\kappa_1 * \mathbf{P}(\bar{v}) \wedge \pi_1 \vdash_{lb} \kappa_2 \wedge \pi_2} \\
\frac{\frac{\text{[PTO-MATCH]}}{(\kappa_1 \wedge \pi_1) \wedge [\bar{v}_1/\bar{v}_2] \vdash_{lb} \kappa_2 \wedge \pi_2}}{\kappa_1 * v \mapsto c(\bar{v}_1) \wedge \pi_1 \vdash_{lb} \kappa_2 * v \mapsto c(\bar{v}_2) \wedge \pi_2}}{\quad} \quad \frac{\text{[PTO-WEAKEN]}}{v \mapsto c(\bar{w}) \notin \kappa_1 \quad (\kappa_1 \wedge \pi_1)[\bar{v}_1/\bar{v}_2] \vdash_{lb} \kappa_2 \wedge \pi_2}}{\kappa_1 * v \mapsto c(v_1) \wedge \pi_1 \vdash_{lb} \kappa_2 * v \mapsto c(v_2) \wedge \pi_2}
\end{array}$$

Fig. 5. Rules for Back-Link.

is formalized as the procedure $\Delta^{bud} \vdash_{lb} \Delta^{comp}$ whose rules are presented in Fig. 5. These rules are applied as follows.

- First, existential variables are refreshed ([EX-L], [EX-R] rules).
- Second, *inductive* variables in Δ^{bud} are substituted ([SUBST] rule). This substitution is based on well-ordering relations R . Let $\mathbf{P}(t)_m^k$ be a predicate instance in Δ^{comp} and its corresponding subformula in Δ^{bud} be $R(s, t)$, then s, t are inductive variables. Two examples of well-founded relations R are structural induction for pointer types where $R(s, t)$ iff s is a subterm of t and natural number induction on integers where $R(s, t)$ iff $0 < s < t$.
- Third, heaps are exhaustively matched ([PRED-MATCH] and [PTO-MATCH] rules) and weakened ([PRED-WEAKEN] and [PTO-WEAKEN] rules). Soundness of these rules directly follows from the frame rule [26,40].
- Last, back-link is decided via the implication between pure formulas ([PURE] rule).

5 Soundness and Termination of S2SAT_{SL}

In the following, we establish the correctness of S2SAT_{SL}.

5.1 Soundness

We show that (i) S2SAT_{SL} is sound and complete for base formulas; and (ii) the functions `UA_test`, `OA_test` and `link_back` in S2SAT_{SL} are sound. These two tasks rely on soundness and completeness of the function `eXPure` over base formulas, soundness of `eXPure` over inductive formulas, and soundness of the function `fcyclic`.

Lemma 1 (Equiv-Satisfiable Reduction). *Let $\Delta \equiv \exists \bar{w} \cdot x_1 \mapsto c_1(\bar{v}_1) * \dots * x_n \mapsto c_n(\bar{v}_n) \wedge \alpha \wedge \phi$ be a base formula. Δ is satisfiable iff $\text{eXPure}(\Delta)$ is satisfiable.*

The proof is based on structural induction on Δ .

Lemma 2 (Over-Approximated Reduction). *Given a formula Δ such that the invariants of user-defined predicates appearing in Δ are sound, then*

$$\forall s, h \cdot s, h \models \Delta \implies s \models \text{eXPure}(\Delta)$$

In the following lemma, we consider the case $\Gamma = \{\}$ at line 8 of Algorithm 1.

Lemma 3. *Given a formula Δ_0 and the matching function $\mathfrak{f}_{\text{cyclic}}$ as presented in the previous section, Δ_0 is UNSAT if $\Gamma = \{\}$ (line 8).*

To prove this Lemma, in [30] we show that there is a “trace manifold” which implies the global infinitary soundness (see [11], ch. 7) when a bud is linked back.

Theorem 2 (Soundness). *Given a formula Δ and a set of user-defined predicates \mathcal{P} ,*

- *Δ is satisfiable if S2SAT_{SL} returns SAT.*
- *if S2SAT_{SL} terminates and returns UNSAT, Δ is unsatisfiable.*

While the soundness of SAT queries follows Lemma 1, the soundness of UNSAT queries follows Lemma 2, and Lemma 3. As satisfiability for SLPA is undecidable [31,30], there is no guarantee that S2SAT_{SL} terminates on all inputs. In the next subsection, we show that S2SAT_{SL} terminates for satisfiable formulas in SLPA and with certain restrictions on the fragment, S2SAT_{SL} always terminates.

5.2 Termination

Termination for SAT In this paragraph, we show that S2SAT_{SL} always terminates when it decides a satisfiable formula. Given a satisfiable formula

$$\Delta \equiv \exists \bar{w} \cdot x_1 \mapsto c_1(\bar{v}_1) * \dots * x_n \mapsto c_n(\bar{v}_n) * P_0(\bar{t}_0)_0^0 * \dots * P_n(\bar{t}_n)_0^n \wedge \pi$$

Hence, there exists a satisfiable base formula Δ_k such as:

$$\Delta_k \equiv x_1 \mapsto c_1(\bar{v}_1) * \dots * x_n \mapsto c_n(\bar{v}_n) * \Delta_{k_0}^{P_0} * \dots * \Delta_{k_n}^{P_n} \wedge \pi$$

where Δ_k^P ($k \geq 0$) denotes a base formula derived by unfolding the predicate P k times and then substituting all predicate instances P by P 's base branch. Let k_m be the maximal number among k_0, \dots, k_n . The breadth-first unfolding manner in the algorithm S2SAT ensures that S2SAT_{SL} identifies Δ_k before it encounters the following leaf:

$$y_1 \mapsto c_1(\bar{t}_1) * \dots * y_i \mapsto c_i(\bar{t}_i) * P_0(\bar{t}_0)_{k_m+1}^{\bar{t}_0} * \dots * P_j(\bar{t}_j)_{k_m+1}^{\bar{t}_j} \wedge \pi$$

We remark that the soundness of cyclic proof ensures that our `link.back` function only considers *infinitely* many unfolding traces. Thus, it never links *finite* many unfolding traces, i.e., traces connecting the root to satisfiable base leaves, like Δ_k .

Decidable Fragment In the following, we describe universal SLPA_{ind} , a fragment of SLPA, for which we prove that S2SAT_{SL} always terminates. Compared to SLPA, universal SLPA_{ind} restricts the set of inductive predicates \mathcal{P} as well as the inputs of S2SAT_{SL} .

Definition 4 (SLPA_{ind}) *An inductive predicate $\text{pred } P(\bar{t}) \equiv \Phi$ is well-founded SLPA_{ind} if it has one induction case with N occurrences of P , and it has the shape as follows.*

$$\Phi \equiv \Phi_0 \vee \exists \bar{w}. x_1 \mapsto c_1(\bar{v}_1) * \dots * x_n \mapsto c_n(\bar{v}_n) * P(\bar{w}_1) * \dots * P(\bar{w}_N) \wedge \pi$$

where Φ_0 is disjunction of base formulas and the two following restrictions.

1. $\forall n \in 1 \dots N \ \bar{w}_n \subseteq \bar{w} \cup \{\text{null}\}$ and \bar{w}_n do not appear in the equalities of π ,
2. if t_i is a numerical parameter and there exists a well-ordering relation R such that $R(s, t_i, w_{1_i}, \dots, w_{m_i})$ ($1 \leq m \leq N$) is a subformula of π , the following conditions hold.
 - t_i is constrained separately (i.e., there does not exist $j \neq i$ and a subformula ϕ of π such that $\{t_i, t_j\} \subseteq \text{FV}(\phi)$ or $\{t_i, w_{n_j}\} \subseteq \text{FV}(\phi)$ or $\{w_{m_i}, w_{n_j}\} \subseteq \text{FV}(\phi)$ $\forall m, n \in 1 \dots N$, and
 - $\forall n \in 1 \dots N, \pi \implies t_i > w_{n_i}$ or $\pi \implies t_i < w_{n_i}$.
 - if $t_i \in \text{FV}(\Phi_0)$ then $\Phi_0 \implies t_i = k$, for some integer k t_i is denoted as inductive parameters.

Restriction 1 guarantees that f_{cyclic} can soundly weaken the heap by discarding irrelevant points-to predicates and $N-1$ occurrences of P (when $N \geq 2$) while it links back. Restriction 2 implies that $t_i > w_i \geq k_1$ or $t_i < w_i \leq k_2$ for some integer k_1, k_2 . This ensures that leaf nodes of unfolding trees of an unsatisfiable input must be UNSAT or linked back.

The above SLPA_{ind} fragment is expressive enough to describe a range of data structures, e.g. sorted lists `sortedll`, lists/trees with size properties, or even AVL trees `avl`.

Definition 5 (Universal SLPA_{ind}) *Given a separation logic formula*

$$\Delta_0 \equiv x_1 \mapsto c_1(\bar{v}_1) * \dots * x_n \mapsto c_n(\bar{v}_n) * P_1(\bar{t}_1) * \dots * P_n(\bar{t}_n) \wedge \phi_0$$

Δ_0 is universal SLPA_{ind} if all predicates P_1, \dots, P_n are well-founded SLPA_{ind} , and if all \bar{x} of free, arithmetical, inductive variables, with $\bar{x} \subseteq (\bar{t}_1 \cup \dots \cup \bar{t}_n)$, ϕ_0 is a conjunction of $\phi_{0,i}$ where $\phi_{0,i}$ is either of the following form: (i) `true`; or (ii) $x_i \geq k_1$ for some integer k_1 ; or (iii) $x_i \leq k_2$ for some integer k_2 .

Theorem 3 (Termination). S2SAT_{SL} terminates for universal SLPA_{ind} formulas.

6 Sound Invariant Inference

In order to perform fully automatic verification without user-provided invariants, S2SAT_{SL} supports automatic invariant inference. In this section, we describe invariant inference from user-defined predicates and predicate branches. While the former is used for over-approximation, the latter is used for context-sensitive predicate unfolding. To infer invariants for a set of user-defined predicates, we first build a dependency graph among the predicates. After that, we process each group of mutual dependent predicates following a bottom-up manner. For simplicity, we present the inference for one directly recursive predicate. The inference for a group of mutual inductive predicates is similar.

Inferring Predicate Invariant Our invariant inference is based on the principle of second-order abduction [28,45]. Given the predicate P defined by m branches as $P(\bar{t}) \equiv \bigvee_{i=1}^m \Delta_i$, we assume a sound invariant of P as an unknown (second-order) variable $I(\bar{t})$. After that we prove the lemma $P(\bar{v}) \vdash I(\bar{v})$ via induction; and simultaneously generate a set of pure relational assumptions using second-order abduction. The steps to prove the above lemma and generate a set of m relational assumptions over I are as follows.

1. Unfold LHS of the lemma to generate a set of m subgoals i.e. $\Delta_i[\bar{v}/\bar{t}] \vdash I(\bar{v})$ where $i \in 1 \dots m$. The original lemma is taken as the induction hypothesis.
2. For each subgoal i , over-approximate its LHS to a pure formula π_i and form an assumption relation $\pi_i \implies I(\bar{v})$. There are two cases to compute π_i .
 - if Δ_i is a base formula, then $\pi_i \equiv \text{eXPure}(\Delta_i)$.
 - if Δ_i includes k instances P such that $\Delta_i \equiv \Delta_{rest_i} * P(\bar{v}_1) * \dots * P(\bar{v}_k)$, then we compute $\pi_{i_0} \equiv \text{eXPure}(\Delta_{rest_i})$, $\pi_{i_j} \equiv I(\bar{v}_j)$, for all $j \in 1 \dots k$, and $\pi_i \equiv \bigwedge_{j=1}^k \pi_{i_j}$.
3. Our system applies a least fixed point analysis to the set of gathered relational assumptions. We use the analyzer LFP presented in [45] to compute these invariants.

We illustrate this procedure to infer an invariant for `sort11`. First, our system introduces an unknown relation $I(\text{root}, n, m)$. Second, it generates the below relational constraints.

$$\begin{aligned} \text{root} \neq \text{null} \wedge n = 1 & \implies I(\text{root}, n, m) \\ \text{root} \neq \text{null} \wedge I(Q, N_1, M_1) \wedge n = N_1 + 1 \wedge m \leq M_1 & \implies I(\text{root}, n, m) \end{aligned}$$

Finally, it analyzes these two constraints and produces the following result:

$$I(\text{root}, n, m) \equiv \text{root} \neq \text{null} \wedge n \geq 1$$

Lemma 4 (Sound Invariant Inference). *Given a predicate $P(\bar{t}) \equiv \Phi$, and \mathcal{R} be a set of relational assumptions generated by the steps above. If \mathcal{R} has a solution, i.e., $I(\bar{v}) \equiv \pi$, then we have $\forall s, h \cdot s, h \models P(\bar{v}), s \models \pi$.*

Proof Sketch: Soundness of Lemma 2 implies that for all $i \in 1 \dots m$, π_i is an over-approximated abstraction of Δ_i . As such, the soundness of this lemma immediately follows from the soundness of second-order abduction [28,45]. \square

Inferring Branch Invariant Given a predicate P defined by m branches as $P(\bar{t}) \equiv \bigvee_{i=1}^m (\exists \bar{w}_i \cdot \Delta_i) \overline{inv} : \pi$, we compute invariants for each branch of P as $\Pi(\text{eXPure}(\Delta_i), \bar{w}_i) \forall i = 1 \dots m$. For example, with the invariant inferred for the predicate `sort11` as above, our system computes its branch invariants π_1^b for the base branch and π_2^b for the inductive branch as below.

$$\begin{aligned} \pi_1^b & \equiv \Pi(\text{eXPure}(\text{root} \mapsto \text{node}(m, \text{null}) \wedge n = 1), \{\}) \equiv \text{root} \neq \text{null} \wedge n = 1 \\ \pi_2^b & \equiv \Pi(\text{eXPure}(\text{root} \mapsto \text{node}(m, q) * \text{sort11}(q, n_1, m_1) \wedge n = n_1 + 1 \wedge m \leq m_1), \\ & \quad \{q, n_1, m_1\}) \equiv \text{root} \neq \text{null} \wedge n \geq 1 \end{aligned}$$

Soundness of `eXPure` implies that the branch invariant over-approximates its branch.

Table 1. Exponential Time and Space Satisfiability Checks.

| succ-circuit (1-20) | | | | | | succ-rec (1-20) | | | | | |
|---------------------|-------|---------------------|----|-------|---------------------|-----------------|---------|---------------------|----|----------|---------------------|
| n | SLSAT | S2SAT _{SL} | n | SLSAT | S2SAT _{SL} | n | SLSAT | S2SAT _{SL} | n | SLSAT | S2SAT _{SL} |
| 1 | 1 ms | 21 ms | 11 | SO | 37.46 s | 1 | 0 ms | 25 ms | 11 | 1796.4 s | 410.92 s |
| 2 | 2 ms | 23 ms | 12 | SO | 170.53s | 2 | 1 ms | 30 ms | 12 | TO | TO |
| 3 | 27 ms | 30 ms | 13 | SO | 988.29s | 3 | 4 ms | 33 ms | 13 | TO | TO |
| 4 | 867ms | 34 ms | 14 | SO | TO | 4 | 21 ms | 39 ms | 14 | X | TO |
| 5 | 30 s | 0.05 s | 15 | SO | TO | 5 | 134 ms | 52 ms | 15 | X | TO |
| 6 | 30 s | 0.09 s | 16 | SO | TO | 6 | 830 ms | 76 ms | 16 | X | TO |
| 7 | SO | 0.20 s | 17 | SO | TO | 7 | 5.0 s | 0.21 s | 17 | X | TO |
| 8 | SO | 0.61 s | 18 | SO | TO | 8 | 29.5 s | 0.87 s | 18 | X | TO |
| 9 | SO | 2.21 s | 19 | SO | TO | 9 | 167.8 s | 4.83 s | 19 | X | TO |
| 10 | SO | 8.49 s | 20 | SO | TO | 10 | 1065 s | 45.28 s | 20 | X | TO |

7 Implementation and Evaluation

We have implemented the proposed solver S2SAT_{SL} and a new interprocedural (top-down) program verification tool, called S2_{td}, which uses S2SAT_{SL}. We make use of Omega Calculator [38] to eliminate existential quantifiers, Z3 [19] as a back-end SMT solver, and FixCalc [37] to find closure form in inferring invariants for user-defined predicates.

In the following, we evaluate S2SAT_{SL} and S2_{td}'s robustness and efficiency on a set of benchmarks from the software verification competition SV-COMP [7]. We also present an evaluation of S2SAT_{SL} in compositional (modular) program verification with the HIP/S2 system [14,28] for a range of data structures.

7.1 Robustness and Efficiency

In [12], Brotherston *et. al.* introduced a new and challenging set of satisfiability benchmarks discussed in Proposition 5.13 of [12]. In this Proposition, Brotherston *et. al.* stated that there exists a family of predicates of size $O(n)$ and that SLSAT runs in $\Omega(2^n)$ time and space regardless of search strategies. Since SLSAT relies on bottom-up and *context-insensitive* fixed point computation, it has to explore all possible models before answering a query. Their approach is designed for computing invariants of shape predicates rather than satisfiability checks. In contrast, S2SAT_{SL} performs top-down and *context-sensitive* searches, as it is dedicated for satisfiability solving. Moreover, it prunes infeasible disjuncts, significantly reduces the search space, and provides better support for model discovery.

We conducted an experiment on comparing SLSAT's and S2SAT_{SL}'s performance on this set of benchmarks. The results are shown in Table 1. The size n of `succ-circuit*` (`succ-rec*`) benchmarks expresses the breadth (depth, resp.) of dependency. This set of benchmarks is a part of the User-Defined Predicate Satisfiability (UDB_{sat}) suite of SL-COMP 2014 [41]. The output is either a definite answer (sat, unsat) with running time (in milliseconds (ms), or seconds (s)), or an error. In particular, SO denotes stack overflow; TO denotes timeout (i.e., tools run longer than 1800 seconds); and X

Table 2. Experimental Results on Complex Data Structures.

| Data Structure (pure props) | #Query | #UNSAT | #SAT | Time |
|-----------------------------------|--------|--------|------|--------|
| Singly llist (size) | 666 | 75 | 591 | 1.25 |
| Even llist (size) | 139 | 125 | 14 | 2.40 |
| Sorted llist (size, sorted) | 217 | 21 | 196 | 0.91 |
| Doubly llist (size) | 452 | 50 | 402 | 2.07 |
| Complete Tree (size, minheight) | 387 | 33 | 354 | 143.98 |
| Heap Trees (size, maxelem) | 467 | 67 | 400 | 13.87 |
| AVL (height, size, near-balanced) | 881 | 64 | 817 | 84.82 |
| BST (height, size, sorted) | 341 | 34 | 307 | 2.28 |
| RBT (size, height, color) | 1741 | 217 | 1524 | 65.54 |
| rose-tree | 55 | 6 | 49 | 0.34 |
| TLL | 128 | 13 | 115 | 0.24 |
| Bubble (size, sorted) | 300 | 20 | 280 | 1.09 |
| Quick sort (size, sorted) | 225 | 29 | 196 | 2.33 |

denotes a fatal error. The experimental results show that $S2SAT_{SL}$ is much more robust and also more efficient than $SLSAT$. While $S2SAT_{SL}$ successfully solved 24 (out of 40) benchmarks, $SLSAT$ was capable of handling 17 benchmarks. Furthermore, on 17 benchmarks that $SLSAT$ discharged successfully, $S2SAT_{SL}$ outperforms $SLSAT$, i.e., about 6.75 (3126seconds/462seconds) times faster. As shown in the table, $S2SAT_{SL}$ ran with neither stack overflow nor fatal errors over all these challenging benchmarks.

7.2 Modular Verification with $S2SAT_{SL}$

In this subsection, we evaluate $S2SAT_{SL}$ in the context of modular program verification. $S2SAT_{SL}$ solver is integrated into the $HIP/S2$ [14,29,28] system to prune infeasible program paths in symbolic execution. Furthermore, $S2SAT_{SL}$ is also used by the entailment procedure $SLEEK$ to discharge verification conditions (VC) generated. In particular, when $SLEEK$ deduces a VC to the following form: $\Delta \vdash \text{emp} \wedge \pi_r$, the error calculus in $SLEEK$ [29] invokes $S2SAT_{SL}$ to discharge the following queries: Δ and $\Delta \wedge \neg \pi_r$ for safety and $\Delta \wedge \pi_r$ for *must* errors. In experiments, we have extracted those VCs generated while $HIP/S2$ verified heap-manipulating programs.

We have evaluated $S2SAT_{SL}$ deciding the VCs discussed above. The experimental results are described in Table 2. Each line shows a test on one program. The first column lists data structures and their pure properties. *rose-trees* are trees with nodes that are allowed to have a variable number of children, stored as doubly-linked lists. *TLL* is a binary tree whose nodes point to their parents and all leaf nodes are linked as a singly-linked list. *#Query* is the number of satisfiability queries sent to $S2SAT_{SL}$ for each data structure. The next two columns report the outputs from $S2SAT_{SL}$. The last column shows the time (in seconds) taken by the $S2SAT_{SL}$ solver. In this experiment, $S2SAT_{SL}$ terminated on all queries. Furthermore it exactly decided all SAT and UNSAT queries. These experimental results affirm the correctness of our algorithm $S2SAT_{SL}$. They also show that $S2SAT_{SL}$ is expressive, effective, and can be integrated into program verification systems for discharging satisfiability problems of separation logic formulas.

7.3 Recursive Program Verification with S2SAT_{SL}

We have evaluated and compared our verification system S2_{td} with state-of-the-art verification tools on a set of SV-COMP benchmarks¹. The results are presented in Table 3.

There are 102 recursive/loop programs taken from *Recursive* and *HeapReach* sub-categories in the benchmark; timeout is set to 180 seconds. In each program, there is at least one user-supplied assertion to model safety properties. The first column identifies the subset of verification systems which

Table 3. Experimental Results on Recursive Programs.

| Tool | #s√ | #e√ | #unk | #s✗ | #e✗ | points | mins |
|------------------|-----------|-----------|-----------|----------|----------|------------|-----------|
| ESBMC [18] | 38 | 40 | 21 | 0 | 3 | 20 | 53 |
| UAutomizer [24] | 17 | 23 | 62 | 0 | 0 | 57 | 23 |
| SeaHorn [22] | 48 | 45 | 5 | 4 | 0 | 77 | 26 |
| CBMC [16] | 33 | 39 | 29 | 1 | 0 | 89 | 90 |
| Smack-Corral [1] | 33 | 37 | 28 | 0 | 0 | 103 | 105 |
| S2 _{td} | 41 | 45 | 16 | 0 | 0 | 127 | 25 |

competed in both the above sub-categories. The next three columns count the instances of correct safe (s√), correct error (e√) and unknown (e.g., timeout). The next two columns capture the number of false positives (s✗) and false negatives (e✗). We rank these tools based on their points. Following the SV-COMP competition, we gave +2 for one s√, +1 for one e√, 0 for unk, -16 for one s✗, and -32 for one e✗. The last column expresses the total time in minutes. The results show that the proposed verification approach is promising; indeed, our system is effective and efficient: it produces the best correctness with *zero* false answers within the nearly-shortest time.

8 Related Work

Close to our work is the SeaHorn verification system [22]. While SeaHorn relies on Z3-PDR to handle inductive predicates on non-heap domains, it is unclear (to us) how SeaHorn supports induction reasoning for heap-based programs (which is one contribution of our present work).

Our S2SAT satisfiability procedure is based on unfolding which is similar to the algorithm in the Leon system [43,44]. Leon, a verifier for functional programs, adds unfolding mechanism for inductive predicates into complete theories. However, Leon only supports classic logic and not structural logic (i.e., separation logic). Neither does Leon support inductive reasoning. Furthermore, our system infers sound invariants for inductive predicates to facilitate over-approximation.

Our work is related to work on developing satisfiability solvers in separation logic. In the following, we summarize the development in this area. Smallfoot [5] has the first implemented decision procedure for a fragment of separation logic. This solver was originally customised to work with spatial formula over list segments. Based on a fixed equality (disequality) constraint branches of the list segment, the proposals presented by [17] and [32] further enhanced decision procedure of this fragment with equality reasoning. They provided normalization rules with graph technique [17] and superposition calculus [32] to infer pointers (dis)equality constraints and used these constraints

¹ <http://sv-comp.sosy-lab.org/2016/>

to prune infeasible branches of predicate instances during unfolding. Although those proposals can decide the formula of that fragment in polynomial time, it is not easy to extend them to a fragment with general inductive predicates (i.e., the fragment SLPA). Decision procedures in [34,36,35] and [33] support decidable fragments of separation logic with inter-reachable data structures using SMT. Our proposal extends these procedures to those fragments with general inductively-defined predicates. Indeed, our decidable fragment can include more complex data structures, such as AVL trees.

$S2SAT_{SL}$ is closely related to the satisfiability solvers [25,12] which are capable of handling separation logic formula with general user-defined predicates. Decision procedures [25] and [12] are able to handle predicates without pure properties. The former described a decidable fragment of user-defined predicates with bounded tree width. The problem of deciding separation logic formulas is then reduced to monadic second-order logic over graph. The latter, $SLSAT$, decides formulas with user-defined predicates via an equi-satisfiable fixed point calculation. The main disadvantage of $SLSAT$ is that it is currently restricted to the domain of pointer equality and disequality, so that it cannot be used to support predicates with pure properties from infinite abstract domains.

Using over-approximation in decision procedures is not new. For example, D’Silva *et. al.* have recently made use of abstract domains inside satisfiability solvers [20,21]. In separation logic, satisfiability procedures in HIP/SLEEK [14] and Dryad [39] decide formulas via a sound reduction that over-approximates predicate instances. HIP/SLEEK and Dryad are capable of proving the validity of a wide range of expressive formulas with arbitrary predicates. However, expressivity comes with cost; as these procedures are incomplete, and they do not address the satisfiability problem. We believe that $S2SAT$ can be integrated into these systems to improve on these two shortcomings.

9 Conclusion and Future Work

We have presented a satisfiability procedure for an expressive fragment of separation logic. Given a formula, our procedure examines both under-approximation (so as to prove SAT) and over-approximation (so as to prove UNSAT). Our procedure was strengthened with invariant generation and cyclic proof detection. We have also implemented a solver and a new verification system for heap-manipulating programs. We have evaluated on a range of competition problems with either complex heap usage patterns or exponential complexity of time and space.

For future work, we might investigate $S2SAT$ -based decision procedures for other complete theories (i.e., Presburger, string, bag/set) augmented with inductive predicates. We would also study a more general decidable fragment of separation logic by relaxing the restrictions for termination. Finally, we would like to improve $S2_{td}$ for array, string and pointer arithmetic reasoning as well as witness generation for erroneous programs.

Acknowledgements. We wish to thank Christopher M. Poskitt for his helpful comments on the manuscript. Quang Loc and Jun Sun are partially supported by NRF grant RGNRF1501 and Wei-Ngan by NRF grant NRF2014NCR-NCR001-040.

References

1. Smack+corral: A modular verifier. In C. Baier and C. Tinelli, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, volume 9035 of *Lecture Notes in Computer Science*, pages 451–454. 2015.
2. M. Barnett, B.-Y. Chang, R. DeLine, B. Jacobs, and K. Leino. Boogie: A modular reusable verifier for object-oriented programs. In *FMCO*, pages 364–387. 2006.
3. C. Barrett, C. Conway, M. Deters, L. Hadarean, D. Jovanović, T. King, A. Reynolds, and C. Tinelli. Cvc4. In *Computer Aided Verification (CAV)*, pages 171–177. 2011.
4. N. E. Beckman, A. V. Nori, S. K. Rajamani, and R. J. Simmons. Proofs from tests. In *ISSTA*, pages 3–14, New York, NY, USA, 2008. ACM.
5. J. Berdine, C. Calcagno, and P. W. O’Hearn. A decidable fragment of separation logic. In *FSTTCS 2004: Foundations of Software Technology and Theoretical Computer Science: 24th International Conference.*, pages 97–109, 2005.
6. J. Berdine, C. Calcagno, and P. W. O’Hearn. Symbolic execution with separation logic. In *Programming Languages and Systems: Third Asian Symposium.*, pages 52–68, 2005.
7. D. Beyer. Reliable and reproducible competition results with benchexec and witnesses (report on SV-COMP 2016). In *TACAS*, pages 887–904, 2016.
8. D. Beyer, A. Cimatti, A. Griggio, M. E. Keremoglu, and R. Sebastiani. Software model checking via large-block encoding. In *Proceedings of 9th International Conference on Formal Methods in Computer-Aided Design, FMCAD 2009.*, pages 25–32, 2009.
9. N. Bjørner, A. Gurfinkel, K. L. McMillan, and A. Rybalchenko. Horn clause solvers for program verification. In *Fields of Logic and Computation II - Essays Dedicated to Yuri Gurevich on the Occasion of His 75th Birthday*, pages 24–51, 2015.
10. N. Bjørner, K. L. McMillan, and A. Rybalchenko. Program verification as satisfiability modulo theories. In *SMT*, pages 3–11, 2012.
11. J. Brotherston. *Sequent Calculus Proof Systems for Inductive Definitions*. PhD thesis, University of Edinburgh, November 2006.
12. J. Brotherston, C. Fuhs, J. A. N. Pérez, and N. Gorogiannis. A decision procedure for satisfiability in separation logic with inductive predicates. In *CSL-LICS ’14*, pages 25:1–25:10, New York, NY, USA, 2014. ACM.
13. J. Brotherston, N. Gorogiannis, and R. L. Petersen. A generic cyclic theorem prover. In *Proceedings of APLAS-12*, LNCS, pages 350–367. Springer, 2012.
14. W. Chin, C. David, H. Nguyen, and S. Qin. Automated verification of shape, size and bag properties via user-defined predicates in separation logic. *SCP*, 77(9):1006–1036, 2012.
15. W.-N. Chin, C. Gherghina, R. Voicu, Q. L. Le, F. Craciun, and S. Qin. A specialization calculus for pruning disjunctive predicates to support verification. In *CAV*. 2011.
16. E. Clarke, D. Kroening, and F. Lerda. A tool for checking ansi-c programs. In *Tools and Algorithms for the Construction and Analysis of Systems*, pages 168–176. 2004.
17. B. Cook, C. Haase, J. Ouaknine, M. Parkinson, and J. Worrell. Tractable reasoning in a fragment of separation logic. In *CONCUR*, volume 6901, pages 235–249. 2011.
18. L. Cordeiro and B. Fischer. Verifying multi-threaded software using smt-based context-bounded model checking. In *Proceedings of the 33rd International Conference on Software Engineering, ICSE ’11*, pages 331–340, New York, NY, USA, 2011. ACM.
19. L. M. de Moura and N. Bjørner. Z3: An Efficient SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages 337–340, 2008.
20. V. D’Silva, L. Haller, and D. Kroening. Satisfiability solvers are static analysers. In *Static Analysis*, volume 7460, pages 317–333. Springer Berlin Heidelberg, 2012.
21. V. D’Silva, L. Haller, and D. Kroening. Abstract satisfaction. In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL ’14*, pages 139–150, New York, NY, USA, 2014. ACM.

22. A. Gurfinkel, T. Kahsai, A. Komuravelli, and J. Navas. The seahorn verification framework. In *Computer Aided Verification*, volume 9206, pages 343–361. 2015.
23. C. Haase, S. Ishtiaq, J. Ouaknine, and M. J. Parkinson. Seloger: A tool for graph-based reasoning in separation logic. In *Computer Aided Verification*, pages 790–795, 2013.
24. M. Heizmann, J. Hoenicke, and A. Podelski. Software model checking for people who love automata. In *Computer Aided Verification*, volume 8044, pages 36–52. 2013.
25. R. Iosif, A. Rogalewicz, and J. Simacek. The tree width of separation logic with recursive definitions. In *Proceedings of the 24th International Conference on Automated Deduction, CADE'13*, pages 21–38, Berlin, Heidelberg, 2013. Springer-Verlag.
26. S. Ishtiaq and P. O'Hearn. BI as an assertion language for mutable data structures. In *ACM POPL*, pages 14–26, London, Jan. 2001.
27. M. Jose and R. Majumdar. Cause clue clauses: error localization using maximum satisfiability. In *PLDI*, pages 437–446, New York, NY, USA, 2011. ACM.
28. Q. L. Le, C. Gherghina, S. Qin, and W.-N. Chin. Shape analysis via second-order bi-abduction. In *Computer Aided Verification (CAV)*, pages 52–68, 2014.
29. Q. L. Le, A. Sharma, F. Craciun, and W.-N. Chin. Towards complete specifications with an error calculus. In *NASA Formal Methods*, pages 291–306, 2013.
30. Q. L. Le, J. Sun, and W.-N. Chin. Satisfiability modulo heap-based programs. 2016. Technical Report, avail. at <http://loc.bitbucket.org/papers/sats1-cav16.pdf>.
31. T. Makoto, Q. L. Le, and W.-N. Chin. Presburger arithmetic and separation logic with inductive definitions. May 2016. Technical Report.
32. J. A. Navarro Pérez and A. Rybalchenko. Separation logic + superposition calculus = heap theorem prover. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pages 556–566, 2011.
33. J. A. N. Pérez and A. Rybalchenko. Separation logic modulo theories. In *Programming Languages and Systems: 11th Asian Symposium, APLAS 2013*, pages 90–106, 2013.
34. R. Piskac, T. Wies, and D. Zufferey. Automating separation logic using smt. In *Computer Aided Verification*, pages 773–789. 2013.
35. R. Piskac, T. Wies, and D. Zufferey. Automating separation logic with trees and data. In *Computer Aided Verification*, pages 711–728. 2014.
36. R. Piskac, T. Wies, and D. Zufferey. Grasshopper: Complete heap verification with mixed specifications. *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages 124–139, 2014.
37. C. Popeea and W.-N. Chin. Inferring disjunctive postconditions. In *ASIAN*, pages 331–345, 2006.
38. W. Pugh. The Omega Test: A fast practical integer programming algorithm for dependence analysis. *Communications of the ACM*, 8:102–114, 1992.
39. X. Qiu, P. Garg, A. Ștefănescu, and P. Madhusudan. Natural proofs for structure, data, and separation. In *PLDI*, pages 231–242, New York, NY, USA, 2013. ACM.
40. J. Reynolds. Separation Logic: A Logic for Shared Mutable Data Structures. In *Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science*, pages 55–74, 2002.
41. M. Sighireanu and D. R. Cok. Report on sl-comp 2014. In *JSAT*, 2016.
42. É.-J. Sims. Extending separation logic with fixpoints and postponed substitution. *Theoretical Computer Science*, 351(2):258–275, 2006.
43. P. Suter, M. Dotta, and V. Kuncak. Decision procedures for algebraic data types with abstractions. In *Proceedings of the 37th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '10*, pages 199–210, New York, NY, USA, 2010. ACM.
44. P. Suter, A. Köksal, and V. Kuncak. Satisfiability modulo recursive programs. In *Proceedings of the 18th International Conference on Static Analysis (SAS)*, pages 298–315. 2011.
45. M.-T. Trinh, Q. L. Le, C. David, and W.-N. Chin. Bi-abduction with pure properties for specification inference. In *APLAS*, pages 107–123, 2013.