

Compositional Entailment Checking for a Fragment of Separation Logic

Constantin Enea¹, Ondřej Lengál², Mihaela Sighireanu¹, and Tomáš Vojnar²

¹ Univ. Paris Diderot, LIAFA CNRS UMR 7089

² FIT, Brno University of Technology, IT4Innovations Centre of Excellence, Czech Republic

Abstract. We present a (semi-)decision procedure for checking entailment between separation logic formulas with inductive predicates specifying complex data structures corresponding to finite nesting of various kinds of linked lists: acyclic or cyclic, singly or doubly linked, skip lists, etc. The decision procedure is compositional in the sense that it reduces the problem of checking entailment between two arbitrary formulas to the problem of checking entailment between a formula and an atom. Subsequently, in case the atom is a predicate, we reduce the entailment to testing membership of a tree derived from the formula in the language of a tree automaton derived from the predicate. We implemented this decision procedure and tested it successfully on verification conditions obtained from programs using singly and doubly linked nested lists as well as skip lists.

1 Introduction

Automatic verification of programs manipulating dynamic linked data structures is highly challenging since it requires one to reason about complex program configurations having the form of graphs of an unbounded size. For that, a highly expressive formalism is needed. Moreover, in order to scale to large programs, the use of such a formalism within program analysis should be highly efficient. In this context, *separation logic* (SL) [14,21] has emerged as one of the most promising formalisms, offering both high expressiveness and scalability. The latter is due to its support of *compositional reasoning* based on the separating conjunction $*$ and the frame rule, which states that if a Hoare triple $\{\phi\}P\{\psi\}$ holds and P does not alter free variables in σ , then $\{\phi * \sigma\}P\{\psi * \sigma\}$ holds too. Therefore, when reasoning about P , one has to manipulate only specifications for the heap region altered by P .

Usually, SL is used together with higher-order *inductive definitions* that describe the data structures manipulated by the program. If we consider general inductive definitions, then SL is undecidable [5]. Various decidable fragments of SL have been introduced in the literature [1,13,18,3] by restricting the syntax of the inductive definitions and the boolean structure of the formulas.

In this work, we focus on a fragment of SL with inductive definitions that allows one to specify program configurations (heaps) containing finite nestings of various kinds of linked lists (acyclic or cyclic, singly or doubly linked, skip lists, etc.), which are common in practice. This fragment contains formulas of the form $\exists \vec{X}. \Pi \wedge \Sigma$ where X is a set of variables, Π is a conjunction of (dis)equalities, and Σ is a set of *spatial*

atoms connected by the separating conjunction. Spatial atoms can be *points-to atoms*, which describe values of pointer fields of a given heap location, or *inductively defined predicates*, which describe data structures of an unbounded size. We propose a novel (semi-)decision procedure for checking the validity of entailments of the form $\varphi \Rightarrow \psi$ where φ may contain existential quantifiers and ψ is a quantifier-free formula. Such a decision procedure can be used in Hoare-style reasoning to check inductive invariants but also in program analysis frameworks to decide termination of fixpoint computations. As usual, checking entailments of the form $\bigvee_i \varphi_i \Rightarrow \bigvee_j \psi_j$ can be soundly reduced to checking that for each i there exists j such that $\varphi_i \Rightarrow \psi_j$.

The key insight of our decision procedure is an idea to use the semantics of the separating conjunction in order to reduce the problem of checking $\varphi \Rightarrow \psi$ to the problem of checking a set of simpler entailments where the right-hand side is an inductively-defined predicate $P(\dots)$. This reduction shows that the compositionality principle holds not only for deciding the validity of Hoare triples but also for deciding the validity of entailments between two formulas. To infer (dis)equalities implied by spatial atoms, our reduction to checking simpler entailments is based on boolean unsatisfiability checking, which is in co-NP but can usually be checked efficiently by current SAT solvers.

Further, to check entailments $\varphi \Rightarrow P(\dots)$ resulting from the above reduction, we define a decision procedure based on the membership problem for tree automata (TA). In particular, we reduce the entailment to testing membership of a tree derived from φ in the language of a TA $\mathcal{A}[P]$ derived from $P(\dots)$. The tree encoding of φ preserves some edges of the graph, called *backbone edges*, while others are re-directed to new nodes, related to the original destination by special symbols. Roughly, such a symbol may be a variable represented by the original destination, or it may show how to reach the original destination using backbone edges only.

Our procedure is complete for formulas speaking about non-nested singly as well as doubly linked lists. Moreover, it runs in polynomial time modulo an oracle for deciding validity of a boolean formula. The procedure is incomplete for nested list structures because it does not consider all possible ways in which targets of inner pointer fields of nested list predicates can be aliased. The construction can be easily extended to become complete even in such cases, but then it becomes exponential. However, even in this case, it is exponential in the size of the definition of the inductive predicates, and not in the size of the formulas, which remains acceptable in practice.

We implemented our decision procedure and tested it successfully on verification conditions obtained from programs using singly and doubly linked nested lists as well as skip lists. The results show that our procedure does not only have a theoretically favorable complexity (for the given context), but it also behaves nicely in practice, at the same time offering the additional benefit of compositionality that can be exploited within larger verification frameworks caching the simpler entailment queries.

Contribution. Overall, the contribution of this paper is a novel (semi-)decision procedure for a rich class of verification conditions with singly as well as doubly linked lists, nested lists, and skip lists. As discussed in more detail in Section 9, existing works that can efficiently deal with fragments of SL capable of expressing verification conditions for programs handling complex dynamic data structures are still rare. Indeed, we are not aware of any technique that could decide the class of verification conditions considered

in this paper at the same level of efficiency as our procedure. In particular, compared with other approaches using TAs [13,12], our procedure is compositional as it uses TAs recognizing models of predicates, not models of entire formulas (further differences are discussed in the related work section). Moreover, our TAs recognize in fact formulas that entail a given predicate, reducing SL entailment to the membership problem for TAs, not the more expensive inclusion problem as in other works.

2 Separation Logic Fragment

Let $Vars$ be a set of *program variables*, ranged over using x, y, z , and $LVars$ a set of *logical variables*, disjoint from $Vars$, ranged over using X, Y, Z . We assume that $Vars$ contains a distinguished variable `nil`. Also, let \mathbb{F} be a set of *fields*.

We consider the fragment of separation logic whose syntax is given below:

$$\begin{array}{llll}
x, y \in Vars \text{ program variables} & X, Y \in LVars \text{ logical variables} & E, F ::= x \mid X & \\
f \in \mathbb{F} \text{ fields} & \rho \subseteq \mathbb{F} \times (Vars \cup LVars) & P \in \mathbb{P} \text{ predicates} & \\
& \vec{B} \in (Vars \cup LVars)^* \text{ vectors of variables} & & \\
\\
\Pi ::= E = F \mid E \neq F \mid \Pi \wedge \Pi & & \text{pure formulas} & \\
\Sigma ::= emp \mid E \mapsto \rho \mid P(E, F, \vec{B}) \mid \Sigma * \Sigma & & \text{spatial formulas} & \\
\varphi \triangleq \exists \vec{X}. \Pi \wedge \Sigma & & \text{formulas} &
\end{array}$$

W.l.o.g., we assume that existentially quantified logical variables have unique names. The set of program variables used in a formula φ is denoted by $pv(\varphi)$. By $\varphi(\vec{E})$ (resp. $\rho(\vec{E})$), we denote a formula (resp. a set of field-variable couples) whose set of free variables is \vec{E} . Given a formula φ , $pure(\varphi)$ denotes its pure part Π . We allow set operations to be applied on vectors. Moreover, $E \neq \vec{B}$ is a shorthand for $\bigwedge_{B_i \in \vec{B}} E \neq B_i$.

The *points-to atom* $E \mapsto \rho$ specifies that the heap contains a location E whose f_i field points to F_i , for all i . W.l.o.g., we assume that each field f_i appears at most once in a set of pairs ρ . The fragment is parameterized by a set \mathbb{P} of *inductively defined predicates*; intuitively, $P(E, F, \vec{B})$ describes a possibly empty *nested list segment* delimited by its arguments, i.e., all the locations it represents are reachable from E and allocated on the heap except the locations in $\{F\} \cup \vec{B}$.

Inductively defined predicates. We consider predicates defined as

$$\begin{aligned}
P(E, F, \vec{B}) \triangleq & (E = F \wedge emp) \vee \\
& \left(E \neq \{F\} \cup \vec{B} \wedge \exists X_{t1}. \Sigma(E, X_{t1}, \vec{B}) * P(X_{t1}, F, \vec{B}) \right)
\end{aligned} \tag{1}$$

where Σ is an existentially-quantified formula, called *the matrix of P* , of the form:

$$\begin{aligned}
\Sigma(E, X_{t1}, \vec{B}) \triangleq & \exists \vec{Z}. E \mapsto \rho(\{X_{t1}\} \cup \vec{V}) * \Sigma' & \text{where } \vec{V} \subseteq \vec{Z} \cup \vec{B} \text{ and} \\
\Sigma' ::= & Q(Z, U, \vec{Y}) \mid \circ^{1+} Q[Z, \vec{Y}] \mid \Sigma' * \Sigma' & \\
& \text{for } Z \in \vec{Z}, U \in \vec{Z} \cup \vec{B} \cup \{E, X_{t1}\}, \vec{Y} \subseteq \vec{B} \cup \{E, X_{t1}\}, \text{ and} & \\
\circ^{1+} Q[Z, \vec{Y}] \triangleq & \exists Z'. \Sigma_Q(Z, Z', \vec{Y}) * Q(Z', Z, \vec{Y}) & \text{where } \Sigma_Q \text{ is the matrix of } Q.
\end{aligned} \tag{2}$$

singly linked lists:

$$\text{ls}(E, F) \triangleq \text{lomp}(E, F) \vee (E \neq F \wedge \exists X_{t1}. E \mapsto \{(f, X_{t1})\} * \text{ls}(X_{t1}, F))$$

lists of acyclic lists:

$$\text{nll}(E, F, B) \triangleq \text{lomp}(E, F) \vee (E \neq \{F, B\} \wedge \exists X_{t1}, Z. E \mapsto \{(s, X_{t1}), (h, Z)\} * \text{ls}(Z, B) * \text{nll}(X_{t1}, F, B))$$

lists of cyclic lists:

$$\text{ncl}(E, F) \triangleq \text{lomp}(E, F) \vee (E \neq F \wedge \exists X_{t1}, Z. E \mapsto \{(s, X_{t1}), (h, Z)\} * \odot^{1+} \text{ls}[Z] * \text{ncl}(X_{t1}, F))$$

skip lists with three levels:

$$\begin{aligned} \text{skl}_3(E, F) &\triangleq \text{lomp}(E, F) \vee (E \neq F \wedge \exists X_{t1}, Z_1, Z_2. E \mapsto \{(f_3, X_{t1}), (f_2, Z_2), \\ &\quad (f_1, Z_1)\} * \text{skl}_1(Z_1, Z_2) * \text{skl}_2(Z_2, X_{t1}) * \text{skl}_3(X_{t1}, F)) \\ \text{skl}_2(E, F) &\triangleq \text{lomp}(E, F) \vee (E \neq F \wedge \exists X_{t1}, Z_1. E \mapsto \{(f_3, \text{nil}), (f_2, X_{t1}), \\ &\quad (f_1, Z_1)\} * \text{skl}_1(Z_1, X_{t1}) * \text{skl}_2(X_{t1}, F)) \\ \text{skl}_1(E, F) &\triangleq \text{lomp}(E, F) \vee (E \neq F \wedge \exists X_{t1}. E \mapsto \{(f_3, \text{nil}), (f_2, \text{nil}), \\ &\quad (f_1, X_{t1})\} * \text{skl}_1(X_{t1}, F)) \end{aligned}$$

Fig. 1. Examples of inductive definitions ($\text{lomp}(E, F) \triangleq E = F \wedge \text{emp}$)

The formula Σ specifies the values of the fields defined in E (using the atom $E \mapsto \rho(\{X_{t1}\} \cup \vec{V})$, where the fields in ρ are constants in \mathbb{F}) and the (possibly cyclic) nested list segments starting at the locations \vec{Z} referenced by fields of E . We assume that Σ contains a single points-to atom in order to simplify the presentation. Notice that the matrix of a predicate P does not contain applications of P .

The macro $\odot^{1+} Q[Z, \vec{Y}]$ is used to represent a *non-empty* cyclic (nested) list segment on Z whose shape is described by the predicate Q .

We consider several restrictions on Σ which are defined using its *Gaifman graph* $Gf[\Sigma]$. The set of vertices of $Gf[\Sigma]$ is given by the set of free and existentially quantified variables in Σ , i.e., $\{E, X_{t1}\} \cup \vec{B} \cup \vec{Z}$. The edges in $Gf[\Sigma]$ represent spatial atoms: for every (f, X) in ρ , $Gf[\Sigma]$ contains an edge from E to X labeled by f ; for every predicate $Q(Z, U, \vec{Y})$, $Gf[\Sigma]$ contains an edge from Z to U labeled by Q ; and for every macro $\odot^{1+} Q[Z, \vec{Y}]$, $Gf[\Sigma]$ contains a self-loop on Z labeled by Q .

The first restriction is that $Gf[\Sigma]$ contains no cycles other than self-loops built solely of edges labeled by predicates. This ensures that the predicate is *precise*, i.e., for any heap, there exists at most one sub-heap on which the predicate holds. Precise assertions are very important for concurrent separation logic [11].

The second restriction requires that all the maximal paths of $Gf[\Sigma]$ start in E and end either in a self-loop or in a node from $\vec{B} \cup \{E, X_{t1}\}$. This restriction ensures that (a) all the heap locations in the interpretation of a predicate are reachable from the head of the list and that (b) only the locations represented by variables in $F \cup \vec{B}$ are dangling. Moreover, for simplicity, we require that every vertex of $Gf[\Sigma]$ has at most one outgoing edge labeled by a predicate.

For example, the predicates defined in Fig. 1 describe singly linked lists, lists of acyclic lists, lists of cyclic lists, and skip lists with three levels.

We define the relation $\prec_{\mathbb{P}}$ on the set of predicates \mathbb{P} by $P_1 \prec_{\mathbb{P}} P_2$ iff P_2 occurs in the matrix of P_1 . The reflexive and transitive closure of $\prec_{\mathbb{P}}$ is denoted by $\prec_{\mathbb{P}}^*$. For example, if $\mathbb{P} = \{\text{skl}_1, \text{skl}_2, \text{skl}_3\}$, then $\text{skl}_3 \prec_{\mathbb{P}} \text{skl}_2$ and $\text{skl}_3 \prec_{\mathbb{P}}^* \text{skl}_1$.

$$\begin{aligned}
(S, H) \models P(E, F, \vec{B}) & \quad \text{iff there exists } k \in \mathbb{N} \text{ s.t. } (S, H) \models P^k(E, F, \vec{B}) \text{ and} \\
& \quad \text{ldom}(H) \cap (\{S(F)\} \cup \{S(B) \mid B \in \vec{B}\}) = \emptyset \\
(S, H) \models P^0(E, F, \vec{B}) & \quad \text{iff } (S, H) \models E = F \wedge \text{emp} \\
(S, H) \models P^{k+1}(E, F, \vec{B}) & \quad \text{iff } (S, H) \models E \neq \{F\} \cup \vec{B} \wedge \\
& \quad \exists X_{t1}. \Sigma(E, X_{t1}, \vec{B}) * P^k(X_{t1}, F, \vec{B})
\end{aligned}$$

Fig. 2. The semantics of predicate atoms

Given a predicate P of the matrix Σ as in (2), let $\mathbb{F}_{\mapsto}(P)$ denote the set of fields f occurring in a pair (f, X) of ρ . For example, $\mathbb{F}_{\mapsto}(\text{n11}) = \{s, h\}$ and $\mathbb{F}_{\mapsto}(\text{sk1}_3) = \mathbb{F}_{\mapsto}(\text{sk1}_1) = \{f_3, f_2, f_1\}$. Also, let $\mathbb{F}_{\mapsto}^*(P)$ denote the union of $\mathbb{F}_{\mapsto}(P')$ for all $P \prec_{\mathbb{P}}^* P'$. For example, $\mathbb{F}_{\mapsto}^*(\text{n11}) = \{s, h, f\}$.

We assume that $\prec_{\mathbb{P}}^*$ is a partial order, i.e., there are no mutually recursive definitions in \mathbb{P} . Moreover, for simplicity, we assume that for any two predicates P_1 and P_2 which are incomparable w.r.t. $\prec_{\mathbb{P}}^*$, it holds that $\mathbb{F}_{\mapsto}(P_1) \cap \mathbb{F}_{\mapsto}(P_2) = \emptyset$. This assumption avoids predicates named differently but having exactly the same set of models.

Semantics. Let $Locs$ be a set of *locations*. A *heap* is a pair (S, H) where $S : Vars \cup LVars \rightarrow Locs$ maps variables to locations and $H : Locs \times \mathbb{F} \rightarrow Locs$ is a partial function that defines values of fields for some of the locations in $Locs$. The domain of H is denoted by $dom(H)$ and the set of locations in the domain of H is denoted by $ldom(H)$. As usual, we assume that `nil` is interpreted to a location $S(\text{nil}) \notin ldom(H)$. We say that a location ℓ (resp., a variable E) is *allocated* in the heap (S, H) or that (S, H) *allocates* ℓ (resp., E) iff ℓ (resp., $S(E)$) belongs to $ldom(H)$.

The set of heaps satisfying a formula φ is defined by the relation $(S, H) \models \varphi$. For brevity, we define in Fig. 2 the relation \models for predicate atoms only. The complete definition of \models can be found in [8]. Note that a heap satisfying a predicate atom $P(E, F, \vec{B})$ doesn't allocate any variable in $F \cup \vec{B}$; the locations represented by these variables don't belong to its domain. A heap satisfying this property is called *well-formed w.r.t. the atom* $P(E, F, \vec{B})$. The set of models of a formula φ is denoted by $\llbracket \varphi \rrbracket$. Given two formulas φ_1 and φ_2 , we say that φ_1 entails φ_2 , denoted by $\varphi_1 \Rightarrow \varphi_2$, iff $\llbracket \varphi_1 \rrbracket \subseteq \llbracket \varphi_2 \rrbracket$. By an abuse of notation, $\varphi_1 \Rightarrow E = F$ (resp., $\varphi_1 \Rightarrow E \neq F$) denotes the fact that E and F are interpreted to the same location (resp., different locations) in all models of φ_1 .

3 Compositional Entailment Checking

We define a procedure for reducing the problem of checking the validity of an entailment between two formulas to the problem of checking the validity of an entailment between a formula and an atom. We assume that the right-hand side of the entailment is a quantifier-free formula (which usually suffices for checking verification conditions in practice). The reduction can be extended to the general case, but it becomes incomplete.

```

 $\varphi_1 \leftarrow \text{norm}(\varphi_1); \varphi_2 \leftarrow \text{norm}(\varphi_2);$            // normalization
if  $\varphi_1 = \text{false}$  then return true;
if  $\varphi_2 = \text{false}$  then return false;
if  $\text{pure}(\varphi_1) \not\Rightarrow \text{pure}(\varphi_2)$  then return false;           // entailment of pure parts
foreach  $a_2 : \text{points-to atom in } \varphi_2$  do                 // entailment of shape parts
     $\varphi_1[a_2] \leftarrow \text{select}(\varphi_1, a_2);$ 
    if  $\varphi_1[a_2] \not\Rightarrow a_2$  then return false;
for  $P_2 \leftarrow \max_{\prec}(\mathbb{P})$  down to  $\min_{\prec}(\mathbb{P})$  do
    forall the  $a_2 = P_2(E, F, \vec{B}) : \text{predicate atom in } \varphi_2 \text{ s.t. } \text{pure}(\varphi_1) \not\Rightarrow E = F$  do
         $\varphi_1[a_2] \leftarrow \text{select}(\varphi_1, a_2);$ 
        if  $\varphi_1[a_2] \not\Rightarrow_{sh} a_2$  then return false;
return isMarked( $\varphi_1$ );

```

Fig. 3. Compositional entailment checking (\prec is any total order compatible with $\prec_{\mathbb{P}}^*$)

3.1 Overview of the Reduction Procedure

We consider the problem of deciding validity of entailments $\varphi_1 \Rightarrow \varphi_2$ with φ_2 quantifier-free. We assume $pv(\varphi_2) \subseteq pv(\varphi_1)$; otherwise, the entailment is not valid.

The main steps of the reduction are given in Fig. 3. The reduction starts by a normalization step (described in Sec. 3.2), which adds to each of the two formulas all (dis-)equalities implied by spatial sub-formulas and removes all atoms $P(E, F, \vec{B})$ representing *empty* list segments, i.e., those where $E = F$ occurs in the pure part. The normalization of a formula outputs *false* iff the input formula is unsatisfiable.

In the second step, the procedure tests the entailment between the pure parts of the normalized formulas. This can be done using any decision procedure for quantifier-free formulas in the first-order theory with equality.

For the spatial parts, the procedure builds a mapping from spatial atoms of φ_2 to sub-formulas of φ_1 . Intuitively, the sub-formula $\varphi_1[a_2]$ associated to an atom a_2 of φ_2 , computed by `select`, describes the region of a heap modeled by φ_1 that should satisfy a_2 . For predicate atoms $a_2 = P_2(E, F, \vec{B})$, `select` is called (in the second loop) only if there exists a model of φ_1 where the heap region that should satisfy a_2 is non-empty, i.e., $E = F$ does not occur in φ_1 . In this case, `select` does also check that for any model of φ_1 , the sub-heap corresponding to the atoms in $\varphi_1[a_2]$ is well-formed w.r.t. a_2 (see Sec. 3.3). This is needed since all heaps described by a_2 are well-formed.

Note that in the well-formedness check above, one cannot speak about $\varphi_1[a_2]$ alone. This is because without the rest of φ_1 , $\varphi_1[a_2]$ may have models which are not well-formed w.r.t. a_2 even if the sub-heap corresponding to $\varphi_1[a_2]$ is well-formed for any model of φ_1 . For example, let $\varphi_1 = \text{ls}(x, y) * \text{ls}(y, z) * z \mapsto \{(f, t)\}$, $a_2 = \text{ls}(x, z)$, and $\varphi_1[a_2] = \text{ls}(x, y) * \text{ls}(y, z)$. If we consider only models of φ_1 , the sub-heaps corresponding to $\varphi_1[a_2]$ are all well-formed w.r.t. a_2 , i.e., the location bound to z is not allocated in these sub-heaps. However, $\varphi_1[a_2]$ alone has lasso-shaped models where the location bound to z is allocated on the path between x and y .

Once $\varphi_1[a_2]$ is obtained, one needs to check that all sub-heaps modeled by $\varphi_1[a_2]$ are also models of a_2 . For points-to atoms a_2 , this boils down to a syntactic identity (modulo some renaming given by the equalities in the pure part of φ_1). For predicate atoms a_2 , a special entailment operator \Rightarrow_{sh} (defined in Sec. 3.5) is used. We cannot use the usual entailment \Rightarrow since, as we have seen in the example above, $\varphi_1[a_2]$ may

have models which are not sub-heaps of models of φ_1 . Thus, $\varphi_1[a_2] \Rightarrow_{sh} a_2$ holds iff all models of $\varphi_1[a_2]$, which are well-formed w.r.t. a_2 , are also models of a_2 .

If there exists an atom a_2 of φ_2 , which is not entailed by the associated sub-formula, then $\varphi_1 \Rightarrow \varphi_2$ is not valid. By the semantics of the separating conjunction, the sub-formulas of φ_1 associated with two different atoms of φ_2 must not share spatial atoms. Due to this, the spatial atoms obtained from each application of `select` are marked and cannot be reused in the future. Note that the mapping is built by enumerating the atoms of φ_2 in a particular order: first, the points-to atoms and then the inductive predicates, in a decreasing order wrt $\prec_{\mathbb{P}}$. This is important for completeness (see Sec. 3.3).

The procedure `select` is detailed in Sec. 3.3. It returns *emp* if the construction of the sub-formula of φ_1 associated with the input atom fails (this implies that also the entailment $\varphi_1 \Rightarrow \varphi_2$ is not valid). If all entailments between formulas and atoms are valid, then $\varphi_1 \Rightarrow \varphi_2$ holds provided that all spatial atoms of φ_1 are marked (tested by `isMarked`). In Sec. 3.5, we introduce a procedure for checking entailments between a formula and a spatial atom.

Graph representations. Some of the sub-procedures mentioned above work on a graph representation of the input formulas, called *SL graphs* (which are different from the Gaifman graphs of Sec. 2). Thus, a formula φ is represented by a directed graph $G[\varphi]$ where each node represents a maximal set of variables equal w.r.t. the pure part of φ , and each edge represents a disequality $E \neq F$ or a spatial atom. Every node n is labeled by the set of variables $\text{Var}(n)$ it represents; for every variable E , $\text{Node}(E)$ denotes the node n s.t. $E \in \text{Var}(n)$. Next, (1) a disequality $E \neq F$ is represented by an undirected edge from $\text{Node}(E)$ to $\text{Node}(F)$, (2) a spatial atom $E \mapsto \{(f_1, E_1), \dots, (f_n, E_n)\}$ is represented by n directed edges from $\text{Node}(E)$ to $\text{Node}(E_i)$ labeled by f_i for each $1 \leq i \leq n$, and (3) a spatial atom $P(E, F, \vec{B})$ is represented by a directed edge from $\text{Node}(E)$ to $\text{Node}(F)$ labeled by $P(\vec{B})$. Edges are referred to as disequality, points-to, or predicate edges, depending on the atom they represent. For simplicity, we may say that the graph representation of a formula is simply a formula.

Running example. In the following, we use as a running example the entailment $\psi_1 \Rightarrow \psi_2$ between the following formulas:

$$\psi_1 \equiv \exists Y_1, Y_2, Y_3, Y_4, Z_1, Z_2, Z_3. x \neq z \wedge Z_2 \neq z \wedge \quad (3)$$

$$x \mapsto \{(s, Z_2), (h, Z_1)\} * Z_2 \mapsto \{(s, y), (h, Z_3)\} * \text{ls}(Z_1, z) * \text{ls}(Z_3, z) * \\ \text{ls}(y, Y_1) * \text{skl}_2(y, Y_3) * \text{ls}(Y_1, Y_2) *$$

$$Y_3 \mapsto \{(f_2, t), (f_1, Y_4)\} * Y_4 \mapsto \{(f_2, \text{nil}), (f_1, t)\} * t \mapsto \{(s, Y_2)\}$$


$$\psi_2 \equiv y \neq t \wedge \text{nil}(x, y, z) * \text{skl}_2(y, t) * t \mapsto \{(s, y)\} \quad (4)$$

The graph representations of these formulas are drawn in the top part of Fig. 4.

3.2 Normalization

To infer the implicit (dis-)equalities in a formula, we adapt the boolean abstraction proposed in [10] for our logic. Therefore, given a formula φ , we define an equisatisfiable boolean formula $\text{BoolAbs}[\varphi]$ in CNF over a set of boolean variables containing the boolean variable $[E = F]$ for every two variables E and F occurring in φ and the boolean variable $[E, a]$ for every variable E and spatial atom a of the form $E \mapsto \rho$ or

Initially: $\psi_1 \Rightarrow \psi_2$



After normalization: $\text{norm}(\psi_1) \Rightarrow \text{norm}(\psi_2)$




Fig. 4. An example of applying compositional entailment checking. Points-to edges are represented by simple lines, predicate edges by double lines, and disequality edges by dashed lines. For readability, we omit the points-to edge from Y_4 to nil , some of the labeling with existentially-quantified variables, and some of the disequality edges in the normalized graphs.

$P(E, F, \vec{B})$ in φ . The variable $[E = F]$ denotes the equality between E and F while $[E, a]$ denotes the fact that the atom a describes a heap where E is allocated.

Given $\varphi \triangleq \exists \vec{X}. \Pi \wedge \Sigma$, $\text{BoolAbs}[\varphi] \triangleq F(\Pi) \wedge F(\Sigma) \wedge F_{=} \wedge F_*$ where $F(\Pi)$ and $F(\Sigma)$ encode the atoms of φ (using \oplus to denote xor), $F_{=}$ encodes reflexivity, symmetry, and transitivity of equality, and F_* encodes the semantics of the separating conjunction:

$$F(\Pi) \triangleq \bigwedge_{E=F \in \Pi} [E = F] \wedge \bigwedge_{E \neq F \in \Pi} \neg[E = F] \quad F(\Sigma) \triangleq \bigwedge_{a=E \mapsto \rho \in \Sigma} [E, a] \wedge \bigwedge_{a=P(E, F, \vec{B}) \in \Sigma} [E, a] \oplus [E = F]$$

$$F_{=} \triangleq \bigwedge_{E_1 \text{ variable in } \varphi} [E_1 = E_1] \wedge \bigwedge_{E_1, E_2 \text{ variables in } \varphi} ([E_1 = E_2] \Leftrightarrow [E_2 = E_1]) \wedge \bigwedge_{E_1, E_2, E_3 \text{ variables in } \varphi} ([E_1 = E_2] \wedge [E_2 = E_3] \Rightarrow [E_1 = E_3]) \quad (5)$$

$$F_* \triangleq \bigwedge_{\substack{E, F \text{ variables in } \varphi \\ a, a' \text{ different atoms in } \Sigma}} ([E = F] \wedge [E, a] \Rightarrow \neg[F, a'])$$

For the formula ψ_1 in our running example (Eq. 3), $\text{BoolAbs}[\psi_1]$ is a conjunction of several formulas including:

1. $[y, \text{skl}_2(y, Y_3)] \oplus [y = Y_3]$, which encodes the atom $\text{skl}_2(y, Y_3)$,
2. $[Y_3, Y_3 \mapsto \{(f_1, Y_4), (f_2, t)\}]$ and $[t, t \mapsto \{(s, Y_2)\}]$, encoding points-to atoms,

3. $([y = t] \wedge [t, t \mapsto \{(s, Y_2)\}]) \Rightarrow \neg[y, \text{skl}_2(y, Y_3)]$, which encodes the separating conjunction between $t \mapsto \{(s, Y_2)\}$ and $\text{skl}_2(y, Y_3)$,
4. $([Y_3 = t] \wedge [t, t \mapsto \{(s, Y_2)\}]) \Rightarrow \neg[Y_3, Y_3 \mapsto \{(f_1, Y_4), (f_2, t)\}]$, which encodes the separating conjunction between $t \mapsto \{(s, Y_2)\}$ and $Y_3 \mapsto \{(f_1, Y_4), (f_2, t)\}$.

Proposition 1. *Let φ be a formula. Then, $\text{BoolAbs}[\varphi]$ is equisatisfiable with φ , and for any variables E and F of φ , $\text{BoolAbs}[\varphi] \Rightarrow [E = F]$ (resp., $\text{BoolAbs}[\varphi] \Rightarrow \neg[E = F]$) iff $\varphi \Rightarrow E = F$ (resp. $\varphi \Rightarrow E \neq F$).*

For example, $\text{BoolAbs}[\psi_1] \Rightarrow \neg[y = t]$, which is a consequence of the sub-formulas we have given above together with F_- .

If $\text{BoolAbs}[\varphi]$ is unsatisfiable, then the output of $\text{norm}(\varphi)$ is *false*. Otherwise, the output of $\text{norm}(\varphi)$ is the formula φ' obtained from φ by (1) adding all (dis-)equalities implied by $\text{BoolAbs}[\varphi]$ and (2) removing all predicates $P(E, F, \vec{B})$ s.t. $E = F$ occurs in the pure part. For example, the normalizations of ψ_1 and ψ_2 are given in the bottom part of Fig. 4. Note that the *ls* atoms reachable from y are removed because $\text{BoolAbs}[\psi_1] \Rightarrow [y = Y_1]$ and $\text{BoolAbs}[\psi_1] \Rightarrow [Y_1 = Y_2]$.

The following result is important for the completeness of the *select* procedure.

Proposition 2. *Let $\text{norm}(\varphi)$ be the normal form of a formula φ . For any two distinct nodes n and n' in the SL graph of $\text{norm}(\varphi)$, there cannot exist two disjoint sets of atoms A and A' in $\text{norm}(\varphi)$ s.t. both A and A' represent paths between n and n' .*

If we assume for contradiction that $\text{norm}(\varphi)$ contains two such sets of atoms, then, by the semantics of the separating conjunction, $\varphi \Rightarrow E = F$ where E and F label n and n' , respectively. Therefore, $\text{norm}(\varphi)$ does not include all equalities implied by φ , which contradicts its definition.

3.3 Selection of Spatial Atoms

Points-to atoms. Let $\varphi_1 \triangleq \exists \vec{X}. \Pi_1 \wedge \Sigma_1$ be a normalized formula. The procedure $\text{select}(\varphi_1, E_2 \mapsto \rho_2)$ outputs the sub-formula $\exists \vec{X}. \Pi_1 \wedge E_1 \mapsto \rho_1$ s.t. $E_1 = E_2$ occurs in Π_1 if it exists, or *emp* otherwise. The procedure *select* is called only if φ_1 is satisfiable and consequently, φ_1 cannot contain two different atoms $E_1 \mapsto \rho_1$ and $E'_1 \mapsto \rho'_1$ such that $E_1 = E'_1 = E_2$. Also, if there exists no such points-to atom, then $\varphi_1 \Rightarrow \varphi_2$ is not valid. Indeed, since φ_2 does not contain existentially quantified variables, a points-to atom in φ_2 could be entailed only by a points-to atom in φ_1 .

In the running example, $\text{select}(\psi_1, t \mapsto \{(s, y)\}) = \exists Y_2. y = Y_2 \wedge \dots \wedge t \mapsto \{(s, Y_2)\}$ (we have omitted some existential variables and pure atoms).

Predicate atoms. Given an atom $a_2 = P_2(E_2, F_2, \vec{B}_2)$, $\text{select}(\varphi_1, a_2)$ builds a sub-graph G' of $G[\varphi_1]$, and then it checks whether the sub-heaps described by G' are well-formed w.r.t. a_2 . If this is not true or if G' is empty, then it outputs *emp*. Otherwise, it outputs the formula $\exists \vec{X}. \Pi_1 \wedge \Sigma'$ where Σ' consists of all atoms represented by edges of the sub-graph G' . Let $\text{Dangling}[a_2] = \text{Node}(F_2) \cup \{\text{Node}(B) \mid B \in \vec{B}_2\}$.

The sub-graph G' is defined as the union of all paths of $G[\varphi_1]$ that (1) consist of edges labeled by fields in $\mathbb{F}_{\mapsto}^*(P_2)$ or predicates Q with $P_2 \prec_{\mathbb{P}}^* Q$, (2) start in the node

labeled by E_2 , and (3) end either in a node from $\text{Dangling}[a_2]$ or in a cycle, in which case they must not traverse nodes in $\text{Dangling}[a_2]$. The paths in G' that end in a node from $\text{Dangling}[a_2]$ must not traverse other nodes from $\text{Dangling}[a_2]$. Therefore, G' does not contain edges that start in a node from $\text{Dangling}[a_2]$. The instances of G' for $\text{select}(\psi_1, \text{nll}(x, y, z))$ and $\text{select}(\psi_1, \text{skl}_2(y, t))$ are emphasized in Fig. 4.

Next, the procedure `select` checks that in every model of φ_1 , the sub-heap described by G' is well-formed w.r.t. a_2 . Intuitively, this means that all the cycles in the sub-heap are explicitly described in the inductive definition of P_2 . For example, if $\varphi_1 = \text{ls}(x, y) * \text{ls}(y, z)$ and $\varphi_2 = a_2 = \text{ls}(x, z)$, then the graph G' corresponds to the entire formula φ_1 and it may have lasso-shaped models (z may belong to the path between x and y) that are not well-formed w.r.t. $\text{ls}(x, z)$ (whose inductive definition describes only acyclic heaps). Therefore, the procedure `select` returns *emp*, which proves that the entailment $\varphi_1 \Rightarrow \varphi_2$ does not hold. For our running example, for any model of ψ_1 , in the sub-heap modeled by the graph $\text{select}(\psi_1, \text{skl}_2(y, t))$ in Fig. 4, t should not be (1) interpreted as an allocated location in the list segment $\text{skl}_2(y, Y_3)$ or (2) aliased to one of nodes labeled by Y_3 and Y_4 .

The well-formedness test is equivalent to the fact that for every variable $V \in \{F_2\} \cup \vec{B}_2$ and every model of φ_1 , the interpretation of V is different from all allocated locations in the sub-heap described by G' . This is in turn equivalent to the fact that for every variable $V \in \{F_2\} \cup \vec{B}_2$, the two following conditions hold:

1. For every predicate edge e included in G' that does not end in $\text{Node}(V)$, V is allocated in all models of $E \neq F \wedge (\varphi_1 \setminus G')$ where E and F are variables labeling the source and the destination of e , respectively, and $\varphi_1 \setminus G'$ is obtained from φ_1 by deleting all *spatial* atoms represented by edges of G' .
2. For every variable V' labeling the source of a points-to edge of G' , $\varphi_1 \Rightarrow V \neq V'$.

The first condition guarantees that V is not interpreted as an allocated location in a list segment described by a predicate edge of G' (this trivially holds for predicate edges ending in $\text{Node}(V)$). If V was not allocated in some model (S, H_1) of $E \neq F \wedge (\varphi_1 \setminus G')$, then one could construct a model (S, H_2) of G' where e would be interpreted to a non-empty list and $S(V)$ would equal an allocated location inside this list. Therefore, there would exist a model of φ_1 , defined as the union of (S, H_1) and (S, H_2) , in which the heap region described by G' would not be well-formed w.r.t. a_2 .

For example, in the graph $\text{select}(\psi_1, \text{skl}_2(y, t))$ in Fig. 4, t is not interpreted as an allocated location in the list segment $\text{skl}_2(y, Y_3)$ since t is allocated (due to the atom $t \mapsto \{(s, Y_2)\}$) in all models of $y \neq Y_3 \wedge (\psi_1 \setminus \text{select}(\psi_1, \text{skl}_2(y, t)))$.

To check that variables are allocated, we use the following property: given a formula $\varphi \triangleq \exists \vec{X}. \Pi \wedge \Sigma$, a variable V is allocated in every model of φ iff $\exists \vec{X}. \Pi \wedge \Sigma * V \mapsto \{(f, V_1)\}$ is unsatisfiable. Here, we assume that f and V_1 are not used in φ . Note that, by Prop. 1, unsatisfiability can be decided using the boolean abstraction `BoolAbs`.

The second condition guarantees that V is different from all allocated locations represented by sources of points-to edges in G' . For the graph $\text{select}(\psi_1, \text{nll}(x, y, z))$ in Fig. 4, the variable z must be different from all existential variables labeling a node which is the source of a points-to edge. These disequalities appear explicitly in the formula. In general, by Prop. 1, $\varphi_1 \Rightarrow V \neq V'$ can be decided using the boolean abstraction.

3.4 Soundness and Completeness

The following theorem states that the procedure given in Fig. 3 is sound and complete. The soundness is a direct consequence of the semantics. The completeness is a consequence of Prop. 1 and 2. In particular, Prop. 2 implies that the sub-formula returned by $\text{select}(\varphi_1, a_2)$ is the only one that can describe a heap region satisfying a_2 .

Theorem 1. *Let φ_1 and φ_2 be two formulas s.t. φ_2 is quantifier-free. Then, $\varphi_1 \Rightarrow \varphi_2$ iff the procedure in Fig. 3 returns true.*

3.5 Checking Entailments between a Formula and an Atom

Given a formula φ and an atom $P(E, F, \vec{B})$, we define a procedure for checking that $\varphi \Rightarrow_{sh} P(E, F, \vec{B})$, which works as follows: (1) $G[\varphi]$ is transformed into a tree $\mathcal{T}[\varphi]$ by splitting nodes that have multiple incoming edges, (2) the inductive definition of $P(E, F, \vec{B})$ is used to define a TA $\mathcal{A}[P]$ s.t. $\mathcal{T}[\varphi]$ belongs to the language of $\mathcal{A}[P]$ only if $\varphi \Rightarrow_{sh} P(E, F, \vec{B})$. Notice that we do not require the reverse implication in order to keep the size of $\mathcal{A}[P]$ polynomial in the size of the inductive definition of P . Thus, $\mathcal{A}[P]$ does not recognize the tree representations of all formulas φ s.t. $\varphi \Rightarrow_{sh} P(E, F, \vec{B})$. The transformation of graphs into trees is presented in Sec. 4 while the definition of the TA is introduced in Sec. 5. In Sec. 6, we also discuss how to obtain a complete method by generating a TA $\mathcal{A}[P]$ of an exponential size.

4 Representing SL Graphs as Trees

We define a canonical representation of SL graphs in the form of trees, which we use for checking \Rightarrow_{sh} . In this representation, the disequality edges are ignored because they have been dealt with previously when checking entailment of pure parts.

We start by explaining the main concepts of the tree encoding using the generic labeled graph in Fig. 5(a). We consider a graph G where all nodes are reachable from a distinguished node called *Root* (this property is satisfied by all SL graphs returned by the select procedure). To construct a tree representation of G , we start with its spanning tree (emphasized using bold edges) and proceed with splitting any node with at least two incoming edges, called a *join node*, into several copies, one for each incoming edge not contained in the spanning tree. The obtained tree is given in Fig. 5(b).

Not to lose any information, the copies of nodes should be labeled with the identity of the original node, which is kept in the spanning tree. However, since the representation does not use node identities, we label every original node with a representation of the path from *Root* to this node in the spanning tree, and we assign every copied node a label describing how it can reach the




Fig. 5. The tree representation of a generic graph

original node in the spanning tree. For example, if a node n has the label $\text{alias} \uparrow[g_1]$, this denotes the fact that n is a copy of some join node, which is the first ancestor of n in the spanning tree that is reachable from $Root$ by a path formed of a (non-empty) sequence of g_1 edges. Further, n labelled by $\text{alias} \uparrow\downarrow[f_1 f_2]$ denotes roughly that (1) the original node is reachable from $Root$ by a path formed by a (non-empty) sequence of f_1 edges followed by a (non-empty) sequence of f_2 edges, and (2) the original node can be reached from n by going up in the tree until the first node that is labelled by a prefix of $f_1 f_2$ and then down until the first node labelled with $f_1 f_2$. The exact definition of these labels can be found later in this section. In general, a label of the form $\text{alias} \uparrow[\dots]$ will be used when breaking loops while a label of the form $\text{alias} \uparrow\downarrow[\dots]$ will be used when breaking parallel paths between nodes. Moreover, if the original node is labeled by a variable, e.g., x , then we will use a label of the form $\text{alias}[x]$. This set of labels is enough to obtain a tree representation from SL graphs that can entail a spatial atom from the considered fragment; for arbitrary graphs, this is not the case.

When applying this construction to an SL graph, the most technical part consists in defining the spanning tree. Based on the inductive definition of predicates, we consider a total order on fields $\prec_{\mathbb{F}}$ that is extended to sequences of fields, $\prec_{\mathbb{F}^*}$, in a lexicographic way. Then, the spanning tree is defined by the set of paths labeled by sequences of fields which are minimum according to the order $\prec_{\mathbb{F}^*}$.

Intuitively, the order $\prec_{\mathbb{F}}$ reflects the order in which the unfolding of the inductive definition of P is done: (1) Fields used in the atom $E \mapsto \rho$ of the matrix of P are ordered before fields of any other predicate called by P . (2) Fields appearing in ρ and going “one-step forward” (i.e., occurring in a pair (f, X_{t1})) are ordered before fields going “down” (i.e., occurring in a pair (f, Z) with $Z \in \vec{Z}$), which are ordered before fields going to the “border” (i.e., occurring in a pair (f, X) with $X \in \vec{B} \setminus \{\text{nil}\}$).

Formally, given a predicate P with the matrix Σ as in (2), we identify in the set $\mathbb{F}_{\mapsto}(P)$ three disjoint sets: (a) $\mathbb{F}_{\mapsto X_{t1}}(P)$ is the set of fields f occurring in a pair (f, X_{t1}) of ρ , (b) $\mathbb{F}_{\mapsto \vec{Z}}(P)$ the set of fields f occurring in a pair (f, Z) of ρ with $Z \in \vec{Z}$, and (c) $\mathbb{F}_{\mapsto \vec{B}}(P)$ the set of fields f occurring in a pair (f, X) of ρ with $X \in \vec{B} \setminus \{\text{nil}\}$. Then, we assume that there exists a total order $\prec_{\mathbb{F}}$ on fields s.t., for all P, P_1, P_2 in \mathbb{P} :

$$\begin{aligned} & \forall f_1 \in \mathbb{F}_{\mapsto X_{t1}}(P) \forall f_2 \in \mathbb{F}_{\mapsto \vec{Z}}(P) \forall f_3 \in \mathbb{F}_{\mapsto \vec{B}}(P). f_1 \prec_{\mathbb{F}} f_2 \prec_{\mathbb{F}} f_3 \text{ and} \\ & (f_1 \in \mathbb{F}_{\mapsto}(P_1) \wedge f_2 \in \mathbb{F}_{\mapsto}(P_2) \wedge f_1 \neq f_2 \wedge P_1 \prec_{\mathbb{P}} P_2) \Rightarrow f_1 \prec_{\mathbb{F}} f_2. \end{aligned}$$

For example, if $\mathbb{P} = \{\text{nil}, \text{ls}\}$ or $\mathbb{P} = \{\text{nlcl}, \text{ls}\}$, then $s \prec_{\mathbb{F}} h \prec_{\mathbb{F}} f$; and if $\mathbb{P} = \{\text{skl}_2, \text{skl}_1\}$, then $f_2 \prec_{\mathbb{F}} f_1$. The order $\prec_{\mathbb{F}}$ is extended to a lexicographic order $\prec_{\mathbb{F}^*}$ over sequences in \mathbb{F}^* . Note that the pointer fields going to nil are not involved in the constraints above (they are not included in neither one of the sets $\mathbb{F}_{\mapsto X_{t1}}(P)$, $\mathbb{F}_{\mapsto \vec{Z}}(P)$, or $\mathbb{F}_{\mapsto \vec{B}}(P)$). They are treated differently because, by definition, there is no pointer field defined in nil . For example, if $\mathbb{P} = \{\text{skl}_1\}$, then $f_2 \prec_{\mathbb{F}} f_1$ and $f_1 \prec_{\mathbb{F}} f_2$ are both valid total orderings on fields.

An f -edge of an SL graph is a points-to edge labeled by f or a predicate edge labeled by $P(\vec{N})$ s.t. the minimum field in $\mathbb{F}_{\mapsto}(P)$ w.r.t. $\prec_{\mathbb{F}}$ is f .

Let G be an SL graph and $P(E, F, \vec{B})$ an atom for which we want to prove that $G \Rightarrow_{sh} P(E, F, \vec{B})$. We assume that all nodes of G are reachable from the node $Root$

labeled by E , which is ensured when G is constructed by `select`. The tree encoding of G is computed by the procedure `toTree($G, P(E, F, \vec{B})$)` that consists of four consecutive steps that are presented below (see also [8]).

Node marking. First, `toTree` computes a mapping \mathbb{M} , called *node marking*, which defines the spanning tree of G . Intuitively, for each node n , $\mathbb{M}(n)$ is the sequence of fields labeling a path reaching n from *Root* that is minimal w.r.t. $\prec_{\mathbb{F}^*}$. Formally, let π be a path in G starting in *Root* and consisting of the sequence of edges $e_1 e_2 \dots e_n$. The *labeling of π* , denoted by $\mathbb{L}(\pi)$, is the sequence of fields $f_1 f_2 \dots f_n$ s.t. for all i , e_i is an f_i -edge. The node marking is defined by

$$\forall n \in G \quad \mathbb{M}(n) \triangleq \text{Reduce}(\min_{\prec_{\mathbb{F}}}(\mathbb{F}_{\mapsto}(P)) \cdot \mathbb{L}_{\min}(n)), \quad (6)$$

$$\mathbb{L}_{\min}(n) \triangleq \min_{\prec_{\mathbb{F}^*}} \{ \mathbb{L}(\pi) \mid \text{Root} \xrightarrow{\pi} n \} \quad (7)$$

where *Reduce* rewrites the sub-words of the form f^+ to f , for any field f . For technical reasons, we add the minimum field (w.r.t. $\prec_{\mathbb{F}}$) in $\mathbb{F}_{\mapsto}(P)$ at the beginning of all $\mathbb{M}(n)$.

Fig. 6(b)–(c) depicts two graphs and the markings of their nodes. (For readability, we omit the markings of the nodes labeled by y and t .)

Splitting join nodes. The join nodes are split in two consecutive steps, denoted as `splitLabeledJoin` and `splitJoin`, depending on whether they are labeled by variables in $\{E, F\} \cup \vec{B}$ or not. In both cases, only the edges of the spanning tree are kept in the tree, the other edges are redirected to fresh copies labeled by some alias $[..]$.

For any join node n , the spanning tree edge is the f -edge (m, n) such that $\text{Reduce}(\mathbb{M}(m) f) = \mathbb{M}(n)$, i.e., (m, n) is at the end of the minimum path leading to n . (For *Root*, all incoming edges are not in the spanning tree.)

In `splitLabeledJoin`, a graph G' is obtained by replacing in G any edge (m, n) such that n is labeled by some $V \in \{E, F\} \cup \vec{B}$ and (m, n) is not in the spanning tree by an edge (m, n') with the same label, where n' is a fresh copy of n labeled by alias $[V]$. Moreover, for uniformity, all (even non-join) nodes labeled by a variable $V \in F \cup \vec{B}$ are labeled by alias $[V]$ in G' . Fig. 6(a) gives the output graph of `splitLabeledJoin` on the SL graphs returned in our running example by `select($\psi_1, \text{nll}(x, y, z)$)` and `select($\psi_1, \text{skl}_2(y, t)$)`.

Subsequently, `splitJoin` builds from G' a tree by splitting unlabeled join nodes as follows. Let n be a join node and (m, n) an edge not in the spanning tree of G' (and G). The edge (m, n) is replaced in the tree by an edge (m, n') with the same edge label, where n' is a fresh copy of n labeled by:

- alias $\uparrow[\mathbb{M}(n)]$ if m is reachable from n and all predecessors of m (by a simple path) marked by $\mathbb{M}(n)$ are also predecessors of n . Intuitively, this label is used to break loops, and it refers to the closest predecessor of n' having the given marking. The use of this labeling is illustrated in Fig. 6(b).
- alias $\uparrow\downarrow[\mathbb{M}(n)]$ if there is a node p which is a predecessor of m s.t. all predecessors of m that have a unique successor marked by $\mathbb{M}(n)$ are also predecessors of p , and n is the unique successor of p marked by $\mathbb{M}(n)$. Intuitively, this transformation is used to break multiple paths between p and n as illustrated in Fig. 6(c).³

³ The combination of up and down arrows in the label corresponds to the need of going up and then down in the resulting tree—whereas in the previous case, it suffices to go up only.

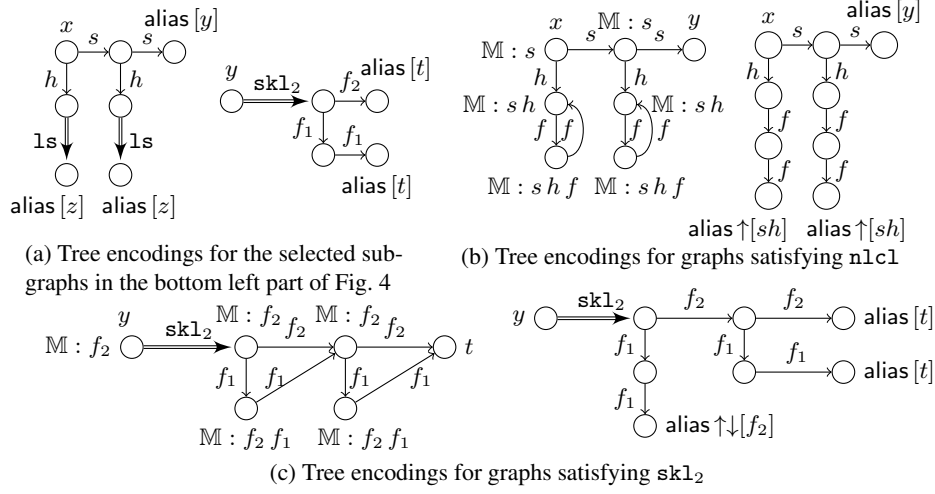


Fig. 6. Tree encodings

If the relation between n and n' does not satisfy the constraints mentioned above, the result of `splitJoin` is an error, i.e., the \perp tree.

At the end of these steps, we obtain a tree with labels on edges (using fields $f \in \mathbb{F}$ or predicates $Q(\vec{B})$) and labels on nodes of the form $alias[...]$.

Updating the labels. In the last step, two transformations are done on the tree. First, the labels of predicate edges are changed in order to replace each argument X different from $\{F\} \cup \vec{B}$ by a label $alias \uparrow[M(n)]$ or $alias \downarrow[M(n)]$, which describes the position of the node n labeled by X w.r.t. *the source node of the predicate edge*.

Finally, as the generated trees will be tested for membership in the language of a TA which accepts node-labelled trees only, the labels of edges are moved to the labels of their source nodes and concatenated in the order given by $\prec_{\mathbb{F}}$ (predicates in the labels are ordered according to the minimum field in their matrix).

The following property ensures the soundness of the entailment procedure:

Proposition 3. *Let $P(E, F, \vec{B})$ be an atom and G an SL graph. If $\text{toTree}(G, P(E, F, \vec{B})) = \perp$, then $G \not\models P(E, F, \vec{B})$.*

5 Tree Automata Recognizing Tree Encodings of SL Graphs

Next, we proceed to the construction of tree automata $\mathcal{A}[P(E, F, \vec{B})]$ that recognize tree encodings of SL graphs that imply atoms of the form $P(E, F, \vec{B})$. Due to space constraints, we cannot provide a full description of the TA construction (which we give in [8]). Instead, we give an intuitive description only and illustrate it on two typical examples (for now, we leave our running examples, TAs for which are given in [8]).

Tree automata. A (non-deterministic) *tree automaton* recognizing tree encodings of SL graphs is a tuple $\mathcal{A} = (Q, q_0, \Delta)$ where Q is a set of states, $q_0 \in Q$ is the initial state, and Δ is a set of transition rules of the form $q \hookrightarrow a_1(q_1), \dots, a_n(q_n)$ or $q \hookrightarrow a$, where

$n > 0$, $q, q_1, \dots, q_n \in Q$, a_i is an SL graph edge label (we assume them to be ordered w.r.t. the ordering of fields as for tree encodings), and a is $\text{alias} \uparrow[m]$, $\text{alias} \uparrow \downarrow[m]$, or $\text{alias}[V]$. The set of trees $L(\mathcal{A})$ recognized by \mathcal{A} is defined as usual.

Definition of $\mathcal{A}[P(E, F, \vec{B})]$. The tree automaton $\mathcal{A}[P(E, F, \vec{B})]$ is defined starting from the inductive definition of P . If P does not call other predicates, the TA simply recognizes the tree encodings of the SL graphs that are obtained by “concatenating” a sequence of Gaifman graphs representing the matrix $\Sigma(E, X_{t1}, \vec{B})$ and predicate edges $P(E, X_{t1}, \vec{B})$. In these sequences, occurrences of the Gaifman graphs representing the matrix and the predicate edges can be mixed in an arbitrary order and in an arbitrary number. Intuitively, this corresponds to a partial unfolding of the predicate P in which there appear concrete segments described by points-to edges as well as (possibly multiple) segments described by predicate edges. Concatenating two Gaifman graphs means that the node labeled by X_{t1} in the first graph is merged with the node labeled by E in the other graph. This is illustrated on the following example.

Consider a predicate $P_1(E, F, B)$ that does not call other predicates and that has the matrix

$$\Sigma_1 \triangleq E \mapsto \{(f_1, X_{t1}), (f_2, X_{t1}), (f_3, B)\}.$$

The tree automaton \mathcal{A}_1 for $P_1(E, F, B)$ has transition rules given in Fig. 7. Rules (1)–(3) recognize the tree encoding of the Gaifman graph of Σ_1 , assuming the following total order on the fields: $f_1 \prec_{\mathbb{F}} f_2 \prec_{\mathbb{F}} f_3$. Rule (4) is used to distinguish the “last” instance of this tree encoding, which ends in the node labeled by $\text{alias}[F]$ accepted by Rule (5). Finally, Rules (6) and (7) recognize predicate edges labeled by $P_1(B)$. As in the previous case, we distinguish the predicate edge that ends in the node labeled by $\text{alias}[F]$.

Note that the TA given above exhibits the simple and generic skeleton of TAs accepting tree encodings of list segments defined in our SL fragment: The initial state q_0 is used in a loop to traverse over an arbitrary number of folded (Rule 6) and unfolded (Rule 1) occurrences of the list segments, and the state q_3 is used to recognize the end of the backbone (Rule 5). The other states (here, q_2) are used to accept alias labels only. The same skeleton can be observed in the TA recognizing tree encodings of singly linked lists, which is given in Fig. 8.

When P calls other predicates, the automaton recognizes tree encodings of concatenations of more general SL graphs, obtained from $Gf[\Sigma]$ by replacing predicate edges with unfoldings of these predicates. On the level of TAs, this operation corresponds to a substitution of transitions labelled by predicates with TAs for the nested predicates. During this substitution, $\text{alias}[\dots]$ labels occurring in the TA for the nested predicate need to be modified. Labels of the form $\text{alias} \uparrow[m]$ and $\text{alias} \uparrow \downarrow[m]$ are adjusted by prefixing m with the marking of the source state of the transition. Moreover, labels of the form $\text{alias}[V]$ are substituted by the marking of $\text{Node}(V)$ w.r.t. the higher-level matrix.




Fig. 7. $\mathcal{A}[P_1(E, F, B)]$




Fig. 8. $\mathcal{A}[\text{ls}(E, F)]$

Let us consider a predicate $P_2(E, F)$ that calls P_1 and that has the matrix

$$\Sigma_2 \triangleq \exists Z. E \mapsto \{(g_1, X_{t1}), (g_2, Z)\} \wedge \odot^{1+} P_1[Z, E].$$

The TA \mathcal{A}_2 for $P_2(E, F)$ includes the transition rules (1') $qq_0 \hookrightarrow g_1(qq_0), g_2(q_0)$ given in Fig. 9. These rules are complemented by (2') $qq_0 \hookrightarrow g_1(qq_1), g_2(q_0)$ the rules of \mathcal{A}_1 where $\text{alias}[F]$ is substituted by (3') $qq_1 \hookrightarrow \text{alias}[F]$ $\text{alias}\uparrow[g_1 g_2]$, $\text{alias}[B]$ by $\text{alias}\uparrow[g_1]$, and $\text{alias}\uparrow\downarrow[f_1]$ is (4') $qq_0 \hookrightarrow P_2(qq_0)$ substituted by $\text{alias}\uparrow\downarrow[g_1 g_2 f_1]$. Rule (1') and the ones (5') $qq_0 \hookrightarrow P_2(qq_1)$ imported (after renaming of the labels) from \mathcal{A}_1 describe trees obtained from the tree encoding of $Gf[\Sigma_2]$ by replacing the edge looping in Z with a tree recognized by \mathcal{A}_1 . According to $Gf[\Sigma_2]$, the node marking of Z is $g_1 g_2$, and so the label $\text{alias}[F]$ shall be substituted by $\text{alias}\uparrow[g_1 g_2]$, and the marking $\text{alias}\uparrow\downarrow[f_1]$ shall be substituted by $\text{alias}\uparrow\downarrow[g_1 g_2 f_1]$.

Fig. 9. $\mathcal{A}[P_2(E, F)]$

The following result states the correctness of the tree automata construction.

Theorem 2. *For any atom $P(E, F, \vec{B})$ and any SL graph G , if the tree generated by $\text{toTree}(G, P(E, F, \vec{B}))$ is recognized by $\mathcal{A}[P(E, F, \vec{B})]$, then $G \Rightarrow P(E, F, \vec{B})$.*

6 Completeness and Complexity

In general, there exist SL graphs that entail $P(E, F, \vec{B})$ whose tree encodings are *not* recognized by $\mathcal{A}[P(E, F, \vec{B})]$. The models of these SL graphs are nested list segments where inner pointer fields specified by the matrix of P are aliased. For example, the TA for skl_2 does not recognize tree encodings of SL graphs modeling heaps where X_{t1} and Z_1 are interpreted to the same location.

The construction of TAs explained above can be easily extended to cover such SL graphs (cf. [8]), but the size of the obtained automata may become exponential in the size of P (defined as the number of symbols in the matrices of all Q with $P \prec_{\mathbb{P}}^* Q$) as the construction considers all possible aliasing scenarios of targets of inner pointer fields permitted by the predicate definition.

For the verification conditions that we have encountered in our experiments, the TAs defined above are precise enough in the vast majority of the cases. In particular, note that the TAs generated for the predicates for `ls` and `d11` (defined below) are precise. We have, however, implemented even the above mentioned extension and realized that it also provides acceptable performance.

In conclusion, the overall complexity of the semi-decision procedure (where aliases between variables in the definition of a predicate are ignored) runs in polynomial time modulo an oracle for deciding validity of a boolean formula (needed in normalization procedure). The complete decision procedure is exponential in the size of the predicates, and not of the formulas, which remains acceptable in practice.

7 Extensions

The procedures presented above can be extended to a larger fragment of SL that uses more general inductively defined predicates. In particular, they can be extended to cover

finite nestings of singly or doubly linked lists (DLL). To describe DLL segments between two locations E and F where P is the predecessor of E and S is the successor of F , one can use the predicate

$$\text{dll}(E, F, P, S) \triangleq (E = S \wedge F = P \wedge \text{emp}) \vee \left(E \neq S \wedge F \neq P \wedge \exists X_{t1}. E \mapsto \{(next, X_{t1}), (prev, P)\} * \text{dll}(X_{t1}, F, E, S) \right). \quad (8)$$

Finite nestings of such list segments can be defined by replacing the matrix $E \mapsto \{(next, X_{t1}), (prev, P)\}$ with more general formulas that include other predicates.

The key point in this extension is the definition of the tree encoding. Basically, one needs to consider two more types of labels for the tree nodes: alias $\uparrow^2[\alpha]$ with $\alpha \in \mathbb{F}^*$, which denotes the fact that the node is a copy of its second predecessor of marking α , and alias $\uparrow_{\text{last}}[\alpha]$ with $\alpha \in \mathbb{F}^*$, which denotes the fact that the node is a copy of the last successor of marking α of its first predecessor that has a successor of marking α . The first label is needed to handle inner nodes of doubly linked lists, which have two incoming edges, one from their successor and one from their predecessor, while the second label is needed to “break” cyclic doubly linked lists. In the latter case, the label is used for the copy of the predecessor of the head of the list (cf. [8] for more details).

8 Implementation and Experimental Results

We implemented our decision procedure in a solver called SPEN (SeParation logic ENtailment). The tool takes as input an entailment problem $\varphi_1 \Rightarrow \varphi_2$ (including the definition of the used predicates) encoded in the SMTLIB2 format. For non-valid entailments, SPEN prints the atom of φ_2 which is not entailed by a sub-formula of φ_1 . The tool is based on the MINISAT solver for deciding unsatisfiability of boolean formulas and the VATA library [15] as the tree automata backend.

We applied SPEN to entailment problems that use various recursive predicates. First, we considered the benchmark provided in [16], which uses only the `ls` predicate. This benchmark has been used in the `ls` division of the first competition of Separation Logic solvers, SL-COMP 2014⁴. It consists of 292 problems split into three classes: the first two classes contain problems generated randomly according to the rules specified in [16], whereas the last class contains problems obtained from the verification conditions generated by the tool SMALLFOOT [2]. SPEN solved the full benchmark in less than 8 seconds (CPU time), which is the second time of the division; the winner of the division was a specialized solver for the `ls` predicate, Asterix [17], which spent less than 4 seconds for the `ls` benchmark. An explanation for this result is that in the current version of SPEN, a new TA has to be built for each `ls` edge, which is time-consuming for problems with several `ls` edges (this issue will be remedied in future versions).

Moreover, the TA for `ls` is quite small, and so the above experiments did not evaluate thoroughly the performance of our procedure for checking entailments between formulas and atoms. For that, we further considered the experiments listed in Table 1,

⁴ The participants in this competition are available at <http://smtcomp.sourceforge.net/2014/participants.shtml>, and the benchmarks for all divisions of the competition are available at <https://github.com/mihasihi/smtcomp14-sl>.

Table 1. Running SPEN on entailments between formulas and atoms.

φ_2	n11			nlc1			skl ₃			dll		
φ_1	tc1	tc2	tc3	tc1	tc2	tc3	tc1	tc2	tc3	tc1	tc2	tc3
Time [ms]	344	335	319	318	316	317	334	349	326	358	324	322
Status	vld	vld	inv	vld	vld	inv	vld	vld	inv	vld	vld	inv
States/Trans. of $\mathcal{A}[\varphi_2]$	6/17			6/15			80/193			9/16		
Nodes/Edges of $T(Gf[\varphi_1])$	7/7	7/7	6/7	10/9	7/7	6/6	7/7	8/8	6/6	7/7	7/7	5/5

among which skl₃ required the extension of our approach to a full decision procedure as discussed in Sec. 6. The full benchmark is available with our tool [9] and it includes the 43 problems of the division “fixed definitions” of SL-COMP 2014. The entailment problems are extracted from verification conditions of operations like adding or deleting an element at the start, in the middle, or at the end of various kinds of list segments. Table 1 gives for each example the running time, the valid/invalid status, and the size of the tree encoding and TA for φ_1 and φ_2 , respectively. SPEN was the winner in this division of SL-COMP 2014 (in front of [4,6]) and it was the only tool that solved all problems of this division.

9 Related Work

Several decision procedures for fragments of SL have been introduced in the literature [1,5,7,10,13,12,16,18,19,4].

Some of these works [1,5,7,16] consider a fragment of SL that uses only one predicate describing singly linked lists, which is a much more restricted setting than what is considered in this paper. In particular, Cook *et al* [7] prove that the satisfiability/entailment problem can be solved in polynomial time. Piskac *et al* [18] show that the boolean closure of this fragment can be translated to a decidable fragment of first-order logic, and this way, they prove that the satisfiability/entailment problem can be decided in NP/co-NP. Furthermore, they consider the problem of combining SL formulas with constraints on data using the Nelson-Oppen theory combination framework. Adding constraints on data to SL formulas is considered also in Qiu *et al* [20].

A fragment of SL covering overlaid nested lists was considered in our previous work [10]. Compared with it, we currently do not consider overlaid lists, but we have enlarged the set of inductively-defined predicates to allow for nesting of cyclic lists and doubly linked lists (DLLs). We also provide a novel and more efficient TA-based procedure for checking simple entailments.

Brotherston *et al* [4] define a generic automated theorem prover relying on the notion of cyclic proofs and instantiate it to prove entailments in a fragment of SL with inductive definitions and disjunctions more general than what we consider here. However, they do not provide a fragment for which completeness is guaranteed. Iosif *et al* [13] also introduce a decidable fragment of SL that can describe more complex data structures than those considered here, including, e.g., trees with parent pointers or trees with linked leaves. However, [13] reduces the entailment problem to MSO on graphs with a bounded tree width, resulting in a multiply-exponential complexity.

The recent work [12] considers a more restricted fragment than [13], incomparable with ours. The work proposes a more practical, purely TA-based decision procedure, which reduces the entailment problem to *language inclusion* on TAs, establishing EXPTIME-completeness of the considered fragment. Our decision procedure deals with the boolean structure of SL formulas using SAT solvers, thus reducing the entailment problem to the problem of entailment between a formula and an atom. Such simpler entailments are then checked using a polynomial semi-decision procedure based on the *membership problem* for TAs. The approach of [12] can deal with various forms of trees and with entailment of structures with skeletons based on different selectors (e.g., DLLs viewed from the beginning and DLLs viewed from the end). On the other hand, it currently cannot deal with structures of zero length and with some forms of structure concatenation (such as concatenation of two DLL segments), which we can handle.

10 Conclusion

We proposed a novel (semi-)decision procedure for a fragment of SL with inductive predicates describing various forms of lists (singly or doubly linked, nested, circular, with skip links, etc.). The procedure is compositional in that it reduces the given entailment query to a set of simpler queries between a formula and an atom. For solving them, we proposed a novel reduction to testing membership of a tree derived from the formula in the language of a TA derived from a predicate. We implemented the procedure, and our experiments show that it not only has a favourable theoretical complexity, but that it also efficiently handles practical verification conditions.

In the future, we plan to investigate extensions of our approach to formulas with a more general boolean structure or using more general inductive definitions. Concerning the latter, we plan to investigate whether some ideas from [12] could be used to extend our decision procedure for entailments between formulas and atoms. From a practical point of view, apart from improving the implementation of our procedure, we plan to integrate it into a complete program analysis framework.

Acknowledgement. This work was supported by the Czech Science Foundation (project 14-11384S), the BUT FIT projects FIT-S-12-1 and FIT-S-14-2486, and the EU/Czech IT4Innovations Centre of Excellence project CZ.1.05/1.1.00/02.0070.

References

1. Josh Berdine, Cristiano Calcagno, and Peter W. O’Hearn. A decidable fragment of separation logic. In *Proc. of FSTTCS’04*, volume 3328 of *LNCS*, pages 97–109. Springer, 2005.
2. Josh Berdine, Cristiano Calcagno, and Peter W. O’Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In *Proc. of FMCO’05*, volume 4111 of *LNCS*, pages 115–137. Springer, 2006.
3. James Brotherston, Carsten Fuhs, Nikos Gorogiannis, and Juan Navarro Pérez. A decision procedure for satisfiability in separation logic with inductive predicates. In *Proceedings of CSL-LICS*. ACM, 2014. To appear.
4. James Brotherston, Nikos Gorogiannis, and Rasmus Lerchedahl Petersen. A generic cyclic theorem prover. In *APLAS*, volume 7705 of *LNCS*, pages 350–367. Springer, 2012.

5. Cristiano Calcagno, Hongseok Yang, and Peter W. O'Hearn. Computability and complexity results for a spatial assertion language for data structures. In *Proc. of FSTTCS'01*, volume 2245 of *LNCS*, pages 108–119, 2001.
6. Wei-Ngan Chin, Cristina David, Huu Hai Nguyen, and Shengchao Qin. Automated verification of shape, size and bag properties via user-defined predicates in separation logic. *Sci. Comput. Program.*, 77(9):1006–1036, 2012.
7. Byron Cook, Christoph Haase, Joël Ouaknine, Matthew J. Parkinson, and James Worrell. Tractable reasoning in a fragment of separation logic. In *Proc. of CONCUR'11*, volume 6901 of *LNCS*, pages 235–249. Springer, 2011.
8. Constantin Enea, Ondřej Lengál, Mihaela Sighireanu, and Tomáš Vojnar. Compositional entailment checking for a fragment of separation logic. Technical Report FIT-TR-2014-01, FIT BUT, 2014. <http://www.fit.vutbr.cz/~ilengal/pub/FIT-TR-2014-01.pdf>.
9. Constantin Enea, Ondřej Lengál, Mihaela Sighireanu, and Tomáš Vojnar. SPEN, 2014. <http://www.liafa.univ-paris-diderot.fr/spen>.
10. Constantin Enea, Vlad Saveluc, and Mihaela Sighireanu. Compositional invariant checking for overlaid and nested linked lists. In *Proc. of ESOP'13*, volume 7792 of *LNCS*, pages 129–148. Springer, 2013.
11. Alexey Gotsman, Josh Berdine, and Byron Cook. Precision and the conjunction rule in concurrent separation logic. *Electronic Notes in Theoretical Computer Science*, 276:171–190, 2011.
12. Radu Iosif, Adam Rogalewicz, and Tomáš Vojnar. Deciding entailments in inductive separation logic with tree automata. In *Proc. of ATVA'14*, *LNCS*, 2014. Technical Report available at <http://arxiv.org/pdf/1402.2127v2.pdf>.
13. Radu Iosif, Adam Rogalewicz, and Jiří Šimáček. The tree width of separation logic with recursive definitions. In *Proc. of CADE-24*, volume 7898 of *LNCS*, pages 21–38. Springer, 2013.
14. S. Ishtiaq and P.W. O'Hearn. BI as an assertion language for mutable data structures. In *POPL*, pages 14–26. ACM, 2001.
15. Ondřej Lengál, Jiří Šimáček, and Tomáš Vojnar. VATA: A library for efficient manipulation of non-deterministic tree automata. In *Proc. of TACAS'12*, volume 7214 of *LNCS*, pages 79–94. Springer, 2012.
16. J.A. Navarro Pérez and A. Rybalchenko. Separation logic + superposition calculus = heap theorem prover. In *Proc. of PLDI'11*, pages 556–566. ACM, 2011.
17. Juan Antonio Navarro Pérez and Andrey Rybalchenko. Separation logic modulo theories. In *APLAS*, volume 8301 of *LNCS*, pages 90–106. Springer, 2013.
18. Ruzica Piskac, Thomas Wies, and Damien Zufferey. Automating separation logic using SMT. In *Proc. of CAV'13*, volume 8044 of *LNCS*, pages 773–789. Springer, 2013.
19. Ruzica Piskac, Thomas Wies, and Damien Zufferey. Automating separation logic with trees and data. In *Proc. of CAV'14*, volume 8559 of *LNCS*, pages 711–728. Springer, 2014.
20. Xiaokang Qiu, Pranav Garg, Andrei Stefanescu, and Parthasarathy Madhusudan. Natural proofs for structure, data, and separation. In *PLDI*, pages 231–242. ACM, 2013.
21. John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proc. of LICS'02*, pages 55–74. IEEE, 2002.