

An Efficient Decision Procedure for Imperative Tree Data Structures ^{*}

Thomas Wies¹, Marco Muñoz², and Viktor Kuncak³

¹ Institute of Science and Technology (IST), Austria
wies@ist.ac.at

² University of Freiburg, Germany
muniz@informatik.uni-freiburg.de

³ EPFL, Switzerland
viktor.kuncak@epfl.ch

Abstract. We present a new decidable logic called TREX for expressing constraints about imperative tree data structures. In particular, TREX supports a transitive closure operator that can express reachability constraints, which often appear in data structure invariants. We show that our logic is closed under weakest precondition computation, which enables its use for automated software verification. We further show that satisfiability of formulas in TREX is decidable in NP. The low complexity makes it an attractive alternative to more expensive logics such as monadic second-order logic (MSOL) over trees, which have been traditionally used for reasoning about tree data structures.

1 Introduction

This paper introduces a new decision procedure for reasoning about imperative manipulations of tree data structures. Our logic of trees with reachability expressions (TREX) supports reasoning about reachability in trees and a form of quantification, which enables its use for expressing invariants of tree data structures, including the tree property itself. Despite the expressive power of the logic, we exhibit a non-deterministic polynomial-time decision procedure for its satisfiability problem, showing that TREX is NP-complete. Our development is directly motivated by our experience with verifying tree data structures in the Jahob verification system [15, 18, 21] in which we used the MONA decision procedure [11] for MSOL over trees. Although MONA contributed great expressive power to our specification language and, in our experience, works well for programs that manipulate lists, there were many tree-manipulating programs whose verification failed due to MONA running out of resources. It was thus a natural goal to identify a logic that suits our needs, but can be decided much more efficiently.

There are other expressive logics supporting reachability but with lower complexity than MSOL [4, 7, 10, 20]. We did not find them suitable as a MONA alternative, for several reasons. First, we faced difficulties in the expressive power: some of the logics can only reason about sets but not individual objects, others have tree model property and

^{*} An extended version of this paper including proofs of the key lemmas is available as a technical report [19].

thus cannot detect violations of the tree invariants. Moreover, the complexity of these logics is still at least EXPTIME, and their decision procedures are given in terms of automata-theoretic techniques or tableaux procedures, which can be difficult to combine efficiently with existing SMT solvers. Similarly, the logic of reachable patterns [20] is decidable through a highly non-trivial construction, but the complexity is at least NEXPTIME, as is the complexity of the Bernays-Schönfinkel Class with Datalog [5]. The logic [2] can express nested list structures of bounded nesting along with constraints on data fields and numerical constraints on paths, but cannot express constraints on arbitrary trees. On the other hand, TREX does not support reasoning on data fields; although such an extension is in principle possible. Other approaches generate induction scheme instances to prove transitive closure properties in general graphs [14]. While this strategy can succeed for certain examples, it gives neither completeness nor complexity guarantees, and suffers from the difficulties of first-order provers in handling transitive relations. Tree automata with size constraints can express properties such as the red-black tree invariant [8]. However, this work does not state the complexity of the reasoning task and the presented automata constructions appear to require running time beyond NP. Regular tree model checking with abstraction has yielded excellent results so far [3] and continues to improve, but has so far not resulted in a logic whose complexity is in NP, which we believe to be an important milestone.

The primary inspiration for our solution came from the efficient SMT-based techniques for reasoning about list structures [13], as well as the idea of viewing single-parent heaps as duals of lists [1]. However, there are several challenges in relying on this immediate inspiration. For integration with other decision procedures, as well as for modular reasoning with preconditions and postconditions, it was essential to obtain a logic and not only a finite-model property for the analysis of systems as in [1]. Furthermore, the need to support imperative updates on *trees* led to technical challenges that are very different than those of [13]. To address these challenges, we introduced a reachability predicate that is parametrized by a carefully chosen class of formulas to control the reachability relation. We show that the resulting logic of trees is closed under weakest preconditions with respect to imperative heap updates, which makes it suitable for expressing verification conditions in imperative programs. We devised a four-step decision procedure that contains formula transformations and ultimately reduces to a Ψ -local theory extension [9, 16]. Consequently, our logic can be encoded using a quantifier instantiation recipe within an SMT solver. We have encoded the axiomatization of TREX in Jahob and used Z3 [6] with a default instantiation strategy to verify tree and list manipulating programs. We have obtained verification times of around 1s, reducing the running times by two orders of magnitude compared to MONA.

Motivating Example. We next show how to use our decision procedure to verify functional correctness of a Java method that manipulates a binary tree data structure.

Fig. 1 shows a fragment of Java code for insertion into a binary search tree, factored out into a separate `insertLeftOf` method. In addition to Java statements, the example in Fig. 1 contains preconditions and postconditions, written in the notation of the Jahob verification system [12, 15, 17, 18, 21].

The search tree has fields `(l, r)` that form a binary tree, and field `p`, which for each node in the tree points to its parent (or null, if the node is the root of the tree). This

```

class Node {Node l, r, p;}
class Tree {
  private Node root;
  invariant "ptree p [l,r]"; invariant "p root = null";
  private specvar content :: objset;
  vardefs "content=={x. root ≠ null ∧ (x,root) ∈ {(x,y). p x = y}*}";
  public void insertLeftOf(Node pos, Node e)
    requires "pos ∈ content ∧ pos ≠ null ∧ l pos = null ∧
             e ∉ content ∧ e ≠ null ∧ p e = null ∧ l e = null ∧ r e = null"
    modifies content,l,p
    ensures "content = old content ∪ {e}"
    { e.p = pos; pos.l = e; }
}

```

Fig. 1. Fragment of insertion into a tree

property is expressed by the first class invariant using the special predicate `ptree`, which takes the parent field and a list of successor fields of the tree structure as arguments. The second invariant expresses that the field `root` points to the root node of the tree. The **vardefs** notation introduces the set `content` denoting the useful content of the tree. Note that if we are given a program that manipulates a tree data structure without explicit parent field then we can always introduce one as a specification variable that is solely used for the purpose of verification. This is possible because the parent field in a tree is uniquely determined by the successor fields.

The `insertLeftOf` method is meant to be invoked when the insertion procedure has traversed the tree and found a node `pos` that has no left child. The node `e` then becomes the new left child of `pos`. Our system checks that after each execution of the method `insertLeftOf` the specified class invariants still hold and that its postcondition is satisfied. The postcondition states that the node `e` has been properly inserted into the tree.

The full verification condition of method `insertLeftOf` can be expressed in our logic. Figure 2 shows one of the subgoals of this verification condition. It expresses that after execution of method `insertLeftOf` the heap graph projected to field `p` is still acyclic. This is a subgoal for checking that the `ptree` invariant is preserved by method `insertLeftOf`. Note that our logic supports field update expressions `upd(p, e, pos)` so that we can express the verification condition directly in the logic. Note further that the precondition stating that the `ptree` invariant holds at entry to the method is not explicitly part of the verification condition. It is implicit in the semantics of our logic.

Our logic also supports reasoning about forward reachability $\langle l, r \rangle^*$ in the trees (i.e., transitive closure of the successor fields rather than the parent field) and quantification over sets of reachable objects. The latter is used, e.g., to prove the postcondition of method `insertLeftOf` stating that the node `e` was properly inserted and that no elements have been removed from the tree.

While we only consider a logic of *binary* trees in this paper; the generalization to trees of arbitrary finite arity is straightforward. In particular, an acyclic doubly-linked list is a special case of a tree with parent pointers, so reasoning about such structures is also supported by our decision procedure.

$$\begin{aligned}
& p(\text{root}) = \text{null} \wedge \text{root} \neq \text{null} \wedge \langle p \rangle^*(\text{pos}, \text{root}) \wedge \neg \langle p \rangle^*(e, \text{root}) \wedge \\
& e \neq \text{null} \wedge p(e) = \text{null} \wedge l(e) = \text{null} \wedge r(e) = \text{null} \\
& \rightarrow (\forall z. \langle \text{upd}(p, e, \text{pos}) \rangle^*(z, \text{null}))
\end{aligned}$$

Fig. 2. Verification condition expressing that, after execution of method `insertLeftOf`, the heap graph projected to field `p` is still acyclic

$$\begin{aligned}
& p(\text{root}) = \text{null} \wedge \text{root} \neq \text{null} \wedge \langle p \rangle^*(\text{pos}, \text{root}) \wedge \neg \langle p \rangle^*(e, \text{root}) \wedge e \neq \text{null} \wedge p(e) = \text{null} \wedge \\
& l(e) = \text{null} \wedge r(e) = \text{null} \wedge \neg (\forall z. \langle p \rangle^*_{(x \neq e)}(z, \text{null}) \vee \langle p \rangle^*(z, e) \wedge \langle p \rangle^*_{(x \neq e)}(\text{pos}, \text{null}))
\end{aligned}$$

Fig. 3. Negated verification condition from Fig. 2 after function update elimination

2 Decision Procedure Through an Example

We consider the negation of the verification condition shown in Figure 2, which is unsatisfiable in tree structures. Our decision procedure is described in Section 5 and proceeds in four steps.

The first step (Section 5.1) is to eliminate all function update expressions in the formula. The result of this step is shown in Figure 3. Our logic supports so called *constrained reachability expressions* of the form $\langle p \rangle_Q^*$ where Q is a binary predicate over dedicated variables x, y . The semantics of this predicate is that $\langle p \rangle_Q^*(u, v)$ holds iff there exists a p -path connecting u and v and between every consecutive nodes w_1, w_2 on this path, $Q(w_1, w_2)$ holds. Using these constrained reachability expressions we can reduce reachability expressions over updated fields to reachability expressions over the non-updated fields, as shown in the example. This elimination even works for updates of successor functions below forward reachability expressions of the form $\langle l, r \rangle^*$.

The second step (Section 5.2) eliminates all forward reachability constraints over fields l, r from the formula and expresses them in terms of the relation $\langle p \rangle^*$. Since there are no such constraints in our formula, we immediately proceed to Step 3.

The third step (Section 5.3) reduces the formula to a formula in first-order logic, whose finite models are exactly the models of the formula from the previous step, which is still expressed in TREX. For the purpose of the reduction, all occurrences of the reachability relation $\langle p \rangle^*$ are replaced by a binary predicate symbol P , which is then axiomatized using universally quantified first-order axioms so that $\langle p \rangle^*$ and P coincide in all finite models. All remaining reachability constraints are of the form $\langle p \rangle_Q^*$. We can express these constraints in terms of P by introducing a unary function bp_Q (called *break point function*) that maps each node u to the first p -reachable node v of u for which $Q(v, p(v))$ does not hold, i.e., $bp_Q(u)$ marks the end of the segment of nodes w that satisfy $\langle p \rangle_Q^*(u, w)$. The function bp_Q can be axiomatized in terms of P and Q . Figure 4 shows the resulting formula (including only the necessary axioms for proving unsatisfiability of the formula).

The fourth step (Section 5.4) computes prenex normal form and skolemizes remaining top-level existential quantifiers. Then we add additional axioms that ensure Ψ -locality of the universally quantified axioms in the formula obtained from Step 3. The key property of the resulting formula is that its universal quantifiers can be instan-

$$\begin{aligned}
 & p(\text{root}) = \text{null} \wedge \text{root} \neq \text{null} \wedge P(\text{pos}, \text{root}) \wedge \neg P(\text{e}, \text{root}) \wedge \\
 & e \neq \text{null} \wedge p(\text{e}) = \text{null} \wedge l(\text{e}) = \text{null} \wedge r(\text{e}) = \text{null} \wedge \\
 & \neg(\forall z. P(z, \text{null}) \wedge P(\text{null}, bp_{(x \neq e)}(z)) \vee P(z, \text{e}) \wedge P(\text{pos}, \text{null}) \wedge P(\text{null}, bp_{(x \neq e)}(\text{pos}))) \wedge \\
 & (\forall z. P(z, \text{null})) \wedge (\forall z. P(z, z)) \wedge (\forall wz. P(w, z) \wedge P(z, w) \rightarrow z = w) \wedge \\
 & (\forall vwz. P(v, w) \wedge P(v, z) \rightarrow P(w, z) \vee P(z, w)) \wedge \\
 & (\forall wz. P(w, z) \rightarrow w = z \vee P(p(w), z)) \wedge \\
 & (\forall z. P(z, bp_{(z \neq e)}(z))) \wedge (\forall z. bp_{(x \neq e)}(z) \neq e \rightarrow bp_{(x \neq e)}(z) = \text{null}) \wedge \\
 & (\forall wz. P(w, z) \wedge P(z, bp_{(x \neq e)}(w)) \rightarrow z \neq e \vee z = bp_{(x \neq e)}(w)) \wedge \dots
 \end{aligned}$$

Fig. 4. Negated verification condition from Figure 2 after the reduction step to first-order logic. Only the axioms that are necessary for proving unsatisfiability of the formula are shown.

tiated finitely many times with terms syntactically derived from the terms within the formula. The result is an equisatisfiable quantifier-free formula, which can be handled by the SMT solver’s congruence closure and the SAT solver.

3 Preliminaries

In the following, we define the syntax and semantics of formulas. We further recall the notions of partial structures and Ψ -local theories as defined in [9].

Sorted logic. We present our problem in sorted logic with equality. A *signature* Σ is a tuple (S, Ω) , where S is a countable set of sorts and Ω is a countable set of function symbols f with associated arity $n \geq 0$ and associated sort $s_1 \times \dots \times s_n \rightarrow s_0$ with $s_i \in S$ for all $i \leq n$. Function symbols of arity 0 are called *constant symbols*. In this paper we will only consider signatures with sorts $S = \{\text{bool}, \text{node}\}$ and the dedicated equality symbol $= \in \Omega$ of sort $\text{node} \times \text{node} \rightarrow \text{bool}$. Note that we generally treat predicate symbols of sort s_1, \dots, s_n as function symbols of sort $s_1 \times \dots \times s_n \rightarrow \text{bool}$. Terms are built as usual from the function symbols in Ω and (sorted) variables taken from a countably infinite set X that is disjoint from Ω . A term t is said to be *ground*, if no variable appears in t . We denote by $\text{Terms}(\Sigma)$ the set of all ground Σ -terms.

A Σ -atom A is a Σ -term of sort bool . We use infix notation for atoms built from the equality symbol. A Σ -formula F is defined via structural recursion as either one of A , $\neg F_1$, $F_1 \wedge F_2$, or $\forall x : s. F_1$, where A is a Σ -atom, F_1 and F_2 are Σ -formulas, and $x \in X$ is a variable of sort $s \in S$. In formulas appearing in this paper we will only ever quantify over variables of sort node , so we typically drop the sort annotation. We use syntactic sugar for Boolean constants (\top , \perp), disjunctions ($F_1 \vee F_2$), implications ($F_1 \rightarrow F_2$), and existential quantification ($\exists x. F_1$). For a finite index set \mathcal{I} and Σ -formulas F_i , for all $i \in \mathcal{I}$, we write $\bigwedge_{i \in \mathcal{I}} F_i$ for the conjunction of the F_i (respectively, \top if \mathcal{I} is empty) and similarly $\bigvee_{i \in \mathcal{I}} F_i$ for their disjunction. We further write $F[x_1 := t_1, \dots, x_n := t_n]$ for the simultaneous substitutions of the free variables x_i appearing in F by the terms t_i . We define literals and clauses as usual. A clause C is called *flat* if no term that occurs in C below a predicate symbol or the symbol $=$ contains nested function symbols. A clause C is called *linear* if (i) whenever a variable occurs in two non-variable terms in

C that do not start with a predicate or the equality symbol, the two terms are identical, and if (ii) no such term contains two occurrences of the same variable.

Total and partial structures. Given a signature $\Sigma = (S, \Omega)$, a *partial Σ -structure* α is a function that maps each sort $s \in S$ to a non-empty set $\alpha(s)$ and each function symbol $f \in \Omega$ of sort $s_1 \times \dots \times s_n \rightarrow s_0$ to a partial function $\alpha(f) : \alpha(s_1) \times \dots \times \alpha(s_n) \rightarrow \alpha(s_0)$. If α is understood, we write just t instead of $\alpha(t)$ whenever this is not ambiguous. We assume that all partial structures interpret the sort `bool` by the two-element set of Booleans $\{0, 1\}$. We therefore call $\alpha(\text{node})$ the *universe* of α and often identify $\alpha(\text{node})$ and α . We further assume that all structures α interpret the symbol $=$ by the equality relation on $\alpha(\text{node})$. A partial structure α is called *total structure* or simply *structure* if it interprets all function symbols by total functions. For a Σ -structure α where Σ extends a signature Σ_0 with additional sorts and function symbols, we write $\alpha|_{\Sigma_0}$ for the Σ_0 -structure obtained by restricting α to Σ_0 .

Given a total structure α and a *variable assignment* $\beta : X \rightarrow \alpha(S)$, the evaluation $\llbracket t \rrbracket_{\alpha, \beta}$ of a term t in α, β is defined as usual. For a ground term t we typically write just $\llbracket t \rrbracket_{\alpha}$. A quantified variable of sort s ranges over all elements of $\alpha(s)$. From the interpretation of terms the notions of satisfiability, validity, and entailment of atoms, formulas, clauses, and sets of clauses in total structures are derived as usual. In particular, we use the standard interpretations for propositional connectives of classical logic. We write $\alpha, \beta \models F$ if α satisfies F under β where F is a formula, a clause, or a set of clauses. Similarly, we write $\alpha \models F$ if F is valid in α . In this case we also call α a *model* of F . The interpretation $\llbracket t \rrbracket_{\alpha, \beta}$ of a term t in a partial structure α is as for total structures, except that if $t = f(t_1, \dots, t_n)$ for $f \in \Omega$ then $\llbracket t \rrbracket_{\alpha, \beta}$ is undefined if either $\llbracket t_i \rrbracket_{\alpha, \beta}$ is undefined for some i , or $(\llbracket t_1 \rrbracket_{\alpha, \beta}, \dots, \llbracket t_n \rrbracket_{\alpha, \beta})$ is not in the domain of $\alpha(f)$. We say that a partial structure α *weakly satisfies* a literal L under β , written $\alpha, \beta \models_w L$, if (i) L is an atom A and either $\llbracket A \rrbracket_{\alpha, \beta} = 1$ or $\llbracket A \rrbracket_{\alpha, \beta}$ is undefined, or (ii) L is a negated atom $\neg A$ and either $\llbracket A \rrbracket_{\alpha, \beta} = 0$ or $\llbracket A \rrbracket_{\alpha, \beta}$ is undefined. The notion of weak satisfiability is extended to clauses and sets of clauses as for total structures. A clause C (respectively, a set of clauses) is *weakly valid* in a partial structure α if α weakly satisfies α for all variable assignments β . We then call α a *weak partial model* of C .

Ψ -local theories. The following definition is a particular special case of the more general notion of Ψ -local theory extensions. For the general definitions of local theory extensions, respectively, Ψ -local theory extensions, we direct the reader to [9, 16].

Let $\Sigma = (S, \Omega)$ be a signature. A *theory* \mathcal{T} for a signature Σ is simply a set of Σ -formulas. We consider theories $\mathcal{T}(\mathcal{K})$ defined as a set of Σ -formulas that are consequences of a given set of clauses \mathcal{K} . We call \mathcal{K} the *axioms* of the theory $\mathcal{T}(\mathcal{K})$ and we often identify \mathcal{K} and $\mathcal{T}(\mathcal{K})$. In the following, when we refer to a set of ground clauses G , we assume they are over the signature $\Sigma^c = (S, \Omega \cup \Omega_c)$ where Ω_c is a set of new constant symbols. For a set of clauses \mathcal{K} , we denote by $\text{st}(\mathcal{K})$ the set of all ground subterms that appear in \mathcal{K} . Let Ψ be a function associating with a set of (universally quantified) clauses \mathcal{K} and a set of ground terms T a set $\Psi(\mathcal{K}, T)$ of ground terms such that (i) all ground subterms in \mathcal{K} and T are in $\Psi(\mathcal{K}, T)$; (ii) for all sets of ground terms T, T' if $T \subseteq T'$ then $\Psi(\mathcal{K}, T) \subseteq \Psi(\mathcal{K}, T')$; (iii) Ψ is a closure operation, i.e., for all sets of ground terms T , $\Psi(\mathcal{K}, \Psi(\mathcal{K}, T)) \subseteq \Psi(\mathcal{K}, T)$. (iv) Ψ is compatible with any map h between constants, i.e., for any map $h : \Omega_c \rightarrow \Omega_c$, $\Psi(\mathcal{K}, \bar{h}(T)) = \bar{h}(\Psi(\mathcal{K}, T))$ where

\bar{h} is the unique extension of h to terms. Let $\mathcal{K}[\Psi(\mathcal{K}, G)]$ be the set of instances of \mathcal{K} in which all terms are in $\Psi(\mathcal{K}, \text{st}(G))$, which here will be denoted by $\Psi(\mathcal{K}, G)$. We say that \mathcal{K} is Ψ -local if it satisfies condition (Loc^Ψ) :

(Loc^Ψ) For every finite set of ground clauses G , $\mathcal{K} \cup G \models \perp$ iff $\mathcal{K}[\Psi(\mathcal{K}, G)] \cup G$ has no weak partial model in which all terms in $\Psi(\mathcal{K}, G)$ are defined.

4 TREX: Logic of Trees with Reachability Expressions

We now formally define the formulas of our logic of trees with reachability expressions (TREX), whose satisfiability we study. For simplifying the exposition in the remainder of this paper, we restrict ourselves to binary trees. The decidability and complexity result carries over to trees of arbitrary finite arity in a straightforward manner.

Syntax of TREX formulas. Figure 5 defines the TREX formulas. A TREX formula is a propositional combination of atomic formulas. An atomic formula is either an equality between terms, a reachability expression, or a restricted quantified formula. A term t is either a constant $c \in \Gamma$ or a function term f applied to a term t . The set of constants Γ is an arbitrary countably infinite set of symbols disjoint from all other symbols used in the syntax of formulas. However, we assume that Γ contains the special constant symbol null. A function term is either one of the function symbols l, r (standing for the two successor functions of a tree), and p (standing for the parent function of a tree), or an update $\text{upd}(f, t_1, t_2)$ of a function term f . In the latter case we call t_1 the *index* of the update and t_2 the *target*. A forward reachability expression relates two terms by a relation $\langle f_l, f_r \rangle_Q^*$ where f_l and f_r are the possibly updated successor functions and Q is a predicate built from boolean combinations of equalities between constants and the dedicated variables x and y . The syntactic restrictions on Q ensure that if one computes the disjunctive normal form of Q then the resulting formula will contain a disjunct that is a conjunction of disequalities between constants and variables. A backward reachability expression is similar but refers to the possibly updated parent function. We call the relations $\langle f_l, f_r \rangle_Q^*$ *descendant relations* and the relations $\langle f_p \rangle_Q^*$ *ancestor relations*. Finally, the formulas below restricted quantified formulas are almost like TREX formulas, except that the quantified variable may only appear at particular positions below function symbols and only as arguments of ancestor relations. For a predicate Q and terms t_1, t_2 , we typically write $Q(t_1, t_2)$ for the formula $Q[x := t_1, y := t_2]$. Finally, we simply write p^* as a shorthand for $\langle p \rangle_\top^*$.

Semantics of TREX formulas. TREX formulas are interpreted over finite forests of finite binary trees. We formally define these forests as first-order structures $\alpha_{\mathcal{F}}$ over the signature $\Sigma_{\mathcal{F}}$ of constant symbols Γ and the unary function symbols l, r and p . To this end define the set of *tree nodes* \mathcal{N} as the set of strings consisting of the empty string ϵ and all strings over alphabet $\mathbb{N} \cup \{L, R\}$ that satisfy the regular expression $\mathbb{N} \cdot (L \mid R)^*$, i.e., we enumerate the trees comprising a forest by attaching a natural number to the nodes in each tree. A *forest* $\alpha_{\mathcal{F}}$ is then a structure whose universe is a finite prefixed-closed subset of tree nodes. The interpretation of the special constant symbol $\text{null} \in \Gamma$ and the function symbols l, r , and p are determined by the universe of $\alpha_{\mathcal{F}}$ as in Figure 6. The remaining constant symbols in Γ may be interpreted by any tree node in $\alpha_{\mathcal{F}}$. Let \mathcal{F} be the set of all forests and let $\mathcal{M}_{\mathcal{F}}$ be the set of all first-order structures over signature

$ \begin{aligned} F &::= A \mid F \wedge F \mid \neg F \\ A &::= t = t \mid \langle f_l, f_r \rangle_Q^*(t, t) \mid \langle f_p \rangle_Q^*(t, t) \mid F \forall \\ t &::= c \mid f(t) \\ f &::= f_l \mid f_r \mid f_p \\ f_l &::= \text{upd}(f_l, t, t) \mid l \\ f_r &::= \text{upd}(f_r, t, t) \mid r \\ f_p &::= \text{upd}(f_p, t, t) \mid p \\ Q &::= v = c \rightarrow R \mid Q \wedge Q \\ R &::= t_R = t_R \mid R \wedge R \mid \neg R \\ t_R &::= v \mid c \end{aligned} $	$ \begin{aligned} F \forall &::= \forall z. G_{\text{in}} \\ G_{\text{in}} &::= f(z) = t \rightarrow G_{\text{in}} \mid F_{\text{in}} \\ F_{\text{in}} &::= A_{\text{in}} \mid F_{\text{in}} \wedge F_{\text{in}} \mid \neg F_{\text{in}} \\ A_{\text{in}} &::= t_{\text{in}} = t_{\text{in}} \mid \langle f_p \rangle_Q^*(t_{\text{in}}, t_{\text{in}}) \\ t_{\text{in}} &::= z \mid t \end{aligned} $ <p>terminals:</p> <ul style="list-style-type: none"> $c \in \Gamma$ - constant symbol l, r, p - function symbols $v \in \{x, y\}$ - dedicated variable $z \in X$ - variable
--	---

Fig. 5. Logic of trees with reachability TREX

$$\begin{aligned}
\alpha_{\mathcal{F}}(\text{null}) &= \epsilon & \alpha_{\mathcal{F}}(l)(n) &= \begin{cases} nL & \text{if } nL \in \alpha_{\mathcal{F}} \\ \epsilon & \text{otherwise} \end{cases} & \alpha_{\mathcal{F}}(r)(n) &= \begin{cases} nR & \text{if } nR \in \alpha_{\mathcal{F}} \\ \epsilon & \text{otherwise} \end{cases} \\
\alpha_{\mathcal{F}}(p)(n) &= \begin{cases} n' & \text{if } n = n's \text{ for some } s \in \mathbb{N} \cup \{L, R\} \text{ and } n' \in \alpha_{\mathcal{F}} \\ \epsilon & \text{otherwise} \end{cases}
\end{aligned}$$

Fig. 6. Semantics of functions and constants in the forest model.

$\Sigma_{\mathcal{F}}$ that are isomorphic to some structure in \mathcal{F} . We extend the term *forest* to all the structures in $\mathcal{M}_{\mathcal{F}}$.

For defining the semantics of TREX formulas, let $\alpha_{\mathcal{F}} \in \mathcal{M}_{\mathcal{F}}$. We only explain the interpretation of terms, function terms, and reachability expressions in detail, the remaining constructs are interpreted as expected. The notions of satisfiability, entailment, etc. for TREX formulas are defined as in Section 3.

The interpretation of terms and function terms in $\alpha_{\mathcal{F}}$ under a variable assignment β recursively extend the interpretation of $\Sigma_{\mathcal{F}}$ -terms as follows:

$$\begin{aligned}
\llbracket f \rrbracket_{\alpha_{\mathcal{F}}, \beta} &\stackrel{\text{def}}{=} \alpha_{\mathcal{F}}(f), \text{ for } f \in \{l, r, p\} \\
\llbracket \text{upd}(f, t_1, t_2) \rrbracket_{\alpha_{\mathcal{F}}, \beta} &\stackrel{\text{def}}{=} \llbracket f \rrbracket_{\alpha_{\mathcal{F}}, \beta} [\llbracket t_1 \rrbracket_{\alpha_{\mathcal{F}}, \beta} \mapsto \llbracket t_2 \rrbracket_{\alpha_{\mathcal{F}}, \beta}] \\
\llbracket f(t) \rrbracket_{\alpha_{\mathcal{F}}, \beta} &\stackrel{\text{def}}{=} \llbracket f \rrbracket_{\alpha_{\mathcal{F}}, \beta} (\llbracket t \rrbracket_{\alpha_{\mathcal{F}}, \beta})
\end{aligned}$$

In order to define the semantics of reachability expressions compactly, we write $\langle Fn \rangle_Q^*(t_1, t_2)$ for either a forward reachability expression $\langle f_l, f_r \rangle_Q^*(t_1, t_2)$ or a backward reachability expression $\langle f_p \rangle_Q^*(t_1, t_2)$. In the first case, the meta variable Fn denotes the set of function terms $\{f_l, f_r\}$ and in the second case the set $\{f_p\}$. We also use the notation $\langle f, Fn \rangle_Q^*(t_1, t_2)$, which denotes: $\langle f_p \rangle_Q^*(t_1, t_2)$ if $f = f_p$ and $Fn = \emptyset$, and denotes $\langle f_l, f_r \rangle_Q^*(t_1, t_2)$ if $Fn = \{f_r\}$ and $f = f_l$ or $Fn = \{f_l\}$ and $f = f_r$. A reachability expression $\langle Fn \rangle_Q^*(t_1, t_2)$ expresses that the node defined by t_2 can be obtained from the node defined by t_1 , by successively applying the functions defined by the function terms in Fn , where at each step Q holds between the current node and its image. Formally, we define the binary predicate $R_{Q, Fn}$ by the formula $(\bigvee_{f \in Fn} f(x) = y) \wedge Q$ and interpret the reachability relation $\langle Fn \rangle_Q^*$ as the reflexive

transitive closure of $R_{Q,Fn}$:

$$\llbracket \langle Fn \rangle_Q^* \rrbracket_{\alpha_{\mathcal{F}}, \beta} \stackrel{\text{def}}{=} \left\{ (u, v) \in \alpha_{\mathcal{F}} \times \alpha_{\mathcal{F}} \mid \llbracket R_{Q,Fn} \rrbracket_{\alpha_{\mathcal{F}}, \beta[x \mapsto u, y \mapsto v]} \right\}^*$$

The interpretation of $\langle Fn \rangle_Q^*(t_1, t_2)$ is then defined as expected.

Definition 1 (Satisfiability Problem for TREX). *The satisfiability problem for TREX asks whether, given a TREX formula F , there exists a forest $\alpha_{\mathcal{F}}$ that satisfies F .*

5 Decision Procedure for TREX

The logic TREX is a proper subset of MSOL over finite trees. Thus, decidability of the satisfiability problem for TREX follows from the decidability of MSOL over trees. In fact TREX formulas can be expressed in terms of MSOL formulas with at most two quantifier alternations, which gives a 2-EXPTIME upper-bound for the complexity. In the following, we show that the satisfiability problem for TREX is actually in NP.

For the remainder of this section we fix a TREX formula F_0 . Our decision procedure proceeds in four steps. The first two steps eliminate function updates and forward reachability expressions from F_0 , resulting in equisatisfiable TREX formulas F_1 and then F_2 . In the third step the formula F_2 is reduced to a first-order formula F_3 that has the same finite models as the original formula F . We then use results on local theories [9, 16] to prove a small model property for the obtained formulas. This allows us to use an existing decision procedure to check satisfiability of F_3 in the final step of our algorithm and obtain NP completeness.

5.1 Elimination of Function Updates

We first describe the elimination of function updates from the input formula F_0 . The algorithm that achieves this is as follows:

1. Flatten the index and target terms of function updates in F_0 by exhaustively applying the following rewrite rule:

$$C[\text{upd}(f, i, t)] \rightsquigarrow C[\text{upd}(f, c_i, c_t)] \wedge c_i = i \wedge c_t = t$$
 where i, t are non-constant terms and $c_i, c_t \in \Gamma$ are fresh constant symbols
2. Eliminate function updates in reachability expressions by exhaustively applying the following rewrite rule:

$$C[\langle \text{upd}(f, c_i, c_t), Fn \rangle_Q^*(t_1, t_2)] \rightsquigarrow C[H] \wedge \bigwedge_{f' \in Fn} c_{f'} = f'(c_i)$$
 where the $c_{f'}$ are fresh constant symbols and

$$H \stackrel{\text{def}}{=} \langle f, Fn \rangle_R^*(t_1, t_2) \vee \langle f, Fn \rangle_Q^*(t_1, c_i) \wedge \langle f, Fn \rangle_R^*(c_i, t_2) \wedge Q(c_i, c_t)$$

$$R \stackrel{\text{def}}{=} Q \wedge (x = c_i \rightarrow \bigvee_{f' \in Fn} y = c_{f'})$$
3. Eliminate all remaining function updates by exhaustively applying the following rewrite rule:

$$t_1 = C[\text{upd}(f, c_i, c_t)(t_2)] \rightsquigarrow t_2 = c_i \wedge t_1 = C[c_i] \vee t_2 \neq c_i \wedge t_1 = C[f(t_2)]$$

Note that the exhaustive application of the rule in each of the steps 1. to 3. is guaranteed to terminate. Thus, let F_1 be any of the possible normal form formulas obtained after exhaustive application of these rules to F_0 .

Lemma 2. F_1 is a TREX formula and is equisatisfiable with F_0 .

5.2 Elimination of Descendant Relations

We next describe the second step of our decision procedure, which eliminates all descendant relations from the formula F_1 . The elimination is performed using the following rewrite rule:

$$\langle l, r \rangle_Q^*(s, t) \rightsquigarrow s = t \vee s \neq \text{null} \wedge (\exists z. (l(z) = t \vee r(z) = t) \wedge \langle p \rangle_{Q^{-1}}^*(z, s) \wedge Q(z, t))$$

where $Q^{-1} \stackrel{\text{def}}{=} Q[x := y, y := x]$. Let F_2 be any of the normal form formulas obtained by exhaustively applying this rewrite rule to F_1 .

Lemma 3. F_2 is a TREX formula and is equisatisfiable with F_1 .

5.3 Reduction to First-Order Logic

In the third step of our decision procedure we reduce the formula F_2 obtained after the second step to a formula F_3 in first-order logic. The idea of the reduction is to provide a first-order axiomatization of the unconstrained ancestor relation p^* whose finite models are precisely the forests $\mathcal{M}_{\mathcal{F}}$ defined in Section 4. For this purpose we introduce a fresh binary predicate symbol P representing p^* . The axioms defining P are given in Figure 7. We can then axiomatize each constrained ancestor relation $\langle p \rangle_Q^*$ in terms of p^* . To achieve this we exploit that the relations $\langle p \rangle_Q^*$ can be characterized as follows:

$$\forall xy. \langle p \rangle_Q^*(x, y) \leftrightarrow p^*(x, y) \wedge p^*(y, bp_Q(x)) \quad (1)$$

where bp_Q is the function that maps a node x to the first ancestor z of x such that $Q(z, p(z))$ does not hold (or null if such a node does not exist). We call bp_Q the *break point* function for $\langle p \rangle_Q^*$. The intuition behind the above definition is that for $\langle p \rangle_Q^*(x, y)$ to be true, the break point for the path of ancestor nodes of x must come after y has been reached (respectively, y itself is the break point of x). Note that this definition exploits the fact that forests are acyclic graphs. The axioms defining bp_Q are given in Figure 8.

Formally, the reduction of F_2 to a first-order logic formula F_3 is defined as follows: Let P be a fresh binary predicate symbol and let $F_{3,1}$ be the formula obtained by conjoining F_2 with the axioms shown in Figure 7. Let \mathcal{Q} be the set of predicates Q appearing in reachability expressions $\langle p \rangle_Q^*(t_1, t_2)$ in F_2 . For each $Q \in \mathcal{Q}$, let bp_Q be a fresh unary function symbol. For each $Q \in \mathcal{Q}$, replace all occurrences of the form $\langle p \rangle_Q^*(t_1, t_2)$ in F_2 by $P(t_1, t_2) \wedge P(t_2, bp_Q(t_1))$. Let the result be $F_{3,2}$. Finally, for each $Q \in \mathcal{Q}$, conjoin $F_{3,2}$ with the axioms shown in Figure 8. Let F_3 be the resulting formula and let Σ_P be the extension of the signature $\Sigma_{\mathcal{F}}$ with the symbols P , and bp_Q , for all $Q \in \mathcal{Q}$.

Lemma 4. For every finite Σ_P -model α of the axioms in Figure 7, $\alpha(P) = \alpha(p)^*$ and $\alpha|_{\Sigma_{\mathcal{F}}} \in \mathcal{M}_{\mathcal{F}}$.

Lemma 5. The TREX formula F_2 has a model in $\mathcal{M}_{\mathcal{F}}$ iff the Σ_P -formula F_3 has a finite Σ_P -model.

l-Child : $\mathfrak{p}(l(x)) = x \vee l(x) = \text{null}$	p-Loop : $\mathfrak{p}(x) = x \rightarrow x = \text{null}$
r-Child : $\mathfrak{p}(r(x)) = x \vee r(x) = \text{null}$	NullTerm : $P(x, \text{null})$
Parent : $l(\mathfrak{p}(x)) = x \vee r(\mathfrak{p}(x)) = x \vee \mathfrak{p}(x) = \text{null}$	Ref1 : $P(x, x)$
lr-Diff : $l(x) = y \wedge r(x) = y \rightarrow y = \text{null}$	Trans : $P(x, y) \wedge P(y, z) \rightarrow P(x, z)$
l-Root : $\mathfrak{p}(x) = \text{null} \wedge l(z) = x \rightarrow x = \text{null}$	AntiSym : $P(x, y) \wedge P(y, x) \rightarrow x = y$
r-Root : $\mathfrak{p}(x) = \text{null} \wedge r(z) = x \rightarrow x = \text{null}$	p-Step : $P(x, \mathfrak{p}(x))$
Total : $P(x, y) \wedge P(x, z) \rightarrow P(z, y) \vee P(y, z)$	p-Unfold : $P(x, y) \rightarrow x = y \vee P(\mathfrak{p}(x), y)$

Fig. 7. First-order axioms for the unconstrained ancestor relation \mathfrak{p}^* (represented by the binary predicate symbol P) and the functions l , r , and \mathfrak{p} in a forest

$$\begin{aligned}
 bp_Q\text{-Def1} : P(x, bp_Q(x)) \quad bp_Q\text{-Def2} : Q(bp_Q(x), \mathfrak{p}(bp_Q(x))) \rightarrow bp_Q(x) = \text{null} \\
 bp_Q\text{-Def3} : P(x, y) \wedge P(y, bp_Q(x)) \rightarrow Q(y, \mathfrak{p}(y)) \vee y = bp_Q(x)
 \end{aligned}$$

Fig. 8. First-order axioms defining the break point functions bp_Q

5.4 Ψ -Locality

Now let F_4 be the formula obtained by transforming F_3 into prenex normal form and skolemizing all existential quantifiers. Note that our syntactic restrictions on TREX formulas ensure that there are no alternating quantifiers appearing in the formulas F_0 , F_1 , F_2 , and hence F_3 . So skolemization only introduces additional Skolem constants, but no additional function symbols.

Let C be the set of clauses obtained by transforming F_4 into clausal normal form. Then partition C into sets of ground clauses G and non-ground clauses \mathcal{K}_P in which all terms have been linearized and flattened. The idea is now to define a closure operator Ψ such that condition (Loc $^\Psi$) from Section 3 holds for the particular pair \mathcal{K}_P, G . To ensure that we can extend finite weak partial models of $\mathcal{K}_P[\Psi(\mathcal{K}_P, G)] \cup G$ to finite total models of $\mathcal{K}_P \cup G$, we have to make sure that $\Psi(\mathcal{K}_P, G)$ contains sufficiently many ground terms.

We will define Ψ such that in every finite weak partial model of $\mathcal{K}_P[\Psi(\mathcal{K}_P, G)] \cup G$, both P and the break point functions are already totally defined. However, for this we have to bound the possible values of the break point functions. In fact, each predicate $Q \in \mathcal{Q}$ bounds the possible values that bp_Q can take. Let $\Gamma(Q)$ be the set of constants appearing in Q and let α be a finite total model of \mathcal{K}_P , then for all $u \in \alpha$, $bp_Q(u)$ is one of null , c , $l(c)$, or $r(c)$ for some $c \in \Gamma(Q)$. Thus, for each predicate $Q \in \mathcal{Q}$ define the set of its potential break points $BP(Q)$ as follows. For sets of ground terms T and a k -ary function symbol f , let $f(T)$ be the set of all (properly sorted) ground terms $f(t_1, \dots, t_k)$ for some $t_1, \dots, t_k \in T$. Then define

$$BP(Q) \stackrel{\text{def}}{=} \Gamma(Q) \cup l(\Gamma(Q)) \cup r(\Gamma(Q)) \cup \{\text{null}\}$$

Let further $BP(\mathcal{Q})$ be the union of all sets $BP(Q)$ for $Q \in \mathcal{Q}$. This leads us to our first approximation Ψ_{bp} of Ψ . To this end let $f^i(T)$ be the set $f(T)$ restricted to the terms in which the function symbol f appears at most i times, and let $bp^-(T)$ be the set of ground terms obtained by removing from each ground term in T all appearances of the

$$\begin{aligned} bp_Q\text{-Def4} &: P(x, y) \wedge P(y, bp_Q(x)) \rightarrow bp_Q(x) = bp_Q(y) \\ bp_Q\text{-Def5} &: \bigvee_{t \in BP(Q)} bp_Q(x) = t \end{aligned}$$

Fig. 9. Additional first-order axioms for bounding the break point functions

$$\begin{aligned} fca\text{-Def1} &: P(x, fca(x, y)) & fca\text{-Def2} &: P(y, fca(x, y)) \\ fca\text{-Def3} &: P(x, z) \wedge P(y, z) \rightarrow P(fca(x, y), z) \\ fca\text{-Def4} &: fca(x, y) = w \wedge fca(x, z) = w \wedge fca(y, z) = w \rightarrow x = y \vee x = z \vee y = z \vee w = \text{null} \end{aligned}$$

Fig. 10. Axioms defining the first common ancestor of two nodes in a forest

function symbols $\{ bp_Q \mid Q \in \mathcal{Q} \}$. Then define

$$\begin{aligned} \Psi_0(T) &\stackrel{\text{def}}{=} T \cup \{ p(t) \mid t \in T, \exists t'. t = l(t') \vee t = r(t') \} \cup BP(\mathcal{Q}) \cup p(BP(\mathcal{Q})) \\ \Psi_4(T) &\stackrel{\text{def}}{=} T \cup \bigcup_{Q \in \mathcal{Q}} bp_Q(bp^-(T)) \\ \Psi_5(T) &\stackrel{\text{def}}{=} T \cup P(T) \\ \Psi_{bp}(\mathcal{K}, T) &\stackrel{\text{def}}{=} \Psi_5 \circ \Psi_4 \circ \Psi_0(\text{st}(\mathcal{K}) \cup T) \end{aligned}$$

Let \mathcal{K}_{bp} be the set of clauses obtained from \mathcal{K}_P by adding the linearized and flattened clauses corresponding to the axioms shown in Figure 9. These additional axioms ensure that the interpretation of the break point functions in weak partial models of \mathcal{K}_P are consistent with those in total models of \mathcal{K}_P .

However, the above definition is not yet sufficient to ensure Ψ -locality. Assume that a clause of the form $z = c \vee z = d$ appears in \mathcal{K}_{bp} that results from a restricted quantified formula $\forall z. z = c \vee z = d$ in F_0 . Then this clause imposes an upper bound of 2 on the cardinality of the models of F_4 . We thus have to make sure that for any weak partial model of $\mathcal{K}_{bp}[\Psi_{bp}(\mathcal{K}_{bp}, G)] \cup G$, we can find a total model of the same cardinality. We can ensure that total models of matching cardinality exist by enforcing that every weak partial model already determines the *first common ancestor* of every pair of nodes. We axiomatize the first common ancestor of two nodes by introducing a fresh binary function symbol fca and then adding the axioms shown in Figure 10. Let Σ_{fca} be the signature Σ_P extended with the binary function symbol fca and let \mathcal{K}_{fca} be the set of clauses obtained by adding to \mathcal{K}_{bp} the linearized and flattened clauses corresponding to the axioms in Figure 10. Our second attempt at defining Ψ is then:

$$\begin{aligned} \Psi_3(T) &\stackrel{\text{def}}{=} T \cup fca^1(T) \cup fca^2(T \cup fca^1(T)) \\ \Psi_{fca}(\mathcal{K}, T) &\stackrel{\text{def}}{=} \Psi_5 \circ \Psi_4 \circ \Psi_3 \circ \Psi_0(\text{st}(\mathcal{K}) \cup T) \end{aligned}$$

Unfortunately, the operator Ψ_{fca} is still not good enough to ensure Ψ -locality. Assume that a clause of the form $f(z) = t \rightarrow H$ appears in \mathcal{K}_{fca} that resulted from a restricted quantified formula in F_0 of the form $\forall z. f(z) = t \rightarrow H$ and where f is either one of p , l , or r . Assume that $f = p$. To ensure that this clause remains valid whenever we complete p to a total function in some weak partial model α , we have to ensure that we never have to define $p(u) = t$, for any $u \in \alpha$ for which p is undefined. Consider first the case that in said model t is not null, then we can guarantee that we never have to define $p(u) = t$ by making sure that α is already defined on the ground terms $p(l(t))$

$$\begin{array}{ll}
 \text{Root1} : P(x, y) \rightarrow P(y, \text{root}(x)) \vee y = \text{null} & \text{Root2} : \text{root}(x) = \text{null} \leftrightarrow x = \text{null} \\
 \text{l-Leaf1} : P(\text{lleaf}(x), x) \vee \text{lleaf}(x) = \text{null} & \text{r-Leaf1} : P(\text{rleaf}(x), x) \vee \text{rleaf}(x) = \text{null} \\
 \text{l-Leaf2} : P(\text{lleaf}(x), \text{l}(x)) & \text{r-Leaf2} : P(\text{rleaf}(x), \text{r}(x)) \\
 \text{l-Leaf3} : \text{lleaf}(\text{lleaf}(x)) = \text{null} & \text{r-Leaf3} : \text{rleaf}(\text{rleaf}(x)) = \text{null} \\
 \text{l-Leaf4} : \text{lleaf}(\text{rleaf}(x)) = \text{null} & \text{r-Leaf4} : \text{rleaf}(\text{rleaf}(x)) = \text{null} \\
 \text{Leaves1} : \text{fca}(\text{lleaf}(x), \text{rleaf}(x)) = x \vee \text{lleaf}(x) = \text{null} \vee \text{rleaf}(x) = \text{null} \\
 \text{Leaves2} : (\text{lleaf}(x) = \text{null} \vee \text{rleaf}(x) = \text{null}) \wedge \text{fca}(y, z) = x \rightarrow x = y \vee x = z \vee x = \text{null} \\
 \text{Leaves3} : \text{lleaf}(x) = \text{null} \wedge \text{rleaf}(x) = \text{null} \wedge P(y, x) \rightarrow y = x \vee x = \text{null}
 \end{array}$$

Fig. 11. Axioms for the auxiliary function symbols root , lleaf , and rleaf

and $\text{p}(r(t))$. This suggests that we should add the following additional ground terms to the set of ground terms generated by $\Psi_0(T)$:

$$\begin{aligned}
 \Psi_1(T) \stackrel{\text{def}}{=} & T \cup \{ \text{l}(t), \text{p}(\text{l}(t)), \text{r}(t), \text{p}(\text{r}(t)) \mid (\text{p}, t) \in \text{Grd} \} \\
 & \cup \{ \text{p}(t), \text{l}(\text{p}(t)), \text{p}(\text{l}(\text{p}(t))) \mid (\text{l}, t) \in \text{Grd} \} \\
 & \cup \{ \text{p}(t), \text{r}(\text{p}(t)), \text{p}(\text{r}(\text{p}(t))) \mid (\text{r}, t) \in \text{Grd} \}
 \end{aligned}$$

where Grd is the set of all pairs (f, t) of function symbols and ground terms appearing in guards of clauses of the form $f(z) = t \rightarrow H$ in \mathcal{K}_{fca} .

If for some $(f, t) \in \text{Grd}$ the weak partial model α satisfies $t = \text{null}$ then the situation is not quite so simple. We have to make sure that α already explicitly determines which nodes $u \in \alpha$ satisfy $f(u) = \text{null}$, even if f is not defined on u . However, there is no finite set of ground terms T over the signature Σ_{fca} such that instantiation of \mathcal{K}_{fca} with the terms in T will ensure this. To enable the construction of such a finite set of terms, we introduce auxiliary functions root , lleaf , and rleaf that determine the root, a left child, and a right child of every node in a forest. More precisely, the semantics of these functions is as follows: for each $u \in \alpha$, $\text{root}(u)$ determines the root of the tree in α to which u belongs (i.e., in all total models α of \mathcal{K}_{fca} and $u \in \alpha$, $\text{p}(u) = \text{null}$ iff $\text{root}(u) = u$). Similarly, $\text{lleaf}(u)$ is some leaf of the tree to which u belongs such that $\text{lleaf}(u)$ is descendant of $\text{l}(u)$, or null if $\text{l}(u)$ is null (i.e., in all total models α of \mathcal{K}_{fca} and $u \in \alpha$, $\text{l}(u) = \text{null}$ holds iff $\text{lleaf}(u) = \text{null}$). The semantics of rleaf is analogous. Let Σ be the signature Σ_{fca} extended with fresh unary function symbols root , lleaf , and rleaf . The axioms capturing this semantics are given in Figure 11. We can then replace every clause $f(z) = t \rightarrow H$ in \mathcal{K}_{fca} by the two clauses

$$f(z) = t \rightarrow t = \text{null} \vee H \quad \text{and} \quad t = \text{null} \wedge N_f(z) \rightarrow H$$

where $N_f(z)$ is $\text{root}(z) = z$ if f is p , $\text{lleaf}(z) = \text{null}$ if f is l , and $\text{rleaf}(z) = \text{null}$ if f is r . Let \mathcal{K} be the resulting set of clauses extended with the linearized and flattened clauses obtained from the axioms in Figure 11. After this final rewriting step no non-ground occurrences of function symbols l , r , p remain in the clauses that resulted from quantified subformulas in the original formula F_0 .

Lemma 6. *The formula F_3 has a finite Σ_P -model iff $\mathcal{K} \cup G$ has a finite Σ -model.*

The final definition of the closure operator Ψ is then as follows:

$$\begin{aligned} \text{Roots}(T) &\stackrel{\text{def}}{=} \text{root}^1(T) \cup \text{root}(\text{root}^1(T)) \\ \text{Leaves}(T) &\stackrel{\text{def}}{=} \text{lleaf}^1(T \cup \text{root}^1(T)) \cup \text{rleaf}^1(T \cup \text{root}^1(T)) \\ \Psi_2(T) &\stackrel{\text{def}}{=} T \cup \text{Roots}(T) \cup \text{Leaves}(T) \cup \text{lleaf}(\text{Leaves}(T)) \cup \text{rleaf}(\text{Leaves}(T)) \\ \Psi(\mathcal{K}, T) &\stackrel{\text{def}}{=} \Psi_5 \circ \Psi_4 \circ \Psi_3 \circ \Psi_2 \circ \Psi_1 \circ \Psi_0(\text{st}(\mathcal{K}) \cup T) \end{aligned}$$

One can easily check that Ψ satisfies the conditions (i) to (iv) on the closure operator of a Ψ -local theory, as defined in Section 3.

Lemma 7. *If there exists a weak partial model of $\mathcal{K}[\Psi(\mathcal{K}, G)] \cup G$ in which all terms in $\Psi(\mathcal{K}, G)$ are defined, then there exists a finite total model of $\mathcal{K} \cup G$.*

Lemma 7 implies that we can decide satisfiability of $\mathcal{K} \cup G$ using the decision procedure described in [9, Section 3.1]. Together with the previous Lemmas we conclude that the combination of the steps described in this section result in a decision procedure for the satisfiability problem of TREX.

Complexity. Note that the number of terms in $\Psi(\mathcal{K}, G)$ is polynomial in the size of $\mathcal{K} \cup G$. From the parametric complexity considerations for Ψ -local theories in [9, 16] follows that satisfiability of $\mathcal{K} \cup G$ can be checked in NP. Further note that all steps of the reduction, except for the elimination of function updates, increase the size of the formula at most by a polynomial factor. The case splits in the rewrite steps 2. and 3. of the function update elimination may cause that the size of the formula increases exponentially in the nesting depth of function updates in the original formula F_0 . However, this exponential blowup can be easily avoided using standard techniques that are used, e.g., for efficient clausal normal form computation.

Theorem 8. *The satisfiability problem for TREX is NP-complete.*

Implementation and experiments. We started implementation of our decision procedure in the Jahob system. Our current prototype implements the first three steps of our decision procedure and already integrates with the verification condition generator of Jahob. Instead of manually instantiating the generated axioms, as described in the fourth step of our decision procedure, we currently give the generated axioms directly to the SMT solver and use triggers to encode some of the instantiation restrictions imposed by Ψ . While this implementation is not yet complete, we already successfully used it to verify implementations of operations on doubly-linked lists and a full insertion method on binary search trees (including the loop traversing the tree). The speedup obtained compared to using the MONA decision procedure is significant. For instance, using our implementation the verification of all 16 subgoals for the insert method takes about 1s in total. Checking the same subgoals using MONA takes 135s. We find these initial results encouraging and consistent with other success stories of using SMT solvers to encode NP decision procedures.

6 Conclusion

This paper introduced the logic TREX for reasoning about imperative tree data structures. The logic supports a transitive closure operator and a form of universal quantification. It is closed under propositional operations and weakest preconditions for heap

manipulating statements. By analyzing the structure of partial and finite models, we exhibited a particular Ψ -local axiomatization of TREX, which implies that the satisfiability problem for TREX is in NP. It also yields algorithms for generating model representations for satisfiable formulas, respectively, proofs of unsatisfiability.

References

1. I. Balaban, A. Pnueli, and L. D. Zuck. Shape analysis of single-parent heaps. In *VMCAI*, Lect. Notes in Comp. Sci. Springer, 2007.
2. A. Bouajjani, C. Dragoi, C. Enea, and M. Sighireanu. A logic-based framework for reasoning about composite data structures. In *CONCUR*, 2009.
3. A. Bouajjani, P. Habermehl, A. Rogalewicz, and T. Vojnar. Abstract regular tree model checking of complex dynamic data structures. In *SAS*, 2006.
4. D. Calvanese, G. di Giacomo, D. Nardi, and M. Lenzerini. Reasoning in expressive description logics. In *Handbook of Automated Reasoning*. Elsevier, 2001.
5. W. Charatonik and P. Witkowski. On the Complexity of the Bernays-Schönfinkel Class with Datalog. In *LPAR (Yogyakarta)*, 2010.
6. L. de Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, pages 337–340, 2008.
7. P. Genevès, N. Layaïda, and A. Schmitt. Efficient static analysis of XML paths and types. In *ACM PLDI*, 2007.
8. P. Habermehl, R. Iosif, and T. Vojnar. Automata-based verification of programs with tree updates. *Acta Inf.*, 47:1–31, January 2010.
9. C. Ihlemann, S. Jacobs, and V. Sofronie-Stokkermans. On local reasoning in verification. In *TACAS*, pages 265–281, 2008.
10. N. Immerman, A. M. Rabinovich, T. W. Reps, S. Sagiv, and G. Yorsh. The boundary between decidability and undecidability for transitive-closure logics. In *Computer Science Logic (CSL)*, pages 160–174, 2004.
11. N. Klarlund and A. Møller. *MONA Version 1.4 User Manual*. BRICS Notes Series NS-01-1, Department of Computer Science, University of Aarhus, January 2001.
12. V. Kuncak. *Modular Data Structure Verification*. PhD thesis, EECS Department, Massachusetts Institute of Technology, February 2007.
13. S. Lahiri and S. Qadeer. Back to the future: revisiting precise program verification using SMT solvers. In *POPL*, 2008.
14. T. Lev-Ami, N. Immerman, T. Reps, M. Sagiv, S. Srivastava, and G. Yorsh. Simulating reachability using first-order logic with applications to verification of linked data structures. In *CADE-20*, 2005.
15. A. Podelski and T. Wies. Counterexample-guided focus. In *ACM Symposium on the Principles of Programming Languages (POPL 2010)*, pages 249–260. ACM, 2010.
16. V. Sofronie-Stokkermans. Hierarchic reasoning in local theory extensions. In *CADE*, pages 219–234, 2005.
17. T. Wies. *Symbolic Shape Analysis*. PhD thesis, University of Freiburg, 2009.
18. T. Wies, V. Kuncak, P. Lam, A. Podelski, and M. Rinard. Field constraint analysis. In *Proc. Int. Conf. Verification, Model Checking, and Abstract Interpretation*, 2006.
19. T. Wies, M. Muñoz, and V. Kuncak. On an efficient decision procedure for imperative tree data structures. Technical Report IST-2011-0005, EPFL-REPORT-165193, IST Austria, EPFL, 2011.
20. G. Yorsh, A. M. Rabinovich, M. Sagiv, A. Meyer, and A. Bouajjani. A logic of reachable patterns in linked data-structures. *J. Log. Algebr. Program.*, 2007.
21. K. Zee, V. Kuncak, and M. Rinard. Full functional verification of linked data structures. In *PLDI*, 2008.