

The Extended Theory of Trees and Algebraic (Co)datatypes

Fabian Zaiser

C.-H. Luke Ong

University of Oxford

fabian.zaiser@cs.ox.ac.uk

luke.ong@cs.ox.ac.uk

The first-order theory of finite and infinite trees has been studied since the eighties, especially by the logic programming community. Following Djelloul, Dao and Frühwirth, we consider an extension of this theory with an additional predicate for finiteness of trees, which is useful for expressing properties about (not just datatypes but also) codatatypes. Based on their work, we present a simplification procedure that determines whether any given (not necessarily closed) formula is satisfiable, returning a simplified formula which enables one to read off all possible models. Our extension makes the algorithm usable for algebraic (co)datatypes, which was impossible in their original work due to restrictive assumptions. We also provide a prototype implementation of our simplification procedure and evaluate it on instances from the SMT-LIB.

1 Introduction

Trees play a fundamental role in computer science: syntactic terms can be regarded as finite trees, and operations like matching and unification, which are essential to functional and logic programming languages, can be viewed as solving certain first-order constraints in the structure of finite trees. Furthermore, trees are a model for program schemes, such as higher-order recursion schemes [12], and more generally as computation trees. The structures of finite and infinite trees are also central to the declarative semantics of logic (e.g. [8]) and functional languages (e.g. [16]). Furthermore, they play a role in the verification of programs [13] and in term rewriting systems [6]. The theory of finite and infinite trees was extensively studied by the logic programming community in the eighties. An axiomatization and a decision procedure for these structures was given by Maher in [9, 10] and another decision procedure independently by Comon and Lescanne in [4].

The structure of trees just consists of what one would normally think of as trees with labeled nodes, except that we allow them to be infinite. Examples of finite and infinite trees are depicted in Fig. 1. The labels for these trees are suggestively named after constructors for common algebraic datatypes because we want to specifically consider applications to the theory of (co)datatypes. In functional programming, two common data structures are natural numbers and linked lists:

```
data nat = zero | succ(pred: nat)
data list = nil | cons(head: nat, tail: list)
```

Inhabitants of these types are naturally viewed as trees: the term `cons(succ(zero), cons(zero, nil))` is shown as a tree in Fig. 1a. In some languages, such as Haskell, datatypes behave in fact more like codatatypes [15], i.e. they can be infinitely nested. For example, the term `let t = cons(zero, cons(succ(zero), t)) in t` corresponds to the infinite tree shown in Fig. 1b.

In this work, we consider the first-order theory of trees, extended with a predicate $\text{fin}(t)$ for stating finiteness of t , and propose it as a tool for reasoning about algebraic datatypes and codatatypes. Why not use the theory of (co)datatypes as implemented in many SMT solvers? First, the theory of (co)datatypes is undecidable because selectors (`head`, `tail`) can be applied to the wrong constructor (`nil`) and the

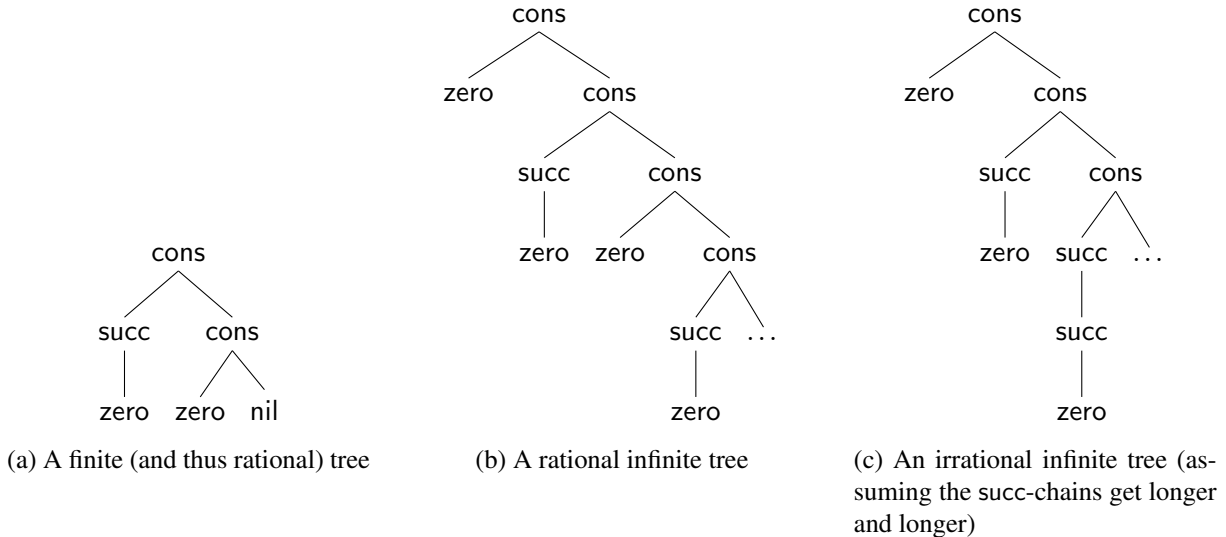


Figure 1: Examples of trees

standard semantics from the SMT-LIB [2] does not specify the result of such an operation (cf. Section 3). Secondly, the theory of trees not only allows us to treat datatypes and codatatypes in a uniform way (both are represented by trees, datatypes just require a fin -predicate), but it is even more expressive: it can explicitly state non-finiteness ($\neg \text{fin}(t)$), as well as finiteness only for a proper subformula (cf. Section 3). Codatatypes are often used in mechanized proofs to represent infinite structures [15] and we believe that this increased flexibility can be useful there as well.

This extended theory of trees (with a fin -predicate) was first presented by Djelloul, Dao and Frühwirth in [5], where they also give a complete axiomatization and a decision procedure. However, one of their core assumptions is that there are infinitely many function symbols (i.e. constructors), which renders it unsuitable for algebraic (co)datatypes found in programming languages since those never have infinitely many constructors. In this work, we lift this restriction and allow sorts with finitely many generators. Note that we require at least two generators, however. Sorts with one generator are not hard to support in principle but require a lot of special casing, so we do not discuss them in this work.

Contributions Our first contribution is a formal description of the relationship between the theory of algebraic (co)datatypes and the extended theory of trees. We also present a decision procedure for first-order formulae (including quantifiers) in the latter, based on [5]. To the best of our knowledge, no decision procedure for this theory allowing finitely generated sorts was known before. Just like the algorithm in [5], it is not only a decision procedure that outputs “satisfiable” or “unsatisfiable” but instead, it simplifies the given input formula as much as possible, which makes it easy to read off all satisfying valuations of the free variables.

We propose this extended theory of trees as an interesting background theory for constrained Horn clauses. Recently, Ong and Wagner [11] proved that satisfiability of higher-order constrained Horn clauses (HoCHCs) is semi-decidable if the background theory is decidable. Since the extended theory of trees is decidable, it is potentially suitable as a background theory for HoCHCs. In addition, we hope that the existence of a simplification procedure instead of a mere decision procedure will also have useful applications to constrained Horn clauses.

Outline The paper is organized as follows. In Section 2, we explain the theory of trees. The following Section 3 introduces (co)datatypes and explores their relationship with the theory of trees. In Section 4, we describe how to check for (finitely generated) sorts with only finitely many finite, respectively infinite, trees. This step is necessary for the extension of the algorithm by Djelloul, Dao and Frühwirth [5] to finitely generated sorts. The extended algorithm is presented in Section 5. Throughout this paper, a sequence of mathematical objects x_1, \dots, x_n is abbreviated as \bar{x} . Unless otherwise stated, u, v, w, x, y, z stand for variables, s for sorts, t for terms, ϕ, ψ for logical formulae, and f, g for constructors and generators.

2 Trees

An *ordered tree* is defined as a (potentially infinite) connected directed acyclic graph with a distinguished node r (the *root*) such that every vertex has exactly one incoming edge, except for r , which has none. Additionally, for each node, its outgoing edges (and the corresponding nodes) are ordered. Furthermore, each node is labeled with an element of a *label set*, also called a *function symbol* or a *constructor*. For a node v , its *subtree* rooted at v is the induced subgraph containing exactly the nodes reachable from v . A tree is called *finite* if it has finitely many nodes, and *rational* if it has finitely many distinct subtrees. Examples of trees can be found in Fig. 1. These use labels suggestively named `zero`, `succ` (representing natural numbers) and `nil`, `cons` (representing lists). This alludes to the connection with algebraic datatypes mentioned in the introduction and further explored in Section 3.

The logical setting for the theory of trees is many-sorted first-order logic. We have a set of *sorts* S , a set of *function symbols* F , a set of *predicate symbols* P and a countable set of *variables* V . Each function symbol f has an *arity* $f : s_1 \times \dots \times s_n \rightarrow s$ with $s_1, \dots, s_n, s \in S$ and we say that f is a *generator* of s . If $n = 0$, we write $f : s$ and say that f is a *constant*. The set of generators of s is called F_s and we assume that each F_s is nonempty. We say that s is *finitely generated* if F_s is finite, and *singular* if F_s is a singleton. Similarly, each predicate symbol p has an arity $p : s_1 \times \dots \times s_n$ with $s_1, \dots, s_n \in S$, and each variable v has a sort $v : s$. For example, the trees in Fig. 1 can be expressed in the first-order language with sorts $S = \{\text{nat}, \text{list}\}$ and function symbols `zero` : nat , `nil` : list , `succ` : $\text{nat} \rightarrow \text{nat}$, and `cons` : $\text{list} \times \text{list} \rightarrow \text{list}$. In the first-order language of trees, the only predicates are $P = \{\text{fin}_s : s \mid s \in S\}$, which state that a given tree of sort s is finite. We will drop the index s if there is no ambiguity.

Given such a *signature* (S, F, P) , a many-sorted *structure* \mathcal{A} consists of non-empty sets $s^{\mathcal{A}}$ for each $s \in S$, functions $f^{\mathcal{A}} : s_1^{\mathcal{A}} \times \dots \times s_n^{\mathcal{A}} \rightarrow s^{\mathcal{A}}$ for each function symbol $f : s_1 \times \dots \times s_n \rightarrow s$ and a predicate $\text{fin}_s^{\mathcal{A}} \subseteq s^{\mathcal{A}}$ for each predicate symbol fin_s . A *valuation* for \mathcal{A} is a family of mappings $V_s \rightarrow s^{\mathcal{A}}$, indexed by S , where V_s denotes the variables of sort s . A *model* of a formula ϕ in \mathcal{A} is a valuation making ϕ true in \mathcal{A} .

The structure \mathcal{T} of trees interprets a signature (S, F, P) as follows. Each sort s is interpreted as the set $s^{\mathcal{T}}$ of *trees of sort* s , meaning the trees where the root is labeled with a generator $f : s_1 \times \dots \times s_n \rightarrow s$ and its children (in order) are roots of subtrees of sorts s_1, \dots, s_n , respectively. Each function symbol $f : s_1 \times \dots \times s_n \rightarrow s$ is interpreted as a function $f^{\mathcal{T}} : s_1^{\mathcal{T}} \times \dots \times s_n^{\mathcal{T}} \rightarrow s^{\mathcal{T}}$ such that $f^{\mathcal{T}}(t_1, \dots, t_n)$ is the tree with a root labeled f and subtrees t_1, \dots, t_n . Each predicate fin_s is interpreted as the subset $\text{fin}_s^{\mathcal{T}} \subseteq s^{\mathcal{T}}$ of finite trees of sort s . Thus $\text{fin}_s(t)$ holds in \mathcal{T} if and only if the interpretation of t is a finite tree. For example, if t_1 is the term `cons(succ(zero), cons(zero, nil))`, depicted in Fig. 1a, then $\text{fin}(t_1)$ is true in \mathcal{T} . On the other hand, if t_2 is the unique tree to make $t_2 = \text{cons}(\text{zero}, \text{cons}(\text{succ}(\text{zero}), t_2))$ true (this tree is shown in Fig. 1b) then $\text{fin}(t_2)$ is false because this tree is infinite.

We call the *theory* of \mathcal{T} , i.e. the set of sentences that are true of \mathcal{T} , the *extended theory of trees* (“extended” because of the additional predicate `fin`). This theory was first presented by Djelloul, Dao

and Frühwirth [5]. However, they require that each sort have at least one constant generator and infinitely many non-constant generators. This assumption simplifies the treatment of the theory significantly but has the serious drawback that it makes their method unsuitable for algebraic (co)datatypes, which typically have only finitely many constructors (i.e. generators). Therefore, we first take a look at the relation between the two theories.

3 Relationship with (Co)Datatypes

The theory of *algebraic (co)datatypes*, also called *(co)inductive datatypes*, is similar to the theory of trees but there are a couple of important differences. For one thing, the set of sorts is partitioned into $S = S_{dt} \cup S_{ct}$ where S_{dt} is the set of datatypes and S_{ct} is the set of codatatypes. The function symbols are partitioned into the constructors F_{ctr} and the selectors F_{sel} , and there are no predicate symbols. Each (co)datatype δ is equipped with $m \geq 1$ constructors $F_{ctr}^\delta = \{C_1, \dots, C_m\}$. Each constructor $C_i \in F_{ctr}^\delta$ has an arity $\delta_1 \times \dots \times \delta_{n_i} \rightarrow \delta$ and is associated with n_i selectors $\mathbf{sel}_{C_i}^j : \delta \rightarrow \delta_j$. Note that for a datatype (resp. codatatype) declaration, all constructor arguments must be datatypes (resp. codatatypes); no mixing is allowed. Such an assumption is common, for example in [15]. Another requirement is that datatypes be well-founded, i.e. one must be able to exhibit a ground term for each datatype. This excludes examples like a datatype *infinite* with a single constructor $\mathbf{next} : \mathit{infinite} \rightarrow \mathit{infinite}$. However, the latter is allowed as a *codatatype*.

Example 3.1. Consider Booleans and lists. Their declaration in many programming languages looks roughly like this:

```
data bool = True | False
data list = Nil | Cons(head : bool, tail : list)
```

where the selectors are called $\mathbf{head} := \mathbf{sel}_{\mathbf{Cons}}^1$ and $\mathbf{tail} := \mathbf{sel}_{\mathbf{Cons}}^2$. The former extracts the first element of a given list, if it is nonempty, and the latter returns the rest of the list. Hence we have $F_{ctr}^{bool} = \{\mathbf{True}, \mathbf{False}\}$, $F_{sel}^{bool} = \emptyset$, $F_{ctr}^{list} = \{\mathbf{Nil}, \mathbf{Cons}\}$, and $F_{sel}^{list} = \{\mathbf{head}, \mathbf{tail}\}$.

Semantics Both datatypes and codatatypes are interpreted as *constructor trees*, i.e. sorted trees labeled only with constructors (not selectors). A structure \mathcal{D} of (co)datatypes interprets a codatatype γ as the set $\gamma^\mathcal{D}$ of constructor trees of sort γ and a datatype δ as the set $\delta^\mathcal{D}$ of finite constructor trees of sort δ . Each constructor $C : s_1 \times \dots \times s_n \rightarrow s$ is interpreted as the function $C^\mathcal{D} : s_1^\mathcal{D} \times \dots \times s_n^\mathcal{D} \rightarrow s^\mathcal{D}$ constructing a new tree out of the given ones, with root C . Each selector $\mathbf{sel}_C^i : s \rightarrow s_i$ is interpreted as a function $(\mathbf{sel}_C^i)^\mathcal{D} : s^\mathcal{D} \rightarrow s_i^\mathcal{D}$, which must satisfy $(\mathbf{sel}_C^i)^\mathcal{D}(C^\mathcal{D}(t_1, \dots, t_n)) = t_i$ but is not specified on inputs built with the wrong constructor. This semantics is very common and what the SMT-LIB standard specifies [2]. Note that other semantics are possible, however, such as returning a fixed default value if a selector is applied to the wrong constructor [3]. We call the latter the *semantics with default values*. The theory of (co)datatypes is the set of sentences that are true in any structure of (co)datatypes satisfying the above.

Theorem 3.2. *The first-order theory of (co)datatypes is undecidable.*

Proof idea. The proof is based on the undecidability of formulae with quantifiers in the theory of uninterpreted functions (EUF). Uninterpreted functions $f : s_1 \times s_2 \rightarrow s$ can be emulated using the following construct:

```
data dummy = c(a1 : s1, a2 : s2) | d(h : s)
```

Then $h(c(x, y))$ acts like an uninterpreted function $f(x, y)$ because the selector h is applied to the wrong constructor c . For the full proof, refer to the appendix. \square

In the semantics with default values, selectors present much less of a problem: they can simply be eliminated.

Theorem 3.3. *In the theory of (co)datatypes with default values, a given formula can be effectively transformed into an equivalent one without selectors.*

Proof idea. The idea is to introduce additional variables such that selectors only occur in equations of the form $x = \mathbf{sel}_C^i(t)$ where t does not contain selectors and then to rewrite such an equation as

$$(\exists \bar{v}. t = C(\bar{v}) \wedge x = v_i) \vee ((\neg \exists \bar{v}. t = C(\bar{v})) \wedge x = T_C^i)$$

where T_C^i is the default value for this selector. For the full proof, refer to the appendix. \square

But even in the standard semantics, quantifier-free formulae in the theory of (co)datatypes are decidable [15]. In fact, we can also eliminate selectors from such formulae.

Theorem 3.4. *In the theory of (co)datatypes with standard semantics, a quantifier-free formula can be effectively transformed into an equisatisfiable one without selectors (but including quantifiers).*

Proof idea. The first step is to introduce additional variables such that selectors only occur in equations of the form $x = \mathbf{sel}_C^i(t)$ where t does not contain selectors. For each such equation, we add the conjunct $\forall \bar{z}. t = C(\bar{z}) \rightarrow z_i = x$, which ensures that the selector correctly extracts the argument when applied to the right constructor. Furthermore, for each pair of such equations $x = \mathbf{sel}_C^i(t)$, $x' = \mathbf{sel}_C^i(t')$, we add the conjunct $t = t' \rightarrow x = x'$, which ensures that selectors behave like functions, i.e. return the same result when applied to the same arguments. For the full proof, refer to the appendix. \square

Finally, we show that the theory of trees is enough for selector-free formulae.

Theorem 3.5. *A selector-free formula in the theory of (co)datatypes can be effectively transformed into an equisatisfiable formula in the extended theory of trees.*

Proof idea. Since in both theories, terms are interpreted as trees, we just have to ensure that datatypes are interpreted as finite trees. Hence, for a datatype d , existential quantification $\exists x : d. \phi$ is replaced by $\exists x : d. \mathbf{fin}(x) \wedge \phi$ and universal quantification $\forall x : d. \phi$ is replaced by $\forall x : d. \mathbf{fin}(x) \rightarrow \phi$. Finally, to ensure equisatisfiability, free variables $x : d$ require adding the conjunct $\mathbf{fin}(x)$ to the whole formula. For the full proof, refer to the appendix. \square

The last result raises the question of how the expressiveness of the extended theory of trees compares to selector-free (co)datatypes. The former is, in fact, more expressive because it allows specifying non-finiteness of individuals, such as $\neg \mathbf{fin}(x)$. This is impossible in the theory of (co)datatypes since datatypes have only finite values and codatatypes can have finite and infinite values. Additionally, it facilitates specifying finiteness only in parts of the formula, such as in $(\mathbf{fin}(t) \rightarrow \phi) \vee (\neg \mathbf{fin}(t) \rightarrow \psi)$, where t is finite in ϕ but infinite in ψ . This shows that the extended theory of trees is more powerful than the (selector-free) theory of (co)datatypes.

4 Analyzing Finitely Generated Sorts

Having shown how formulae involving (co)datatypes can often be reduced to formulae involving trees, we want to find a decision procedure for the latter based on the work by Djelloul, Dao and Frühwirth [5]. Their algorithm, however, makes the assumption that each sort contains infinitely many non-constant generators and one constant generator. As a consequence, each sort contains infinitely many finite and

false : <i>bool</i>	zero : <i>nat</i>	nil : <i>list</i>
true : <i>bool</i>	succ : <i>nat</i> → <i>nat</i>	cons : <i>nat</i> × <i>list</i> → <i>list</i>
tree1 : <i>inftree</i> → <i>inftree</i>	c1 : <i>bool</i> → <i>d</i>	g1 : <i>bool</i> × <i>bool</i> → <i>t</i>
tree2 : <i>inftree</i> × <i>inftree</i> → <i>inftree</i>	c2 : <i>nat</i> × <i>inftree</i> → <i>d</i>	g2 : <i>bool</i> × <i>nat</i> → <i>t</i>

Figure 2: Generators for the sorts $S = \{bool, nat, list, inftree, d, t\}$.

infinitely many infinite trees. This simplifies solving logical formulae: the predicate $\text{fin}(x)$ can always be made true or false for an appropriate valuation of x . However, their assumption is obviously not satisfied for sorts arising from (co)datatypes.

Therefore, we consider the setting with finitely generated sorts, where the situation is more complicated. For instance, if x is of, say, a Boolean sort with only constant generators then the predicate $\text{fin}(x)$ is always true. Due to these complications, we need to analyze the set of sorts and check for sorts with only finitely many finite or infinite trees.

In the following, we allow sorts with finitely many generators but assume that any sort has at least two generators. As mentioned before, this restriction is not hard to lift in principle but saves us a lot of technical details and space in this paper. Note that sorts with a single non-recursive generator can just be unfolded in the place that they are used.

For a sort s , denote by s_{fin} , respectively s_{infin} , the set of finite, respectively infinite, trees of sort s . Denote by S_{0F} , S_{FF} , S_{0I} , S_{1I} , $S_{FI} \subseteq S$ the sets of sorts with no finite trees, finitely many finite trees, no infinite trees, exactly one infinite tree, and finitely many infinite trees, respectively. In the following, we present algorithms for computing these sets.

Algorithm 1 Algorithm computing the sets of sorts containing no finite ($S_{0F} \subseteq S$) or no infinite trees ($S_{0I} \subseteq S$), when given a signature (S, F, P) as input.

```

 $S_{0I} \leftarrow \emptyset$ 
 $S_{0F} \leftarrow S$ 
repeat
   $S_{0I} \leftarrow S_{0I} \cup \{s \in S \mid \forall (g : s_1 \times \dots \times s_n \rightarrow s) \in F_s : \forall i \in \{1, \dots, n\} : s_i \in S_{0I}\}$ 
   $S_{0F} \leftarrow S_{0F} \setminus \{s \in S \mid \exists (g : s_1 \times \dots \times s_n \rightarrow s) \in F_s : \forall i \in \{1, \dots, n\} : s_i \notin S_{0F}\}$ 
until no changes in the last iteration

```

Theorem 4.1. *Given a signature (S, F, P) , Algorithm 1 correctly computes the sets S_{0F} and S_{0I} .*

Proof idea. A sort contains no infinite trees if every generator only takes arguments of sorts containing no infinite trees. A sort contains no finite trees unless some generator takes only arguments of sorts with finite trees. The sets S_{0F} and S_{0I} can thus be computed as fixed points, the former a least fixed point, the latter a greatest fixed point. For details, refer to the full proof in the appendix. \square

Example 4.2. Consider the sorts and generators in Fig. 2. How would Algorithm 1 act on this input? In the first iteration, it would add *bool* to S_{0I} because each generator is a constant. At the same time, *bool*, *list*, and *nat* are removed from S_{0F} because each one has a constant generator. In the next iteration, S_{0I}

Algorithm 2 Algorithm computing the sets of sorts containing only finitely many finite (S_{FF}), respectively infinite (S_{FI}), trees; and their finite (s_{fin}), respectively infinite (s_{infin}), inhabitants.

Compute S_{0I} and S_{0F} as in Algorithm 1

$S_{FF} \leftarrow S_{0F}$

$s_{\text{fin}} \leftarrow \emptyset$ for each $s \in S$

$S_{1I} \leftarrow S \setminus S_{0I}$

$U_s \leftarrow \emptyset$ for each $s \in S$

repeat

for $s \in S$ **do**

$F_s^{\text{infin}} \leftarrow \{(g : s_1 \times \dots \times s_n \rightarrow s) \in F_s \mid \exists i \in \{1, \dots, n\} : s_i \in S_{0F}\}$

if $|F_s \setminus F_s^{\text{infin}}| < \infty$ and $\forall (g : s_1 \times \dots \times s_n \rightarrow s) \in F_s \setminus F_s^{\text{infin}} : \forall i \in \{1, \dots, n\} : s_i \in S_{FF}$ **then**

$S_{FF} \leftarrow S_{FF} \cup \{s\}$

$s_{\text{fin}} \leftarrow \{g(r_1, \dots, r_n) \mid (g : s_1 \times \dots \times s_n \rightarrow s) \in F_s \setminus F_s^{\text{infin}}, r_i \in (s_i)_{\text{fin}}\}$

if $\exists (g : s_1 \rightarrow s) \in F_s : s_1 \in S_{1I} \wedge (\forall (g' : s'_1 \times \dots \times s'_n \rightarrow s) \in F_s \setminus \{g\} : \forall i : s'_i \in S_{0I})$ **then**

$U_s \leftarrow \{u_s = g(u_{s_1})\} \cup U_{s_1}$

else

$S_{1I} \leftarrow S_{1I} \setminus \{s\}$

until no changes in the last iteration

$S_{FI} \leftarrow S_{0I} \cup S_{1I}$

$s_{\text{infin}} \leftarrow \emptyset$ for each $s \in S_{0I}$

$s_{\text{infin}} \leftarrow \{u_s\}$ for each $s \in S_{1I}$

repeat

for $s \in S$ **do**

$F_s^{\text{infin}} \leftarrow \{(g : s_1 \times \dots \times s_n \rightarrow s) \in F_s \mid \exists i \in \{1, \dots, n\} : s_i \notin S_{0I}\}$

if $|F_s^{\text{infin}}| < \infty$ and $\forall (g : s_1 \times \dots \times s_n \rightarrow s) \in F_s^{\text{infin}} : \forall i \in \{1, \dots, n\} :$

$s_i \in S_{0I} \vee (s_i \in S_{FI} \wedge (\forall j \in \{1, \dots, n\} \setminus \{i\} : s_j \in S_{FF} \cap S_{FI}))$ **then**

$S_{FI} \leftarrow S_{FI} \cup \{s\}$

$s_{\text{infin}} \leftarrow \{g(r_1, \dots, r_n) \mid (g : s_1 \times \dots \times s_n \rightarrow s) \in F_s^{\text{infin}}, r_j \in (s_j)_{\text{fin}} \cup (s_j)_{\text{infin}} \text{ for } j = 1, \dots, n \text{ such that } \exists i \in \{1, \dots, n\} : r_i \notin (s_i)_{\text{fin}}\}$

until no changes in the last iteration

stays unchanged but d is removed from S_{0F} because it has the generator c1 whose parameter sort bool is not in S_{0F} anymore. For a similar reason, t is removed from S_{0F} . After this point, no more changes happen and we obtain $S_{0I} = \{\text{bool}\}$ and $S_{0F} = \{\text{inf tree}\}$.

Next, we consider the sets S_{FF}, S_{1I}, S_{FI} . Note that the sort nat from the above example has exactly one infinite tree, namely $\text{succ}(\text{succ}(\dots))$. For such sorts $s \in S_{1I}$, we introduce variables u_s for their unique infinite tree. For instance, $u_{\text{nat}} = \text{succ}(u_{\text{nat}})$ describes the unique infinite tree of nat . The sort t has two infinite trees $g2(\text{false}, u_{\text{nat}})$ and $g2(\text{true}, u_{\text{nat}})$. Hence to describe all infinite trees of sorts $s \in S_{FI}$, we need the variables u_s for $s \in S_{1I}$ and their equations, like $u_{\text{nat}} = \text{succ}(u_{\text{nat}})$. Algorithm 2 computes all this.

Theorem 4.3. *Given a signature (S, F, P) , Algorithm 2 correctly computes the sets S_{FF} , S_{1I} , and S_{FI} . Furthermore it computes the set s_{fin} (the terms for the finite trees of sort s for $s \in S_{FF}$), and the set s_{infin} (the terms for the infinite trees of sort s for $s \in S_{FI}$). The latter makes use of the variables u_s (for $s \in S_{1I}$), standing for the unique infinite tree of s . The equations that uniquely determine these u_s are output in U_s .*

Proof idea. Similarly to the previous algorithm, these sets are computed as fixed points. A sort s has only finitely many finite trees if there is a finite number of generators that only take sorts with finitely many finite trees as arguments ($F_s \setminus F_s^{\text{infin}}$), and the remaining generators (F_s^{infin}) take at least one argument of a sort that contains no finite trees (because such a generator cannot create finite trees). Along the way, the algorithm builds up the set s_{fin} from the generators of the former category.

Constructing the set S_{FI} works similarly, except for the fact that we start the fixed point iteration with $S_{0I} \cup S_{1I}$ instead of the empty set. The reason is that for every sort with finitely many infinite trees, it can be shown (Lemma B.5 in the appendix) that the infinite parts of each such tree are built from the unique infinite trees of the sorts S_{1I} . These sorts with a unique infinite tree are also constructed by fixed point iteration. They can only have a single generator g that constructs infinite trees and it can only take one argument because otherwise we would have at least two infinite trees since each sort is assumed to have at least two generators.

A sort s only has finitely many infinite trees if the set of generators constructing infinite trees (F_s^{infin}) is finite, and when picking an arbitrary argument i of it, this argument allows no infinite trees; or it allows finitely many infinite trees and all the other arguments allow only finitely many trees. This explains the fixed point iteration for S_{FI} and s_{infin} . For details, refer to the full proof in the appendix. \square

Example 4.4. Consider again the signature from Fig. 2. How does Algorithm 2 act on it? At the start of the first loop, we have $S_{FF} = \{\text{inf tree}\}$ and $S_{1I} = S$. In the first iteration, bool is added to S_{FF} because all its generators are constants, and $\text{bool}_{\text{fin}} = \{\text{false}, \text{true}\}$. Additionally, nat stays in S_{1I} because its generator succ satisfies $\text{nat} \in S_{1I}$ and the other generator is constant. Therefore $U_{\text{nat}} = \{u_{\text{nat}} = \text{succ}(u_{\text{nat}})\}$. All the other sorts are removed from S_{1I} , either because they don't have a unary generator (bool , list) or there is another generator that allows infinite trees, destroying uniqueness (tree2 for inf tree , c2 for d , and g2 for t). In the second iteration, d is added to S_{FF} because $\text{c1} : \text{bool} \rightarrow d$ only constructs finitely many finite trees since $\text{bool} \in S_{FF}$ and its other generator c2 constructs only infinite trees. Therefore, d_{fin} is set to $\{\text{c1}(\text{true}), \text{c1}(\text{false})\}$. After this point, no more changes happen.

At the start of the second loop, we have $S_{FI} = \{\text{bool}, \text{nat}\}$ and $\text{nat}_{\text{infin}} = \{u_{\text{nat}}\}$. In the loop iteration, t is added to S_{FI} because we have $F_t^{\text{infin}} = \{\text{g2}\}$, which is finite, and its only generator g2 has the property that its first parameter is $\text{bool} \in S_{0I}$ and its second parameter is $\text{nat} \in S_{FI}$ with the additional property that all remaining parameters, i.e. bool , are in $S_{FF} \cap S_{FI}$. Therefore t_{infin} is set to $\{\text{g2}(\text{false}, u_{\text{nat}}), \text{g2}(\text{true}, u_{\text{nat}})\}$. After this point, no more changes happen. The algorithm has computed $S_{FF} = \{\text{inf tree}, \text{bool}, d\}$ and $S_{FI} = \{\text{nat}, t\}$.

5 Simplification Procedure for the Extended Theory of Trees

Having explained how to analyze finitely generated sorts, we can now describe how the simplification procedure from [5] is extended to finitely generated sorts. Before going into detail, we provide a brief outline of this algorithm. The procedure works on special formulae, called *normal formulae*. Any formula can be transformed into an equivalent normal formula, so this is not a restriction. Roughly speaking, the output of our algorithm is a disjunction of *fully simplified* formulae that is equivalent to the original formula. A fully simplified formula makes it easy to read off all its models. The simplification algorithm works similarly to [5], except for the fact that finitely generated sorts sometimes require case splits (also called instantiations) for certain variables (called *instantiable*). These case splits can be on the finitely many generators of a sort, or on the finitely many (finite or infinite) inhabitants of a sort if it is in S_{FF} or S_{FI} . In this section, we focus on these instantiable variables and case splits because that

is the novel part of our extension of [5]. The full algorithm is described in the appendix (Algorithm 5). Before we can start with the concept of normal formulae, we first need to define *basic formulae*.

Definition 5.1. A *basic formula* is of the form $(\bigwedge_i v_i = t_i) \wedge (\bigwedge_j \text{fin}(u_j))$ where \bar{u}, \bar{v} are variables and each t_i is a variable or a term of the form $f(\bar{z})$ for a function symbol f and variables \bar{z} . Such a formula will be abbreviated as $\bar{v} = \bar{t} \wedge \overline{\text{fin}(u)}$. Given a total order on its variables, it is called *solved* if (1) the variables \bar{u}, \bar{v} are distinct and for each equation $x = y$, we have $x > y$, and (2) if $\text{fin}(v)$ occurs then the sort of v contains both finite and infinite trees. A variable x_n is *reachable* from a variable x_0 if the basic formula contains $x_0 = t_0 \wedge x_1 = t_1 \wedge \dots \wedge x_{n-1} = t_{n-1}$ where each t_i contains x_{i+1} . It is *properly reachable* if $n > 0$. The subformulae $u = t$ and $\text{fin}(u)$ are considered reachable if u is.

The variable ordering is important when we consider basic subformulae of larger formulae. Then this ordering ensures that in solved basic formulae, variables bound deeper inside the whole formula occur on the left-hand side of equations, which is important for the correctness proof. Intuitively, reachability means the following: if y is reachable from x then y is a subtree of x . Djelloul, Dao and Frühwirth describe an algorithm to solve a basic formula (rules 1–10 in [5, Section 4.6]). In our extended setting, two things have to be changed: if $\text{fin}(u)$ occurs in the basic formula where $u : s$ with $s \in S_{0I}$, or $s \in S_{0F}$, then $\text{fin}(u)$ is always satisfied and can be removed, or is never satisfied and the basic formula is unsolvable, respectively. This is summarized by the following theorem.

Theorem 5.2. *There is an algorithm SOLVEBASIC($(v_0, \dots, v_n), \alpha$) (Algorithm 4 in the appendix) that correctly solves basic formulae α containing the variables v_0, \dots, v_n , i.e. it turns α into an equivalent solved formula (with respect to the variable ordering $v_0 < \dots < v_n$) or returns false if none exists.*

Basic formulae are insufficient for the general case but they are an important building block for the concept of *normal formulae*, which can express any first-order formula.

Definition 5.3. A *normal formula* ϕ of depth $d \geq 1$ takes the form $\neg(\exists \bar{x}. \alpha \wedge \bigwedge_{i=1}^n \phi_i)$ where α is a basic formula, and each ϕ_i is a normal formula of depth d_i with $d = 1 + \max(0, d_1, \dots, d_n)$.

The simplest normal formula is $\neg \text{true}$. As normal formulae allow expressing negation, conjunction, existential quantification, and nesting, the next theorem is straightforward to prove [5, Property 4.3.3].

Theorem 5.4. *There is an algorithm NORMALIZE(ϕ) which turns any first-order formula ϕ into a normal one that is equivalent in the theory of trees.*

Example 5.5. Consider the formula $\forall x : \text{nat}. \neg \text{fin}(x) \rightarrow x = \text{succ}(x)$. It can be rewritten as $\neg(\exists x : \text{nat}. \neg \text{fin}(x) \wedge \neg(x = \text{succ}(x)))$, which is a normal formula of depth 2.

Now we come to the main difference from the original algorithm in [5]: our more general setting necessitates case splits (or instantiations) for certain variables. For instance, consider the normal formula $\phi_1 \equiv \neg(\exists x : \text{list}. \neg(x = \text{nil}) \wedge \neg(\exists y, z. x = \text{cons}(y, z)))$. If *list* had infinitely many generators, it would always be possible to find a value for x that is neither *nil* nor *cons*. However, since *list* only has those two generators, no such x exists and the formula is true. Here our extended algorithm will do a case split on both constructors of *list* (described later in more detail) and realize that neither works.

As another example, consider $\phi_2 \equiv \neg(\exists x : t. \neg \text{fin}(x) \wedge \neg(x = y) \wedge \neg(x = z))$. If t had infinitely many infinite trees, then this would be true because we could always choose a valuation for x that is different from the free variables y and z . Since t has only two infinite trees, our extended algorithm does a case split on all two infinite trees of t , instantiating x with $\text{g2}(\text{true}, u_{\text{nat}})$ and $\text{g2}(\text{false}, u_{\text{nat}})$ where u_{nat} is the unique tree of sort *nat*, namely $\text{succ}(\text{succ}(\dots))$. A similar case occurs with a constraint $\text{fin}(x)$ where x only has finitely many finite trees or, in general, if x has only finitely many trees. This leads us to the definition of an *instantiable variable*, i.e. a variable that requires a case split.

Note that every normal formula can be transformed into an equivalent conjunction of normal formulae of depth at most 2 by repeatedly applying rule 16 (depth reduction) from [5, Section 4.6]. Therefore we can limit our attention to such formulae in the following.

Definition 5.6 (instantiable variable). Let $\neg(\exists \bar{x}. \alpha \wedge \bigwedge_i \neg(\exists \bar{y}_i. \beta_i))$ be a normal formula of depth at most 2 such that α and each β_i are solved basic formulae. Let β_i^* be β_i with all conjuncts also occurring in α removed. Then a variable $v : s$ that is free in the formula or occurs in \bar{x} is called *instantiable* if one of the following conditions is satisfied:

1. s has finitely many generators F_s and some β_i^* contains $v = f(\bar{w})$ and v is not properly reachable from v in β_i , or
2. $s \in S_{FF} \cap S_{FI}$, some β_i^* contains v , and α contains no equation $v = t$ for any term t , or
3. $s \in S_{FF}$, α contains $\text{fin}(v)$, and some β_i^* contains v , or
4. $s \in S_{FI}$ and some β_i^* contains only $\text{fin}()$ -constraints (no equations), among them $\text{fin}(v)$.

Algorithm 3 Algorithm for finding instantiable variables and their instantiations.

function FINDINSTANTIATION($\bar{v}, \neg(\exists \bar{x}. \alpha \wedge \bigwedge_i \neg(\exists \bar{y}_i. \beta_i))$)
 $\beta_i^* \leftarrow \beta_i$ without the conjuncts occurring in α
for $u : s \in \bar{v}\bar{x}$ **do**
 if s has finitely many generators, and $u = f(\bar{z})$ occurs in some β_i^* ,
 and u is not properly reachable from u in β_i **then**
 return $\{\exists \bar{z}. u = g(\bar{z}) \mid g \in F_s\}$
 if $s \in S_{FF} \cap S_{FI}$ and u occurs in some β_i^* and α contains no $u = t$ **then**
 return $\{\exists u_{s_1}, \dots, u_{s_n}. u = t \wedge \bigwedge_{s \in S_{1I}} U_s \mid t \in s_{\text{fin}} \cup s_{\text{infin}}\}$ where $S_{1I} = \{s_1, \dots, s_n\}$
 if $s \in S_{FF}$ and u occurs in some β_i^* , and $\text{fin}(u)$ in α **then**
 return $\{u = t \mid t \in s_{\text{fin}}\}$
 if $s \in S_{FI}$ and $\text{fin}(u)$ occurs in some β_i^* that contains only $\text{fin}()$ -constraints **then**
 return $\{\text{fin}(u)\} \cup \{\exists u_{s_1}, \dots, u_{s_n}. u = t \wedge \bigwedge_{s \in S_{1I}} U_s \mid t \in s_{\text{infin}}\}$ where $S_{1I} = \{s_1, \dots, s_n\}$
return *none*

Algorithm 3 looks for an instantiable variable (if any) in a normal formula $\neg(\exists \bar{x}. \alpha \wedge \bigwedge_i \phi_i)$ with free variables \bar{v} by checking exactly the four conditions from above. If it finds an instantiable variable u , it returns a set I of formulae, called *instantiations*. Note that while we write “ $u = t$ ” in the return values for simplicity, we actually mean an equivalent formula $\exists \bar{z}. \gamma$ where γ is a basic formula. For instance, by $u = \text{c1}(\text{true})$, we mean $\exists z. u = \text{c1}(z) \wedge z = \text{true}$ for a fresh variable z . We can use these instantiations to get rid of instantiable variables, as the following theorem explains.

Theorem 5.7. Let $\phi \equiv \neg(\exists \bar{x}. \alpha \wedge \bigwedge_i \phi_i)$ be a normal formula of depth at most 2 with free variables \bar{v} . Let I be the result of FINDINSTANTIATION(\bar{v}, ϕ) from Algorithm 3. If I is “*none*”, then there is no instantiable variable. Otherwise, let u be the first instantiable variable found in FINDINSTANTIATION. Then ϕ is equivalent to the following conjunction of normal formulae, in which the variable u is no longer instantiable:

$$\bigwedge_{(\exists \bar{z}. \psi) \in I} \neg(\exists \bar{x} \bar{z}. \alpha \wedge \psi \wedge \bigwedge_i \phi_i).$$

Example 5.8. In the formula $\phi_1 \equiv \neg(\exists x : \text{list}. \neg(x = \text{nil}) \wedge \neg(\exists y, z. x = \text{cons}(y, z)))$ from above, x is instantiable because of condition 1. (Note that the reachability check in this condition is required to

avoid infinite loops for recursive equations like $x = \text{cons}(y, x)$.) Here `FINDINSTANTIATION` returns $I = \{x = \text{nil}; \exists y, z. x = \text{cons}(y, z)\}$. By the above theorem, ϕ_1 is equivalent to

$$\neg(\exists x : \text{list}. x = \text{nil} \wedge \neg(x = \text{nil}) \wedge \neg(\exists y, z. x = \text{cons}(y, z))) \\ \wedge \neg(\exists x : \text{list}, y, z. x = \text{cons}(y, z) \wedge \neg(x = \text{nil}) \wedge \neg(\exists y, z. x = \text{cons}(y, z))).$$

Both existential subformulae obviously contain a contradiction, so the whole formula simplifies to $\neg(\text{false}) \wedge \neg(\text{false})$ and thus true.

Example 5.9. In the other formula $\phi_2 \equiv \neg(\exists x : t. \neg \text{fin}(x) \wedge \neg(x = y) \wedge \neg(x = z))$ from above, t is instantiabile because of condition 4. Algorithm 3 returns the instantiations

$I = \{\text{fin}(x); \exists u_{\text{nat}}. x = \text{g2}(\text{true}, u_{\text{nat}}) \wedge u_{\text{nat}} = \text{succ}(u_{\text{nat}}); \exists u_{\text{nat}}. x = \text{g2}(\text{true}, u_{\text{nat}}) \wedge u_{\text{nat}} = \text{succ}(u_{\text{nat}})\}$, which means x is either a finite tree or one of the two infinite trees $\text{g2}(\text{false}, u_{\text{nat}})$, $\text{g2}(\text{true}, u_{\text{nat}})$ where u_{nat} is the unique tree with $u_{\text{nat}} = \text{succ}(u_{\text{nat}})$. By the above theorem, ϕ_2 is equivalent to

$$\neg(\exists x : t. \text{fin}(x) \wedge \neg \text{fin}(x) \wedge \neg(x = y) \wedge \neg(x = z)) \\ \wedge \neg(\exists x, u_{\text{nat}} : t. x = \text{g2}(\text{false}, u_{\text{nat}}) \wedge u_{\text{nat}} = \text{succ}(u_{\text{nat}}) \wedge \neg \text{fin}(x) \wedge \neg(x = y) \wedge \neg(x = z)) \\ \wedge \neg(\exists x, u_{\text{nat}} : t. x = \text{g2}(\text{true}, u_{\text{nat}}) \wedge u_{\text{nat}} = \text{succ}(u_{\text{nat}}) \wedge \neg \text{fin}(x) \wedge \neg(x = y) \wedge \neg(x = z)).$$

The other parts of the simplification procedure (unchanged from [5]) simplify this to

$$\text{true} \\ \wedge \neg(\exists u_{\text{nat}} : t. u_{\text{nat}} = \text{succ}(u_{\text{nat}}) \wedge \neg(y = \text{g2}(\text{false}, u_{\text{nat}})) \wedge \neg(z = \text{g2}(\text{false}, u_{\text{nat}}))) \\ \wedge \neg(\exists u_{\text{nat}} : t. u_{\text{nat}} = \text{succ}(u_{\text{nat}}) \wedge \neg(y = \text{g2}(\text{true}, u_{\text{nat}})) \wedge \neg(z = \text{g2}(\text{true}, u_{\text{nat}}))),$$

where the variable x is removed because it is unreachable from the free variables. The resulting formula essentially expresses that y or z has to be equal to $\text{g2}(\text{false}, u_{\text{nat}})$; and that y or z has to be equal to $\text{g2}(\text{true}, u_{\text{nat}})$. In other words, they can only take on those two infinite values and have to be different. Note that the algorithm has not completed at this point yet because y and z are now instantiabile by condition 1. We skip the following (less interesting) instantiations for space reasons.

At this point, we can introduce the notions of *solved* and *fully simplified* formulae, which make up the output of our extended simplification procedure.

Definition 5.10. A normal formula $\phi \equiv \neg(\exists \bar{x}. \alpha \wedge \bigwedge_i \neg(\exists \bar{y}_i. \beta_i))$, of depth at most 2, is called *solved* if it satisfies the following properties.

1. Each β_i and α are solved basic formulae with respect to a variable ordering where $u < v$ if the binding of v is more deeply nested than u , i.e. u is free where v is bound.
2. The equations of α are included in every β_i .
3. Each β_i contains at least one conjunct that does not occur in α .
4. There are no instantiabile variables.
5. All the variables \bar{x} and \bar{y}_i are reachable from the free variables of $\exists \bar{x}. \alpha$ and $\exists \bar{y}_i. \beta_i$, respectively.

A formula ψ is called *fully simplified* if $\neg\psi$ is a solved normal formula. (This is an extension of the definition of “explicit solved form” in [5, Definition 4.4.6].)

Example 5.11. The point of fully simplified formulae is that they are easy to interpret, i.e. it is easy to read off all possible models from them. For instance, consider the fully simplified formula

$$\exists v. x = \text{succ}(v) \wedge v = y \wedge \text{fin}(y) \wedge \neg(\exists w. y = \text{succ}(w) \wedge \text{fin}(w) \wedge \text{fin}(z))$$

Any model has to satisfy $x = \text{succ}(y)$ and y has to be finite. To falsify the other part $\exists w. y = \text{succ}(w) \wedge \text{fin}(w) \wedge \text{fin}(z)$, there are two options for the free variables y and z : (1) instantiate y with any finite tree with a root other than succ and z with any tree, or (2) instantiate y with any finite tree and z with any infinite tree. These are the only two classes of models for the above fully simplified formula. In general, the following holds about fully simplified formulae.

Theorem 5.12. *Let ϕ be a fully simplified formula. If ϕ has no free variables then $\phi \equiv \text{true}$. Otherwise both ϕ and $\neg\phi$ are satisfiable in the theory of trees.*

Given a formula ϕ , the main algorithm returns an equivalent disjunction of fully simplified formulae. Since each disjunct allows an easy description of its models, we can describe all possible models of ϕ .

Theorem 5.13. *There is an algorithm $\text{SOLVE}(\phi)$ (Algorithm 5 in the appendix) that, given a formula ϕ , returns true, false, or a disjunction of fully simplified formulae that is equivalent to ϕ in the extended theory of trees. In particular, if ϕ is closed, it returns true or false.*

Proof idea. The simplification procedure from [5] does not have to be changed a lot. We use the function SOLVEBASIC from Theorem 5.2 to solve basic formulae. Afterwards, the rules 12–14 and 16 from [5, Section 4.6] ensure that the result is a disjunction of formulae satisfying conditions (1–3) of Definition 5.10. At this point, we make use of FINDINSTANTIATION and Theorem 5.7, to ensure that condition (4) is satisfied. After each such instantiation, the previous rules have to be applied again because conditions (1–3) may have been invalidated. These steps will not, however, invalidate condition (4). It is nontrivial to prove that these instantiations terminate (Lemma C.4 in the appendix). Avoiding infinite loops of instantiations is the reason for the complicated condition 1 in Definition 5.6. Finally, Rule 15 of the original algorithm [5, Section 4.6] ensures that condition (5) is satisfied as well. It is again nontrivial to show that this part is still correct in our more general setting (Lemma C.6 in the appendix). \square

Time complexity Regarding the performance of our extended algorithm, note that the original algorithm has non-elementary time complexity [5]. In fact, Vorobyov proved that deciding first-order formulae in the (ordinary) theory of trees already has non-elementary time complexity [17], so we cannot hope for an efficient algorithm for the extended theory of trees in the worst case.

Implementation In order to evaluate the performance in practice, we created a prototype implementation [18]. Due to a lack of benchmarks involving formulae of trees, we took the `tests` set of the `QF_DT` (quantifier-free datatypes) suite of the `SMT-LIB` [2] and transformed each instance into a formula in the extended theory of trees, using the results from Section 3. Then we ran our extended simplification procedure on the transformed instances. Over 90% of them completed in less than 1 second and about 5% timed out after 10 seconds (more data in Table 1 in the appendix). While state-of-the-art SMT solvers decide each `QF_DT` instance in a few milliseconds, our transformed instances are considerably harder because they contain quantifiers and can be significantly larger than the original ones. Furthermore, our prototype implementation obviously cannot compete with heavily optimized SMT solvers and leaves a lot of room for improvements: for instance, the normalization of formulae can be optimized, and heuristics for choosing which variable to instantiate (instead of picking the first one) could make a big difference. It nevertheless demonstrates that our algorithm has a reasonable performance on many practical instances. Our implementation can be found at <https://github.com/fzaiser/tree-theory-solver/> or [18]. Additionally, there is a web interface at <http://mjolnir.cs.ox.ac.uk/trees-codata/>.

6 Conclusion

We believe that the extended theory of trees is an interesting theory to study because of its decidability and connections with algebraic (co)datatypes. We have explained the complications arising from finitely generated sorts, which are necessary to apply it to (co)datatypes. The fact that we not only provide a decision procedure but a simplification procedure should make it easier to conduct further research on the theory of trees, such as investigating Craig interpolation.

References

- [1] W. Ackermann (1954): *Solvable cases of the decision problem*. Studies in logic and the foundations of mathematics, North-Holland Pub. Co.
- [2] Clark Barrett, Pascal Fontaine & Cesare Tinelli (2017): *The SMT-LIB Standard: Version 2.6*. Technical Report, Department of Computer Science, The University of Iowa. Available at <http://www.smt-lib.org>.
- [3] Clark Barrett, Igor Shikanian & Cesare Tinelli (2007): *An Abstract Decision Procedure for a Theory of Inductive Data Types*. *Journal on Satisfiability, Boolean Modeling and Computation* 3(1–2), pp. 21–46, doi:10.3233/SAT190028.
- [4] Hubert Comon & Pierre Lescanne (1989): *Equational Problems and Disunification*. *J. Symb. Comput.* 7(3/4), pp. 371–425, doi:10.1016/S0747-7171(89)80017-3.
- [5] Khalil Djelloul, Thi-Bich-Hanh Dao & Thom W. Frühwirth (2008): *Theory of finite or infinite trees revisited*. *TPLP* 8(4), pp. 431–489, doi:10.1017/S1471068407003171.
- [6] Jean-Pierre Jouannaud & Emmanuel Kounalis (1986): *Automatic Proofs by Induction in Equational Theories Without Constructors*. In: *Proceedings of the Symposium on Logic in Computer Science (LICS '86), Cambridge, Massachusetts, USA, June 16-18, 1986*, pp. 358–366.
- [7] Daniel Kroening & Ofer Strichman (2016): *Decision Procedures – An Algorithmic Point of View, Second Edition*. Texts in Theoretical Computer Science. An EATCS Series, Springer, doi:10.1007/978-3-662-50497-0.
- [8] John W. Lloyd (1987): *Foundations of Logic Programming, 2nd Edition*. Springer, doi:10.1007/978-3-642-83189-8.
- [9] Michael J. Maher (1988): *Complete Axiomatizations of the Algebras of Finite, Rational and Infinite Trees*. In: *Proceedings of the Third Annual Symposium on Logic in Computer Science (LICS '88), Edinburgh, Scotland, UK, July 5-8, 1988*, pp. 348–357, doi:10.1109/LICS.1988.5132.
- [10] M.J. Maher (1988): *Complete Axiomatizations of the Algebras of Finite, Rational and Infinite Trees*. Technical Report, IBM – T.J. Watson Research Center.
- [11] C.-H. Luke Ong & Dominik Wagner (2019): *HoCHC: A Refutationally Complete and Semantically Invariant System of Higher-order Logic Modulo Theories*. In: *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24-27, 2019*, IEEE, pp. 1–14, doi:10.1109/LICS.2019.8785784.
- [12] Luke Ong (2015): *Higher-Order Model Checking: An Overview*. In: *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015*, pp. 1–15, doi:10.1109/LICS.2015.9.
- [13] Derek C. Oppen (1980): *Reasoning About Recursively Defined Data Structures*. *J. ACM* 27(3), pp. 403–411, doi:10.1145/322203.322204.
- [14] Emil L. Post (1946): *A variant of a recursively unsolvable problem*. *Bulletin of the American Mathematical Society* 52(4), pp. 264–269, doi:10.1090/s0002-9904-1946-08555-9.
- [15] Andrew Reynolds & Jasmin Christian Blanchette (2017): *A Decision Procedure for (Co)datatypes in SMT Solvers*. *J. Autom. Reasoning* 58(3), pp. 341–362, doi:10.1007/s10817-016-9372-6.
- [16] David Turner (1987): *An overview of Miranda*. *Bulletin of the EATCS* 33, pp. 103–114.
- [17] Sergei G. Vorobyov (1996): *An Improved Lower Bound for the Elementary Theories of Trees*. In: *Automated Deduction - CADE-13, 13th International Conference on Automated Deduction, New Brunswick, NJ, USA, July 30 - August 3, 1996, Proceedings*, pp. 275–287, doi:10.1007/3-540-61511-3_91.
- [18] Fabian Zaiser (2020): *Artifact for "The Extended Theory of Trees and Algebraic (Co)datatypes"*, doi:10.5281/zenodo.3828522.

A Omitted proofs from Section 3

Theorem (Theorem 3.2, repeated). *The first-order theory of (co)datatypes is undecidable.*

Proof of Theorem 3.2. The proof works by reduction from Post's correspondence problem. An instance of the problem is given by a finite set of pairs of bit strings $\{(x_1, y_1), \dots, (x_n, y_n)\}$, i.e. each $x_i, y_i \in \{0, 1\}^*$. A solution to such an instance is a nonempty finite sequence of indices i_1, \dots, i_k with each $i_j \in \{1, \dots, n\}$ such that

$$x_{i_1} \dots x_{i_k} = y_{i_1} \dots y_{i_k}.$$

The decision problem is to decide whether such a solution exists, and is famously undecidable [14]. Let $\{(x_1, y_1), \dots, (x_n, y_n)\}$ be an instance of the problem. In the following, we construct a signature and formula ϕ , which is satisfiable in the theory of datatypes if and only if the instance has a solution.

For the encoding, consider the following datatype declarations.

```
data bool = true | false
data bitstring = e | 0(tail0 : bitstring) | 1(tail1 : bitstring)
data dummy = f(b1 : bitstring, b2 : bitstring) | g(h : bool)
```

We have the usual definition of Booleans (true or false), bit strings are empty (e) or start with a 0 or with a 1, and the dummy datatype is just used to emulate uninterpreted functions: $h(f(s, s'))$ acts like an unspecified function on the bit strings s, s' because h is applied to the wrong constructor and the standard semantics leaves this case unspecified. In this proof, ϕ is constructed in such a way that $h(f(s, s')) = \text{true}$ holds if s, s' is a valid pair of bit strings that can be built out of the (x_i, y_i) .

First, for a bit string x , we write $p_x(t)$ for the bit string that prepends x to t . For example, if $x = 101$, $p_x(t)$ stands for $1(0(1(t)))$. Consider the following formulae.

$$\begin{aligned} \phi_1 &\equiv \bigwedge_{i=1}^n h(f(p_{x_i}(e), p_{y_i}(e))) = \text{true} \\ \phi_2 &\equiv \forall u, v. h(f(u, v)) = \text{true} \rightarrow \bigwedge_{i=1}^n h(f(p_{x_i}(u), p_{y_i}(v))) = \text{true} \\ \phi_3 &\equiv \exists u. h(f(u, u)) = \text{true} \end{aligned}$$

First, ϕ_1 and ϕ_2 specify what pairs of bit strings can be constructed out of x_i and y_i , and finally ϕ_3 specifies that a solution to the instance exists. Hence we claim that $\phi \equiv \phi_1 \wedge \phi_2 \rightarrow \phi_3$ is valid if and only if the given instance of Post's correspondence problem is solvable.

First, suppose ϕ is valid. Consider the model \mathcal{M} where h is interpreted as follows:

$$h^{\mathcal{M}}(z) = \begin{cases} y & \text{if } z = g^{\mathcal{M}}(y) \\ \text{true}^{\mathcal{M}} & \text{if } z = \mathcal{M}(f(p_x(e), p_y(e))) \\ & \text{where } \exists k \geq 1, i_1, \dots, i_k : x = x_{i_1} \dots x_{i_k} \wedge y = y_{i_1} \dots y_{i_k} \\ \text{false}^{\mathcal{M}} & \text{otherwise.} \end{cases}$$

In this model, ϕ_1 and ϕ_2 are satisfied by construction, hence ϕ_3 is satisfied as well. This means there is an element $u \in \text{bitstring}^{\mathcal{M}}$ such that $h^{\mathcal{M}}(f^{\mathcal{M}}(u, u)) = \text{true}^{\mathcal{M}}$. By the choice of \mathcal{M} , this means that $u = \mathcal{M}(p_x(e))$ such that there are i_1, \dots, i_k with $x = x_{i_1} \dots x_{i_k}$ and $y = y_{i_1} \dots y_{i_k}$, which means the given instance is solvable.

Conversely, suppose the instance of Post's correspondence problem has a solution i_1, \dots, i_k . Let \mathcal{M} be a structure where $\phi_1 \wedge \phi_2$ is true. Let $s_j = x_{i_1} \dots x_{i_j}$ and $t_j = y_{i_1} \dots y_{i_j}$. Then ϕ_1 ensures that $h(f(p_{s_j}(e), p_{t_j}(e))) = \text{true}$ in \mathcal{M} and ϕ_2 ensures that the induction step works and hence that $h(f(p_{s_j}(e),$

$p_{t_j}(e)) = \text{true}$ holds for all j , in particular for $j = k$. Set $u = p_{s_k}(e)$. Since $s_k = t_k$ is a solution to the instance, we have $h(f(u, u)) = \text{true}$ in \mathcal{M} . Hence ϕ_3 also holds in \mathcal{M} . Therefore ϕ is valid.

Since a formula is valid if and only if its negation is unsatisfiable, this proof also shows that the satisfiability problem of first-order formulae in the theory of datatypes is undecidable.

Note that in the proof, we never used the fact that a *bitstring* is finite, so the same proof works when replacing the above datatype declarations by codatatype declarations. \square

Theorem (Theorem 3.3, repeated). *In the theory of (co)datatypes with default values, a given formula can be effectively transformed into an equivalent one without selectors.*

Proof of Theorem 3.3. The transformation works in two steps. First, it ensures that selectors don't occur nested but instead only occur in equations of the form $x = \text{sel}_C^i(t)$ for a variable x . Afterwards, it replaces such equations by an equivalent formula that doesn't contain selectors.

Step 1. The first step isolates selectors to simple equations. For every equation $t = t'$ where one side contains a selector, do the following. If $t = t'$ is of the form $x = \text{sel}_C^i(s)$ for a variable x and a selector-free term s then there is nothing to do. Otherwise, without loss of generality, assume that t contains a (nested) selector, so t can be written as $r[\text{sel}_C^i(s)/x]$ for a term r where x is a fresh variable. Then rewrite $r[\text{sel}_C^i(s)/x] = t'$ to $\exists x. r = t' \wedge x = \text{sel}_C^i(s)$, which is always equivalent in first-order logic. This process is repeated until no more changes can be made.

Step 2. Now the formula only contains selectors as part of equations of the form $x = \text{sel}_C^i(s)$ where s is selector-free. Each such equation can be rewritten as

$$(\exists \bar{v}. s = C(\bar{v}) \wedge x = v_i) \vee ((\neg \exists \bar{v}. s = C(\bar{v})) \wedge x = T_C^i)$$

where \bar{v} are fresh variables and T_C^i is the default value for this selector. This is equivalent because

$$\begin{aligned} x = \text{sel}_C^i(s) &\leftrightarrow ((\exists \bar{v}. s = C(\bar{v})) \wedge x = \text{sel}_C^i(s)) \vee (\neg(\exists \bar{v}. s = C(\bar{v})) \wedge x = \text{sel}_C^i(s)) \\ &\leftrightarrow (\exists \bar{v}. s = C(\bar{v}) \wedge x = \text{sel}_C^i(s)) \vee (\neg(\exists \bar{v}. s = C(\bar{v})) \wedge x = \text{sel}_C^i(s)) \\ &\leftrightarrow (\exists \bar{v}. s = C(\bar{v}) \wedge x = v_i) \vee (\neg(\exists \bar{v}. s = C(\bar{v})) \wedge x = T_C^i). \end{aligned}$$

The first equivalence is simply a case split on $\exists \bar{v}. s = C(\bar{v})$, the second one just moves a quantifier outward and the last one uses the definition of selectors. \square

Theorem (Theorem 3.4, repeated). *In the theory of (co)datatypes with standard semantics, a quantifier-free formula can be effectively transformed into an equisatisfiable one without selectors (but including quantifiers).*

Proof of Theorem 3.4. The transformation works in two steps. The first step is similar to the previous proof and replaces every selector term by an existentially quantified fresh variable. The second step is different from before obviously, but is similar to Ackermann's reduction for uninterpreted functions [1] (see [7, Section 11.2.1] for an exposition).

Step 1. Let ϕ be the formula to transform. Let S be a set of equations, initially empty. As long as ϕ contains a selector term, pick one that doesn't contain any nested selector terms. Let that term be $t \equiv \text{sel}_C^i(s)$. Replace t by a fresh variable v in ϕ and add the equation $v = \text{sel}_C^i(s)$ to S . This is repeated until ϕ contains no more selector terms.

Step 2. Let ϕ' be the formula after the first step and $S = \{v_j = \mathbf{sel}_{C_j}^{i_j}(s_j) \mid j \in \{1, \dots, n\}\}$ the set of equations of selectors. For each $j = 1, \dots, n$ define the following two formulae

$$\begin{aligned}\psi_j &\equiv \forall \bar{z}. s_j = C_j(\bar{z}) \rightarrow z_{i_j} = v_j \\ \chi_j &\equiv \bigwedge_{j'=1}^n s_j = s_{j'} \rightarrow v_j = v_{j'}\end{aligned}$$

where \bar{z} are fresh variables for each j . The idea is that the ψ_j 's specify that $\mathbf{sel}_{C_j}^i(C(\bar{z})) = z_i$ and the χ_j 's (which is also part of Ackermann's reduction) ensure functional consistency, i.e. that if the same selector is applied to equal terms then the results should be equal. The latter property is important if the selector $\mathbf{sel}_{C_j}^i$ is applied to a term that is not of the form $C(\dots)$ because then the result is unspecified but it still has to be consistent. Finally define ψ to be the following formula

$$\psi \equiv \bigwedge_{j=1}^n (\psi_j \wedge \chi_j).$$

Let ϕ be the original formula and ϕ' be the result of step 1. We claim that $\psi \wedge \phi'$, which does not contain selectors anymore, is equisatisfiable to ϕ .

First, suppose that ϕ is satisfiable. Let \mathcal{M} be a model of ϕ . We modify it to a model \mathcal{M}' by just changing the valuation of some variables, leaving the interpretation of selector functions untouched:

$$\mathcal{M}'(v) = \begin{cases} \mathcal{M}'(\mathbf{sel}_{C_j}^{i_j}(s_j)) & \text{if } v \equiv v_j \\ \mathcal{M}(v) & \text{otherwise.} \end{cases}$$

Note that the first case is well-defined since there are no cyclic dependencies between the variables \bar{v} by the design of Step 1. Then \mathcal{M}' satisfies each ψ_j since the interpretation of $\mathbf{sel}_{C_j}^i$ in \mathcal{M} (and hence \mathcal{M}') has to satisfy the selector axioms. It also satisfies each χ_j because $\mathcal{M}'(\mathbf{sel}_{C_j}^i)$ is a function that has to give the same result if applied to the same object. Hence if $\mathcal{M}'(s_j) = \mathcal{M}'(s_{j'})$ then $\mathcal{M}'(v_j) = \mathcal{M}'(v_{j'})$. This shows that \mathcal{M}' satisfies ψ .

Furthermore, it is not hard to see that \mathcal{M}' satisfies ϕ' as well. First of all, \mathcal{M}' satisfies ϕ because the variables v_j don't occur in ϕ since they were assumed to be fresh and \mathcal{M}' agrees with \mathcal{M} on everything else. Also note that by construction of \mathcal{M}' , every equation $v_j = \mathbf{sel}_{C_j}^{i_j}(s_j)$ in S is true of \mathcal{M}' . Since each $\mathbf{sel}_{C_j}^{i_j}(s_j)$ in ϕ is replaced by v_j in ϕ' , the interpretations of ϕ and ϕ' are the same in \mathcal{M}' . Hence ϕ' is true of \mathcal{M}' . Altogether, $\psi \wedge \phi'$ is satisfiable.

Next, suppose that $\psi \wedge \phi'$ is satisfiable and let \mathcal{M} be a model of it. We define a new model \mathcal{M}' modifying the interpretation of the selectors such that it satisfies:

$$\mathcal{M}'(\mathbf{sel}_{C_j}^i)(x) = \begin{cases} y_i & \text{if } x = C^{\mathcal{M}}(y_1, \dots, y_n) \\ \mathcal{M}(v_j) & \text{if } x = \mathcal{M}(s_j) \text{ for some } j \in \{1, \dots, n\} \\ \mathcal{M}(\mathbf{sel}_{C_j}^i)(x) & \text{otherwise.} \end{cases}$$

It is not immediately clear that this is well-defined. First, we check what happens if there is more than one j such that $x = \mathcal{M}(s_j)$. Suppose $x = \mathcal{M}(s_j) = \mathcal{M}(s_{j'})$. Since \mathcal{M} satisfies ψ and in particular χ_j , we have $\mathcal{M}(v_j) = \mathcal{M}(v_{j'})$. Hence it does not matter which j is chosen.

Next, we check that the first and second case are compatible. Suppose $x = C^{\mathcal{M}}(y_1, \dots, y_n) = \mathcal{M}(s_j)$. Since \mathcal{M} satisfies ψ and in particular ψ_j , we have $y_i = \mathcal{M}(v_j)$. So the first and second case are not in conflict. Altogether, this shows that \mathcal{M}' is well-defined.

Furthermore, it is not hard to see that \mathcal{M}' is a model of ϕ . First of all, \mathcal{M}' satisfies ϕ' because \mathcal{M} does, ϕ' does not contain any selectors, and \mathcal{M}' agrees with \mathcal{M} on everything else. Also note that by construction of \mathcal{M}' , every equation $v_j = \mathbf{sel}_{C_j}^{i_j}(s_j)$ in S is true of \mathcal{M}' . Since each $\mathbf{sel}_{C_j}^{i_j}(s_j)$ in ϕ is replaced

by v_j in ϕ' , the interpretations of ϕ and ϕ' are the same in \mathcal{M}' . Hence ϕ is true of \mathcal{M}' , and thus ϕ is satisfiable, which finishes the proof. \square

Theorem (Theorem 3.5, repeated). *A selector-free formula in the theory of (co)datatypes can be effectively transformed into an equisatisfiable formula in the extended theory of trees.*

Proof of Theorem 3.5. Without loss of generality, we can assume that the given formula ϕ contains only existential quantifiers, no universal ones, because $\forall x.\psi$ can be replaced by $\neg\exists x.\neg\psi$. Let ϕ' be the result of replacing each $\exists v_1, \dots, v_n.\psi$ occurring in ϕ with

$$\exists v_1 : s_1, \dots, v_n : s_n. \left(\bigwedge_{i \in \{1, \dots, n\}, s_i \text{ is a datatype}} \text{fin}(v_i) \right) \wedge \psi.$$

Finally let $u_1 : s_1, \dots, u_n : s_n$ be the free variables of ϕ' . Set

$$\phi'' \equiv \left(\bigwedge_{i \in \{1, \dots, n\}, s_i \text{ is a datatype}} \text{fin}(u_i) \right) \wedge \phi'.$$

Then it is clear that each model of ϕ can be reduced to a model of ϕ'' by forgetting the interpretation of the selector functions because the interpretation of datatypes are finite trees, thus satisfying all the additional $\text{fin}()$ -constraints. Conversely, each model of ϕ'' in the theory of trees can be extended to a model of ϕ by picking arbitrary selector functions. The $\text{fin}()$ -constraints ensure that the interpretation of each datatype variable is a finite tree, so it is, in fact, a model of ϕ . \square

B Omitted proofs from Section 4

Note that we make the standard assumption that the first-order language of trees is well-founded, in the sense that there is no infinite sequence s_0, s_1, \dots of sorts such that for all $i \in \mathbb{N}$, there is a function symbol f_i of arity $f_i : \dots \times s_{i+1} \times \dots \rightarrow s_i$.

Theorem (Theorem 4.1, repeated). *Given a signature (S, F, P) , Algorithm 1 correctly computes the sets S_{0F} and S_{0I} .*

Proof of Theorem 4.1. The validity of the fixed point computations is implied by the following two lemmas. \square

Lemma B.1. *The set S_{0I} of sorts without infinite trees is the least fixed point of the following function $f : P(S) \rightarrow P(S)$ where $P(\cdot)$ denotes the power set:*

$$f(X) := X \cup \{s \in S \mid \forall (\mathbf{g} : s_1 \times \dots \times s_n \rightarrow s) \in F_s : \forall i \in \{1, \dots, n\} : s_i \in X\}$$

Proof. Note that f is monotonic, so the least fixed point is guaranteed to exist by the Knaster-Tarski theorem. It is clear that S_{0I} is a fixed point because for each generator \mathbf{g} each parameter sort must only allow finite trees, i.e. be in S_{0I} , yielding $f(S_{0I}) = S_{0I}$. We claim that any sort s containing only finite trees of depth $\leq d$ is included in the set $f^{d+1}(\emptyset)$. This is proved inductively. For $d = 0$, this means s contains only constant symbols. Then the condition $\forall \mathbf{g} : s_1 \times \dots \times s_n \rightarrow s \in F_s : \forall i \in \{1, \dots, n\} : s_i \in \emptyset$ is vacuously true and $s \in f(\emptyset)$. For $d \geq 1$, it means that each argument of each generator has depth at most $d - 1$. Hence by the induction hypothesis, all generator arguments have sorts in $f^d(\emptyset)$. By the definition of f , this means that then $s \in f(f^d(\emptyset))$, proving the claim. Hence $S_{0I} = \cup_{d \in \omega} f^{d+1}(\emptyset)$, in other words, S_{0I} is the least fixed point of f . \square

Lemma B.2. *The set S_{0F} of sorts without finite trees is the greatest fixed point of the following function $f : P(S) \rightarrow P(S)$:*

$$f(X) := X \setminus \{s \in S \mid \exists (\mathbf{g} : s_1 \times \cdots \times s_n \rightarrow s) \in F_s : \forall i \in \{1, \dots, n\} : s_i \notin X\}$$

Proof. Note again that f is monotonic, so the greatest fixed point is guaranteed to exist by the Knaster-Tarski theorem. It is clear that S_{0F} is a fixed point because a sort with a generator \mathbf{g} where all parameter sorts of \mathbf{g} allow finite trees cannot contain only infinite trees. Hence all such sorts must be excluded, which is what f does. Therefore $S_{0F} = f(S_{0F})$. Similarly to the proof of Lemma B.1, it is easy to see by induction that any sort s containing a finite tree of depth d is excluded from the set $f^{d+1}(S)$. Hence $S_{0F} = \bigcap_{d \in \omega} f^{d+1}(S)$, in other words, S_{0F} is the greatest fixed point of f . \square

Theorem (Theorem 4.3, repeated). *Given a signature (S, F, P) , Algorithm 2 correctly computes the sets S_{FF} , S_{1I} , and S_{FI} . Furthermore it computes the set s_{fin} (the terms for the finite trees of sort s for $s \in S_{FF}$), and the set s_{infin} (the terms for the infinite trees of sort s for $s \in S_{FI}$). The latter makes use of the variables u_s (for $s \in S_{1I}$), standing for the unique infinite tree of s . The equations that uniquely determine these u_s are output in U_s .*

Proof of Theorem 4.3. The correctness of the fixed point computation of S_{FF} is implied by Lemma B.3. An analogous argument verifies the computation of s_{fin} , for $s \in S_{FF}$. Lemma B.4 shows the correctness of the fixed point computation of S_{1I} and the U_s . For the correctness proof of the fixed point computation of S_{FI} , we need Lemma B.5, which states that the finitely many infinite trees of sorts $s \in S_{FI}$ are all built from the unique infinite trees u_s with $s \in S_{1I}$. Using this result, Lemma B.6 proves the fixed point computation of S_{FI} correct. An analogous argument works for the sets s_{infin} for $s \in S_{FI}$. \square

Lemma B.3. *Let $S_{FF} \subseteq S$ be the set of sorts such that each $s \in S_{FF}$ has only finitely many finite trees. Let $F_s^{\text{infin}} = \{(\mathbf{g} : s_1 \times \cdots \times s_n \rightarrow s) \in F_s \mid \exists i \in \{1, \dots, n\} : s_i \in S_{0F}\}$ be the set of generators building only infinite trees. Then S_{FF} is the least fixed point of the following function $f : P(S) \rightarrow P(S)$:*

$$f(X) := X \cup S_{0F} \cup \{s \in S \mid |F_s \setminus F_s^{\text{infin}}| < \infty \wedge \forall (\mathbf{g} : s_1 \times \cdots \times s_n \rightarrow s) \in (F_s \setminus F_s^{\text{infin}}) : \forall i \in \{1, \dots, n\} : s_i \in X\}$$

where S_{0F} denotes the set of sorts with only infinite trees.

Proof. Note that f is monotonic, so the least fixed point is guaranteed to exist by the Knaster-Tarski theorem. Why is S_{FF} a fixed point? First, F_s^{infin} are generators that can only construct infinite trees. So any other generator (in $F_s \setminus F_s^{\text{infin}}$) can construct at least one finite tree. For there to be only finitely many finite trees in s , there have to be finitely many of the latter generators and for each such generator, each parameter sort s_i must have only finitely many finite trees. This explains the definition of the function f .

In fact, similarly to the proof of Lemma B.1, it is easy to see inductively that if a sort s has finitely many finite trees of depth at most d then $s \in f^{d+1}(\emptyset)$. Hence $S_{FF} = \bigcup_{d \in \omega} f^{d+1}(\emptyset)$, in other words, S_{FF} is in fact the least fixed point of f . \square

Lemma B.4. *Let $S_{1I} \subseteq S$ be the set of sorts such that each $s \in S_{1I}$ has exactly one infinite tree. Then S_{1I} is the greatest fixed point of the following function $f : P(S \setminus S_{0I}) \rightarrow P(S \setminus S_{0I})$:*

$$f(X) := \{s \in X \mid \exists (\mathbf{g} : s_1 \rightarrow s) \in F_s : s_1 \in X \wedge (\forall (\mathbf{g}' : s'_1 \times \cdots \times s'_n \rightarrow s) \in F_s \setminus \{\mathbf{g}\} : \forall i \in \{1, \dots, n\} : s'_i \in S_{0I})\}$$

Furthermore, the equations U_s , for $s \in S_{1I}$, that uniquely determine the unique infinite inhabitant u_s of $s \in S_{1I}$ are given by the least fixed point of the function f' , a mapping between families of sets of equations, indexed by $s \in S_{1I}$, which for each such s , is given by

$$(f'((X_s)_{s \in S_{1I}}))_s = \{u_s = \mathbf{g}_s(u_{s_1})\} \cup X_{s_1}$$

where $\mathbf{g}_s : s_1 \rightarrow s$ is the unique generator with $s_1 \in S_{1I}$.

Proof. Note that f is monotonic, so the greatest fixed point is guaranteed to exist by the Knaster-Tarski theorem. Why is S_{1I} a fixed point? If s has a unique infinite tree then it must start with some generator $g \in F_s$. If g had more than one parameter then the choice of the other parameter would create at least two infinite inhabitants, contradiction. So g has only one parameter. Furthermore every other generator g' can only create finite trees because otherwise we would lose uniqueness of the infinite tree. The function f removes all sorts from X that do not satisfy these criteria. Hence S_{1I} is a fixed point.

Conversely, if a sort $s \in S \setminus S_{0I}$ has two distinct infinite trees then they have to differ at some finite depth d . We claim that $s \notin f^{d+1}(S \setminus S_{0I})$ because f removes sorts that have more than one infinite inhabitant and if those two inhabitants differ at depth d , this is detected after at most $d + 1$ applications of f . This can be proved by induction, similarly to the proof of Lemma B.1. Hence $S_{1I} = \bigcap_{d \in \omega} f^{d+1}(S \setminus S_{0I})$, in other words, S_{1I} is the greatest fixed point of f .

Why is U_s a fixed point of f' ? The equation $u_s = g_s(u_{s_1})$ must be true by the above arguments. In order to describe u_{s_1} uniquely, we need the equations U_{s_1} as well. Thus the U_s are a fixed point of f' . They are, in fact, the least fixed point because we are interested in the smallest set of equations describing the u_s . \square

Lemma B.5. *Let s be a sort with at least one but only finitely many infinite trees. Then each infinite tree of s can be described by a term containing only variables $u_s : s$ with $s \in S_{1I}$, each representing the unique infinite tree of sort s .*

Proof. Proof by induction on the number $\#s_{\text{infin}}$ of infinite trees of sort s . If $\#s_{\text{infin}} = 1$ then s has a unique infinite tree represented by u_s and the statement is trivial. Hence suppose $\#s_{\text{infin}} \geq 2$. Let $a = f(b_1, \dots, b_n)$ be an infinite tree with subtrees $b_1 : s_1, \dots, b_n : s_n$.

Suppose $n \geq 2$. Then without loss of generality, assume that b_1 is an infinite subtree. Since s_2 has at least two generators, the number of infinite inhabitants of s_1 is at most $\#s_{\text{infin}}/2 < \#s_{\text{infin}}$. By induction hypothesis, b_1 has the desired form. The same argument works for other infinite subtrees of a . For each finite subtree, there is a ground term describing it. Hence a has the desired form.

Next, suppose $n = 1$, i.e. $a = f(a_1)$. If s has another infinite tree $a' = f'(a'_1)$ starting with a different function symbol f' then the sort of a_1 has less than $\#s_{\text{infin}}$ infinite trees, and the induction hypothesis gives us the desired form for a_1 and thus for a . Otherwise, all infinite trees of s start with the same function symbol f . We can apply the same argument to a_1 and see that we can either proceed as above or all the infinite trees of the sort of a_1 must have the form $a_1 = f_1(a_2)$. If we keep repeating this argument, there are three cases.

Case 1. There are b_1, \dots, b_n with $n \geq 2$ such that $a = f(f_1(\dots f_m(b_1, \dots, b_n) \dots))$. Then the first argument from above provides a term for a .

Case 2. There is an a' of sort s' such that $a = f(f_1(\dots f_m(a') \dots))$ and s' has two infinite trees starting with different generators. Then the second argument from above provides a term for a .

Case 3. There is no such a' , meaning that the tree a is uniquely determined, as an infinite path of unary function symbols. But then s only contains one infinite tree, contradiction. So this case cannot occur. \square

Lemma B.6. *Let $S_{FI} \subseteq S$ be the set of sorts such that each $s \in S_{FI}$ has only finitely many infinite trees. Let $F_s^{\text{infin}} = \{(g : s_1 \times \dots \times s_n \rightarrow s) \mid \exists i \in \{1, \dots, n\} : s_i \notin S_{0I}\}$ be the set of generators that can construct infinite trees. Then S_{FI} is the least fixed point of the following function $f : P(S) \rightarrow P(S)$:*

$$f(X) := X \cup S_{0I} \cup S_{1I} \cup \{s \in S \mid |F_s^{\text{infin}}| < \infty \wedge \forall g : s_1 \times \dots \times s_n \rightarrow s \in F_s : \\ \forall i \in \{1, \dots, n\} : s_i \in S_{0I} \vee (s_i \in X \wedge (\forall j \in \{1, \dots, n\} \setminus \{i\} : s_j \in S_{FI} \cap X))\}.$$

Proof. Note that f is monotonic, so the least fixed point is guaranteed to exist by the Knaster-Tarski theorem. Why is S_{FI} a fixed point? First of all, it is clear that $S_{0I} \cup S_{1I} \subseteq S_{FI}$. Furthermore, for the sort s to have finitely many infinite trees, there have to be finitely many generators F_s^{infin} that can construct infinite trees. Additionally, for each such generator $g : s_1 \times \cdots \times s_n \rightarrow s \in F_s^{\text{infin}}$, there have to be finitely many infinite trees starting with g .

It is easier to describe the negation of this: If a generator g starts infinitely many infinite trees, there must be a parameter i such that s_i contains infinite trees and one of the following: (1) s_i containing infinitely many infinite trees or (2) one of the other s_j containing infinitely many trees. In either case, this leads to infinitely many infinite trees starting with g . This can be formulated as

$$\exists i \in \{1, \dots, n\} : s_i \notin S_{0I} \wedge (s_i \notin S_{FI} \vee \exists j \in \{1, \dots, n\} \setminus \{i\} : s_j \notin S_{FF} \cap S_{FI}).$$

The negation of this is what is written in the above function definition. Hence S_{FI} is a fixed point of f .

Next, we show that S_{FI} is the least fixed point. Let s be a sort with finitely many infinite trees and a such a tree. By Lemma B.5, there is a term t_a describing a , containing only variables $u_i : s_i$ representing the unique infinite tree of s_i . We always choose t_a to be of minimal depth among those terms. By definition of f , each $s_i \in f(\emptyset)$. Let a be the infinite tree in S_{FI} such that its corresponding t_a has maximal depth d . Then one can see inductively, as in the proof of Lemma B.1, that $s \in f^{d+1}(\emptyset)$. Hence $S_{FI} = \bigcup_{d \in \omega} f^{d+1}(\emptyset)$, in other words, S_{FI} is in fact the least fixed point of f . \square

C Supplementary material for Section 5

Throughout this section, we assume the variable convention that bound variables of terms occurring in a certain mathematical context (like definitions and proofs) are assumed to be distinct and different from the free variables. Furthermore, when talking about reachability in a formula $\exists \bar{x}. \alpha$ where α is a basic formula, we mean reachability in α from the free variables of the whole formula.

We are also going to need the *Unique Solution Axiom* [5, Axiom 3 in Section 3.2], which states that for any sequence of distinct variables \bar{z} and non-variable terms t_i containing only the variables \bar{x} and \bar{z} , we have

$$\forall \bar{x}. \exists ! \bar{z}. \bigwedge_i z_i = t_i$$

in the extended theory of trees. This is proved in [5, Theorem 3.3.1].

Theorem (Theorem 5.2, repeated). *The function SOLVEBASIC($(v_0, \dots, v_n), \alpha$) from Algorithm 4 correctly solves basic formulae α containing the variables v_0, \dots, v_n , i.e. it turns α into an equivalent solved formula (with respect to the variable ordering $v_0 < \dots < v_n$) or returns false if none exists.*

Proof of Theorem 5.2. Most parts of the algorithm (the numbered rules) are taken from [5, Section 4.6] and the fact that the result satisfies property (1) of Definition 5.1 is proven there. (Note that Rule 6 is used only for bookkeeping in [5], which is why it is not needed in our formulation of the algorithm.) The two additional rules (*) involve variables $u : s$ where s is a sort without infinite, respectively finite, trees. Obviously, $\text{fin}(u)$ is always, respectively never, satisfied in those cases. Therefore, the result satisfies property (2) of Definition 5.1 as well. \square

Theorem (Theorem 5.7, repeated). *Let $\phi \equiv \neg(\exists \bar{x}. \alpha \wedge \bigwedge_i \phi_i)$ be a normal formula of depth at most 2 with free variables \bar{v} . Let I be the result of FINDINSTANTIATION(\bar{v}, ϕ) from Algorithm 3. If I is “none”, then there is no instantiable variable. Otherwise, let u be the first instantiable variable found in*

Algorithm 4 Algorithm for solving a basic formula α with free variables $v_0 < \dots < v_n$. The rules 1–10 are taken from [5]. The two rules in blue at the end are new.

function SOLVEBASIC($(v_0, \dots, v_n), \alpha$)
 let $<$ be the ordering where $v_0 < \dots < v_n$
repeat
 if α is $u = u \wedge \alpha'$ **then** $\alpha \leftarrow \alpha'$ ▷ Rule 1 (numbering as in [5, Section 4.6])
 if α is $u = v \wedge \alpha'$ and $u < v$ **then** $\alpha \leftarrow v = u \wedge \alpha'$ ▷ Rule 2
 if α is $v = u \wedge v = t \wedge \alpha'$ and $u < v$ **then** $\alpha \leftarrow v = u \wedge u = t \wedge \alpha'$ ▷ Rule 3
 if α is $u = f(\bar{y}) \wedge u = g(\bar{z}) \wedge \alpha'$ and $f \not\equiv g$ **then return false** ▷ Rule 4
 if α is $u = f(\bar{y}) \wedge u = f(\bar{z}) \wedge \alpha'$ **then** $\alpha \leftarrow u = f(\bar{y}) \wedge \bar{y} \equiv \bar{z} \wedge \alpha'$ ▷ Rule 5
until no changes in the last iteration
repeat
 if α is $\text{fin}(u) \wedge \text{fin}(u) \wedge \alpha'$ **then** $\alpha \leftarrow \text{fin}(u) \wedge \alpha'$ ▷ Rule 7
 if α is $v = u \wedge \text{fin}(v) \wedge \alpha'$ and $u < v$ **then** $\alpha \leftarrow v = u \wedge \text{fin}(u) \wedge \alpha'$ ▷ Rule 8
 if α is $\text{fin}(u) \wedge \alpha'$ and u is properly reachable from u **then return false** ▷ Rule 9
 if α is $u = f(\bar{y}) \wedge \text{fin}(u) \wedge \alpha'$ **then** $\alpha \leftarrow u = f(\bar{y}) \wedge \overline{\text{fin}(\bar{y})} \wedge \alpha'$ ▷ Rule 10
 if α is $\text{fin}(u) \wedge \alpha'$ and $u : s$ with $s \in S_{0I}$ **then** $\alpha \leftarrow \alpha'$ ▷ (*)
 if α is $\text{fin}(u) \wedge \alpha'$ and $u : s$ with $s \in S_{0F}$ **then return false** ▷ (*)
until no changes in the last iteration
return α

FINDINSTANTIATION. Then ϕ is equivalent to the following conjunction of normal formulae, in which the variable u is no longer instantiable:

$$\bigwedge_{(\exists \bar{z}. \psi) \in I} \neg(\exists \bar{x} \bar{z}. \alpha \wedge \psi \wedge \bigwedge_i \phi_i).$$

Proof of Theorem 5.7. We first show that in each case, the result I of the call to FINDINSTANTIATION satisfies $\alpha \rightarrow \bigvee_{\psi \in I} \psi$. For the first return statement in FINDINSTANTIATION, this is clear because if there are finitely many generators of s then one of them has to be used to construct a tree of sort s . For the second return statement, it is clear because u can have only finitely many values, so if I contains formulae describing each possible value then the disjunction over all of them must be true. For the third return statement, note that $\text{fin}(u)$ occurs in α , so u has to be finite. Hence I only contains formulae describing each finite value of s , and we have $\text{fin}(u) \rightarrow \bigvee_{\psi \in I} \psi$. Finally, consider the fourth return statement. The variable u has to represent either a finite tree, meaning $\text{fin}(u)$ or one of the finitely many infinite trees in s . Again, we find that $\bigvee_{\psi \in I} \psi$ holds. Since $\alpha \rightarrow \bigvee_{\psi \in I} \psi$ holds in each case, ϕ is equivalent to:

$$\begin{aligned} & \neg \left(\exists \bar{x}. \alpha \wedge \left(\bigvee_{\psi \in I} \psi \right) \wedge \bigwedge_i \neg(\exists \bar{y}_i. \beta_i) \right) \\ & \leftrightarrow \neg \left(\exists \bar{x}. \alpha \wedge \left(\bigvee_{(\exists \bar{z}. \psi') \in I} (\exists \bar{z}. \psi') \right) \wedge \bigwedge_i \neg(\exists \bar{y}_i. \beta_i) \right) \\ & \leftrightarrow \neg \left(\bigvee_{(\exists \bar{z}. \psi') \in I} \exists \bar{x} \bar{z}. \alpha \wedge \psi' \wedge \bigwedge_i \neg(\exists \bar{y}_i. \beta_i) \right) \end{aligned}$$

$$\leftrightarrow \bigwedge_{(\exists \bar{z}. \psi') \in I} \neg \left(\exists \bar{x} \bar{z}. \alpha \wedge \psi' \wedge \bigwedge_i \neg (\exists \bar{y}_i. \beta_i) \right)$$

By the construction of I , the instantiable variable u found in the algorithm is no longer instantiable in this transformed formula. \square

In the following, we will often have to show that for sorts s with infinitely many trees of sort s , there is a tree that contradicts a certain set of finitely many equations. The following lemma formalizes this.

Lemma C.1. *Let v be a variable of sort $s \notin S_{FI} \cap S_{FF}$. Let β_i be a family of solved basic formula, indexed by $i = 1, \dots, m$, where v is properly reachable from itself in each β_i . Let T be a finite set of trees of sort s . Then there is a tree v^* of sort s such that $v^* \notin T$ and each β_i is made false by any valuation with value v^* for v . (Roughly speaking, β_i are “forbidden recursive equations”, T are “forbidden values” and v^* avoids both.) Furthermore, v^* can be chosen to be finite if $s \notin S_{FF}$ and infinite if $s \notin S_{FI}$.*

Proof. If $s \notin S_{FF}$ then there is a finite tree v^* not in the finite set T . This value for v also makes each β_i false because it is finite but v is properly reachable from itself in β_i and thus infinite.

For the other case, $s \notin S_{FI}$, we need the concept of *contradicting a formula at a certain depth*. Let β be a basic formula containing a subformula of the form

$$x_0 = \mathbf{g}_0(\dots, x_1, \dots) \wedge x_1 = \mathbf{g}_1(\dots, x_2, \dots) \wedge \dots \wedge x_{n-1} = \mathbf{g}_{n-1}(\dots, x_0, \dots).$$

Let \mathcal{V} be a valuation of the variables of β that maps x_i to a tree t and let \bar{y}, \bar{z} be variables such that the equation with left-hand side x_i in β is $x_i = \mathbf{g}_i(\bar{y}, x_{i+1}, \bar{z})$. (We view the indices of x modulo n .) Next, we define

$$\text{contra}(x_i, t, \beta, \mathcal{V}) = \begin{cases} 0 & \text{if } \mathbf{g}_i \text{ is not the root of } t \\ \text{contra}(x_{i+1}, t', \beta) + 1 & \text{if } t' \text{ is the subtree of } t \text{ with } t = \mathbf{g}_i^{\mathcal{V}}(\overline{\mathcal{V}(\bar{y})}, t', \overline{\mathcal{V}(\bar{z})}) \\ \infty & \text{otherwise} \end{cases}$$

and say that *the value t for x_i contradicts β at depth $\text{contra}(x_i, t, \beta, \mathcal{V})$ under the valuation \mathcal{V}* . Intuitively, this means that when picking the value t for x_i and checking the equations in β , we notice a problem at depth d of the tree. If no valuation is specified, we define

$$\text{contra}(x_i, t, \beta) := \max_{\mathcal{V}} \text{contra}(x_i, t, \beta, \mathcal{V})$$

where \mathcal{V} ranges over valuations sending x_i to t , and say that *the value t for x_i contradicts β at depth $\text{contra}(x_i, t, \beta)$* . If $\text{contra}(x_i, t, \beta) < \infty$, we say that *the value t for x_i contradicts β* . Note that only the nodes up to depth d are relevant for contradicting β at depth d .

Phrased in this new terminology, our goal is to prove that there is a tree $v^* \notin T$ of sort s such that the value v^* for v contradicts β_1, \dots, β_m . In the following, we will iteratively construct a sequence of injections $f_0, \dots, f_m : s^{\mathcal{T}} \rightarrow s^{\mathcal{T}}$, and of infinite sets $W_0, \dots, W_m \subseteq s^{\mathcal{T}}$, such that for each $i \in \{1, \dots, m\}$ and all $w \in W_i$, the value $f_i(w)$ for v contradicts all the formulae β_1, \dots, β_i .

The base case is easy, simply define f_0 as the identity function, and W_0 as the set of infinite trees of sort s . Next, suppose f_i and W_i are defined with the desired property. If the value $f_i(w)$ for v contradicts β_{i+1} for all $w \in W_i$, we can simply use $W_{i+1} = W_i$, and $f_{i+1} = f_i$. Otherwise there is a $w \in W_i$ such that there is a valuation \mathcal{V} sending v to $f_i(w)$ that makes β_{i+1} true. Since v is reachable from itself, β_{i+1} contains a subformula

$$v = \mathbf{g}_0(\dots, x_1, \dots) \wedge x_1 = \mathbf{g}_1(\dots, x_2, \dots) \wedge \dots \wedge x_{n-1} = \mathbf{g}_{n-1}(\dots, v, \dots).$$

Let $t = f_i(w)$. We are going to label the nodes of t with the corresponding variables of β_{i+1} . That is to say, we label the root of t with v , the child node corresponding to x_1 with x_1 , and so on, such that the

labeled nodes form an infinite path labeled $v, x_1, \dots, x_{n-1}, v, x_1, \dots$ in the tree t . Let d be an integer such that the value t for v contradicts each β_1, \dots, β_i at depth at most d . It exists because each the value t for v contradicts each β_1, \dots, β_i at some finite depth by the induction hypothesis. Let n be a node in t labeled v , at a depth $> d$. The subtree rooted at n must be t again because otherwise, β_{i+1} would not be true. If we replace this subtree by a different subtree, β_{i+1} cannot be satisfied under any valuation of the variables x_1, \dots, x_n because the values of the latter are determined by other subtrees of t . Hence let $W_{i+1} = W_i \setminus \{t\}$ (which is also infinite) and $f_{i+1}(w)$ be the function returning t but with the subtree rooted at n replaced by w . By construction, the value $f_{i+1}(w)$ for v contradicts β_{i+1} . Since for each $w \in W_{i+1}$, all nodes of t and $f_{i+1}(w)$ agree up to depth d , the value $f_{i+1}(w)$ for v also contradicts β_1, \dots, β_i at depth at most d , as desired.

At the end of this iterative process, we obtain an infinite set W_m and an injection $f_m : s^{\mathcal{T}} \rightarrow s^{\mathcal{T}}$ such that for all $w \in W_m$, the value $f_m(w)$ for v contradicts β_1, \dots, β_m . The set $W := \{f_m(w) \mid w \in W_i\}$ is infinite because f_m is injective. Furthermore, each tree $t \in W$ contradicts all β_i . Since T is finite, there is a tree $v^* \in W \setminus T$. \square

Theorem (Theorem 5.12, repeated). *Let ϕ be a fully simplified formula. If ϕ has no free variables then $\phi \equiv \text{true}$. Otherwise both ϕ and $\neg\phi$ are satisfiable in the theory of trees.*

Proof of Theorem 5.12. The formula ϕ has the form

$$\exists \bar{x}. \alpha \wedge \bigwedge_{i \in I} \neg(\exists \bar{y}_i. \beta_i).$$

First consider the case of no free variables. Then no variable can be reachable in $\exists \bar{x}. \alpha$, hence by the reachability condition (5) of Definition 5.10, \bar{x} is empty. This implies that α is just true because it cannot mention any variables. The same argument applied to each $\exists \bar{y}_i. \beta_i$ means that \bar{y}_i is empty and $\beta_i \equiv \text{true}$. By condition (3) of Definition 5.10, each β_i must include a conjunct not occurring in α , hence $I = \emptyset$. Altogether, we have $\phi \equiv \text{true}$

If ϕ contains free variables, it is enough to find a valuation for the free variables such that ϕ is true in the theory of trees and another one such that ϕ is false in the theory of trees. To find a valuation that makes ϕ false, consider the following: If α contains a free variable z , it can be made false like this.

- If α contains $z = w$, then $z > w$ according to the variable ordering since α is solved. Hence w is also a free variable and α can be made false by instantiating z and w with different trees.
- If $z = f(\bar{w})$ occurs in α , it is enough to instantiate z with a tree not starting with f to make α false, which is always possible because each sort has at least two generators.
- If $w = t$ with t containing z occurs in α , this equation must be reachable in $\exists \bar{x}. \alpha$ by condition (5) of Definition 5.10. This means that there is an equation of the form $z' = \dots$ in α , with z' free and w reachable from z' . This situation was already handled in one of the previous two cases.
- If $\text{fin}(z)$ occurs in α , simply instantiate z to an infinite tree (which is possible by condition (2) of solved basic formulae) to make α false.

Otherwise, α contains no free variables, so \bar{x} is empty and α is true by the same argument as before. Since ϕ contains a free variable, there must be a β_i that contains a free variable, so is nonempty. Since β_i is a solved basic formula, it is satisfiable by Lemma C.2. Hence there is a valuation of free variables that makes $\neg\exists \bar{y}_i. \beta_i$ false. Then the same valuation makes ϕ false.

Next, we want to find a valuation making ϕ true. Let β_i^* be β_i with all conjuncts occurring in α removed. Our goal is to find a valuation of the free variables and \bar{x} that makes α true and every $\exists \bar{y}_i. \beta_i^*$ false (since we cannot make the parts of β_i that also occur in α false). Let \bar{x}_{lhs} me the variables from \bar{x} that occur on the left-hand side of an equation in α . The valuation for these variables will be picked

last because it is uniquely determined by the Unique Solutions Axiom, once the valuation for the other variables is chosen. So the equations of α are taken care of.

If $\text{fin}(v)$ occurs in α then any equation $v = f(\bar{w})$ occurring in any β_i^* is automatically false because v has to be properly reachable from itself (otherwise v would be instantiable), but then v cannot be finite. So the only equations with v on the left-hand side in β_i that we care about are $v = w_i$ for other variables w_i . In this case, each w_i is also a free variable because $v > w_i$ by the variable ordering, and the sort of v has infinitely many finite trees because otherwise v would be instantiable. Thus it is always possible to find a valuation that contradicts all these finitely many equations of the form $v = w_i$ by picking a value for v that is different from the values picked for all the w_i . This proves that we can always make α true.

Next, we do a case analysis on the β_i^* that have not been made false yet. By reachability, each β_i^* has to contain $\text{fin}(v)$ or $v = t$ for a free variable v . Then $v \notin \bar{x}_{lhs}$ because β_i is solved. We can assume that $\text{fin}(v)$ does not occur in α because this case was already discussed above. For each such free variable v , we do the following case analysis:

- Suppose there is a β_i^* that contains $\text{fin}(v)$. If β_i^* also contains an equation, then the following cases apply and suffice to make it false. So suppose β_i^* only contains fin-constraints. Then the sort of v has infinitely many infinite trees because otherwise v would be instantiable. This makes the following cases work, by restricting the set of possible values for v to the set of infinite trees. Using such a value also makes $\text{fin}(v)$, and thus β_i^* , false as desired.
- Suppose $v = f(\bar{w})$ occurs in some β_i^* . Then v must be properly reachable from itself in β_i because otherwise, it would be instantiable. If the sort of v had only finitely many trees then v would be instantiable, contradiction. Hence the sort of v has infinitely many trees. Since the previous cases are already handled, we can assume that the only constraints on v in all the β_j^* 's are of the form $v = f(\bar{w})$ with v properly reachable from itself in β_j^* or $v = w$. Since the sort of v has infinitely many trees, it is possible to contradict all these constraints by Lemma C.1
- Suppose $v = w$ occurs in some β_i . Since the previous cases are already handled, we can assume that the only constraints on v from the β_i^* 's are of the form $v = w_i$ for variables w_i . Then each w_i is also a free variable because $v > w_i$ by the variable ordering, and the sort of v has infinitely many trees because otherwise v would be instantiable. Thus it is always possible to find a valuation that contradicts all these finitely many equations of the form $v = w_i$ by picking a value for v that is different from the values picked for all the w_i .

This case analysis shows that we can make all the β_i^* false. Thus it is always possible to find a valuation that makes ϕ true, as desired. \square

Lemma C.2. *Any solved basic formula is satisfiable.*

Proof. Let the basic formula be given by $\overline{v = t} \wedge \overline{\text{fin}(u)}$. By condition (2) of solved basic formulae (Definition 5.1), each \bar{u} can be given the value of some finite tree. Since the variables \bar{v} and \bar{u} are disjoint, the Unique Solution Axiom tells us that $\exists \bar{v}. \overline{v = t}$ is satisfiable for this valuation of \bar{u} . \square

Theorem (Theorem 5.13, repeated). *Given a formula ϕ , the function $\text{SOLVE}(\phi)$ from Algorithm 5 returns true, false, or a disjunction of fully simplified formulae that is equivalent to ϕ in the extended theory of trees. In particular, if ϕ is closed, it returns true or false.*

Proof of Theorem 5.13. The proof of this is quite involved and will take up the rest of this section. The function $\text{SOLVE}(\phi)$ first normalizes $\neg\phi$ and then solves its basic formula. If the latter contains a contradiction, ϕ is unsatisfiable. Otherwise, the function SOLVENESTED recursively solves $\neg\phi$: it returns a set of solved normal formulae $\{\psi_1, \dots, \psi_n\}$ such that $\neg\phi$ is equivalent to $\bigwedge_{i=1}^n \psi_i$. It works very similarly to

Algorithm 5 Extension of Djelloul, Dao, and Frühwirth’s algorithm [5, Section 4.6] for transforming a normal formula into an equivalent conjunction of solved formulae. The added part is in blue.

function SOLVE(ϕ)
 $\tilde{\phi} \leftarrow \text{NORMALIZE}(\neg\phi)$ ▷ cf. Theorem 5.4
 $\bar{v} \leftarrow$ the free variables of $\tilde{\phi}$ in some fixed order
let $\neg(\exists\bar{x}. \alpha \wedge \bigwedge_i \phi_i) \equiv \tilde{\phi}$ where α is a basic formula and ϕ_i are normal formulae
 $\alpha \leftarrow \text{SOLVEBASIC}(\bar{v}\bar{x}, \alpha)$ ▷ cf. Algorithm 4
if $\alpha \equiv \text{false}$ **then return** false
 $\{\psi_1, \dots, \psi_n\} \leftarrow \text{SOLVENESTED}(\bar{v}, \neg(\exists\bar{x}. \alpha \wedge \bigwedge_i \phi_i))$
if $n = 0$ **then return** false
if each ψ_i is of the form $\neg(\text{true})$ **then return** true
let $\neg(\exists\bar{x}_i. \alpha_i \wedge \bigwedge_{j \in J_i} \neg(\exists\bar{y}_{ij}. \beta_{ij})) \equiv \psi_i$ for each i
remove all conjuncts (of the form $u = t$ or $\text{fin}(u)$) from each β_{ij} that already occur in α_i
return $\bigvee_{i=1}^n (\exists\bar{x}_i. \alpha_i \wedge \bigwedge_{j \in J_i} \neg(\exists\bar{y}_{ij}. \beta_{ij}))$

function SOLVENESTED(\bar{v}, ϕ)
let $\neg(\exists\bar{x}. \alpha \wedge \bigwedge_i \neg(\exists\bar{y}_i. \phi_i)) \equiv \phi$
for each i **do**
let $\beta_i \wedge \bigwedge_j \psi_{ij} \equiv \phi_i$ where β_i is a basic formula and ψ_{ij} are normal formulae
 $\beta_i \leftarrow \alpha \wedge \beta_i$ ▷ Rule 12 in [5, Section 4.6]
 $\beta_i \leftarrow \text{SOLVEBASIC}(\bar{v}\bar{x}\bar{y}_i, \beta_i)$
if $\beta_i \equiv \text{false}$ **then** $\Psi_i \leftarrow \emptyset$
else
replace each $u = t$ in β_i by $u = s$ if $u = s$ occurs in α ▷ Rule 13
 $\Psi_i \leftarrow \text{SOLVENESTED}(\bar{v}\bar{x}, \neg(\exists\bar{y}_i. \beta_i \wedge \bigwedge_j \psi_{ij}))$
return SOLVEFINAL($\bar{v}, \neg(\exists\bar{x}. \alpha \wedge \bigwedge_i (\bigcup_j \Psi_j))$)

function SOLVEFINAL(\bar{v}, ϕ)
let $\neg(\exists\bar{x}. \alpha \wedge \bigwedge_i \phi_i) \equiv \phi$
if there is an i such that $\phi_i \equiv \neg(\exists\bar{y}. \alpha)$ **then return** \emptyset ▷ Rule 14
if depth of ϕ is 3 **then** ▷ Rule 16 (depth reduction)
choose a j such that ϕ_j has depth 2
let $\neg(\exists\bar{y}. \beta \wedge \bigwedge_k \neg(\exists\bar{z}_k. \gamma_k)) \equiv \phi_j$
 $\psi \leftarrow \neg(\exists\bar{x}. \alpha \wedge \neg(\exists\bar{y}. \beta) \wedge \bigwedge_{i \neq j} \phi_i)$
 $\chi_k \leftarrow \neg(\exists\bar{x}\bar{y}\bar{z}_i. \gamma_i \wedge \bigwedge_{i \neq j} \phi_i)$
return SOLVEFINAL(\bar{v}, ψ) $\cup \bigcup_k \text{SOLVENESTED}(\bar{v}, \chi_k)$

$I \leftarrow \text{FINDINSTANTIATION}(\bar{v}, \phi)$ ▷ cf. Algorithm 3
if $I \neq \text{none}$ **then**
return $\bigcup \{ \text{SOLVENESTED}(\bar{v}, \neg(\exists\bar{x}\bar{z}. \alpha \wedge \psi \wedge \bigwedge_i \phi_i)) \mid (\exists\bar{z}. \psi) \in I \}$ ▷ cf. Theorem 5.7
else
return $\{ \text{REMOVEUNREACHABLEPARTS}(\bar{v}, \phi) \}$ ▷ Rule 15 (cf. Algorithm 6)

Algorithm 6 Rule 15 from [5, Section 4.6], which removes unreachable variables and subformulae of a normal formula of depth at most 2.

function REMOVEUNREACHABLEPARTS($\bar{v}, \neg(\exists \bar{x}. \alpha \wedge \bigwedge_i \psi_i)$) ▷ Rule 15

$\bar{x}' \leftarrow$ the variables from \bar{x} reachable in α from the free variables \bar{v}

$\alpha' \leftarrow$ the conjuncts of $\exists \bar{x}. \alpha$ reachable from the free variables \bar{v}

$\alpha'' \leftarrow$ the fin-subformulae of $\exists \bar{x}. \alpha$ unreachable from the free variables \bar{v}

$\alpha''' \leftarrow$ the equations of $\exists \bar{x}. \alpha$ unreachable from the free variables \bar{v}

$\bar{x}''' \leftarrow$ the variables of $\exists \bar{x}. \alpha$ occurring on the LHS of an equation in α , and unreachable from \bar{v}

$\bar{x}'' \leftarrow \bar{x}$ without \bar{x}' and \bar{x}'''

for i **do**

Let $\neg(\exists y_i. \beta_i) \equiv \phi_i$

$\beta_i^* \leftarrow \beta_i$ with α'' removed

$\bar{y}'_i \leftarrow$ the variables of $\bar{x}''' \bar{y}_i$ in $\exists \bar{x}''' \bar{y}_i. \beta_i^*$ reachable from its free variables

$\beta'_i \leftarrow$ the conjuncts of $\exists \bar{x}''' \bar{y}_i. \beta_i^*$ reachable from its free variables

$K \leftarrow$ the set of indices i where no variable of \bar{x}'' occurs in β'_i

return $\{\neg \exists \bar{x}'. \alpha' \wedge \bigwedge_{i \in K} \neg(\exists \bar{y}'_i. \beta'_i)\}$

the original algorithm in [5, rules 12–16 in Section 4.6]. (Note that Rule 11 is used only for bookkeeping in [5], which is why it is not needed in our formulation of the algorithm.) The only change is the instantiation step, highlighted in Algorithm 5. The unchanged parts are proven correct in [5, Property 4.6.3].

The following lemmas prove the correctness of our change. Lemma C.4 establishes the termination of repeated instantiation steps. The termination of the unchanged parts of the original algorithm is shown in [5, Property 4.6.3]. The fact that the instantiation step is correct was proven in Theorem 5.7 already. Next, Lemma C.5 proves that the properties (1) to (4) of a solved formula (Definition 5.10) are satisfied when REMOVEUNREACHABLEPARTS is called. Lemma C.6 proves that the return value of REMOVEUNREACHABLEPARTS is correct. By construction, it satisfies property (5) as well, thus it is solved.

Since the return value of SOLVENESTED is a set of solved normal formulae $\{\psi_1, \dots, \psi_n\}$ such that $\neg\phi$ is equivalent to $\bigwedge_{i=1}^n \psi_i$, the original formula ϕ is equivalent to $\bigvee_{i=1}^n \neg\psi_i$, a disjunction of fully simplified formulae. In particular, if $n = 0$ then ϕ is always false. Conversely, if each ψ_i is \neg -true then ϕ is always true. In all other cases, we remove subformulae that were duplicated by Rule 12 in SOLVENESTED. This last simplification step is not strictly necessary: even without it, the results would be fully simplified formulae. Finally, we return the whole disjunction. \square

In order to prove the termination of repeated instantiations, we need the following concept.

Definition C.3 (depth of a variable). Let α be a solved basic formula. The *depth* of a variable v in α , denoted by $\text{depth}_\alpha(v)$, is defined as follows. If v is properly reachable from itself or doesn't occur on the left-hand side of an equation in α , its depth is 0. Else if $v = f(\bar{w})$ occurs in α , its depth is $\text{depth}_\alpha(v) := 1 + \max_i(\text{depth}_\alpha(w_i))$. Else if $v = w$ occurs in α , its depth is $\text{depth}_\alpha(v) := \text{depth}_\alpha(w)$.

Note that this is well-defined because of the ‘‘cycle check’’ using reachability in the definition.

Lemma C.4. *There are only finitely many instantiations (calls to FINDINSTANTIATION that do not return ‘‘none’’) happening in Algorithm 5. Hence the algorithm terminates.*

Proof. For a given normal formula $\phi \equiv \neg(\exists \bar{x}. \alpha \wedge \bigwedge_i \neg(\exists \bar{y}_i. \beta_i))$ of depth 2 with free variables \bar{v} , let X be the set of instantiable variables and define

$$N_j(\phi) = |\{v \in X \mid \max_i(\text{depth}_{\beta_i}(v)) = j\}|,$$

in other words, the number of instantiable variables with maximum depth j . Let k be the maximum integer such that $N_k(\phi) > 0$ and define

$$N(\phi) = (N_k(\phi), \dots, N_1(\phi), N_0(\phi)).$$

By the definition of FINDINSTANTIATION, the Instantiation Rule is only applied if $N(\phi) \neq (0, \dots, 0)$. We claim that the value of $N(\phi)$ decreases with respect to lexicographical order in each recursive call of SOLVENORMALIZED after every application of the Instantiation Rule. Note that it was proved in [5, Property 4.6.3] that when FINDINSTANTIATION is called, the normal formula ϕ satisfies conditions (1) to (3) of Definition 5.10.

Suppose the variable u returned by FINDINSTANTIATION was selected because there is an equation $u = f(\bar{w})$ in β_i^* where u is not properly reachable from u . Then u does not occur on a LHS in α because of the variable ordering: If $u = v$ occurred in α , it would also occur in β_i by condition (2) of solved basic formulae (Definition 5.1) and β_i would not be solved, violating condition (1). After instantiating u , i.e. adding the equation $u = g(\bar{z})$, the resulting basic formula $\alpha \wedge u = g(\bar{z})$ is therefore solved, so SOLVEBASIC does not change it at all. Next, Rule 12 copies α into each β_j . What can happen in Algorithm 4 now? If β_i contains $u = f(\bar{w})$, this leads to the situation $u = g(\bar{z}) \wedge u = f(\bar{w})$. If $f \neq g$, this is a conflict and β_i is removed from ϕ . Otherwise, that part of β_i is replaced with $\bar{z} \equiv \bar{w}$. Given that β_i was a solved basic formula before adding $\bar{z} \equiv \bar{w}$, the only applicable rule in β_i is Rule 2, switching the ordering of $z_k = w_k$ to $w_k = z_k$ if $w_k > z_k$. If there is another equation $w_k = t$, Rule 3 will change it to $w_k = z_k$ and $z_k = t$. Afterward, no more rules are applicable, and the resulting formula is solved. (In fact, Rule 8 or 10 could also be applied but this is irrelevant for the depths.) Denote the formulae resulting from the original ϕ and β_i by ϕ' and β'_i , respectively. By the above discussion, we have $\text{depth}_{\beta'_i}(z_k) = 0$ or $\text{depth}_{\beta'_i}(z_k) = \text{depth}_{\beta_i}(w_k) < \text{depth}_{\beta_i}(u)$. In either case, we have $\text{depth}_{\beta'_i}(z_k) \leq d - 1$ where $d = \max_i(\text{depth}_{\beta_i}(u))$. In other words, the depths of the newly introduced variables are smaller than the maximal depth of u . Thus $N_j(\phi) = N_j(\phi')$ for $j > d$ and $N_d(\phi') < N_d(\phi)$ since u is no longer instantiable. Therefore $N(\phi') < N(\phi)$, as desired.

Next, suppose the variable u returned by FINDINSTANTIATION was selected because u occurs in β_i^* and the sort s of u has only finitely many trees. Then by the same arguments as before, new variables \bar{z} are introduced in ϕ after adding $\exists \bar{z}. \gamma$ to α . However, since γ describes a single value for u , every variable among \bar{z}, u occurs on the left-hand side of an equation in the new α . Hence u is no longer instantiable and none of the newly introduced variables \bar{z} are. Hence the number of instantiable variables decreases and thus $N(\phi') < N(\phi)$ for the new formula ϕ' as desired.

Next, suppose the variable u returned by FINDINSTANTIATION was selected because $\text{fin}(u)$ occurs in α , u occurs in β_i^* and $s \in S_{FF}$. Then the same argument as in the previous case can be applied.

Next, suppose the variable u returned by FINDINSTANTIATION was selected because $s \in S_{FI}$ and there is a β_j^* consisting only of $\text{fin}(\cdot)$ -constraints, including $\text{fin}(u)$. After an instantiation of the form $\text{fin}(u)$ and subsequent simplification, $\text{fin}(u)$ will be removed from each β_j^* since α is contained in β_j by condition (2) of Definition 5.10, so u is no longer instantiable. Hence the number of instantiable variables, and thus $N(\phi)$, decreases. After an instantiation of the form $\exists \bar{z}. \gamma$ describing an infinite value for u , the variable u is also not instantiable anymore. Since all the additional variables \bar{z} occur on the left-hand side of an equation in $\exists \bar{z}. \gamma$, they are not instantiable either. Hence the number of instantiable variables, and thus $N(\phi)$, decreases.

Altogether, $N(\phi)$ decreases after each instantiation step. Therefore, only finitely many instantiations can happen. \square

Lemma C.5. *In Algorithm 5, when REMOVEUNREACHABLEPARTS is called, ϕ satisfies conditions (1) to (4) of a solved formula from Definition 5.10.*

Proof. From the proof of correctness of the unmodified algorithm [5, Property 4.6.3], which works the same until the instantiation step, it follows that up until that point, ϕ satisfies conditions (1) to (3). As soon as REMOVEUNREACHABLEPARTS is called, (4) is satisfied because otherwise FINDINSTANTIATION would find a variable violating (4). \square

Lemma C.6. *The function REMOVEUNREACHABLEPARTS from Algorithm 6 (Rule 15 in [5, Section 4.6]) is still correct in the context of the extended algorithm.*

Proof. As the previous lemma states, at the point where Rule 15 is applied, $\phi \equiv \neg(\exists \bar{x}. \alpha \wedge \bigwedge_{i=1}^n \neg(\exists \bar{y}_i. \beta_i))$ satisfies conditions (1) to (4) of Definition 5.10. Repeating what was stated at the beginning of this section, when talking about reachability in a formula $\exists \bar{x}. \alpha$ where α is a basic formula, we mean reachability in α from the free variables of the whole formula. As in the algorithm pseudocode, let

- \bar{x}' be the reachable variables of $\exists \bar{x}. \alpha$,
- \bar{x}''' the unreachable variables from \bar{x} that occur on the LHS of an equation in α ,
- \bar{x}'' the variables from \bar{x} that are not in $\bar{x}'\bar{x}'''$,
- α' be the reachable conjuncts of $\exists \bar{x}. \alpha$,
- α'' the unreachable $\text{fin}()$ -subformulae of $\exists \bar{x}. \alpha$,
- α''' the unreachable equations of $\exists \bar{x}. \alpha$,
- β_i^* the result of removing α'' from β_i ,
- \bar{y}'_i the reachable variables among $\bar{x}''\bar{y}_i$ in $\exists \bar{x}''\bar{y}_i. \beta_i^*$,
- β'_i the reachable conjuncts in $\exists \bar{x}''\bar{y}_i. \beta_i^*$
- $K \subseteq \{1, \dots, n\}$ the set of indices i such that $i \in K$ if and only if no variable of \bar{x}'' occurs in β'_i .

Then the claim is that $\neg(\exists \bar{x}. \alpha \wedge \bigwedge_{i=1}^n \neg(\exists \bar{y}_i. \beta_i))$ is equivalent to $\neg(\exists \bar{x}'. \alpha' \wedge \bigwedge_{i \in K} \neg(\exists \bar{y}'_i. \beta'_i))$.

First note that $\neg(\exists \bar{x}. \alpha \wedge \bigwedge_{i=1}^n \neg(\exists \bar{y}_i. \beta_i))$ is equivalent to

$$\neg(\exists \bar{x}'. \alpha' \wedge (\exists \bar{x}'' . \alpha'' \wedge (\exists \bar{x}''' . \alpha''' \wedge \bigwedge_{i=1}^n \neg(\exists \bar{y}_i. \beta_i))))$$

because the variables \bar{x}'' can only occur in α'' and the variables \bar{x}''' can only occur in α''' . By the Unique Solution Axiom and since α''' is a solved formula, we have $\exists! \bar{x}''' . \alpha'''$ in the extended theory of trees. According to [5, Property 3.1.11], the previous formula is equivalent to

$$\neg(\exists \bar{x}'. \alpha' \wedge (\exists \bar{x}'' . \alpha'' \wedge \bigwedge_{i=1}^n \neg(\exists \bar{x}''' . \alpha''' \wedge \exists \bar{y}_i. \beta_i))).$$

By our variable convention, no variable names conflict, so the innermost existential can be pulled outside:

$$\neg(\exists \bar{x}'. \alpha' \wedge (\exists \bar{x}'' . \alpha'' \wedge \bigwedge_{i=1}^n \neg(\exists \bar{x}''' \bar{y}_i. \alpha''' \wedge \beta_i))).$$

By condition (2) of Definition 5.10, the equations of α are included in each β_i . In particular, α''' is part of each β_i , which simplifies the formula to

$$\neg(\exists \bar{x}'. \alpha' \wedge (\exists \bar{x}'' . \alpha'' \wedge \bigwedge_{i=1}^n \neg(\exists \bar{x}''' \bar{y}_i. \beta_i))).$$

Note that $\beta_i^* \wedge \alpha'' \leftrightarrow \beta_i \wedge \alpha''$ by definition, so we can propagate α'' into the innermost existential formulae: $\alpha'' \wedge \bigwedge_{i=1}^n \neg(\exists \bar{x}'' \bar{y}_i. \alpha'' \wedge \beta_i) \leftrightarrow \alpha'' \wedge \bigwedge_{i=1}^n \neg(\exists \bar{x}'' \bar{y}_i. \alpha'' \wedge \beta_i^*)$; and back out, yielding:

$$\neg(\exists \bar{x}'. \alpha' \wedge (\exists \bar{x}'' . \alpha'' \wedge \bigwedge_{i=1}^n \neg(\exists \bar{x}'' \bar{y}_i. \beta_i^*))).$$

Since unreachable parts of a solved basic formula can be removed by the following Lemma C.7, this is equivalent to

$$\neg(\exists \bar{x}'. \alpha' \wedge (\exists \bar{x}'' . \alpha'' \wedge \bigwedge_{i=1}^n \neg(\exists \bar{y}'_i. \beta'_i))).$$

Since a variable from \bar{x}'' can only occur in β'_i if $i \notin K$, this is equivalent to

$$\neg \left(\exists \bar{x}'. \alpha' \wedge \left(\bigwedge_{i \in K} \neg(\exists \bar{y}'_i. \beta'_i) \right) \wedge \left(\exists \bar{x}'' . \alpha'' \wedge \bigwedge_{i \notin K} \neg(\exists \bar{y}'_i. \beta'_i) \right) \right).$$

To complete the proof, we show that the last conjunct $\exists \bar{x}'' . \alpha'' \wedge \bigwedge_{i \notin K} \neg(\exists \bar{y}'_i. \beta'_i)$ is always true. For this, it suffices to find valuations for \bar{x}'' satisfying α'' but none of $\exists \bar{y}'_i. \beta'_i$ for $i \notin K$. Note that since a variable from \bar{x}'' occurs in β'_i for all $i \notin K$, each such β'_i contains a conjunct of one of the following forms:

- $\text{fin}(v)$ for $v \in \bar{x}''$ and by the construction of β_i^* , $\text{fin}(v)$ does not occur in α'' ,
- $v = f(\bar{w})$ for $v \in \bar{x}''$,
- $v = w$ where $v \in \bar{x}''$ and $v > w$, implying $w \notin \bar{y}'_i$,
- $u = t$ where $v \in \bar{x}''$ occurs in t . Since it has to be reachable, that means that β'_i contains the conjunction $\bigwedge_{j=1}^k w_j = t_j$ with t_j containing w_{j+1} , $w_{k+1} \equiv v$, and $w_1 \notin \bar{y}'_i$. Since the case $w_1 \in \bar{x}''$ was already handled in a previous case, we can assume without loss of generality that w_1 is a free variable.

The goal now is to find a valuation of \bar{x}'' that satisfies α'' but that makes each of the above cases false, thus making $\neg \exists \bar{y}'_i. \beta'_i$ true. Fix a valuation for the free variables of the formula. Let $v : s$ be a variable from \bar{x}'' .

- If $\text{fin}(v)$ occurs in α'' then no β'_i can contain $v = f(\bar{w})$ because v is not instantiable and thus v would have to be properly reachable from itself, contradicting finiteness. If s only contains finitely many finite trees then v occurs in no β'_i because v is not instantiable. Then v can be given any finite value to make α'' true. Otherwise, v occurs only in equations of the form $u = t$ (reachable from some free variable w_1 as seen above) or $v = w$ in the β'_i . In the former case, to make the equation false, we pick a value for v that is different from the one that is determined by the fixed value of w_1 . In the latter case, we pick a value v different from the value of w . Since s contains infinitely many finite trees, it is possible to pick one as the value for v that contradicts all those finitely many equations.
- If $\text{fin}(v)$ does not occur in α'' and there is a β'_j containing only $\text{fin}()$ -constraints, among them $\text{fin}(v)$, then since v is not instantiable, we have $s \notin S_{FI}$. Thus there are infinitely many infinite trees of sort s . Since there are only finitely many equations of the form $v = w$, or $u = t$ with t containing v (reachable from some free variable as above), or $v = t$ with v properly reachable from itself in the β'_i , it is possible to find a value for v that contradicts all of them by Lemma C.1.
- If $\text{fin}(v)$ does not occur in α'' and there is no β'_j containing only $\text{fin}()$ -constraints, among them $\text{fin}(v)$, then there are two cases. If $s \in S_{FF} \cap S_{FI}$ then since v is not instantiable, no β'_i contains v , and there are no constraints to contradict, or α contains an equation $v = t$, in which case each β'_i also contains the same equation. Hence all the β'_i can be contradicted by picking a value different from t for v . Otherwise ($s \notin S_{FF} \cap S_{FI}$), there are infinitely many possible valuations for v while

Time to solve	selector semantics			
	standard		default values	
< 1 ms	534	13.35%	197	4.93%
< 10 ms	2241	56.04%	1415	35.38%
< 100 ms	3247	81.20%	3224	80.62%
< 1 s	3659	91.50%	3779	94.50%
< 10 s	3816	95.42%	3929	98.25%
timed out (> 10 s)	183	4.58%	70	1.75%
total	3999	100%	3999	100%

Table 1: Results of the SMT-LIB QF_DT benchmark suite: the number of benchmarks solved in the specified time limit (wall-clock time). The measurements were made on a notebook computer with an Intel® Core™ i5-8250U CPU and 16 GB of RAM.

there are only finitely many constraints of the form $u = t$ with t containing v (reachable from some free variable as above), or $v = w$, or $v = f(\bar{w})$ with v properly reachable from itself. Again, Lemma C.1 shows that it is possible to find a value for v that contradicts all of these constraints. We have shown above that by picking valuations for the variables from \bar{x}' as described above, each β'_i containing an equation is contradicted by the above valuation. If a β'_i contains only $\text{fin}()$ -constraints then at least one of those $\text{fin}(v)$ is contradicted as described above. This means that the above valuations for v make all the β'_i false, while satisfying α'' , independently of the values of the free variables.

This means that the formula

$$\exists \bar{x}'' . \alpha'' \wedge \bigwedge_{i \notin K} \neg(\exists \bar{y}'_i . \beta'_i)$$

is valid in the extended theory of trees. □

The above proof made use of the following lemma.

Lemma C.7. *Let \bar{x} be a vector of variables and α a solved basic formula. Let \bar{x}' be reachable variables in $\exists \bar{x} . \alpha$ and α' be the conjunction of equations and $\text{fin}()$ -formulae that are reachable in $\exists \bar{x} . \alpha$. Then in the theory of trees, $\exists \bar{x} . \alpha$ is equivalent to $\exists \bar{x}' . \alpha'$.*

Proof. Let \bar{x}'' be the unreachable variables in $\exists \bar{x} . \alpha$ that do not occur on the LHS of an equation of α and \bar{x}''' be the unreachable variables which do. Similarly, let α'' be the conjunction of unreachable $\text{fin}()$ -formulae and α''' be the conjunction of unreachable equations in $\exists \bar{x} . \alpha$. By the definition reachability, \bar{x}'' and \bar{x}''' do not occur in α' . Hence $\exists \bar{x} . \alpha$ is equivalent to

$$\exists \bar{x}' . \alpha' \wedge (\exists \bar{x}'' . \alpha'' \wedge (\exists \bar{x}''' . \alpha''')).$$

By the Unique Solution Axiom, $\exists \bar{x}''' . \alpha'''$ holds for any valuation of the free variables in the extended theory of trees. Hence the formula simplifies to

$$\exists \bar{x}' . \alpha' \wedge (\exists \bar{x}'' . \alpha'').$$

Since α'' contains only $\text{fin}()$ -formulae and since by condition (2) of Definition 5.1, they are all satisfiable, $\exists \bar{x}'' . \alpha''$ is true in the theory of trees as well. Hence the original formula is equivalent to $\exists \bar{x}' . \alpha'$, as desired. □