Inclusion Constraints over Non-empty Sets of Trees*

Martin Müller[1], Joachim Niehren[1] and Andreas Podelski[2]

[1] Programming System Lab,
Universität des Saarlandes, 66041 Saarbrücken, Germany
{mmueller,niehren}@ps.uni-sb.de
[2] Max-Planck-Institut für Informatik,
Im Stadtwald, 66123 Saarbrücken, Germany
podelski@mpi-sb.mpg.de

Abstract. We present a new constraint system called INES. Its constraints are conjunctions of inclusions $t_1 \subseteq t_2$ between first-order terms (without set operators) which are interpreted over non-empty sets of trees. The existing systems of set constraints can express INES constraints only if they include negation. Their satisfiability problem is NEXPTIME-complete. We present an incremental algorithm that solves the satisfiability problem of INES constraints in cubic time. We intend to apply INES constraints for type analysis for a concurrent constraint programming language.

1 Introduction

We propose a new constraint system called INES (Inclusions over Non-Empty Sets) and present an incremental algorithm to decide the satisfiability of INES constraints in time $O(n^3)$. INES constraints are conjunctions of inclusions $t_1 \subseteq t_2$ between first-order terms (without set operators) which are interpreted over the domain of non-empty sets of trees. In this paper we focus on sets of possibly infinite trees. All given results can be easily adapted to finite trees.

An INES-constraint $t_1 \subseteq t_2$ is satisfiable over non-empty sets if and only if $t_1 \not\subseteq \emptyset \wedge t_1 \subseteq t_2$ is satisfiable over arbitrary sets. Note that the constraint $t \not\subseteq \emptyset$ cannot be expressed by positive set constraints only [16]. The expressiveness of INES constraints is subsumed by that of set constraints with negation [9, 16]. In the case of finite trees, the satisfiability problem of set constraints with negation is known to be decidable [1, 13]; it is complete for nondeterministic exponential time [9, 10]. This result implies that the satisfiability problem of INES constraints over sets of finite trees is decidable. The corresponding problem for infinite trees has not been considered before.

We characterize the satisfiability of INES constraints by a set of axioms such that an INES constraint is satisfiable over non-empty sets if and only if it is satisfiable

* A summary has appeared in: Max Dauchet, ed., Proc. of CAAP'97 as part of TAPSOFT'97, Theory and Practice of Software Development. April 1997, Lille, France.

in some model of these axioms. These axioms define a fixpoint algorithm that closes a given input constraint under its consequences with respect to the axioms. We prove that a constraint φ is satisfiable if and only if the algorithm with input φ does not derive $\bot$ as a consequence of φ. All axioms (for infinite trees) will be discussed later in this introduction.

Sets versus Trees. The satisfiability problems of several classes of first-order formulae interpreted over trees and over non-empty sets of trees are closely related. The following two instances of this observation have inspired our choice of axioms or underly our proofs.

Equality constraints are conjunctions of equations $t_1{=}t_2$ between first-order terms. Over sets, they can be expressed by inclusion constraints due to anti-symmetry of set inclusion ($t_1{=}t_2 \leftrightarrow t_1{\subseteq}t_2 \wedge t_2{\subseteq}t_1$). Actually, even the first-order theories of equality constraints over trees and of equality constraints over non-empty sets of trees coincide. This follows from the complete axiomatization of the first-order theory of equality constraints over trees [18, 19, 12] since its axioms also hold over non-empty sets of trees (but don't over possibly empty sets).

There exists a natural interpretation of Ines-constraint over tree like structures that we call tree prefixes. In a different context [6] tree prefixes are called Böhm trees (without λ-binders). Tree prefixes come with a natural ordering relation where the empty tree prefix is the greatest element. We prove that an Ines constraint is satisfiable over non-empty sets of trees if and only if it is satisfiable over tree prefixes (where the inclusion symbol is interpreted as the inverse of the prefix ordering on tree prefixes).

Axioms. The first two axioms we need postulate the reflexivity and transitivity of the inclusion relation. We also assume the following decomposition axiom (here formulated for a binary function symbol f).

$$f(x,y){\subseteq}f(x',y') \;\rightarrow\; x{\subseteq}x' \wedge y{\subseteq}y'$$

This axiom holds over non-empty sets of trees but not over possibly empty sets, since every variable assignment α with $\alpha(x) = \emptyset$ or $\alpha(y) = \emptyset$ is a solution of $f(x,y){\subseteq}f(x',y')$ but not necessarily of $x{\subseteq}x' \wedge y{\subseteq}y'$. An analogous statement holds for the following clash axiom.

$$f(x,y){\subseteq}g(x',y') \;\rightarrow\; \bot \qquad \text{for } f \neq g$$

These axioms do not suffice to characterize the satisfiability of Ines constraints. For instance, the unsatisfiability of the constraint φ given by $x{\subseteq}g(x) \wedge x{\subseteq}g(y) \wedge y{\subseteq}z \wedge z{\subseteq}a$ is not derivable with these axioms alone. We need further axioms that use non-disjointness constraints $t_1 \between t_2$ defined as $t_1{\cap}t_2{\not\subseteq}\emptyset$. For the nondisjointness relation we require reflexivity and symmetry and a decomposition axiom as for the inclusion relation.

$$f(y,z) \between f(y',z') \rightarrow y \between y' \wedge z \between z'$$

Finally, we assume a clash axiom similar to the one for inclusion and require nondisjointness to be compatible with inclusion in the following sense.

$$x \not| z \wedge x \subseteq y \rightarrow y \not| z$$

Now reconsider the constraint φ given above and observe that we can derive $x \not| x$ by reflexivity, then $x \not| y$ by decomposition, and $x \not| z$ by compatibility. This yields a clash with $x \subseteq g(x) \wedge z \subseteq a$.

Algorithm and Complexity. The above axioms yield an algorithm that adds constraints of the form $x \subseteq y$, $x \not| y$ to a given input constraint φ until φ is closed under all axioms or implies $\bot$. The INES constraint $x \subseteq t_1 \wedge \ldots \wedge x \subseteq t_n$ expresses the n sets denoted by the terms $t_1, \ldots, t_n$ have a non-empty intersection. Fortunately, it is not necessary to add k-ary non-disjointness constraints of the form $x_1 \cap \ldots \cap x_k \not\subseteq \emptyset$ (which can be expressed by the formula $\exists y(y \subseteq x_1 \wedge \ldots \wedge y \subseteq x_k)$) of which there are exponentially many. Instead, our algorithm adds at most $O(n^2)$ constraints to the input constraint φ, where n is the number of variables in φ. The addition of a single constraint can be implemented such that it costs time $O(n)$. This yields an implementation of our algorithm with time complexity $O(n^3)$. This implementation can be organized incrementally.

Type Analysis. One application for INES constraints which we are investigating in [23] is type analysis for concurrent constraint programming [17, 27], in particular OZ [28]. As formal foundations we intend to use the calculi in [24, 25]. There, INES constraints are used to approximate the set of run-time values for program variables. Since values in OZ include infinite trees, it is important that INES allows an interpretation over sets of possibly infinite trees. It is considered an error if the set of possible run-time values is empty for some variable. This fact was our initial motivation for the choice of non-empty sets of trees as the interpretation domain for INES constraints.

Plan of the Paper. In Section 2, we discuss relate work. In Section 3, we define the syntax and semantics of INES constraints and in Section 4, we present the axioms and the algorithm. In Section 5, we prove the completeness of our algorithm. In Section 6, we compare the interpretations of INES constraints over tree prefixes and over non-empty sets of trees. Due to space limitations, we omit the details of the proofs in the conference version of the paper.

Appendix A gives an example illustrating program analysis for OZ with INES constraints. Appendix B contains the omitted proofs. Appendix C details how to implement the algorithm with incremental $O(n^3)$ complexity. In Appendix D, we adapt the algorithm to the finite-tree case, and in Appendix E to a subclass of standard set constraints (interpreted over possibly empty sets of finite trees) with explicit non-emptiness constraints $x \not\subseteq \emptyset$. We also prove that satisfiability of atomic set constraints (standard set constraints without set operators and negation) is invariant with respect to the choice of finite or infinite trees.

2 Related Work

Standard Set Constraints. Set constraints as in [2, 5, 10, 15] are inclusions between first-order terms with set operators interpreted over sets of finite trees. Our algorithm can be adapted such that it solves a subclass of set constraints without set operators in cubic time (see Appendix E). The general case is non-deterministically exponential time complete as proved in [1, 13]. The subclass that we can solve in cubic time syntactically extends the INES constraints with explicit non-emptiness constraint $x \not\subseteq \emptyset$ (see Appendix E). Note that the satisfiability of these set constraints depends on the choice of finite or infinite trees (consider $x \subseteq f(x) \wedge x \not\subseteq \emptyset$), which is in contrast to standard set constraints without negation. Our algorithm accounts for finiteness through the occur check.

Atomic Set Constraints. Heintze and Jaffar consider so-called *atomic set constraints* [15] which syntactically coincide with INES constraints but are interpreted over possibly empty sets of finite trees. The satisfiability problem for atomic set constraints is also $O(n^3)$. This result is implicit in the combined results of [14] and [15]. An explicit proof is given in Appendix E of this paper.

Set Constraints for Type Analysis. Aiken *et al.* [3, 4] use constraints over specific sets of trees called "types" for the type analysis of FL. There is a minimal type $\mathsf{0}$ which – in terms of constraint solving – behaves just like the empty set in standard set constraints (although it is not an empty set from the types point of view but contains a value denoting non-termination). In contrast to the constraints of this paper, their set constraints provide for union and intersection. One of the optimizations used by Aiken *et al.* is to strengthen the following constraint simplification rule by dropping the disjuncts in brackets [4].

$$f(x,y) \subseteq f(x',y') \rightarrow x \subseteq x' \wedge y \subseteq y' \; [\, \vee x \subseteq \mathsf{0} \vee y \subseteq \mathsf{0} \,]$$

As stated in [4], this optimization does not preserve soundness ($f(a,\mathsf{0}) \subseteq f(b,\mathsf{0})$ holds but $a \subseteq b \wedge \mathsf{0} \subseteq \mathsf{0}$ does not). It might be possible to justify it by using non-empty sets as interpretation domain. This is left to further research.

Entailment and Independence for Ines Constraints. Charatonik and Podelski [11] give an algorithm which decides the entailment problem between INES constraints when interpreted over sets of finite trees. They also decide the satisfiability of INES constraints with negation in the finite tree case. The results in [11] do not include any of the results presented here since they use as an explicit prerequisite the fact that satisfiability of INES constraints is decidable.

Tarskian Set Constraints. MacAllester and Givan [21] give a cubic algorithm which decides satisfiability for a class of Tarskian set constraints [22], and which also contains a non-disjointness constraint. Apart from this syntactic similarity, the two satisfiability problems are rather different problems since Tarskian set constraints are not interpreted over the domain of trees (this is also observed in [22]). A related open question is whether our axioms define a local theory [20, 8], which would also proof the cubic complexity bound of our algorithm.

3 Syntax and Semantics of Ines Constraints

We assume a set of *variables* ranged over by x, y, z and a signature Σ that defines a set of *function symbols* f, g and their respective arity $n \geq 0$. *Constants* (i.e. function symbols of arity 0) are denoted with a and b.

Trees. We base the definition of trees on the notion of paths since we wish to include infinite trees. Paths will turn out central for our proofs in Section 5. A *path* p is a sequence of positive integers ranged over by i, j, n, m. The *empty path* is denoted by ε. We write the free-monoid concatenation of paths p and q as pq; we have $\varepsilon p = p\varepsilon = p$. Given paths p and q, q is called a *prefix of* p if $p = qp'$ for some path p'.

Let τ be a set of pairs $(p,\ f)$ of paths p and function symbols f. We say that τ is *prefix closed*, if $(p,\ f) \in \tau$ and q is a prefix of p implies that there is a g such that $(q,\ g) \in \tau$. It is *path consistent*, if $(p,\ f) \in \tau$ and $(p,\ g) \in \tau$ implies $f{=}g$. We call τ *arity consistent*, if $(p,\ f) \in \tau$, $(pi,\ g) \in \tau$ implies that $i \in \{1, \ldots, n\}$ provided the arity of f is n. Finally, τ is called *arity complete*, if $(p,\ f) \in \tau$, where the arity of f is n, implies for all $i \in \{1, \ldots, n\}$ the existence of a g with $(pi,\ g) \in \tau$.

A (possibly infinite) *tree* τ is a set of pairs $(p,\ f)$ that is non-empty, prefix closed, arity complete, path consistent, and arity consistent. The set of all (possibly infinite) trees over Σ is denoted by Tree and the set of all non-empty sets of trees by $\mathsf{P}^+(\mathrm{Tree})$.

Ines Constraints. An Ines *constraint* $t_1{\subseteq}t'_1 \wedge \ldots \wedge t_n{\subseteq}t'_n$ is a conjunction of inclusions between first-order terms t defined by the following abstract syntax.

$$t \ ::= \quad x \ \mid \ f(\bar{t})$$

Here and throughout the paper, $\bar{t}$ stands for a sequence of terms and we assume implicitly that the length of $\bar{t}$ coincides with the arity of f. We interpret Ines constraints over the structure $\mathsf{P}^+(\mathrm{Tree})$ of non-empty sets of trees. In this structure, a function symbol f of Σ is interpreted as elementwise tree constructor and the relation symbol $\subseteq$ as subset relation. We call a first-order formula over Ines constraint *satisfiable* if it is satisfiable in the structure $\mathsf{P}^+(\mathrm{Tree})$. Two first-order formulae over Ines constraints are called *equivalent* if they are equivalently interpreted in $\mathsf{P}^+(\mathrm{Tree})$.

Flat Ines Constraints. For algorithmic reasons, we use an alternative constraint syntax in the sequel. First, we restrict ourselves to flat terms $f(\overline{x})$ and x instead of possibly deep terms t. Second, we use equalities $x{=}f(\overline{y})$ rather than inclusions $x{\subseteq}f(\overline{y})$ and $f(\overline{y}){\subseteq}x$ (this is a matter of taste). And third, we need binary non-disjointness constraints $x \not\parallel y$. Their semantics is given by the equivalence to the formula $x{\cap}y{\not\subseteq}\emptyset$ over sets of trees. Over *non-empty* sets of trees, $x \not\parallel y$

is equivalent to $\exists z(z{\subseteq}x \wedge z{\subseteq}y)$. Crucially, however, nondisjointness constraints $x \not\parallel y$ avoid explicit existential quantification in our algorithm.

These three steps lead us to *flat* INES *constraints* φ defined as follows.

$$\varphi \ ::= \quad \varphi_1 \wedge \varphi_2 \ \mid \ x{\subseteq}y \ \mid \ x{=}f(\overline{y}) \ \mid \ x \not\parallel y$$

We identify flat INES constraints φ up to associativity and commutativity of conjunction, *i.e.*, we consider φ as a multiset of inclusions $x{\subseteq}y$, equalities $x{=}f(\overline{y})$, and non-disjointness constraints $x \not\parallel y$.

From now on, we will consider only flat INES constraints and call them *constraints* for short. This is justified by the following Proposition. Let the *size of a constraint* φ be the number of function symbol occurrences plus variable occurrences in φ.

Proposition 1. *The satisfiability problems of* INES *constraints and of flat* INES *constraints have the same time complexity up to a linear transformation.*

4 Axioms and Algorithm

We present a set of axioms valid for INES-constraints interpreted over non-empty sets of trees. In a second step, we interpret these axioms as an algorithm that solves the satisfiability problem of INES constraints. The correctness and the complexity of this algorithm will be proved in Section 5.

A1. $x{\subseteq}x$ and $x{\subseteq}y \wedge y{\subseteq}z \rightarrow x{\subseteq}z$

A2. $x{=}f(\overline{y}) \wedge x{\subseteq}x' \wedge x'{=}f(\overline{z}) \rightarrow \overline{y}{\subseteq}\overline{z}$

A3. $x{\subseteq}y \rightarrow x \not\parallel y$ and $x{\subseteq}y \wedge x \not\parallel z \rightarrow y \not\parallel z$ and $x \not\parallel y \rightarrow y \not\parallel x$

A4. $x{=}f(\overline{y}) \wedge x \not\parallel x' \wedge x'{=}g(\overline{z}) \rightarrow \bot$ for $f \neq g$

A5. $x{=}f(\overline{y}) \wedge x \not\parallel x' \wedge x'{=}f(\overline{z}) \rightarrow \overline{y} \not\parallel \overline{z}$

Table 1. Axioms of INES constraints over non-empty sets of infinite trees

Table 1 contains five rules A1-A5 representing sets of axioms.[1] The union of these sets is denoted by A. For instance, a rule $x{\subseteq}x$ represents the infinite set

[1] Note that these axioms differ from the ones given in the introduction. The constraints used there are not flat and the variable-variable case $x{\subseteq}y$ and $x \not\parallel y$ are omitted. Indeed, the axioms in the introduction are semantically complete, although this is non-trivial to see and depends on the correctness of the algorithm presented here.

of axioms that is obtained by instantiation of the meta variable x with concrete variables. Note that an axiom is either a constraint φ, an implication between constraints $\varphi \to \psi$, or an implication $\varphi \to \bot$.

Proposition 2. *The structure* $\mathsf{P}^+(\mathrm{Tree})$ *is a model of the axioms in* A.

Proof. By a routine check. We note that the non-emptiness assumption of $\mathsf{P}^+(\mathrm{Tree})$ is essential for axioms A2 and A3.1. □

The Algorithm. The set of axioms A can be considered as a (naïve) fixed point algorithm A that, given an input constraint φ, iteratively adds logical consequences of $\mathsf{A} \cup \{\varphi\}$ to φ. More precisely, in every step A inputs a constraint φ and either terminates with $\bot$ or outputs a constraint $\varphi \wedge \psi$. Termination with $\bot$ takes place if there exists $\psi' \in \varphi$ such that $\psi' \to \bot \in \mathsf{A}$. Output of $\varphi \wedge \psi$ is possible if $\psi \in \mathsf{A}$ or there exists ψ' *in* φ with $\psi' \to \psi \in \mathsf{A}$.

Example 1. A first type of inconsistency depends on the transitivity of set inclusion. Here is a typical example:

$$x{=}a \wedge x{\subseteq}y \wedge y{\subseteq}z \wedge z{=}b \to \bot \qquad \text{for } a \neq b$$

Algorithm A may add $x{\subseteq}z$ by A1.2, then $x \not\parallel z$ with A3.1, and then terminate with $\bot$ by A4.

Example 2. A second type of inconsistency comes with implicit or explicit non-disjointness requirements. For illustration, we consider:

$$x{=}a \wedge z{\subseteq}x \wedge z{\subseteq}y \wedge y{=}b \to \bot \text{ for} \qquad \text{for } a \neq b$$

Algorithm A may add $z \not\parallel x$ by A3.1, then $x \not\parallel z$ via A3.3, then $x \not\parallel y$ with A3.2, and finally terminate with $\bot$ via A4.

Example 3. Inconsistencies of the above two types may be detected by structural reasoning with A2. Consider:

$$x{=}f(x) \wedge x{=}f(z) \wedge z{=}a \to \bot$$

Algorithm A may add $x{\subseteq}x$ by A1.1, then $x{\subseteq}z$ with A2, then $x \not\parallel z$ by A3.1, and finally terminate with $\bot$ with A4.

Example 4. We need another structural argument based on A5 for deriving the unsatisfiability of the following constraint.

$$x{=}f(x) \wedge z{\subseteq}x \wedge z{\subseteq}y \wedge y{=}f(x') \wedge x'{=}a \to \bot$$

Algorithm A may add $x \not\parallel y$ after several steps as shown in Example 2. Then it may proceed with $x \not\parallel x'$ via A5 and terminate with $\bot$ via A4.

Termination. Algorithm A can be organized in a terminating manner by adding a simple *control*. Given an input constraint φ, we add only such constraints $x \not| y$ and $x \subseteq y$ to φ which are not contained in φ. We also restrict reflexivity of inclusion $x \subseteq x$ to such variables x occurring in φ. Given a subset S of A, a constraint φ is called A'*-closed*, if algorithm A under the given control and restricted to the axioms in A' cannot proceed. (Note that constraints do not contain $\bot$ by definition.) This defines the notion of A-closedness but also of $\mathsf{A1}$-closedness, $\mathsf{A2}$-closedness, *etc.*, which will be needed later on.

Example 5. Our control takes care of termination in presence of cycles like $x{=}f(x)$. For instance, the following constraint is A-closed.

$$x{=}f(x) \wedge x \subseteq y \wedge y{=}f(x) \wedge x \subseteq x \wedge y \subseteq y \wedge x \not| x \wedge y \not| y \wedge x \not| y \wedge y \not| x$$

In particular, $\mathsf{A2}$ and $\mathsf{A5}$ do not loop through the cycle $x{=}f(x)$ infinitely often.

Proposition 3. *If φ is a constraint with m variables then algorithm A with input φ terminates under the above control in at most $2 \cdot m^2$ steps.* □

Proof. Since A does not introduce new variables, it may add at most m^2 non-disjointness constraints $x \not| y$ and m^2 inclusions $x \subseteq y$. □

Proposition 4. *Every A-closed constraint φ is satisfiable over $\mathsf{P}^+(\mathrm{Tree})$.*

The proof of this statement is the subject of Section 5. There, we construct the greatest solution for a satisfiable constraint (Lemma 9). Note that constraints in general do not have a smallest solution (consider $x \subseteq f(x\ y)$).

Theorem 5. *The satisfiability of* Ines *constraints can be decided in time $O(n^3)$ (offline and online) where n is the constraint size.*

Proof. Proposition 2 shows that φ is unsatisfiable if A started with φ terminates with $\bot$. Proposition 4 proves that φ is satisfiable if A started with φ terminates with a constraint. Since A terminates for all input constraints under the above control (Proposition 3), this yields a effective decision procedure. The complexity statement is proved in Proposition 14. The main idea is that every step of algorithm A can be implemented in time $O(n)$ and that there are $O(n^2)$ steps (Proposition 3). [2] In the proof of Proposition 14, we present an incremental implementation of algorithm A. It exploits that algorithm A leaves the order unspecified in which axioms in A are applied. □

There is a class of constraints on which algorithm A indeed takes cubic time, namely the inclusions cycles $x_1 \subseteq x_2 \wedge \ldots \wedge x_{n-1} \subseteq x_n \wedge x_n \subseteq x_1$ where $n \geq 1$. The closure under A is the full transitive closure $\bigwedge\{x_i \subseteq x_j \mid i, j \in \{1 \ldots n\}\}$ plus the corresponding non-disjointness constraints.

[2] Every step of algorithm A costs time $O(n)$ only with respect to an amortized time analysis, which we do not make explicit in our complexity proof.

5 Completeness

The goal of this Section is to prove the completeness of our algorithm as stated in Proposition 4. We have to construct a solution for every A-closed constraint. The idea is to construct solution in a substructure of $\mathsf{P}^+(\text{Tree})$ the structure of tree prefixes.

Tree Prefixes. A *tree prefix* τ is a set of pairs (p, f) that is prefix closed, path consistent, and arity consistent. Note that every tree is a tree prefix. The set of all tree prefixes is denoted by Prefix. We can naturally interpret INES constraints over tree prefixes such that Prefix becomes a structure. Function symbols $f \in \Sigma$ are interpreted as tree prefix constructors (generalizing tree constructors). The inclusion symbol $\subseteq$ is interpreted as the *inverted* subset relation on tree prefixes that we denote with $\leq$ (*i.e.*, $\tau_1 \leq \tau_2$ iff $\tau_1 \supseteq \tau_2$). The relation $\tau_1 \not\parallel \tau_2$ holds over Prefix iff $\tau_1 \cup \tau_2$ is path consistent (and hence a tree prefix).

Proposition 6. Prefix *is a substructure of* $\mathsf{P}^+(\text{Tree})$ *with respect to the embedding Trees* : Prefix $\to \mathsf{P}^+(\text{Tree})$ *given by:*

$$\mathit{Trees}(\tau) = \{\tau' \mid \tau' \text{ is a tree such that } \tau' \leq \tau\}$$

Proof. The mapping *Trees* is a homomorphism with respect to function symbols $f \in \Sigma$ and the relation symbols $\subseteq$ and $\not\parallel$. □

Corollary 7. *If a constraint is satisfiable over* Prefix *then it is satisfiable over* $\mathsf{P}^+(\text{Tree})$.

Proof. For constraints $x{\subseteq}y$, $x{=}f(\overline{y})$, and $x \not\parallel y$, this follows from Proposition 6. A conjunction of such constraints is satisfiable if all conjuncts are satisfiable. □

Path Reachability. We introduce the path reachability relations $\overset{\varphi}{\rightsquigarrow}_p$ and the notion of path consistency with respect to constraints. For all paths p and constraint φ, we define a binary relation $\overset{\varphi}{\rightsquigarrow}_p$, where $x \overset{\varphi}{\rightsquigarrow}_p y$ reads as "y is reachable from x over path p in φ":

$$\begin{array}{ll} x \overset{\varphi}{\rightsquigarrow}_\varepsilon y & \text{if } x{\subseteq}y \text{ in } \varphi \\ x \overset{\varphi}{\rightsquigarrow}_i y_i & \text{if } x{=}f(y_1 \ldots y_i \ldots y_n) \text{ in } \varphi, \\ x \overset{\varphi}{\rightsquigarrow}_{pq} y & \text{if } x \overset{\varphi}{\rightsquigarrow}_p u \text{ and } u \overset{\varphi}{\rightsquigarrow}_q y. \end{array}$$

We define relations $x \overset{\varphi}{\rightsquigarrow}_p f$ meaning "f can be reached from x via path p in φ":

$$x \overset{\varphi}{\rightsquigarrow}_p f \text{ if } x \overset{\varphi}{\rightsquigarrow}_p y \text{ and } y{=}f(\bar{u}) \text{ in } \varphi,$$

For example, if φ is the constraint $x{\subseteq}y \wedge y{=}f(u,z) \wedge z{=}g(x)$ then the following reachability from x relationships hold: $x \overset{\varphi}{\rightsquigarrow}_\varepsilon y$, $x \overset{\varphi}{\rightsquigarrow}_2 z$, $x \overset{\varphi}{\rightsquigarrow}_{21} x$, $x \overset{\varphi}{\rightsquigarrow}_{21} y$, etc., as well as $x \overset{\varphi}{\rightsquigarrow}_\varepsilon f$, $x \overset{\varphi}{\rightsquigarrow}_2 g$, $x \overset{\varphi}{\rightsquigarrow}_{21} f$, etc.

Definition 8 Path Consistency. We call a constraint φ *path consistent* if the following two conditions hold for all x, y, p, f, and g.

1. If $x \overset{\varphi}{\rightsquigarrow}_p g$, $x{\subseteq}x$, and $x \overset{\varphi}{\rightsquigarrow}_p f$ then $f = g$.
2. If $x \overset{\varphi}{\rightsquigarrow}_p g$, $x \not\parallel y$, and $y \overset{\varphi}{\rightsquigarrow}_p f$ then $f = g$.

Lemma 9. *Every* A1-A2*-closed and path consistent constraint is satisfiable over* Prefix.

Lemma 10. *Every* A3-A5*-closed constraint is path consistent.*

Proof of Proposition 4. We have to show that every A-closed constraint φ is satisfiable. φ is path consistent by Lemma 10, satisfiable in Prefix by Lemma 9, and hence satisfiable in P^+(Tree) by Corollary 7. □

6 Non-Empty Sets versus Trees

We discuss interpretations of INES constraints over tree prefixes and over non-empty sets of trees. For the fragment of equality constraints we also consider an interpretation over trees.

Theorem 11. *Given an* INES *constraints φ, the following three statements are equivalent:*

1. *φ is satisfiable (over P^+(Tree)).*
2. *φ is satisfiable over Prefix.*
3. *φ is satisfiable in some model of the axioms in A.*

Proof. 1) to 3). If φ is satisfiable over P^+(Tree), then it is satisfiable in some model of A, since P^+(Tree) is a model of A by Proposition 2.

3) to 2). Let φ be satisfiable in some model of A. Algorithm A terminates when started with φ by Proposition 3. It outputs a constraint ψ (and not $\bot$) that is equivalent to φ in all models of A. ψ is A-closed and hence satisfiable over Prefix by Lemmata 9 and 10.

2) to 1). If φ is satisfiable over Prefix then it is satisfiable by Corollary 7. □

An *equality constraint* is a conjunction of equalities $x{=}y$ and $x{=}f(\overline{y})$. Over P^+(Tree), equalities can be expressed by inclusions since the inclusion ordering is antisymmetric ($x{=}y \;\leftrightarrow\; x{\subseteq}y \wedge y{\subseteq}x$).

Theorem 12. *The three first-order theories of equality constraints over non-empty sets of trees, over tree prefixes, and over trees coincide (i.e., of the structures* P^+(Tree), Prefix *and* Tree*).*[3]

[3] Independently, A. Colmerauer observed this for P^+(Tree) and Tree (pers. comm.).

Proof. This follows from the fact that all axioms of the complete axiomatization of trees [18, 19, 12] are valid for non-empty sets of trees. This holds for the axioms of the form $\forall \overline{y} \exists! \overline{x}(x_1{=}f_1(\overline{x}\ \overline{y}) \wedge \ldots \wedge x_n{=}f_n(\overline{x}\ \overline{y}))$. Validity of the other axioms is immediate since they are already contained in A with inclusion replaced for equality. □

In contrast, first-order formulae over inclusion constraints can distinguish the structures P^+(Tree) and Prefix. A formula that holds over Prefix but not over P^+(Tree) is given by

$$\forall x(a{\subseteq}x \wedge b{\subseteq}x \rightarrow \forall y(y{\subseteq}x))$$

where $a \neq b$. Another formula distinguishing both structures comes with a constraint-based reformulation of the coherence property (defined for complete partial orders in [6]).

We say that an ordering relation satisfies the *coherence property* if it satisfies the following formulae for all finite sets I (where inclusion symbol is interpreted as the given ordering).

$$\textstyle\bigwedge_{i,j \in I} \exists z\,(z{\subseteq}x_i \wedge z{\subseteq}x_j) \rightarrow \exists z(\bigwedge_{i \in I} z{\subseteq}x_i)$$

This formula states that for all variable assignment α the elements from the finite set $\{\alpha(x_i) \mid i \in I\}$ have a common lower bound if every two of its elements $\alpha(x_i), \alpha(x_j)$ have $(i, j, \in \{1, \ldots, n\})$. For inclusion over non-empty sets this property does not hold. There it states the non-emptiness of an n-intersection $t_1 \cap \ldots \cap t_n$ if all pairwise intersections $t_i \cap t_j$ are non-empty $(i, j \in \{1 \ldots n\})$, which is refuted by the example $I = \{1, 2, 3\}$ and $\alpha(x_1) = \{a, b\}$, $\alpha(x_2) = \{a, c\}$, $\alpha(x_3) = \{b, c\}$ for distinct constants a, b, c.

Proposition 13. *The tree prefix ordering $\leq$ satisfies the coherence property.*

Proof. For some finite index set $J \subseteq I$ and variable assignment α into Prefix, note that α is a solution of $\exists z(\bigwedge_{i \in J} z{\subseteq}x_i)$ iff $\bigcup_{i \in J} \alpha(x_i)$ is path consistent. If α is a solution of all $\exists z\,(z{\subseteq}x_i \wedge z{\subseteq}x_j)$ then all pairwise unions $\alpha(x_i) \cup \alpha(x_j)$ are path consistent such that the union $\bigcup_{i \in I} \alpha(x_i)$ is path consistent. Hence α is a solution of $\exists z(\bigwedge_{i \in I} z{\subseteq}x_i)$. □

Acknowledgements. We would like to thank David Basin, Denys Duchier, Witold Charatonik, Harald Ganzinger, Gert Smolka, Ralf Treinen and Uwe Waldmann, as well as the anonymous referees for valuable comments on drafts of this paper. The research reported in this paper has been supported by the the Esprit Working Group CCL II (EP 22457) and the Deutsche Forschungsgemeinschaft through the Graduiertenkolleg Kognitionswissenschaft and the SFB 378 at the Universität des Saarlandes.

References

1. A. Aiken, D. Kozen, and E. Wimmers. Decidability of Systems of Set Constraints with Negative Constraints. *Information and Computation*, 1995.

2. A. Aiken and E. Wimmers. Solving Systems of Set Constraints. In *Proc. 7^{th} LICS*, pp. 329–340. IEEE, 1992.
3. A. Aiken and E. Wimmers. Type Inclusion Constraints and Type Inference. In *Proc. 6^{th} FPCA*, pp. 31–41. 1993.
4. A. Aiken, E. Wimmers, and T. Lakshman. Soft Typing with Conditional Types. In *Proc. 21^{st} POPL*. ACM, 1994.
5. L. Bachmair, H. Ganzinger, and U. Waldmann. Set Constraints are the Monadic Class. In *Proc. 8^{th} LICS*, pp. 75–83. IEEE, 1993.
6. H. P. Barendregt. *The Type Free Lambda Calculus*. In Barwise [7], 1977.
7. J. Barwise, ed. *Handbook of Mathematical Logic*. Number 90 in Studies in Logic. North–Holland, 1977.
8. D. Basin and H. Ganzinger. Automated Complexity Analysis Based on Ordered Resolution. In *11^{th} LICS*. IEEE, 1996.
9. W. Charatonik and L. Pacholski. Negative set constraints with equality. In *Proc. 9^{th} LICS*, pp. 128–136. 1994.
10. W. Charatonik and L. Pacholski. Set constraints with projections are in NEXPTIME. In *Proc. 35^{th} FOCS*, pp. 642–653. 1994.
11. W. Charatonik and A. Podelski. The Independence Property of a Class of Set Constraints. In *Proc. 2nd CP*. LNCS 1118, Springer, 1996.
12. H. Comon and P. Lescanne. Equational problems and disunification. *Journal of Symbolic Computation*, 7:371–425. 1989.
13. R. Gilleron, S. Tison, and M. Tommasi. Solving Systems of Set Constraints with Negated Subset Relationships. In *Proc. 34^{nd} FOCS*, pp. 372–380. 1993.
14. N. Heintze. Set Based Analysis of ML Programs. Technical Report CMU–CS–93–193, School of Computer Science, Carnegie Mellon University. July 1993.
15. N. Heintze and J. Jaffar. A Decision Procedure for a Class of Set Constraints (Extended Abstract). In *Proc. 5^{th} LICS*, pp. 42–51. IEEE, 1990.
16. D. Kozen. Logical aspects of set constraints. In *Proc. CSL*, pp. 175–188. 1993.
17. M. J. Maher. Logic semantics for a class of committed-choice programs. In J.-L. Lassez, ed., *Proc. 4^{th} ICLP*, pp. 858–876. The MIT Press, 1987.
18. M. J. Maher. Complete Axiomatizations of the Algebras of Finite, Rational and Infinite Trees. In *Proc. 3^{rd} LICS*, pp. 348–457. IEEE, 1988.
19. A. I. Malc'ev. Axiomatizable Classes of Locally Free Algebras of Various Type. In *The Metamathematics of Algebraic Systens: Collected Papers 1936-1967*, ch. 23, pp. 262–281. North–Holland, 1971.
20. D. McAllester. Automatic Recognition of Tractability in Inference Relations. *Journal of the ACM*, 40(2), Apr. 1993.
21. D. McAllester and R. Givan. Taxonomic Syntax for First-Order Inference. *Journal of the ACM*, 40(2), Apr. 1993.
22. D. McAllester, R. Givan, D. Kozen, and C. Witty. Tarskian Set Constraints. In *Proc. 11^{th} LICS*. IEEE, 1996.
23. M. Müller. *Type Analysis for a Higher-Order Concurrent Constraint Language*. Doctoral Dissertation. Universität des Saarlandes, Technische Fakultät, 66041 Saarbrücken, Germany. In preparation.
24. J. Niehren. Functional Computation as Concurrent Computation. In *23^{rd} POPL*, pp. 333–343. ACM, 1996.
25. J. Niehren and M. Müller. Constraints for Free in Concurrent Computation. In *Proc. 1^{st} ASIAN*, LNCS 1023, pp. 171–186. Springer, 1995.
26. *The Oz Programming System*. Programming Systems Lab, Universität des Saarlandes. Available at `http://www.ps.uni-sb.de/www/oz/`.

27. V. A. Saraswat. *Concurrent Constraint Programming.* The MIT Press, 1993.
28. G. Smolka. The Oz Programming Model. In J. van Leeuwen, ed., *Computer Science Today*, LNCS 1000, pp. 324–343. Springer, 1995.

A Ines-Constraints for Program Analysis

We are investigating the application of Ines constraints for program analysis. More specifically, we intend to construct a type analysis system for concurrent constraint programming languages [17, 27] such as Oz [28] (see [24, 25] for formal foundations of Oz). During the execution of programs in these languages, the possible values of program variables are approximated by constraints. For programs without search (backtracking), it is considered a programming error if the set of possible values is empty for some program variable.

For illustration, consider the following Oz program with its constraint-based analysis added in comments (using the special function symbol proc).[4]

```
proc {P X} X=a end          % ∃x(p=proc(x) ∧ x=a) ∧
proc {Q Y} Y=b end          % ∃y(q=proc(y) ∧ y=b) ∧
{P Z} {Q Z}                 % proc(z)⊆p ∧ proc(z)⊆q
```

The program contains the definition of two procedures `P` and `Q` with formal arguments `X` and `Y`, respectively, as well as two procedure applications with the same actual argument `Z`. On execution of these applications, the constraints `Z=a` and `Z=b` will be emitted which are inconsistent with each other.

A program analysis in terms of Ines-constraints can detect this error as follows. The program variables `P`, `Q`, `X`, `Y`, and `Z` are mapped to constraint variables p, q, x, y, and z, and the program subexpressions are mapped to constraints as indicated in the comments. The conjunction of these constraints is checked for satisfiability and the program is rejected if this test fails. The above program is rejected since its analysis implies $z{\subseteq}a \wedge z{\subseteq}b$ which is unsatisfiable.

We have implemented a type analysis system based on Ines-constraints and use it experimentally for Oz programs. The full description of the type analysis system is out of the scope of this paper and will be reported in [23].

B Omitted Proofs

Proposition 1. The satisfiability problems of Ines constraints and of flat Ines constraints have the the same time complexity up to a linear transformation.

[4] This example also appeared in the follow-up paper [11] with the explicit statement that it is borrowed from here.

Proof. With respect to the structure $\mathsf{P}^+(\mathrm{Tree})$, every flat Ines constraint is equivalent to a first-order formula over Ines constraints.

$$x{=}f(\overline{y}) \leftrightarrow x{\subseteq}f(\overline{y}) \wedge f(\overline{y}){\subseteq}x \qquad x \not\parallel y \leftrightarrow \exists z(z{\subseteq}x \wedge z{\subseteq}y).$$

Conversely, every Ines constraint is equivalent to a first-order formula over flat Ines constraints.

$$\begin{array}{ll} x{\subseteq}f(\bar{t}) \leftrightarrow \exists y \exists \overline{z}(x{\subseteq}y \wedge y{=}f(\overline{z}) \wedge \overline{z}{\subseteq}\bar{t}) & \qquad t{\subseteq}t' \leftrightarrow \exists x(t{\subseteq}x \wedge x{\subseteq}t') \\ f(\bar{t}){\subseteq}x \leftrightarrow \exists y \exists \overline{z}(\bar{t}{\subseteq}\overline{z} \wedge f(\overline{z}){=}y \wedge y{\subseteq}x) & \end{array}$$

These equivalences can be interpreted as constraint transformers from Ines constraints into flat Ines constraints and vice versa. Hence, for every Ines constraint there exists a satisfaction equivalent constraint and vice versa. It is easy to organize the transformations such that they preserve the size of constraints up to a factor of 2. Hence, the complexity of the satisfiability problems is preserved. □

Lemma 9. Every A1-A2-closed and path consistent constraint is satisfiable over Prefix.

Proof. Let φ be A1-A2-closed and path consistent. We define a variable assignment prefix_φ into Prefix as follows:

$$\mathsf{prefix}_\varphi(x) = \{(p, f) \mid x \overset{\varphi}{\rightsquigarrow}_p f\}$$

The path consistency of φ (condition 1) implies the path consistency of $\mathsf{prefix}_\varphi(x)$. Thus $\mathsf{prefix}_\varphi(x)$ is a tree prefix (one can show this by induction over p). We now verify that prefix_φ is a solution of φ.

- Let $x{\subseteq}y$ *in* φ. If $y \overset{\varphi}{\rightsquigarrow}_p g$ then $x \overset{\varphi}{\rightsquigarrow}_p g$ by the definition of path reachability. Thus, $\mathsf{prefix}_\varphi(y) \subseteq \mathsf{prefix}_\varphi(x)$.
- Consider $x{=}f(y_1 \ldots y_n)$ *in* φ. If $i \in \{1 \ldots n\}$ and $y_i \overset{\varphi}{\rightsquigarrow}_p g$ then $x \overset{\varphi}{\rightsquigarrow}_{ip} g$. Thus, $f(\mathsf{prefix}_\varphi(y_1) \ldots \mathsf{prefix}_\varphi(y_n)) \subseteq \mathsf{prefix}_\varphi(x)$. For the converse inclusion, we first show that φ satisfies the following two properties for all g and i:

P1	if $x \overset{\varphi}{\rightsquigarrow}_\varepsilon g$ then $f = g$.
P2	if $i \in \{1 \ldots n\}$ and $x \overset{\varphi}{\rightsquigarrow}_{ip} g$ then $y_i \overset{\varphi}{\rightsquigarrow}_p g$.

  For proving P1 we assume $x \overset{\varphi}{\rightsquigarrow}_\varepsilon g$. Since $x{=}f(\bar{u})$ *in* φ we have $x{\subseteq}x$ *in* φ by A1.1-closedness. Thus $x \overset{\varphi}{\rightsquigarrow}_\varepsilon f$ which implies $f = g$ since φ is path consistent (condition 1) and A1.1-closed, i.e. P1 holds.

  For proving P2, we assume $i \in \{1 \ldots n\}$ and $x \overset{\varphi}{\rightsquigarrow}_{ip} g$. By definition of path reachability there exists x', f', and $\bar{v}$ such that

  $$x \overset{\varphi}{\rightsquigarrow}_\varepsilon x', \qquad x'{=}f'(y'_1 \ldots y'_i \ldots y'_n), \qquad y'_i \overset{\varphi}{\rightsquigarrow}_p g.$$

The A1.2-closedness of φ and $x \overset{\varphi}{\rightsquigarrow}_\varepsilon x'$ imply $x \subseteq x'$ *in* φ. The path consistency of φ (condition 1) and the A1.1-closedness of φ implies $f = f'$. Hence, A2-closedness ensures $y_i \subseteq y_i'$ *in* φ such that $y_i \overset{\varphi}{\rightsquigarrow}_p g$ holds. This proves P2.

We finally show $\mathsf{prefix}_\varphi(x) \subseteq f(\mathsf{prefix}_\varphi(y_1) \ldots \mathsf{prefix}_\varphi(y_n))$. Given $(p, g) \in \mathsf{prefix}_\varphi(x)$, we distinguish two cases. If $p{=}\varepsilon$, then $x \overset{\varphi}{\rightsquigarrow}_\varepsilon g$ such that P1 implies $f = g$ and hence $(\varepsilon, g) \in f(\mathsf{prefix}_\varphi(y_1) \ldots \mathsf{prefix}_\varphi(y_n))$. If $p = iq$ then $x \overset{\varphi}{\rightsquigarrow}_{iq} g$ such that P2 yields $y_i \overset{\varphi}{\rightsquigarrow}_q g$ and hence $(p, g) \in f(\mathsf{prefix}_\varphi(y_1) \ldots \mathsf{prefix}_\varphi(y_n))$.

– Let $x \not\parallel y$ in φ. We have to show that the set $\mathsf{prefix}_\varphi(x) \cup \mathsf{prefix}_\varphi(y)$ is path consistent. If $(p, g) \in \mathsf{prefix}_\varphi(x)$ and $(p, f) \in \mathsf{prefix}_\varphi(y)$ then $x \overset{\varphi}{\rightsquigarrow}_p g$ and $y \overset{\varphi}{\rightsquigarrow}_p f$. The path consistency of φ (condition 2) implies $f = g$. □

Lemma 10. Every A3-A5-closed constraint is path consistent.

Proof. Let φ be A3−A5-closed. Condition 1 of Definition 8 follows from condition 2 of Definition 8 and A3.1-closedness. The proof of condition 2 in Definition 8 is by induction on paths p. We assume x, y, f, and g such that $x \overset{\varphi}{\rightsquigarrow}_p f$, $x \not\parallel y$ *in* φ, and $x \overset{\varphi}{\rightsquigarrow}_p g$.

If $p = \varepsilon$, then there exist $n, m \geq 0$, $x_1, \ldots, x_n$, $y_1, \ldots y_m$, $\bar{u}$, and $\bar{v}$ such that:

$$x \subseteq x_1 \wedge \ldots \wedge x_{n-1} \subseteq x_n \wedge x_n = f(\overline{x'}) \text{ in } \varphi\,,$$
$$y \subseteq y_1 \wedge \ldots \wedge y_{m-1} \subseteq y_m \wedge y_m = g(\overline{y'}) \text{ in } \varphi\,.$$

A3-closedness implies that $x_n \not\parallel y_m$ *in* φ (A3.2 yields $x \not\parallel y_1$ *in* φ, …, $x \not\parallel y_m$ *in* φ. Thus $y_m \not\parallel x$ *in* φ by A3.3-closedness such that A3.2-closedness yields $y_m \not\parallel x_1$ *in* φ, …, $y_m \not\parallel x_n$ *in* φ). Hence, A4-closedness implies $f = g$. If $p = iq$, then there exist f', g', x', y', $\bar{u}$, $\bar{v}$ with:

$$x \overset{\varphi}{\rightsquigarrow}_\varepsilon x', \quad x' = f'(x_1' \ldots x_i' \ldots x_n') \text{ in } \varphi\,, \quad x_i' \overset{\varphi}{\rightsquigarrow}_p f\,,$$
$$y \overset{\varphi}{\rightsquigarrow}_\varepsilon y', \quad y' = g'(y_1' \ldots y_i' \ldots y_n') \text{ in } \varphi\,, \quad y_i' \overset{\varphi}{\rightsquigarrow}_p g\,.$$

Since $x \not\parallel x'$ *in* φ, we have $x' \not\parallel y'$ *in* φ by A3-closedness (this has been proved for the case $p = \varepsilon$). Thus, A4-closedness yields $f' = g'$ such that A5-closedness implies $x_i' \not\parallel y_i'$ *in* φ, and hence $f = g$ holds by induction assumption. □

C Complexity

We elaborate the proof of the complexity and incrementality statement in Theorem 5 by presenting an implementation of algorithm A.

Proposition 14. *Algorithm* A *can be implemented (online and offline) such that it terminates in time* $O(n^3)$ *where* n *is the size of the input constraint.*

Proof. We organize algorithm A as a reduction relation on pairs (φ, ψ) or $\bot$, where φ is called *pool* and ψ *store*. The store and the pool are either constraints or empty multisets represented by $\top$. Initially, the pool φ is the input constraint called φ_0 (which may be inputed incrementally in the online case) and the store ψ is empty.

Reduction preserves the invariant that $\varphi \wedge \psi$ contains all one-step consequences of ψ with respect to algorithm A (and restricted to variables occuring in φ_0). If a pair (φ, ψ) reduces to $\bot$ then $\varphi \wedge \psi$ is equivalent to $\bot$. If (φ, ψ) reduces to (φ', ψ') then $\varphi \wedge \psi$ is equivalent to $\varphi' \wedge \psi'$. Reduction either terminates with $\bot$ or with an empty pool. In the latter case, the above invariants ensure that the final store is A-closed and equivalent to the initial constraint φ_0.

Let a *basic constraint* be of the form $x \subseteq y$, $x \not\downarrow y$, or $x = f(\overline{y})$. Reduction can be implemented by recursively executing the following sequence of instructions:

1. Select a basic constraint φ' from the pool. If φ' is contained in the store delete if from the pool and finish.
2. Else, for all axioms in A of the form $\varphi' \wedge \psi' \to \varphi''$ with ψ' in the pool add φ'' to the pool. If there exists an axiom of the form $\varphi' \wedge \psi' \to \bot$ in A with ψ' in the pool then reduce to $\bot$. If φ' contains a variable x such that $x \subseteq x$ is not contained in the store then add it to the store.
3. Add φ' to the store and delete it from the pool.

We first discuss the necessary data structures for implementing the reduction in a restricted case. In a second step we show that these restrictions can be omitted.

R1 The algorithm is offline, i.e. the input constraint φ_0 is statically known.

R2 The arity of constructors in φ_0 is bounded by a constant, say k.

R3 φ_0 contains at most one equality per variable.

Let m be the number of variables in φ_0. The pool can be implemented such that it provides for the following operations (for instance as a queue).

- select and delete a basic constraint from the pool in $O(1)$.
- add a basic constraint to the pool in $O(1)$.

The store can be implemented as an array of size m for the equalities $x = f(\overline{y})$ (at most one per variable) and a table of size $2 * m^2$ for the constraints $x \subseteq y$ and $x \not\downarrow y$ for all occuring variables. The store can support the the following operations:

- test the presence of an equality for x in $O(1)$.
- test the membership $x \subseteq y \in \psi$ and $x \not\downarrow y \in \psi$ in time $O(1)$.
- given a variable x with $x = f(\overline{y}) \in \varphi$, retrieve the function symbol f and the sequence $\overline{y}$ in time $O(1)$.

- add a basic constraint in time $O(1)$.
- given a variable x, retrieve the set of all y such that $x \subseteq y \in \varphi$ in time $O(m)$ (analogously for $x \not\parallel y$).

As shown in the next paragraph, the reduction relation can be implemented such that all operations on the store and the pool are invoked ad most $O(m^2)$ times. Since every operation costs at most $O(m)$ time and $m \leq n$, this yields an $O(n^3)$ implementation.

There are at most $O(m^2)$ distinct basic constraints that may be added to the store and every basic constraint may be added at most once. Hence there are at most $O(m^2)$ add operations on the store. Constraints are added to the pool only if some basic constraint is added to the store. In this case, at most $O(k)$ basic constraints are added to the pool by R2. Hence, there are at most $O(k * m^2)$ add operations on the pool.

We finally discuss how to get rid of the above restrictions.

R2 If the arity of constructors is unbounded then we still know that every operation cost at most $O(n)$ where n is the size of φ_0. The only problem is that the number of basic constraints that may be added to the pool is no more bounded by $O(n^2)$. This can be circumvented by adding constraints to the pool at most once, i.e. by remembering those constraints that have been added to pool (and possibly deleted) before. This can be done with a quadratic table as for the store.

R3 If we replace all equalities $x{=}f(\overline{y})$ in φ_0 by $x \subseteq x' \wedge x' \subseteq x \wedge x'{=}f(\overline{y})$ where x' is a fresh variable respectively then the resulting constraint does not contain two equations for the same variable.

R1 For an online algorithm, we can add the input constraint φ_0 incrementally to the pool. The problem is that the number of variables in φ_0 is not known statically. We have to replace our static tables and arrays by dynamic hash tables such that new variables can be inserted. □

D Finite Trees

The satisfiability of INES constraints depends on the interpretation over sets of finite or infinite trees.

Example 6. For instance, the constraint $x \subseteq f(x)$ is satisfiable over sets of infinite trees by $x \mapsto \{f(f(f(\ldots)))\}$, but non-satisfiable over sets of finite trees.

The results of Section C carry over to the finite tree case when we add the "occurs-check" axiom A6 from Table 2 to axiom set A. In particular, Lemma 9 and Theorem 5 can be adapted. Call $\mathrm{Tree}_{\mathrm{fin}}$ the set of finite trees.

A6.	$\varphi \to \bot$	if $x \overset{\varphi}{\rightsquigarrow}_p x$ for some path $p \neq \varepsilon$

Table 2. The occurs check axiom

Lemma 15. *A path consistent constraint φ closed under* A1-A3 *and* A6 *is satisfiable in* $\mathsf{P}^+(\mathrm{Tree}_{\mathrm{fin}})$.

Proof. To show a A1-A3 and A6-closed and path consistent constraint φ satisfiable in $\mathsf{P}^+(\mathrm{Tree})$ we have defined the prefix $\mathsf{prefix}_\varphi(x) = \{(p, f) \mid x \overset{\varphi}{\rightsquigarrow}_p f\}$. Since φ is finite and the assumption about axiom A6 excludes cyclic paths, $\mathsf{prefix}_\varphi(x)$ must be a finite prefix for all x. Hence, φ is satisfiable in $\mathsf{P}^+(\mathrm{Tree}_{\mathrm{fin}})$. □

Theorem 16. *The satisfiability of* Ines *constraints over non-empty sets of finite trees can be decided (offline or online) in time $O(n^3)$.*

Proof. The offline version of our algorithm may perform the occurs-check upon termination. This is linear in the size of the final constraint and cubic in the size of the start constraint. The online version must schedule the occurs-check after every step. This is constant if the closure of the reachability relation between variables is (just like $\subseteq$) implemented by a table of size quadratic in the number of variables. □

E Standard Set Constraints

In this section, we take a alternative approach to achieve the expressiveness of Ines constraints. We consider a class of standard set constraints by interpreting Ines constraints over possibly empty sets of trees and allowing for explicit non-emptiness constraints $x \not\subseteq \emptyset$ ("x denotes a non-empty set"). We show that the cubic algorithm for Ines constraints can be adapted to this fragment of standard set constraints at the cost of additional axioms.

We extend our constraint syntax with explicit non-emptiness constraints as follows.

$$\varphi \ ::= \quad \varphi_1 \wedge \varphi_2 \ \mid \ x{=}f(\overline{y}) \ \mid \ x{\subseteq}y \ \mid \ x \not| y \ \mid \ x \not\subseteq \emptyset \tag{1}$$

We interprete these constraint either in the structure of sets of trees $\mathsf{P}(\mathrm{Tree})$ or in the structure of sets of finite trees $\mathsf{P}(\mathrm{Tree}_{\mathrm{fin}})$. Due to the constraint $x \not\subseteq \emptyset$, satisfiability of set constraints differs depending on the choice of finite or infinite trees. This is not the case without $x \not\subseteq \emptyset$ as we will show below (Corollary 20). Example 6 adapts as follows.

Example 7. The constraint $x \not\subseteq \emptyset \wedge x \subseteq f(x)$ is satisfiable over sets of infinite trees by the variable assignment $x \mapsto \{f(f(f(\ldots)))\}$, but non-satisfiable over sets of finite trees.

A1.	$x \subseteq x$ and $x \subseteq y \wedge y \subseteq z \rightarrow x \subseteq z$
A2'.	$x \not\subseteq \emptyset \wedge x = f(\overline{y}) \wedge x \subseteq x' \wedge x' = f(\overline{z}) \rightarrow \overline{y} \subseteq \overline{z}$
A3'.	$x \not\subseteq \emptyset \wedge x \subseteq y \rightarrow x \not\parallel y$ and $x \subseteq y \wedge x \not\parallel z \rightarrow y \not\parallel z$ and $x \not\parallel y \rightarrow y \not\parallel x$
A4.	$f(\overline{y}) \subseteq x \wedge x \not\parallel x' \wedge x' \subseteq g(\overline{z}) \rightarrow \bot$ for $f \neq g$
A5.	$x = f(\overline{y}) \wedge x \not\parallel x' \wedge x' = f(\overline{z}) \rightarrow \overline{y} \not\parallel \overline{z}$
A6'.	$x \not\subseteq \emptyset \wedge \varphi \rightarrow \bot$ if $x \overset{\varphi}{\rightsquigarrow}_p x$ for some path $p \neq \varepsilon$
B7.	$x \not\subseteq \emptyset \wedge x = f(\overline{y}) \rightarrow \overline{y} \not\subseteq \emptyset$ and $\overline{y} \not\subseteq \emptyset \wedge x = f(\overline{y}) \rightarrow x \not\subseteq \emptyset$
B8.	$x \not\parallel y \rightarrow y \not\subseteq \emptyset$

Table 3. Axioms for inclusion constraints over (possibly empty) sets of finite trees

In Table 3, we present the set of axioms B, which adapts the set A for the new constraints. The axiom sets A2′, A3′, and A6′ are changed to make implicit non-emptiness premises explicit, and B7 and B8 have been added. B7 propagates non-emptiness through terms. For every constant symbol $a \in \Sigma$, B7.2 postulates that $x = a \rightarrow x \not\subseteq \emptyset$. B8 states that variables involved in a non-disjointness constraint must have a non-empty denotation themselves. It is easily checked that all these axioms are valid in $\mathsf{P}(\mathrm{Tree}_{\mathrm{fin}})$ and that all axioms apart from A6 are valid in $\mathsf{P}(\mathrm{Tree})$.

Proposition 17. *Every* B*-closed constraint is satisfiable over* $\mathsf{P}(\mathrm{Tree}_{\mathrm{fin}})$. *Every constraint that is* B*-closed apart from the occurs-check axiom* A6*' is satisfiable over* $\mathsf{P}(\mathrm{Tree})$.

Proof. Given a B-closed constraint φ, we define the set of variables in φ which are constrained to be non-empty.

$$\mathsf{Var}^{\varphi}_{\not\subseteq\emptyset} \overset{def}{=} \{x \mid x \not\subseteq \emptyset \text{ in } \varphi\}$$

The part ψ of φ containing only variables $\mathsf{Var}^{\varphi}_{\not\subseteq\emptyset}$ and no non-emptiness constraints is a closed Ines constraint.

For the finite tree case, assume φ to be B-closed. By Proposition 15 there exists a variable assignment α into $\mathsf{P}^+(\mathrm{Tree}_{\mathrm{fin}})$ which satisfies ψ in $\mathsf{P}(\mathrm{Tree}_{\mathrm{fin}})$.

Define the variable assignment β by $\beta(x) = \emptyset$ for $x \notin \mathsf{Var}^{\varphi}_{\not\subseteq\emptyset}$ and $\beta(x) = \alpha(x)$ elsewhere. We show that β satisfies φ. Let $x \notin \mathsf{Var}^{\varphi}_{\not\subseteq\emptyset}$. We consider the inclusions in φ containing x.

- Constraints $x \subseteq y$ are trivially satisfied by β.
- If $y = f(\ldots x \ldots)$ *in* φ, then $y \not\subseteq \emptyset$ cannot be in φ due to B7.2. Hence $\beta(y) = \emptyset \subseteq \beta(f(\ldots x \ldots))$.
- If $x = f(y_1 \ldots y_n)$ *in* φ, then $y_i \not\subseteq \emptyset$ cannot be in φ for some y_i due to B7.1. Hence $\beta(x) = \emptyset = \beta(f(\ldots y_i \ldots))$.
- If $y \subseteq x$ *in* φ, then $y \not\subseteq \emptyset$ cannot be in φ due to A3′ and B8. Hence $\beta(y) = \emptyset \subseteq \beta(x)$.

For the infinite tree case assume φ to be B-closed with the exception of A6'. Then by Lemma 9, there exists a satisfying variable assignment α into $\mathsf{P}^+(\mathrm{Tree})$. Apart from that, the above argument is unchanged. □

Theorem 18. *The satisfiability of conjunctions of inclusion constraints and non-emptiness constraints over sets of finite trees can be tested in* $O(n^3)$.

Proof. The axioms in Table 3 again induce a fixed point algorithm for the satisfiability test. By carrying over the techniques for the complexity results from Section C, we obtain the same complexity bound. □

Atomic Set Constraints. INES constraints interpreted over all sets of finite trees $\mathsf{P}(\mathrm{Tree}_{\mathrm{fin}})$ are also called *atomic set constraints* [15]. Theorem 18 implies time complexity $O(n^3)$ for their satisfiability problem. Furthermore, we show that the occurs check axiom A6′ is not needed to decide satisfiability of atomic set constraints.

Lemma 19. *Let the constraint* φ *be* B*-closed with the exception of* A6′. *Then* φ *does not imply* $\bot$ *according to axiom* A6′.

Proof. If $x \overset{\varphi}{\rightsquigarrow}_p x$ for some $p \neq \varepsilon$, then also $x \overset{\varphi}{\rightsquigarrow}_{p^n} x$ for every path $p^n = pp \ldots p$ (n-fold concatenation). Thus, for every prefix q of such a path p^n, there exists a non-constant function symbol $f \in \Sigma$ and a term $f(\overline{y})$ such that $x \overset{\varphi}{\rightsquigarrow}_q f(\overline{y})$.

If $x \not\subseteq \emptyset \in \varphi$ then φ contains a conjunction expressing that $t \subseteq x$ for some ground term t. (2)

But then there exist $n \geq 1$ and a prefix q of p^n leading to a leaf in t. Thus, $(q, a) \in t$ for some constant symbol $a \in \Sigma$. If $x \overset{\varphi}{\rightsquigarrow}_q f(\overline{y})$, we can show by induction over q that there exist z, z' such that $z = a$, $z \subseteq z'$, and $z' = f(\overline{y})$ *in* φ. From B7.2, and A3′ we obtain $z \not\parallel z'$ and hence, $\bot$ is a consequence of Axiom A4 which contradicts the assumption. □

From Lemma 19 and Theorem 18 we have the following Corollary[. Note that this is in contrast to Examples 6 and 7.

Corollary 20. *The satisfiability of atomic set constraints is invariant with respect to the interpretation over sets of finite or infinite trees.*