

Metalevel Algorithms For Variant Satisfiability

Stephen Skeirik and José Meseguer

Department of Computer Science
University of Illinois at Urbana-Champaign, USA

Abstract. Variant satisfiability is a theory-generic algorithm to decide quantifier-free satisfiability in an initial algebra $T_{\Sigma/E}$ when the theory (Σ, E) has the finite variant property and its constructors satisfy a compactness condition. This paper: (i) gives a precise definition of several *meta-level sub-algorithms* needed for variant satisfiability; (ii) proves them correct; and (iii) presents a *reflective implementation* in Maude 2.7.1 of variant satisfiability checking using these sub-algorithms.

Keywords: finite variant property (FVP), folding variant narrowing, satisfiability in initial algebras, metalevel algorithms, reflection, Maude.

1 Introduction

SMT solving is at the heart of some of the most effective theorem proving and infinite-state model checking formal verification methods that can scale up to impressive verification tasks. A current limitation, however, is its *lack of extensibility*: current SMT solvers support a (typically small) library of decidable theories. Although these theories can be combined by the Nelson-Oppen (NO) [32, 33] or Shostak [35] methods under some conditions, only the theories in the SMT solver library and their combinations are available to the user: any other theories extending the tool must be implemented by the tool builders.

In practice, of course, the problem a user has to solve may not be expressible by the theories available in an SMT solver's library. Therefore, the goal of making SMT solvers *user-extensible*, so that a *user* can easily *define* new decidable theories and use them in the verification process is highly desirable.


For a well-known *subproblem* of SMT solving, such user extensibility has recently been achieved: *E-unifiability* is the subproblem of *satisfiability* defined by: (i) considering theories of the form $th(T_{\Sigma/E}(X))$, associated to equational theories (Σ, E) , where $th(T_{\Sigma/E}(X))$ denotes the theory of the free (Σ, E) -algebra $T_{\Sigma/E}(X)$ on countably many variables X , and (ii) restricting ourselves to *positive* (i.e., negation-free) quantifier-free (QF) formulas. Lack of extensibility was the same: a unification tool supports a usually small library of theories (Σ, E) , which can be combined by methods similar to the NO one (the paper [2] explicitly relates the NO algorithm and combination algorithms for unification). Again, the *user* could not extend such *decidable* unifiability/unification algorithms by defining new theories and using a *theory-generic* algorithm. This is now possible for theories (Σ, E) satisfying the *finite variant property* (FVP) [13] thanks

to *variant unification* based on *folding variant narrowing* [18]. In fact, variant unification for user-definable FVP theories is already supported by Maude 2.7.1.

This suggests an obvious question: could variant unification be generalized to *variant satisfiability*, so that, under suitable conditions on an FVP theory (Σ, E) , satisfiability of QF formulas in the initial algebra $T_{\Sigma/E}$ becomes *decidable* by a *theory-generic* satisfiability algorithm? This would then make satisfiability *user-extensible* as desired. This question has been positively answered in [28–30] by giving general conditions under which satisfiability of QF formulas in the initial algebra $T_{\Sigma/E}$ of an FVP theory (Σ, E) is decidable. Section 3 summarizes the main results from [28–30]; but the punchline is easy to summarize: Suppose that: (i) the convergent rewrite theory $\mathcal{R} = (\Sigma, B, R)$ is a so-called FVP decomposition of (Σ, E) (which is what it means for (Σ, E) to be FVP), (ii) B has a finitary B -unification algorithm, and (iii) \mathcal{R} has an *OS-compact* constructor decomposition \mathcal{R}_Ω (definition in Section 3). Then satisfiability of QF formulas in $T_{\Sigma/E}$ is decidable by a *theory-generic* algorithm called variant satisfiability.

What this paper is about. The results in [28–30] do not really provide an *algorithm* in the full sense of the word, but rather a theoretical *skeleton* on which such an algorithm can be fleshed out. Specifically, they *assume* that the constructor decomposition \mathcal{R}_Ω is *OS-compact*, but do not provide a way to *automate* both the checking of OS-compactness and the implementation of the various *auxiliary functions* needed for variant satisfiability based on OS-compactness. They also use the notions of *constructor variant* and *constructor unifier* (see Section 3), but give only their theoretical definitions instead of algorithms to compute them.

Main Contributions. A *theory-generic* algorithm such as variant satisfiability manipulates *metalevel* data structures such as theories, signatures, equations, disequations, rewrite rules, and the like. In this paper we provide for the first time: (i) a full-fledged algorithm for variant satisfiability with its sub-algorithms; (ii) a proof of its correctness; and (iii) a reflective Maude implementation of it. The algorithm uses the following *auxiliary functions*:



These functions automate solutions for the two main subproblems already mentioned: (a) checking and satisfiability in OS-compact theories; and (b) computing constructor variants and constructor unifiers. These sub-algorithms are defined and proved correct *at the metalevel of rewriting logic*. Since rewriting logic is *reflective* [10], the correctness-preserving passage from the metalevel description of the sub-algorithms to their implementations is very direct: we just *meta-represent* them at the logic’s object level as suitable meta-level theories extending Maude’s META-LEVEL module [8].

This paper is a substantially extended version of the conference paper [36]. In comparison with the conference version, the following are totally new additions:

1. Proofs of all results are given.
2. More examples are given throughout the paper to illustrate key notions and point out counterexamples, including several substantial new examples in Sections 4.3 and 5.1.
3. A description of important new optimizations, now supported by the tool, of the algorithms for generating constructor variants and constructor unifiers has been added in Section 4.4.
4. A new section describing several *descent maps* in the sense of [28, 30] now supported by the tool to both widen the scope of theories it can handle and increase efficiency has also been added (Section 4.5).

2 Preliminaries on Order-Sorted Algebra and Rewriting

The material is adapted from [26, 18, 29]. Due to space limitations the following elementary notions, which can be found in [26], are assumed known: (i) order-sorted (OS) signature Σ ; (ii) set $\hat{S} = (S/\equiv_{\leq})$ of connected components (denoted $[s] \in \hat{S}$) of a poset of sorts (S, \leq) ; (iii) sensible OS signature; (iv) order-sorted Σ -algebras and homomorphisms, and its associated category \mathbf{OSAlg}_{Σ} ; and (v) the construction of the term algebra T_{Σ} and its initiality in \mathbf{OSAlg}_{Σ} when Σ is sensible. Furthermore, for connected components $[s_1], \dots, [s_n], [s] \in \hat{S}$,

$$f_{[s]}^{[s_1] \dots [s_n]} = \{f : s'_1 \dots s'_n \rightarrow s' \in \Sigma \mid s'_i \in [s_i], 1 \leq i \leq n, s' \in [s]\}$$

denotes the family of “subsort polymorphic” operators f .

T_{Σ} will (ambiguously) denote: (i) the term algebra; (ii) its underlying S -sorted set; and (iii) the set $T_{\Sigma} = \bigcup_{s \in S} T_{\Sigma, s}$. For $[s] \in \hat{S}$, $T_{\Sigma, [s]} = \bigcup_{s' \in [s]} T_{\Sigma, s'}$. An OS signature Σ is said to *have non-empty sorts* iff for each $s \in S$, $T_{\Sigma, s} \neq \emptyset$. We will assume throughout that Σ has non-empty sorts. An OS signature Σ is called *preregular* [19] iff for each $t \in T_{\Sigma}$ the set $\{s \in S \mid t \in T_{\Sigma, s}\}$ has a least element, denoted $ls(t)$. We will assume throughout that Σ is *preregular*.

An S -sorted set $X = \{X_s\}_{s \in S}$ of *variables*, satisfies $s \neq s' \Rightarrow X_s \cap X_{s'} = \emptyset$, and the variables in X are always assumed disjoint from all constants in Σ . The Σ -*term algebra* on variables X , $T_{\Sigma}(X)$, is the *initial algebra* for the signature $\Sigma(X)$ obtained by adding to Σ the variables X as *extra constants*. Since a $\Sigma(X)$ -algebra is just a pair (A, α) , with A a Σ -algebra, and α an *interpretation of the constants* in X , i.e., an S -sorted function $\alpha \in [X \rightarrow A]$, the $\Sigma(X)$ -initiality of $T_{\Sigma}(X)$ can be expressed as the following theorem:

Theorem 1. (*Freeness Theorem*). *If Σ is sensible, for each $A \in \mathbf{OSAlg}_{\Sigma}$ and $\alpha \in [X \rightarrow A]$, there exists a unique Σ -homomorphism, $\lrcorner \alpha : T_{\Sigma}(X) \rightarrow A$ extending α , i.e., such that for each $s \in S$ and $x \in X_s$ we have $x\alpha_s = \alpha_s(x)$.*

In particular, when $A = T_\Sigma(X)$, an interpretation of the constants in X , i.e., an S -sorted function $\sigma \in [X \rightarrow T_\Sigma(X)]$ is called a *substitution*, and its unique homomorphic extension $\cdot\sigma : T_\Sigma(X) \rightarrow T_\Sigma(X)$ is also called a substitution. Define $\text{dom}(\sigma) = \{x \in X \mid x \neq x\sigma\}$, and $\text{ran}(\sigma) = \bigcup_{x \in \text{dom}(\sigma)} \text{vars}(x\sigma)$. A *variable specialization* is a substitution ρ that just renames a few variables and may lower their sort. More precisely, $\text{dom}(\rho)$ is a finite set of variables $\{x_1, \dots, x_n\}$, with respective sorts s_1, \dots, s_n , and ρ injectively maps the x_1, \dots, x_n to variables x'_1, \dots, x'_n with respective sorts s'_1, \dots, s'_n such that $s'_i \leq s_i$, $1 \leq i \leq n$.

The first-order language of *equational Σ -formulas* is defined in the usual way: its atoms are Σ -equations $t = t'$, where $t, t' \in T_\Sigma(X)_{[s]}$ for some $[s] \in \widehat{S}$ and each X_s is assumed countably infinite. The set $\text{Form}(\Sigma)$ of *equational Σ -formulas* is then inductively built from atoms by: conjunction (\wedge), disjunction (\vee), negation (\neg), and universal ($\forall x:s$) and existential ($\exists x:s$) quantification with sorted variables $x:s \in X_s$ for some $s \in S$. The literal $\neg(t = t')$ is denoted $t \neq t'$. Given a Σ -algebra A , a formula $\varphi \in \text{Form}(\Sigma)$, and an assignment $\alpha \in [Y \rightarrow A]$, with $Y = \text{fvars}(\varphi)$ the free variables of φ , the *satisfaction relation* $A, \alpha \models \varphi$ is defined inductively as usual: for atoms, $A, \alpha \models t = t'$ iff $t\alpha = t'\alpha$; for Boolean connectives it is the corresponding Boolean combination of the satisfaction relations for subformulas; and for quantifiers: $A, \alpha \models (\forall x:s) \varphi$ (resp. $A, \alpha \models (\exists x:s) \varphi$) holds iff for all $a \in A_s$ (resp. some $a \in A_s$) we have $A, \alpha \uplus \{(x:s, a)\} \models \varphi$, where the assignment $\alpha \uplus \{(x:s, a)\}$ extends α by mapping $x:s$ to a . Finally, $A \models \varphi$ holds iff $A, \alpha \models \varphi$ holds for each $\alpha \in [Y \rightarrow A]$, where $Y = \text{fvars}(\varphi)$. We say that φ is *valid* (or *true*) in A iff $A \models \varphi$. We say that φ is *satisfiable* in A iff $\exists \alpha \in [Y \rightarrow A]$ such that $A, \alpha \models \varphi$, where $Y = \text{fvars}(\varphi)$. For a subsignature $\Omega \subseteq \Sigma$ and $A \in \mathbf{OSAlg}_\Sigma$, the *reduct* $A|_\Omega \in \mathbf{OSAlg}_\Omega$ agrees with A in the interpretation of all sorts and operations in Ω and discards everything in $\Sigma - \Omega$. If $\varphi \in \text{Form}(\Omega)$ we have the equivalence $A \models \varphi \Leftrightarrow A|_\Omega \models \varphi$.

An OS *equational theory* is a pair $T = (\Sigma, E)$, with E a set of Σ -equations. $\mathbf{OSAlg}_{(\Sigma, E)}$ denotes the full subcategory of \mathbf{OSAlg}_Σ with objects those $A \in \mathbf{OSAlg}_\Sigma$ such that $A \models E$, called the *(Σ, E) -algebras*. $\mathbf{OSAlg}_{(\Sigma, E)}$ has an *initial algebra* $T_{\Sigma/E}$ [26]. Given $T = (\Sigma, E)$ and $\varphi \in \text{Form}(\Sigma)$, we call φ *T -valid*, written $E \models \varphi$, iff $A \models \varphi$ for each $A \in \mathbf{OSAlg}_{(\Sigma, E)}$. We call φ *T -satisfiable* iff there exists $A \in \mathbf{OSAlg}_{(\Sigma, E)}$ with φ satisfiable in A . Note that φ is *T -valid* iff $\neg\varphi$ is *T -unsatisfiable*. The inference system in [26] is *sound and complete* for OS equational deduction, i.e., for any OS equational theory (Σ, E) , and Σ -equation $u = v$ we have an equivalence $E \vdash u = v \Leftrightarrow E \models u = v$. Deducibility $E \vdash u = v$ is abbreviated as $u =_E v$, called *E -equality*. An *E -unifier* of a system of Σ -equations, i.e., a conjunction $\phi = u_1 = v_1 \wedge \dots \wedge u_n = v_n$ of Σ -equations is a substitution σ such that $u_i\sigma =_E v_i\sigma$, $1 \leq i \leq n$. An *E -unification algorithm* for (Σ, E) is an algorithm generating a *complete set* of E -unifiers $\text{Unif}_E(\phi)$ for any system of Σ equations ϕ , where “complete” means that for any E -unifier σ of ϕ there is a $\tau \in \text{Unif}_E(\phi)$ and a substitution ρ such that $\sigma =_E \tau\rho$, where $=_E$ here means that for any variable x we have $x\sigma =_E x\tau\rho$. The algorithm is *finitary* if it always terminates with a *finite set* $\text{Unif}_E(\phi)$ for any ϕ .

Given a set of equations B used for deduction modulo B , a preregular OS signature Σ is called B -preregular¹ iff for each $u = v \in B$ and variable specialization ρ , $ls(u\rho) = ls(v\rho)$.

In the above logical notions the lack of predicate symbols is only *apparent*: full order-sorted first-order logic can be *reduced* to order-sorted algebra and equational formulas. The essential idea is to view a predicate $p(x_1:s_1, \dots, x_n:s_n)$ as a function symbol $p : s_1 \dots s_n \rightarrow Pred$, with $Pred$, a new sort having a constant tt . An atomic formula $p(t_1, \dots, t_n)$ is then expressed as the equation $p(t_1, \dots, t_n) = tt$. We refer the reader to [28, 29] for a detailed account of this reduction of predicate symbols to function symbols.

Recall the notation for term positions, subterms, and term replacement from [14]: (i) positions in a term viewed as a tree are marked by strings $p \in \mathbb{N}^*$ specifying a path from the root, (ii) $t|_p$ denotes the subterm of term t at position p , and (iii) $t[u]_p$ denotes the result of *replacing* subterm $t|_p$ at position p by u .

Definition 1. A rewrite theory is a triple $\mathcal{R} = (\Sigma, B, R)$ with (Σ, B) an order-sorted equational theory and R a set of Σ -rewrite rules, i.e., sequents $l \rightarrow r$, with $l, r \in T_\Sigma(X)_{[s]}$ for some $[s] \in \hat{S}$. In what follows it is always assumed that:

1. For each $l \rightarrow r \in R$, $l \notin X$ and $\text{vars}(r) \subseteq \text{vars}(l)$.
2. Each rule $l \rightarrow r \in R$ is sort-decreasing, i.e., for each variable specialization ρ , $ls(l\rho) \geq ls(r\rho)$.
3. Σ is B -preregular (if $B = B_0 \uplus U$, in the broader sense of Footnote 1).
4. Each equation $u = v \in B$ is regular, i.e., $\text{vars}(u) = \text{vars}(v)$, and linear, i.e., there are no repeated variables in u , and no repeated variables in v .

The one-step R, B -rewrite relation $t \rightarrow_{R,B} t'$, holds between $t, t' \in T_\Sigma(X)_{[s]}$, $[s] \in \hat{S}$, iff there is a rewrite rule $l \rightarrow r \in R$, a substitution $\sigma \in [X \rightarrow T_\Sigma(X)]$, and a term position p in t such that $t|_p =_B l\sigma$, and $t' = t[r\sigma]_p$. Note that, by assumptions (2)–(3) above, $t[r\sigma]_p$ is always a well-formed Σ -term.

\mathcal{R} is called: (i) *terminating* iff the relation $\rightarrow_{R,B}$ is well-founded; (ii) *strictly B -coherent* [27] iff whenever $u \rightarrow_{R,B} v$ and $u =_B u'$ there is a v' such that $u' \rightarrow_{R,B} v'$ and $v =_B v'$; (iii) *confluent* iff $u \rightarrow_{R,B}^* v_1$ and $u \rightarrow_{R,B}^* v_2$ imply that there are w_1, w_2 such that $v_1 \rightarrow_{R,B}^* w_1$, $v_2 \rightarrow_{R,B}^* w_2$, and $w_1 =_B w_2$ (where $\rightarrow_{R,B}^*$ denotes the reflexive-transitive closure of $\rightarrow_{R,B}$); and (iv) *convergent* if (i)–(iii) hold. If \mathcal{R} is convergent, for each Σ -term t there is a term u such that $t \rightarrow_{R,B}^* u$ and $(\nexists v) u \rightarrow_{R,B} v$. We then write $u = t!_{R,B}$, and call $t!_{R,B}$ the R, B -normal form of t , which, by confluence, is unique up to B -equality.

¹ When the axioms B consist of a combination of associativity, commutativity, and (left and/or right) identity axioms, we can decompose B into the disjoint union $B = B_0 \uplus U$, where B_0 are associativity and/or commutativity axioms, and U are left and/or right identity axioms. The equations in U , of the general form $f(e, x) = x$ and/or $f(x, e) = x$, can be oriented as rewrite rules $R(U)$ of the form $f(e, x) \rightarrow x$ and/or $f(x, e) \rightarrow x$ to be applied *modulo* B_0 . The B -preregularity notion can then be *broadened* by requiring only that: (i) Σ is preregular; (ii) Σ is B_0 -preregular in the standard sense that $ls(u\rho) = ls(v\rho)$ for all $u = v \in B_0$ and sort specializations ρ ; and (iii) the rules $R(U)$ are *sort-decreasing* in the sense of Definition 1. Maude automatically checks B -preregularity of an OS signature Σ in this broader sense [8].

Given a set E of Σ -equations, let $R(E) = \{u \rightarrow v \mid u = v \in E\}$. A *decomposition* of an order-sorted equational theory (Σ, E) is a convergent rewrite theory $\mathcal{R} = (\Sigma, B, R)$ such that $E = E_0 \uplus B$ and $R = R(E_0)$. The key property of a decomposition is the following:

Theorem 2. (*Church-Rosser Theorem*) [22, 27] *Let $\mathcal{R} = (\Sigma, B, R)$ be a decomposition of (Σ, E) . Then we have an equivalence:*

$$E \vdash u = v \Leftrightarrow u!_{R,B} =_B v!_{R,B}.$$

If $\mathcal{R} = (\Sigma, B, R)$ is a decomposition of (Σ, E) , and X an S -sorted set of variables, the *canonical term algebra* $C_{\mathcal{R}}(X)$ has $C_{\mathcal{R}}(X)_s = \{[t!_{R,B}]_B \mid t \in T_{\Sigma}(X)_s\}$, and interprets each $f : s_1 \dots s_n \rightarrow s$ as the function $C_{\mathcal{R}}(X)_f : ([u_1]_B, \dots, [u_n]_B) \mapsto [f(u_1, \dots, u_n)!_{R,B}]_B$. By the Church-Rosser Theorem we then have an isomorphism $h : T_{\Sigma/E}(X) \cong C_{\mathcal{R}}(X)$, where $h : [t]_E \mapsto [t!_{R,B}]_B$. In particular, when X is the empty family of variables, the canonical term algebra $C_{\mathcal{R}}$ is an initial algebra, and is the most intuitive possible model for $T_{\Sigma/E}$ as an algebra of *values* computed by R, B -simplification.

Quite often, the signature Σ on which $T_{\Sigma/E}$ is defined has a natural decomposition as a disjoint union $\Sigma = \Omega \uplus \Delta$, where the elements of $C_{\mathcal{R}}$, that is, the *values* computed by R, B -simplification, are Ω -terms, whereas the function symbols $f \in \Delta$ are viewed as *defined functions* which are *evaluated away* by R, B -simplification. Ω (with same poset of sorts as Σ) is then called a *constructor subsignature* of Σ . Call a decomposition $\mathcal{R} = (\Sigma, B, R)$ of (Σ, E) *sufficiently complete* with respect to the *constructor subsignature* Ω iff for each $t \in T_{\Sigma}$ we have: (i) $t!_{R,B} \in T_{\Omega}$, and (ii) if $u \in T_{\Omega}$ and $u =_B v$, then $v \in T_{\Omega}$. This ensures that for each $[u]_B \in C_{\mathcal{R}}$ we have $[u]_B \subseteq T_{\Omega}$. Of course, we want Ω *as small as possible* with these properties. In Example 1 below, $\Omega = \{\top, \perp\}$ and $\Delta = \{- \wedge -, - \vee -\}$. Tools based on tree automata [11], equational tree automata [21], or narrowing [20], can be used to automatically check sufficient completeness of a decomposition \mathcal{R} with respect to constructors Ω under some assumptions.

Sufficient completeness is closely related to the notion of a *protecting* theory inclusion.

Definition 2. *An equational theory (Σ, E) protects another theory (Ω, E_{Ω}) iff $(\Omega, E_{\Omega}) \subseteq (\Sigma, E)$ and the unique Ω -homomorphism $h : T_{\Omega/E_{\Omega}} \rightarrow T_{\Sigma/E}|_{\Omega}$ is an isomorphism $h : T_{\Omega/E_{\Omega}} \cong T_{\Sigma/E}|_{\Omega}$.*

A decomposition $\mathcal{R} = (\Sigma, B, R)$ protects another decomposition $\mathcal{R}_0 = (\Sigma_0, B_0, R_0)$ iff $\mathcal{R}_0 \subseteq \mathcal{R}$, i.e., $\Sigma_0 \subseteq \Sigma$, $B_0 \subseteq B$, and $R_0 \subseteq R$, and for all $t, t' \in T_{\Sigma_0}(X)$ we have: (i) $t =_{B_0} t' \Leftrightarrow t =_B t'$, (ii) $t = t!_{R_0, B_0} \Leftrightarrow t = t!_{R, B}$, and (iii) $C_{\mathcal{R}_0} = C_{\mathcal{R}}|_{\Sigma_0}$.

*$\mathcal{R}_{\Omega} = (\Omega, B_{\Omega}, R_{\Omega})$ is a constructor decomposition of $\mathcal{R} = (\Sigma, B, R)$ iff \mathcal{R} protects \mathcal{R}_{Ω} and Σ and Ω have the same poset of sorts, so that by (iii) above \mathcal{R} is sufficiently complete with respect to Ω . Furthermore, Ω is called a *subsignature of free constructors modulo B_{Ω}* iff $R_{\Omega} = \emptyset$, so that $C_{\mathcal{R}_0} = T_{\Omega/B_{\Omega}}$.*

3 Variants and Variant Satisfiability

The notion of *variant* answers two questions: (i) how can we best describe symbolically the elements of $C_{\mathcal{R}}(X)$ that are *reduced substitution instances* of a given *pattern term* t ? and (ii) when is such a symbolic description *finite*?

Definition 3. Given a decomposition $\mathcal{R} = (\Sigma, B, R)$ of an OS equational theory (Σ, E) and a Σ -term t , a variant² [13, 18] of t is a pair (u, θ) such that: (i) $u =_B (t\theta)!$, (ii) if $x \notin \text{vars}(t)$, then $x\theta = x$, and (iii) $\theta = \theta!$, that is, $x\theta = (x\theta)!$ for all variables x . (u, θ) is called a ground variant iff $u \in T_{\Sigma}$. Note that if (u, θ) is a ground variant of some t , then $[u]_B \in C_{\mathcal{R}}$. Given variants (u, θ) and (v, γ) of t , (u, θ) is called more general than (v, γ) , denoted $(u, \theta) \supseteq_{R,B} (v, \gamma)$, iff there is a substitution ρ such that: (i) $\theta\rho =_B \gamma$, and (ii) $u\rho =_B v$. Let $\llbracket t \rrbracket_{R,B} = \{(u_i, \theta_i) \mid i \in I\}$ denote a most general complete set of variants of t , that is, a set of variants such that: (i) for any variant (v, γ) of t there is an $i \in I$, such that $(u_i, \theta_i) \supseteq_{R,B} (v, \gamma)$; and (ii) for $i, j \in I$, $i \neq j \Rightarrow ((u_i, \theta_i) \not\supseteq_{R,B} (u_j, \theta_j) \wedge (u_j, \theta_j) \not\supseteq_{R,B} (u_i, \theta_i))$. A decomposition $\mathcal{R} = (\Sigma, B, R)$ of (Σ, E) has the finite variant property [13] (FVP) iff for each Σ -term t there is a finite most general complete set of variants $\llbracket t \rrbracket_{R,B} = \{(u_1, \theta_1), \dots, (u_n, \theta_n)\}$.

If B has a finitary unification algorithm, the *folding variant narrowing* strategy described in [18] provides an effective method to generate $\llbracket t \rrbracket_{R,B}$. Furthermore, $\llbracket t \rrbracket_{R,B}$ is finite for each t , so that the strategy *terminates* iff \mathcal{R} is FVP.

Example 1. Let $\mathcal{B} = (\Sigma, B, R)$ with Σ having a single sort, say *Bool*, constants \top, \perp , and binary operators $-\wedge-$ and $-\vee-$, B the associativity and commutativity (AC) axioms for both $-\wedge-$ and $-\vee-$, and R the rules: $x \wedge \top \rightarrow x$, $x \wedge \perp \rightarrow \perp$, $x \vee \perp \rightarrow x$, and $x \vee \top \rightarrow \top$. Then \mathcal{B} is FVP. For example, $\llbracket x \wedge y \rrbracket_{R,B} = \{(x \wedge y, id), (y, \{x \mapsto \top\}), (x, \{y \mapsto \top\}), (\perp, \{x \mapsto \perp\}), (\perp, \{y \mapsto \perp\})\}$.

FVP is a *semi-decidable* property [7], which can be easily verified (when it holds) by checking, using folding variant narrowing, that for each function symbol f the term $f(x_1, \dots, x_n)$, with the sorts of the x_1, \dots, x_n those of f , has a finite number of most general variants.

Folding variant narrowing provides also a method for generating a *complete set of E-unifiers* when (Σ, E) has a decomposition $\mathcal{R} = (\Sigma, B, R)$ with B having a finitary B -unification algorithm [18]. To express systems of equations, say, $u_1 = v_1 \wedge \dots \wedge u_n = v_n$, as *terms*, we can extend Σ to a signature Σ^\wedge by adding:

1. for each connected component $[s]$ that does not already have a top element, a fresh new sort $\top_{[s]}$ with $\top_{[s]} > s'$ for each $s' \in [s]$. In this way we obtain a (possibly extended) poset of sorts (S_{\top}, \supseteq) ;
2. fresh new sorts *Lit* and *Conj* with a subsort inclusion $\text{Lit} < \text{Conj}$, with a binary conjunction operator $-\wedge- : \text{Lit Conj} \rightarrow \text{Conj}$, and

² For a discussion of similar but not exactly equivalent versions of the variant notion see [7]. Here we follow the formulation in [18].

3. for each connected component $[s] \in \widehat{S}_\top$ with top sort $\top_{[s]}$, binary operators $- = - : \top_{[s]} \top_{[s]} \rightarrow Lit$ and $- \neq - : \top_{[s]} \top_{[s]} \rightarrow Lit$.

Theorem 3. [29] *Under the above assumptions on \mathcal{R} , let $\phi = u_1 = v_1 \wedge \dots \wedge u_n = v_n$ be a system of Σ -equations viewed as a Σ^\wedge -term of sort *Conj*. Then*

$$\{\theta\gamma \mid (\phi', \theta) \in \llbracket \phi \rrbracket_{R,B} \wedge \gamma \in \text{Unif}_B(\phi') \wedge (\phi'\gamma, \theta\gamma) \text{ is a variant of } \phi\}$$

is a complete set of E -unifiers for ϕ , where $\text{Unif}_B(\phi')$ denotes a complete set of most general B -unifiers for each variant $\phi' = u'_1 = v'_1 \wedge \dots \wedge u'_n = v'_n$.

Since if $\mathcal{R} = (\Sigma, B, R)$ is FVP, then $\mathcal{R}^\wedge = (\Sigma^\wedge, B, R)$ is also FVP, Theorem 3 shows that if a finitary B -unification algorithm exists and \mathcal{R} is an FVP decomposition of (Σ, E) , then E has a finitary E -unification algorithm.

The key question asked and answered in [28, 29] is: given an FVP decomposition $\mathcal{R} = (\Sigma, B, R)$ of an equational theory (Σ, E) , under what conditions is satisfiability of QF equational Σ -formulas in the canonical term algebra $C_{\mathcal{R}}$ decidable? It turns out that: (i) \mathcal{R} having a constructor decomposition $\mathcal{R}_\Omega = (\Omega, B_\Omega, R_\Omega)$, and (ii) the associated notions of *constructor variant* and *constructor unifier* [29] play a crucial role in answering this question.

Definition 4. *Let $\mathcal{R} = (\Sigma, B, R)$ be a decomposition of (Σ, E) , and let $\mathcal{R}_\Omega = (\Omega, B_\Omega, R_\Omega)$ be a constructor decomposition of \mathcal{R} . Then an R, B -variant (u, θ) of a Σ -term t is called a constructor R, B -variant of t iff $u \in T_\Omega(X)$.*

Suppose, furthermore, that B has a finitary B -unification algorithm, so that, given a unification problem $\phi = u_1 = v_1 \wedge \dots \wedge u_n = v_n$, Theorem 3 allows us to generate the complete set of E -unifiers

$$\{\theta\gamma \mid (\phi', \theta) \in \llbracket \phi \rrbracket_{R,B} \wedge \gamma \in \text{Unif}_B(\phi') \wedge (\phi'\gamma, \theta\gamma) \text{ is a variant of } \phi\}$$

Then a constructor E -unifier³ of ϕ is either: (1) a unifier $\theta\gamma$ in the above set with $\phi'\gamma \in T_{\Omega^\wedge}(X)$; or otherwise, (2) a unifier $\theta\gamma\alpha$ such that: (i) $\theta\gamma$ belongs the above set, (ii) α is a substitution of the variables in $\text{ran}(\theta\gamma)$ such that $\phi'\gamma\alpha \in T_{\Omega^\wedge}(X)$, and (iii) $(\phi'\gamma\alpha, \theta\gamma\alpha)$ is a variant of ϕ . $\text{mgu}_{\mathcal{R}}^\Omega(\phi)$ denotes a set of most general constructor E -unifiers of ϕ , i.e., for any constructor E -unifier μ of ϕ there is another one $\eta \in \text{mgu}_{\mathcal{R}}^\Omega(\phi)$ and a substitution ν such that $\mu =_B \eta\nu$.

Note that if (v, δ) is a ground variant of t , then $[v]_B \in C_{\mathcal{R}}$, so that v is an Ω -term. Therefore, any ground variant (v, δ) of t is “covered” by some constructor variant (u, θ) of t , i.e., $(u, \theta) \supseteq_{R,B} (v, \delta)$. If (Σ, E) has a decomposition $\mathcal{R} = (\Sigma, B, R)$, B has a finitary B -unification algorithm and we are only interested in characterizing the *ground solutions* of an equation in the initial algebra $T_{\Sigma/E}$, only constructor E -unifiers are needed, since they completely cover all such solutions. Likewise, if we are only interested in *unifiability* of a system of equations only constructor E -unifiers are needed.

Theorem 4. [28, 29] *Let (Σ, E) have a decomposition $\mathcal{R} = (\Sigma, B, R)$ with B having a finitary B -unification algorithm. Then, for each system of Σ -equations $\phi = u_1 = v_1 \wedge \dots \wedge u_n = v_n$, where $Y = \text{vars}(\phi)$, we have:*

³ [28, 29] give examples of constructor variants and constructor unifiers.

1. (Completeness for Ground Unifiers). If $\delta \in [Y \rightarrow T_\Sigma]$ is a ground E -unifier of ϕ , then there is a constructor E -unifier $\eta \in \text{mgu}_{\mathcal{R}}^Q(\phi)$ and a substitution β such that $\delta =_E \eta\beta$, i.e., $x\delta =_E x\eta\beta$ for each variable $x \in Y$.
2. (Unifiability). $T_{\Sigma/E} \models (\exists Y) \phi$ iff ϕ has a constructor E -unifier.

Given an OS equational theory (Σ, E) , call a Σ -equality $u = v$ E -trivial iff $u =_E v$, and a Σ -disequality $u \neq v$ E -consistent iff $u \neq_E v$. Likewise, call a conjunction $\bigwedge D$ of Σ -disequalities E -consistent iff each $u \neq v$ in D is so.

Theorem 4 is a key step to find conditions for the decidable satisfiability of QF equational Σ -formulas in $C_{\mathcal{R}}$ for $\mathcal{R} = (\Sigma, B, R)$ an FVP decomposition of (Σ, E) , where B has a finitary B -unification algorithm and \mathcal{R} has a constructor decomposition $\mathcal{R}_\Omega = (\Omega, B_\Omega, R_\Omega)$. The key idea is to reduce the problem to one of satisfiability of a conjunction of Ω -disequalities in the simpler canonical term algebra $C_{\mathcal{R}_\Omega}$. By $C_{\mathcal{R}}|_\Omega = C_{\mathcal{R}_\Omega}$, Theorem 4, and the Descent Theorems in [28, 29] (see [28, 29] for full details), we can apply the following algorithm to a conjunction of literals $\phi = \bigwedge G \wedge \bigwedge D$, with G equations and D disequations:

1. Thanks to Theorem 4 we need only compute the constructor E -unifiers $\text{mgu}_{\mathcal{R}}^Q(\bigwedge G)$, and reduce to the case of deciding the satisfiability of some conjunction of disequalities $(\bigwedge D\alpha)!_{R,B}$, for some $\alpha \in \text{mgu}_{\mathcal{R}}^Q(\bigwedge G)$, discarding any $(\bigwedge D\alpha)!_{R,B}$ containing a B -inconsistent disequality.
2. For each remaining $(\bigwedge D\alpha)!_{R,B}$ we can then compute a finite, complete set of most general R, B -variants $\llbracket (\bigwedge D\alpha)!_{R,B} \rrbracket_{R,B}$ by folding variant narrowing, and obtain for each of them its B_Ω -consistent constructor variants $\bigwedge D'$.
3. Then by the Descent Theorems in [28, 29], ϕ will be satisfiable in $C_{\mathcal{R}}$ iff $\bigwedge D'$ is satisfiable in $C_{\mathcal{R}_\Omega}$ for some such $\bigwedge D'$ and some such α .

Therefore, the method hinges upon being able to decide when a conjunction of Ω -disequalities $\bigwedge D'$ is satisfiable in $C_{\mathcal{R}_\Omega}$. This is decidable if \mathcal{R}_Ω is the decomposition of an OS-compact theory, which generalizes the notion of compact theory in [12]:

Definition 5. [28, 29] An equational theory (Σ, E) is called OS-compact iff: (i) for each sort s in Σ we can effectively determine whether $T_{\Sigma/E, s}$ is finite or infinite, and, if finite, can effectively compute a representative ground term $\text{rep}([u]) \in [u]$ for each $[u] \in T_{\Sigma/E, s}$ (ii) $=_E$ is decidable and E has a finitary unification algorithm; and (iii) any E -consistent finite conjunction $\bigwedge D$ of Σ -disequalities whose variables all have infinite sorts is satisfiable in $T_{\Sigma/E}$.

The reason why satisfiability of a conjunction of disequalities in the initial algebra of an OS-compact theory is decidable [28, 29] is fairly obvious: by (iii) it is decidable when all variables have infinite sorts; and we can always reduce to a disjunction of formulas in that case by instantiating each variable with a finite sort s by all the possible representatives in $T_{\Sigma/E, s}$. Therefore we have:

Corollary 1. For $\mathcal{R} = (\Sigma, B, R)$ an FVP decomposition of (Σ, E) , where B has a finitary B -unification algorithm and \mathcal{R} has an OS-compact constructor decomposition \mathcal{R}_Ω , satisfiability of QF equational Σ -formulas in $C_{\mathcal{R}}$ is decidable.

The papers [28, 29] contain many examples of commonly used theories that have FVP specifications whose constructor decompositions are OS-compact. This can be established by one of the two methods discussed below.

A first method to show OS-compactness is both very simple and widely applicable to constructor decompositions of FVP theories. It applies to OS equational theories of the form $(\Omega, ACCU)$, where $ACCU$ stands for any combination of associativity and/or commutativity and/or left- or right-identity axioms, except combinations where the same operator is associative but not commutative. We also assume that if any typing for a binary operator f in a subsort-polymorphic family $f_{[s]}^{[s]}$ satisfies some axioms in $ACCU$, then any other typing in $f_{[s]}^{[s]}$ satisfies the *same* axioms. The following theorem generalizes to the order-sorted and $ACCU$ case a similar result in [12] for the unsorted and AC case:

Theorem 5. [28, 29] *Under the above assumptions $(\Omega, ACCU)$ is OS-compact. Furthermore, satisfiability of QF Ω -formulas in $T_{\Omega/ACCU}$ is decidable.*

The range of FVP theories whose initial algebras have decidable QF satisfiability is greatly increased by a second method of *satisfiability-preserving FVP parameterized theories*. For our present purposes it suffices to summarize the basic general facts and assumptions for the case of FVP parameterized data types with a *single parameter* X . That is, we can focus on parameterized FVP theories of the form $\mathcal{R}[X] = (\mathcal{R}, X)$, where $\mathcal{R} = (\Sigma, B, R)$ is an FVP decomposition of an OS equational theory (Σ, E) , and X is a sort in Σ (called the *parameter sort*) such that: (i) is empty, i.e., $T_{\Sigma, X} = \emptyset$; and (ii) X is a minimal element in the sort order, i.e., there is no other sort s' with $s' < X$.

Consider an FVP decomposition $\mathcal{G} = (\Sigma', B', R')$ of a finitary OS equational theory (Σ', E') , which we can assume without loss of generality is disjoint from (Σ, E) , and additionally let s be a sort in Σ' . Then the *instantiation* $\mathcal{R}[\mathcal{G}, X \mapsto s] = (\Sigma[\Sigma', X \mapsto s], B \cup B', R \cup R')$ is the decomposition of a theory $(\Sigma[\Sigma', X \mapsto s], E \cup E')$, extending (Σ', E') , where the signature $\Sigma[\Sigma', X \mapsto s]$ is defined as the union $\Sigma[X \mapsto s] \cup \Sigma'$, with $\Sigma[X \mapsto s]$ just like Σ , except for X renamed to s . Its set of sorts is $(S - \{X\}) \uplus S'$, and the poset ordering combines those of $\Sigma[X \mapsto s]$ and Σ' . Furthermore, $\mathcal{R}[\mathcal{G}, X \mapsto s]$ is also FVP under mild assumptions [28].

Suppose B , B' and $B \cup B'$ have finitary unification algorithms and both $\mathcal{R}[X] = (\mathcal{R}, X)$ and \mathcal{G} protect, respectively, the two constructor theories, say $\mathcal{R}_{\Omega}[X] = (\Omega, B_{\Omega}, R_{\Omega})$ and $\mathcal{G}_{\Omega'} = (\Omega', B_{\Omega'}, R_{\Omega'})$. Then $\mathcal{R}[\mathcal{G}, X \mapsto s]$ will protect $\mathcal{R}_{\Omega}[\mathcal{G}_{\Omega'}, X \mapsto s]$. Suppose, further, that B_{Ω} , $B_{\Omega'}$, and $B_{\Omega} \cup B_{\Omega'}$ have decidable equality. The general satisfiability-preserving method of interest is then as follows: (i) assuming that $\mathcal{G}_{\Omega'}$ is the decomposition of an OS-compact theory, then (ii) under some assumptions about the cardinality of the sort s , prove the OS-compactness of $\mathcal{R}_{\Omega}[\mathcal{G}_{\Omega'}, X \mapsto s]$. It then follows from our earlier reduction of satisfiability in initial FVP algebras to their constructor decompositions that satisfiability of QF formulas in the initial model of the instantiation $\mathcal{R}[\mathcal{G}, X \mapsto s]$ is decidable.

In [28] the following parameterized data types have been proved satisfiability-preserving following the just-described pattern of proof: (i) $\mathcal{L}[X]$, *parameterized*

lists, which is just an example illustrating the general case of any constructor-selector-based [31] parameterized data type; (ii) $\mathcal{L}^e[X]$, *parameterized compact lists*, where any two identical contiguous list elements are identified [16, 15]; (iii) $\mathcal{M}[X]$, *parameterized multisets*; (iv) $\mathcal{S}[X]$, *parameterized sets*; and (v) $\mathcal{H}[X]$, *parameterized hereditarily finite sets*.

4 Metalevel Algorithms for Variant Satisfiability

For $\mathcal{R} = (\Sigma, B, R)$ an FVP decomposition of (Σ, E) , where B has a finitary B -unification algorithm and \mathcal{R} has a constructor decomposition \mathcal{R}_Ω , the issue of the decidable satisfiability of QF equational Σ -formulas in $C_{\mathcal{R}}$ has been condensed in Section 3 to two key sub-issues: (i) steps (1)–(3) in the high-level algorithm, which reduce satisfiability of a conjunction of Σ -literals in $C_{\mathcal{R}}$ to satisfiability of a conjunction of Ω -disequalities in $C_{\mathcal{R}_\Omega}$; and (ii) decidable satisfiability of conjunctions of Ω -disequalities in $C_{\mathcal{R}_\Omega}$ when \mathcal{R}_Ω is OS-compact (Corollary 1).

At a theoretical level this gives the *skeleton* of a high-level algorithm for variant satisfiability. But at a concrete, algorithmic level several important questions, essential for having an actual satisfiability *algorithm*, remain unresolved, including: (1) how can we *automatically check* that the constructor decomposition \mathcal{R}_Ω is OS-compact using the two methods for OS-compactness outlined in Section 3? (2) how can we *compute* constructor variants and constructor unifiers? (3) how can we *prove* that the auxiliary algorithms answering questions (1) and (2) are *correct*? and (4) how can we *implement* both the main algorithm and the auxiliary algorithms in a correctness-preserving manner?

Let us begin with question (3). The algorithm skeleton sketched in Section 3 manipulates metalevel entities like operators, signatures, terms, equations, and theories. Likewise, the checks for OS-compactness and the computation of constructor variants and constructor unifiers (questions (1)–(2)) are problems fully expressible in terms of such metalevel entities. Therefore, both for mathematical clarity and for simplicity of the needed correctness proofs, the definitions of the auxiliary algorithms should be carried out at the metalevel of rewriting logic.

This brings us to question (4), which has a simple answer: since rewriting logic is *reflective* [10], once we have defined and proved correct at the metalevel the auxiliary algorithms solving questions (1) and (2), we can derive correct implementations for them by *meta-representing* them at the logic’s object level as equational or rewrite theories. In fact, this can be carried out in Maude by defining suitable meta-level theories extending the **META-LEVEL** module [8].

The previous paragraphs lead us to the main contributions of the present paper. We answer questions (1) and part of (3) by defining and proving correct at the metalevel a method to check OS-compactness, including: (a) checking which sorts s satisfy $|T_{\Omega/B_\Omega, s}| < \aleph_0$, and (b) computing for each such s a unique representative $rep([t]_{B_\Omega})$ for each $[t]_{B_\Omega} \in T_{\Omega/B_\Omega, s}$. We answer question (2) and the other part of (3) by defining and proving correct at the meta-level a method to compute constructor unifiers and constructor variants. Furthermore, for improved efficiency we also provide an optimized version of constructor variant and

unifier generation in Section 4.4; and discuss also the method of *descent maps* in the sense of [28, 30] —which can both increase efficiency and widen the scope of decidable theories— and some specific descent maps currently supported in Section 4.5. Finally, we answer question (4) by meta-representing both the auxiliary algorithms (proved correct in this section), and the main algorithm (already proved correct in [28–30]) in Section 5.

To help guide the discussion, the reader may refer to the tree diagram in the Introduction, which describes the dependencies among different subalgorithms.

4.1 Checking OS-Compactness

In this section we present a high-level description of the algorithms needed to check that a constructor decomposition $\mathcal{R}_\Omega = (\Omega, B_\Omega, R_\Omega)$ is OS-compact. Since these checks are auxiliary to the main functionality needed, namely, computing constructor variants and constructor unifiers, some details are omitted to ease readability; all remaining details, together with full proofs of correctness, can be found in Appendix A.

E_Ω-consistency of a conjunction of Ω -disequalities $\bigwedge D'$ in a constructor decomposition $\mathcal{R}_\Omega = (\Omega, B_\Omega, R_\Omega)$ is easy to check: we may assume $\bigwedge D'$ in R_Ω, B_Ω -normal form and just need to check that $u \neq_{B_\Omega} v$ for each $u \neq v$ in $\bigwedge D'$.

Checking that the constructor subtheory \mathcal{R}_Ω of \mathcal{R} is *OS-compact* breaks into two cases: (1) when \mathcal{R} is an *unparameterized theory*; and (2) when \mathcal{R} is the instantiation of a possibly *nested* collection of satisfiability-preserving *parameterized theories* such as, for, example, *sets of lists of natural numbers*. In case (2) it is enough (for the parameterized theories described in Section 3) to check that: (i) the unparameterized theory \mathcal{G} in the innermost instantiation (in our example the theory \mathcal{N}_+ of naturals with addition) is *OS-compact*, and the chosen sort (in our example the sort *Nat*) is *infinite*; and (ii) that the sorts chosen to instantiate each remaining parameter is the *principal sort* of the parameterized module immediately below in the nesting. In the above example this is just checking that the parameter sort X for the *set* parameterized module is instantiated to the principal sort, namely *List*, of the *list* parameterized module immediately below. In this way, checking OS-compactness of \mathcal{R}_Ω in the, nested, parameterized case is reduced to checking OS-compactness of the unparameterized inner argument, plus a check of an infinite sort. All checks for the unparameterized case (1), including the two needed in case (2), are described below.

OS-Compactness Check (Unparameterized Case). As shown in Theorem 5, a sufficient condition for an unparameterized constructor decomposition $\mathcal{R}_\Omega = (\Omega, B_\Omega, R_\Omega)$ to be OS-compact is for \mathcal{R}_Ω to be of the form $\mathcal{R}_\Omega = (\Omega, ACCU, \emptyset)$. Thus, a sufficient condition is to require: (1) B_Ω to be a set of ACCU axioms, and (2) Ω to be a signature of *free constructors* modulo B_Ω . Fortunately, both of these subgoals are quite simple to check. Goal (1) can be solved by iterating over each axiom and applying a case analysis against its structure. Goal (2) can be

solved by an application of propositional tree automata (PTA). In particular, if the rules R in \mathcal{R} are linear and unconditional, then constructor freeness modulo B is translatable into a PTA emptiness problem; see [34] for further details.

Finite Sort Classification. Another needed algorithm takes as input a signature Ω and a sort s and checks if $|T_{\Omega/B_{\Omega},s}| < \aleph_0$. We solve this problem in two phases: (1) we devise an algorithm to check $|T_{\Omega,s}| < \aleph_0$ (2) we use this as a subroutine in an *approximate* algorithm to check $|T_{\Omega/B_{\Omega},s}| < \aleph_0$ when $B_{\Omega} = ACCU$. If the approximate algorithm fails to classify some s as either infinite or finite, s returned to the user as a proof obligation (Appendix A, Corollary 6).

If Ω is finite and has non-empty sorts, we show that $|T_{\Omega,s}| = \aleph_0$ iff there exists a *cycle* in the relation $(<) \subseteq S^2$ reachable from s where $s < s'$ iff the formula $\exists f : s_1 \cdots s_n \rightarrow s'' \in \Omega \exists i \in \mathbb{N}[s'' \leq s \wedge s \leq s_i] \vee [s' < s]$ holds. We construct a rewrite theory R_F over S such that $s \rightarrow_{R_F} s'$ iff $s < s'$. If $cy(S) = \{s \in S \mid s \rightarrow_{R_F}^+ s\}$, then $s \rightarrow_{R_F}^* s'$ with $s' \in cy(S)$ implies $|T_{\Omega,s}| = \aleph_0$. Then $\bigvee_{s' \in cy(S \setminus \emptyset)} R_F \vdash s \rightarrow^* s'$ holds iff there is a cycle in the relation $(<)$ reachable from s (Appendix A, Theorem 10).

We now lift the algorithm above to phase (2). We can show that for ACC axioms B_{Ω} there is an exact correspondence $|T_{\Omega/B_{\Omega},s}| < \aleph_0$ iff $|T_{\Omega,s}| < \aleph_0$. The tricky case is when B_{Ω} contains unit axioms, since they may break this happy correspondence. For example, consider the unsorted signature $\Omega = (0, _ + _)$ where 0 is a unit element for $_ + _$. For the ACCU case, two simple checks apply in most cases (Appendix A, Lemmas 11, 12, and 13). Failing that, the classification of sort s is returned to the user as a proof obligation. Note that, thanks to the results in Appendix B of [30], any remaining proof obligations for checking $|T_{\Omega/B_{\Omega},s}| < \aleph_0$ when $B_{\Omega} = ACCU$ are always *decidable*, but additional algorithms beyond those presented here are needed to discharge such remaining proof obligations. Extending the current algorithms and their meta-level implementation to check finiteness of sorts in all cases when $B_{\Omega} = ACCU$ is left for future research.

Finite Sort Representative Generation. Here we require a method to do two things: (1) when $|T_{\Omega/B_{\Omega},s}| < \aleph_0$, we can compute each $[t]_{B_{\Omega}} \in T_{\Omega/B_{\Omega},s}$ (2) for each such $[t]_{B_{\Omega}}$, we can compute a unique representative $rep([t]_{B_{\Omega}})$. We first show how to generate $T_{\Omega,s}$. Recall that any order-sorted signature Ω can be viewed as a tree automaton such that the tree automaton accepts a term t in final state s iff $t \in T_{\Omega,s}$. Note also that tree automata are very simple *ground* rewrite theories. Let R_P be the ground rewrite rules for Ω 's tree automaton over $T_{\Omega \cup S}$, so that $t \in T_{\Omega,s}$ iff $t \rightarrow_{R_P}^+ s$. Let $R_G = R_P^{-1}$ then $T_{\Omega,s} = \{t \in T_{\Omega} \mid s \rightarrow_{R_G}^! t\}$ (Appendix A, Corollary 5). Furthermore, if $|T_{\Omega,s}| < \aleph_0$ and Ω has no empty sorts, this process will always terminate. Note that we can apply the rules R_G modulo B_{Ω} . Then the set $Rep(T_{\Omega/B_{\Omega},s}) = \{rep([t]) \mid [t] \in T_{\Omega/B_{\Omega},s}\}$ is exactly the set $Rep(T_{\Omega/B_{\Omega},s}) = \{t \mid s \rightarrow_{R_G, B_{\Omega}}^! t\}$.

4.2 Constructor Variants and Constructor Unifiers

We first show how to compute a set of *most general* constructor variants of a term t (i.e. a set of constructor variants $\llbracket t \rrbracket_{R,B}^\Omega$ such that for any constructor variant (t', θ) , we have $\exists (t'', \psi) \in \llbracket t \rrbracket_{R,B}^\Omega [(t'', \psi) \cong_{R,B} (t', \phi)]$) and then show how to use this method to compute a set of *most general* constructor unifiers $mgu_R^\Omega(\phi)$. Recall that a constructor variant is just an variant (t, θ) such that $t \in T_\Omega(X)$. Thus, $\llbracket t \rrbracket_{R,B}^\Omega$ can be computed in two steps: (1) computing a set of most general variants $\llbracket t \rrbracket_{R,B}$ (2) for each most general variant (t', θ) , compute the set of its *most general* constructor instances, i.e. a set of instances $mgi_B(t') = \{t'\eta_1, \dots, t'\eta_n\}$ where for any other instance $t'\alpha$, there exists a substitution γ and η_i with $\alpha =_B \eta_i \gamma$. Note that (1) can be solved via folding variant narrowing, so we tackle (2) by a reduction to a B -unification problem via a signature transformation $\Sigma \mapsto \Sigma^c$. In this transformed signature, the instances $mgi_B(t')$ correspond exactly to the solutions of a *single* B -unification problem.

The signature transformation $\Sigma \mapsto \Sigma^c$ splits into two steps: (i) we extend the sort poset $(S, <)$ of Σ and Ω and (ii) likewise extend the operator sets F and F_Ω , as specified by the definitions below, respectively. Recall we assume Σ (and thus Ω) are finite; otherwise these transformations would not be effective.

Definition 6. A constructor sort refinement of $(S, <)$ is defined by the following: (a) a set $S^c = S \uplus S^\downarrow$ with $c : S \rightarrow S^\downarrow$ a bijection, (b) a relation $(<^c)$ the smallest strict order where: (i) $\forall s, s' \in S [s < s' \Leftrightarrow [s <^c s' \wedge c(s) <^c c(s')]]$ and (ii) $\forall s \in S [c(s) <^c s]$, and (c) functions $(\bullet) : S^c \rightarrow S$ and $(\bullet) : S^c \rightarrow S^\downarrow$ defined by $s^\bullet = s$ if $s \in S$ else $c^{-1}(s)$ and $s_\bullet = s$ if $s \in S^\downarrow$ else $c(s)$.

We let $(<^c)$ also ambiguously denote its extension to strings of sorts $(S^c)^*$. Also, note that $(<) \subseteq (<^c)$ by definition and functions (\bullet) and (\bullet) have unique homomorphic extensions to free monoid homomorphisms denoted by: $(\bullet) : (S^c)^* \rightarrow S^*$ and $(\bullet) : (S^c)^* \rightarrow (S^\downarrow)^*$. Likewise, (\bullet) and (\bullet) have unique extensions to powersets, $(\bullet) : \mathcal{P}(S^c) \rightarrow \mathcal{P}(S)$ and $(\bullet) : \mathcal{P}(S^c) \rightarrow \mathcal{P}(S^\downarrow)$. Lastly, $(\bullet)_{(S^\downarrow)^*}$ and $(\bullet)_{S^*}$ are bijective by definition and lift into poset and powerset isomorphisms.

Definition 7. Given $\Sigma = ((S, <), F)$ and $\Omega = ((S, <), F_\Omega)$ where $\Omega \subseteq \Sigma$ and $(S^c, <^c, (\bullet), (\bullet))$ is a constructor sort refinement of $(S, <)$, we define:

1. $\Sigma^+ = ((S^c, <^c), F)$ and $\Omega^+ = ((S^c, <^c), F_\Omega)$
2. $\Sigma^\downarrow = ((S^c, <^c), F^\downarrow)$ and $\Omega^\downarrow = ((S^c, <^c), F_\Omega^\downarrow)$
3. $\Sigma^c = ((S^c, <^c), F^c)$ and $\Omega^c = ((S^c, <^c), F_\Omega^c)$
4. $\Omega_\bullet^\downarrow = ((S^\downarrow, <^c|_{S^\downarrow}), F_\Omega^\downarrow)$

where $F^\downarrow = (F/F_\Omega) \uplus F_\Omega^\downarrow$, $F_\Omega^\downarrow = \{f : w_\bullet \rightarrow s_\bullet \mid f : w \rightarrow s \in F_\Omega\}$, $F^c = F \uplus F_\Omega^\downarrow$, and $F_\Omega^c = F_\Omega \uplus F_\Omega^\downarrow$. Similarly, we also define $X^\downarrow = \{X_s\}_{s \in S^\downarrow}$. Then $X^c = X \uplus X^\downarrow$.

We can summarize the definition above with the figure below:

$$\begin{array}{ccccccccc}
\Sigma(X) & \hookrightarrow & \Sigma^+(X^c) & \hookrightarrow & \Sigma^c(X^c) & \hookleftarrow & \Sigma^\downarrow(X^c) & \hookleftarrow & \Sigma^\downarrow(X^\downarrow) \\
\uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow \\
\Omega(X) & \hookrightarrow & \Omega^+(X^c) & \hookrightarrow & \Omega^c(X^c) & \hookleftarrow & \Omega^\downarrow(X^c) & \hookleftarrow & \Omega^\downarrow(X^\downarrow)
\end{array}$$

where each arrow is a signature inclusion. The signature decorations are intended to be suggestive of the transformation: Σ^+ extends the subsort relation; Σ^c copies each constructor; Σ^\downarrow shifts constructors below; and finally Ω^\downarrow shifts constructors below and discards sorts S by applying (\bullet) . In this section, we will primarily consider $\Sigma^c(X^c)$ and $\Omega^\downarrow(X^c)$ which we refer to as the *constructor sort refinements* of Σ and Ω . The other signatures will be referenced as needed.

Note that (\bullet) and (\bullet) naturally extend into signature morphisms. The sort mapping is either (\bullet) or (\bullet) . If $t \in T_{\Sigma^c}(X^c)$, then the term mapping is given by: (a) if $t = x : s \in X^c$, then $(x : s)^\bullet = x : (s^\bullet)$ and $(x : s)_\bullet = x : (s_\bullet)$, (b) if $t = a : \rightarrow s \in F^c$, then $a^\bullet = a_\bullet = a$ (c) if $t = f(t_1, \dots, t_n)$, then $t^\bullet = f(t_1^\bullet, \dots, t_n^\bullet)$ and $t_\bullet = f(t_{1_\bullet}, \dots, t_{n_\bullet})$. The term mappings (\bullet) and (\bullet) also naturally extend to substitutions $\theta \in [X^c \rightarrow T_{\Sigma^c}(X^c)]$. Then for each $(x, t) \in \theta$, we have $(x^\bullet, t^\bullet) \in \theta^\bullet$ and $(x_\bullet, t_\bullet) \in \theta_\bullet$. In particular, we note three facts: (i) $(\bullet) : \Omega(X) \rightarrow \Omega^\downarrow(X^\downarrow)$ is a signature isomorphism with inverse (\bullet) (ii) $(\bullet) : \Sigma^c(X^c) \rightarrow \Sigma(X)$ is a signature morphism (iii) as sets of terms, $T_{\Omega^\downarrow}(X^\downarrow) = T_{\Omega^\downarrow}(X^\downarrow)$ and $T_\Omega = T_{\Omega^\downarrow} = T_{\Omega^\downarrow}$.

Our first goal in this subsection is to show that term sorting, sensibility, and prerogularity are all preserved by constructor sort refinement, i.e., refinement in the sense that all existing sort information is preserved and only new sort information is added. Note that we trivially have preservation of term sorts by facts (i)-(iii) above since $\forall s \in S^c \forall t \in T_{\Sigma^c}(X^c)_s [t^\bullet \in T_\Sigma(X)_{s^\bullet} \wedge s \leq^c s^\bullet]$, (\bullet) specializes to the identity when $t \in T_\Sigma(X)$, and $\forall s \in S [t \in T_{\Omega^\downarrow, s_\bullet} \Leftrightarrow t \in T_{\Omega, s}]$. Thus, it is enough to prove preservation of sensibility and prerogularity. However, the example below shows our current assumptions are not strong enough.

Example 2. Consider sort poset $(S, <) = (\{a, b\}, \{(a, b)\})$ and signatures $\Sigma = ((S, <), \{f : a \rightarrow a, f : b \rightarrow b\})$ and $\Omega = ((S, <), \{f : b \rightarrow b\})$. The ctor sort refinement $(S^c, <^c) = (S \uplus \{a_\bullet, b_\bullet\}, (<) \uplus \{(a_\bullet, a), (b_\bullet, b), (a_\bullet, b_\bullet), (a_\bullet, b)\})$ where $\Sigma^c = ((S^c, <^c), \{f : a \rightarrow a, f : b \rightarrow b, f : b_\bullet \rightarrow b_\bullet\})$ violates prerogularity for sort a_\bullet where $(a_\bullet \leq^c a \wedge a_\bullet \leq^c b_\bullet)$ but $(a \not\leq^c b_\bullet \wedge b_\bullet \not\leq^c a)$ even though Σ and Ω are both prerogular by construction.

Note in the previous example the violation occurred when a constructor had a subsort-overloaded defined operator below. However, just restricting subsort-overloading does not fix the problem.

Example 3. Let $(S, <) = (\{a, b, c\}, \{(a, b), (a, c)\})$, $\Sigma = ((S, <), \{f : b \rightarrow a, f : c \rightarrow c\})$, and $\Omega = ((S, <), \{f : c \rightarrow c\})$. Then $(S^c, <^c) = (S \uplus \{a_\bullet, b_\bullet\}, (<) \uplus \{(a_\bullet, a), (b_\bullet, b), (c_\bullet, c), (a_\bullet, b_\bullet), (a_\bullet, c_\bullet), (a_\bullet, b), (a_\bullet, c)\})$. But now note $\Sigma^c = ((S^c, <^c), \{f : b \rightarrow a, f : c \rightarrow c, f : c_\bullet \rightarrow c_\bullet\})$ violates prerogularity for sort a_\bullet where $(a_\bullet \leq^c b \wedge a_\bullet \leq^c c_\bullet)$ holds but $(a \not\leq^c c_\bullet \wedge c_\bullet \not\leq^c a)$.

Essentially, the invariant violated by both examples was Ω was not *prerogular below* Σ , in the sense that, given a symbol and arity with a constructor typing,

it's minimal typing was not a constructor. In order to formally specify this invariant, we will need some auxiliary notation.

Let $\Sigma = ((S, <), F)$ be an arbitrary signature and (P, \triangleleft) an arbitrary poset. Let $ty_\Sigma : T_\Sigma \rightarrow F$ be defined by the two equations $ty_\Sigma(c) = \{c : \rightarrow s \in F\}$ and $ty_\Sigma(f(t_1, \dots, t_n)) = \{f : s_1 \cdots s_n \rightarrow s \in F \mid t_i \in T_{\Sigma_{s_i}}\}$. Also let ty_Σ denote the function $ty_\Sigma(f, w) = \{f : w' \rightarrow s \in F \mid w \leq w'\}$. Further let $min_{\triangleleft} : \mathcal{P}(P) \rightarrow P \uplus \{\emptyset\}$ be $min_{\triangleleft}(I) = \bigwedge I$ if $(\exists \bigwedge I) \wedge \bigwedge I \in I$ else \emptyset where $\bigwedge I$ denotes the greatest lower bound of I in (P, \triangleleft) if it exists.

Definition 8. Let $\Sigma = ((S, <), F)$ have subsignature $\Omega = ((S, <), F_\Omega)$. Then Ω is prerregular below Σ (written $\Omega < \Sigma$) iff Ω and Σ are prerregular and for any f we have $\forall w \in S^*[ty_\Omega(f, w) \neq \emptyset \Rightarrow min_{<}(ty_\Sigma(f, w)) \in ty_\Omega(f, w)]$ where $(F, <)$ is the poset where $f : w \rightarrow s < g : w' \rightarrow s' \Leftrightarrow s < s'$.

We now prove constructor sort refinements $\Omega^\downarrow(X^c)$ and $\Sigma^c(X^c)$ preserve sensibility and prerregularity iff Ω and Σ are sensible and $\Omega < \Sigma$. Note, by definition, for any signature Σ , we have $ls_\Sigma(t) = min_{<}(ty_\Sigma(t))$ for the poset $(F, <)$ and to prove Σ is prerregular it is enough to show $\forall t \in T_\Sigma[ls_\Sigma(t) \neq \emptyset]$. To complete the proof, we will need four lemmas. To preserve the logical flow of the argument, we state them here as assumptions to be used in the main argument and give a detailed proof of each of them in Appendix B.

Lemma 1. $\forall t \in T_\Sigma [t \in T_\Omega \Rightarrow ls_\Omega(t) = ls_\Sigma(t)]$

Lemma 2. $\forall t \in T_{\Omega^\downarrow(X^c)}/X^c [ty_{\Omega^\downarrow(X^c)}(t) = ty_{\Omega^\downarrow(X^\downarrow)}(t) = ty_{\Omega^\downarrow(X^\downarrow)}(t)]$

Lemma 3. $\forall t \in T_{\Sigma^c(X^c)}/X^c [ty_{\Sigma^c(X^c)}(t) = ty_{\Sigma(X)}(t^\bullet)]$

Lemma 4. $\Sigma = ((S, <), F)$ is sensible iff $\widehat{\Sigma} = ((\widehat{S}, \emptyset), \widehat{F})$ is sensible where $f : [s_1] \cdots [s_n] \rightarrow [s_0] \in \widehat{F}$ iff $\exists f : s'_1 \cdots s'_n \rightarrow s'_0 \in F$ with $s'_i \in [s_i]$ for $0 \leq i \leq n$.

Theorem 6. If $\Omega < \Sigma$ and Ω and Σ are sensible, then the constructor sort refinements $\Sigma^c(X^c) = ((S^c, <^c), F^c \uplus X^c)$ and $\Omega^\downarrow(X^c) = ((S^c, <^c), F_\Omega^\downarrow \uplus X^c)$ are both sensible and prerregular.

Proof.

Note proving Σ^c is sensible implies $\Sigma^c(X^c)$ is sensible which implies $\Omega^\downarrow(X^c)$ is sensible. Then note that $\widehat{\Sigma^c} \cong \widehat{\Sigma}$ and signature isomorphism preserves sensibility, and finally apply Lemma 4.

We now prove that $\Omega^\downarrow(X^c)$ is prerregular. By abuse of language, let X also denote the signature $((S, <), X)$. Then note $\forall t \in T_{\Omega^\downarrow(X^c)}[ty_{\Omega^\downarrow(X^c)}(t) = ty_{\Omega^\downarrow(X^\downarrow)}(t) \uplus ty_X(t)]$ and $\Omega^\downarrow(X^\downarrow) \cap X = \emptyset$. Thus, by Lemma 2, we obtain that $\forall t \in T_{\Omega^\downarrow(X^c)}/X^c [ty_{\Omega^\downarrow(X^c)}(t) = ty_{\Omega^\downarrow(X^\downarrow)}(t)]$. Thanks to the facts above, $ls_{\Omega^\downarrow(X^c)} = ls_{\Omega^\downarrow(X^\downarrow)} \uplus ls_X$. By signature isomorphism $\Omega^\downarrow(X^\downarrow) \cong \Omega(X)$, this is equivalent to $ls_{\Omega^\downarrow(X^c)} = (\bullet); ls_{\Omega(X)}; (\bullet) \uplus ls_X$ where semicolon denotes in-order function composition. Since X is prerregular by definition and $\Omega(X)$ by assumption, $ls_{\Omega^\downarrow(X^c)}$ satisfies $\forall t \in T_{\Omega^\downarrow(X^c)}[ls_{\Omega^\downarrow(X^c)}(t) \neq \emptyset]$, as required.

We now prove $\Sigma^c(X^c)$ is preregular. First let $t \in X^c$. Then $t \in X \uplus X^\downarrow$. If $t = x : s \in X$ then $ls_{\Sigma^c(X^c)}(x : s) = ls_{\Sigma(X)}(x : s^\bullet) = s$. Similarly, if $t = x : s \in X^\downarrow$, $ls_{\Sigma^c(X^c)}(x : s) = ls_{\Omega(X)}(x : s^\bullet)_\bullet = s$.

Now let $t \in T_{\Sigma^c(X^c)}/X^c$. Note $ty_{\Sigma^c(X^c)}(t) = ty_{\Omega^\downarrow(X^c)}(t) \uplus ty_{\Sigma^+(X^c)}(t)$, i.e., the type of non-variable t is from F_Ω^\downarrow or F and $ls_{\Sigma^c(X^c)}(t) = \min_{<}(ty_{\Omega^\downarrow(X^c)}(t) \uplus ty_{\Sigma^+(X^c)}(t))$. Suppose $t \in T_{\Omega^\downarrow(X^\downarrow)}/X^\downarrow$. By Lemma 2 and $\Omega^\downarrow(X^\downarrow) \cong \Omega(X)$, we obtain $ty_{\Omega^\downarrow(X^c)}(t) = ty_{\Omega^\downarrow(X^\downarrow)}(t) = ty_{\Omega(X)}(t^\bullet)_\bullet$. By Lemmas 1 and 3, we have $\min_{<}(ty_{\Omega(X)}(t^\bullet)) = \min_{<}(ty_{\Sigma(X)}(t^\bullet)) = \min_{<}(ty_{\Sigma^+(X^c)}(t))$. Then note $ls_{\Sigma^c(X^c)}(t) = \min_{<}(ty_{\Omega(X)}(t^\bullet)_\bullet \uplus ty_{\Sigma(X)}(t^\bullet)) = ls_{\Omega(X)}(t^\bullet)_\bullet$. Finally, assume that $t \in T_{\Sigma^c(X^c)}/T_{\Omega^\downarrow(X^c)}$. Then we obtain $ty_{\Omega^\downarrow(X^c)}(t) = \emptyset$ and $ls_{\Sigma^c(X^c)}(t) = \min_{<}(ty_{\Sigma^+(X^c)}(t)) = \min_{<}(ty_{\Sigma(X)}(t^\bullet)) = ls_{\Sigma(X)}(t^\bullet)$ by Lemma 3. Thus, we have $\forall t \in T_{\Sigma^c(X^c)}[ls_{\Sigma^c(X^c)}(t) \neq \emptyset]$, as required. \square

Corollary 2. *The functions $ls_{\Omega^\downarrow(X^c)}$ and $ls_{\Sigma^c(X^c)}$ are defined by:*

- (a) $\forall t \in T_{\Omega^c(X^c)} \quad ls_{\Omega^\downarrow(X^c)}(t) = ls_{\Omega(X)}(t^\bullet)_\bullet$ if $t \in T_{\Omega^\downarrow(X^\downarrow)}$ else $ls_{\Sigma(X)}(t^\bullet)$
- (b) $\forall t \in T_{\Sigma^c(X^c)} \quad ls_{\Sigma^c(X^c)}(t) = ls_{\Omega(X)}(t^\bullet)_\bullet$ if $t \in T_{\Omega^\downarrow(X^\downarrow)}$ else $ls_{\Sigma(X)}(t^\bullet)$

We now extend the result above to show that B -preregularity is preserved under a weak assumption that is often satisfied in practice. We first state the required condition and then give the proof.

Definition 9. *Let B be a set of axioms, $t = t' \in B$ with $\text{vars}(t = t') = Y$ and $\alpha \in [Y \rightarrow X]$. We say B respects constructors iff $t\alpha \in T_\Omega(X) \Leftrightarrow t'\alpha \in T_\Omega(X)$.*

Theorem 7. *Assume $\Sigma(X)$ and $\Omega(X)$ are sensible and B -preregular and that $\Omega(X) < \Sigma(X)$ and B respects constructors. Then their respective constructor sort refinements $\Sigma^c(X^c)$ and $\Omega^\downarrow(X^c)$ are also B -preregular.*

Proof. We apply Theorem 6 to immediately show that $\Sigma^c(X^c)$ and $\Omega^\downarrow(X^c)$ are sensible and preregular. We first prove $\Sigma^c(X^c)$ is B -preregular. Thus, let $t = t' \in B$ with $Y = \text{vars}(t) = \text{vars}(t')$ and $\alpha \in [Y \rightarrow X^c]$. Note the value of functions $ls_{\Omega^\downarrow(X^c)}$ and $ls_{\Sigma^c(X^c)}$ is completely determined by the input term t and functions $ls_{\Omega(X)}$ and $ls_{\Sigma(X)}$. In particular, if $ls_{\Sigma(X)}(t\alpha) = ls_{\Sigma(X)}(t'\alpha)$, then $ls_{\Sigma^c(X^c)}(t\alpha) = ls_{\Sigma^c(X^c)}(t'\alpha)$ iff $t\alpha \in T_{\Omega^\downarrow(X^\downarrow)} \Leftrightarrow t'\alpha \in T_{\Omega^\downarrow(X^\downarrow)}$ by Corollary 2 (the same holds true for $ls_{\Omega^\downarrow(X^c)}$). Since B protects constructors, it is enough to show $\forall t \in T_\Omega(X)[t\alpha \in T_{\Omega^\downarrow(X^\downarrow)} \Leftrightarrow \alpha \in [Y \rightarrow T_{\Omega^\downarrow(X^\downarrow)}]]$ where $Y = \text{vars}(t) \neq \emptyset$. The base case where $t = x : s$ is trivial, so assume $t = f(t_1, \dots, t_n)$. Then $t\alpha = f(t_1\alpha, \dots, t_n\alpha)$ and $t_i\alpha \in T_{\Omega^\downarrow(X^\downarrow)} \Leftrightarrow \alpha \in [Y \rightarrow T_{\Omega^\downarrow(X^\downarrow)}]$ for $1 \leq i \leq n$ by induction hypothesis. But then $f : s_1 \cdot \dots \cdot s_n \rightarrow s \in F_\Omega$ with $t_i \in T_\Omega(X)_{s_i}$ iff $f : s_{1\bullet} \cdot \dots \cdot s_{n\bullet} \rightarrow s_\bullet \in F_\Omega^\downarrow$ and $f(t_1\alpha, \dots, t_n\alpha) = t\alpha \in T_{\Omega^\downarrow(X^\downarrow)}$, as required. \square

The following corollary lifts the result above to decompositions.

Corollary 3. *Let $R = (\Sigma, B, R)$ be convergent with constructor decomposition $R_\Omega = (\Omega, B_\Omega, R_\Omega)$ and $\Omega < \Sigma$. Then Σ^c and Ω^\downarrow are sensible and B -preregular.*

Proof. Note that protecting a constructor decomposition implies B respects constructors (see Def. 2). Then apply Theorem 7.

We have now shown that our construction, under mild conditions, preserves sensibility and B -preregularity. Thus, B -unification will be well-defined in our new signature. We now move to prove the main theorem of this section which shows how most general constructor instances of a term modulo B may be obtained by a single unification problem in $\Sigma^c(X^c)$. We first collect a number of essential facts which relate $T_\Omega(X)$ to $T_{\Omega^\downarrow}(X^\downarrow)$ we will use in the proof.

Lemma 5. *Suppose that $\alpha, \beta \in [X \rightarrow T_\Omega(X)]$, $\alpha', \beta' \in [X^\downarrow \rightarrow T_{\Omega^\downarrow}(X^\downarrow)]$, and $\theta, \gamma \in [X^c \rightarrow T_{\Sigma^c}(X^c)]$. Let $id^\downarrow \in [X^c \rightarrow X^\downarrow]$ where $id^\downarrow(x : s) = x : s_\bullet$. Then:*

- (a) $(\alpha_\bullet)^\bullet = \alpha \wedge (\alpha^\bullet)_\bullet = \alpha$
- (b) $\forall t, t' \in T_\Omega(X)[t =_B t' \Leftrightarrow t_\bullet =_B t'_\bullet] \wedge \forall t, t' \in T_{\Omega^\downarrow}(X^\downarrow)[t =_B t' \Leftrightarrow t^\bullet =_B t'^\bullet]$
- (c) $[\alpha =_B \beta \Leftrightarrow \alpha_\bullet =_B \beta_\bullet] \wedge [\alpha' =_B \beta' \Leftrightarrow \alpha'^\bullet =_B \beta'^\bullet]$
- (d) $\forall t \in T_{\Sigma^c}(X^c)[t_\bullet = t(id^\downarrow)] \wedge (id^\downarrow)^\bullet = id$
- (e) $\forall t \in T_{\Sigma^c}(X^c)[(t\theta)_\bullet = t_\bullet(\theta_\bullet) \wedge (t\theta)^\bullet = t^\bullet(\theta^\bullet) \wedge (\theta\gamma)_\bullet = \theta_\bullet(\gamma_\bullet) \wedge (\theta\gamma)^\bullet = \theta^\bullet(\gamma^\bullet)]$

Proof. Both (a) and (b) follow immediately since $T_{\Omega^\downarrow}(X^\downarrow) = T_{\Omega^\downarrow}(X^\downarrow)$ and by isomorphism $(\bullet) : \Omega^\downarrow(X^\downarrow) \rightarrow \Omega(X)$. Then (c) is an immediate application of (b). Finally, (d) and (e) are easy structural induction proofs. \square

We now give a precise construction of $mgci_B^\Omega$ using B -unification in $\Sigma^c(X^c)$.

Theorem 8. *Suppose $\Sigma(X)$ and $\Omega(X)$ are sensible and B -preregular, $\Omega < \Sigma$, and B respects constructors. Then (a) $\forall t \in T_\Sigma(X)_s \forall t' \in T_\Omega(X)_{s'}$ with $s \equiv_{<} s'$ and $x \notin \text{vars}(t)$, $t\alpha =_B t'$ iff there are $\eta \in mgu_B(t = x : c(s'))$ and θ such that $\eta^\bullet\theta|_{\text{vars}(t)} =_B \alpha$ where $\alpha \in [\text{vars}(t) \rightarrow T_\Omega(X)]$ and $\theta \in [X \rightarrow T_\Omega(X)]$ and (b) the set of most general constructor instances of t modulo B is defined by $mgci_B^\Omega(t) = \{t(\eta^\bullet) \mid \eta \in mgu_B(t = x : ls_{\Sigma(X)}(t)_\bullet)\}$.*

Proof. We first prove (a). Let $\beta = \alpha_\bullet \uplus \{(x : s'_\bullet, t'_\bullet)\}$. Then observe:

$$\begin{aligned} t\alpha =_B t' &\Leftrightarrow (t\alpha)_\bullet =_B t'_\bullet \\ &\Leftrightarrow t_\bullet(\alpha_\bullet) =_B t'_\bullet \\ &\Leftrightarrow t_\bullet\beta =_B x : s'_\bullet\beta \\ &\Leftrightarrow \exists \eta' \in mgu_B(t_\bullet = x : s'_\bullet) \exists \theta' \in [X^\downarrow \rightarrow T_{\Omega^\downarrow}(X^\downarrow)] [\eta'\theta' =_B \beta] \end{aligned}$$

which follow by Lemma 5 and the fact B respects constructors so $t\alpha \in T_\Omega(X)$. Let id be the identity substitution and note $x : (s'_\bullet)_\bullet = x : s'_\bullet$. Then we obtain:

$$\begin{array}{ll}
 \eta' \in \text{mgu}_B(t_\bullet = x : s'_\bullet) & \eta' \theta' =_B \beta \\
 \Leftrightarrow \eta' \in \text{mgu}_B(t_\bullet = x : (s'_\bullet)_\bullet) & \Leftrightarrow (\eta' \theta')^\bullet =_B \beta^\bullet \\
 \Leftrightarrow \eta' \in \text{mgu}_B(t(id^\downarrow) = x : s'_\bullet(id^\downarrow)) & \Leftrightarrow \eta'^\bullet(\theta'^\bullet) =_B \beta^\bullet \\
 \Leftrightarrow id^\downarrow \eta' \in \text{mgu}_B(t = x : s'_\bullet) & \Leftrightarrow \eta'^\bullet(\theta'^\bullet)|_{\text{vars}(t)} =_B \alpha \wedge \\
 & \eta'^\bullet(\theta'^\bullet)(x) =_B t'
 \end{array}$$

by Lemma 5. Now let $\eta = id^\downarrow \eta'$ and $\theta = \theta'^\bullet$. Then we can derive equalities $\eta^\bullet \theta = (id^\downarrow \eta')^\bullet \theta = (id^\downarrow)^\bullet (\eta'^\bullet) \theta = id(\eta'^\bullet) \theta = \eta'^\bullet(\theta'^\bullet)$ as required. Finally (b) is an immediate application of (a). \square

In case the constructor decomposition has no rules (constructors are free modulo B_Ω), Theorem 8 provides a method to compute constructor variants.

Corollary 4. *Let (Σ, B, R) be convergent and protect constructor decomposition $(\Omega, B_\Omega, \emptyset)$ and $\Omega < \Sigma$. The most general constructor variants of $t \in T_\Sigma(X)$ are $\llbracket t \rrbracket_{R,B}^\Omega = \{(t'(\eta^\bullet), \theta \eta^\bullet) \mid (t', \theta) \in \llbracket t \rrbracket_{R,B} \wedge \eta \in \text{mgu}_B(t', x : \text{ls}_{\Sigma(X)}(t'))\}$.*

Proof. Apply Corollary 3. It is sufficient to prove: (a) each $(t' \eta \theta \eta) \in \llbracket t \rrbracket_{R,B}^\Omega$ is a constructor variant (b) for any constructor variant (t'', ψ) , we obtain that $\exists (t' \eta, \theta \eta) \in \llbracket t \rrbracket_{R,B}^\Omega [(t' \eta, \theta \eta) \supseteq_{R,B} (t'', \psi)]$. To see (a), suppose $(t', \theta) \in \llbracket t \rrbracket_{R,B}$. By definition of most general unifier and Theorem 8, $\text{mgu}_B(t', x : c(\text{ls}_{\Sigma(X)}(t')))$ is the set of most general substitutions η modulo B such that $t' \eta^\bullet \in T_{\Omega(X)}$. Since (Σ, B, R) protects $(\Omega, B_\Omega, \emptyset)$ and Ω is a signature of free constructors modulo B , we obtain $t' \eta^\bullet!_{R,B} = t' \eta^\bullet$, and $(t' \eta^\bullet, \theta \eta^\bullet)$ is a constructor variant. To see (b), note, by definition, $\llbracket t \rrbracket$ covers every variant, and $\text{mgi}_B^\Omega(t')$ covers every constructor instance, as required. \square

The reduction of constructor unifiers to constructor variants is simple. Recall any unification problem ϕ is a Σ^\wedge -term $\phi \in T_{\Sigma^\wedge(X)} \text{Conj}$. Let $\{\alpha_i\}_{i \in I}$ denote the finite set of most general R, B -variant unifiers of ϕ obtained as explained in Theorem 3. Then the set of most general constructor unifiers of ϕ is the set $\{\alpha_i \eta^\bullet \mid \eta \in \text{mgu}_B((\phi \alpha_i)!_{R,B}, x : \text{Conj}_\bullet)\}$.

4.3 Constructor Variants and Unifiers: An Example

The notions of constructor variant and constructor unifier become more subtle when, due to order-sortedness, a same subsort-polymorphic operator f has some typings that are constructors and some other typings that are defined functions. The following examples illustrates the issues involved.

Example 4. (Integers with Addition). The FVP decomposition \mathcal{Z}_+ for integers with addition has sorts Nat , NzNat , NzNeg , and Int , and subsorts $\text{NzNat} < \text{Nat}$ and $\text{Nat NzNeg} < \text{Int}$, where NzNat (resp. NzNeg) denotes the non-zero naturals (resp. negatives). The constructor signature Ω has constants 0 of sort Nat and 1 of sort NzNat , and operators $- + - : \text{Nat Nat} \rightarrow \text{Nat}$, $- + - : \text{NzNat NzNat} \rightarrow$

$NzNat$, and $- : NzNat \rightarrow NzNeg$. The only defined function symbol is: $_+ : Int\ Int \rightarrow Int$, also ACU . The rewrite rules R defining $+$ and making (Ω, ACU, \emptyset) an ACU -free constructor decomposition of \mathcal{Z}_+ are the following (with i a variable of sort Int , and n, m variables of sort $NzNat$): $i + n + -(n) \rightarrow i$, $i + -(n) + -(m) \rightarrow i + -(n + m)$, $i + n + -(n + m) \rightarrow i + -(m)$, and $i + n + m + -(n) \rightarrow i + m$.

Note \mathcal{Z}_+ is FVP and protects its constructor subtheory, so we can already compute variants as usual. To compute constructor variants/unifiers in \mathcal{Z}_+ , we generate its refinement according to Definition 7. Figure 1 below illustrates how this is done, where for each sort s , we let s_\bullet denote its lowered sort.

Consider now the term $x + y$ with x, y variables of sort Int . By folding variant narrowing, it is easy to show $x + y$ has twelve variants in general, but to simplify the example, we focus on its most simple variant, i.e. $u = (x + y, id)$ with id the identity substitution. Note that u is *not* a constructor variant in \mathcal{Z}_+ , and there are variants that are *less general* than $(x + y, id)$ and are constructor variants. The *most general* constructor variants that are less general than $(x + y, id)$ are: (i) $(x, \{y \mapsto 0\})$, (ii) $(y, \{x \mapsto 0\})$, and (iii) $(x' + y', \{x \mapsto x' : Nat, y \mapsto y' : Nat\})$. In Section 5, we show that these constructor variants are all generated by our Maude implementation. Likewise, let ϕ be the equation $z = x + y$, with x, y, z of sort Int . Then $\{z \mapsto x + y\}$ is a trivial \mathcal{Z}_+ -unifier of ϕ , but *not* a constructor unifier. A complete set $mgu_{\mathcal{R}}^{\Omega}(\phi)$ of most general constructor \mathcal{Z}_+ -unifiers of ϕ is given by the unifiers: (i) $\{z \mapsto x, y \mapsto 0\}$, (ii) $\{z \mapsto y, x \mapsto 0\}$, and (iii) $\{z \mapsto x' + y', x \mapsto x' : Nat, y \mapsto y' : Nat\}$. Similarly, these can be shown to be the solution to the corresponding unification problem in the signature $(\Sigma^\wedge)^c$ as described in the end of Section 4.2.

For other examples of constructor variants and constructor unifiers we refer the reader to Examples 3–4 in [29].




Fig. 1. INT signature Σ and its refinement Σ^c

4.4 Optimizing Constructor Variant and Unifier Generation

As the example in Section 4.3 illustrates, the subtleties involved in generating constructor variants and constructor unifiers are all related to the fact that we can have a subsort-overloaded operator f (+ in the above example), which is a *constructor* for some typings and a *defined symbol* for other typings. But this of course is impossible in a *many-sorted* signature, because there all subsort-overloaded typings must coincide. Also, in many order-sorted examples, the distinction between constructors and defined symbols applies to *entire families* of subsort-overloaded symbols: either all typings in the family are constructors, or they are all defined symbols.

This suggests an obvious optimization: before solving a constructor variant or a constructor unifier problem, check first whether any subsort-overloaded function symbol has both constructor and defined symbol typings. If it does, apply the general algorithms developed in Section 4.2. If it does not (and an additional property on $B - B_\Omega$ holds, see below) apply instead the considerably simpler and more efficient algorithms described below.

First of all, the above check can be made more precise as follows. Recall the notation used in Section 2 for a subsort-overloaded family of operators f :

$$f_{[s]}^{[s_1] \dots [s_n]} = \{f : s'_1 \dots s'_n \rightarrow s' \in \Sigma \mid s'_i \in [s_i], 1 \leq i \leq n, s' \in [s]\}$$

for connected components $[s_1], \dots, [s_n], [s] \in \widehat{S}$. Given a constructor subsignature $\Omega \subseteq \Sigma$, we have *two* different subsort-overloaded families for f : (i) we can further qualify the above family by $f_{[s], \Sigma}^{[s_1] \dots [s_n]}$, and (ii) we can denote instead by $f_{[s], \Omega}^{[s_1] \dots [s_n]}$ the subset $f_{[s], \Omega}^{[s_1] \dots [s_n]} = \{f : s'_1 \dots s'_n \rightarrow s' \in \Omega \mid s'_i \in [s_i], 1 \leq i \leq n, s' \in [s]\}$. The above-mentioned check that there is no subsort-overloaded symbol with *both* constructor and defined symbol typings can now be easily expressed. We just need to check that for each *non-empty* family $f_{[s], \Omega}^{[s_1] \dots [s_n]}$ the following set-theoretic equality holds:

$$f_{[s], \Sigma}^{[s_1] \dots [s_n]} - f_{[s], \Omega}^{[s_1] \dots [s_n]} = \emptyset.$$

This check has an easy implementation as a meta-level function in Maude. In their metalevel representations the operators in Σ (resp. Ω) are represented as *sets* with associative-commutative union. We just need to extract for each f in Ω its subsort-overloaded family $f_{[s], \Omega}^{[s_1] \dots [s_n]}$ and likewise compute in Σ the family $f_{[s], \Sigma}^{[s_1] \dots [s_n]}$. The check is then a simple emptiness check for the corresponding set difference.

As mentioned above, an additional property must be checked. If B , resp. B_Ω , denotes the equational axioms of Σ , resp. Ω , then we must also check that for each $u = v \in B - B_\Omega$ we have $u \notin T_\Omega(X)$ and $v \notin T_\Omega(X)$. This ensures that if $t \in T_\Sigma(X) - T_\Omega(X)$ and $t =_B t'$, then $t' \in T_\Sigma(X) - T_\Omega(X)$. For example, if all axioms in $B - B_\Omega$ are combinations of associativity and/or commutativity

axioms, the condition holds. But if $B - B_\Omega$ contains any identity axioms, the condition would fail.

The two algorithms for the case when no subsort-overloaded symbol is both a constructor and defined symbol, and where the above condition on $B - B_\Omega$ holds, are both very simple forms of *filtering* the corresponding variants (resp. variant unifiers):

1. $\llbracket t \rrbracket_{R,B}^\Omega = \{(u, \theta) \in \llbracket t \rrbracket_{R,B} \mid u \in T_\Omega(X)\}$.
2. Given a system of Σ -equations $\phi = u_1 = v_1 \wedge \dots \wedge u_n = v_n$, the set $mgu_{\mathcal{R}}^\Omega(\phi)$ of its most general constructor unifiers is defined by:

$$mgu_{\mathcal{R}}^\Omega(\phi) = \{\alpha \in mgu_{\mathcal{R}}(\phi) \mid (u_i \alpha)!_{R,B} \in T_\Omega(X), 1 \leq i \leq n\},$$

where $mgu_{\mathcal{R}}(\phi)$ denotes the set of most general variant unifiers of ϕ .

The Maude meta-level functions needed for the filtering performed in (1) and (2) are also very simple. In case (1), the check that $u \in T_\Omega(X)$ is a simple call to Maude's `wellFormed` predicate, that takes the meta-representations of a module and a term and returns `true` iff the term is well-formed in that module. In case (2), we can incrementally compute each variant unifier α using Maude 2.7.1 `metaVariantUnify` function; then we can compute each $(u_i \alpha)!_{R,B}$ by calling the `metaReduce` function; and we finally check $(u_i \alpha)!_{R,B} \in T_\Omega(X)$ for each $1 \leq i \leq n$ by calling the `wellFormed` predicate.

4.5 Descent Maps

There are two ways in which the methods presented in this paper may be insufficient to prove satisfiability of QF formulas in the initial algebra of an order-sorted equational theory having an FVP decomposition \mathcal{R} :

1. At the *theoretical* level, \mathcal{R} may lack an OS-compact constructor decomposition, so that the methods presented here cannot be applied to \mathcal{R} .
2. At the *practical* level, even if \mathcal{R} has an OS-compact constructor decomposition amenable to the methods and algorithms presented here, directly checking satisfiability of QF formulas in \mathcal{R} may be quite inefficient. This can happen because: (i) B -unification itself may generate a large number of unifiers; and (ii) there may also be a large number of variant R, B -unifiers of a given term t .

Faced with any of these theoretical and/or practical limitations, the following notion of a *descent map*, presented in [28, 30], may provide a way out of such limitations:

Definition 10. A descent map is a triple $(\mathcal{R}, \bullet, \mathcal{D})$ where \mathcal{R} and \mathcal{D} are decompositions of order-sorted equational theories, and \mathcal{R} conservatively extends \mathcal{D} , and where \bullet is a total computable function, $\varphi \mapsto \varphi^\bullet$, mapping each QF formula φ in the theory decomposed by \mathcal{R} into a corresponding QF formula φ^\bullet in the theory decomposed by \mathcal{D} and such that $C_{\mathcal{R}} \models \exists \varphi \Leftrightarrow C_{\mathcal{D}} \models \exists \varphi^\bullet$, where $\exists \varphi$ denotes the existential closure of φ .

Limitation (1) can be overcome when \mathcal{R} lacks an OS-compact constructor decomposition but \mathcal{D} has one. And limitation (2) can be overcome because solving satisfiability in $C_{\mathcal{D}}$ of the QF formula φ^\bullet may be considerably more efficient than solving satisfiability in $C_{\mathcal{R}}$ of the original formula φ . Since descent maps form a *category* and therefore can be *composed*, suitable compositions of such maps can greatly help in solving limitations (1) and (2). Furthermore, they can substantially extend the theoretical and practical reach of the variant satisfiability methods presented in this paper.

In experimenting with the current implementation of the variant-satisfiability algorithms described in Section 5 for solving SMT problems for various automated deduction applications, we have found descent maps to be quite helpful in overcoming type (1) and (2) limitations, specifically in the context of Presburger arithmetic,⁴ for the following reasons: (i) the simplest FVP specifications $\mathcal{Z}_{+,>,\geq}$, resp. $\mathcal{N}_{+,>,\geq}$, of Presburger arithmetic for the integers (resp. the naturals) fail to have OS-compact constructor decompositions [30]; (ii) solving satisfiability of QF formulas by variant satisfiability in the initial algebras of $\mathcal{Z}_{+,>,\geq}$, or even just in that of \mathcal{Z}_+ (the Abelian group of the integers) is quite inefficient due to a usually large number of variants modulo *ACU*; whereas (iii) solving satisfiability of QF formulas by variant satisfiability in the initial algebra of \mathcal{N}_+ (the Abelian monoid of the integers) is much more efficient, since, being free modulo *ACU* and OS-compact, it essentially reduces to computing *ACU* unifiers, which, although expensive for large terms, is efficiently supported by Maude 2.7.1.

In [30] three descent maps are defined: (i) $\mathcal{N}_{+,>,\geq} \xrightarrow{\text{lit2at } \delta_0} \mathcal{N}_+$, reducing natural Presburger arithmetic satisfiability to satisfiability in the Abelian monoid of the naturals; (ii) an entirely similar map $\mathcal{Z}_{+,>,\geq} \xrightarrow{\text{lit2at } \delta_0} \mathcal{Z}_+$, reducing integer Presburger arithmetic satisfiability to satisfiability in the Abelian group of the integers; and (iii) $\mathcal{Z}_+ \xrightarrow{v^-} \mathcal{N}_+$, reducing satisfiability in the Abelian group of the integers to satisfiability in the Abelian monoid of the naturals. These three maps ease limitations of type (1) and/or (2). Furthermore, it is shown in Theorems 14–15 of [30] that the descent maps (i)–(ii) can be modularly extended to FVP theory combinations where $\mathcal{N}_{+,>,\geq}$, resp. $\mathcal{Z}_{+,>,\geq}$, is a subspecification of a larger FVP theory. Although not explicitly treated in [30], descent map (iii) has a natural extension to a descent map (iv) $\mathcal{Z}_{+,>,\geq} \xrightarrow{v^-} \mathcal{N}_{+,>,\geq}$, so that we get the following diagram of descent maps:

⁴ This does not exclude the possibility of using not only descent maps, but also well-known *domain-specific* SMT solving algorithms for Presburger arithmetic. However, the applications we have experimented with, in which variant satisfiability algorithms are used, almost never involve just Presburger arithmetic alone. For such applications, the less efficient use of *theory-generic* variant satisfiability algorithms is compensated for by the trivial way in which various FVP theories can be combined by *theory union*, as opposed to by a more complex Nelson-Oppen theory combination infrastructure [32, 33]. Experimenting with such tradeoffs between domain-specific and theory-generic algorithms and their various forms of composition is an important topic for future research.

$$\begin{array}{ccc}
\mathcal{Z}_{+,>,\geq} & \xrightarrow{\text{lit2at } \delta_0} & \mathcal{Z}_+ \\
v-\downarrow & & \downarrow v- \\
\mathcal{N}_{+,>,\geq} & \xrightarrow{\text{lit2at } \delta_0} & \mathcal{N}_+
\end{array}$$

We have implemented in Maude maps (i) and (iv) as meta-level functions and, as further explained in Section 6, have used those maps effectively in a considerable number of reachability logic verification tasks.

5 Implementation and Examples

Here we describe our implementation of all the above metalevel algorithms using Maude. The complete codebase, including binaries and examples, can be downloaded from our website: <http://maude.cs.illinois.edu/tools/var-sat/>.

Thanks to the reflective nature of rewriting logic and the fact that Maude directly implements rewriting logic, we can directly represent metalevel concepts in Maude as terms in a theory. In fact, such a library already exists in Maude's META-LEVEL module. By using META-LEVEL, we can directly write functions over meta-level constructs to implement our algorithms. Essentially, the algorithm follows the outline sketched in Section 4 and shown in the diagram in the Introduction, except that the finite sort checks for theories with unit axioms have not been implemented yet. The algorithm takes as input a reflected theory M and a formula $\phi = \bigwedge G \wedge \bigwedge D$ and returns a boolean indicating if the formula is satisfiable in M . Thanks to mixfix parsing, we can use a more natural notation to write ϕ as:

$$\bar{u}_1 ==? \bar{v}_1 /\wedge \dots /\wedge \bar{u}_k ==? \bar{v}_k /\wedge \bar{u}'_1 =!? \bar{v}'_1 /\wedge \dots /\wedge \bar{u}'_l =!? \bar{v}'_l$$

where each \bar{u}_i, \bar{v}_i and \bar{u}'_j, \bar{v}'_j for $1 \leq i \leq k$ and $1 \leq j \leq l$ is a meta-term. Though in the tool we always operate at the meta-level, for readability, in this section we render all of our examples in object level notation.

5.1 Examples

We illustrate our implementation by means of a few examples. As a first example, we consider Example 4 in Section 4.3, the theory \mathcal{Z}_+ . Recall that we wanted to compute all of the *most general* constructor variants that are *less general* than variant $(x + y, id)$. Conveniently, our codebase provides a function `ctor-variants` which takes a theory and a variant and generates a complete set of most general constructor variants. When we run `ctor-variants`($\mathcal{Z}_+, (x + y, id)$) we obtain $(x' + y', \{x \mapsto x' : Nat, y \mapsto y' : Nat\})$, $(x, \{y \mapsto 0\})$, $(y, \{x \mapsto 0\})$, and $(0, \{x \mapsto 0, y \mapsto 0\})$. Note that an extra variant was generated beyond what the theory requires; in general, our algorithms only generate a complete, but not necessarily *minimal*, set of constructor variants/variant unifiers.

Example 5. (Lists of Natural Numbers) The FVP decomposition of the theory NATLIST has four sorts: *Bool*, *Nat*, *NeList*, and *List* such that $NeList < List$,

seven constructors $0 : \rightarrow Nat$, $1 : \rightarrow Nat$, $_{+} : Nat Nat \rightarrow Nat$, $_{-} : Nat List \rightarrow NeList$, $nil : \rightarrow List$, $true : \rightarrow Bool$, and $false : \rightarrow Bool$, and three defined operators $_{<} : Nat Nat \rightarrow Bool$, $hd : NeList \rightarrow Nat$, and $tl : NeList \rightarrow List$ where $_{+}$ satisfies associativity, commutativity, and identity axioms for element 0. The theory has four equations: $m + 1 + n > n = true$, $n > n + m = false$, $hd(n : l) = n$ and $tl(n : l) = l$ where $n, m : Nat$ and $l : List$.

Suppose we want to show $\phi = hd(l) > hd(l') = true \Rightarrow l \neq l'$ is a theorem of the initial algebra of NATLIST. Usually, to solve equations in this combined theory, we would need a separate solver for each subtheory and use the Nelson-Open combination method to reason in the combined theory, but here, since the theory NATLIST is FVP and protects an OS-compact subtheory, we can directly reason in the combined theory. Thus, we proceed by proving $\neg\phi$, i.e. the formula $hd(l) > hd(l') = true \wedge l = l'$, is unsatisfiable. In this case, to answer this question, it is enough to show this unification problem has no variant unifiers, since then it trivially also has no constructor variant unifiers, and thus is unsatisfiable. We provide a function `var-unifiers` which takes a theory and a unification problem and returns its set of variant unifiers. In this case, `var-unifiers(NATLIST, $\neg\phi$)` returns the empty set of unifiers, as expected.

Example 6. (Zero Predicate) The FVP decomposition of the theory ZEROPRED has sorts *Nat* and *Bool* and four constructors: $0 : \rightarrow Nat$, $s : Nat \rightarrow Nat$, $true : \rightarrow Bool$, and $false : \rightarrow Bool$ and one defined symbol $zero? : Nat \rightarrow Bool$. Finally, it satisfies two equations: $zero?(s(N)) = false$ and $zero?(0) = true$.

Suppose we want to check if $zero?(n) = x \wedge x \neq true \wedge x \neq false$ is satisfiable in ZEROPRED where $n, m : Nat$ and $x : Bool$. This example was originally used in [28] to show that variants and variant unifiers are in general insufficient to reduce the satisfiability problem from one theory into its subtheory. To see why, we can compute the variants/constructor variants of $zero?(N)$ by the aptly named functions, `variants` and `ctor-variants` which take a theory and a term and compute its variants (constructor variants) respectively. Then the function call `variants(ZEROPRED, $zero?(n)$)` gives variants $(true, \{n \mapsto 0\})$, $(false, \{n \mapsto s(m)\})$, and $(zero?(n), id)$. Obviously, letting $x = zero?(n)$, the formula above is *satisfiable* since $zero(n) \neq true \wedge zero(n) \neq false$ are both *consistent* with the empty theory; this is clearly *not* what we want. However, computing `ctor-variants(ZEROPRED, $zero?(n)$)` gives only the two variants $(true, \{n \mapsto 0\})$ and $(false, \{n \mapsto s(m)\})$ since $zero?(n)$ is not a constructor term. Substituting x by these constructor variants, the disequations are trivially inconsistent, as we expected.

6 Conclusions and Related Work

We have presented the meta-level sub-algorithms needed to obtain a full-fledged variant satisfiability algorithm, proved them correct, and derived a Maude reflective implementation. Correctness has been the main concern, but efficiency

has also been taken into account. Much work remains ahead. A crucial next step is experimentation. We have initiated such an experimentation by using the Maude reflective implementation of variant satisfiability as a key component to mechanize a new version of reachability logic for rewrite theories developed in [37], which further advances reachability logic ideas in [38, 25]. We have been able to verify various reachability properties for a substantial number of examples using the variant satisfiability algorithm as a backend procedure. This is already helping us optimize the performance of the main algorithm and its sub-algorithms, which has been an explicit theme in Sections 4.4–4.5. Furthermore, we also plan to use the variant satisfiability algorithm in other theorem proving and infinite-state model checking applications in the near future. Further work is needed to experimentally evaluate our algorithm in a more systematic way. As pointed out in Footnote 4, this should also involve comparison with domain-specific algorithms when those are available, including a comparison of the tradeoffs between different kinds of theory combination methods. Such comparisons will require developing new theory combination infrastructure not yet available in our implementation (besides of course theory unions for FVP theories, which are fully supported already).

The most closely-related work is [28–30], for which it provides the first full-fledged algorithm and implementation. Other related topics include folding variant narrowing [18], the FVP [13], and unsorted compactness [12]. Of course, this work occurs in the larger context of decidable satisfiability algorithms and the vast literature on SMT solving, e.g., [6, 23, 3, 5, 4, 6, 24, 1, 17], and additional references in [28, 29]. Finally, the literature on Maude’s reflective algorithms and tools, e.g., [9, 8] is also closely related.

Acknowledgements. Partially supported by NSF Grant CNS 13-19109.

References

1. Armando, A., Ranise, S., Rusinowitch, M.: A rewriting approach to satisfiability procedures. *Inf. Comput.* 183(2), 140–164 (2003)
2. Baader, F., Schulz, K.U.: Combining constraint solving. In: *Constraints in Computational Logics CCL’99, International Summer School*. vol. 2002, pp. 104–158. Springer LNCS (1999)
3. Barrett, C., Sebastiani, R., Seshia, S., Tinelli, C.: Satisfiability modulo theories. In: Biere, A., Heule, M.J.H., van Maaren, H., Walsh, T. (eds.) *Handbook of Satisfiability*, vol. 185, chap. 26, pp. 825–885. IOS Press (February 2009)
4. Barrett, C., Shikanian, I., Tinelli, C.: An abstract decision procedure for satisfiability in the theory of inductive data types. *Journal on Satisfiability, Boolean Modeling and Computation* 3, 21–46 (2007)
5. Barrett, C., Tinelli, C.: Satisfiability modulo theories. In: Clarke, E., Henzinger, T., Veith, H. (eds.) *Handbook of Model Checking*. Springer (2014), (to appear)
6. Bradley, A.R., Manna, Z.: *The calculus of computation - decision procedures with applications to verification*. Springer (2007)
7. Cholewa, A., Meseguer, J., Escobar, S.: Variants of variants and the finite variant property. Tech. rep., CS Dept. University of Illinois at Urbana-Champaign (February 2014), available at <http://hdl.handle.net/2142/47117>

8. Clavel, M., Durán, F., Eker, S., Meseguer, J., Lincoln, P., Martí-Oliet, N., Talcott, C.: All About Maude. Springer LNCS Vol. 4350 (2007)
9. Clavel, M., Durán, F., Eker, S., Meseguer, J., Stehr, M.O.: Maude as a formal meta-tool. In: Wing, J., Woodcock, J. (eds.) FM'99 — Formal Methods. Springer LNCS, vol. 1709, pp. 1684–1703. Springer-Verlag (1999)
10. Clavel, M., Meseguer, J., Palomino, M.: Reflection in membership equational logic, many-sorted equational logic, Horn logic with equality, and rewriting logic. *Theoretical Computer Science* 373, 70–91 (2007)
11. Comon, H., Dauchet, M., Gilleron, R., Löding, C., Jacquemard, F., Lugiez, D., Tison, S., Tommasi, M.: Tree automata techniques and applications. Available on: <http://www.grappa.univ-lille3.fr/tata> (2007), Release October, 12th 2007
12. Comon, H.: Complete axiomatizations of some quotient term algebras. *Theor. Comput. Sci.* 118(2), 167–191 (1993)
13. Comon-Lundth, H., Delaune, S.: The finite variant property: how to get rid of some algebraic properties, in *Proc RTA'05*, Springer LNCS 3467, 294–307, 2005
14. Dershowitz, N., Jouannaud, J.P.: Rewrite systems. In: van Leeuwen, J. (ed.) *Handbook of Theoretical Computer Science*, Vol. B, pp. 243–320. North-Holland (1990)
15. Dovier, A., Piazza, C., Rossi, G.: A uniform approach to constraint-solving for lists, multisets, compact lists, and sets. *ACM Trans. Comput. Log.* 9(3) (2008)
16. Dovier, A., Policriti, A., Rossi, G.: A uniform axiomatic view of lists, multisets, and sets, and the relevant unification algorithms. *Fundam. Inform.* 36(2-3), 201–234 (1998)
17. Dross, C., Conchon, S., Kanig, J., Paskevich, A.: Adding Decision Procedures to SMT Solvers using Axioms with Triggers. *Journal of Automated Reasoning* (2016), <https://hal.archives-ouvertes.fr/hal-01221066>, accepted for publication
18. Escobar, S., Sasse, R., Meseguer, J.: Folding variant narrowing and optimal variant termination. *J. Algebraic and Logic Programming* 81, 898–928 (2012)
19. Goguen, J., Meseguer, J.: Order-sorted algebra I. *Theoretical Computer Science* 105, 217–273 (1992)
20. Hendrix, J., Clavel, M., Meseguer, J.: A sufficient completeness reasoning tool for partial specifications. In: *Proc. RTA 2005*. vol. 3467, pp. 165–174. Springer LNCS (2005)
21. Hendrix, J., Meseguer, J., Ohsaki, H.: A sufficient completeness checker for linear order-sorted specifications modulo axioms. In: *Automated Reasoning, Third International Joint Conference, IJCAR 2006*. pp. 151–155 (2006)
22. Jouannaud, J.P., Kirchner, H.: Completion of a set of rules modulo a set of equations. *SIAM Journal of Computing* 15, 1155–1194 (November 1986)
23. Kroening, D., Strichman, O.: *Decision Procedures - An Algorithmic Point of View*. Texts in Theoretical Computer Science. An EATCS Series, Springer (2008)
24. Krstic, S., Goel, A., Grundy, J., Tinelli, C.: Combined satisfiability modulo parametric theories. In: *Proc. TACAS 2007*. vol. 4424, pp. 602–617. Springer LNCS (2007)
25. Lucanu, D., Rusu, V., Arusoae, A., Nowak, D.: Verifying reachability-logic properties on rewriting-logic specifications. In: *Logic, Rewriting, and Concurrency - Essays dedicated to José Meseguer on the Occasion of His 65th Birthday*. vol. 9200, pp. 451–474. Springer LNCS (2015)
26. Meseguer, J.: Membership algebra as a logical framework for equational specification. In: *Proc. WADT'97*. pp. 18–61. Springer LNCS 1376 (1998)
27. Meseguer, J.: Strict coherence of conditional rewriting modulo axioms. Tech. Rep. <http://hdl.handle.net/2142/50288>, C.S. Department, University of Illinois at Urbana-Champaign (August 2014)

28. Meseguer, J.: Variant-based satisfiability in initial algebras. Tech. Rep. <http://hdl.handle.net/2142/88408>, University of Illinois at Urbana-Champaign (November 2015)
29. Meseguer, J.: Variant-based satisfiability in initial algebras. In: Artho, C., Ölveczky, P. (eds.) Proc. FTSCS 2015. pp. 1–32. Springer CCIS 596 (2016), in press.
30. Meseguer, J.: Variant-based satisfiability in initial algebras. [28], <http://maude.cs.illinois.edu/links/varsat16.html>, revised version of Variant-Based Satisfiability in Initial Algebras
31. Meseguer, J., Goguen, J.: Order-sorted algebra solves the constructor-selector, multiple representation and coercion problems. *Information and Computation* 103(1), 114–158 (1993)
32. Nelson, G., Oppen, D.C.: Simplification by cooperating decision procedures. *ACM Trans. Program. Lang. Syst.* 1(2), 245–257 (1979)
33. Oppen, D.C.: Complexity, convexity and combinations of theories. *Theor. Comput. Sci.* 12, 291–302 (1980)
34. Rocha, C., Meseguer, J.: Constructors, sufficient completeness, and deadlock freedom of rewrite theories. In: Proc. LPAR 2010. Lecture Notes in Computer Science, vol. 6397, pp. 594–609. Springer (2010)
35. Shostak, R.E.: Deciding combinations of theories. *Journal of the ACM* 31(1), 1–12 (Jan 1984)
36. Skeirik, S., Meseguer, J.: Metalevel algorithms for variant-based satisfiability. In: Lucanu, D. (ed.) Proc. WRLA 2016. vol. 9942, pp. 167–184. Springer LNCS (2016)
37. Skeirik, S., Stefanescu, A., Meseguer, J.: A constructor-based reachability logic for rewrite theories. Tech. rep., University of Illinois Computer Science Department (December 2016), <http://maude.cs.illinois.edu/links/reachlogic16.html>
38. Stefanescu, A., Ștefan Ciobăcă, Mereuta, R., Moore, B.M., Serbanuta, T., Rosu, G.: All-path reachability logic. In: Proc. RTA-TLCA 2014. vol. 8560, pp. 425–440. Springer LNCS (2014)

A Empty and Finite Sort Constructions

In this section, we present three algorithms and prove their correctness. Given an order-sorted signature, possibly with axioms, we define rewrite theories and sentences in rewriting logic which represent solutions to the: (i) sort emptiness, (ii) sort finiteness, and (iii) term generation problems by rewrite theories implementable in the Maude rewrite engine. In the following definitions we always assume that we are reasoning over an order-sorted, kind-complete⁵ signature $\Sigma = ((S, <), F)$ where B is a set of associative/commutative/unit axioms over Σ . Before proceeding, we define some notation. For $f : s_1 \cdots s_n \rightarrow s$, let $\text{rags}(f) = \{s_1, \dots, s_n\}$ and $\text{ran}(f) = s$. Let $S_{\supset \emptyset} = \{s \in S \mid T_{\Sigma/B, s} \neq \emptyset\}$, $F_{\supset \emptyset} = \{f \in F \mid \text{args}(f) \subseteq S_{\supset \emptyset}\}$, and $\Sigma_{\supset \emptyset} = ((S_{\supset \emptyset}, <|_{S_{\supset \emptyset}}), F_{\supset \emptyset})$. Given $F' \subseteq F$, let $\Sigma|_{F'} = ((S, <), F')$. Given binary relations $R_1 \subseteq S_1 \times S_1$, and

⁵ Any signature can be easily extended to a *kind-complete* one by: (i) adding a top sort, named $[s]$, above each connected component $[s]$; and (ii) adding for each operator $f : s_1 \dots s_n \rightarrow s$ in the original signature a new typing $f : [s_1] \dots [s_n] \rightarrow [s]$. For the original sorts $s \in S$, the terms in the original signature and in its kind-completion are the *same*. Maude always perform this kind completion for any user-given signature.

$R_2 \subseteq S_2 \times S_2$, we write $R_1 \cong R_2$ iff R_1 and R_2 are bisimilar. Given $S \subseteq S_1 \cap S_2$, $R_1 \xleftrightarrow{S} R_2$ holds iff for all $s \in S$, (R_1, s) terminates iff (R_2, s) terminates where, by definition, (R, s) *terminates* iff there is no infinite R -path starting from s .

A.1 Sort Emptiness Check for General Signatures.

Here we develop an algorithm that checks if a sort $s \in S$ satisfies $T_{\Sigma, s} = \emptyset$ by performing unsorted rewriting over $\mathcal{P}(S)$. The initial state is the sort we wish to check for non-emptiness. We trace the operator declarations in reverse to see which sorts are needed to build operators inhabiting the argument sort.

Definition 11. Let $\mathcal{R}_M(\Sigma) = (\Sigma_M, ACI, R_M)$ where:

- (1) $\Sigma_M = S \uplus \{\ast\} \uplus \{-, -\}$ (an unsorted signature)
- (2) $ACI = \{\mathbf{x}, \mathbf{y} = \mathbf{y}, \mathbf{x}\} \cup \{(\mathbf{x}, \mathbf{y}), \mathbf{z} = \mathbf{x}, (\mathbf{y}, \mathbf{z})\} \cup \{\mathbf{x}, \mathbf{x} = \mathbf{x}\}$
- (3) R_M is the smallest rewrite relation such that:
 - (a) $(s, s') \in (<) \Rightarrow s' \rightarrow s \in R_M$
 - (b) $c : \rightarrow s \in F \Rightarrow s \rightarrow \ast \in R_M$
 - (c) $f : s_1 \cdots s_k \rightarrow s \in F \wedge k \geq 1 \Rightarrow s \rightarrow s_1, \dots, s_k \in R_M$

In the text below, let $(\rightarrow) \subseteq T_{\Sigma_M} \times T_{\Sigma_M}$ abbreviate $(=_{ACI}; \rightarrow_{R_M}; =_{ACI})$. We further let $(\rightarrow^0) = (=_{ACI})$, $(\rightarrow^{n+1}) = (\rightarrow); (\rightarrow^n)$, $(\rightarrow^\ast) = \bigcup_{n \geq 0} (\rightarrow^n)$, and also $(\rightarrow^+) = \bigcup_{n > 0} (\rightarrow^n)$.

Lemma 6. Let a_1, \dots, a_k , $k \geq 1$ be a ground Σ_M -term, so that $a_i \in S \uplus \{\ast\}$, i.e., a_1, \dots, a_k is a multiset. If $a_1, \dots, a_k \rightarrow^n \ast$, then for each nonempty submultiset $B \subseteq a_1, \dots, a_k$ there is an $m \leq n$ such that $B \rightarrow^m \ast$.

Proof. By induction on n .

Base Case. If $n = 0$ we must have $a_i = \ast$, $1 \leq i \leq k$, and the result follows trivially.

Induction Step. Suppose the result true for n and let $a_1, \dots, a_k \rightarrow^{n+1} \ast$. Since rewriting takes place modulo ACI we may assume without loss of generality that $i \neq j \Rightarrow a_i \neq a_j$. Then we must have some $a_i \in S$, a rule $a_i \rightarrow D$ in R_M , and rewrites

$$a_1, \dots, a_k \rightarrow a_1, \dots, a_{i-1}, D, a_{i+1}, \dots, a_n \rightarrow^n \ast.$$

Note that $a_1, \dots, a_{i-1}, D, a_{i+1}, \dots, a_n$ may have repeated elements. We now reason by cases on $B \subseteq a_1, \dots, a_k$. If $a_i \notin B$, then $B \subseteq a_1, \dots, a_{i-1}, D, a_{i+1}, \dots, a_n$ and the result follows trivially by the induction hypothesis. If $B = a_i, B'$ (where by convention B' could be empty), then $B \rightarrow D, B'$ and we have an inclusion $D, B' \subseteq a_1, \dots, a_{i-1}, D, a_{i+1}, \dots, a_n$ so the result follows again trivially by the induction hypothesis. \square

Lemma 7. $\forall s \in S [T_{\Sigma, s} \neq \emptyset \Leftrightarrow s \rightarrow^+ \ast]$

Proof. (\Rightarrow). Let $s \in S$ with $T_{\Sigma,s} \neq \emptyset$. Pick any $t \in T_{\Sigma,s}$ and proceed by structural induction on t .

Base case. [$t = c$]: Suppose $c : \rightarrow s'' \in F$ is a constant. Since $c \in T_{\Sigma,s}$, we know $s'' \leq s$. If $s'' = s$, then directly apply rule $s \rightarrow *$ generated by declaration $c : \rightarrow s'' \in F$. If $s'' < s$, we will have an additional rule $s \rightarrow s''$, which we can apply followed by $s \rightarrow *$. In either case, obtain $s \rightarrow^+ *$.

Induction Step. [$t = f(t_1, \dots, t_n)$]: Since $t = f(t_1, \dots, t_n) \in T_{\Sigma,s}$, we have $\exists f : s_1 \cdots s_k \rightarrow s'' \in F$ with $s'' \leq s$ where $t_i \in T_{\Sigma,s_i}$ for $i \in k$. If $s'' = s$, then directly apply rule $s \rightarrow s_1, \dots, s_k$ generated by declaration $f : s_1 \cdots s_k \rightarrow s'' \in F$. Since $t_i \in T_{\Sigma,s_i}$ for $i \in k$, we know that $T_{\Sigma,s_i} \neq \emptyset$. Thus, by inductive hypothesis, obtain that $s_i \rightarrow^+ *$ for $i \in k$. By transitivity, we have $s'' \rightarrow^+ *, \dots, *$. By idempotency, obtain $s'' \rightarrow^+ *$. If $s'' < s$, we will have an additional rule $s \rightarrow s''$ we can apply followed by $s'' \rightarrow^+ *$. In either case, obtain $s \rightarrow^+ *$.

(\Leftarrow). Suppose towards a contradiction the set $S' = \{s \in S \mid T_{\Sigma,s} = \emptyset \wedge s \rightarrow^+ *\}$ is non-empty. For each $s \in S'$ there is an $m(s) \in \mathbb{N}$ with $s \rightarrow^{m(s)} *$ and $m(s)$ smallest possible with that property. Pick $s_0 \in S'$ with $m(s_0)$ smallest among such $m(s)$. We now have two cases to consider: $m(s_0) = 1$ or $m(s_0) > 1$. Suppose $m(s_0) = 1$. Then $s_0 \rightarrow *$. But this can only happen if there is a $c : \rightarrow s_0 \in F$. But then $c \in T_{\Sigma,s_0}$ and $T_{\Sigma,s_0} \neq \emptyset$, a contradiction. Thus, assume $m(s_0) > 1$. Again, there are two possibilities: $s_0 \rightarrow s' \rightarrow^{m(s_0)-1} *$ or $s_0 \rightarrow s_1, \dots, s_k \rightarrow^{m(s_0)-1} *$. If $s_0 \rightarrow s' \rightarrow^{m(s_0)-1} *$, since $m(s_0)$ is smallest possible in S' , we must have $s' \notin S'$ and therefore $T_{\Sigma,s'} \neq \emptyset$. But this rewrite can only occur if $s' < s_0$. Thus, $T_{\Sigma,s'} \subseteq T_{\Sigma,s_0}$, so that $T_{\Sigma,s_0} \neq \emptyset$, a contradiction. If $s_0 \rightarrow s_1, \dots, s_k \rightarrow^{m(s_0)-1} *$, by Lemma 6 for each $1 \leq i \leq k$ we have $s_i \rightarrow^{m_i} *$ for some $m_i \leq m(s_0) - 1$. Therefore, $T_{\Sigma,s_i} \neq \emptyset$, $1 \leq i \leq k$. But the rewrite $s_0 \rightarrow s_1, \dots, s_k$ can only occur if there is an $f : s_1 \cdots s_k \rightarrow s_0 \in F$. But given any $t_i \in T_{\Sigma,s_i}$, $1 \leq i \leq k$, we can construct $f(t_1, \dots, t_k) \in T_{\Sigma,s_0}$. Thus, $T_{\Sigma,s_0} \neq \emptyset$, a contradiction. \square

There are two remaining questions: (i) is checking the sentence $s \rightarrow^+ *$ decidable? and (ii) can this approach compute emptiness of equivalence classes of terms $T_{\Sigma/E}$ defined by a theory (Σ, E) ? Fortunately, in this case, there is no extra work to be done. To answer (i), note that whenever $|S| + |F| < \aleph_0$, then $|\mathcal{P}(S)| + |R_M| < \aleph_0$ by construction. Thus, we have a finite number of states and rules, rendering the search problem decidable. To answer (ii), note that, $T_{\Sigma/E,s}$ is just an equivalence relation over $T_{\Sigma,s}$. Thus, $T_{\Sigma/E,s} = \emptyset$ iff $T_{\Sigma,s} = \emptyset$. As a result of this section, note that the set of sorts $S_{\supset \emptyset} \subseteq S$ is computable; thus, we obtain that $F_{\supset \emptyset}$ and $\Sigma_{\supset \emptyset}$ are computable as well.

A.2 Term Generation for General Signatures.

In this section, we present an algorithm which, given an order-sorted signature Σ and a sort s , will generate all terms in $T_{\Sigma,s}$. We begin with a few opening remarks. Note that: (i) an order-sorted signature Σ can be modeled as a tree automaton so that $t \in T_{\Sigma,s}$ iff t is accepted by the corresponding automaton when the accepting state is s ; and (ii) any tree automaton and its computations can be modeled as an unsorted ground rewrite theory. Clearly, an order-sorted ground

rewrite theory will also work; here we prefer an order-sorted theory because it gives a simpler definition that preserves the original signature. Throughout this section, we let S_Σ denote the signature of constants s associated to sorts $s \in S$, where each sort s is declared a constant whose sort is the top sort $[s]$: $S_\Sigma = ((S, <), \{s : \rightarrow [s] \mid s \in S\})$.

Definition 12. Let $\mathcal{R}_P(\Sigma) = (\Sigma_{\supset \emptyset} \uplus S_\Sigma, \emptyset, R_P)$ where R_P is the smallest rewrite relation $R_P = R_{P,S} \uplus R_{P,NC} \uplus R_{P,C}$ such that:

- (a) $(s, s') \in (<) \Rightarrow s \rightarrow s' \in R_{P,S}$
- (b) $f : s_1 \cdots s_k \rightarrow s \in F_{\supset \emptyset} \wedge k \geq 1 \Rightarrow f(s_1, \dots, s_k) \rightarrow s \in R_{P,NC}$
- (c) $c : \rightarrow s \in F_{\supset \emptyset} \Rightarrow c \rightarrow s \in R_{P,C}$

Note that, even though $\Sigma_{\supset \emptyset} \subseteq \Sigma$, we do not lose completeness for parsing, since any sort in $s \in S/S_{\supset \emptyset}$ necessarily satisfies $T_{\Sigma,s} = \emptyset$. Furthermore, it is straightforward to show that $\Sigma \uplus S_\Sigma$ is sensible and preregular iff Σ is sensible and preregular and $\forall s \in S_{\supset \emptyset} [t \in T_{\Sigma,s} \Leftrightarrow t \xrightarrow{+}_{R_P} s]$. We now turn to term generation.

Definition 13. Let $\mathcal{R}_G(\Sigma) = (\Sigma_{\supset \emptyset} \uplus S_\Sigma, \emptyset, R_G)$ with $R_G = R_P^{-1}$. Since $R_P = R_{P,S} \uplus R_{P,NC} \uplus R_{P,C}$ we will use the notation: $R_{G,S} = R_{P,S}^{-1}$, $R_{G,NC} = R_{P,NC}^{-1}$, and $R_{G,C} = R_{P,C}^{-1}$.

Again, by only considering $\Sigma_{\supset \emptyset} \subseteq \Sigma$, we do not lose completeness for term generation. We immediately obtain the following corollary.

Corollary 5. $\forall s \in S_{\supset \emptyset} [t \in T_{\Sigma \uplus S_\Sigma} \Leftrightarrow s \xrightarrow{!}_{R_G} t]$

A.3 Finite Sort Detection for Finite Signatures.

Here we develop an algorithm which, given $s \in S$, checks if $|T_{\Sigma,s}| < \aleph_0$. Note that using \mathcal{R}_G we already trivially obtain a semi-decidable algorithm for sort finiteness: compute $S_{\supset \emptyset}$ via \mathcal{R}_M ; if $s \notin S_{\supset \emptyset}$, then return yes; otherwise compute $\{t \in T_{\Sigma,s} \mid s \xrightarrow{!}_{R_G} t\}$; if the process terminates, then return yes. Of course, an efficient, decidable algorithm would be preferable. Nevertheless, \mathcal{R}_G is not too far from our desired decidable solution.

Our strategy is as follows: (i) give sufficient conditions so that termination of \mathcal{R}_G corresponds to sort finiteness in Σ , (ii) define a rewrite system \mathcal{R}_F and give sufficient conditions to prove termination of \mathcal{R}_F , (iii) show \mathcal{R}_F terminates if and only if \mathcal{R}_G terminates, (iv) and finally, present a decidable algorithm using LTL model checking to characterize when \mathcal{R}_F terminates.

Lemma 8. If $|S| + |F| < \aleph_0$ then (\mathcal{R}_G, s) is non-terminating iff $|T_{\Sigma,s}| = \aleph_0$

Proof. By construction of R_G , $|R_G| = |(<)| + |F| < |S|^2 + |F| < \aleph_0$. Viewing possible rewrite paths starting from s as forming a tree, observe that the tree branches finitely, since each term has finite positions and possible rewrites. Suppose (\mathcal{R}_G, s) is terminating. Then, by König's Lemma, the tree of rewrites

must be finite and therefore there is a finite number of final states, so that $|T_{\Sigma,s}| < \aleph_0$. Otherwise, if (\mathcal{R}_G, s) is non-terminating, we have an infinite path $s \rightarrow_{R_G} t_1 \rightarrow_{R_G} t_2 \rightarrow_{R_G} \cdots t_n \rightarrow_{R_G} \cdots$. Since $|R_G| < \aleph_0$, $\exists R \subseteq R_G$ that repeats infinitely often. Since $R_G = R_{G,S} \uplus R_{G,C} \uplus R_{G,NC}$ and $R_{G,S} \uplus R_{G,C}$ terminates (because acyclicity/finiteness of $<$ and only S -terms can be rewritten), we must have $R \cap R_{G,NC} \neq \emptyset$. But note that, if $|t|$ is the of t as viewed as a tree, then if $t \rightarrow_{R_{G,S} \uplus R_{G,C}} t'$, we must have $|t| = |t'|$, whereas if $t \rightarrow_{R_{G,NC}} t'$, we must have $|t| < |t'|$, so that $\{|t_i|\}_{i \in \mathbb{N}}$ is a sequence such that $|t_i| \rightarrow \infty$. Also note that by the definition of R_G , all sorts s' occurring as a subterm of t_i belong to $S_{\supset \emptyset} = \{s_1, \dots, s_m\}$, so that we can choose terms $u_1 \in T_{\Sigma,s_1}, \dots, u_m \in T_{\Sigma,s_m}$. We can then regard $S_{\supset \emptyset}$ as a set of variables and view $\sigma = \{s_1 \mapsto u_1, \dots, s_m \mapsto u_m\}$ as a substitution. But, by definition of R_G , this gives us an infinite sequence $\{t_i \sigma\}_{i \in \mathbb{N}}$ of terms where for each $i \in \mathbb{N}$, $t_i \sigma \in T_{\Sigma,s}$ and $|t_i \sigma| \geq |t_i|$. Therefore, $|t_i \sigma| \rightarrow \infty$, and since $T_{\Sigma,s}$ contains terms of unbounded size, we have $|T_{\Sigma,s}| = \aleph_0$. \square

Definition 14. Let $\mathcal{R}_F(\Sigma) = (S_{\supset \emptyset}, \emptyset, R_F)$ where $R_F = R_{F,S} \cup R_{F,NC}$ is the smallest rewrite relation such that:

- (a) $s < s' \Rightarrow s' \rightarrow s \in R_{F,S}$
- (b) $f : s_1 \cdots s_n \rightarrow s' \in F_{\supset \emptyset} \wedge \{s\} \subseteq \{s_1, \dots, s_n\} \Rightarrow s' \rightarrow s \in R_{F,NC}$

Note that we only consider $S_{\supset \emptyset}$ and $F_{\supset \emptyset}$, because, implicitly, any sort $s \in S/S_{\supset \emptyset}$ trivially satisfies $|T_{\Sigma,s}| < \aleph_0$ and any operator $f \in F/F_{\supset \emptyset}$ cannot contribute meaningfully to building a term $t \in T_{\Sigma,s}$. Before we complete the main proof, we prove a lemma and add an additional definition.

Lemma 9. Given $|S_{\supset \emptyset}| < \aleph_0$ and $s \in S_{\supset \emptyset}$, then the following are equivalent:

1. (\mathcal{R}_F, s) is non-terminating
2. $\exists s' \in S_{\supset \emptyset} [s \xrightarrow{*}_{R_F} s' \xrightarrow{+}_{R_F} s']$
3. there is an infinite R_F -rewrite path $s \rightarrow_{R_F} s_1 \rightarrow_{R_F} s_2 \cdots \rightarrow_{R_F} s_n \rightarrow_{R_F} \cdots$ and $s' \in S_{\supset \emptyset}$ occurring infinitely often in the sequence

Proof. Obviously, (3) implies (2), since if s' occurs infinitely often, we must have $s \xrightarrow{*}_{R_F} s' \xrightarrow{+}_{R_F} s'$. Also, (2) implies (1) since $s \xrightarrow{*}_{R_F} s' \xrightarrow{+}_{R_F} s' \xrightarrow{+}_{R_F} s' \xrightarrow{+}_{R_F} \cdots$ is a non-terminating sequence. Finally, (1) implies (3), since $|S_{\supset \emptyset}| \leq \aleph_0$, which forces some $s' \in S_{\supset \emptyset}$ to occur infinitely often in any infinite sequence. \square

Definition 15. Given $\Sigma = ((S, <), NC \uplus C)$ with non-constants and constants NC and C respectively, let $\mathcal{R}_G^{\star}(\Sigma) = (\Sigma_{\supset \emptyset} |_{NC \uplus S_{\Sigma}}, \emptyset, R_{G,\star})$ such that $R_{G,\star} = R_{G,S} \uplus R_{G,NC}$.

Observe that \mathcal{R}_G^{\star} is identical to \mathcal{R}_G except that \mathcal{R}_G^{\star} contains neither constants nor rewrite rules over constants. Now we are ready to prove the main theorem.

Theorem 9. $\mathcal{R}_F \cong \mathcal{R}_G^{\star}$ and $\mathcal{R}_G^{\star} \xleftrightarrow{S_{\supset \emptyset}} \mathcal{R}_G$

Proof. We first prove $\mathcal{R}_F \cong \mathcal{R}_G^\star$. Define a relation $H \subseteq (S_{\supset\emptyset} \times T_{\Sigma|_{NC \cup \hat{S}}})$ where $(s, t) \in H$ iff $s \leq t$. To prove $\mathcal{R}_F \cong \mathcal{R}_G^\star$, we show that given two arrows, we can find another two arrows to make the diagrams below commute.

$$\begin{array}{ccc} s & \xrightarrow{R_F} & s' \\ H \downarrow & & \downarrow H \\ t & \xrightarrow{R_{G,\star}} & t' \end{array} \qquad \begin{array}{ccc} s & \xrightarrow{R_F} & s' \\ H \downarrow & & \downarrow H \\ t & \xrightarrow{R_{G,\star}} & t' \end{array}$$

Suppose $s \leq t$. If $(s, s') \in R_F$ then $(s, s') \in R_{F,S}$ or $(s, s') \in R_{F,NC}$. Assume $(s, s') \in R_{F,S}$. Then $s' < s$ in $\Sigma_{\supset\emptyset}$. But then, by definition, $(s, s') \in R_{G,S}$. Thus, $t[s] \rightarrow_{R_{G,\star}} t[s']$ and $s' \leq t[s']$, as required. Alternatively, assume $(s, s') \in R_{F,NC}$. Then $\exists f : s_1 \cdots s_n \rightarrow s' \in F_{\supset\emptyset}$ with $\{s\} \subseteq \text{args}(f)$. But then, by definition, $(s', f(s_1, \dots, s_n)) \in R_{G,NC}$. Thus, $t[s] \rightarrow_{R_G^{NC}} t[f(s_1, \dots, s_n)]$ and $s' \leq t[f(s_1, \dots, s_n)]$. Since we used only definitional equivalences, the other direction follows symmetrically.

To prove $\mathcal{R}_G^\star \xleftrightarrow{S_{\supset\emptyset}} \mathcal{R}_G$, given $s \in S_{\supset\emptyset}$, we must show (\mathcal{R}_G^\star, s) terminates iff (\mathcal{R}_G, s) terminates. To begin, note $R_G = R_{G,\star} \uplus R_{G,C}$. Thus, if $R_{G,\star}$ is non-terminating, R_G must also be non-terminating. To see the other direction, note $R_{G,C}$ always terminates since each rule has the form $s \rightarrow c \in C$ and constants cannot be rewritten. We proceed by proving the contrapositive. Thus, assume $R_{G,\star}$ terminates. By Lemma 10, $s \rightarrow_{R_G}^n t$ iff $s \rightarrow_{R_{G,\star}}^i t' \rightarrow_{R_{G,C}}^j t$ with $n = i + j$. Since $R_{G,\star}$ and $R_{G,C}$ are terminating and finitely branching, there are maximum bounds on the size of i and j , say, i_{max} and j_{max} respectively. But then any rewrite path $s \rightarrow_{R_G}^n t$ necessarily has $n \leq i_{max} + j_{max}$; thus (R_G, s) is terminating. \square

Lemma 10. $\forall n \in \mathbb{N} [[s \rightarrow_{R_G}^n t] \Leftrightarrow [\exists i, j \in \mathbb{N} [s \rightarrow_{R_{G,\star}}^i t' \rightarrow_{R_{G,C}}^j t \wedge n = i + j]]]$

Proof. To begin, recall $R_G = R_{G,\star} \uplus R_{G,C}$ and note the following equivalence for $s \in S_{\supset\emptyset}, n \in \mathbb{N}$, and $t \in T_\Sigma$:

$$\begin{aligned} & s \rightarrow_{R_G}^n t \\ & \Leftrightarrow \\ & \exists l_1, l_2, m_1, m_2 \in \mathbb{N} \exists t', t'', t''', t^{iv} \in T_\Sigma \\ & \quad [[[s \rightarrow_{R_{G,\star}}^{l_1} t' \rightarrow_{R_{G,C}}^{l_2} t] \vee [s \rightarrow_{R_{G,\star}}^{m_1} t'' \rightarrow_{R_{G,C}} t''' \rightarrow_{R_{G,\star}} t^{iv} \rightarrow_{R_G}^{m_2} t]] \wedge \\ & \quad l_1 + l_2 = m_1 + m_2 + 2 = n] \end{aligned}$$

That is, either all the applications of rules in $R_{G,C}$ occur at the end, or there is at least one such application *before* a rule in $R_{G,\star}$. Since the first case already fits the desired form, we need only consider the second case. Note all rules in R_G have the form $S \ni s \rightarrow t \in T_{\Sigma \cup S_\Sigma}$. $R_{G,C}$ rules in particular have the form $s \rightarrow c$ for $c \in F$. Thus, if a $R_{G,C}$ rule is applied to $t[s]_p$ at position p , a $R_{G,\star}$ rule cannot later also be applied at p . Now suppose $s \rightarrow_{R_{G,\star}}^{m_1} t'' \rightarrow_{R_{G,C}} t''' \rightarrow_{R_{G,\star}} t^{iv} \rightarrow_{R_G}^{m_2} t$. Then, $t'' = t''[s', s'']_{p,q}$ with p, q disjoint positions and:

$$\begin{array}{ccc}
s & \xrightarrow{R_{G,\star}}^* t''[s', s''] & \xrightarrow{R_{G,C}} t''[c, s''] \\
& \downarrow R_{G,\star} & \downarrow R_{G,\star} \\
& t''[s', u] & \xrightarrow{R_{G,C}} t''[c, u]
\end{array}$$

for any $c \in C$ and $u \in T_{\Sigma \uplus S \Sigma}$, the diagram above commutes. We complete the proof by induction on m_2 , the number of rewrites occurring after the first $R_{G,C}$ rule followed by a $R_{G,\star}$ rule. Suppose $m_2 = 0$. Then we can commute the $R_{G,\star}$ and $R_{G,C}$ arrows as above, to obtain a rewrite chain of the form $s \xrightarrow{R_{G,\star}}^{m_1+1} v \xrightarrow{R_{G,C}} t$, for some $v \in T_{\Sigma \uplus S \Sigma}$, as required. Now suppose $m_2 > 0$. Again, we commute the two arrows to obtain $s \xrightarrow{R_{G,\star}}^{m_1+1} v_1 \xrightarrow{R_{G,C}} v_2 \xrightarrow{R_{G,C}}^{m_2} t$. We apply our induction hypothesis to obtain $s \xrightarrow{R_{G,\star}}^{m_1+1} v_1 \xrightarrow{R_{G,\star}}^{k_1} v_3 \xrightarrow{R_{G,C}}^{k_2} t$ with $k_1+k_2 = m_2$ which is equivalent to $s \xrightarrow{R_{G,\star}}^{m_1+k_1+1} v_3 \xrightarrow{R_{G,C}}^{k_2} t$ and $m_1+k_1+k_2+1 = m_1 + m_2 + 1 = n$, as required. \square

Thus, according to Lemmas 8 and 9 and Theorem 9, (\mathcal{R}_F, s) will generate a rewrite path containing a cycle iff $|T_{\Sigma, s}| = \aleph_0$. To complete the proof, for any $s \in S$, we just to characterize when $\exists s' \in S_{\supset \emptyset} [s \xrightarrow{R_F}^* s' \xrightarrow{R_F}^+ s']$ holds. Thus, define the set of *cycle sorts* by $cy(S_{\supset \emptyset}) = \{s \in S_{\supset \emptyset} \mid s \xrightarrow{R_F}^+ s\}$. This set can be computed by search, since the sort set and rules are both finite. Then, we immediately obtain the following theorem.

Theorem 10. $\forall s \in S_{\supset \emptyset} \ |T_{\Sigma, s}| = \aleph_0$ iff $\bigvee_{s' \in cy(S_{\supset \emptyset})} R_F \vdash s \xrightarrow{*} s'$

Proof. By Lemmas 8 and 9 and Theorem 9, obtain $|T_{\Sigma, s}| = \aleph_0$ iff the formula $\exists s' \in S_{\supset \emptyset} [s \xrightarrow{R_F}^* s' \xrightarrow{R_F}^+ s']$ holds. But by definition, any s' which satisfies the formula satisfies $s' \in cy(S_{\supset \emptyset})$, so reduce to $\exists s' \in cy(S_{\supset \emptyset}) [s \xrightarrow{R_F}^* s']$. Since S is finite by assumption, $cy(S_{\supset \emptyset})$ is finite. So, reduce to $\bigvee_{s' \in cy(S_{\supset \emptyset})} s \xrightarrow{R_F}^* s'$, which holds iff $\bigvee_{s' \in cy(S_{\supset \emptyset})} R_F \vdash s \xrightarrow{*} s'$ holds, as required. \square

A final consideration is how to check, for a theory (Σ, B) , whether equivalence classes of terms $T_{\Sigma/B, s}$ are finite, given that $T_{\Sigma, s}$ is finite. Since $T_{\Sigma/B, s}$ is a set of B -equivalence classes $[t]$, each containing at least one $t' \in [t]$ with $t' \in T_{\Sigma, s}$, if $|T_{\Sigma, s}| < \aleph_0$, then $T_{\Sigma/B, s} < \aleph_0$. Nevertheless, in general, it may be the case that $|T_{\Sigma/B, s}| < \aleph_0$ but $|T_{\Sigma, s}| = \aleph_0$.

Example 7. $\Sigma = ((\{a, b\}, \{(a, b)\}), 0 := a, 1 := b, - + - : aa \rightarrow a, - + - : bb \rightarrow b)$. Let B contain a unit axiom for 0 over $(+)$. Then $|T_{\Sigma, a}| = |T_{\Sigma, b}| = \aleph_0$ but $|T_{\Sigma/B, a}| = 1$ and $|T_{\Sigma/B, b}| = \aleph_0$.

However, under some conditions on B , finiteness of $T_{\Sigma/B, s}$ can still be checked.

Lemma 11. *Suppose B is a set of associativity and/or commutativity axioms, $|\Sigma| < \aleph_0$, and that Σ is B -preregular. Then $|T_{\Sigma/B, s}| < \aleph_0$ iff $|T_{\Sigma, s}| < \aleph_0$.*

Proof. Since Σ is B -preregular, all axioms in B are sort preserving. Then obtain $[u]_B \in T_{\Sigma/AC, s}$ iff $[u]_B \subseteq T_{\Sigma, s}$, proving (\Leftarrow) . To show (\Rightarrow) , note that for any combination of associativity and/or commutativity axioms, $[u]_B$ is a *finite* set. Since $T_{\Sigma/B, s}$ is finite, then $T_{\Sigma, s}$ is a finite union of finite sets and thus finite. \square

Let U be a set of unit axioms for unit elements $e_1 := s_1, \dots, e_n := s_n$ in Σ . Then define $\Sigma - U = \Sigma - \{e_1 := s_1, \dots, e_n := s_n\}$.

Lemma 12. *Let B_0 be a set of associative and/or commutative axioms and U a set of unit axioms in Σ , $B = B_0 \uplus U$, $|\Sigma| < \aleph_0$, and $\Sigma = ((S, <), F)$ be B -preregular according to Footnote 1. If $|T_{\Sigma-U,s}| = \aleph_0$, then $|T_{\Sigma/B,s}| = \aleph_0$.*

Proof. We can orient a unit axiom $f(x, e) = x$ as a rewrite rule $f(x, e) \rightarrow x$, so that the set U becomes a set of rewrite rules $R(U)$. In this way the theory $(\Sigma, B_0 \uplus U)$ can be decomposed as a convergent rewrite theory $(\Sigma, B_0, R(U))$. Observe $T_{\Sigma-U/B_0} \subseteq C_{\mathcal{R}_U}$ and $C_{\mathcal{R}_U} \cong T_{\Sigma/B}$. By Lemma 11, $|T_{\Sigma-U,s}| = \aleph_0$ iff $|T_{\Sigma-U/B_0,s}| = \aleph_0$. Thus, $\aleph_0 = |T_{\Sigma-U,s}| = |T_{\Sigma-U/B_0,s}| \leq |C_{\mathcal{R}_B,s}| = |T_{\Sigma/B,s}|$. Since $|T_{\Sigma/B,s}| \leq \aleph_0$, obtain $|T_{\Sigma/B,s}| = \aleph_0$, as required. \square

The following lemma gives sufficient conditions such that $|T_{\Sigma,s}| = \aleph_0$ but $|T_{\Sigma/B,s}| < \aleph_0$ when B is a combination of associativity and/or commutativity and/or unit axioms.

Lemma 13. *Let B_0 be a set of associative and/or commutative axioms and U a set of unit axioms in Σ , $B = B_0 \uplus U$, $|\Sigma| < \aleph_0$, and $\Sigma = ((S, <), F)$ be B -preregular according to Footnote 1. Let $f : s_1 s_2 \rightarrow s'$ with $ls(e) \leq s_1, s_2 \leq s' \leq s$ and let e be a unit element satisfying either a left-unit, right-unit, or left- and right-unit axiom(s) for f with $s \in S$. If $\nexists g : w \rightarrow s'' \in F/\{f, e\}[s'' \leq s]$ then $|T_{\Sigma,s}| = \aleph_0$ and $T_{\Sigma/B,s} = \{\{e\}\}$.*

Proof. By an easy structural induction, $\forall u \in T_{\Sigma,s}[u!_{R(U),B_0} = e]$. \square

A.4 Decidable Sort Classifications

Here, we present a summary of the results of the previous sections by illustrating how our methods can be used to compute a partitioning of S that respects sort classifications.

Corollary 6. *Let B be a set of associative and/or commutative axioms, $|\Sigma| < \aleph_0$, and Σ be B -preregular. Then S has the following computable partitioning:*

$$S = S_{\supset \emptyset} \uplus S_{\emptyset} = S_{\infty} \uplus S_F \uplus S_{\emptyset}$$

where $S_{\infty} = \{s \in S_{\supset \emptyset} \mid |T_{\Sigma/B,s}| = \aleph_0\}$ and $S_F = S_{\supset \emptyset}/S_{\infty}$.

Proof. First apply Lemma 11 to reduce to the case with no axioms. By Lemma 7, $s \mapsto_{R_M, ACI} * \text{ iff } s \in S_{\emptyset}$, and $S_{\supset \emptyset} = S/S_{\emptyset}$. Thus, obtain $\Sigma_{\supset \emptyset}$. By Theorem 10, if $s \in S_{\supset \emptyset}$ then $s \in S_F$ iff $\neg(\bigvee_{s' \in cy(S_{\supset \emptyset})} R_F \vdash s \rightarrow^* s')$. Otherwise, by definition, $s \in S_{\infty}$. Since each step—performing search via $(=_{ACI}; \rightarrow_{R_M}; =_{ACI})$, filtering $F_{\supset \emptyset}$, computing $cy(S_{\supset \emptyset})$, and search over R_F —is decidable, the entire sort classification algorithm is decidable, as required. \square

In the more general ACU case, this partitioning can no longer be computed by the methods we have presented. However, in many cases we can still compute such a partition, for example if all sorts s for which $|T_{\Sigma,s}| = \aleph_0$ fall into one of the cases laid out in Lemmas 12 and 13. Otherwise, the partitioning algorithm will fail to classify some sorts, leaving some proof obligations for the user.

B Auxiliary Lemmas for Section 4.2

In these proofs, we always assume $(S^c, <^c)$ is a constructor sort refinement of $(S, <)$. In Lemma 1, we require two simple lemmas which are left as an exercise to the reader. Let Σ be an arbitrary signature. Then (1) if Σ is preregular and $f(t_1, \dots, t_n) \in T_\Sigma$ then $ty_\Sigma(f(t_1, \dots, t_n)) = ty_\Sigma(f, ls_\Sigma(t_1) \cdots ls_\Sigma(t_n))$ with $n \geq 0$ and (2) $t \in T_\Sigma \Leftrightarrow ty_\Sigma(t) \neq \emptyset$.

Lemma 1. *If $\Omega < \Sigma$ then $\forall t \in T_\Sigma [t \in T_\Omega \Rightarrow ls_\Omega(t) = ls_\Sigma(t)]$.*

Proof. Assume $\Omega < \Sigma$ and $t \in T_\Omega$. Suppose that $t = c \in T_\Omega$ is a constant. Then $ty_\Omega(c, nil) \neq \emptyset$ and $min_{<}(ty_\Sigma(c, nil)) \in ty_\Omega(c, nil)$. Since we have $ty_\Omega(c, nil) \subseteq ty_\Sigma(c, nil)$ then $min_{<}(ty_\Omega(c, nil)) = min_{<}(ty_\Sigma(c, nil))$ and $ls_\Omega(t) = ls_\Sigma(t)$. Now suppose $t = f(t_1, \dots, t_n)$. Then $ty_\Omega(f(t_1, \dots, t_n)) \neq \emptyset$ and $ty_\Omega(f, w) \neq \emptyset$ where $w = ls_\Omega(t_1) \cdots ls_\Omega(t_n)$. But $t_1 \cdots t_n \in T_\Omega$, so by induction hypothesis, $w = ls_\Sigma(t_1) \cdots ls_\Sigma(t_n)$. Since $min_{<}(ty_\Sigma(f, w)) \in ty_\Omega(f, w)$ and $ty_\Omega(f, w) \subseteq ty_\Sigma(f, w)$, then we have $min_{<}(ty_\Omega(f, w)) = min_{<}(ty_\Sigma(f, w))$ and $ls_\Omega(t) = ls_\Sigma(t)$, as required. \square

Lemma 2. *$\forall t \in T_{\Omega^\downarrow}(X^c)/X^c [ty_{\Omega^\downarrow(X^c)}(t) = ty_{\Omega^\downarrow(X^\downarrow)}(t) = ty_{\Omega^\downarrow(X^\downarrow)}(t)]$*

Proof. The base case where $t = c \in T_{\Omega^\downarrow}(X^c)/X^c$, a constant, is trivial, so suppose $t = f(t_1, \dots, t_n)$. There are two cases: either for each $1 \leq i \leq n$, we have $vars(t_i) \subseteq X^\downarrow$ or not. If not, $ty_{\Omega^\downarrow(X^c)}(t) = ty_{\Omega^\downarrow(X^\downarrow)}(t) = ty_{\Omega^\downarrow(X^\downarrow)}(t) = \emptyset$ since these three signatures share the same non-variable operators F_Ω^\downarrow whose arity is contained in $(S^\downarrow)^*$. Otherwise, by induction hypothesis, for $1 \leq i \leq n$, we have $ty_{\Omega^\downarrow(X^c)}(t_i) = ty_{\Omega^\downarrow(X^\downarrow)}(t_i) = ty_{\Omega^\downarrow(X^\downarrow)}(t_i)$, and since operators F_Ω^\downarrow are shared, we have $ty_{\Omega^\downarrow(X^c)}(t) = ty_{\Omega^\downarrow(X^\downarrow)}(t) = ty_{\Omega^\downarrow(X^\downarrow)}(t)$. \square

Lemma 3. *$\forall t \in T_{\Sigma^c}(X^c)/X^c [ty_{\Sigma^+(X^c)}(t) = ty_{\Sigma(X)}(t^\bullet)]$*

Proof. The case where $t = c \in T_{\Sigma^c}(X^c)/T_{\Omega^\downarrow}(X^c)$, a constant, is trivial, so suppose $t = f(t_1, \dots, t_n)$. By definition, $\exists f : s_1 \cdots s_n \rightarrow s \in F$ with $s_i \in S$, $t_i : s'_i$, and $s'_i \leq^c s_i$ for $1 \leq i \leq n$. But $(\bullet) : (S, <^c) \rightarrow (S, <)$ —also $(\bullet) : \Sigma^+(X^c) \rightarrow \Sigma(X) \subseteq \Sigma^+(X^c)$ —is a poset/signature morphism, so $s'_i \bullet \leq s_i \bullet = s_i$, $t_i^\bullet \in T_{\Sigma(X)}$, and $ty_{\Sigma^+(X^c)}(t) = ty_{\Sigma^+(X^c)}(t^\bullet)$. Also note $ty_{\Sigma^+(X^c)}|_{T_{\Sigma(X)}} = ty_{\Sigma(X)}$, since $f : s_1 \cdots s_n \rightarrow s \in F \cup X^c$ with $s_1 \cdots s_n \in S^*$ iff $f : s_1 \cdots s_n \rightarrow s \in F \cup X$. But $t^\bullet \in T_{\Sigma(X)}$, thus $ty_{\Sigma^+(X^c)}(t) = ty_{\Sigma^+(X^c)}(t^\bullet) = ty_{\Sigma(X)}(t^\bullet)$, as required. \square

Lemma 4. *$\Sigma = ((S, <), F)$ is sensible iff $\widehat{\Sigma} = ((\widehat{S}, \emptyset), \widehat{F})$ is sensible where $f : [s_1] \cdots [s_n] \rightarrow [s_0] \in \widehat{F}$ iff $\exists f : s'_1 \cdots s'_n \rightarrow s'_0 \in F$ with $s'_i \in [s_i]$ for $0 \leq i \leq n$.*

Proof. Given a tuple of sorts $w = s_1 \cdots s_n$, let $[w] = [s_1] \cdots [s_n]$. To see (\Rightarrow) , suppose if $f : w \rightarrow s, f : w' \rightarrow s' \in F$ and $w \equiv_{\leq} w'$ then $s \equiv_{\leq} s'$. But note $w \equiv_{\leq} w'$ iff $w, w' \in [w]$ and $s \equiv_{\leq} s'$ iff $s, s' \in [s]$. To see (\Leftarrow) , assume $f : [w] \rightarrow [s], f : [w'] \rightarrow [s'] \in \widehat{F}$ and $[w] \equiv_{\leq} [w']$ then $[s] \equiv_{\leq} [s']$. But note $[w] \equiv_{\leq} [w']$ iff $[w] = [w']$ and $[s] \equiv_{\leq} [s']$ iff $[s] = [s']$ since in $\widehat{\Sigma}$, $(\leq) = \emptyset$.

Then assume towards a contradiction that $\exists f : w_1 \rightarrow s_1, f : w_2 \rightarrow s_2 \in F$ with $w_1 \equiv_{\leq} w_2$ and $s_1 \not\equiv_{\leq} s_2$. But then $w_1, w_2 \in [w_1] = [w_2]$ and $s_1 \in [s_1]$ and $s_2 \in [s_2]$ with $[s_1] \neq [s_2]$, a contradiction.