

Polite Theories Revisited*

Dejan Jovanović and Clark Barrett

New York University

dejan@cs.nyu.edu, barrett@cs.nyu.edu

Abstract. The classic method of Nelson and Oppen for combining decision procedures requires the theories to be stably-infinite. Unfortunately, some important theories do not fall into this category (e.g. the theory of bit-vectors). To remedy this problem, previous work introduced the notion of *polite* theories. Polite theories can be combined with any other theory using an extension of the Nelson-Oppen approach. In this paper we revisit the notion of polite theories, fixing a subtle flaw in the original definition. We give a new combination theorem which specifies the degree to which politeness is preserved when combining polite theories. We also give conditions under which politeness is preserved when instantiating theories by identifying two sorts. These results lead to a more general variant of the theorem for combining multiple polite theories.

1 Introduction

The seminal paper of Nelson and Oppen [5] introduced a general combination framework that allows the creation of a decision procedure for the combination of two first-order theories in a *modular* fashion. Using the Nelson-Oppen framework, decision procedures for two individual theories can be used as black boxes to create a decision procedure for the combined theory.

Although very general and widely-used in practice, the Nelson-Oppen approach is not applicable to all theories encountered in practical applications. A significant restriction of Nelson-Oppen is the requirement that theories be *stably-infinite*. While many important theories are stably-infinite, some are not, including those with inherently finite domains such as the theory of bit-vectors. As bit-precise reasoning about both programs and hardware is becoming more important and more feasible, it is desirable to find ways of overcoming this restriction.

As a possible remedy for this problem, the notion of *shiny theories* and an appropriate combination algorithm was introduced in [13]. The requirements on a shiny theory are stronger than just stable-infiniteness, but this allows it to be combined with an arbitrary other (possibly non-stably-infinite) theory. The main drawback to this approach is the requirement that a shiny theory T has to be equipped with a function $mincard_T$. This function, given a set of constraints, must be able to compute the minimal cardinality of a T -interpretation that satisfies the constraints.

* This work was funded in part by SRC contract 2008-TJ-1850.

A related approach for combining theories is presented in [4]. The authors start from a framework of parametrically polymorphic logics to devise a Nelson-Oppen-style combination procedure for theories that are *flexible*. Flexibility is a property similar to the ability to move to a bigger or a smaller (infinite) model via the Löwenheim-Skolem theorem in first-order logic. Most commonly-used theories can be represented in this framework and are shown to be flexible. Reasoning about cardinality also plays a major role in this approach—a solver for a parametric theory (called a strong solver) is required to process not only the formula being checked, but also a set of cardinality constraints over the domain sizes. Although this direction is promising, particularly because of the advantages of parametricity, the approach as developed thus far would be cumbersome to implement in a practical system. In particular, while reasoning about cardinalities is possible for a wide class of important theories, it can be computationally expensive, and theory decision procedures are typically not designed with this additional requirement in mind.

An alternative approach uses the notion of *polite theories* introduced in [8]. Polite theories can also be combined with an arbitrary other theory. However, this approach does not require the computation of the *mincard* function. Instead, a decision procedure for a polite theory must be able to generate explicitly a *witness* formula that enumerates any required domain elements using additional variables. The authors show that many commonly-used theories are polite (including theories of lists, arrays, sets, and multi-sets). This approach seems more practical than those that require reasoning about cardinalities explicitly. And, while proving that a theory is polite can be difficult and needs to be done on a per-theory basis, once this is done, the combination method can be easily implemented.

In this paper, we revisit and extend the results on polite theories from [8]. Section 2 gives definitions and background on many-sorted logic and theory combination. Section 3 begins by introducing polite theories, making a small but needed modification to the definition of finite witnessability (one of the two properties that make up politeness), and then goes on to show that when combining polite theories, the resulting theory is also polite (with respect to a possibly reduced set of sorts). Section 4 addresses *theory instantiation*, the construction of a new theory by identifying two sorts in an existing theory; we prove that instantiation preserves politeness. Finally, Section 3.4 discusses the combination of multiple polite theories, culminating in a combination result that is more general than the ones presented in [8].

Due to space limitations, proofs of the results in this paper are omitted. The full paper with proofs is available from the authors as a technical report [3].

2 Preliminaries

2.1 Many-Sorted First-Order Logic

We start with a brief overview of the syntax and semantics of many-sorted first-order logic. For a more detailed exposition, we refer the reader to [2, 11].

Syntax. A *signature* Σ is a triple (S, F, P) where S is a set of *sorts*, F is a set of *function symbols*, and P is a set of *predicate symbols*. For a signature $\Sigma = (S, F, P)$, we write $\Sigma^{\mathbb{S}}$ for the set S of sorts, $\Sigma^{\mathbb{F}}$ for the set F of function symbols, and $\Sigma^{\mathbb{P}}$ for the set P of predicates. Each predicate and function symbol is associated with an *arity*, a tuple constructed from the sorts in S . We write $\Sigma_1 \cup \Sigma_2 = (S_1 \cup S_2, F_1 \cup F_2, P_1 \cup P_2)$ for the union¹ of signatures $\Sigma_1 = (S_1, F_1, P_1)$ and $\Sigma_2 = (S_2, F_2, P_2)$. Additionally, we write $\Sigma_1 \subseteq \Sigma_2$ if $S_1 \subseteq S_2$, $F_1 \subseteq F_2$, $P_1 \subseteq P_2$, and the symbols of Σ_1 have the same arity as those in Σ_2 .

For a signature Σ , we assume the logic (but not the signature) includes an equality symbol $=_{\sigma}$, for each sort $\sigma \in \Sigma^{\mathbb{S}}$. We will frequently omit the subscript on equality when the sort of the equation is not relevant to the discussion. We assume the standard notions of a Σ -*term*, Σ -*literal*, and Σ -*formula*. In the following, we assume that all formulas are quantifier-free, if not explicitly stated otherwise. A literal is called *flat* if it is of the form $x = y$, $x \neq y$, $x = f(y_1, \dots, y_n)$, $p(y_1, \dots, y_n)$, or $\neg p(y_1, \dots, y_n)$, where x, y, y_1, \dots, y_n are variables, f is a function symbol, and p is a predicate symbol.

If ϕ is a term or a formula, we will denote by $\text{vars}_{\sigma}(\phi)$ the set of variables of sort σ that occur (free) in ϕ . We overload this function in the usual way, $\text{vars}_S(\phi)$ denoting variables in ϕ of the sorts in S , and $\text{vars}(\phi)$ denoting all variables in ϕ . We also sometimes refer to a set Φ of formulas as if it were a single formula, in which case the intended meaning is the conjunction $\bigwedge \Phi$ of the formulas in the set.

Semantics. Let Σ be a signature, and let X be a set of variables whose sorts are in $\Sigma^{\mathbb{S}}$. A Σ -*interpretation* \mathcal{A} over X is a map that interprets: each sort $\sigma \in \Sigma^{\mathbb{S}}$ as a non-empty domain A_{σ} ,² each variable $x \in X$ of sort σ as an element $x^{\mathcal{A}} \in A_{\sigma}$, each function symbol $f \in \Sigma^{\mathbb{F}}$ of arity $\sigma_1 \times \dots \times \sigma_n \times \tau$ as a function $f^{\mathcal{A}} : A_{\sigma_1} \times \dots \times A_{\sigma_n} \rightarrow A_{\tau}$, each predicate symbol $p \in \Sigma^{\mathbb{P}}$ of arity $\sigma_1 \times \dots \times \sigma_n$ as a subset $p^{\mathcal{A}}$ of $A_{\sigma_1} \times \dots \times A_{\sigma_n}$. A Σ -*structure* is a Σ -interpretation over an empty set of variables. As usual, the interpretations of terms and formulas in an interpretation \mathcal{A} are defined inductively over their structure (with equality, Boolean operations, and quantifiers interpreted as usual). For a term t , we denote with $t^{\mathcal{A}}$ the evaluation of t under the interpretation \mathcal{A} . Likewise, for a formula ϕ , we denote with $\phi^{\mathcal{A}}$ the truth-value (true or false) of ϕ under interpretation \mathcal{A} . A Σ -formula ϕ is *satisfiable* iff it evaluates to true in some Σ -interpretation over $\text{vars}(\phi)$.

Given a Σ -interpretation \mathcal{A} , a vector of variables \vec{x} , and a vector of domain elements of \mathcal{A} , \vec{a} , we denote by $\mathcal{A}\{\vec{x} \leftarrow \vec{a}\}$ the Σ -interpretation with the same domains as \mathcal{A} that interprets each variable in \vec{x} as the corresponding element

¹ In this paper, when combining two signatures, we always assume that function and predicate symbols from the signatures do not overlap, so that the union operation is well-defined. On the other hand, the signatures are allowed to have non-disjoint sets of sorts.

² In the rest of the paper we will use the calligraphic letters $\mathcal{A}, \mathcal{B}, \dots$ to denote interpretations, and the corresponding subscripted Roman letters $A_{\sigma}, B_{\sigma}, \dots$ to denote the domains of the interpretations.

in \vec{a} and all other symbols as in \mathcal{A} (note that to be well-defined, we require that for each corresponding pair (x_i, a_i) in \vec{x} and \vec{a} , we must have $a_i \in A_{\sigma_i}$ where σ_i is the sort of x_i).

Let \mathcal{A} be an Ω -interpretation over some set V of variables. For a signature $\Sigma \subseteq \Omega$, and a set of variables $U \subseteq V$, we denote with $\mathcal{A}^{\Sigma, U}$ the interpretation obtained from \mathcal{A} by restricting it to interpret only the symbols in Σ and the variables in U .

Theories. We will use the definition of theories as classes of structures, rather than sets of sentences. We define a theory formally as follows (see e.g. [12] and Definition 2 in [8]).

Definition 1 (Theory). *Given a set of Σ -sentences \mathbf{Ax} a Σ -theory $T_{\mathbf{Ax}}$ is a pair (Σ, \mathbf{A}) where Σ is a signature and \mathbf{A} is the class of Σ -structures that satisfy \mathbf{Ax} .*

Given a theory $T = (\Sigma, \mathbf{A})$, a T -interpretation is a Σ -interpretation \mathcal{A} such that $\mathcal{A}^{\Sigma, \emptyset} \in \mathbf{A}$. A Σ -formula ϕ is T -satisfiable iff it is satisfiable in some T -interpretation \mathcal{A} . This is denoted as $\mathcal{A} \models_T \phi$, or just $\mathcal{A} \models \phi$ if the theory is clear from the context. Given a Σ -theory T , two Σ -formulas ϕ and ψ are T -equivalent if they evaluate to the same truth value in every T -interpretation.

2.2 Combination of Theories

As theories in our formalism are represented by classes of structures, a combination of two theories is represented by those structures that can interpret both theories (Definition 3 in [8]).

Definition 2 (Combination). *Let $T_1 = (\Sigma_1, \mathbf{A}_1)$ and $T_2 = (\Sigma_2, \mathbf{A}_2)$ be two theories. The combination of T_1 and T_2 is the theory $T_1 \oplus T_2 = (\Sigma, \mathbf{A})$ where $\Sigma = \Sigma_1 \cup \Sigma_2$ and $\mathbf{A} = \{\Sigma\text{-structures } \mathcal{A} \mid \mathcal{A}^{\Sigma_1, \emptyset} \in \mathbf{A}_1 \text{ and } \mathcal{A}^{\Sigma_2, \emptyset} \in \mathbf{A}_2\}$.*

The set of Σ -structures resulting from the combination of two theories is indeed a theory in the sense of Definition 1. If \mathbf{Ax}_1 is the set of sentences defining theory T_1 , and \mathbf{Ax}_2 is the set of sentences defining theory T_2 , then \mathbf{A} is the set of Σ -structures that satisfy the set $\mathbf{Ax} = \mathbf{Ax}_1 \cup \mathbf{Ax}_2$ (see Proposition 4 in [8]).

Given decision procedures for the satisfiability of formulas in theories T_1 and T_2 , we are interested in constructing a decision procedure for satisfiability in $T_1 \oplus T_2$ using as black boxes the known procedures for T_1 and T_2 . The Nelson-Oppen combination method [5, 10, 11] gives a general mechanism for doing this. Given a formula ϕ over the combined signature $\Sigma_1 \cup \Sigma_2$, the first step is to *purify* ϕ by constructing an equisatisfiable set of formulas $\phi_1 \cup \phi_2$ such that each ϕ_i consists of only Σ_i -formulas. This can easily be done by finding a pure (i.e. Σ_i -for some i) subterm t , replacing it with a new variable v , adding the equation $v = t$, and then repeating this process until all formulas are pure. The next step is to force the decision procedures for the individual theories to agree on whether

variables appearing in both ϕ_1 and ϕ_2 (called *shared* variables) are equal. This is done by introducing an *arrangement* over the shared variables [8, 10].

Definition 3 (Arrangement). *Given a set of variables V over a set of sorts S , with $V_\sigma = \text{vars}_\sigma(V)$ so that $V = \bigcup_{\sigma \in S} V_\sigma$, we call a formula δ_V an arrangement of V if there exists a family of equivalence relations $E = \{ E_\sigma \subseteq V_\sigma \times V_\sigma \mid \sigma \in S \}$, such that the equivalence relations induce δ_V , i.e. $\delta_V = \bigwedge_{\sigma \in S} \delta_\sigma$, where each δ_σ is determined by E_σ as follows:*

$$\delta_\sigma = \bigwedge_{(x,y) \in E_\sigma} (x = y) \wedge \bigwedge_{(x,y) \in \overline{E}_\sigma} (x \neq y) .$$

In the above definition, \overline{E}_σ denotes the complement of the equivalence relation E_σ , i.e. $V_\sigma \times V_\sigma \setminus E_\sigma$. When the family of equivalence relations is not clear from the context, we will denote the arrangement as $\delta_V(E)$.

The Nelson-Oppen method is only complete when the theories satisfy certain conditions. Sufficient conditions for completeness are signature-disjointness and *stable-infiniteness*. Stable-infiniteness was originally introduced in a single-sorted setting [6]. In the many-sorted setting, stable-infiniteness is defined with respect to a subset of the signature sorts (Definition 6 from [11]).

Definition 4 (Stable-Infiniteness). *Let Σ be a signature, let $S \subseteq \Sigma^{\mathbb{S}}$ be a set of sorts, and let T be a Σ -theory. We say that T is stably-infinite with respect to S if for every T -satisfiable quantifier-free Σ -formula ϕ , there exists a T -interpretation \mathcal{A} satisfying ϕ , such that A_σ is infinite for each sort $\sigma \in S$.*

The Nelson-Oppen combination theorem states that, given two theories T_1 and T_2 , stably-infinite over (at least) the set of common sorts $\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}}$, and whose signatures are otherwise disjoint, ϕ is satisfiable in $T_1 \oplus T_2$ iff there exists an arrangement δ_V of the shared variables $V = \text{vars}(\phi_1) \cap \text{vars}(\phi_2)$ such that $\phi_i \cup \delta_V$ is satisfiable in T_i , for $i = 1, 2$.

It is interesting to note that stable-infiniteness is preserved when combining theories, a fact that follows easily from known results.

Proposition 1. *Let Σ_1 and Σ_2 be signatures. If*

- T_1 is a Σ_1 -theory stably-infinite with respect to $S_1 \subseteq \Sigma_1^{\mathbb{S}}$,
- T_2 is a Σ_2 -theory stably-infinite with respect to $S_2 \subseteq \Sigma_2^{\mathbb{S}}$,
- $\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}} = S_1 \cap S_2$,

then $T_1 \oplus T_2$ is a $(\Sigma_1 \cup \Sigma_2)$ -theory and is stably-infinite with respect to $S_1 \cup S_2$.

Although many interesting theories are stably-infinite, some important theories are not. For example, the theory of fixed-size bit-vectors contains sorts whose domains are all finite. Hence, this theory cannot be stably-infinite. The Nelson-Oppen method may be incomplete for combinations involving this theory as shown by the following example.

Example 1. Consider the theory of arrays T_{array} where both indices and elements are of the same sort bv , so that the sorts of T_{array} are $\{\text{array}, \text{bv}\}$, and a theory T_{bv} that requires the sort bv to be interpreted as bit-vectors of size 1. Both theories are decidable and we would like to decide the combination theory in a Nelson-Oppen-like framework. Let a_1, \dots, a_5 be array variables and consider the following constraints:

$$a_i \neq a_j, \text{ for } 1 \leq i < j \leq 5 .$$

These constraints are entirely within the language of T_{array} (i.e. no purification is necessary), there are no shared variables, and there are no constraints over bit-vectors. Thus, the array theory decision procedure is given all of the constraints and the bit-vector decision procedure is given an empty set of constraints. Any decision procedure for the theory of arrays will tell us that these constraints are satisfiable. But, there are only four possible different arrays with elements and indices over bit-vectors of size 1, so this set of constraints is unsatisfiable.

The notion of politeness, which we define in the following section allows us to overcome this problem.

3 Polite Theories

Polite theories were introduced in [8] to extend the Nelson-Oppen method to allow combinations with non-stably-infinite theories. A theory can be combined with any other theory (with no common function or predicate symbols) if it is *polite* with respect to the set of shared sorts. The notion of politeness depends on two other important properties: *smoothness* and *finite witnessability*. In this section, we define these terms, noting that our definition of finite witnessability differs slightly from that given in [8] in order to fix a correctness problem in that paper (as we explain below). We then give a new theorem showing that the combination of two theories preserves politeness with respect to some of the sorts.

3.1 Definitions

First we define the smoothness property of a theory (Definition 7 from [8]).

Definition 5 (Smoothness). *Let Σ be a signature, let $S \subseteq \Sigma^{\mathbb{S}}$ be a set of sorts, and let T be a Σ -theory. We say that T is smooth with respect to S if:*

- for every T -satisfiable quantifier-free Σ -formula ϕ ,
- for every T -interpretation \mathcal{A} satisfying ϕ ,
- for all choices of cardinal numbers κ_σ , such that $\kappa_\sigma \geq |A_\sigma|$ for all $\sigma \in S$,

there exists a T -interpretation \mathcal{B} satisfying ϕ such that $|B_\sigma| = \kappa_\sigma$, for all $\sigma \in S$.

Recall that when a theory T is stably-infinite with respect to a sort σ and a T -interpretation exists, we can always find another T -interpretation in which the domain of σ is infinite. On the other hand, if T is smooth with respect to σ and we have a T -interpretation, then there exist interpretations in which the domain of σ can be chosen to be any larger size. Hence every theory that is smooth with respect to a set of sorts S is also stably-infinite with respect to S .

Being able to combine two interpretations from different theories mainly depends on the ability to bring the domains of the shared sorts to the same size. This is where stable-infiniteness helps in the Nelson-Oppen framework: it ensures that the domains of the shared sorts can have the same infinite cardinalities. Since we are interested in combining theories that may require finite domains, we need more flexibility than that afforded by stable-infiniteness. Smoothness gives us more flexibility in resizing structures upwards. This is not quite enough as we also need to ensure that the structures are small enough. Rather than attempting to resize structures downwards, we rely on the notion of *finite witnessability* which allows us to find a kind of “minimal” structure for a theory.

Definition 6 (Finite Witnessability). *Let Σ be a signature, let $S \subseteq \Sigma^{\mathbb{S}}$ be a set of sorts, and let T be a Σ -theory. We say that T is finitely witnessable with respect to S if there exists a computable function, witness, which, for every quantifier-free Σ -formula ϕ , returns a quantifier-free Σ -formula $\psi = \text{witness}(\phi)$ such that*

- ϕ and $(\exists \vec{w})\psi$ are T -equivalent, where $\vec{w} = \text{vars}(\psi) \setminus \text{vars}(\phi)$ are fresh variables;
- if $\psi \wedge \delta_V$ is T -satisfiable, for an arrangement δ_V , where V is a set of variables of sorts in S , then there exists a T -interpretation \mathcal{A} satisfying $\psi \wedge \delta_V$ such that $A_\sigma = [\text{vars}_\sigma(\psi \wedge \delta_V)]^{\mathcal{A}}$, for all $\sigma \in S$,

where the notation $[U]^{\mathcal{A}}$ indicates the set $\{v^{\mathcal{A}} \mid v \in U\}$.

Both of the definitions above use an arbitrary quantifier-free formula ϕ in the definition. As shown by Proposition 11 and Proposition 12 in [7], it is enough to restrict ourselves to conjunctions of flat literals in the definitions. This follows in a straightforward fashion from the fact that we can always construct an equisatisfiable formula in disjunctive normal form over flat literals.

It is important to note that our definition of finite witnessability differs from the definition given in [8]. Their definition is equivalent to ours except that there is no mention of an arrangement (i.e. the formula ψ appears alone everywhere $\psi \wedge \delta_V$ appears in the definition above). The reason for this is explained and illustrated in Section 3.2 below.

Finally, a theory that is both smooth and finitely witnessable is *polite* (Definition 9 in [8]).

Definition 7 (Politeness). *Let Σ be a signature, let $S \subseteq \Sigma^{\mathbb{S}}$ be a set of sorts, and let T be a Σ -theory. We say that T is polite with respect to S if it is both smooth and finitely witnessable with respect to S .*

Note that any theory is polite (stably-infinite, smooth, finitely witnessable) with respect to an empty set of sorts.

Example 2. The extensional theory of arrays T_{array} has a signature Σ_{array} that contains a sort `elem` for elements, a sort `index` for indices, and a sort `array` for arrays, as well as the two function symbols `read` : `array` \times `index` \mapsto `elem` and `write` : `array` \times `index` \times `elem` \mapsto `array`. Semantics of the array function symbols can be axiomatized as usual, and we refer the reader to [9] for more detail.

It is not hard to see that T_{array} is smooth with respect to the sorts $\{\text{index}, \text{elem}\}$ – any interpretation satisfying a quantifier-free formula ϕ can be extended to arbitrary cardinalities over indices and elements by adding as many additional indices and elements as we need while keeping the satisfiability of ϕ .

As for finite witnessability, it is enough to use a witness transformation that works over conjunctions of flat literals and replaces each array disequality $a \neq b$ with the conjunction of literals $e_1 = \text{read}(a, i) \wedge e_2 = \text{read}(b, i) \wedge e_1 \neq e_2$, where i is a fresh variable of sort `index` and e_1, e_2 are fresh variables of sort `elem`. The witness function creates a fresh witness index i , to witness the position where a and b are different, and names those different elements e_1 and e_2 . For the detailed proof of politeness for the theory T_{array} we refer the reader to [8].

3.2 Finite Witnessability Revisited

A main result of [8] is a combination method for two theories, one of which is polite over the shared sorts.

Proposition 2 (Proposition 12 of [8]). *Let T_i be a Σ_i -theory for $i = 1, 2$ such that the two theories have no function or predicate symbols in common. Assume that T_2 is polite with respect to $S = \Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}}$. Also, let Γ_i be a set of Σ_i literals for $i = 1, 2$, and let $\psi_2 = \text{witness}_{T_2}(\Gamma_2)$. Finally, let $V_\sigma = \text{vars}_\sigma(\psi_2)$, for each $\sigma \in S$, and let $V = \bigcup_{\sigma \in S} V_\sigma$. Then the following are equivalent:*

1. $\Gamma_1 \cup \Gamma_2$ is $(T_1 \oplus T_2)$ -satisfiable;
2. There exists an arrangement δ_V such that $\Gamma_1 \cup \delta_V$ is T_1 -satisfiable and $\{\psi_2\} \cup \delta_V$ is T_2 -satisfiable.

Proposition 2 differs from the standard Nelson-Oppen theorem in its application of the witness function to Γ_2 and in that the arrangement is over *all* the variables with shared sorts in ψ_2 rather than just over the shared variables.

As mentioned above, our definition of finite witnessability (Definition 6 above) differs from the definition given in [8]. Without the change, Proposition 2 does not hold, as demonstrated by the following example.

Example 3. Let Σ be a signature containing no function or predicate symbols and a single sort σ . Let T_1 be a Σ -theory containing all structures such that the domain of σ has exactly one element (i.e. the structures of T_1 are those satisfying $\forall x y. x = y$). Similarly, let T_2 be a Σ -theory over the same sort σ containing all structures such that the domain of σ has at least two elements (i.e. axiomatized

by $\exists x y. x \neq y$). Note that the combination of these two theories contains no structures, and hence no formula is satisfiable in $T_1 \oplus T_2$.

Theory T_2 is clearly smooth with respect to σ . To be polite, T_2 must also be finitely witnessable with respect to σ . Consider the following candidate witness function:

$$\text{witness}(\phi) \triangleq \phi \wedge w_1 = w_1 \wedge w_2 = w_2 \quad ,$$

where w_1 and w_2 are fresh variables of sort σ not appearing in ϕ .

Let ϕ be a conjunction of flat Σ -literals, let $\psi = \text{witness}(\phi)$, and let $V = \text{vars}(\psi)$. It is easy to see that the first condition for finite witnessability holds: ϕ is satisfied in a T_2 model iff $\exists w_1 w_2. \psi$ is. Now, consider the second condition according to [8] (i.e. without the arrangement). We must show that if ψ is T_2 -satisfiable (in interpretation \mathcal{B} , say), then there exists a T_2 -interpretation \mathcal{A} satisfying ψ such that $A_\sigma = [V]^\mathcal{A}$. The obvious candidate for \mathcal{A} is obtained by setting $A_\sigma = [V]^\mathcal{B}$ and by letting \mathcal{A} interpret only those variables in V (interpreting them as in \mathcal{B}). Clearly \mathcal{A} satisfies ψ . However, if $[V]^\mathcal{B}$ contains only one element, then \mathcal{A} is not a T_2 -interpretation. But in this case, we can always first modify the way variables are interpreted in \mathcal{B} to ensure that $w_2^\mathcal{B}$ is different from $w_1^\mathcal{B}$ (\mathcal{B} is a T_2 -interpretation, so B_σ must contain at least two different elements). Since w_2 does not appear in ϕ , this change cannot affect the satisfiability of ψ in \mathcal{B} . After making this change, $[V]^\mathcal{B}$ is guaranteed to contain at least two elements, so we can always construct \mathcal{A} as described above. Thus, the second condition for finite witnessability is satisfied and the candidate witness function is indeed a witness function according to [8].

As we will see below, however, this witness function leads to problems. Notice that according to the definition of finite witnessability in this paper, the candidate witness function is not acceptable. To see why, consider again the second condition. Let δ_V be an arrangement of V . According to our definition, we must show that if $\psi \wedge \delta_V$ is satisfied by T_2 -interpretation \mathcal{B} , then there exists a T_2 -interpretation \mathcal{A} satisfying $\psi \wedge \delta_V$ such that $A_\sigma = [V]^\mathcal{A}$. We can consider the same construction as above, but this time, the case when $[V]^\mathcal{B}$ contains only one element cannot be handled as before. This is because δ_V requires \mathcal{A} to preserve equalities and disequalities in V . In particular, δ_V may include $w_1 = w_2$. In this case, there is no way to construct an appropriate interpretation \mathcal{A} .

Now, we show what happens if the candidate witness function given above is allowed. Consider using Proposition 2 to check the satisfiability of $x = x$ (where x is a variable of sort σ). Although this is trivially satisfiable in any theory that has at least one structure, it is not satisfiable in $T_1 \oplus T_2$ since there are no structures to satisfy it. To apply the proposition we let $\Gamma_1 = \emptyset$, $\Gamma_2 = \{x = x\}$,

$$\psi_2 = \text{witness}(\Gamma_2) = (x = x \wedge w_1 = w_1 \wedge w_2 = w_2) \quad ,$$

and $V = \text{vars}(\psi_2) = \{x, w_1, w_2\}$. Proposition 2 allows us to choose an arrangement over the variables of V . Let $\delta_V = \{x = w_1, x = w_2, w_1 = w_2\}$ be an arrangement over the variables in V . It is easy to see that $\Gamma_1 \cup \delta_V$ is satisfiable

in a T_1 -interpretation \mathcal{A} and $\psi_2 \cup \delta_V$ is satisfiable in a T_2 -interpretation \mathcal{B} , where \mathcal{A} and \mathcal{B} interpret the domains and variables as follows:

$$\sigma^{\mathcal{A}} = \{a_1\}, \sigma^{\mathcal{B}} = \{b_1, b_2\}, x^{\mathcal{A}} = w_1^{\mathcal{A}} = w_2^{\mathcal{A}} = a_1, x^{\mathcal{B}} = w_1^{\mathcal{B}} = w_2^{\mathcal{B}} = b_1 .$$

Thus, according to Proposition 2, $T_1 \cup T_2$ should be $T_1 \oplus T_2$ -satisfiable, but we know that this is impossible. Finally, consider what happens if we use a witness function for T_2 that is acceptable according to our new definition:

$$\text{witness}(\phi) \triangleq \phi \wedge w_1 \neq w_2 .$$

If we look at the same example using this witness function, we can verify that for every arrangement δ_V , either $w_1 \neq w_2 \in \delta_V$, in which case $T_1 \cup \delta_V$ is not T_1 -satisfiable, or else $w_1 = w_2 \in \delta_V$, in which case $\text{witness}(T_2) \cup \delta_V$ is not T_2 -satisfiable.

As shown by the example above, the definition of finite witnessability in [8] is not strong enough. It allows witness functions that can falsify Proposition 2. The changes in Definition 6 remedy the problem.

In the same paper, the authors also prove that a number of theories are polite. We are confident that the proofs of politeness for the theories of equality, arrays, sets, and multi-sets are still correct, given the new definition. Other results in the paper (in particular the proof of politeness for the theory of lists and the proof that shiny theories are polite) have some problems in their current form. We hope to address these in future work.

3.3 A New Combination Theorem for Polite Theories

Proposition 2 shows how to combine two theories, one of which is polite. However, the theorem tells us nothing about the politeness of the resulting (combined) theory. In particular, if we want to combine more than two theories by iterating the combination method, we cannot assume that the result of applying Proposition 2 is a theory that is polite with respect to any (non-empty) set of sorts.

In this section, we show that the combination described in Proposition 2 does preserve politeness over some of the sorts. This lays the foundation for the more general combination theorem described in Section 3.4.

Theorem 1. *Let Σ_1 and Σ_2 be signatures and let $S = \Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}}$. If*

1. T_1 is a Σ_1 -theory polite with respect to $S_1 \subseteq \Sigma_1^{\mathbb{S}}$,
2. T_2 is a Σ_2 -theory polite with respect to $S_2 \subseteq \Sigma_2^{\mathbb{S}}$,
3. $S \subseteq S_2$,

then $T_1 \oplus T_2$ is polite with respect to $S^ = S_1 \cup (S_2 \setminus \Sigma_1^{\mathbb{S}})$.*

We illustrate the application of the theorem with an example using two theories of arrays.




Fig. 1. Diagram for Theorem 1.

Example 4. Let $T_{\text{array},1}$ and $T_{\text{array},2}$ be two theories of arrays over the sets of sorts $S_1 = \{\text{array}_1, \text{index}_1, \text{elem}_1\}$ and $S_2 = \{\text{array}_2, \text{index}_2, \text{array}_1\}$ respectively. These two theories together model two-dimensional arrays with indices in index_1 and index_2 , and elements in elem_1 .

We know that the theory $T_{\text{array},1}$ is polite with respect to $S_1^* = \{\text{index}_1, \text{elem}_1\}$, and the theory $T_{\text{array},2}$ is polite with respect to $S_2^* = \{\text{index}_2, \text{array}_1\}$. Using Theorem 1, we know that we can combine them into a theory T_{array} that is polite with respect to the set $S_1^* \cup (S_2^* \setminus \{\text{array}_1\}) = \{\text{index}_1, \text{index}_2, \text{elem}_1\}$. This means that we can combine the theory of two-dimensional arrays with any other theories that operate over the elements and indices, even if they are not stably-infinite (such as bit-vectors for example).

An interesting corollary of Theorem 1 is that, if both theories are polite with respect to the shared sorts then, analogously to Proposition 1, we get a theory that is polite with respect to the union of the sorts.

Corollary 1. *Let Σ_1 and Σ_2 be signatures. If*

- T_1 is a Σ_1 -theory polite with respect to $S_1 \subseteq \Sigma_1^S$,
- T_2 is a Σ_2 -theory polite with respect to $S_2 \subseteq \Sigma_2^S$,
- $\Sigma_1^S \cap \Sigma_2^S = S_1 \cap S_2$,

then $T_1 \oplus T_2$ is polite with respect to $S_1 \cup S_2$.

3.4 Combining Multiple Polite Theories

Now we give a general theorem for combining multiple theories in a sequential manner.

Theorem 2. *Let T_i be a Σ_i -theory, for $1 \leq i \leq n$. Assume that*

- theories T_i have no function or predicate symbols in common;
- the quantifier-free satisfiability problem of T_i is decidable, for $1 \leq i \leq n$;
- T_i is polite with respect to S_i , for $1 \leq i \leq n$;
- $\Sigma_i^S \cap \Sigma_j^S \subseteq S_j$, for $1 \leq i < j \leq n$.

Then the quantifier-free satisfiability problem for $T = T_1 \oplus \dots \oplus T_n$ is decidable. Moreover, the resulting theory T is polite with respect to the set of sorts $S = \cup_{j=1}^n (S_j \setminus (\cup_{i < j} \Sigma_i^{\mathbb{S}}))$.

Example 5. Assume we have a theory of arrays $T_{\text{array},1}$ over the sorts $\Sigma_{\text{array},1}^{\mathbb{S}} = \{\text{array}_1, \text{index}_1, \text{elem}\}$, as well as theories of arrays $T_{\text{array},k}$ over the sorts $\Sigma_{\text{array},k}^{\mathbb{S}} = \{\text{array}_k, \text{index}_k, \text{array}_{k-1}\}$, for $k \geq 2$. These theories represent different layers in the theory of n -dimensional arrays. The theories satisfy the assumption of Theorem 2 and thus we can combine them into the full theory $T_{\text{array}} = T_{\text{array},1} \oplus \dots \oplus T_{\text{array},n}$. This theory is polite with respect to the union of all indices and elements $S = \{\text{index}_1, \text{index}_2, \dots, \text{index}_n, \text{elem}\}$.

Note that, although we are combining theories in a straightforward fashion, we could not have used Theorem 14 from [7] to achieve this combination, since the common intersection of the polite sets of sorts is empty, and the pairwise intersection of sorts is not. More importantly, we are able to easily deduce the politeness of the resulting theory. We finish this section with a theorem that gives an easy complete method for checking whether we can combine a set of theories in the framework of multiple polite theories.

Theorem 3. *Let T_1, T_2, \dots, T_n be pairwise signature-disjoint theories such that individual quantifier-free T_i -satisfiability problems are decidable. The quantifier-free satisfiability problem of $T = T_1 \oplus \dots \oplus T_n$ is decidable by iterating the polite combination method for two theories if and only if there is a reordering of the theories T_i that satisfies the conditions of Theorem 2.*

4 Theory Instantiations

The way theories are defined in Definition 1 is meant to be general, i.e. the sorts can be interpreted in any domain. But, sometimes we are interested in a variant of a theory obtained by identifying some of the sorts. For example, consider a theory of arrays with elements and indices, i.e. $\Sigma_{\text{array}}^{\mathbb{S}} = \{\text{array}, \text{elem}, \text{index}\}$. In practice, we often deal with a closely related theory of arrays in which the indices and the elements are from the same sort. Note that these two theories are indeed different – in the general theory of arrays, the well-sortedness prevents us from comparing indices with elements (the term $\text{read}(a, i) \neq i$ is not well-sorted, for example). We will call this merging of sorts *theory instantiation by sort equality*.

Definition 8 (Signature Instantiation). *Let $\Sigma = (S, F, P)$ be a signature. We call $\Sigma_s^{\sigma_1 = \sigma_2} = (S', F', P')$ a signature instantiation by sort equality $\sigma_1 = \sigma_2$, for sorts $\sigma_1, \sigma_2 \in S$ and $s \notin S$, if the following holds:*

- $S' = (S \setminus \{\sigma_1, \sigma_2\}) \cup \{s\}$;
- F' contains the same function symbols as F except that we replace σ_1 and σ_2 with s in every arity;
- P' contains the same predicate symbols as P except that we replace σ_1 and σ_2 with s in every arity.

To enable the translation of formulas from the instantiated signature to the original signature and vice versa, we will use the satisfiability-preserving syntactic formula transformation α that maps conjunctions of flat $\Sigma_s^{\sigma_1=\sigma_2}$ -literals into formulas from the signature Σ . Given such a conjunction $\phi = \bigwedge_{1 \leq k \leq m} l_k$, with $\text{vars}_s(\phi) = \{v_1, v_2, \dots, v_n\}$, we first introduce fresh variables $v_i^{\sigma_1}$ of sort σ_1 , and $v_i^{\sigma_2}$ of sort σ_2 , for $i = 1, \dots, n$. The function α transforms the formula ϕ into

$$\alpha(\phi) \triangleq \bigwedge_{1 \leq k \leq m} \alpha_l(l_k) \wedge \bigwedge_{1 \leq i < j \leq n} (v_i^{\sigma_1} =_{\sigma_1} v_j^{\sigma_1} \leftrightarrow v_i^{\sigma_2} =_{\sigma_2} v_j^{\sigma_2}) \quad ,$$

The transformation α_l acts on the individual literals as follows:

- Literals of the form $x =_{\sigma} y$ and $x \neq_{\sigma} y$, where $\sigma \neq s$, are left unchanged.
- Literals of the form $x =_s y$ and $x \neq_s y$ are transformed into $x^{\sigma_1} =_{\sigma_1} y^{\sigma_1}$ and $x^{\sigma_1} \neq_{\sigma_1} y^{\sigma_1}$ respectively.³
- Literals of the form $x =_{\sigma} f(y_1, \dots, y_n)$, where $\sigma \neq s$, are transformed into $x =_{\sigma} f(y_1^*, \dots, y_n^*)$. The variables y_i^* are taken to comply with the original arity of f in Σ , i.e.

$$y_i^* = \begin{cases} y_i^{\sigma_1} & \text{if } y_i \text{ should be of sort } \sigma_1 \text{ in the arity of } f \text{ in } \Sigma, \\ y_i^{\sigma_2} & \text{if } y_i \text{ should be of sort } \sigma_2 \text{ in the arity of } f \text{ in } \Sigma, \\ y_i & \text{otherwise.} \end{cases}$$

- Literals of the form $x =_s f(y_1, \dots, y_n)$ are transformed into either $x^{\sigma_1} =_{\sigma_1} f(y_1^*, \dots, y_n^*)$ or $x^{\sigma_2} =_{\sigma_2} f(y_1^*, \dots, y_n^*)$, depending on the sort of the codomain of f in Σ .
- Literals of the form $p(y_1, \dots, y_n)$ and $\neg p(y_1, \dots, y_n)$ are transformed in a similar manner.

In the other direction, we define a transformation γ_V , where V is a set of variables of sort s , from Σ -formulas to $\Sigma_s^{\sigma_1=\sigma_2}$ -formulas, as follows

$$\gamma_V(\phi) = \phi \wedge \bigwedge_{v \in V} (v^{\sigma_1} = v \wedge v^{\sigma_2} = v) \quad .$$

In the new formula variables formerly of sort σ_1 or σ_2 are now of sort s .

Definition 9 (Theory Instantiation). *Let Σ be a signature and $T = (\Sigma, \mathbf{A})$ be a Σ -theory. We call a theory $T_s^{\sigma_1=\sigma_2} = (\Sigma_s^{\sigma_1=\sigma_2}, \mathbf{B})$ the theory instantiated by sort equality $\sigma_1 = \sigma_2$, for sorts $\sigma_1, \sigma_2 \in \Sigma^{\mathbb{S}}$ and $s \notin \Sigma^{\mathbb{S}}$, when $\mathbf{B} \in \mathbf{B}$ iff*

- there exists an $\mathcal{A} \in \mathbf{A}$ such that $B_s = A_{\sigma_1} = A_{\sigma_2}$, and $B_{\sigma} = A_{\sigma}$ for $\sigma \neq s$;
and
- all the predicate and function symbols in $\Sigma_s^{\sigma_1=\sigma_2}$ are interpreted in \mathbf{B} exactly the same as they are interpreted in \mathcal{A} .

³ The choice of σ_1 over σ_2 is arbitrary, as the right part of $\alpha(\phi)$ will force the same on the dual variables.

The above definition simply restricts the original theory structures to those in which the sorts σ_1 and σ_2 are interpreted by the same domain. The lemma below shows that the result, $T_s^{\sigma_1=\sigma_2}$, is indeed a theory.

Lemma 1. *Let T and \mathbf{B} be as in Definition 9, and let \mathbf{Ax} be the set of closed Σ -formulas that defines T . The class \mathbf{B} is exactly the set of $\Sigma_s^{\sigma_1=\sigma_2}$ -structures that satisfies the set of formulas $\gamma_\emptyset(\mathbf{Ax}) = \{\gamma_\emptyset(\phi) \mid \phi \in \mathbf{Ax}\}$.*

Our motivating example is the theory of arrays where we restrict the sorts `elem` and `index` to be equal to each other and to `bv`, i.e. we are interested in the theory $T_{\text{array}}^{\text{bv}} = (T_{\text{array}})_{\text{bv}}^{\text{elem}=\text{index}}$. We know that T_{array} is polite with respect to the sorts `elem` and `index`. We want to know whether it is also the case that $T_{\text{array}}^{\text{bv}}$ is polite with respect to the sort `bv`.

The main result of this section is to show that by merging two sorts σ_1 and σ_2 in a theory, we preserve the politeness of the theory: the new theory will be polite with respect to the same set of sorts as the original theory, modulo renaming of the instantiated sorts σ_1 and σ_2 .

Theorem 4. *Let Σ be a signature, $\sigma_1, \sigma_2 \in \Sigma^{\mathbb{S}}$, and $s \notin \Sigma^{\mathbb{S}}$. If Σ -theory T is polite with respect to S , where $\sigma_1, \sigma_2 \in S$ and $s \notin S$, then $T_s^{\sigma_1=\sigma_2}$ is polite with respect to $S' = S \setminus \{\sigma_1, \sigma_2\} \cup \{s\}$. Furthermore, if witness is a witness function for theory T , then an acceptable witness function for $T_s^{\sigma_1=\sigma_2}$ is*

$$\text{witness}_s^{\sigma_1=\sigma_2}(\phi) = (\gamma_{\text{vars}_s(\phi)} \circ \text{witness} \circ \alpha)(\phi) .$$

Example 6. Consider again example 5, i.e. we have a theory of arrays T_{array} that operates over the sorts $\Sigma^{\mathbb{S}} = \{\text{array}_1, \dots, \text{array}_n, \text{index}_1, \dots, \text{index}_n, \text{elem}_1\}$ and is polite with respect to the index and element sorts $\Sigma^{\mathbb{S}} = \{\text{index}_1, \dots, \text{index}_n, \text{elem}_1\}$. Using Theorem 4, we can now safely replace the sorts `index1`, `index2` and `elem1` with the sort of bit-vectors `bv`, obtaining a theory $T_{\text{array}(\text{bv})}$ of n -dimensional arrays where the elements and the indices are of the same bit-vector sort. This theory $T_{\text{array}(\text{bv})}$ of arrays over bit-vectors is polite with respect to the sort `bv`, and therefore we can safely combine it with the theory of bit-vectors T_{bv} .

Using the combination method for polite theories, we can therefore get a *sound and complete decision procedure for deciding the theory of n -dimensional arrays over bit-vectors*, given a decision procedure and witness function for the theory of arrays T_{array} and a decision procedure for the theory of bit-vectors T_{bv} .

5 Conclusion

One of the crucial issues in the development of verification systems is the problem of combining decision procedures. Nelson and Oppen laid the foundation for the most commonly used framework, but their approach is limited by the requirement that the theories involved be stably-infinite. In this paper we revisited the problem of modular combination of non-stably-infinite theories in a many-sorted setting, using the previously introduced [8] notion of polite theories. We corrected the definition of polite theories that made the combination

method incomplete. Then we gave several new results that can be used to construct new polite theories from existing ones. These results led to a general combination result for multiple polite theories. Our result is not only applicable to a broader class of theories, but also precisely describes the politeness of the resulting theory. In future work, we plan to investigate the politeness of other common theories including general theories of inductive data-types [1]. We also are interested in finding efficient witness functions that minimize the number of variables that need to be considered in the arrangement shared by all theories.

Acknowledgments. We would like to thank the anonymous reviewers as well as Cesare Tinelli who provided valuable feedback on this work.

References

1. C. Barrett, I. Shikanian, and C. Tinelli. An abstract decision procedure for a theory of inductive data types. *Journal on Satisfiability, Boolean Modeling and Computation*, 3:21–46, 2007.
2. H. B. Enderton. *A mathematical introduction to logic*. Academic press New York, 1972.
3. D. Jovanović and C. Barrett. Polite theories revisited. Technical Report TR2010-922, Department of Computer Science, New York University, Jan. 2010.
4. S. Krstić, A. Goel, J. Grundy, and C. Tinelli. Combined Satisfiability Modulo Parametric Theories. In *TACAS 2007*, volume 4424 of *LNCS*, pages 602–617. Springer, 2007.
5. G. Nelson and D. C. Oppen. Simplification by cooperating decision procedures. *ACM Transactions on Programming Languages and Systems*, 1(2):245–257, 1979.
6. D. C. Oppen. Complexity, convexity and combinations of theories. *Theoretical Computer Science*, 12(3):291–302, 1980.
7. S. Ranise, C. Ringeissen, and C. Zarba. Combining Data Structures with Nonstably Infinite Theories using Many-Sorted Logic. Research Report RR-5678, INRIA, 2005.
8. S. Ranise, C. Ringeissen, and C. G. Zarba. Combining Data Structures with Nonstably Infinite Theories Using Many-Sorted Logic. In *FroCoS 2005*, volume 3717 of *LNCS*, pages 48–64. Springer, 2005.
9. A. Stump, D. L. Dill, C. W. Barrett, and J. Levitt. A decision procedure for an extensional theory of arrays. In *Proceedings of the 16th IEEE Symposium on Logic in Computer Science (LICS '01)*, pages 29–37. IEEE Computer Society, June 2001. Boston, Massachusetts.
10. C. Tinelli and M. T. Harandi. A new correctness proof of the Nelson–Oppen combination procedure. In *Frontiers of Combining Systems, Applied Logic*, pages 103–120. Kluwer Academic Publishers, 1996.
11. C. Tinelli and C. Zarba. Combining decision procedures for sorted theories. In *9th European Conference on Logic in Artificial Intelligence (JELIA '04)*, volume 3229 of *LNAI*, pages 641–653. Springer, 2004.
12. C. Tinelli and C. Zarba. Combining decision procedures for theories in sorted logics. Technical Report 04-01, Department of Computer Science, The University of Iowa, Feb. 2004.
13. C. Tinelli and C. G. Zarba. Combining nonstably infinite theories. *Journal of Automated Reasoning*, 34(3):209–238, 2005.