Polite Theories Revisited*

Dejan Jovanović and Clark Barrett
New York University
dejan@cs.nyu.edu, barrett@cs.nyu.edu

January 26, 2010

New York University Technical Report: TR2010-922

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 [4] 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 [11]. 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.

A related approach for combining theories is presented in [3]. 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

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

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 [7]. 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 [7]. 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 5 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 4 discusses the combination of multiple polite theories, culminating in a combination result that is more general than the ones presented in [7].

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, 9].

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 union[1] 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

[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.

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 $x = y$, $x \neq y$, $x = f(y_1, \ldots, y_n)$, $p(y_1, \ldots, y_n)$, or $\neg p(y_1, \ldots, y_n)$, where $x, y, y_1, \ldots, y_n$ 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 $vars_\sigma(\phi)$ the set of variables of sort σ that occur (free) in ϕ. We overload this function in the usual way, $vars_S(\phi)$ denoting variables in ϕ of the sorts in S, and $vars(\phi)$ 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 $\sigma_1 \times \cdots \times \sigma_n \times \tau$ as a function $f^{\mathcal{A}} : A_{\sigma_1} \times \cdots \times A_{\sigma_n} \to A_\tau$,
- each predicate symbol $p \in \Sigma^{\mathbb{P}}$ of arity $\sigma_1 \times \cdots \times \sigma_n$ as a subset $p^{\mathcal{A}}$ of $A_{\sigma_1} \times \cdots \times A_{\sigma_n}$.

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 $vars(\phi)$.

Given a Σ-interpretation $\mathcal{A}$, a vector of variables $\overrightarrow{x}$, and a vector of domain elements of $\mathcal{A}$, $\overrightarrow{a}$, we denote by $\mathcal{A}\{\overrightarrow{x} \leftarrow \overrightarrow{a}\}$ the Σ-interpretation with the same domains as $\mathcal{A}$ that interprets each variable in $\overrightarrow{x}$ as the corresponding element in $\overrightarrow{a}$ 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 $\overrightarrow{x}$ and $\overrightarrow{a}$, 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. [10] and Definition 2 in [7]).

Definition 2.1 (Theory)**.** *Given a set of* Σ*-sentences* $\mathbf{Ax}$ *a* Σ-theory $T_{\mathbf{Ax}}$ *is a pair* $(\Sigma, \mathbf{A})$ *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 $\mathcal{A} \vDash_T \phi$, or just $\mathcal{A} \vDash \phi$ 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] In the rest of the paper we will use the calligraphic letters $\mathcal{A}, \mathcal{B}, \ldots$ to denote interpretations, and the corresponding subscripted Roman letters A_σ, B_σ, ... to denote the domains of the interpretations.

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 [7]).

Definition 2.2 (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* $\mathbf{A} = \{\Sigma$*-structures* $\mathcal{A} \mid \mathcal{A}^{\Sigma_1,\emptyset} \in \mathbf{A_1}$ *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 2.1. 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 [7]).

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 [4, 8, 9] 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 [7, 8].

Definition 2.3 (Arrangement)**.** *Given a set of variables* V *over a set of sorts* S*, with* $V_\sigma = vars_\sigma(V)$ *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, $\overline{E}_\sigma$ 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 combination theorem states that ϕ is satisfiable in $T_1 \oplus T_2$ iff there exists an arrangement δ_V of the shared variables $V = vars(\phi_1) \cap vars(\phi_2)$ such that $\phi_i \cup \delta_V$ is satisfiable in T_i, for $i = 1, 2$. However, as mentioned earlier, some restrictions on the theories are necessary in order for the Nelson-Oppen method to be complete. Sufficient conditions for completeness are:

- the two signatures have no function or predicate symbols in common, and
- the two theories are *stably-infinite* over (at least) the set of common sorts $\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}}$.

Stable-infiniteness was originally introduced in a single-sorted setting [5]. In the many-sorted setting, stable-infiniteness is defined with respect to a subset of the signature sorts (Definition 6 from [9]).

Definition 2.4 (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* $\sigma \in S$*.*

It is interesting to note that stable-infiniteness is preserved when combining theories, a fact that follows easily from known results. For completeness, we give the proof here. First, we need the following theorem, obtained by adapting Theorems 10 and 11 of [9].

Theorem 2.5. *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. Let $\Sigma = \Sigma_1 \cup \Sigma_2$, $T = T_1 \oplus T_2$, and let $S = \Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}}$ be the set of shared sorts. Let Γ_i be a set of Σ_i-literals for $i = 1, 2$, and let $V = vars(\Gamma_1) \cap vars(\Gamma_2)$ be the set of variables shared between Γ_1 and Γ_2. If there exists a T_1-interpretation $\mathcal{A}$ and a T_2-interpretation $\mathcal{B}$ and an arrangement δ_V of V such that:*

- $\mathcal{A} \models \Gamma_1 \cup \delta_V$,
- $\mathcal{B} \models \Gamma_2 \cup \delta_V$, *and*
- $|A_\sigma| = |B_\sigma|$, *for all $\sigma \in S$,*

then there exists a T-interpretation $\mathcal{C}$ such that:

- $\mathcal{C} \models \Gamma_1 \cup \Gamma_2 \cup \delta_V$,
- $C_\sigma = A_\sigma$ *for all $\sigma \in \Sigma_1^{\mathbb{S}}$, and*
- $C_\sigma = B_\sigma$ *for all $\sigma \in \Sigma_2^{\mathbb{S}} \setminus S$.*

Proof. Let $V_\sigma = vars_\sigma(\Gamma_1) \cap vars_\sigma(\Gamma_2)$, for $\sigma \in S$. Define a family of functions $h = \{h_\sigma : V_\sigma^{\mathcal{B}} \mapsto V_\sigma^{\mathcal{A}} \mid \sigma \in S\}$ such that $h_\sigma(v^{\mathcal{B}}) = v^{\mathcal{A}}$ for each $v \in V_\sigma$. Since the interpretations $\mathcal{B}$ and $\mathcal{A}$ agree on equalities over V (by satisfying the same arrangement δ_V), the functions h_σ are well-defined and bijective. This implies that $|V_\sigma^{\mathcal{B}}| = |V_\sigma^{\mathcal{A}}|$ and, since $|B_\sigma| = |A_\sigma|$ for $\sigma \in S$, we can extend each function to a bijection $h'_\sigma : B_\sigma \mapsto A_\sigma$. Let h'_σ be the identity function for each $\sigma \in \Sigma_2^{\mathbb{S}} \setminus S$. Now, we can define a new Σ_2-interpretation $\mathcal{B}'$ (over the same set of variables as in $\mathcal{B}$) in such a way that $h' = \bigcup_{\sigma \in S} h'_\sigma$ is an isomorphism from $\mathcal{B}$ to $\mathcal{B}'$:

$$
\begin{aligned}
B'_\sigma &= \begin{cases} A_\sigma & \text{if } \sigma \in S \\ B_\sigma & \text{if } \sigma \in \Sigma_2^{\mathbb{S}} \setminus S \end{cases} \\
v^{\mathcal{B}'} &= h'(v^{\mathcal{B}}) \\
f^{\mathcal{B}'}(b_1, \ldots, b_n) &= h'(f^{\mathcal{B}}(h'^{-1}(b_1), \ldots, h'^{-1}(b_n))), \text{ and} \\
\langle b_1, \ldots, b_n \rangle \in p^{\mathcal{B}'} &\text{ iff } \langle h'^{-1}(b_1), \ldots, h'^{-1}(b_n) \rangle \in p^{\mathcal{B}}.
\end{aligned}
$$

Because h' is an isomorphism, we have $\mathcal{B}' \models \Gamma_2 \cup \delta_V$. We can now define the Σ-interpretation $\mathcal{C}$ as follows:

$$
C_\sigma = \begin{cases} A_\sigma & \text{if } \sigma \in \Sigma_1^{\mathbb{S}} \setminus S \\ A_\sigma = B'_\sigma & \text{if } \sigma \in S \\ B'_\sigma = B_\sigma & \text{if } \sigma \in \Sigma_2^{\mathbb{S}} \setminus S \end{cases}
\qquad
v^{\mathcal{C}} = \begin{cases} v^{\mathcal{A}} & \text{if } v \text{ is of sort } \sigma \in \Sigma_1^{\mathbb{S}} \setminus S \\ v^{\mathcal{A}} = v^{\mathcal{B}'} & \text{if } v \text{ is of sort } \sigma \in S \\ v^{\mathcal{B}'} & \text{if } v \text{ is of sort } \sigma \in \Sigma_2^{\mathbb{S}} \setminus S \end{cases}
$$

$$
f^{\mathcal{C}} = \begin{cases} f^{\mathcal{A}} & \text{if } f \in \Sigma_1^F \\ f^{\mathcal{B}'} & \text{if } f \in \Sigma_2^F \end{cases}
\qquad
p^{\mathcal{C}} = \begin{cases} p^{\mathcal{A}} & \text{if } p \in \Sigma_1^P \\ p^{\mathcal{B}'} & \text{if } p \in \Sigma_2^P \end{cases}
$$

Clearly, $C_\sigma = A_\sigma$ for all $\sigma \in \Sigma_1^{\mathbb{S}}$ and $C_\sigma = B_\sigma$ for all $\sigma \in \Sigma_2^{\mathbb{S}} \setminus S$. It is also easy to see from the definition above that $\mathcal{C}^{\Sigma_1, vars(\Gamma_1)} = \mathcal{A}$ and $\mathcal{C}^{\Sigma_2, vars(\Gamma_2)} = \mathcal{B}'$, and thus $\mathcal{C} \models \Gamma_1 \cup \Gamma_2 \cup \delta_V$. □

Now we show that stable-infiniteness is preserved when combining theories.

Proposition 2.6. *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 $S_1 \cup S_2$.

Proof. Let $\Sigma = \Sigma_1 \cup \Sigma_2$, $T = T_1 \oplus T_2$, $S = S_1 \cap S_2$. Assume ϕ is a Σ-formula satisfied in a T-interpretation $\mathcal{A}$. As in the Nelson-Oppen algorithm, we can separate ϕ into the Σ_1-part ϕ_1 and the Σ_2-part ϕ_2. We can assume wlog that $\mathcal{A}$ is an interpretation over the variables $vars(\phi_1) \cup vars(\phi_2)$ and that $\mathcal{A} \models \phi_1 \wedge \phi_2$. Let $V = vars_S(\phi_1) \cap vars_S(\phi_2)$. Let δ_V be the arrangement over these variables that agrees with $\mathcal{A}$. We have that $\mathcal{A}^{\Sigma_1, vars(\phi_1)} \models \phi_1 \cup \delta_V$ and $\mathcal{A}^{\Sigma_2, vars(\phi_2)} \models \phi_2 \cup \delta_V$.

Since T_1 is stably-infinite with respect to S_1, we can conclude that there is a T_1-interpretation $\mathcal{B}$ such that the cardinalities $|B_\sigma|$ are infinite for $\sigma \in S_1$ and $\mathcal{B} \models \phi_1 \cup \delta_V$. Similarly, there is a T_2-interpretation $\mathcal{C}$ with infinite cardinalities $|C_\sigma|$, $\sigma \in S_2$ and $\mathcal{C} \models \phi_2 \cup \delta_V$. By the downward Löwenheim-Skolem theorem in the many-sorted setting,[3] we can assume that the cardinalities $|B_\sigma|$ and $|C_\sigma|$ are the same for $\sigma \in S$.

It is easy to see that $\mathcal{B}$ and $\mathcal{C}$ satisfy all the conditions of Theorem 2.5, so we can conclude that there is a T-interpretation $\mathcal{D}$ that satisfies $\phi_1 \wedge \phi_2$ (and hence ϕ) such that the cardinalities $|D_\sigma|$ are infinite for $\sigma \in S_1 \cup S_2$. □

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 2.7. *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 $\{\mathsf{array}, \mathsf{bv}\}$, 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 $a_1, \ldots, a_5$ 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] See Theorem 9 in [9] for the statement, or [10] for a full proof. The theorem there is in the more general setting of order-sorted logic, of which ordinary many-sorted logic is a simple instance.

3 Polite Theories

Polite theories were introduced in [7] 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 [7] 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 [7]).

Definition 3.1 (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 3.2 (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 $\psi = witness(\phi)$ such that*

- *ϕ and $(\exists \overrightarrow{w})\psi$ are T-equivalent, where $\overrightarrow{w} = vars(\psi) \setminus vars(\phi)$ 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 $A_\sigma = [vars_\sigma(\psi \wedge \delta_V)]^{\mathcal{A}}$, 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 [6] (see Lemmas A.1, A.2 in Appendix A), 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 [7]. 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).[4] 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 [7]).

Definition 3.3 (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 3.4. *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 following two function symbols.*

$$\mathsf{read} : \mathsf{array} \times \mathsf{index} \mapsto \mathsf{elem}$$
$$\mathsf{write} : \mathsf{array} \times \mathsf{index} \times \mathsf{elem} \mapsto \mathsf{array}$$

The theory T_{array} *is axiomatized by*

$$(\forall a : \mathsf{array})(\forall i : \mathsf{index})(\forall e : \mathsf{elem})(\mathsf{read}(\mathsf{write}(a, i, e), i) = e) \tag{1}$$
$$(\forall a : \mathsf{array})(\forall i, j : \mathsf{index})(\forall e : \mathsf{elem})(i \neq j \rightarrow \mathsf{read}(\mathsf{write}(a, i, e), j) = \mathsf{read}(a, j)) \tag{2}$$
$$(\forall a, b : \mathsf{array})\,[a \neq b \rightarrow (\exists i : \mathsf{index})(\mathsf{read}(a, i) \neq \mathsf{read}(b, i))] \tag{3}$$

It is not hard to see that T_{array} *is smooth with respect to the sorts* $\{\mathsf{index}, \mathsf{elem}\}$ *– 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*

$$e_1 = \mathsf{read}(a, i) \wedge e_2 = \mathsf{read}(b, i) \wedge e_1 \neq e_2 \ ,$$

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 [7].*

[4] It is worth noting that in order to prove Proposition 3.5 and Theorem 5.5, below, it is sufficient to require V to be equal to $vars_S(\psi_2)$ rather than letting it be an arbitrary set of variables with sorts in S. However, this more general flexibility is needed for the proofs of Lemma A.2 and Theorem 3.7.

3.2 Finite Witnessability Revisited

A main result of [7] is a combination method for two theories, one of which is polite over the shared sorts.

Proposition 3.5 (Proposition 12 of [7]). *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 $\psi_2 = witness_{T_2}(\Gamma_2)$. Finally, let $V_\sigma = vars_\sigma(\psi_2)$, 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 3.5 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 3.2 above) differs from the definition given in [7]. Without the change, Proposition 3.5 does not hold, as demonstrated by the following example.

Example 3.6. *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:

$$witness(\phi) \triangleq \phi \wedge w_1 = w_1 \wedge w_2 = w_2 \ ,$$

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 $\psi = witness(\phi)$, and let $V = vars(\psi)$. 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 [7] (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 [7].

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 3.5 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

- $\Gamma_1 = \emptyset, \Gamma_2 = \{x = x\}$,
- $\psi_2 = \mathit{witness}(\Gamma_2) = (x = x \wedge w_1 = w_1 \wedge w_2 = w_2)$, *and*
- $V = \mathit{vars}(\psi_2) = \{x, w_1, w_2\}$.

Proposition 3.5 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 3.5, $\Gamma_1 \cup \Gamma_2$ 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:

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

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 $\Gamma_1 \cup \delta_V$ is not T_1-satisfiable, or else $w_1 = w_2 \in \delta_V$, in which case $\mathit{witness}(\Gamma_2) \cup \delta_V$ is not T_2-satisfiable.

As shown by the example above, the definition of finite witnessability in [7] is not strong enough. It allows witness functions that can falsify Proposition 3.5. The changes in Definition 3.2 remedy the problem. For completeness, we include the proof of Proposition 3.5 below, adapted from [7], indicating where the proof fails if the weaker definition of finite witnessability is used.

Proof. $(1 \Rightarrow 2)$: Assume that $\Gamma_1 \cup \Gamma_2$ is $(T_1 \oplus T_2)$-satisfiable and let $\overrightarrow{v} = \mathit{vars}(\psi_2) \setminus \mathit{vars}(\Gamma_2)$. Since Γ_2 and $(\exists \overrightarrow{v})\psi_2$ are T_2-equivalent, it follows that $\Gamma_1 \cup \{\psi_2\}$ is also $(T_1 \oplus T_2)$-satisfiable. We can therefore fix a $(T_1 \oplus T_2)$-interpretation $\mathcal{A}$ satisfying $\Gamma_1 \cup \{\psi_2\}$. Next, let δ_V be the arrangement of V induced by $\mathcal{A}$, that is the arrangement determined by the equivalence classes $E_\sigma = \{(x, y) \mid x, y \in V_\sigma \text{ and } x^{\mathcal{A}} = y^{\mathcal{A}}\}$, for $\sigma \in S$. By construction we have an interpretation $\mathcal{A}$ such that both $\Gamma_1 \cup \delta_V$ is T_1 satisfied and $\{\psi_2\} \cup \delta_V$ is T_2-satisfied.

$(2 \Rightarrow 1)$: Let $\mathcal{A}$ be a T_1-interpretation satisfying $\Gamma_1 \cup \delta_V$, and let $\mathcal{B}$ be a T_2-interpretation satisfying $\{\psi_2\} \cup \delta_V$. Since T_2 is finitely witnessable, we can assume without loss of generality that

$B_\sigma = V_\sigma^{\mathcal{B}}$.[5] For each $\sigma \in S$ we now have that

$$\begin{aligned} |B_\sigma| &= |V_\sigma^{\mathcal{B}}| && \text{since } B_\sigma = V_\sigma^{\mathcal{B}} \ , \\ &= |V_\sigma^{\mathcal{A}}| && \text{since both } \mathcal{A} \text{ and } \mathcal{B} \text{ satisfy } \delta_V \ , \\ &\leq |A_\sigma| && \text{since } V_\sigma^{\mathcal{A}} \subseteq A_\sigma \ . \end{aligned}$$

Now we can use the smoothness of T_2 to obtain a T_2-interpretation $\mathcal{C}$ that satisfies $\{\psi_2\} \cup \delta_V$ such that $|C_\sigma| = |A_\sigma|$, for each $\sigma \in S$. Now we have all the conditions necessary to combine T_1-interpretation $\mathcal{A}$ and T_2-interpretation $\mathcal{C}$ into a $(T_1 \oplus T_2)$-interpretation $\mathcal{D}$, via Theorem 2.5,[6] $\mathcal{D}$ satisfies $\Gamma_1 \cup \{\psi_2\} \cup \delta_V$. Since Γ_2 and $(\exists \overrightarrow{v})\psi_2$ are T_2-equivalent, it follows that $\mathcal{D}$ also satisfies $\Gamma_1 \cup \Gamma_2$. □

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 3.5 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 3.5 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 3.5 does preserve politeness over some of the sorts. This lays the foundation for the more general combination theorem described in Section 4.

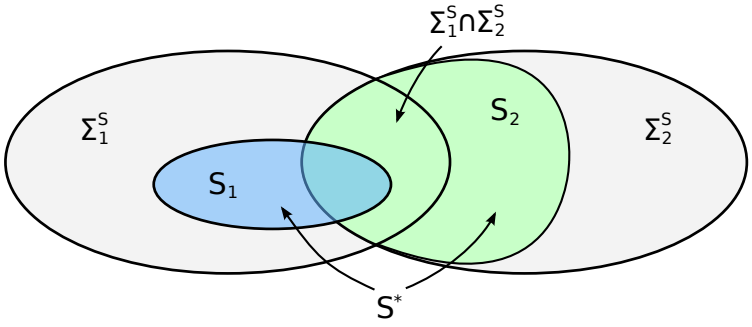

Figure 1: Diagram for Theorem 3.7.

Theorem 3.7. *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 $S_1 \subseteq \Sigma_1^{\mathbb{S}}$,*

[5] This is where the proof breaks with the original definition of finite witnessability–it is clear that in order to make this assumption and keep the satisfiability of $\{\psi_2\} \cup \delta_V$ we need to include δ_V in the definition of finite witnessability.

[6] Note that δ_V may contain more variables than those shared between Γ_1 and ψ_2, but we can still apply the theorem simply by assuming that δ_V is included among the literals of both theories.

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 $S^* = S_1 \cup \left(S_2 \setminus \Sigma_1^{\mathbb{S}}\right)$.

Proof. First we prove that the combined theory is smooth with respect to S^*. Let $T = T_1 \oplus T_2$, $\Sigma = \Sigma_1 \cup \Sigma_2$, and assume that ϕ is a conjunction of flat Σ-literals satisfiable in a T-interpretation $\mathcal{A}$. We are given cardinalities $\kappa_\sigma \geq |A_\sigma|$, for all $\sigma \in S^*$, and we must construct a new T-interpretation satisfying ϕ that obeys the given cardinalities.

We can separate ϕ into a Σ_1-part ϕ_1 and a Σ_2-part ϕ_2 such that $\phi = \phi_1 \wedge \phi_2$. Let $V_i = \mathit{vars}(\phi_i)$. We know that:

$$\mathcal{A}^{\Sigma_1, V_1} \models_{T_1} \phi_1 \ ,$$
$$\mathcal{A}^{\Sigma_2, V_2} \models_{T_2} \phi_2 \ .$$

Since T_2 is finitely witnessable, we know there is a witness function $\mathit{witness}_2$ such that ϕ_2 is T_2-equivalent to $(\exists \overrightarrow{w})\psi_2$ for $\psi_2 = \mathit{witness}_2(\phi_2)$. Since the variables in $\overrightarrow{w}$ are fresh, we can assume without loss of generality that $\mathcal{A} \models \psi_2$. If we then let $V_2' = V_2 \cup \mathit{vars}(\overrightarrow{w})$, we have:

$$\mathcal{A}^{\Sigma_1, V_1} \models_{T_1} \phi_1 \ ,$$
$$\mathcal{A}^{\Sigma_2, V_2'} \models_{T_2} \psi_2 \ .$$

Now, let $V_S = \mathit{vars}_S(\psi_2)$ and $V_{\overline{S}} = \mathit{vars}_{S_2 \setminus S}(\psi_2)$ and let δ_{V_S} and $\delta_{V_{\overline{S}}}$ be the (unique) arrangements of these sets of variables that are satisfied by $\mathcal{A}$. We can add these arrangements (letting $V_1' = V_1 \cup \mathit{vars}_S(\psi_2)$) to obtain:

$$\mathcal{A}^{\Sigma_1, V_1'} \models_{T_1} \phi_1 \wedge \delta_{V_S} \ ,$$
$$\mathcal{A}^{\Sigma_2, V_2'} \models_{T_2} \psi_2 \wedge \delta_{V_S} \wedge \delta_{V_{\overline{S}}} \ .$$

Because T_1 is smooth with respect to S_1, we can lift the domains of sorts $\sigma \in S_1$ to cardinalities κ_σ, i.e. there is a T_1-interpretation $\mathcal{B}$ that satisfies $\phi_1 \wedge \delta_{V_S}$ with $|B_\sigma| = \kappa_\sigma$ for $\kappa \in S_1$. We can't assume anything about the rest of the sorts, so let $\kappa'_\sigma = |B_\sigma|$ for $\sigma \in S \setminus S_1$.

Next, because T_2 is finitely witnessable with respect to S_2, we know that there is a T_2-interpretation $\mathcal{C}$ that satisfies $\psi_2 \wedge \delta_{V_S} \wedge \delta_{V_{\overline{S}}}$ such that for $\sigma \in S_2$ we have

$$C_\sigma = \left[\mathit{vars}_\sigma(\psi_2 \wedge \delta_{V_S} \wedge \delta_{V_{\overline{S}}})\right]^{\mathcal{C}} = \left[\mathit{vars}_\sigma(V_S \cup V_{\overline{S}})\right]^{\mathcal{C}} \ .$$

Consider the sizes of the domains in $\mathcal{C}$:

- for $\sigma \in S \cap S_1$, we have that $|C_\sigma| \leq |B_\sigma| = \kappa_\sigma$, since both $\mathcal{C}$ and $\mathcal{B}$ agree on δ_{V_S};
- for $\sigma \in S \setminus S_1$, we have that $|C_\sigma| \leq |B_\sigma| = \kappa'_\sigma$, since both $\mathcal{C}$ and $\mathcal{B}$ agree on δ_{V_S};
- for $\sigma \in S_2 \setminus S$, we have that $|C_\sigma| \leq |A_\sigma| \leq \kappa_\sigma$, since both $\mathcal{C}$ and $\mathcal{A}$ agree on $\delta_{V_{\overline{S}}}$.

Finally, because T_2 is smooth with respect to S_2, we know there is a T_2-structure $\mathcal{D}$ that satisfies $\psi_2 \wedge \delta_{V_S} \wedge \delta_{V_{\overline{S}}}$ such that:

- for $\sigma \in S \cap S_1$, we have that $|D_\sigma| = |B_\sigma| = \kappa_\sigma$;
- for $\sigma \in S \setminus S_1$, we have that $|D_\sigma| = |B_\sigma| = \kappa'_\sigma$;

- for $\sigma \in S_2 \setminus S$, we have $|D_\sigma| = \kappa_\sigma$.

Since the structures $\mathcal{B}$ and $\mathcal{D}$ agree on the sizes of the shared sorts, and they agree on the arrangement of the common variables, we know from Theorem 2.5 that we can combine them into a T-interpretation $\mathcal{F}$ that satisfies $\phi_1 \wedge \psi_2$ and has the required cardinalities. This interpretation also satisfies ϕ, so T is smooth with respect to S^*.

The second part of the proof requires showing that T is finitely witnessable with respect to S^*. Let ϕ be a conjunction of flat Σ-literals. Again, we can separate ϕ into $\phi_1 \wedge \phi_2$ such that ϕ_1 is a Σ_1-formula and ϕ_2 is a Σ_2-formula. Since T_1 and T_2 are both finitely witnessable (with respect to S_1 and S_2 respectively), there are computable witness functions $\mathit{witness}_1$ and $\mathit{witness}_2$. We define the witness function $\mathit{witness}$ for T using the following steps:

1. compute the T_2 witness function $\psi_2 = \mathit{witness}_2(\phi_2)$;
2. let $\mathcal{E}$ be the set of all equivalence relations over $V_S = \mathit{vars}_S(\psi_2)$;
3. compute the T_1-part of the witness function
$$\psi_1 = \bigvee_{E \in \mathcal{E}} \mathit{witness}_1(\phi_1 \wedge \delta(E)) \ ;$$
4. let $\psi = \mathit{witness}(\phi) = \psi_1 \wedge \psi_2$.

To show the first requirement of Definition 3.2, suppose we have a T-interpretation $\mathcal{A}$ such that

$$\mathcal{A} \vDash \phi_1 \wedge \phi_2 \ .$$

We can use the T_2-equivalence of applying $\mathit{witness}_2$ to obtain

$$\mathcal{A} \vDash \phi_1 \wedge \exists \overrightarrow{w}_2 . \psi_2 \ ,$$

where $\overrightarrow{w}_2 = \mathit{vars}(\psi_2) \setminus \mathit{vars}(\phi_2)$. It follows that we can find a suitable vector of elements $\overrightarrow{a}_2$ such that

$$\mathcal{A}\{\overrightarrow{w}_2 \leftarrow \overrightarrow{a}_2\} \vDash \phi_1 \wedge \psi_2 \ .$$

Now, let E_S be the unique equivalence relation over $V_S = \mathit{vars}_S(\psi_2)$ compatible with $\mathcal{A}\{\overrightarrow{w}_2 \leftarrow \overrightarrow{a}_2\}$. Adding the arrangement $\delta(E_S)$ and using the T_1-equivalence of applying $\mathit{witness}_1$ we further obtain

$$\begin{aligned} &\mathcal{A}\{\overrightarrow{w}_2 \leftarrow \overrightarrow{a}_2\} \vDash \phi_1 \wedge \delta(E_S) \wedge \psi_2 \ , \\ &\mathcal{A}\{\overrightarrow{w}_2 \leftarrow \overrightarrow{a}_2\} \vDash \exists \overrightarrow{w}_1 . \mathit{witness}_1(\phi_1 \wedge \delta(E_S)) \wedge \psi_2 \ , \end{aligned}$$

where $\overrightarrow{w}_1 = \mathit{vars}(\mathit{witness}_1(\phi_1 \wedge \delta(E_S))) \setminus \mathit{vars}(\phi_1 \wedge \delta(E_S))$. Since we always assume that variables introduced by witness functions are fresh, we can safely conclude that $\overrightarrow{w}_1$ and $\overrightarrow{w}_2$ are disjoint and thus there is a suitable $\overrightarrow{a}_1$ such that

$$\mathcal{A}\{\overrightarrow{w}_1 \leftarrow \overrightarrow{a}_1, \overrightarrow{w}_2 \leftarrow \overrightarrow{a}_2\} \vDash \mathit{witness}_1(\phi_1 \wedge \delta(E_S)) \wedge \psi_2 \ .$$

Let $\overrightarrow{w}_3$ be the variables from $\mathit{vars}(\psi) \setminus \mathit{vars}(\phi)$ not already in $\overrightarrow{w}_1$ or $\overrightarrow{w}_2$. Clearly there is an $\overrightarrow{a}_3$ such that

$$\mathcal{A}\{\overrightarrow{w}_1 \leftarrow \overrightarrow{a}_1, \overrightarrow{w}_2 \leftarrow \overrightarrow{a}_2, \overrightarrow{w}_3 \leftarrow \overrightarrow{a}_3\} \vDash \mathit{witness}_1(\phi_1 \wedge \delta(E_S)) \wedge \psi_2 \ . \tag{4}$$

Finally, since $\mathit{witness}_1(\phi_1 \wedge \delta(E_S))$ entails ψ_1, we can conclude

$$\mathcal{A}\{\overrightarrow{w}_1 \leftarrow \overrightarrow{a}_1, \overrightarrow{w}_2 \leftarrow \overrightarrow{a}_2, \overrightarrow{w}_3 \leftarrow \overrightarrow{a}_3\} \vDash \bigvee_{E \in \mathcal{E}} \mathit{witness}_1(\phi_1 \wedge \delta(E)) \wedge \psi_2 \ , \tag{5}$$

and thus

$$\mathcal{A} \vDash \exists \overrightarrow{w}_1 \exists \overrightarrow{w}_2 \exists \overrightarrow{w}_3 . \psi \ .$$

To show the implication in the other direction, each step is straightforward except the step from equation 5 to equation 4. Notice however, that because of the first property of witness functions, if a T_1 interpretation satisfies $witness_1(\phi_1 \wedge \delta(E))$, then it also satisfies $\delta(E)$. Now, since exactly one arrangement $\delta(E)$ is true in a particular interpretation, this means that exactly one of the disjuncts holds.

To see that the second requirement of Definition 3.2 is also satisfied, let $\mathcal{A}$ be a T-interpretation satisfying $\psi \wedge \delta_{S^*}$, where δ_{S^*} is an arrangement over a set V of variables with sorts in S^*. We will assume that $vars_{S^*}(\psi) \subseteq V$, as we can always add the extra variables from ψ to the arrangement while keeping compatibility with $\mathcal{A}$. This does not affect the correctness of our argument: we will show that there is an interpretation $\mathcal{E}$ such that $\mathcal{E} \models_T \psi \wedge \delta_{S^*}$ and $E_\sigma = [vars_\sigma(\psi \wedge \delta_{S^*})]^{\mathcal{E}}$; notice that if extra variables from ψ were included in δ_{S^*}, we can remove them and the same interpretation $\mathcal{E}$ still has the desired properties.

In the following, for any set U of sorts, we will abbreviate $\delta_{vars_U(V)}$ as δ_U. Note that δ_{S^*} can be decomposed into:

$$\delta_{S^*} = \delta_{S_1 \setminus S} \wedge \delta_{S_1 \cap S} \wedge \delta_{S_2 \setminus S} \ .$$

We can construct an additional variable arrangement $\delta_{S \setminus S_1}$ over the variables $vars_{S \setminus S_1}(\psi_2)$ that is compatible with $\mathcal{A}$. These arrangements are all true in $\mathcal{A}$, so letting $V_i = vars_{\Sigma_i^{\mathbb{S}}}(\psi \wedge \delta_{S^*})$ we have:

$$\mathcal{A}^{\Sigma_1, V_1} \models_{T_1} \left(\psi_1 \wedge \delta_{S_1 \setminus S} \wedge \delta_{S_1 \cap S} \wedge \delta_{S \setminus S_1}\right) = \Psi_1 \ ,$$
$$\mathcal{A}^{\Sigma_2, V_2} \models_{T_2} \left(\psi_2 \wedge \delta_{S_2 \setminus S} \wedge \delta_{S_1 \cap S} \wedge \delta_{S \setminus S_1}\right) = \Psi_2 \ .$$

Expanding the first equation (and dropping the last arrangement) we get that

$$\mathcal{A}^{\Sigma_1, V_1} \models_{T_1} \bigvee_{E \in \mathcal{E}} witness_1(\phi_1 \wedge \delta(E)) \wedge \delta_{S_1 \setminus S} \wedge \delta_{S_1 \cap S} \ .$$

Note that exactly one of the arrangements $\delta(E)$ is satisfied by $\mathcal{A}^{\Sigma_1, V_1}$. Call this arrangement $\delta(E_S)$. Because of the T_1-equivalence of applying $witness_1$, we have

$$\mathcal{A}^{\Sigma_1, V_1} \models_{T_1} witness_1(\phi_1 \wedge \delta(E_S)) \wedge \delta_{S_1 \setminus S} \wedge \delta_{S_1 \cap S} = \Psi_1' \ .$$

Now, because T_1 is finitely witnessable over S_1, we can obtain a T_1-interpretation $\mathcal{B}$ such that

$$\mathcal{B} \models_{T_1} \Psi_1' \ ,$$

and for all $\sigma \in S_1$ we have $B_\sigma = [vars_\sigma(\Psi_1')]^{\mathcal{B}}$. Note that though Ψ_1' and Ψ_1 differ, we have that $vars(\Psi_1') \subseteq vars(\Psi_1)$. We can thus extend $\mathcal{B}$ arbitrarily to interpret all of the variables in $vars(\Psi_1)$ so that $B_\sigma = [vars_\sigma(\Psi_1)]^{\mathcal{B}}$ for $\sigma \in S_1$.

Because $\mathcal{B} \models \Psi_1'$, we know that $\mathcal{B}$ will also satisfy $\delta(E_S)$ (by the first property of $witness_1$). Now, since $\delta(E_S)$ includes all the variables in $\delta_{S \setminus S_1}$ by definition (they both only arrange the variables from ψ_2), and because both $\delta(E_S)$ and $\delta_{S \setminus S_1}$ are satisfied by the same interpretation $\mathcal{A}$, we know that $\mathcal{B}$ also satisfies $\delta_{S \setminus S_1}$:

$$\mathcal{B} \models_{T_1} witness_1(\phi_1 \wedge \delta(E_S)) \wedge \delta_{S_1 \setminus S} \wedge \delta_{S_1 \cap S} \wedge \delta_{S \setminus S_1} \ .$$

Since one disjunct of ψ_1 is satisfied, we can conclude that

$$\mathcal{B} \models_{T_1} \Psi_1 \ .$$

Now, let's consider Ψ_2. Because T_2 is finitely witnessable over S_2, we can obtain a T_2-interpretation $\mathcal{C}$ satisfying Ψ_2 such that for $\sigma \in S_2$, we have $C_\sigma = [vars_\sigma(\Psi_2)]^{\mathcal{C}}$.

Since both $\mathcal{B}$ and $\mathcal{C}$ satisfy the arrangement $\delta_{S_1 \cap S}$ and this arrangement contains all the variables of sorts $S_1 \cap S$ from ψ_1 and ψ_2, it follows that for $\sigma \in S_1 \cap S$, we have $|B_\sigma| = |C_\sigma|$. For the other shared sorts $\sigma \in S \setminus S_1$, we have that $|C_\sigma| \leq |B_\sigma|$ because Ψ_1 and Ψ_2 agree on $\delta_{\{\sigma\}}$, and we know that C_σ does not interpret any elements beyond those named by variables in $\delta_{\{\sigma\}}$ (since we chose $\delta_{S \setminus S_1}$ to include $vars_{S \setminus S_1}(\psi_2)$).

As in the proof of smoothness, we now proceed to combine the two structures. By smoothness of T_2 with respect to S_2 we can lift the structure $\mathcal{C}$ to a structure $\mathcal{D}$ that satisfies Ψ_2, such that

- $|D_\sigma| = |B_\sigma| = |C_\sigma|$ for $\sigma \in S_1 \cap S_2$,
- $|D_\sigma| = |B_\sigma| \geq |C_\sigma|$ for $\sigma \in S \setminus S_1$,
- $|D_\sigma| = |C_\sigma|$ for $\sigma \in S_2 \setminus S$.

Interpretations $\mathcal{B}$ and $\mathcal{D}$ agree on the arrangements $\delta_{S \cap S_1} \wedge \delta_{S \setminus S_1}$. These arrangements include all of the shared variables of Ψ_1 and Ψ_2. For sorts in $S \cap S_1$, this follows by our assumption that δ_{S^*} includes all the variables in ψ of sorts in S^*. For sorts in $S \setminus S_1$, this follows from the fact that $\delta_{S \setminus S_1}$ includes all the variables in $vars_{S \setminus S_1}(\psi_2)$.

Finally, by Theorem 2.5, given interpretations $\mathcal{B}$ and $\mathcal{D}$, we can find an interpretation $\mathcal{E}$ satisfying $\Psi_1 \wedge \Psi_2$, because they agree on the arrangement over the shared variables of Ψ_1 and Ψ_2, and have the same cardinalities over the shared sorts. Moreover, since we are keeping the cardinalities of $\mathcal{B}$ over S_1 and $\mathcal{C}$ over $S_2 \setminus S$, and these cardinalities are determined by the arrangements in δ^*_S, and $vars_{S^*}(\Psi_1 \wedge \Psi_2) = vars(\delta_{S^*})$, we will also have that for all $\sigma \in S^*$, $E_\sigma = [vars_\sigma(\Psi_1 \wedge \Psi_2)]^{\mathcal{E}} = [vars_\sigma(\psi \wedge \delta_{S^*})]^{\mathcal{E}}$, as required. This concludes the proof of finite witnessability and shows that $T_1 \oplus T_2$ is polite with respect to S^*. □

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

Example 3.8. *Let $T_{\mathsf{array},1}$ and $T_{\mathsf{array},2}$ be two theories of arrays over the following sets of sorts respectively*

$$S_1 = \{\mathsf{array}_1, \mathsf{index}_1, \mathsf{elem}_1\} \ ,$$
$$S_2 = \{\mathsf{array}_2, \mathsf{index}_2, \mathsf{array}_1\} \ .$$

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 $T_{\mathsf{array},1}$ is polite with respect to $S_1^* = \{\mathsf{index}_1, \mathsf{elem}_1\}$, and the theory $T_{\mathsf{array},2}$ is polite with respect to $S_2^* = \{\mathsf{index}_2, \mathsf{array}_1\}$. Using Theorem 3.7, we know that we can combine them into a theory T_{array} that is polite with respect to the set

$$S_1^* \cup (S_2^* \setminus \{\mathsf{array}_1\}) = \{\mathsf{index}_1, \mathsf{index}_2, \mathsf{elem}_1\} \ .$$

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 3.7 is that, if both theories are polite with respect to the shared sorts then, analogously to Proposition 2.6, we get a theory that is polite with respect to the union of the sorts.

Corollary 3.9. *Let Σ_1 and Σ_2 be signatures. If*

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

4 Combining Multiple Polite Theories

Given a Σ_1-theory T_1, polite with respect to sorts S_1, and a Σ_2-theory T_2, polite with respect to sorts S_2, we will denote their combination using the combination framework for polite theories as $T_1 \oplus_{\mathsf{p}} T_2$. Here, $\oplus_{\mathsf{p}}$ is a partial, asymmetric operator: $T_1 \oplus_{\mathsf{p}} T_2$ is defined as $T_1 \oplus T_2$ if $\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}} \subseteq S_2$ and is undefined otherwise. Note that if defined, $T_1 \oplus_{\mathsf{p}} T_2$ is polite with respect to $S_1 \cup (S_2 \setminus \Sigma_1^{\mathbb{S}})$ by Theorem 3.7.

Because of the asymmetry in its definition, it is not obvious whether $\oplus_{\mathsf{p}}$ is associative or commutative. When dealing with several theories there might be several ways we could try to combine them: given theories T_1, T_2 and T_3, we could first combine T_1 and T_2 into $T_1 \oplus_{\mathsf{p}} T_2$ and then combine the result with T_3 to obtain $(T_1 \oplus_{\mathsf{p}} T_2) \oplus_{\mathsf{p}} T_3$. Or we might opt to combine T_2 and T_3 first and then combine T_1 with $T_2 \oplus_{\mathsf{p}} T_3$ to get the same theory $T_1 \oplus_{\mathsf{p}} (T_2 \oplus_{\mathsf{p}} T_3)$. Some of these operations might not be defined, and if they are, it is not obvious whether Theorem 3.7 ensures that the resulting theories are polite with respect to the same set of sorts.

Since, as explained above, there are several ways of obtaining a combined theory using the combination framework, we will write $T_1 \leftrightarrow_p T_2$ to denote that T_1 and T_2 are either both undefined or both defined, and in the latter case that T_1 and T_2 are polite with respect to the same sets of sorts.

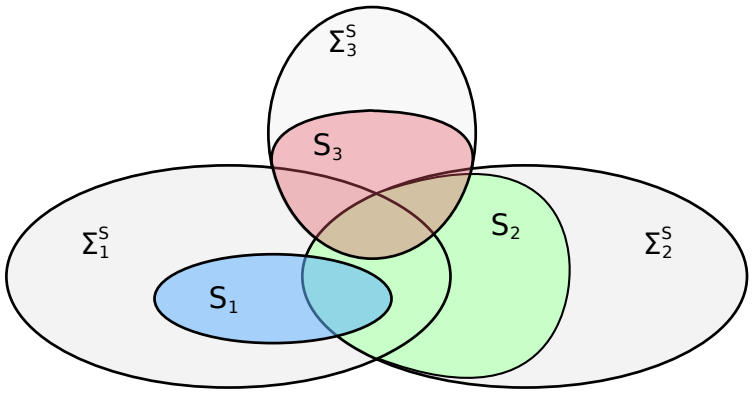

Figure 2: Diagram for Lemma 4.1.

Lemma 4.1. *Let T_i be a Σ_i-theory, polite with respect to sorts S_i, for $i = 1, 2, 3$. Then*

$$T_1 \oplus_{\mathsf{p}} (T_2 \oplus_{\mathsf{p}} T_3) \leftrightarrow_p (T_1 \oplus_{\mathsf{p}} T_2) \oplus_{\mathsf{p}} T_3 \ . \tag{6}$$

Proof. The proof of this statement primarily relies on simple manipulations in basic set theory. For convenience, it might be easier to understand the result by looking at Figure 2.

We first note that the theory combination operator $\oplus$ is clearly associative. Thus, it suffices to show that if one side of (6) is defined, then the other is also, and that they are polite with respect to the same sets of sorts.

Assume that the right-associative combination $T_1 \oplus_{\mathsf{p}} (T_2 \oplus_{\mathsf{p}} T_3)$ is defined. This implies that we have

$$\Sigma_2^{\mathbb{S}} \cap \Sigma_3^{\mathbb{S}} \subseteq S_3 \ . \tag{7}$$

Using Theorem 3.7 we know that $T_2 \oplus T_3$ is polite with respect to $S_2 \cup (S_3 \setminus \Sigma_2^{\mathbb{S}})$. Then, since we can combine T_1 with $T_2 \oplus T_3$, we must have $\Sigma_1^{\mathbb{S}} \cap (\Sigma_2^{\mathbb{S}} \cup \Sigma_3^{\mathbb{S}}) \subset S_2 \cup (S_3 \setminus \Sigma_2^{\mathbb{S}})$ which is equivalent to the following

$$\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}} \subset S_2 \cup (S_3 \setminus \Sigma_2^{\mathbb{S}}) \ , \tag{8}$$
$$\Sigma_1^{\mathbb{S}} \cap \Sigma_3^{\mathbb{S}} \subset S_2 \cup (S_3 \setminus \Sigma_2^{\mathbb{S}}) \ . \tag{9}$$

It follows from (8) (intersecting both sides with $\Sigma_2^{\mathbb{S}}$) that

$$\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}} \subseteq S_2 \ ,$$

which is enough to conclude that we can combine T_1 and T_2 into $T_1 \oplus T_2$. To be able to combine $T_1 \oplus T_2$ with T_3 we must show that

$$(\Sigma_1^{\mathbb{S}} \cup \Sigma_2^{\mathbb{S}}) \cap \Sigma_3^{\mathbb{S}} \subseteq S_3,$$

which is equivalent to

$$\Sigma_1^{\mathbb{S}} \cap \Sigma_3^{\mathbb{S}} \subseteq S_3 \ , \tag{10}$$
$$\Sigma_2^{\mathbb{S}} \cap \Sigma_3^{\mathbb{S}} \subseteq S_3 \ . \tag{11}$$

We have that (7) and (11) are the same. To show (10), it is sufficient to show both of the following

$$(\Sigma_1^{\mathbb{S}} \cap \Sigma_3^{\mathbb{S}}) \cap \Sigma_2^{\mathbb{S}} \subseteq S_3 \cap \Sigma_2^{\mathbb{S}} \ , \tag{12}$$
$$(\Sigma_1^{\mathbb{S}} \cap \Sigma_3^{\mathbb{S}}) \setminus \Sigma_2^{\mathbb{S}} \subseteq S_3 \setminus \Sigma_2^{\mathbb{S}} \ . \tag{13}$$

Equation (12) follows directly from (11) (intersect both sides with $\Sigma_2^{\mathbb{S}}$ and then intersect just the left side with $\Sigma_1^{\mathbb{S}}$). Equation (13) is obtained by subtracting $\Sigma_2^{\mathbb{S}}$ from both sides of (9). This shows that the left-associative combination $(T_1 \oplus_{\mathsf{p}} T_2) \oplus_{\mathsf{p}} T_3$ is defined.

In the opposite direction, assume that the left-associative combination $(T_1 \oplus_{\mathsf{p}} T_2) \oplus_{\mathsf{p}} T_3$ is defined. For this to be possible we need to combine T_1 and T_2 first, so it must be the case that

$$\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}} \subseteq S_2 \ . \tag{14}$$

Next, to combine $T_1 \oplus T_2$ with T_3, we must have

$$(\Sigma_1^{\mathbb{S}} \cup \Sigma_2^{\mathbb{S}}) \cap \Sigma_3^{\mathbb{S}} \subseteq S_3 \ . \tag{15}$$

This is equivalent to

$$\Sigma_1^{\mathbb{S}} \cap \Sigma_3^{\mathbb{S}} \subseteq S_3 \ , \tag{16}$$
$$\Sigma_2^{\mathbb{S}} \cap \Sigma_3^{\mathbb{S}} \subseteq S_3 \ . \tag{17}$$

From (17) we immediately get that we can combine theories T_2 and T_3 into $T_2 \oplus T_3$ which is polite with respect to $S_2 \cup (S_3 \setminus \Sigma_2^{\mathbb{S}})$. To be able to combine T_1 with $T_2 \oplus T_3$ we need to show that

$$\Sigma_1^{\mathbb{S}} \cap (\Sigma_2^{\mathbb{S}} \cup \Sigma_3^{\mathbb{S}}) \subseteq S_2 \cup (S_3 \setminus \Sigma_2^{\mathbb{S}}) \ .$$

This is in turn equivalent to

$$\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}} \subseteq S_2 \cup (S_3 \setminus \Sigma_2^{\mathbb{S}}) \ , \tag{18}$$
$$\Sigma_1^{\mathbb{S}} \cap \Sigma_3^{\mathbb{S}} \subseteq S_2 \cup (S_3 \setminus \Sigma_2^{\mathbb{S}}) \ . \tag{19}$$

From (14) we immediately get (18). To show (19), it is sufficient to show both of the following

$$(\Sigma_1^{\mathbb{S}} \cap \Sigma_3^{\mathbb{S}}) \cap \Sigma_2^{\mathbb{S}} \subseteq S_2 \ , \tag{20}$$
$$(\Sigma_1^{\mathbb{S}} \cap \Sigma_3^{\mathbb{S}}) \setminus \Sigma_2^{\mathbb{S}} \subseteq S_3 \setminus \Sigma_2^{\mathbb{S}} \ . \tag{21}$$

Equation (20) follows directly from (14). Equation (21) is obtained by subtracting $\Sigma_2^{\mathbb{S}}$ from both sides of (16). This proves that the right-associative combination $T_1 \oplus_{\mathsf{p}} (T_2 \oplus_{\mathsf{p}} T_3)$ is defined.

To show that the order of combination has no impact on the resulting sets of polite sorts, we compute the sets for both cases. If we consider the combination $T_1 \oplus_{\mathsf{p}} (T_2 \oplus_{\mathsf{p}} T_3)$, we would first get that $T_2 \oplus T_3$ is polite with respect to $S_2 \cup (S_3 \setminus \Sigma_2^{\mathbb{S}})$. Combining the resulting theory with T_1 gives the final set of polite sorts

$$S_1 \cup (S_2 \cup (S_3 \setminus \Sigma_2^{\mathbb{S}})) \setminus \Sigma_1^{\mathbb{S}} = S_1 \cup (S_2 \setminus \Sigma_1^{\mathbb{S}}) \cup (S_3 \setminus (\Sigma_1^{\mathbb{S}} \cup \Sigma_2^{\mathbb{S}})) \ . \tag{22}$$

Combining in the other direction, we first get that $T_1 \oplus T_2$ is polite with respect to $S_1 \cup S_2 \setminus \Sigma_1^{\mathbb{S}}$. Combining the result with T_3 gives the set of sorts (22). □

Lemma 4.1 gives us the associativity of $\oplus_{\mathsf{p}}$. The next lemma shows that we can also achieve commutativity if both theories are polite with respect to at least the shared sorts.

Lemma 4.2. *Let T_i be a Σ_i-theory polite with respect to the set of sorts $S_i \subseteq \Sigma_i^{\mathbb{S}}$, for $i = 1, 2$. Then the following are equivalent*

1. $T_1 \oplus_{\mathsf{p}} T_2 \leftrightarrow_p T_2 \oplus_{\mathsf{p}} T_1$;
2. $\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}} = S_1 \cap S_2$.

Proof. If both $T_1 \oplus_{\mathsf{p}} T_2$ and $T_2 \oplus_{\mathsf{p}} T_1$ are defined, then we can use either T_1 or T_2 as the polite theory in the combination framework. This means that both $\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}} \subseteq S_1$ and $\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}} \subseteq S_2$, which implies that $\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}} = S_1 \cap S_2$. In the other direction, if $\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}} = S_1 \cap S_2$ holds, then (1) is a consequence of Corollary 3.9. □

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

Theorem 4.3. *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 $1 \leq i \leq n$;*
- *$\Sigma_i^{\mathbb{S}} \cap \Sigma_j^{\mathbb{S}} \subseteq S_j$, for $1 \leq i < j \leq n$.*

Then the quantifier-free satisfiability problem for $T = T_1 \oplus \cdots \oplus T_n$ is decidable. Moreover, the resulting theory T is polite with respect to the set of sorts

$$S = \bigcup_{j=1}^{n} \left(S_j \setminus (\bigcup_{i<j} \Sigma_i^{\mathbb{S}}) \right) \ .$$

Proof. We prove the statement by induction on the number of theories n. In the base case, when $n = 2$, this directly follows from Proposition 3.5 and Theorem 3.7, i.e. if we have that $\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}} \subseteq S_2$, then we know how to devise the decision procedure for $T_1 \oplus T_2$ using the algorithm from [7]. Moreover, the resulting theory is polite with respect to $S_1 \cup (S_2 \setminus \Sigma_1^{\mathbb{S}})$.

Assume that the statement holds for $n > 1$ and consider the case for $n + 1$. By the inductive hypothesis, we have that the theory $T = T_1 \oplus \cdots \oplus T_n$ over the signature $\Sigma = \Sigma_1 \cup \ldots \cup \Sigma_n$ is decidable and polite with respect to

$$S = \bigcup_{j=1}^{n} \left(S_j \setminus (\bigcup_{i<j} \Sigma_i^{\mathbb{S}}) \right) \quad .$$

We have that $\Sigma_i^{\mathbb{S}} \cap \Sigma_{n+1}^{\mathbb{S}} \subseteq S_{n+1}$, for $1 \leq i \leq n$. Taking the union of these we get that

$$(\Sigma_1^{\mathbb{S}} \cup \ldots \cup \Sigma_n^{\mathbb{S}}) \cap \Sigma_{n+1}^{\mathbb{S}} = \Sigma^{\mathbb{S}} \cap \Sigma_{n+1}^{\mathbb{S}} \subseteq S_{n+1}$$

Since quantifier-free satisfiability in both T and T_{n+1} are decidable and the theories satisfy the conditions of Proposition 3.5 and Theorem 3.7, we know that quantifier-free satisfiability is decidable in the combination $T \oplus T_{n+1} = T_1 \oplus \cdots \oplus T_{n+1}$. Furthermore, the combination is polite with respect to the set

$$\begin{aligned} S &= \bigcup_{j=1}^{n} \left(S_j \setminus (\bigcup_{i<j} \Sigma_i^{\mathbb{S}}) \right) \cup (S_{n+1} \setminus \bigcup_{k=1}^{n} \Sigma_k^{\mathbb{S}}) \\ &= \bigcup_{j=1}^{n+1} \left(S_j \setminus (\bigcup_{i<j} \Sigma_i^{\mathbb{S}}) \right) \quad . \end{aligned}$$

This concludes the proof. □

Example 4.4. *Assume we have a theory of arrays $T_{\mathsf{array},1}$ over the sorts*

$$\Sigma_{\mathsf{array},1}^{\mathbb{S}} = \{\mathsf{array}_1, \mathsf{index}_1, \mathsf{elem}\} \ ,$$

as well as theories of arrays $T_{\mathsf{array},k}$ over the sorts

$$\Sigma_{\mathsf{array},k}^{\mathbb{S}} = \{\mathsf{array}_k, \mathsf{index}_k, \mathsf{array}_{k-1}\} \ ,$$

for $k \geq 2$. These theories represent different layers in the theory of n-dimensional arrays. The theories satisfy the assumption of Theorem 4.3 and thus we can combine them into the full theory

$$T_{\mathsf{array}} = T_{\mathsf{array},1} \oplus T_{\mathsf{array},2} \oplus \cdots \oplus T_{\mathsf{array},n} \ .$$

This theory is polite with respect to the union of all indices and elements

$$S = \{\mathsf{index}_1, \mathsf{index}_2, \ldots, \mathsf{index}_n, \mathsf{elem}\} \ .$$

Note that, although we are combining theories in a straightforward fashion, we could not have used Theorem 14 from [6] 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 4.5. *Let T_1, T_2, ..., T_n be pairwise signature-disjoint theories such that individual quantifier-free T_i-satisfiability problems are decidable. The quantifier-free satisfiability problem of $T = T_1 \oplus \cdots \oplus T_n$ 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 4.3.*

Proof. The if direction is obvious. In the other direction, assume there is a way to combine the theories T_i using the framework. Then there exists some expression combining the T_i's using $\oplus_{\mathsf{p}}$ that is defined. Using the associativity of $\oplus_{\mathsf{p}}$ (from Lemma 4.1), we can transform the expression into a sequential combination $(\cdots(T_{p_1} \oplus_{\mathsf{p}} T_{p_2}) \oplus_{\mathsf{p}} T_{p_3}) \oplus_{\mathsf{p}} \cdots \oplus_{\mathsf{p}} T_{n-1}) \oplus_{\mathsf{p}} T_n$ that satisfies the requirements of Theorem 4.3. □

5 Theory Instantiations

The way theories are defined in Definition 2.1 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. $\Sigma^{\mathbb{S}}_{\mathsf{array}} = \{\mathsf{array}, \mathsf{elem}, \mathsf{index}\}$. 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 $\mathsf{read}(a, i) \neq i$ is not well-sorted, for example). We will call this merging of sorts *theory instantiation by sort equality*.

Definition 5.1 (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 (see Lemma 5.4) 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 $vars_s(\phi) = \{v_1, v_2, \ldots, v_n\}$, we first introduce fresh variables $v_i^{\sigma_1}$ of sort σ_1, and $v_i^{\sigma_2}$ of sort σ_2, for $i = 1, \ldots, n$. The function α transforms the formula ϕ into

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

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.[7]

[7] 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.

- Literals of the form $x =_\sigma f(y_1, \ldots, y_n)$, where $\sigma \neq s$, are transformed into $x =_\sigma f(y_1^*, \ldots, y_n^*)$. 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 $x =_s f(y_1, \ldots, y_n)$ are transformed into either $x^{\sigma_1} =_{\sigma_1} f(y_1^*, \ldots, y_n^*)$ or $x^{\sigma_2} =_{\sigma_2} f(y_1^*, \ldots, y_n^*)$, depending on the sort of the co-domain of f in Σ.

- Literals of the form $p(y_1, \ldots, y_n)$ and $\neg p(y_1, \ldots, y_n)$ 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 5.2 (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 $\mathcal{B} \in \mathbf{B}$ 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 $\mathcal{B}$ exactly the same as they are interpreted in $\mathcal{A}$.*

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.

As we did with formulas, we define a transformation on structures (which we will also call α) that maps $\Sigma_s^{\sigma_1=\sigma_2}$-interpretations into Σ-interpretations. Given a $\Sigma_s^{\sigma_1=\sigma_2}$-interpretation $\mathcal{A}$, we construct the transformed structure $\mathcal{B} = \alpha(\mathcal{A})$ as follows. For sorts $\sigma \in \Sigma^{\mathbb{S}} \setminus \{\sigma_1, \sigma_2\}$, we define $B_\sigma = A_\sigma$. For the sorts σ_1 and σ_2, we define $B_{\sigma_1} = B_{\sigma_2} = A_s$. The set of variables interpreted by $\mathcal{B}$ includes all of those interpreted by $\mathcal{A}$, without the variables of sort s, and we define $v^{\mathcal{B}} = v^{\mathcal{A}}$. Finally, since $B_{\sigma_1} = B_{\sigma_2} = A_s$, we can simply define $f^{\mathcal{B}} = f^{\mathcal{A}}$ and $p^{\mathcal{B}} = p^{\mathcal{A}}$ for each function symbol f and predicate symbol p. Additionally, it is clear that if $\mathcal{A}$ is a $T_s^{\sigma_1=\sigma_2}$-interpretation, then $\alpha(\mathcal{A})$ will be a T-interpretation.

Lemma 5.3. *Let T and $\mathbf{B}$ be as in Definition 5.2, 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 $\gamma_\emptyset(\mathbf{Ax}) = \{\gamma_\emptyset(\phi) \mid \phi \in \mathbf{Ax}\}$.*

Proof. First, for every $\mathcal{B} \in \mathbf{B}$ there is a Σ-structure $\mathcal{A} \in \mathbf{A}$ such that $\mathcal{A} \vDash \mathbf{Ax}$. By the definition of γ, we also have $\mathcal{B} \vDash \gamma_\emptyset(\mathbf{Ax})$. In the other direction, let $\mathcal{B}$ be a $\Sigma_s^{\sigma_1=\sigma_2}$-structure satisfying $\gamma_\emptyset(\mathbf{Ax})$. We define a Σ-structure $\mathcal{A} = \alpha(B)$. It follows that $\mathcal{A} \vDash \mathbf{Ax}$. This implies that $\mathcal{A} \in \mathbf{A}$ and hence $\mathcal{B} \in \mathbf{B}$. □

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 $T^{\mathsf{bv}}_{\mathsf{array}} = (T_{\mathsf{array}})^{\mathsf{elem}=\mathsf{index}}_{\mathsf{bv}}$. 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 $T^{\mathsf{bv}}_{\mathsf{array}}$ 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. Before proving this, we need the following lemma.

Lemma 5.4. *Let Σ be a signature such that $\sigma_1, \sigma_2 \in \Sigma^{\mathbb{S}}$ and $s \notin \Sigma^{\mathbb{S}}$, and ϕ be a conjunction of flat $\Sigma_s^{\sigma_1=\sigma_2}$-literals. Furthermore, let $S \subseteq \Sigma^{\mathbb{S}}$ be such that $\sigma_1, \sigma_2 \in S$ and $S' = S \setminus \{\sigma_1, \sigma_2\} \cup \{s\}$. Then the following are equivalent:*

1. *ϕ is satisfiable in a $T_s^{\sigma_1=\sigma_2}$-interpretation $\mathcal{A}$ with $|A_\sigma| = \kappa_\sigma$ for $\sigma \in S'$;*
2. *$\alpha(\phi)$ is satisfiable in a T-interpretation $\mathcal{B}$ with $|B_{\sigma_1}| = |B_{\sigma_2}| = \kappa_s$, and $|B_\sigma| = \kappa_\sigma$ for $\sigma \in S \setminus \{\sigma_1, \sigma_2\}$.*

Proof. Assume that ϕ is satisfiable in a $T_s^{\sigma_1=\sigma_2}$-interpretation $\mathcal{A}$ with $|A_\sigma| = \kappa_\sigma$ for $\sigma \in S'$. Let $\mathcal{B} = \alpha(\mathcal{A})$. Then it is easy to see that the domains of B have the required sizes and $\exists \overrightarrow{v}.\alpha(\phi)$ will be satisfied by $\mathcal{B}$, where v is the vector of fresh variables introduced by α. Hence there is an interpretation $\mathcal{B}'$ that satisfies $\alpha(\phi)$ such that $|B'_{\sigma_1}| = |B'_{\sigma_2}| = \kappa_s$ and $|B_{sigma}| = \kappa_\sigma, \sigma \in S \setminus \{\sigma_1, \sigma_2\}$.

In the other direction, assume that $\alpha(\phi)$ is satisfiable in a T-interpretation $\mathcal{B}$ with $|B_\sigma| = \kappa_\sigma$, for $\sigma \in S \setminus \{\sigma_1, \sigma_2\}$, and $|B_{\sigma_1}| = |B_{\sigma_2}| = \kappa_s$. The domains B_{σ_1} and B_{σ_2} are of the same size κ_s. They also agree on the arrangement of the dual variables of sorts σ_1 and σ_2 as $\alpha(\phi)$ enforces it. Let $V_{\sigma_1} = vars_{\sigma_1}(\alpha(\phi))$ and $V_{\sigma_2} = vars_{\sigma_2}(\alpha(\phi))$. Because α introduced these variables, and because α enforces the same arrangement on on the dual variables, we have that $[V_{\sigma_1}]^{\mathcal{B}} = [V_{\sigma_2}]^{\mathcal{B}}$.

Now, let $h : V_{\sigma_1}^{\mathcal{B}} \mapsto V_{\sigma_2}^{\mathcal{B}}$ be defined as follows

$$h((v^{\sigma_1})^{\mathcal{B}}) \triangleq (v^{\sigma_2})^{\mathcal{B}} \quad .$$

This function is a bijection and is well-defined since $\mathcal{B}$ satisfies $\alpha(\phi)$. Because $|B_{\sigma_1}| = |B_{\sigma_2}|$, we can extend h to a full bijection $h_{\sigma_1} : B_{\sigma_1} \mapsto B_{\sigma_2}$. Let h_σ be the identity function for $\sigma \neq \sigma_1$.

We use this family of functions to define an interpretation $\mathcal{B}'$ isomorphic to $\mathcal{B}$ as follows. $\mathcal{B}'$ interprets all the domains of the sorts $\sigma \neq \sigma_1$, as $B'_\sigma = B_\sigma$, and the domain of the sort σ_1 as $B'_{\sigma_1} = B_{\sigma_2}$. For each variable v of sort σ, $v^{\mathcal{B}'} \triangleq h_\sigma(v^{\mathcal{B}})$. For each function symbol f, we define $f^{\mathcal{B}'}(b_1, \ldots, b_n) \triangleq h_{\tau_{n+1}}(f^{\mathcal{B}}(h^{-1}_{\tau_1}(b_1), \ldots, h^{-1}_{\tau_n}(b_n)))$ where τ_i is chosen to match the i^{th} sort in the arity of f. Similarly, we define $p^{\mathcal{B}'}(b_1, \ldots, b_n)$ iff $p^{\mathcal{B}}(h^{-1}_{\tau_1}(b_1), \ldots, h^{-1}_{\tau_n}(b_n))$. It is easy to see that the resulting interpretation $\mathcal{B}'$ is indeed isomorphic to $\mathcal{B}$, and as a result, $\mathcal{B}'$ is also a T-interpretation and satisfies $\alpha(\phi)$.

Finally, let $\mathcal{A}$ be a $T_s^{\sigma_1=\sigma_2}$-interpretation obtained from $\mathcal{B}'$ as in Definition 5.2 (i.e. $A_\sigma = B'_\sigma$ for $\sigma \in S' \setminus \{s\}$, $A_s = B'_{\sigma_1} = B'_{\sigma_2}$, and the function and predicate symbols are interpreted the same in $\mathcal{A}$ as in $\mathcal{B}$). It is easy to see that we have $|A_\sigma| = \kappa_\sigma$ for $\sigma \in S'$. It remains to say how variables are interpreted in $\mathcal{A}$. For variables v of sort $\sigma \in S' \setminus \{s\}$, we let $v^{\mathcal{A}} = v^{\mathcal{B}'}$. In addition, for each variable $v \in vars_s(\phi)$, we let $v^{\mathcal{A}} = (v^{\sigma_1})^{\mathcal{B}'}$. Note that because of the way $\mathcal{B}'$ was constructed, we also have $v^{\mathcal{A}} = (v^{\sigma_2})^{\mathcal{B}'}$. Because $\mathcal{A}$ interprets the variables v of sort s in ϕ the same as both v^{σ_1} and v^{σ_2} in $\mathcal{B}'$, $\mathcal{A}$ interprets everything else exactly the same as in $\mathcal{B}'$, and because $\mathcal{B}'$ satisfies $\alpha(\phi)$, it follows that $\mathcal{A}$ satisfies ϕ. □

Now we can prove the main theorem.

Theorem 5.5. *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*

$$witness_s^{\sigma_1=\sigma_2}(\phi) = (\gamma_{vars_s(\phi)} \circ witness \circ \alpha)(\phi) \ .$$

Proof. First we show that $T_s^{\sigma_1=\sigma_2}$ is smooth with respect to S'. Let ϕ be a conjunction of flat $\Sigma_s^{\sigma_1=\sigma_2}$-literals satisfiable in a $T_s^{\sigma_1=\sigma_2}$-structure $\mathcal{A}$. We are given cardinalities $\kappa_\sigma \geq |A_\sigma|$, for $\sigma \in S'$. By Lemma 5.4 we know that $\alpha(\phi)$ is satisfiable in a T-interpretation $\mathcal{B}$ such that $|B_\sigma| = |A_\sigma|, \sigma \in S' \setminus \{s\}$, and $|B_{\sigma_1}| = |B_{\sigma_2}| = |A_s|$. By smoothness of T there is a T-interpretation $\mathcal{B}'$ that satisfies $\alpha(\phi)$, such that $|B'_\sigma| = \kappa_\sigma$, for $\sigma \in S' \setminus \{s\}$, and $|B'_{\sigma_1}| = |B'_{\sigma_2}| = \kappa_s$. Then, applying Lemma 5.4 one more time (in the other direction), we get that ϕ is satisfiable in a $T_s^{\sigma_1=\sigma_2}$-interpretation $\mathcal{A}'$ such that $|A'_\sigma| = \kappa_\sigma$, for $\sigma \in S'$, which proves smoothness.

Next, we need to show that $T_s^{\sigma_1=\sigma_2}$ is also finitely witnessable. Let ϕ be a conjunction of flat $T_s^{\sigma_1=\sigma_2}$-literals. Because T is finitely witnessable with respect to S, it has a witness function *witness*. We define the witness function of the instantiated theory as

$$witness_s^{\sigma_1=\sigma_2}(\phi) = (\gamma_{vars_s(\phi)} \circ witness \circ \alpha)(\phi) \ .$$

Among the fresh variables introduced by $witness_s^{\sigma_1=\sigma_2}$ we will distinguish $\overrightarrow{w}$, the fresh variables introduced by *witness*, and $\overrightarrow{v}^{\sigma_1}$ and $\overrightarrow{v}^{\sigma_2}$, the fresh variables introduced by transformation α (i.e. the variables v^{σ_1} and v^{σ_2}, corresponding to variables $v \in vars_s(\phi)$).

First we need to show that if $\psi = witness_s^{\sigma_1=\sigma_2}(\phi)$, then $\exists \overrightarrow{v}^{\sigma_1} \exists \overrightarrow{v}^{\sigma_2} \exists \overrightarrow{w}.\psi$ and ϕ are $T_s^{\sigma_1=\sigma_2}$-equivalent. This follows from the equivalence of the following statements:

$$\begin{aligned}
&\mathcal{A} \models \phi && \\
&\alpha(\mathcal{A})\{\overrightarrow{v}^{\sigma_1} \leftarrow \overrightarrow{v}^{\mathcal{A}}, \overrightarrow{v}^{\sigma_2} \leftarrow \overrightarrow{v}^{\mathcal{A}}\} \models \alpha(\phi) && \text{(definition of } \alpha\text{)} \\
&\alpha(\mathcal{A})\{\overrightarrow{v}^{\sigma_1} \leftarrow \overrightarrow{v}^{\mathcal{A}}, \overrightarrow{v}^{\sigma_2} \leftarrow \overrightarrow{v}^{\mathcal{A}}\} \models \exists \overrightarrow{w}.witness(\alpha(\phi)) && (T \text{ finitely witnessable}) \\
&\alpha(\mathcal{A})\{\overrightarrow{v}^{\sigma_1} \leftarrow \overrightarrow{v}^{\mathcal{A}}, \overrightarrow{v}^{\sigma_2} \leftarrow \overrightarrow{v}^{\mathcal{A}}, \overrightarrow{w} \leftarrow \overrightarrow{a}\} \models witness(\alpha(\phi)) && \\
&\mathcal{A}\{\overrightarrow{v}^{\sigma_1} \leftarrow \overrightarrow{v}^{\mathcal{A}}, \overrightarrow{v}^{\sigma_2} \leftarrow \overrightarrow{v}^{\mathcal{A}}, \overrightarrow{w} \leftarrow \overrightarrow{a}\} \models \gamma_{vars_s(\phi)}(witness(\alpha(\phi))) && \text{(definition of } \gamma\text{)} \\
&\mathcal{A} \models \exists \overrightarrow{v}^{\sigma_1} \exists \overrightarrow{v}^{\sigma_2} \exists \overrightarrow{w}.\psi &&
\end{aligned}$$

To show that the defined witness function satisfies the second requirement of Definition 3.2, let

$$\mathcal{E} = \{ \ E_\sigma \ \mid \ \sigma \in S' \ \}$$

be a family of equivalence relations over a set V of variables with sorts in S', and $\delta_V(\mathcal{E})$ be the arrangement induced by $\mathcal{E}$. Now, assume that there is a $T_s^{\sigma_1=\sigma_2}$-interpretation $\mathcal{A}$ such that

$$\mathcal{A} \models \overbrace{(\gamma_{vars_s(\phi)} \circ witness \circ \alpha)(\phi)}^{\psi} \wedge \delta_V(\mathcal{E}) \ .$$

For convenience, we will let $\psi' = (witness \circ \alpha)(\phi)$ (so that $\psi = \gamma_{vars_s(\phi)}(\psi')$). As in the proof of Theorem 3.7 (page 14), we can assume wlog that $vars_{S'}(\psi) \subseteq V$, i.e. the arrangement $\delta_V(\mathcal{E})$ covers all the variables of ψ with sorts in S' . We will make use of the following sets of variables:

- $V_s = vars_s(V)$
- $V_\phi = vars_s(\phi)$

- V_α = superscripted variables in $\alpha(\phi)$ introduced by α (i.e. v^{σ_1} and v^{σ_2} for each $v \in V_\phi$).
- W_s = variables of sort σ_1 or σ_2 in ψ' introduced by the *witness* function for T.
- $\Delta_s = V_s \setminus (V_\phi \cup V_\alpha \cup W_s)$

Note that because γ reinterprets variables of sorts σ_1 and σ_2 as variables of sort s, all of the variables in the above sets are of sort s in the formula $\psi \wedge \delta_V$.

The definition of γ also guarantees that, for all $v \in V_\phi$, the variables v, v^{σ_1}, and v^{σ_2} must be equal in $\mathcal{A}$ and (since $\mathcal{A}$ also satisfies δ_V) belong to the same equivalence class in E_s. We can thus construct a new family of equivalence relations $\mathcal{E}'$ in which the variables from V_α do not appear, while keeping the same number of equivalence classes for each sort. Concretely, let

$$\begin{aligned} V'_s &= V_s \setminus V_\alpha \ , \\ E'_\sigma &= \begin{cases} E_\sigma & \text{for } \sigma \neq s \ , \\ E_s \cap V'_s \times V'_s & \text{for } \sigma = s \ , \end{cases} \\ \mathcal{E}' &= \{\ E'_\sigma \ \mid \ \sigma \in S' \ \} \ . \end{aligned}$$

Also, let $V' = V \setminus V_\alpha$, and for an equivalence relation E, let $Q(E)$ denote the quotient set of E (i.e. the set of all equivalence classes in E). It is clear that $\mathcal{A} \vDash \psi \wedge \delta_{V'}(\mathcal{E}')$ and $|Q(E_\sigma)| = |Q(E'_\sigma)|$ for each $\sigma \in S'$.

In order to switch to reasoning in the signature Σ (as opposed to $\Sigma_s^{\sigma_1=\sigma_2}$), we need to modify the equivalence relations so that variables of different sorts (when considered in the signature Σ) are not in the same equivalence class (so that the induced arrangement is well-sorted). To this end, we define the variable mappings β_{σ_1} and β_{σ_2} as follows. For $v \in V'_s$,

$$\beta_{\sigma_1}(v) = \begin{cases} v^{\sigma_1} \text{ if } v \in V_\phi \ , \\ v \text{ if } v \in W_s \text{ and } v \text{ is of sort } \sigma_1 \ , \\ v' \text{ if } v \in W_s \text{ and } v \text{ is of sort } \sigma_2 \ , \\ v \text{ if } v \in \Delta_s \end{cases} \qquad \beta_{\sigma_2}(v) = \begin{cases} v^{\sigma_2} \text{ if } v \in V_\phi \ , \\ v' \text{ if } v \in W_s \text{ and } v \text{ is of sort } \sigma_1 \ , \\ v \text{ if } v \in W_s \text{ and } v \text{ is of sort } \sigma_2 \ , \\ v' \text{ if } v \in \Delta_s \end{cases}$$

In the above, the primed variables v' are to be understood as fresh variables of the appropriate sort. In addition, we (arbitrarily) choose to interpret variables in Δ_s to be of sort σ_1 when working in Σ. Note that both functions are injective. The purpose of these mappings is to ensure that every variable in V' has some corresponding variable of sorts σ_1 and σ_2 when working in Σ. We can now construct a new family of equivalence relations as follows. First, let

$$\begin{aligned} E''_{\sigma_1} &= \{\ \ (\beta_{\sigma_1}(v_1), \beta_{\sigma_1}(v_2)) \ \mid \ (v_1, v_2) \in E'_s \ \ \} \ , \\ E''_{\sigma_2} &= \{\ \ (\beta_{\sigma_2}(v_1), \beta_{\sigma_2}(v_2)) \ \mid \ (v_1, v_2) \in E'_s \ \ \} \ , \end{aligned}$$

For the other sorts $\sigma \in S \setminus \{\sigma_1, \sigma_2\}$, we simply let $E''_\sigma = E'_\sigma = E_\sigma$. We then set $\mathcal{E}'' = \{\ E''_\sigma \ \mid \ \sigma \in S \ \}$. Let V'' be the set of variables appearing in $\mathcal{E}''$. Note that as desired, variables in the same equivalence class have the same sort. In addition, with the exception of the variables in V_ϕ, all variables appearing in $\mathcal{E}'$ also appear in $\mathcal{E}''$. The fresh primed variables are used just as temporary placeholders of the appropriate sort. It is easy to see that the number of equivalence classes is preserved, i.e. $|Q(E''_{\sigma_1})| = |Q(E''_{\sigma_2})| = |Q(E'_s)| = |Q(E_s)|$.

Now, it is not hard to construct a T-interpretation $\mathcal{B}$ starting from $\mathcal{A}$ such that

$$\mathcal{B} \vDash \psi' \wedge \delta_{V''}(\mathcal{E}'') \ .$$

We do this as follows. The domains of the Σ-structure $\mathcal{B}$ will mimic those in $\mathcal{A}$, except that $B_{\sigma_1} = B_{\sigma_2} = A_s$. We also keep the same interpretations of all function and predicate symbols, while moving back to the original signature, and hence use the new domains where necessary. It follows that $\mathcal{B}$ is a T-structure. We interpret the variables of sorts in $S \setminus \{\sigma_1, \sigma_2\}$ as they were, and the variables of sorts σ_1 and σ_2 as $(\beta_{\sigma_1}(v))^{\mathcal{B}} = (\beta_{\sigma_2}(v))^{\mathcal{B}} = (v)^{\mathcal{A}}$. By the definition of γ and due to the way we constructed $\mathcal{E}''$, it is clear that $\psi' \wedge \delta_{V''}(\mathcal{E}'')$ is indeed satisfied by $\mathcal{B}$.

Now we can apply the finite witnessability of T to obtain a T-interpretation $\mathcal{C}$ satisfying $\psi' \wedge \delta_{V''}(\mathcal{E}'')$ such that for all $\sigma \in S$ we have

$$C_\sigma = \left[vars_\sigma(\psi' \wedge \delta_{V''}(\mathcal{E}''))\right]^{\mathcal{C}} \ .$$

Since all of the variables in ψ' are also in V'', we have that

$$|C_{\sigma_1}| = |C_{\sigma_2}| = |Q(E''_{\sigma_1})| = |Q(E''_{\sigma_2})| = |Q(E'_s)| = |Q(E_s)| \ .$$

Similarly, for $\sigma \in S \setminus \{\sigma_1, \sigma_2\}$, $|C_\sigma| = |Q(E_\sigma)|$. Now, define $g_{\sigma_1} : Q(E'_s) \mapsto C_{\sigma_1}$ as $g_{\sigma_1}([v]) = (\beta_{\sigma_1}(v))^{\mathcal{C}}$. This is well-defined since $\mathcal{C}$ satisfies $\delta_{V''}(\mathcal{E}'')$. For the same reason, g_{σ_1} is injective. Finally, it must be surjective because $|Q(E'_s)| = |C_{\sigma_1}|$. Define the bijection g_{σ_2} similarly.

Now, let $h : C_{\sigma_1} \mapsto C_{\sigma_2} = g_{\sigma_2} \circ g_{\sigma_1}^{-1}$. Clearly, h is a bijection. As in the proof of Lemma 5.4, we can extend h to a family of bijections that forms an isomorphism into a T-interpretation $\mathcal{D}$ such that:

- for $\sigma \in S \setminus \{\sigma_1, \sigma_2\}$, $D_\sigma = C_\sigma$,
- $D_{\sigma_1} = D_{\sigma_2} = C_{\sigma_2}$,
- for $v \in vars_{\sigma_1}(V'')$, $v^{\mathcal{D}} = h(v^{\mathcal{C}})$,
- for $v \in V'' \setminus vars_{\sigma_1}(V'')$, $v^{\mathcal{D}} = v^{\mathcal{C}}$,
- $\mathcal{D} \vDash \psi' \wedge \delta_{V''}(\mathcal{E}'')$.

Note that for $v \in V'_s$, we have:

$$(\beta_{\sigma_1}(v))^{\mathcal{D}} = h((\beta_{\sigma_1}(v))^{\mathcal{C}}) = (\beta_{\sigma_2}(v))^{\mathcal{C}} = (\beta_{\sigma_2}(v))^{\mathcal{D}} \tag{23}$$

Finally, we construct a $T_s^{\sigma_1=\sigma_2}$-structure $\mathcal{F}$ from $\mathcal{D}$ by using the construction of Definition 5.2. We interpret in $\mathcal{F}$ only the variables $v \in V$ as follows:

$$v^{\mathcal{F}} = \begin{cases} (v^{\sigma_2})^{\mathcal{D}} \text{ if } v \in V_\phi \ , \\ v^{\mathcal{D}} \text{ otherwise} \ . \end{cases}$$

We also claim that for $v \in V'_s$,

$$v^{\mathcal{F}} = (\beta_{\sigma_2}(v))^{\mathcal{D}} \, . \tag{24}$$

This follows easily from the definition for $v \in V_\phi$. Otherwise, we have $v^{\mathcal{F}} = v^{\mathcal{D}}$ with $v \in W_s \cup \Delta_s$. But notice that in this case we know that either β_{σ_1} or β_{σ_2} is the identity function. The claim (24) then follows from (23).

Since interpretations in $\mathcal{F}$ follow those in $\mathcal{D}$ except for variables in V_ϕ, and since no variables from V_ϕ appear in $\psi' \wedge \delta_{V''}(\mathcal{E}'')$, it should be clear that $\mathcal{F} \vDash \gamma_\emptyset(\psi' \wedge \delta_{V''}(\mathcal{E}''))$. Furthermore, for $v \in V_\phi$, we have

$$\begin{aligned} v^{\mathcal{F}} &= (v^{\sigma_2})^{\mathcal{D}} \ , \\ (v^{\sigma_2})^{\mathcal{F}} &= (v^{\sigma_2})^{\mathcal{D}} \ , \\ (v^{\sigma_1})^{\mathcal{F}} &= (v^{\sigma_1})^{\mathcal{D}} = (\beta_{\sigma_1}(v))^{\mathcal{D}} = (\beta_{\sigma_2}(v))^{\mathcal{D}} = (v^{\sigma_2})^{\mathcal{D}}, \end{aligned}$$

and thus

$$\mathcal{F} \vDash v^{\sigma_1} = v \wedge v^{\sigma_2} = v \ . \tag{25}$$

It follows that $\mathcal{F} \vDash \psi$. It remains to show that $\mathcal{F} \vDash \delta_V(\mathcal{E})$. Recall that for each $v \in V_\phi$, we know that v, v^{σ_1}, and v^{σ_2} must all be in the same equivalence class in $\mathcal{E}$. It follows that if $\mathcal{F} \vDash \delta_{V'}(\mathcal{E}')$, then by (25), we will have $\mathcal{F} \vDash \delta_V(\mathcal{E})$.

To show $\mathcal{F} \vDash \delta_{V'}(\mathcal{E}')$, consider a pair of variables $v_1, v_2 \in V'$. Suppose the sorts of v_1 and v_2 are $\sigma \neq s$. We know that $E'_\sigma = E''_\sigma$, so $(v_1, v_2) \in \mathcal{E}'$ iff $(v_1, v_2) \in \mathcal{E}''$ iff $\mathcal{F} \vDash v_1 = v_2$ (since $\mathcal{F} \vDash \gamma_\emptyset(\delta_{V''}(\mathcal{E}''))$ and γ has no effect in this case).

Finally, suppose that v_1, v_2 have sort s, so that $v_1, v_2 \in V'_s$. We have

$$\begin{array}{lll} (v_1, v_2) \in \mathcal{E}' & \text{iff} & (\beta_{\sigma_2}(v_1), \beta_{\sigma_2}(v_2)) \in \mathcal{E}'' \quad \text{by def of } \mathcal{E}'' \text{ and since } \beta_{\sigma_2} \text{ is injective} \\ & \text{iff} & ((\beta_{\sigma_2}(v_1))^{\mathcal{D}} = (\beta_{\sigma_2}(v_2))^{\mathcal{D}} \quad \text{since } \mathcal{D} \vDash \delta_{V''}(\mathcal{E}'') \\ & \text{iff} & v_1^{\mathcal{F}} = v_2^{\mathcal{F}} \quad \text{by (24)} \end{array}$$

Thus, $\mathcal{F} \vDash \psi \wedge \delta_V(\mathcal{E})$.

The last step is to show that $F_\sigma = [vars_\sigma(\psi \wedge \delta_V(\mathcal{E}))]^{\mathcal{F}}$, for $\sigma \in S'$. Since $vars_{S'}(\psi) \subseteq V$, it suffices to show that $F_\sigma = [vars_\sigma(\delta_V(\mathcal{E}))]^{\mathcal{F}}$. For this to hold, it suffices to know that $|F_\sigma| = |Q(E_\sigma)|$. For $\sigma \neq s$, we have that $|F_\sigma| = |D_\sigma| = |C_\sigma| = |Q(E_\sigma)|$. Similarly, we have $|F_s| = |D_{\sigma_2}| = |C_{\sigma_2}| = |Q(E_s)|$. This concludes the proof. □

Example 5.6. *Consider again example 4.4, i.e. we have a theory of arrays* T_{array} *that operates over the sorts*

$$\Sigma^{\mathbb{S}} = \{\mathsf{array}_1, \ldots, \mathsf{array}_n, \mathsf{index}_1, \ldots \mathsf{index}_n, \mathsf{elem}_1\}$$

and is polite with respect to the index and element sorts

$$\Sigma^{\mathbb{S}} = \{\mathsf{index}_1, \ldots \mathsf{index}_n, \mathsf{elem}_1\} \ .$$

Using Theorem 5.5, we can now safely replace the sorts index_1, index_2 *and* elem_1 *with the sort of bit-vectors* bv*, obtaining a theory* $T_{\mathsf{array(bv)}}$ *of n-dimensional arrays where the elements and the indices are of the same bit-vector sort. This theory* $T_{\mathsf{array(bv)}}$ *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 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}*.*

Theorem 4.5 together with Theorem 5.5 give a practical modular approach for reasoning about and deciding combinations of polite theories.

6 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 [7] 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] Clark Barrett, Igor Shikanian, and Cesare Tinelli. An abstract decision procedure for a theory of inductive data types. *Journal on Satisfiability, Boolean Modeling and Computation*, 3:21–46, 2007.

[2] Herbert B. Enderton. *A mathematical introduction to logic*. Academic press New York, 1972.

[3] Sava Krstić, Amit Goel, Jim Grundy, and Cesare Tinelli. Combined Satisfiability Modulo Parametric Theories. In Orna Grumberg and Michael Huth, editors, *Tools and Algorithms for the Construction and Analysis of Systems, 13th International Conference, TACAS 2007*, volume 4424 of *Lecture Notes in Computer Science*, pages 602–617. Springer, 2007.

[4] Greg Nelson and Derek C. Oppen. Simplification by cooperating decision procedures. *ACM Transactions on Programming Languages and Systems*, 1(2):245–257, October 1979.

[5] Derek C. Oppen. Complexity, convexity and combinations of theories. *Theoretical Computer Science*, 12(3):291–302, 1980.

[6] Silvio Ranise, Christophe Ringeissen, and Calogero Zarba. Combining Data Structures with Nonstably Infinite Theories using Many-Sorted Logic. Research Report RR-5678, INRIA, 2005.

[7] Silvio Ranise, Christophe Ringeissen, and Calogero G. Zarba. Combining Data Structures with Nonstably Infinite Theories Using Many-Sorted Logic. In Bernhard Gramlich, editor, *Frontiers of Combining Systems, 5th International Workshop, FroCoS 2005, Vienna, Austria, September 19-21, 2005, Proceedings*, volume 3717 of *Lecture Notes in Computer Science*, pages 48–64. Springer, 2005.

[8] Cesare Tinelli and Mehdi T. Harandi. A new correctness proof of the Nelson–Oppen combination procedure. In Franz Baader and Klaus Ulrich Schulz, editors, *Frontiers of Combining Systems: Proceedings of the 1st International Workshop (Munich, Germany)*, Applied Logic, pages 103–120. Kluwer Academic Publishers, March 1996.

[9] Cesare Tinelli and Calogero Zarba. Combining decision procedures for sorted theories. In José Alferes and João Leite, editors, *Proceedings of the 9th European Conference on Logic in Artificial Intelligence (JELIA'04), Lisbon, Portugal*, volume 3229 of *Lecture Notes in Artificial Intelligence*, pages 641–653. Springer, 2004.

[10] Cesare Tinelli and Calogero Zarba. Combining decision procedures for theories in sorted logics. Technical Report 04-01, Department of Computer Science, The University of Iowa, February 2004.

[11] Cesare Tinelli and Calogero G. Zarba. Combining nonstably infinite theories. *Journal of Automated Reasoning*, 34(3):209–238, 2005.

A Flat Literals

When proving that a Σ-theory is smooth or finitely witnessable with respect to a set of sorts S, we can restrict ourselves to conjunctions of flat Σ-literals. The following two lemmas show that this can indeed be done without loss of generality. The proofs are simple and already presented in [7], but we reiterate them here since they are affected by the change in the definition of finite witnessability.

Lemma A.1. *Let Σ be a signature, let $S \subseteq \Sigma^{\mathbb{S}}$ be a set of sorts, and let T be a Σ-theory. Assume that:*

- *for every T-satisfiable conjunction of flat Σ-literals ψ,*
- *for every T-interpretation $\mathcal{A}$ satisfying ψ,*
- *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 ψ such that $|B_\sigma| = \kappa_\sigma$, for all $\sigma \in S$. Then T is smooth with respect to S.

Proof. Assume that a quantifier-free Σ-formula ϕ is satisfiable in a T-interpretation $\mathcal{A}$ and we are given cardinal numbers κ_σ, such that $\kappa_\sigma \geq |A_\sigma|$ for all $\sigma \in S$. We can transform ϕ into its disjunctive normal form $DNF(\phi) = \psi_1 \vee \ldots \vee \psi_m$. Since ϕ and $DNF(\phi)$ are equivalent, T-interpretation $\mathcal{A}$ will satisfy one of the disjuncts ψ_k, for some $1 \leq k \leq m$. We can transform ψ_k into a conjunction of flat literals ψ by introducing fresh variables $\overrightarrow{v}$, such that ψ_k is logically equivalent to $\exists \overrightarrow{v}.\psi$. It follows that there exists a T-interpretation $\mathcal{A}'$ equivalent to $\mathcal{A}$ except in its interpretation of $\overrightarrow{v}$ such that $\mathcal{A}'$ satisfies ψ.

From here, we use the assumptions to obtain a new T-interpretation $\mathcal{B}$ satisfying ψ such that $|B_\sigma| = \kappa_\sigma$, for all $\sigma \in S$. $\mathcal{B}$ will also satisfy $\exists \overrightarrow{v}.\psi$ and, by equivalence, also ψ_k, $DNF(\phi)$ and ϕ. This shows that T is smooth with to respect to S. □

Lemma A.2. *Let Σ be a signature, let $S \subseteq \Sigma^{\mathbb{S}}$ be a set of sorts, and let T be a Σ-theory. Assume there exists a computable function, $witness_F$, which, for every conjunction of flat Σ-literals ϕ, returns a quantifier-free Σ-formula $\psi = witness_F(\phi)$ such that*

- *ϕ and $(\exists \overrightarrow{w})\psi$ are T-equivalent, where $\overrightarrow{w} = vars(\psi) \setminus vars(\phi)$ 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 $A_\sigma = [vars_\sigma(\psi \wedge \delta_V)]^{\mathcal{A}}$, for all $\sigma \in S$.*

Then T is finitely witnessable with respect to S.

Proof. We want to define a witness function *witness* on all quantifier-free Σ-formulas by using the function $witness_F$ as a black box. Let ϕ be a quantifier-free Σ-formula, we compute *witness* using the following steps

1. convert ϕ into a T-equivalent disjunctive normal form $DNF(\phi) = \psi_1 \vee \ldots \psi_m$;

2. transform each disjunct ψ_i into a conjunction of flat Σ-literals ψ'_i by introducing fresh variables;

3. let $witness(\phi) = witness_F(\psi'_1) \vee \ldots \vee witness_F(\psi'_m)$.

If $\overrightarrow{v}_i$ are the fresh variables introduced by flattening ψ_i, we know that ψ_i and $\exists \overrightarrow{v}_i.\psi'_i$ are logically equivalent. Since ψ'_i is T-equivalent to $\exists \overrightarrow{w}_i.witness_F(\psi'_i)$, where $\overrightarrow{w}_i$ are the fresh variables introduced by applying the witness function, we can conclude that $\exists \overrightarrow{v}_i \exists \overrightarrow{w}_i.witness_F(\psi'_i)$ is T-equivalent to $\exists \overrightarrow{v}_i.\psi'_i$, and hence also T-equivalent to ψ_i. Since we can move existential quantifiers over disjunctions (maintaining logical equivalence), we can also conclude that ϕ is T-equivalent to $\exists \overrightarrow{v}_1 \exists \overrightarrow{w}_1 \ldots \exists \overrightarrow{v}_m \exists \overrightarrow{w}_m.witness(\phi)$. This proves the first requirement of the witness function.

For the second requirement, let $\psi = witness(\phi)$ and assume that $\psi \wedge \delta_V$ is T-satisfiable in a T-interpretation $\mathcal{A}$, for an arrangement δ_V, where V is a set of variables of sorts in S. This implies that one of the disjuncts, say $witness_F(\psi'_k)$, together with δ_V, is satisfied in $\mathcal{A}$, for some $1 \leq k \leq m$. Of course, it is likely the case that $vars_S(witness_F(\psi'_k) \wedge \delta_V)$ does not include all the variables present in $vars_S(\psi \wedge \delta_V)$, but we can add the missing variables to our arrangement δ_V [8], while keeping compatibility with $\mathcal{A}$, thus obtaining a stronger arrangement δ'.

Using the assumptions we can therefore get a T-interpretation $\mathcal{B}$ that satisfies $witness_F(\psi'_k) \wedge \delta'$ such that

$$B_\sigma = \left[vars_\sigma(witness_F(\psi'_k) \wedge \delta')\right]^{\mathcal{B}} = \left[vars_\sigma(witness(\phi) \wedge \delta_V)\right]^{\mathcal{B}} \quad ,$$

for all $\sigma \in S$. Since δ' includes δ_V, $\mathcal{B}$ will also satisfy $witness_F(\psi'_k) \wedge \delta_V$, and hence also $witness(\phi) \wedge \delta_V$. This proves that T is indeed finitely witnessable. □

[8] This is the reason why the definition of finite witnessability includes the arrangement over an *arbitrary* set of variables V instead of just an arrangement over $vars_S(\psi_2)$.