Politeness and Stable Infiniteness: Stronger Together

Ying Sheng[1], Yoni Zohar[1], Christophe Ringeissen[2], Andrew Reynolds[3], Clark Barrett[1], and Cesare Tinelli[3]

[1] Stanford University
[2] Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France
[3] The University of Iowa

Abstract. We make two contributions to the study of polite combination in satisfiability modulo theories. The first contribution is a separation between politeness and strong politeness, by presenting a polite theory that is not strongly polite. This result shows that proving strong politeness (which is often harder than proving politeness) is sometimes necessary. The second contribution is an optimization to the polite combination method, obtained by borrowing from the Nelson-Oppen method. In its non-deterministic form, the Nelson-Oppen method is based on guessing arrangements over shared variables. In contrast, polite combination requires an arrangement over *all* variables of the shared sort (not just the shared variables). We show that when using polite combination, if the other theory is stably infinite with respect to a shared sort, only the shared variables of that sort need be considered in arrangements, as in the Nelson-Oppen method. Reasoning about arrangements of variables is exponential in the worst case, so reducing the number of variables that are considered has the potential to improve performance significantly. We show preliminary evidence for this in practice by demonstrating a speed-up on a smart contract verification benchmark.

1 Introduction

Solvers for satisfiability modulo theories (SMT) [6] are used in a wide variety of applications. Many of these applications require determining the satisfiability of formulas with respect to a *combination* of background theories. In order to make reasoning about combinations of theories modular and easily extensible, a combination framework is essential. Combination frameworks provide mechanisms for automatically deriving a decision procedure for the combined theories by using the decision procedures for the individual theories as black boxes. To integrate a new theory into such a framework, it then suffices to focus on the decoupled decision procedure for the new theory alone, together with its interface to the generic combination framework.

In 1979, Nelson and Oppen [19] proposed a general framework for combining theories with disjoint signatures. In this framework, a quantifier-free formula in the combined theory is purified to a conjunction of formulas, one for each theory.

Each pure formula is then sent to a dedicated theory solver, along with a guessed arrangement (a set of equalities and disequalities that capture an equivalence relation) of the variables shared among the pure formulas. For completeness [18], this method requires all component theories to be stably infinite. While many important theories are stably infinite, some are not, including the widely-used theory of fixed-length bit-vectors. To address this issue and support more general kinds of theory combination, the polite combination method was introduced by Ranise et al. [20], and later refined by Jovanovic and Barrett [15]. In polite combination, one theory must be *polite*, a stronger requirement than stable-infiniteness, but the requirement on the other theory is relaxed: specifically, it need not be stably infinite. The price for this generality is that unlike the Nelson-Oppen method, polite combination requires guessing arrangements over *all* variables of certain sorts, not just the shared ones. At a high level, polite theories have two properties: smoothness and finite witnessability (these are explained in Section 2). The polite combination theorem in [20] contained an error, which was identified in [15]. A fix was also proposed in [15], which relies on stronger requirements for finite witnessability. Following Casal and Rasga [11], we call this strengthened version *strong finite witnessability*. A theory that is both smooth and strongly finitely witnessable is called *strongly polite*.

This paper makes two contributions. First, we give an affirmative answer to the question of whether the strengthening of the definition of finite witnessability is necessary by giving an example of a theory that is polite but not strongly polite. The given theory is over an empty signature and has two sorts, and was originally studied in [11] in the context of shiny theories. Though not explicitly mentioned in [11], this theory could be shown to separate politeness from strong politeness using elements already available in [11,20]. Here we state this result explicitly and provide a direct proof, without using shiny theories. We show that for empty signatures, at least two sorts are needed to present a polite theory that is not strongly polite and also show that this is not the case for finite witnessability. Second, we explore different polite combination scenarios, where additional information is known about the theories being combined. In particular, we improve the polite combination method for the case where one theory is strongly polite w.r.t. a set S of sorts and the other is stably infinite w.r.t. a subset $S' \subseteq S$ of the sorts. For such cases, we show that it is possible to perform Nelson-Oppen combination for S' and polite combination for $S \setminus S'$. This means that for the sorts in S', only shared variables need to be considered for the guessed arrangement, which can considerably reduce its size. We also show that the set of shared variables can be reduced for a couple of other variations of conditions on the theories. The proofs are based on arguments similar to those of the original politeness proofs from [15,20]. Finally, we present a preliminary case study using a challenge benchmark from a smart contract verification application. For this benchmark, we show that the reduction of shared variables is evident and significantly improves the solving time.

Related Work: Polite combination is part of a more general effort to replace the symmetric condition in the Nelson-Oppen approach with an asymmetric condition. Other examples of this effort include the notions of *shiny* [24], *parametric* [16], and *gentle* [14] theories. Shiny theories were introduced by Tinelli and Zarba [24] for mono-sorted signatures, and were extended to many-sorted signatures by Ranise et al. [20], who also provided a sufficient condition for their equivalence with polite theories. For the mono-sorted case, a sufficient condition for the equivalence of shiny theories and strongly polite theories was given by Casal and Rasga [10]. In later work [11], the same authors proposed a generalization of shiny theories to many-sorted signatures different from the one in [20], and proved that it is equivalent to strongly polite theories with a decidable quantifier-free fragment. Restricting the guessing of shared formulas to those built over the shared variables is also known to be sufficient in other combination methods such the ones introduced by Baader and Schulz for satisfiability problems in freely-generated structures [2,3].

The paper is organized as follows. Section 2 provides the necessary notions from first-order logic and polite theories. Section 3 discusses the difference between politeness and strong politeness and shows they are not equivalent. Section 4 gives the improvements for the combination process under certain conditions, and Section 5 demonstrates the effectiveness of these improvements for a challenge benchmark. Section 6 concludes with directions for further research. [Due to space constraints, some proofs are omitted. They can be found in the appendix.]

2 Preliminaries

2.1 Signatures and Structures

We briefly review the usual definitions of many-sorted first-order logic with equality (see [13,22] for more details). A *signature* Σ consists of a set $\mathcal{S}_\Sigma$ (of *sorts*), a set $\mathcal{F}_\Sigma$ of function symbols, and a set $\mathcal{P}_\Sigma$ of predicate symbols. We assume $\mathcal{S}_\Sigma$, $\mathcal{F}_\Sigma$ and $\mathcal{P}_\Sigma$ are countable. Function symbols have arities of the form $\sigma_1 \times \ldots \times \sigma_n \rightarrow \sigma$, and predicate symbols have arities of the form $\sigma_1 \times \ldots \times \sigma_n$, with $\sigma_1, \ldots, \sigma_n, \sigma \in \mathcal{S}_\Sigma$. For each sort $\sigma \in \mathcal{S}_\Sigma$, $\mathcal{P}_\Sigma$ includes an *equality symbol* $=_\sigma$ of arity $\sigma \times \sigma$. We denote it by $=$ when σ is clear from context. We assume an underlying countably infinite set of variables for each sort. Terms, formulas, and literals are defined in the usual way. For a Σ-formula ϕ and a sort σ, we denote the set of free variables in ϕ of sort σ by $vars_\sigma(\phi)$. This notation naturally extends to $vars_S(\phi)$ when S is a set of sorts. $vars(\phi)$ is the set of all free variables in ϕ. We denote by $QF(\Sigma)$ the set of quantifier-free Σ-formulas.

A *Σ-structure* is a many-sorted structure that provides semantics for the symbols in Σ (but not for variables). It consists of a *domain* $\sigma^{\mathcal{A}}$ for each sort $\sigma \in \mathcal{S}_\Sigma$, an interpretation $f^{\mathcal{A}}$ for every $f \in \mathcal{F}_\Sigma$, as well as an interpretation $P^{\mathcal{A}}$ for every $P \in \mathcal{P}_\Sigma$. We further require that $=_\sigma$ be interpreted as the identity relation over $\sigma^{\mathcal{A}}$ for every $\sigma \in \mathcal{S}_\Sigma$. A *$\Sigma$-interpretation* $\mathcal{A}$ is an extension of a Σ-structure with interpretations for some set of variables. For any Σ-term α, $\alpha^{\mathcal{A}}$ denotes the interpretation of α in $\mathcal{A}$. When α is a set of Σ-terms, $\alpha^{\mathcal{A}} = \left\{x^{\mathcal{A}} \mid x \in \alpha\right\}$. Satisfaction is defined as usual. $\mathcal{A} \models \varphi$ denotes that $\mathcal{A}$ satisfies φ.

A Σ-theory $\mathcal{T}$ is a class of all Σ-structures that satisfy some set Ax of Σ-sentences. For each such set Ax, we then say that $\mathcal{T}$ is *axiomatized* by Ax. A Σ-interpretation whose variable-free part is in $\mathcal{T}$ is called a $\mathcal{T}$-interpretation. A Σ-formula ϕ is $\mathcal{T}$-satisfiable if $\mathcal{A} \models \phi$ for some $\mathcal{T}$-interpretation $\mathcal{A}$. A set A of Σ-formulas is $\mathcal{T}$-satisfiable if $\mathcal{A} \models \phi$ for every $\phi \in A$. Two formulas ϕ and ψ are $\mathcal{T}$-*equivalent* if they are satisfied by the same $\mathcal{T}$-interpretations.

Note that for any class $\mathcal{C}$ of Σ-structures there is a natural theory $\mathcal{T}_{\mathcal{C}}$ that *corresponds* to it, and induces the same set of satisfiable formulas: the Σ-theory axiomatized by the set Ax of Σ-sentences that are satisfied in every structure of $\mathcal{C}$. In the examples that follow, we define theories $\mathcal{T}_{\mathcal{C}}$ implicitly by specifying only the class $\mathcal{C}$, as done in the SMT-LIB 2 standard. This can be done without loss of generality.[4]

Example 1. Let Σ_{List} be a signature of finite lists containing the sorts elem_1, elem_2, and list, as well as the function symbols cons of arity $\mathsf{elem}_1 \times \mathsf{elem}_2 \times \mathsf{list} \to \mathsf{list}$, car_1 of arity $\mathsf{list} \to \mathsf{elem}_1$, car_2 of arity $\mathsf{list} \to \mathsf{elem}_2$, cdr of arity $\mathsf{list} \to \mathsf{list}$, and nil of arity list. The Σ_{List}-theory T_{List} corresponds to an SMT-LIB 2 theory of algebraic datatypes [5,7], where elem_1 and elem_2 are interpreted as some sets (of "elements"), and list is interpreted as finite lists of pairs of elements, one from elem_1 and the other from elem_2. cons is a list constructor that takes two elements and a list, and inserts the two elements at the head of the list. The pair $(\mathsf{car}_1(l), \mathsf{car}_2(l))$ is the first entry in l, and $\mathsf{cdr}(l)$ is the list obtained from l by removing its first entry. nil is the empty list. □

Example 2. The signature Σ_{Int} includes a single sort int, all numerals $0, 1, \ldots$, the function symbols $+$, $-$ and $\cdot$ of arity $\mathsf{int} \times \mathsf{int} \to \mathsf{int}$ and the predicate symbols $<$ and $\leq$ of arity $\mathsf{int} \times \mathsf{int}$. The Σ_{Int}-theory T_{Int} corresponds to integer arithmetic in SMT-LIB 2, and the interpretation of the symbols is the same as in the standard structure of the integers. The signature Σ_{BV4} includes a single sort Σ_{BV4} and various function and predicate symbols for reasoning about bit-vectors (finite sequences of bits) of length 4 (such as & for bit-wise *and*, constants of the form 0110, etc.). The Σ_{BV4}-theory T_{BV4} corresponds to SMT-LIB 2 bit-vectors of size 4, with the expected semantics of constants and operators. □

Let Σ_1 and Σ_2 be signatures, $\mathcal{T}_1$ a Σ_1-theory, and $\mathcal{T}_2$ a Σ_2-theory. The *combination* of $\mathcal{T}_1$ and $\mathcal{T}_2$, denoted $\mathcal{T}_1 \oplus T_2$, is the class of all $\Sigma_1 \cup \Sigma_2$-structures $\mathcal{A}$, such that $\mathcal{A}^{\Sigma_1}$ is in $\mathcal{T}_1$ and $\mathcal{A}^{\Sigma_2}$ is in $\mathcal{T}_2$, where $\mathcal{A}^{\Sigma_i}$ is the reduct of $\mathcal{A}$ to Σ_i for $i \in \{1, 2\}$.

Example 3. Let T_{IntBV4} be $T_{\text{Int}} \oplus T_{\text{BV4}}$. It is the combined theory of integers and bit-vectors. It has all the sorts and operators from both theories. If we rename the sorts elem_1 and elem_2 of Σ_{List} to int and $\mathsf{BV4}$, respectively, we can obtain a theory $T_{\text{ListIntBV4}}$ defined as $T_{\text{IntBV4}} \oplus T_{\text{List}}$. This is the theory of lists of pairs, where each pair consists of an integer and a bit-vector of size 4. □

The following theorems will be useful in the sequel.

[4] For further discussion on this point, see Appendix A.1.

Theorem 1 (Theorem 9 of [22]). *Let Σ be a signature, and A a set of Σ-formulas that is satisfiable. Then there exists an interpretation $\mathcal{A}$ that satisfies A, in which $\sigma^{\mathcal{A}}$ is countable whenever it is infinite.*[5]

The following theorem from [15] is a variant of a theorem from [23].

Theorem 2 (Theorem 2.5 of [15]). *For $i = 1, 2$, let Σ_i be disjoint signatures, $S_i = \mathcal{S}_{\Sigma_i}$ with $S = S_1 \cap S_2$, $\mathcal{T}_i$ be a Σ_i-theory, Γ_i be a set of Σ_i-literals, and $V = vars(\Gamma_1) \cap vars(\Gamma_2)$. If there exist a $\mathcal{T}_1$-interpretation $\mathcal{A}$, a $\mathcal{T}_2$ interpretation $\mathcal{B}$, and an arrangement δ_V of V such that: 1. $\mathcal{A} \models \Gamma_1 \cup \delta_V$; 2. $\mathcal{B} \models \Gamma_2 \cup \delta_V$; and 3. $|A_\sigma| = |B_\sigma|$ for every $\sigma \in S$, then $\Gamma_1 \cup \Gamma_2$ is $\mathcal{T}_1 \oplus \mathcal{T}_2$-satisfiable.*

2.2 Polite Theories

We now give the background definitions necessary for both Nelson-Oppen and polite combination. In what follows, Σ is an arbitrary (many-sorted) signature, $S \subseteq \mathcal{S}_\Sigma$, and $\mathcal{T}$ is a Σ-theory. We start with the basic notions of stable infiniteness, smoothness, and arrangements.

Definition 1 (Stably Infinite). *$\mathcal{T}$ is* stably infinite with respect to S *if every quantifier-free Σ-formula that is $\mathcal{T}$-satisfiable is also satisfiable in a $\mathcal{T}$-interpretation $\mathcal{A}$ in which $\sigma^{\mathcal{A}}$ is infinite for every $\sigma \in S$.*

Definition 2 (Smooth). *$\mathcal{T}$ is* smooth *w.r.t. S if for every quantifier-free formula ϕ, $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies ϕ, and function κ from S to the class of cardinals such that $\kappa(\sigma) \geq \left|\sigma^{\mathcal{A}}\right|$ for every $\sigma \in S$, there exists a $\mathcal{T}$-interpretation $\mathcal{A}'$ that satisfies ϕ with $\left|\sigma^{\mathcal{A}'}\right| = \kappa(\sigma)$ for every $\sigma \in S$.*

Definition 3 (Arrangement). *Let V be a finite set of variables whose sorts are in S and let $\{V_\sigma \mid \sigma \in S\}$ be a partition of V such that V_σ is the set of variables of sort σ in V. A formula δ is an* arrangement of V *if*

$$\delta = \bigwedge_{\sigma \in S} \Big(\bigwedge_{(x,y) \in E_\sigma} (x = y) \ \wedge \bigwedge_{x,y \in V_\sigma, (x,y) \notin E_\sigma} (x \neq y) \Big) \ ,$$

where E_σ is some equivalence relation over V_σ for each $\sigma \in S$.

We identify singleton sets with their single elements when there is no ambiguity (e.g., when saying that a theory is smooth w.r.t. a sort σ).

We next define finite witnessability and related concepts, following the presentation in [21]. Let ϕ be a quantifier-free Σ-formula. A Σ-interpretation $\mathcal{A}$ *finitely witnesses ϕ for $\mathcal{T}$ w.r.t. S* (or, is a *finite witness of ϕ for $\mathcal{T}$ w.r.t. S*), if $\mathcal{A} \models \phi$ and $\sigma^{\mathcal{A}} = vars_\sigma(\phi)^{\mathcal{A}}$ for every $\sigma \in S$. We say that ϕ is *finitely witnessed for $\mathcal{T}$ w.r.t. S* if it is either $\mathcal{T}$-unsatisfiable or has a finite witness for $\mathcal{T}$ w.r.t. S. We say that ϕ is *strongly finitely witnessed for $\mathcal{T}$ w.r.t. S* if $\phi \wedge \delta_V$ is finitely witnessed for $\mathcal{T}$ w.r.t. S for every arrangement δ_V of V, where V is any set of

[5] In [22] this was proven more generally, for ordered sorted logics.

Statement	[15]	This Paper
f.w. formulas $\neq$ s.f.w. formulas	Example 3 of [15]	Example 4
witness $\neq$ strong witness	Example 3 of [15]	Example 5
polite $\neq$ strongly polite	—	Section 3.1
1-sort, empty sig: polite = strongly polite	—	Section 3.2
f.w. theories $\neq$ s.f.w. theories	—	Section 3.3

Fig. 1. A summary of the results regarding politeness and strong politeness. The abbreviation (s.f.w) f.w. stands for (strong) finite witnessability.

variables whose sorts are in S. A function $wit : QF(\Sigma) \rightarrow QF(\Sigma)$ is a *(strong) witness for $\mathcal{T}$ w.r.t. S* if for every $\phi \in QF(\Sigma)$ we have that: 1. ϕ and $\exists \overrightarrow{w}.\, wit(\phi)$ are $\mathcal{T}$-equivalent for $\overrightarrow{w} = vars(wit(\phi)) \setminus vars(\phi)$; and 2. $wit(\phi)$ is (strongly) finitely witnessed for $\mathcal{T}$ w.r.t. S.[6]

Definition 4 (Finitely Witnessable). *$\mathcal{T}$ is* (strongly) finitely witnessable *w.r.t. S if there exists a computable (strong) witness for $\mathcal{T}$ w.r.t. S.*

Definition 5 (Polite). *$\mathcal{T}$ is* (strongly) polite w.r.t. S *if it is smooth and (strongly) finitely witnessable w.r.t. S.*

3 Politeness and Strong Politeness

In this section we study the difference between politeness and strong politeness. Since the introduction of strong politeness in [15], it has been unclear whether it is indeed strictly stronger than standard politeness, that is, whether there exists a theory that is polite but not strongly polite. We present an example of such a theory below, answering the open question affirmatively. This result is followed by some further analysis of notions related to politeness.

Figure 1 summarizes the results of this section and compares them to what was already known in [15]. We first distinguish between ordinary and strong finite witnessability of formulas, as well as ordinary and strong witness functions. In [15], this distinction between formulas and functions is not made, with a single example used to explain both (Example 3 of [15]). We elaborate on this example in Examples 4 and 5 using the definitions in this paper and distinguishing between formulas and functions. We then present a polite theory that is not strongly polite in Section 3.1. The theory is over a signature that has two sorts but is otherwise empty. We show in Section 3.2 that for the case of signatures containing only sorts, two sorts are actually needed, since every polite theory is strongly polite when its has no symbols and only one sort. Finally, we show in Section 3.3 that this equivalence does not hold for finite witnessability alone.

[6] The new variables in $wit(\phi)$ are assumed to be fresh not only with respect to ϕ, but also with respect to the formula from the second theory being combined.

We start by showing, with an example adapted from [15], that there are finitely witnessed formulas for $\mathcal{T}$ w.r.t. S that are not strongly finitely witnessed for $\mathcal{T}$ w.r.t. S.

Example 4. Let Σ_0 be a signature with a single sort σ and no function or predicate symbols, and let $\mathcal{T}_0$ be a Σ_0-theory consisting of all Σ_0-structures with at least two elements. Let ϕ to be the formula $x = x \wedge w = w$. This formula is finitely witnessed for $\mathcal{T}_0$ w.r.t. σ, but not strongly. Indeed, for $\delta_V \equiv (x = w)$, $\phi \wedge \delta_V$ is not finitely witnessed for $\mathcal{T}_0$ w.r.t. σ: a finite witness would be required to have only a single element and would therefore not be a $\mathcal{T}_0$-interpretation. □

The next example shows that the notions of a witness and a strong witness are not equivalent.

Example 5. Take Σ_0, σ, and $\mathcal{T}_0$ as in Example 4, and define $wit(\phi)$ as the function $(\phi \wedge w_1 = w_1 \wedge w_2 = w_2)$ for fresh w_1, w_2. The function is a witness for $\mathcal{T}_0$ w.r.t. σ. However, it is not a strong witness for $\mathcal{T}$ w.r.t. σ; the argument is similar to that in Example 4. □

Although the theory $\mathcal{T}_0$ in the above examples does serve to distinguish formulas and functions that are and are not strong, it cannot be used to do the same for theories themselves. This is because $\mathcal{T}_0$ is, in fact, strongly polite, via a different witness function.

Example 6. The function $wit'(\phi) = (\phi \wedge w_1 \neq w_2)$, for some $w_1, w_2 \notin vars_\sigma(\phi)$, *is* a strong witness for $\mathcal{T}_0$ w.r.t. S, as proved in [15]. □

A natural question, then, is whether there is a theory that can separate the two notions of politeness. The following subsection provides an affirmative answer.

3.1 A Polite Theory that is not Strongly Polite

Let Σ_2 be a signature with two sorts σ_1 and σ_2 and no function or predicate symbols (except =). Let $\mathcal{T}_{2,3}$ be the Σ_2-theory consisting of all Σ_2-structures $\mathcal{A}$ such that either $\left|\sigma_1^{\mathcal{A}}\right| = 2 \wedge \left|\sigma_2^{\mathcal{A}}\right| \geq \aleph_0$ or $\left|\sigma_1^{\mathcal{A}}\right| \geq 3 \wedge \left|\sigma_2^{\mathcal{A}}\right| \geq 3$ [11].[7]

$\mathcal{T}_{2,3}$ is polite, but is not strongly polite. Its smoothness is shown by extending any given structure with new elements as much as necessary.

Lemma 1. *$\mathcal{T}_{2,3}$ is smooth w.r.t. $\{\sigma_1, \sigma_2\}$.*

For finite witnessability, consider the function *wit* defined as follows:

$$wit(\phi) := \phi \wedge x_1 = x_1 \wedge x_2 = x_2 \wedge x_3 = x_3 \wedge y_1 = y_1 \wedge y_2 = y_2 \wedge y_3 = y_3, \quad (1)$$

for fresh variables x_1, x_2, and x_3 of sort σ_1 and y_1, y_2, and y_3 of sort σ_2. It can be shown that *wit* is a witness for $\mathcal{T}_{2,3}$ but is not strongly finitely witnessable.

Lemma 2. *$\mathcal{T}_{2,3}$ is finitely witnessable w.r.t. $\{\sigma_1, \sigma_2\}$.*

$$distinct(x_1, \ldots, x_n) := \bigwedge_{1 \le i < j <= n} x_i \neq x_j$$

$$\psi^{\sigma}_{\geq n} := \exists x_1, \ldots, x_n . distinct(x_1, \ldots, x_n)$$

$$\psi^{\sigma}_{\leq n} := \exists x_1, \ldots, x_n . \forall y. \bigvee_{i=1}^{n} y = x_i$$

$$\psi^{\sigma}_{=n} := \psi^{\sigma}_{\geq n} \wedge \psi^{\sigma}_{\leq n}$$

Fig. 2. Cardinality formulas for sort σ. All variables are assumed to have sort σ.

Lemma 3. *$\mathcal{T}_{2,3}$ is not strongly finitely witnessable w.r.t. $\{\sigma_1, \sigma_2\}$.*

Notice that $\mathcal{T}_{2,3}$ can be axiomatized using the following set of axioms:

$$\left\{ \psi^{\sigma_1}_{\geq 2}, \psi^{\sigma_2}_{\geq 3} \right\} \cup \{ \psi^{\sigma_1}_{=2} \rightarrow \neg \psi^{\sigma_2}_{=n} \mid n \geq 3 \}$$

where $\psi^{\sigma}_{\geq k}$ and $\psi^{\sigma}_{=k}$ are defined in Figure 2.

Remark 1. An alternative way to separate politeness from strong politeness uses $\mathcal{T}_{2,3}$ and relies on the two different notions of shininess used in [11] and [20], as well as their respective connections to politeness and strong politeness: $\mathcal{T}_{2,3}$ is shiny, and therefore polite [20], but is not shiny according to the stronger definition from [11] and therefore not strongly polite [11]. However, we have (and prefer) a direct proof based only on politeness, without a detour through shininess.

3.2 The Case of Mono-sorted Polite Theories

Theory $\mathcal{T}_{2,3}$ includes two sorts, but is otherwise empty. In this section we show that requiring two sorts is essential for separating politeness from strong politeness in otherwise empty signatures. That is, we prove that politeness implies strong politeness otherwise. Let Σ_0 be the signature with a single sort σ and no function or predicate symbols (except =), We start by showing that smooth Σ_0-theories have a certain form.

Lemma 4. *Let $\mathcal{T}$ be a Σ_0-theory. If $\mathcal{T}$ is smooth w.r.t. σ and includes at least one finite structure, then there exists $n > 0$, such that $\mathcal{T}$ is axiomatized by $\psi^{\sigma}_{\geq n}$ from Figure 2.*

Proposition 1. *If $\mathcal{T}$ is a Σ_0-theory that is polite w.r.t. σ, then it is strongly polite w.r.t. σ.*

[7] In [11], the first condition is written $\left|\sigma_1^{\mathcal{A}}\right| \geq 2$. We use equality as this is equivalent and we believe it makes things clearer.

Remark 2. We again note (as we did in Remark 1) that an alternative way to obtain this result is via a detour through shiny theories, using [20], which introduced polite theories, as well as [9], which compared strongly polite theories to shiny theories in the mono-sorted case. Given Lemma 4 and finite witnessability, we get that it is decidable whether a finite structure is in $\mathcal{T}$. Proposition 19 of [20] then gives us that $\mathcal{T}$ is shiny. Similar arguments can show that $\mathcal{T}$ is decidable. Then, Proposition 1 of [9] gives us that $\mathcal{T}$ is strongly polite.

3.3 Mono-sorted Finite witnessability

We have seen that for Σ_0-theories, politeness and strong politeness are the same. Now we show that smoothness is crucial for this equivalence, i.e., that there is no such equivalence between finite witnessability and strong finite witnessability. Let $\mathcal{T}^{\infty}_{\text{Even}}$ be the Σ_0-theory of all Σ_0-structures $\mathcal{A}$ such that $\left|\sigma^{\mathcal{A}}\right|$ is even or infinite.[8] Clearly, this theory is not smooth.

Lemma 5. *$\mathcal{T}^{\infty}_{\text{Even}}$ is not smooth w.r.t. σ.*

We can construct a witness *wit* for $\mathcal{T}^{\infty}_{\text{Even}}$ as follows. Let ϕ be a quantifier-free Σ_0-formula, and let E be the set of all equivalence relations over $vars(\phi) \cup \{w\}$ for some fresh variable w. Let $even(E)$ be the set of all equivalence relations in E with an even number of equivalence classes. Then, $wit(\phi)$ is $\phi \wedge \bigvee_{e \in even(E)} \delta_e$, where for each $e \in even(E)$, δ_e is the arrangement induced by e:

$$\bigwedge_{(x,y)\in e} x = y \ \wedge \bigwedge_{x,y \in vars(\phi)\cup\{w\} \wedge (x,y)\notin e} x \neq y$$

It can be shown that *wit* is indeed a witness.

Lemma 6. *$\mathcal{T}^{\infty}_{\text{Even}}$ is finitely witnessable w.r.t. σ.*

Finally, we can show that $\mathcal{T}^{\infty}_{\text{Even}}$ has no strong witness, with a proof similar to the one for Lemma 3.

Lemma 7. *$\mathcal{T}^{\infty}_{\text{Even}}$ is not strongly finitely witnessable w.r.t. σ.*

4 A Blend of Polite and Stably-Infinite Theories

In this section, we show that the polite combination method can be optimized to reduce the search space of possible arrangements. In what follows, Σ_1 and Σ_2 are disjoint signatures, $S = \mathcal{S}_{\Sigma_1} \cap \mathcal{S}_{\Sigma_2}$, $\mathcal{T}_1$ is a Σ_1-theory, $\mathcal{T}_2$ is a Σ_2-theory, Γ_1 is a set of Σ_1-literals, and Γ_2 is a set of Σ_2-literals.

The Nelson-Oppen procedure reduces the $\mathcal{T}_1 \oplus \mathcal{T}_2$-satisfiability of $\Gamma_1 \cup \Gamma_2$ to the existence of an arrangement δ over the set $V = vars_S(\Gamma_1) \cap vars_S(\Gamma_2)$, such that $\Gamma_1 \cup \delta$ is $\mathcal{T}_1$-satisfiable and $\Gamma_2 \cup \delta$ is $\mathcal{T}_2$-satisfiable. The correctness of this reduction relies on the fact that both theories are stably infinite w.r.t. S. In

[8] Notice that $\mathcal{T}^{\infty}_{\text{Even}}$ can be axiomatized using the set $\{\neg\psi^{\sigma}_{=2n+1} \mid n \in \mathbb{N}\}$.

contrast, the polite combination method only requires a condition (namely strong politeness) from one of the theories, while the other theory is unrestricted and, in particular, not necessarily stably infinite. In polite combination, the $\mathcal{T}_1 \oplus \mathcal{T}_2$-satisfiability of $\Gamma_1 \cup \Gamma_2$ is again reduced to the existence of an arrangement δ, but over a different set $V' = vars_S(wit(\Gamma_2))$, such that $\Gamma_1 \cup \delta$ is $\mathcal{T}_1$-satisfiable and $wit(\Gamma_2) \cup \delta$ is $\mathcal{T}_2$-satisfiable, where wit is a strong witness for $\mathcal{T}_2$ w.r.t. S. Thus, the flexibility offered by polite combination comes with a price. The set V' is potentially larger than V as it contains *all* variables with sorts in S that occur in $wit(\Gamma_2)$, not just those that also occur in Γ_1. Since the search space of arrangements over a set grows exponentially with its size, this difference can become crucial. If $\mathcal{T}_1$ happens to be stably infinite w.r.t. S, however, we can fall back to Nelson-Oppen combination and only consider variables that are shared by the two sets. But what if $\mathcal{T}_1$ is stably infinite only w.r.t. to some proper subset $S' \subset S$? Can this knowledge about $\mathcal{T}_1$ help in finding some set V'' of variables between V and V', such that we need only consider arrangements of V''? In this section we prove that this is possible by taking V'' to include only the variables of sorts in S' that are shared between Γ_1 and $wit(\Gamma_2)$, and all the variables of sorts in $S \setminus S'$ that occur in $wit(\Gamma_2)$. We also identify several weaker conditions on $\mathcal{T}_2$ that are sufficient for the combination theorem to hold.

4.1 Refined Combination Theorem

To put the discussion above in formal terms, we recall the following theorem.

Theorem 3 ([15]). *If $\mathcal{T}_2$ is strongly polite w.r.t. S with a witness wit, 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 over V, such that $\Gamma_1 \cup \delta_V$ is $\mathcal{T}_1$-satisfiable and $wit(\Gamma_2) \cup \delta_V$ is $\mathcal{T}_2$-satisfiable, where $V = \bigcup_{\sigma \in S} V_\sigma$, and $V_\sigma = vars_\sigma(wit(\Gamma_2))$ for each $\sigma \in S$.*

Our goal is to identify general cases in which information regarding $\mathcal{T}_1$ can help reduce the size of the set V. We extend the definitions of stably infinite, smooth, and strongly finitely witnessable to two sets of sorts rather than one. Roughly speaking, in this extension, the usual definition is taken for the first set, and some cardinality-preserving constraints are enforced on the second set.

Definition 6. *Let Σ be a signature, S_1, S_2 two disjoint subsets of $\mathcal{S}_\Sigma$, and $\mathcal{T}$ a Σ-theory.*

$\mathcal{T}$ is (strongly) stably infinite w.r.t. (S_1, S_2) *if for every quantifier-free Σ-formula ϕ and $\mathcal{T}$-interpretation $\mathcal{A}$ satisfying ϕ, there exists a $\mathcal{T}$-interpretation $\mathcal{B}$ such that $\mathcal{B} \models \phi$, $|\sigma^{\mathcal{B}}|$ is infinite for every $\sigma \in S_1$, and $|\sigma^{\mathcal{B}}| \leq |\sigma^{\mathcal{A}}|$ ($|\sigma^{\mathcal{B}}| = |\sigma^{\mathcal{A}}|$) for every $\sigma \in S_2$.*

$\mathcal{T}$ is smooth w.r.t. (S_1, S_2) *if for every quantifier-free Σ-formula ϕ, $\mathcal{T}$-interpretation $\mathcal{A}$ satisfying ϕ, and function κ from S_1 to the class of cardinals such that $\kappa(\sigma) \geq \left|\sigma^{A}\right|$ for each $\sigma \in S_1$, there exists a $\mathcal{T}$-interpretation $\mathcal{B}$ that satisfies ϕ, with $\left|\sigma^{B}\right| = \kappa(\sigma)$ for each $\sigma \in S_1$, and with $\left|\sigma^{B}\right|$ infinite whenever $\left|\sigma^{\mathcal{A}}\right|$ is infinite for each $\sigma \in S_2$.*

$\mathcal{T}$ *is* strongly finitely witnessable w.r.t. (S_1, S_2) *if there exists a computable function wit* : $QF(\Sigma) \to QF(\Sigma)$ *such that for every quantifier-free* Σ*-formula* ϕ*: 1.* ϕ *and* $\exists\, \overrightarrow{w}.\ wit(\phi)$ *are* $\mathcal{T}$*-equivalent for* $\overrightarrow{w} = vars(wit(\phi)) \setminus vars(\phi)$*; and 2. for every* $\mathcal{T}$*-interpretation* $\mathcal{A}$ *and arrangement* δ *of any set of variables whose sorts are in* S_1*, if* $\mathcal{A}$ *satisfies* $wit(\phi) \wedge \delta$*, then there exists a* $\mathcal{T}$*-interpretation* $\mathcal{B}$ *that finitely witnesses* $wit(\phi) \wedge \delta$ *w.r.t.* S_1 *and for which* $\left|\sigma^{\mathcal{B}}\right|$ *is infinite whenever* $\left|\sigma^{\mathcal{A}}\right|$ *is infinite, for each* $\sigma \in S_2$*.*

Our main result is the following.

Theorem 4. *Let* $S^{si} \subseteq S$ *and* $S^{nsi} = S \setminus S^{si}$*. Suppose* $\mathcal{T}_1$ *is stably infinite w.r.t.* S^{si} *and one of the following holds:*

1. $\mathcal{T}_2$ *is strongly stably infinite w.r.t.* (S^{si}, S^{nsi}) *and strongly polite w.r.t.* S^{nsi} *with a witness wit.*
2. $\mathcal{T}_2$ *is stably infinite w.r.t.* (S^{si}, S^{nsi})*, smooth w.r.t.* (S^{nsi}, S^{si})*, and strongly finitely witnessable w.r.t.* S^{nsi} *with a witness wit.*
3. $\mathcal{T}_2$ *is stably infinite w.r.t.* S^{si} *and both smooth and strongly finitely-witnessable w.r.t.* (S^{nsi}, S^{si}) *with a witness wit.*

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 *over* V *such that* $\Gamma_1 \cup \delta_V$ *is* $\mathcal{T}_1$*-satisfiable, and* $wit(\Gamma_2) \cup \delta_V$ *is* $\mathcal{T}_2$*-satisfiable, where* $V = \bigcup_{\sigma \in S} V_\sigma$*, with* $V_\sigma = vars_\sigma(wit(\Gamma_2))$ *for every* $\sigma \in S^{nsi}$ *and* $V_\sigma = vars_\sigma(\Gamma_1) \cap vars_\sigma(wit(\Gamma_2))$ *for every* $\sigma \in S^{si}$*.*

The proof of this theorem is provided in Section 4.2, below. Figure 3 is a visualization of the claims in Theorem 4. The theorem considers two variants of strong finite witnessability, two variants of smoothness, and three variants of stable infiniteness. For each of the three cases of Theorem 4, Figure 3 shows which variant of each property is assumed. The height of each bar corresponds to the strength of the property. In the first case, we use ordinary strong finite witnessability and smoothness, but the strongest variant of stable infiniteness; in the second, we use ordinary strong finite witnessability with the new variants of stable infiniteness and smoothness; and for the third, we use ordinary stable infiniteness and the stronger variants of strong finite witnessability and smoothness. The order of the bars corresponds to the order of their usage in the proof of each case. The stage at which stable infiniteness is used determines the required strength of the other properties: whatever is used before is taken in ordinary form, and whatever is used after requires a stronger form.

Going back to the standard definitions of stable infiniteness, smoothness, and strong finite witnessability, we get the following corollary by using case 1 of the theorem and noticing that smoothness w.r.t. S implies strong stable infiniteness w.r.t. any partition of S.

Corollary 1. *Let* $S^{si} \subseteq S$ *and* $S^{nsi} = S \setminus S^{si}$*. Suppose* $\mathcal{T}_1$ *is stably infinite w.r.t.* S^{si} *and* $\mathcal{T}_2$ *is strongly finitely witnessable w.r.t.* S^{nsi} *with a witness wit and smooth w.r.t.* $S^{si} \cup S^{nsi}$*. 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 *over* V *such that* $\Gamma_1 \cup \delta_V$

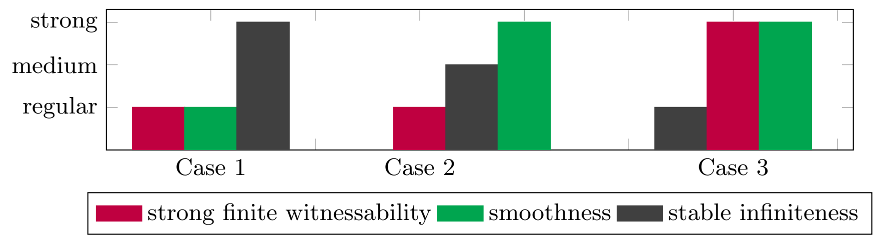

Fig. 3. Theorem 4. The height of each bar corresponds to the strength of the property. The bars are ordered according to their usage in the proof.

is $\mathcal{T}_1$-satisfiable and $wit(\Gamma_2) \cup \delta_V$ is $\mathcal{T}_2$-satisfiable, where $V = \bigcup_{\sigma \in S} V_\sigma$, with $V_\sigma = vars_\sigma(wit(\Gamma_2))$ for each $\sigma \in S^{nsi}$ and $V_\sigma = vars_\sigma(\Gamma_1) \cap vars_\sigma(wit(\Gamma_2))$ for each $\sigma \in S^{si}$.

Finally, the following result, which is closest to Theorem 3, is directly obtained from Corollary 1, since the strong politeness of $\mathcal{T}_2$ w.r.t. $S^{si} \cup S^{nsi}$ implies that it is strongly finitely witnessable w.r.t. S^{nsi} and smooth w.r.t. $S^{si} \cup S^{nsi}$.

Corollary 2. *Let $S^{si} \subseteq S$ and $S^{nsi} = S \setminus S^{si}$. If $\mathcal{T}_1$ is stably infinite w.r.t. S^{si} and $\mathcal{T}_2$ is strongly polite w.r.t. S with a witness wit, then the following are equivalent: 1. $\Gamma_1 \cup \Gamma_2$ is $(\mathcal{T}_1 \oplus \mathcal{T}_2)$-satisfiable; 2. there exists an arrangement δ_V over V such that $\Gamma_1 \cup \delta_V$ is $\mathcal{T}_1$-satisfiable and $wit(\Gamma_2) \cup \delta_V$ is $\mathcal{T}_2$-satisfiable, where $V = \bigcup_{\sigma \in S} V_\sigma$, with $V_\sigma = vars_\sigma(wit(\Gamma_2))$ for each $\sigma \in S^{nsi}$ and $V_\sigma = vars_\sigma(\Gamma_1) \cap vars_\sigma(wit(\Gamma_2))$ for each $\sigma \in S^{si}$.*

Note that a direct proof of Corollary 2 can be obtained by revising the proof of Theorem 3 in [15,20], taking into account the stable infiniteness of $\mathcal{T}_1$. Here, we obtain it as a corollary of the more general Theorem 4 and Corollary 1.

Compared to Theorem 3, Corollary 2 partitions S into S^{si} and S^{nsi} and requires that $\mathcal{T}_1$ be stably infinite w.r.t. S^{si}. The gain from this requirement is that the set V_σ is potentially reduced for $\sigma \in S^{si}$. Note that unlike Theorem 4 and Corollary 1, Corollary 2 has the same assumptions regarding $\mathcal{T}_2$ as the original Theorem 3 from [15]. We show its potential impact in the next example.

Example 7. Consider the theory $T_{\text{ListIntBV4}}$ from Example 3. Let Γ_1 be $x = 5 \wedge v = 0000 \wedge w = w \,\&\, v$, and let Γ_2 be $a_0 = cons(x, v, a_1) \wedge \bigwedge_{i=1}^{n} a_i = cons(y_i, w, a_{i+1})$. Using the witness function wit from [21], $wit(\Gamma_2) = \Gamma_2$. The polite combination approach reduces the $T_{\text{ListIntBV4}}$-satisfiability of $\Gamma_1 \wedge \Gamma_2$ to the existence of an arrangement δ over $\{x, v, w\} \cup \{y_1, \ldots, y_n\}$, such that $\Gamma_1 \wedge \delta$ is T_{IntBV4}-satisfiable and $wit(\Gamma_2) \wedge \delta$ is T_{List}-satisfiable. Corollary 2 shows that we can do better. Since T_{IntBV4} is stably infinite w.r.t. $\{\mathsf{int}\}$, it is enough to check the existence of an arrangement over the variables of sort $\mathsf{BV4}$ that occur in $wit(\Gamma_2)$, together with the variables of sort int that are shared between Γ_1 and

Γ_2. This means that arrangements over $\{x, v, w\}$ are considered, instead of over $\{x, v, w\} \cup \{y_1, \ldots, y_n\}$. As n becomes large, standard polite combination requires considering exponentially more arrangements, while the number of arrangements considered by our combination method remains the same. □

4.2 Proof of Theorem 4

The left-to-right direction is straightforward, using the reducts of the satisfying interpretation of $\Gamma_1 \cup \Gamma_2$ to Σ_1 and Σ_2. We now focus on the right-to-left direction, and begin with the following lemma, which strengthens Theorem 1, obtaining a many-sorted Löwenheim-Skolem Theoerm, where the cardinality of the finite sorts remains the same.

Lemma 8. *Let Σ be a signature, $\mathcal{T}$ a Σ-theory, φ a Σ-formula, and $\mathcal{A}$ a $\mathcal{T}$-interpretation that satisfies ϕ. Let $\mathcal{S}_\Sigma = S_{\mathcal{A}}^{fin} \uplus S_{\mathcal{A}}^{inf}$, where $\sigma^{\mathcal{A}}$ is finite for every $\sigma \in S_{\mathcal{A}}^{fin}$ and $\sigma^{\mathcal{A}}$ is infinite for every $\sigma \in S_{\mathcal{A}}^{inf}$. Then there exists a $\mathcal{T}$-interpretation $\mathcal{B}$ that satisfies φ such that $\left|\sigma^{\mathcal{B}}\right| = \left|\sigma^{\mathcal{A}}\right|$ for every $\sigma \in S_{\mathcal{A}}^{fin}$ and $\sigma^{\mathcal{B}}$ is countable for every $\sigma \in S_{\mathcal{A}}^{inf}$.*

The proof of Theorem 4 continues with the following main lemma.

Lemma 9 (Main Lemma). *Let $S^{si} \subseteq S$ and $S^{nsi} = S \setminus S^{si}$, Suppose $\mathcal{T}_1$ is stably infinite w.r.t. S^{si} and that one of the three cases of Theorem 4 holds. Further, assume there exists an arrangement δ_V over V such that $\Gamma_1 \cup \delta_V$ is $\mathcal{T}_1$-satisfiable, and $wit(\Gamma_2) \cup \delta_V$ is $\mathcal{T}_2$-satisfiable, where $V = \bigcup_{\sigma \in S} V_\sigma$, with $V_\sigma = vars_\sigma(wit(\Gamma_2))$ for each $\sigma \in S^{nsi}$ and $V_\sigma = vars_\sigma(\Gamma_1) \cap vars_\sigma(wit(\Gamma_2))$ for each $\sigma \in S^{si}$. Then, there is a $\mathcal{T}_1$-interpretation $\mathcal{A}$ that satisfies $\Gamma_1 \cup \delta_V$ and a $\mathcal{T}_2$-interpretation $\mathcal{B}$ that satisfies $wit(\Gamma_2) \cup \delta_V$ such that $\left|\sigma^{\mathcal{A}}\right| = \left|\sigma^{\mathcal{B}}\right|$ for all $\sigma \in S$.*

Proof: Let $\psi_2 := wit(\Gamma_2)$. Since $\mathcal{T}_1$ is stably infinite w.r.t. S^{si}, there is a $\mathcal{T}_1$-interpretation $\mathcal{A}$ satisfying $\Gamma_1 \cup \delta_V$ in which $\sigma^{\mathcal{A}}$ is infinite for each $\sigma \in S^{si}$. By Theorem 1, we may assume that $\sigma^{\mathcal{A}}$ is countable for each $\sigma \in S^{si}$. We consider the first case of Theorem 4 (the others are omitted due to space constraints). Suppose $\mathcal{T}_2$ is strongly stably infinite w.r.t. (S^{si}, S^{nsi}) and strongly polite w.r.t. S^{nsi}. Since $\mathcal{T}_2$ is strongly finitely-witnessable w.r.t. S^{nsi}, there exists a $\mathcal{T}_2$-interpretation $\mathcal{B}$ that satisfies $\psi_2 \cup \delta_V$ such that $\sigma^{\mathcal{B}} = V_\sigma^{\mathcal{B}}$ for each $\sigma \in S^{nsi}$. Since $\mathcal{A}$ and $\mathcal{B}$ satisfy δ_V, we have that for every $\sigma \in S^{nsi}$, $\left|\sigma^{\mathcal{B}}\right| = \left|V_\sigma^{\mathcal{B}}\right| = \left|V_\sigma^{\mathcal{A}}\right| \leq \left|\sigma^{\mathcal{A}}\right|$. $\mathcal{T}_2$ is also smooth w.r.t. S^{nsi}, and so there exists a $\mathcal{T}_2$-interpretation $\mathcal{B}'$ satisfying $\psi_2 \cup \delta_V$ such that $\left|\sigma^{\mathcal{B}'}\right| = \left|\sigma^{\mathcal{A}}\right|$ for each $\sigma \in S^{nsi}$. Finally, $\mathcal{T}_2$ is strongly stably infinite w.r.t. (S^{si}, S^{nsi}), so there exists a $\mathcal{T}_2$-interpretation $\mathcal{B}''$ that satisfies $\psi_2 \cup \delta_V$ such that $\sigma^{\mathcal{B}''}$ is infinite for each $\sigma \in S^{si}$ and $\left|\sigma^{\mathcal{B}''}\right| = \left|\sigma^{\mathcal{B}'}\right| = \left|\sigma^{\mathcal{A}}\right|$ for each $\sigma \in S^{nsi}$. By Lemma 8, we may assume that $\sigma^{\mathcal{B}''}$ is countable for each $\sigma \in S^{si}$. Thus, $\left|\sigma^{\mathcal{B}''}\right| = \left|\sigma^{\mathcal{A}}\right|$ for each $\sigma \in S$. □

We now conclude the proof of Theorem 4. Lemma 9 gives us a $\mathcal{T}_1$ interpretation $\mathcal{A}$ such that $\mathcal{A} \models \Gamma_1 \cup \delta_V$ and a $\mathcal{T}_2$ interpretation $\mathcal{B}$ with $\mathcal{B} \models \psi_2 \cup \delta_V$, and $\left|\sigma^{\mathcal{A}}\right| = \left|\sigma^{\mathcal{B}}\right|$ for each $\sigma \in S$. Now, take $\Gamma_1' := \Gamma_1 \cup \delta_V$ and $\Gamma_2' := \psi_2 \cup \delta_V$. Then, $V_\sigma = vars_\sigma(\Gamma_1') \cap vars_\sigma(\Gamma_2')$ for each $\sigma \in S$. Clearly, $\mathcal{A} \models \Gamma_1' \cup \delta_V$ and $\mathcal{B} \models \Gamma_2' \cup \delta_V$. Also, $|\sigma^{\mathcal{A}}| = |\sigma^{\mathcal{B}}|$ for each $\sigma \in S$. By Theorem 2, $\Gamma_1' \cup \Gamma_2'$ is $\mathcal{T}_1 \oplus \mathcal{T}_2$-satisfiable. In particular, $\Gamma_1 \cup \{\psi_2\}$ is $\mathcal{T}_1 \oplus \mathcal{T}_2$-satisfiable, and hence also $\Gamma_1 \cup \{\exists \overline{w} \psi_2\}$, where $\overline{w} = vars(wit(\Gamma_2)) \setminus vars(\Gamma_2)$. Finally, recall that $\exists \overline{w} wit(\Gamma_2)$ is $\mathcal{T}_2$-equivalent to Γ_2, and hence $\Gamma_1 \cup \Gamma_2$ is $\mathcal{T}_1 \oplus \mathcal{T}_2$-satisfiable. □

5 Preliminary Case Study

The optimization to the polite combination method presented in Section 4 was motivated by a set of smart contract verification benchmarks. We obtained these benchmarks by applying the open-source Move Prover verifier [25] to smart contracts found in the open-source Diem project [12]. The Move prover is a formal verifier for smart contracts written in the Move language [8] and was designed to target smart contracts used in the Diem blockchain [1]. It works via a translation to the Boogie verification framework [17], which in turn produces SMT-LIB 2 benchmarks that are dispatched to SMT solvers. The benchmarks we obtained involve datatypes, integers, Booleans, and quantifiers. Our case study began by running CVC4 [4] on the benchmarks. For most of the benchmarks that were solved by CVC4, theory combination took a small percentage of the overall runtime of the solver, accounting for 10% or less in all but 1 benchmark. However, solving that benchmark took 81 seconds, of which 20 seconds was dedicated to theory combination.

We implemented an optimization to the datatype solver of CVC4 based on Corollary 2. With the original polite combination method, every term that originates from the theory of datatypes with an integer, Boolean, or uninterpreted sort is shared with the other theories, triggering an analysis of the arrangements of these terms. In our optimization, we limit the sharing of such terms to those of Boolean sort. In the language of Corollary 2, $\mathcal{T}_1$ is the combined theory of Booleans, uninterpreted functions, and integers, which is stably infinite w.r.t. the uninterpreted sorts and the integer sort. $\mathcal{T}_2$ is an instance of the theory of datatypes, which is strongly polite w.r.t its element sorts, which in this case include all the sorts of $\mathcal{T}_1$.

A comparison of an original and optimized run on the difficult benchmark is shown in Figure 4. As shown, the optimization reduces the total running time by 75%, and the time spent on theory combination in particular by 91%, from 20 seconds to 3 seconds. To further isolate the effectiveness of our optimization on this benchmark, we report metrics on the number of terms that each theory solver involved in this problem considered. In CVC4, constraints are not flattened, and so CVC4 deals with shared *terms* instead of shared variables. Also, each theory solver maintains its own data structure for tracking equality and congruence information. These data structures contain terms belonging to the

	total (s)	comb (s)	DT	INT	UFB	shared
optimized	34.9	3.4	236.1	212.1	78.4	125.8
original	81.5	20.3	116.0	281.0	123.9	163.5

Fig. 4. Running times (in seconds) and number of terms (in thousands) added to the equality data structures of different theories (DT, INT, UFB), as well as the number of shared terms (shared).

theory that either come from the input assertions or are shared (i.e., sent from another theory). A data structure is also maintained that contains all shared terms belonging to any theory. The last 4 columns of Figure 4 count the number of times (in thousands) a term was added to the equality data structure for the theory of datatypes (DT), integers (INT), and uninterpreted functions and Booleans (UFB), as well as to the the shared term data structure (shared). With the optimization, the datatype theory solver keeps more inferred assertions internally, which leads to an increase in the number of additions of terms to its internal data structure. However, by sharing fewer terms, the number of terms in the data structures for the other theories is reduced. Moreover, while the total number of terms considered by the three theories remains roughly the same (526.7 for optimized and 521.0 for original), the number of shared terms decreases by 24%, from 163.5 to 125.8. This suggests that in the two runs of this benchmark, although the workload on the individual theory solvers is roughly similar, a decrease in the number of shared terms in the optimized run results in a significant improvement in the overall runtime. Although our evidence is only anecdotal at the moment, we believe this benchmark is highly representative of the potential benefits of our optimization.

6 Conclusion

This paper makes two contributions: first, we separated politeness and strong politeness, which shows that sometimes, the (typically harder) task of finding a strong witness is necessary. Then, we provided an optimization to the polite combination method, which is available when one of the theories in the combination is stably infinite. This optimization has the potential to improve the performance of SMT solvers that implement polite combination, such as CVC4.

We envision several directions for future work. First, the existence of a polite theory that is not strongly polite demonstrates a need to identify sufficient criteria for the equivalence of these notions — such as, for instance, the *additivity* criterion introduced by Sheng et al. [21]. Second, polite combination could be optimized in a different way: by applying the witness function only to part of the purified input formula (perhaps the one that involves sorts for which the other theory is not stably infinite). We plan to pursue both of these directions. Finally, we plan to extend the initial implementation of this approach in CVC4 and evaluate its impact in practice based on more benchmarks.

References

1. Amsden, Z., Arora, R., Bano, S., Baudet, M., Blackshear, S., Bothra, A., Cabrera, G., Catalini, C., Chalkias, K., Cheng, E., Ching, A., Chursin, A., Danezis, G., Giacomo, G.D., Dill, D.L., Ding, H., Doudchenko, N., Gao, V., Gao, Z., Garillot, F., Gorven, M., Hayes, P., Hou, J.M., Hu, Y., Hurley, K., Lewi, K., Li, C., Li, Z., Malkhi, D., Margulis, S., Maurer, B., Mohassel, P., de Naurois, L., Nikolaenko, V., Nowacki, T., Orlov, O., Perelman, D., Pott, A., Proctor, B., Qadeer, S., Rain, Russi, D., Schwab, B., Sezer, S., Sonnino, A., Venter, H., Wei, L., Wernerfelt, N., Williams, B., Wu, Q., Yan, X., Zakian, T., Zhou, R.: The Libra Blockchain. https://developers.diem.com/docs/technical-papers/the-diem-blockchain-paper/ (2019)
2. Baader, F., Schulz, K.U.: Combination techniques and decision problems for disunification. Theor. Comput. Sci. **142**(2), 229–255 (1995)
3. Baader, F., Schulz, K.U.: Unification in the union of disjoint equational theories: Combining decision procedures. J. Symb. Comput. **21**(2), 211–243 (1996)
4. Barrett, C., Conway, C.L., Deters, M., Hadarean, L., Jovanović, D., King, T., Reynolds, A., Tinelli, C.: CVC4. In: Proceedings of the 23rd International Conference on Computer Aided Verification. pp. 171–177. CAV'11, Springer-Verlag (2011), http://dl.acm.org/citation.cfm?id=2032305.2032319
5. Barrett, C., Fontaine, P., Tinelli, C.: The SMT-LIB Standard: Version 2.6. Tech. rep., Department of Computer Science, The University of Iowa (2017), available at `www.SMT-LIB.org`
6. Barrett, C., Tinelli, C.: Satisfiability modulo theories. In: Handbook of Model Checking, pp. 305–343. Springer (2018)
7. Barrett, C.W., Shikanian, I., Tinelli, C.: An abstract decision procedure for a theory of inductive data types. Journal on Satisfiability, Boolean Modeling and Computation **3**(1-2), 21–46 (2007)
8. Blackshear, S., Cheng, E., Dill, D.L., Gao, V., Maurer, B., Nowacki, T., Pott, A., Qadeer, S., Rain, Russi, D., Sezer, S., Zakian, T., Zhou, R.: Move: A language with programmable resources. https://developers.diem.com/docs/technical-papers/move-paper/ (2019)
9. Casal, F., Rasga, J.: Revisiting the equivalence of shininess and politeness. In: McMillan, K., Middeldorp, A., Voronkov, A. (eds.) Logic for Programming, Artificial Intelligence, and Reasoning. pp. 198–212. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)
10. Casal, F., Rasga, J.: Equivalence of shiny and strongly polite theories (2014), https://pdfs.semanticscholar.org/109f/d4f188663a28076b7eb7931f69c5ffc7b3c5.pdf, unpublished
11. Casal, F., Rasga, J.: Many-sorted equivalence of shiny and strongly polite theories. J. Autom. Reasoning **60**(2), 221–236 (2018)
12. diem: https://github.com/diem/diem
13. Enderton, H.B.: A mathematical introduction to logic. Academic Press (2001)
14. Fontaine, P.: Combinations of theories for decidable fragments of first-order logic. In: Ghilardi, S., Sebastiani, R. (eds.) Frontiers of Combining Systems, 7th International Symposium, FroCoS 2009, Trento, Italy, September 16-18, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5749, pp. 263–278. Springer (2009)
15. Jovanovic, D., Barrett, C.W.: Polite theories revisited. In: Fermüller, C.G., Voronkov, A. (eds.) Logic for Programming, Artificial Intelligence, and Reasoning - 17th International Conference, LPAR-17, Yogyakarta, Indonesia, October 10-15, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6397, pp. 402–416.

Springer (2010), extended technical report is available at http://theory.stanford.edu/~barrett/pubs/JB10-TR.pdf
16. Krstic, S., Goel, A., Grundy, J., Tinelli, C.: Combined satisfiability modulo parametric theories. In: Grumberg, O., Huth, M. (eds.) Tools and Algorithms for the Construction and Analysis of Systems, 13th International Conference, TACAS 2007, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2007 Braga, Portugal, March 24 - April 1, 2007, Proceedings. Lecture Notes in Computer Science, vol. 4424, pp. 602–617. Springer (2007)
17. Leino, K.R.M.: This is boogie 2 (2008), https://www.microsoft.com/en-us/research/publication/this-is-boogie-2-2/, manuscript KRML 178
18. Nelson, G.: Techniques for program verification. Tech. Rep. CSL-81-10, Xerox, Palo Alto Research Center (1981)
19. Nelson, G., Oppen, D.C.: Simplification by cooperating decision procedures. ACM Trans. Program. Lang. Syst. **1**(2), 245–257 (1979)
20. Ranise, S., Ringeissen, C., Zarba, C.G.: Combining data structures with nonstably infinite theories using many-sorted logic. In: Gramlich, B. (ed.) Frontiers of Combining Systems, 5th International Workshop, FroCoS 2005, Vienna, Austria, September 19-21, 2005, Proceedings. Lecture Notes in Computer Science, vol. 3717, pp. 48–64. Springer (2005), extended technical report is available at https://hal.inria.fr/inria-00070335/
21. Sheng, Y., Zohar, Y., Ringeissen, C., Lange, J., Fontaine, P., Barrett, C.: Politeness for the theory of algebraic datatypes. In: Peltier, N., Sofronie-Stokkermans, V. (eds.) Automated Reasoning. pp. 238–255. Springer International Publishing, Cham (2020)
22. Tinelli, C., Zarba, C.G.: Combining decision procedures for sorted theories. In: Alferes, J.J., Leite, J.A. (eds.) Logics in Artificial Intelligence, 9th European Conference, JELIA 2004, Lisbon, Portugal, September 27-30, 2004, Proceedings. Lecture Notes in Computer Science, vol. 3229, pp. 641–653. Springer (2004)
23. Tinelli, C., Zarba, C.G.: Combining decision procedures for sorted theories. In: JELIA. Lecture Notes in Computer Science, vol. 3229, pp. 641–653. Springer (2004)
24. Tinelli, C., Zarba, C.G.: Combining nonstably infinite theories. J. Autom. Reasoning **34**(3), 209–238 (2005)
25. Zhong, J.E., Cheang, K., Qadeer, S., Grieskamp, W., Blackshear, S., Park, J., Zohar, Y., Barrett, C.W., Dill, D.L.: The move prover. In: CAV (1). Lecture Notes in Computer Science, vol. 12224, pp. 137–150. Springer (2020)

A Appendix

A.1 Theories vs. Classes of Structures

In papers about theory combination, theories are often defined in terms of some set Ax of sentences (*axioms*) (see, e.g., [11,22,15]). Specifically, a theory is defined as the set of all sentences entailed by Ax or, interchangeably, as the class of all structures that satisfy Ax. This is the approach we take in this paper. The main reason for this is that the combination theorems we prove and cite here rely on some forms of the Löwenheim-Skolem theorem, which do not hold for arbitrary classes of structures, but do hold when defining theories this way. On the other hand, theories in the SMT-LIB 2 standard, as well as in many SMT papers about individual theories, are defined more generally as classes of structures without reference to a set of axioms.

However, this discrepancy is not substantial since the two notions of a theory as a class of structures are easily interreducible; as mentioned in the introduction, every theory T in the second, more general sense induces a theory in the first sense that is equivalent to T for all of our intents and purposes since it entails exactly the same sentences as T. To be more precise, the combination theorems that we prove and cite only regard satisfiability of formulas in a theory (though their proofs may analyze the structures of a theory). The important thing is that the transformation between the two notions preserves satisfiability, and therefore interchanging these notions can be done without loss of generality. For completeness, we prove this fact below:

Lemma 10. *Let Σ be a signature, $\mathcal{C}$ a class of Σ-structures, Ax the set of Σ-sentences satisfied by all structures of $\mathcal{C}$, and $\mathcal{T}_{\mathcal{C}}$ the class of all Σ-structures that satisfy all sentences of Ax. Then, for every Σ-formula φ, φ is $\mathcal{T}_{\mathcal{C}}$-satisfiable iff φ is satisfied by some Σ-interpretation whose variable-free part is in $\mathcal{C}$.*

Proof: Every interpretation whose variable-free part is in $\mathcal{C}$ is a $\mathcal{T}_{\mathcal{C}}$-interpretation, and so the right-to-left direction trivially holds. Now, suppose φ is not satisfied by any Σ-interpretation whose variable-free part is in $\mathcal{C}$. Then its existential closure $\exists\overline{x}\varphi$ is not satisfied by any structure of $\mathcal{C}$, and hence $\neg\exists\overline{x}\varphi \in Ax$. *Ad absurdum*, suppose that φ is $\mathcal{T}_{\mathcal{C}}$-satisfiable. Then there is a $\mathcal{T}_{\mathcal{C}}$-interpretation $\mathcal{A}$ such that $\mathcal{A} \models \varphi$. In particular, $\mathcal{A} \models \exists\overline{x}\varphi$. But since $\mathcal{A}$ is a $\mathcal{T}_{\mathcal{C}}$-interpretation, we must also have $\mathcal{A} \models \neg\exists\overline{x}\varphi$, which is a contradiction. $\square$

A.2 Proof of Lemma 1

Let ϕ be a quantifier-free Σ_2-formula, $\mathcal{A}$ a $\mathcal{T}_{2,3}$-interpretation that satisfies ϕ and κ a function from $\{\sigma_1, \sigma_2\}$ to the class of cardinals such that $\kappa(\sigma_1) \geq \left|\sigma_1^{\mathcal{A}}\right|$ and $\kappa(\sigma_2) \geq \left|\sigma_2^{\mathcal{A}}\right|$. We construct a Σ_2-interpretation $\mathcal{A}'$ as follows. For $i \in \{1, 2\}$, we let $\sigma_i^{\mathcal{A}'} := \sigma_i^{\mathcal{A}} \uplus B$ for some set $\mathcal{B}$ of countable cardinality if $\kappa(\sigma_i)$ is infinite or of cardinality $\kappa(\sigma_i) - |\sigma_i|^{\mathcal{A}}$ otherwise. Notice that this is well defined because

$\kappa(\sigma_i) \geq \left|\sigma_i^{\mathcal{A}}\right|$. As for variables, $x^{\mathcal{A}'} := x^{\mathcal{A}}$ for each variable in $vars(\phi)$. This is well defined because the domains of σ_1 and σ_2 were only possibly extended, not reduced. First, we prove that $\mathcal{A}'$ is a $\mathcal{T}_{2,3}$-interpretation. If $\kappa(\sigma_1) = 2$, then since $\kappa(\sigma_1) \geq \left|\sigma_1^{\mathcal{A}}\right|$, we must have that $\left|\sigma_1^{\mathcal{A}}\right| = 2$, which means that $|\sigma_2|^{\mathcal{A}}$ is infinite, which in turn means that $\kappa(\sigma_2)$ is infinite as well. Hence in this case we have $\left|\sigma_1^{\mathcal{A}'}\right| = \kappa(\sigma_1) = 2$ and $\left|\sigma_2^{\mathcal{A}'}\right| = \kappa(\sigma_2) = \infty$. Otherwise, $\kappa(\sigma_1) \geq 3$, and hence $\left|\sigma_1^{\mathcal{A}'}\right| = \kappa(\sigma_1) \geq 3$ and also $\left|\sigma_2^{\mathcal{A}'}\right| = \kappa(\sigma_2) \geq \left|\sigma_2^{\mathcal{A}}\right| \geq 3$. Clearly, $\mathcal{A}'$ satisfies ϕ as the interpretations of variables did not change. Finally, $\left|\sigma_1^{\mathcal{A}'}\right| = \kappa(\sigma_1)$ and $\left|\sigma_2^{\mathcal{A}'}\right| = \kappa(\sigma_2)$ by construction.

□

A.3 Proof of Lemma 2

Define a function *wit* by $wit(phi) := phi \wedge x_1 = x_1 \wedge x_2 = x_2 \wedge x_3 = x_3 \wedge y_1 = y_1 \wedge y_2 = y_2 \wedge y_3 = y_3$ for fresh variables x_1, x_2 and x_3 of sort σ_1 and y_1, y_2 and y_3 of sort σ_2. We prove that *wit* is a witness for $\mathcal{T}_{2,3}$ w.r.t. $\{\sigma_1, \sigma_2\}$. ϕ and $\exists x_1, x_2, x_3, y_1, y_2, y_3, wit(\phi)$ are trivially logically equivalent and in particular $\mathcal{T}_{2,3}$-equivalent. We prove that $wit(\phi)$ is finitely witnessed for $\mathcal{T}_{2,3}$ w.r.t. $\{\sigma_1, \sigma_2\}$. Suppose that $wit(\phi)$ is $\mathcal{T}_{2,3}$-satisfiable and let $\mathcal{A}$ be a satisfying $\mathcal{T}_{2,3}$-interpretation. Define a Σ_2-interpretation $\mathcal{B}$ simply by $\sigma_1^{\mathcal{B}} = vars_{\sigma_1}(\phi)^{\mathcal{A}} \uplus \{a_1, a_2, a_3\}$ and $\sigma_2^{\mathcal{B}} = vars_{\sigma_2}(\phi)^{\mathcal{A}} \uplus \{b_1, b_2, b_3\}$ for $a_1, a_2, a_3 \notin \sigma_1^{\mathcal{A}}$ and $b_1, b_2, b_3 \notin \sigma_2^{\mathcal{A}}$. The interpretations of variables from ϕ are the same as in $\mathcal{A}$. As for the fresh variables $x_i^{\mathcal{B}} := a_i$ and $y_i^{\mathcal{B}} := b_i$ for $i \in \{1, 2, 3\}$. We prove that $\mathcal{B}$ finitely witnesses $wit(\phi)$ for $\mathcal{T}_{2,3}$ w.r.t. $\{\sigma_1, \sigma_2\}$. First, $\mathcal{B}$ is a $\mathcal{T}_{2,3}$-interpretation, as by construction $\left|\sigma_1^{\mathcal{B}}\right|, \left|\sigma_2^{\mathcal{B}}\right| \geq 3$. Second, $\mathcal{B} \models \phi$ as the interpretations of variables from ϕ did not change, and trivially satisfies the new identities, and so $\mathcal{B} \models wit(\phi)$. Third, by construction $\sigma_1^{\mathcal{B}} = vars_{\sigma_1}(\phi)^{\mathcal{A}} \uplus \{a_1, a_2, a_3\} = vars_{\sigma_1}(\phi)^{\mathcal{B}} \uplus \left\{x_1^{\mathcal{B}}, x_2^{\mathcal{B}}, x_3^{\mathcal{B}}\right\} = vars_{\sigma_1}(wit(\phi))^{\mathcal{B}}$, and similarly for σ_2.

□

A.4 Proof of Lemma 3

Let *wit* be a witness for $\mathcal{T}_{2,3}$ w.r.t. $\{\sigma_1, \sigma_2\}$. We show that it is not strong. In particular, we show that $wit(v = v)$ is not strongly finitely witnessed for $\mathcal{T}_{2,3}$ w.r.t. $\{\sigma_1, \sigma_2\}$. Consider a $\mathcal{T}_{2,3}$-interpretation $\mathcal{A}$ with $\left|\sigma_1^{\mathcal{A}}\right| = 2$ and $\left|\sigma_2^{\mathcal{A}}\right| = \aleph_0$. Clearly, $\mathcal{A} \models v = v$, and so $\mathcal{A} \models \exists\, \overline{w}.\ wit(v = v)$, with $\overline{w}$ being the variables in $wit(v = v)$ other than v. This in turn means that there is a $\mathcal{T}_{2,3}$-interpretation $\mathcal{A}'$ that satisfies $wit(v = v)$, different from $\mathcal{A}$ only in the interpretations of $\overline{w}$, if anywhere. Let δ be the arrangement over $vars(wit(v = v))$ induced by $\mathcal{A}'$. Then, δ either asserts that all variables in $vars_{\sigma_1}(wit(v = v))$ are identical, or it partitions them into two equivalence classes. $\mathcal{A}' \models wit(v = v) \wedge \delta$, and so $wit(v = v) \wedge \delta$ is $\mathcal{T}_{2,3}$-satisfiable. We show that it does not have a finite witness for $\mathcal{T}_{2,3}$ w.r.t. S. Suppose for contradiction that $\mathcal{B}$ is a finite witness of $wit(v = v) \wedge \delta$

for $\mathcal{T}_{2,3}$ w.r.t. S. Then $\left|\sigma_1^{\mathcal{B}}\right| = \left|vars_{\sigma_1}(wit(v = v) \wedge \delta)^{\mathcal{B}}\right|$. Now, $\mathcal{B} \models \delta$ and $\mathcal{B}$ is a $\mathcal{T}_{2,3}$-interpretation, meaning $\left|\sigma_1^{\mathcal{B}}\right| \geq 2$, so if δ requires all variables of sort σ_1 to be equal, we already have a contradiction. On the other hand, if δ partitions the variables into two equivalence classes, we get that $\left|\sigma_1^{\mathcal{B}}\right| = 2$. But since $\mathcal{B}$ finitely witnesses $wit(v = v) \wedge \delta$ for $\mathcal{T}_{2,3}$ w.r.t. $\{\sigma_1, \sigma_2\}$, we also get that $\sigma_2^{\mathcal{B}}$ is finite, meaning $\mathcal{B}$ is not a $\mathcal{T}_{2,3}$-interpretation. □

A.5 Proof of Lemma 4

Let $\mathcal{A}$ be the $\mathcal{T}$-structure with a minimal number of elements, and let $n = \left|\sigma^{\mathcal{A}}\right|$. To show that every Σ_0-structure that satisfies $\psi_{\geq n}^{\sigma}$ belongs to $\mathcal{T}$, let $\mathcal{B}$ be a Σ_0-structure that satisfies $\psi_{\geq n}^{\sigma}$ and let m be the cardinality of $\sigma^{\mathcal{B}}$. Then $m \geq n$. Clearly, $\mathcal{A} \models x = x$ and has n elements. Since $\mathcal{T}$ is smooth w.r.t. σ, there exists a $\mathcal{T}$-interpretation (that satisfies $x = x$) with cardinality m. This interpretation must be $\mathcal{B}$, as the lack of any symbols means that the only thing that distinguishes between Σ_0-structures is their cardinality (modulo isomorphism). For the converse, note that by the choice of n as minimal, every $\mathcal{T}$-structure satisfies $\psi_{\geq n}^{\sigma}$. □

A.6 Proof of Proposition 1

$x = x$ is clearly $\mathcal{T}$-satisfiable. Since $\mathcal{T}$ is finitely witnessable (say with witness wit), there is a $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies $wit(x = x)$ such that $\sigma^{\mathcal{A}}$ is finite. $\mathcal{T}$ is smooth, and hence, by Lemma 4, it is axiomatized by $\psi_{\geq n}^{\sigma}$ for some n. Define $wit'(\phi) := \phi \wedge distinct(x_1, \ldots, x_n)$ for fresh $x_1, \ldots, x_n$. Since $\mathcal{T}$ is axiomatized by $\psi_{\geq n}^{\sigma}$, ϕ is $\mathcal{T}$-equivalent to $\exists \overline{x}.wit'(\phi)$. Further, for any arrangement δ over some set of variables, and any $\mathcal{T}$-interpretation $\mathcal{A}'$ that satisfies $wit'(\phi) \wedge \delta$, if the domain of $\mathcal{A}'$ is reduced to contain only the elements in $vars(wit'(\phi) \wedge \delta)^{\mathcal{A}'}$, the result is still a $\mathcal{T}$-interpretation since $wit'(\phi)$ contains $distinct(x_1, \ldots, x_n)$. We therefore get that wit' is a strong witness for $\mathcal{T}$ w.r.t. σ. □

A.7 Proof of Lemma 5

Let ϕ be $x = x$ and $\mathcal{A}$ be a Σ-interpretation with $\sigma^{\mathcal{A}} = \{1, 2\}$ and $x^{\mathcal{A}} = 1$. Then $\mathcal{A}$ is a $\mathcal{T}_{\text{Even}}^{\infty}$-interpretation that satisfies ϕ. Let κ defined by $\kappa(s) = 3$. Then $3 = \kappa(s) \geq \left|\sigma^{\mathcal{A}}\right| = 2$. However, there is no Σ-interpretation $\mathcal{A}'$ with $\left|\sigma^{\mathcal{A}'}\right| = 3$. □

A.8 Proof of Lemma 6

Define $wit(\phi)$ as follows. Let E be the set of all equivalence relations over $vars(\phi) \cup \{w\}$ for some fresh variable w. Let $even(E)$ be the set of all equivalence relations in E for which the number of equivalence classes is even. Then, $wit(\phi)$

is $\phi \wedge \bigvee_{e \in even(E)} \delta_e$, where for an equivalence relation $e \in even(E)$, δ_e is the arrangement induced by e:

$$\bigwedge_{(x,y) \in e} x = y \wedge \bigwedge_{x,y \in vars(\phi) \cup \{w\} \wedge (x,y) \notin e} x \neq y$$

We prove that wit is a witness. Let ϕ be a Σ-formula. We first prove that it is $\mathcal{T}^{\infty}_{\text{Even}}$-equivalent to $\exists\, w.\ wit(\phi)$. Since ϕ is a conjunct of $wit(\phi)$ that does not include w, every $\mathcal{A}$-interpretation that satisfies $wit(\phi)$ also satisfies ϕ. For the other direction, let $\mathcal{A}$ be a $\mathcal{T}^{\infty}_{\text{Even}}$-interpretation satisfying ϕ. Even though $\mathcal{A}$ may have infinitely many elements, the number of elements in $vars(\phi)^{\mathcal{A}}$ must be finite. If the number of elements in $vars(\phi)^{\mathcal{A}}$ is even, then let a be some arbitrary element of $vars(\phi)^{\mathcal{A}}$. Otherwise, let a be an element in $\mathcal{A}$ different from all the elements in $vars(\phi)^{\mathcal{A}}$ (there must be such an element since $\mathcal{A}$ has an even or infinite number of elements). In either case, the number of elements in $(vars(\phi) \cup \{w\})^{\mathcal{A}}$ is even. Thus, if we modify $\mathcal{A}$ to map w to a, then it must satisfy one of the disjuncts in $wit(\phi)$. Hence, $\mathcal{A}$ satisfies $\exists\, w.\ wit(\phi)$.

Next, if $wit(\phi)$ is $\mathcal{T}^{\infty}_{\text{Even}}$-satisfiable, then there is a satisfying $\mathcal{T}^{\infty}_{\text{Even}}$-interpretation $\mathcal{A}$ satisfying it. $\mathcal{A}$ must satisfy one of the disjuncts in $wit(\phi)$, which means $\left|vars(wit(\phi))^{\mathcal{A}}\right|$ is even. The restriction of $\mathcal{A}$ to $vars(wit(\phi))^{\mathcal{A}}$ is a $\mathcal{T}^{\infty}_{\text{Even}}$-interpretation that finitely witnesses $wit(\phi)$. □

A.9 Proof of Lemma 7

Let $wit : QF(\Sigma_0) \rightarrow QF(\Sigma_0)$. We prove that wit is not a strong witness for $\mathcal{T}^{\infty}_{\text{Even}}$ w.r.t. σ, by showing that $wit(x = x)$ is not strongly finitely witnessed for $\mathcal{T}^{\infty}_{\text{Even}}$ w.r.t. σ. Consider a $\mathcal{T}^{\infty}_{\text{Even}}$-interpretation $\mathcal{A}$ with 2 elements, which interprets all the variables in $vars(wit(x = x))$. Clearly, $\mathcal{A} \models x = x$, and therefore, $\mathcal{A} \models \exists\, \overline{w}.\ wit(x = x)$, where $\overline{w}$ is $vars(wit(x = x)) \setminus \{x\}$. Hence, there exists a $\mathcal{T}^{\infty}_{\text{Even}}$-interpretation $\mathcal{A}'$, identical to $\mathcal{A}$, except possibly in its interpretation of variables in $vars(wit(x = x)) \setminus \{x\}$, that satisfies $wit(x = x)$. In particular, $\mathcal{A}'$ has two elements. Let $\delta_{\mathcal{A}'}$ be the arrangement over $vars(wit(x = x))$ satisfied by $\mathcal{A}'$. Then $\delta_{\mathcal{A}'}$ induces an equivalence relation with either 1 or 2 equivalence classes. Let v be a variable not in $vars(wit(x = x))$. Define an arrangement δ over $vars(wit(x = x)) \cup \{v\}$ as follows: If $\delta_{\mathcal{A}'}$ induces one equivalence class, $\delta := \delta_{\mathcal{A}'} \wedge \bigwedge_{u \in vars(wit(x=x))} v = u$. Otherwise, $\delta := \delta_{\mathcal{A}'} \wedge \bigwedge_{u \in vars(wit(x=x))} v \neq u$. In the first case, δ induces one equivalence class, and in the second, three. $wit(x = x) \wedge \delta$ does not have a finite witness for $\mathcal{T}^{\infty}_{\text{Even}}$ w.r.t. σ, as any interpretation $\mathcal{B}$ that finitely witnesses it has either 1 or 3 elements, and hence it is not in $\mathcal{T}^{\infty}_{\text{Even}}$.

A.10 Proof of Corollary 1

$\mathcal{T}_2$ is smooth w.r.t. $S^{si} \cup S^{nsi}$. In particular, it is smooth w.r.t. S^{nsi}. We show that it is also strongly stably infinite w.r.t. (S^{si}, S^{nsi}), and then the result follows from case 1 of Theorem 4. Let ϕ be a Σ-formula, $\mathcal{A}$ a $\mathcal{T}$-interpretation that satisfies

ϕ. Define $\kappa(\sigma)$ to be $\aleph_0$ for every $\sigma \in S^{si}$ such that $\sigma^{\mathcal{A}}$ is finite, $\kappa(\sigma) = |\sigma^{\mathcal{A}}|$ for every $\sigma \in S^{si}$ such that $\sigma^{\mathcal{A}}$ is infinite, and $\kappa(\sigma) = |\sigma^{\mathcal{A}}|$ for every $\sigma \in S^{nsi}$. Since $\mathcal{T}$ is smooth w.r.t. $S^{si} \cup S^{nsi}$, there exists a $\mathcal{T}$-interpretation $\mathcal{B}$ that satisfies ϕ with $|\sigma^{\mathcal{B}}| = \kappa(\sigma)$ for every $\sigma \in S^{si}$ and $|\sigma^{\mathcal{B}}| = \kappa(\sigma) = |\sigma^{\mathcal{A}}|$ for every $\sigma \in S^{nsi}$.

□

A.11 Proof of Lemma 8

Let Ax be the set of sentences that are satisfied by every $\mathcal{T}$-structure. Define the following sets, based on formulas that are defined in Figure 2:

$$\mathit{fin}_{\mathcal{A}} := \left\{ \psi^{\sigma}_{=|\sigma^{\mathcal{A}}|} \mid \sigma \in S^{\mathit{fin}}_{\mathcal{A}} \right\}$$
$$\mathit{inf}_{\mathcal{A}} := \left\{ \neg\psi^{\sigma}_{=n} \mid \sigma \in S^{\mathit{inf}}_{\mathcal{A}}, n \in \mathbb{N} \right\}$$
$$A := Ax \cup \mathit{fin}_{\mathcal{A}} \cup \mathit{inf}_{\mathcal{A}} \cup \{\phi\}$$

Clearly, $\mathcal{A} \models A$. By Theorem 1, there exists a Σ-interpretation $\mathcal{B}$ that satisfies A in which $\sigma^{\mathcal{B}}$ is countable whenever it is infinite, for every $\sigma \in \mathcal{S}_{\Sigma}$. This in particular holds for every $\sigma \in S^{\mathit{inf}}_{\mathcal{A}}$. Now let $\sigma \in S^{\mathit{fin}}_{\mathcal{A}}$, then since $\mathcal{B} \models \mathit{fin}_{\mathcal{A}}$, $|\sigma^{\mathcal{B}}| = |\sigma^{\mathcal{A}}|$. Finally, $\mathcal{B} \models \phi$ and it is a $\mathcal{T}$-interpretation.

□

A.12 Remaining Cases in The Proof of Lemma 9

Let $\psi_2 := \mathit{wit}(\Gamma_2)$. Since $\mathcal{T}_1$ is stably infinite w.r.t. S^{si}, there is a $\mathcal{T}_1$-interpretation $\mathcal{A}$ satisfying $\Gamma_1 \cup \delta_V$ in which $\sigma^{\mathcal{A}}$ is infinite for each $\sigma \in S^{si}$. By Theorem 1, we may assume that $\sigma^{\mathcal{A}}$ is countable for each $\sigma \in S^{si}$.

Case 2 : Suppose $\mathcal{T}_2$ is stably infinite w.r.t (S^{si}, S^{nsi}), smooth w.r.t. (S^{nsi}, S^{si}), and strongly finitely witnessable w.r.t. S^{nsi}. Then, there exists a $\mathcal{T}_2$-interpretation $\mathcal{B}$ that satisfies $\psi_2 \cup \delta_V$ such that $\sigma^{\mathcal{B}} = V^{\mathcal{B}}_{\sigma}$ for every $\sigma \in S^{nsi}$. Since $\mathcal{A}$ and $\mathcal{B}$ satisfy δ_V, we have that for every $\sigma \in S^{nsi}$, $|\sigma^{\mathcal{B}}| = |V^{\mathcal{B}}_{\sigma}| = |V^{\mathcal{A}}_{\sigma}| \leq |\sigma^{\mathcal{A}}|$. $\mathcal{T}_2$ is stably infinite w.r.t. (S^{si}, S^{nsi}), and so there exists a $\mathcal{T}_2$-interpretation $\mathcal{B}'$ that satisfies $\psi_2 \cup \delta_V$ such that $\sigma^{\mathcal{B}'}$ is infinite for every $\sigma \in S^{si}$ and $\left|\sigma^{\mathcal{B}'}\right| \leq |\sigma^{\mathcal{B}}| \leq |\sigma^{\mathcal{A}}|$ for every $\sigma \in S^{nsi}$. $\mathcal{T}_2$ is smooth w.r.t. (S^{nsi}, S^{si}) and so there is a $\mathcal{T}_2$-interpretation $\mathcal{B}''$ satisfying $\psi_2 \cup \delta_V$ such that $\left|\sigma^{\mathcal{B}''}\right| = |\sigma^{\mathcal{A}}|$ for every $\sigma \in S^{nsi}$ and $\left|\sigma^{\mathcal{B}''}\right|$ is infinite for every $\sigma \in S^{si}$. Using lemma 8, we may assume $\sigma^{\mathcal{B}''}$ is countable and hence $\left|\sigma^{\mathcal{B}''}\right| = |\sigma^{\mathcal{A}}|$ for every $\sigma \in S$.

Case 3 : Suppose $\mathcal{T}_2$ is stably infinite w.r.t. S^{si}, smooth w.r.t. (S^{nsi}, S^{si}), and strongly finitely witnessable w.r.t. (S^{nsi}, S^{si}). Since it is stably infinite w.r.t. S^{si}, there exists a $\mathcal{T}_2$-interpretation $\mathcal{B}$ that satisfies $\psi_2 \cup \delta_V$ such that $\sigma^{\mathcal{B}}$ is infinite for every $\sigma \in S^{si}$. $\mathcal{T}_2$ is strongly finitely-witnessable w.r.t. (S^{nsi}, S^{si}), and hence there exists a $\mathcal{T}_2$-interpretation $\mathcal{B}'$ that satisfies $\psi_2 \cup \delta_V$ such that $\sigma^{\mathcal{B}'} = V_\sigma^{\mathcal{B}'}$ for every $\sigma \in S^{nsi}$ and $\left|\sigma^{\mathcal{B}'}\right|$ is infinite for every $\sigma \in S^{si}$. Since $\mathcal{A}$ and $\mathcal{B}'$ satisfy δ_V, we have that for every $\sigma \in S^{nsi}$, $\left|\sigma^{\mathcal{B}'}\right| = \left|V_\sigma^{\mathcal{B}'}\right| = \left|V_\sigma^{\mathcal{A}}\right| \leq \left|\sigma^{\mathcal{A}}\right|$. $\mathcal{T}_2$ is smooth w.r.t. (S^{nsi}, S^{si}), and so there exists a $\mathcal{T}_2$-interpretation $\mathcal{B}''$ that satisfies $\psi_2 \cup \delta_V$ such that $\left|\sigma^{\mathcal{B}''}\right| = \left|\sigma^{\mathcal{A}}\right|$ for every $\sigma \in S^{nsi}$ and $\left|\sigma^{\mathcal{B}''}\right|$ is infinite for every $\sigma \in S^{si}$. By Lemma 8, we may assume that $\sigma^{\mathcal{B}''}$ is countable for every $\sigma \in S^{si}$, with the same cardinalities for sorts of S^{nsi}, and so we have $\left|\sigma^{\mathcal{B}''}\right| = \left|\sigma^{\mathcal{A}}\right|$ also for every $\sigma \in S$.