

A NEW CORRECTNESS PROOF OF THE NELSON-OPPEN COMBINATION PROCEDURE

CESARE TINELLI AND MEHDI HARANDI
University of Illinois at Urbana-Champaign
1304 W. Springfield Ave.
Urbana, IL 61801 – USA
{tinelli,harandi}@cs.uiuc.edu

Abstract.

The Nelson-Oppen combination procedure, which combines satisfiability procedures for a class of first-order theories by propagation of equalities between variables, is one of the most general combination methods in the field of theory combination. We describe a new non-deterministic version of the procedure that has been used to extend the Constraint Logic Programming Scheme to unions of constraint theories. The correctness proof of the procedure that we give in this paper not only constitutes a novel and easier proof of Nelson and Oppen's original results, but also shows that equality sharing between the satisfiability procedures of the component theories, the main idea of the method, can be confined to a restricted set of variables.

While working on the new correctness proof, we also found a new characterization of the consistency of the union of first-order theories. We discuss and give a proof of such characterization as well.

1. Introduction

Nelson and Oppen were among the first to provide a fairly general method to combine logical theories and relative satisfiability procedures ([17]). Since then, almost all the effort in the field of combination has been concentrated on equational theories and unification algorithms ([2, 3, 5, 6, 7, 9, 10, 14, 19, 20, 21, 25, 28, 27]). Others have worked on combinations of more general theories as well (see [15, 23, 24], for instance) but to date the Nelson-Oppen method appears to be still one of the most general in the field.

The need of extending the focus from unification to more general satisfiability problems is well felt in the combination literature after the emergence and consolidation of several constraint-based computational paradigms that operate on more general constraint domains than those described by equational theories.

We have shown in [26] how the generality of the Nelson-Oppen method allows us to easily incorporate a combination procedure based on that method into a constraint-based computation framework, namely, the CLP scheme of Jaffar and Lassez [11, 12, 13], with few modifications of the scheme itself. In that work, we first describe a non-deterministic version of the combination procedure originally devised by Nelson and Oppen. Then we integrate the procedure into a modified version of the CLP scheme to obtain a new scheme that operates with unions of constraint theories by using constraint solvers for the single component theories.

For space limitations, we are not able to describe here the integration in CLP nor prove that the main properties of CLP lift to the new scheme. For this we refer the reader to [26]. Instead, in this paper, we describe and discuss the combination procedure used in [26] and provide a novel proof of Nelson and Oppen’s original results. In addition, we provide a characterization of the consistency of the union of first-order theories.

1.1. NOTATION AND CONVENTIONS

In general, we will adhere to the notation and definitions given in [22]. The most notable notational conventions followed are given below with the understanding that other notations that may appear in the paper follow the common conventions of the field.

The letters v, x, y, z denote logical variables, φ, ψ first order formulas, and ϑ a value assignment, or valuation, to a set of variables.

Some of the above symbols may be subscripted or have an over-tilde which will represent a finite sequence. For instance, \tilde{x} stands for a sequence of the form (x_1, x_2, \dots, x_n) for some natural number n . When convenient, we will use the tilde notation to denote just sets of symbols—as opposed to sequences. The notation $\varphi(\tilde{x})$ is used to indicate that the free variables of φ are *exactly* the ones in \tilde{x} . In general, $var(\varphi)$ is the set of all the free variables of φ . The shorthand $\exists \tilde{x} \varphi$ stands for the existential closure of φ .

Where \mathcal{M} is a structure and φ a sentence, that is, a closed formula, the notation $\mathcal{M} \models \varphi$ means that \mathcal{M} satisfies φ or, equivalently, that φ is true in \mathcal{M} . If φ is in general a formula and ϑ a valuation on \mathcal{M} of φ ’s free variables, if any, the notation $\mathcal{M} \models \varphi^\vartheta$ means that ϑ satisfies φ in \mathcal{M} . Notice that, in analogy with substitutions, we write valuation applications

in postfix form. For convenience, where ϑ is a valuation of \tilde{x} , we will also indicate with ϑ the reduction of ϑ to \tilde{y} if $\tilde{y} \subset \tilde{x}$, or an arbitrary expansion of it to \tilde{y} if $\tilde{x} \subset \tilde{y}$.

We will generally identify first-order theories with the set of their theorems. We will also identify union of multi-sets of formulas with logical conjunction.

1.2. ORGANIZATION OF THE PAPER

In Sect. 2, we first recall and briefly discuss the Nelson-Oppen method; then we describe our non-deterministic version of their combination procedure. In Sect. 3, we formalize the combination problem more rigorously and discuss the conditions under which union of theories are consistent. Then we prove the correctness of the combination procedure introduced in Sect. 2. We conclude in Sect. 4 with some remarks on the possible extension of the procedure and give direction for further development.

2. The Nelson and Oppen Combination Method

In [17], Nelson and Oppen show how a satisfaction procedure for a theory built by combining several first-order theories can be derived as a combination of the satisfaction procedures for each of these theories. The main idea is to combine the satisfiability procedures by means of equality sharing. We will clarify this in a moment, but first let us set the problem more rigorously.

We say that a formula is in *simple Conjunctive Normal Form* if it is a conjunction of literals. Consider n first-order theories with equality, $\mathcal{T}_1, \dots, \mathcal{T}_n$, with respective signatures $\Sigma_1, \dots, \Sigma_n$. Assume that no two signatures have any non-logical symbols in common.¹ For each theory, let $sCNF(\mathcal{T}_i)$ be the class of simple Conjunctive Normal Form formulas built with the symbols of Σ_i . Where $\Sigma := \bigcup_{i=1..n} \Sigma_i$, let \mathcal{T} the Σ -theory defined as the (deductive closure of) the union of all the above theories and let $sCNF(\mathcal{T})$ be the class of sCNF Σ -formulas.

If for each \mathcal{T}_i we have a procedure that decides the satisfiability in \mathcal{T}_i of formulas of $sCNF(\mathcal{T}_i)$, we can easily derive a procedure that decides the satisfiability in \mathcal{T} of any formula of $sCNF(\mathcal{T})$. Because of the possible presence of “mixed” terms and predicates (that is, expressions built with symbols from different signatures), φ cannot be processed directly by any of the satisfiability procedures unless it is of the form $\varphi_1 \wedge \varphi_2 \wedge \dots \wedge \varphi_m$ —call it *separate form*—where each sub-formula φ_i is a formula of some $sCNF(\mathcal{T}_j)$. If that is the case, we know that φ is unsatisfiable in \mathcal{T} if and only if

¹Where we treat the equality symbol “=” as a logical constant.

any φ_i is. Now, if $\varphi_i \in \text{sCNF}(\mathcal{T}_j)$ say, by construction it is unsatisfiable in \mathcal{T} exactly when it is unsatisfiable in \mathcal{T}_j and so we can use directly the satisfiability procedure for $\text{sCNF}(\mathcal{T}_j)$ to verify the unsatisfiability of φ .

If φ is not already in separate form, we can apply a conversion procedure that, given φ , returns an equivalent separate form. The separation procedure (and its correctness proof) is straightforward but to describe it we need some definitions and notation first. We have adapted these from those in [2], among others, which appear to be well established in the field.

Consider the theories described above. For $i = 1, \dots, n$, a member of Σ_i is an *i-symbol*. A Σ -term t is an *i-term* if it is a variable or it has the form $f\bar{s}$ and f is an *i-symbol*. An *i-predicate* is defined analogously. A sub-term of an *i-term* t is an *alien* sub-term of t if it is a *j-term*, with $j \neq i$, and all of its super-terms in t are *i-terms*. An *i-term* is *pure* (or also, *i-pure*) if it only contains *i-symbols*. Analogously we can define alien predicate arguments. An *i-predicate* (including equations) is pure if all of its arguments are *i-terms*. Pure formulas are defined as obvious. Observe that, given our assumption on the various signatures, a variable is an *i-term* for any i and so an equation is always pure if one of its members is a variable and the other is a pure term.

The separation procedure consists of the following steps. Consider a formula $\varphi \in \text{sCNF}(\mathcal{T})$ and see it as a multi-set of literals.

1. Variable Abstraction

In φ , recursively replace each alien predicate argument or sub-term t with a newly generated variable z and add the equation $z = t$ to φ . Purify each equations of the form $t_1 = t_2$, where neither of t_1 and t_2 is a variable, by replacing it with the equations $z = t_1, z = t_2$ (where z is a new variable) and purifying t_1 and t_2 in turn.

2. Partition

Partition the new multi-set in $m \leq n$ blocks containing only *i-pure* literals.²

The resulting partition can be seen as a sCNF formula of the form $\varphi_1 \wedge \dots \wedge \varphi_m$ where each φ_i is a *j-pure* sCNF formula for some $j \in \{1, \dots, n\}$.

A formula may have many separate forms, but they are all equivalent modulo variable renaming and the standard properties of logical conjunction and equality, hence it is appropriate to speak of *the* separate form of a formula. We indicate the separate form of a formula $\varphi \in \text{sCNF}(\mathcal{T})$ with $\check{\varphi}$. For notational convenience, we will always think of $\check{\varphi}$ as a conjunction of the form $\varphi_1 \wedge \dots \wedge \varphi_n$, with n being the number of component theories, where for each $i \in \{1, \dots, n\}$, φ_i is an *i-pure* sCNF formula, even if

²Equations between variables are partitioned arbitrarily.

φ may not contain any i -symbol for some i . In that case, φ_i is defined as the identically true formula—which can be thought of as belonging to all $sCNF(\mathcal{T}_i)$'s.

It is immediate that any $\varphi \in sCNF(\mathcal{T})$ is logically equivalent to $\exists \bar{z} \tilde{\varphi}$ where \bar{z} is the set of new variables introduced by the separation procedure. This entails the following

Proposition 2.1 *A sCNF formula is satisfiable iff its separate form is.*

Clearly, the problem with deciding the satisfiability of a formula φ by analyzing its separate form is that, in general, each sub-formula φ_i could be singly satisfiable without their conjunction being satisfiable. Therefore, to be able to apply distinct satisfiability procedures to each φ_i and correctly decide the satisfiability of φ we need to establish some sort of communication between the various procedures. In the Nelson-Oppen method, such communication is achieved by propagating from one procedure to the others any implied equalities between the variables of $\tilde{\varphi}$.

Actually, the method is a little more complex because, in general, it is possible that at a certain step, not one, but a proper disjunction of variable equalities is implied and so reasoning by cases becomes necessary. Such complication does not arise though if the component theories are *convex*, that is, such that their formulas never imply proper *splits*³.

The issue of theory convexity is important to assessing the time and space complexity of a combination procedure based on the above method because of the case reasoning required with non-convex theories. Computational complexities issues related to the implementation of the Nelson-Oppen method have been extensively investigated in [18] and we refer the reader to that work. We will ignore those issues here by considering a non-deterministic combination procedure that we have adapted from those in [18] itself and [2] and which applies to convex as well as non-convex theories. Essentially, instead of propagating variable equalities back and forth and having to reason by cases when splits are implied, the non-deterministic procedure guesses in advance all the equalities that hold between the variables of the input formula.

We describe our version of the procedure in the next section where we consider only the simple case of two component theories, the general case being an easy generalization.

³A *split* of a formula φ is a disjunction of variable equalities which is implied by φ and is such that none of its proper sub-formulas are implied by φ . A proper split is a split with at least two disjuncts.

2.1. THE COMBINATION PROCEDURE

If P is any partition on a set S of words and R is the corresponding equivalence relation, we call the *arrangement* of S given by P the set

$$ar(S) := \{x = y \mid x, y \in S \text{ and } xRy\} \cup \{x \neq y \mid x, y \in S \text{ and not } xRy\}$$

containing, modulo reflexivity and symmetry of “=”, all the equations between any two equivalent words and all the disequations between any two non-equivalent words of S . For instance, if $S := \{x_0, x_1, x_2, x_3\}$ and $P := \{\{x_0, x_1, x_2\}, \{x_3\}\}$,

$$ar(S) := \{x_0 = x_1, x_0 = x_2, x_1 = x_2, x_0 \neq x_3, x_1 \neq x_3, x_2 \neq x_3\}.$$

In the following, we will make use of arrangements over sets of variables. Since, where \tilde{z} is a set of variables, $ar(\tilde{z})$ is a set of formulas, we will also treat it when convenient as the conjunction of all its equations and disequations.

Given two theories \mathcal{T}_1 and \mathcal{T}_2 , for $i = 1, 2$, we assume the availability of a procedure, Sat_i , that decides the satisfiability in \mathcal{T}_i for the formulas in $sCNF(\mathcal{T}_i)$.

The combination procedure given below decides the satisfiability in the union of \mathcal{T}_1 and \mathcal{T}_2 in two phases. In the first phase, which is non-deterministic, two i -pure formulas $\langle \psi_1, \psi_2 \rangle$ are generated from the input formula⁴. In the second phase, each ψ_i is tested for satisfiability in \mathcal{T}_i . The combination procedure succeeds when both ψ_i 's are satisfiable and fails otherwise.

More specifically, let L be the (multi)set of literals of the input formula. The procedure is composed of the following steps:

– **Decomposition Phase**

1. For $i = 1, 2$, let L_i be the i -pure part of L 's separate form.
2. Where \tilde{x} is the set of all variables shared by some literal in L_1 and some in L_2 , choose an arrangement $ar(\tilde{x})$.
3. Pass the pair $\langle L_1 \cup ar(\tilde{x}), L_2 \cup ar(\tilde{x}) \rangle$ to the next phase.

– **Check Phase**

1. Run Sat_1 on $L_1 \cup ar(\tilde{x})$.
2. Run Sat_2 on $L_2 \cup ar(\tilde{x})$.

⁴This phase corresponds to the decomposition algorithm of Baader and Schulz (see, [3] or [4]).

3. Succeed if both Sat_1 and Sat_2 succeed. Fail if either of Sat_1 or Sat_2 fails.

Observe that, the choice of an arrangement corresponds to the *variable identification* step of most combination methods ([2, 4, 7, 14, 19]) and to the equality propagation mechanism of the original Nelson-Oppen procedure. In our case though, variable identification is only performed on the *shared* variables of the two *i*-pure halves of the input instead of on *all* of them. That this is enough is also recognized in [3], for a similar combination method over structures, on the basis of model-theoretic justifications analogous to the ones we give later.

This combination procedure is provably sound and complete for a restricted class of first-order theories. In addition, it allows incremental implementations when incremental satisfiability procedures for the component theories are available. A simple incremental implementation is the following.

Let L be any permutation of the literals in the separate form of the input formula. We feed the two satisfiability procedures by picking one literal at a time from L and passing it to Sat_1 or Sat_2 according to whether it is a 1-literal or a 2-literal. In addition, if the literal shares some variables \bar{v} with the literals already passed to the other satisfiability procedure, we choose some $ar(\bar{v})$ and pass it to both procedures.⁵

It is an incremental version of the combination procedure that it is actually used in [26] to extend the CLP scheme.

Finally, notice that, although we have restricted our attention to sCNF formulas, the above procedure can also be used to decide the satisfiability of quantifier-free formulas in general, since a quantifier-free formula is satisfiable iff some disjunct—which is in turn a sCNF formula—of its disjunctive normal form is satisfiable.

3. Correctness of the Combination Procedure

All the theories we consider in the following are first-order theories incorporating the theory of equality. For convenience, we will follow the common convention that considers the theory of equality as an integral part of the logical machinery of First Order Predicate Logic and so, from now on, when we say “theory” we will mean an theory in FOPL with equality.

Earlier, we defined the combination, or union, of two or more theories as the deductive closure of the union of the component theories. In the fol-

⁵It should be obvious that since implementations are necessarily deterministic, some sort of backtracking mechanism is required in this case to recover from *wrong* choices of arrangements.

lowing, where \mathcal{T}_1 and \mathcal{T}_2 are two theories, we will denote their combination simply as $\mathcal{T}_1 \cup \mathcal{T}_2$.

3.1. CONSISTENT UNIONS OF THEORIES

Before proving the correctness of the combination procedure, we make a brief digression on the combination of theories that will later justify some of the restrictions to the application of the procedure.

Combining theories is, in general, a non-trivial operation because most of the model and proof-theoretic properties of theories are not modular, including the most important one: consistency. Some papers in the combination literature do provide a proof of consistency for their combined theories, or structures, but their proofs are often ad hoc. Others (including Nelson and Oppen's) seem to ignore the issue by either assuming consistency or giving it for granted in their case.

Craig and Robinson have identified a while ago a *local* criterion for the consistency of combined theories which justifies the choice of signature-disjoint component theories if some conditions on the cardinality of the theories' models are met.

We formalize this in the following theorem by giving a necessary and sufficient condition for combining signature-disjoint theories meaningfully. The theorem and its following extension to the general case of non-signature-disjoint theories essentially subsume all previous similar results in the literature.

In our proofs, we will use the class of formulas defined below.

Definition 3.1 *A first order formula is called an equational formula iff all its atomic sub-formulas are equalities between variables.*⁶

Proposition 3.1 *Let \mathcal{T}_1 and \mathcal{T}_2 be two theories. Assume that they are consistent and their respective signatures, Σ_1 and Σ_2 , are disjoint. Then, their union is consistent iff there is a cardinal κ such that both \mathcal{T}_1 and \mathcal{T}_2 have a model of cardinality κ .*

Proof: (\Rightarrow) Let $\mathcal{T} := \mathcal{T}_1 \cup \mathcal{T}_2$ and consider any $\mathcal{M} \in \text{Mod}(\mathcal{T})$ —where $\text{Mod}(\mathcal{T}) \neq \emptyset$ as \mathcal{T} is consistent by assumption. By construction of \mathcal{T} , the reduct of \mathcal{M} to Σ_i is a model of \mathcal{T}_i , for $i = 1, 2$. Obviously, both reducts have the same cardinality.

(\Leftarrow) Let \mathcal{M}_1 and \mathcal{M}_2 be models of \mathcal{T}_1 and \mathcal{T}_2 , respectively, and assume they have the same cardinality. By the Craig-Robinson Theorem (see [22]), \mathcal{T} is inconsistent iff there is a sentence φ , whose non-logical symbols are in $\Sigma_1 \cap \Sigma_2$, such that $\mathcal{T}_1 \models \varphi$ and $\mathcal{T}_2 \models \neg\varphi$. If such φ exists, we have that

$$\mathcal{M}_1 \models \varphi \quad \text{and} \quad \mathcal{M}_2 \models \neg\varphi. \quad (1)$$

⁶Nelson-Oppen call such formulas *simple* formulas.

Now, as \mathcal{T}_1 and \mathcal{T}_2 are signature-disjoint, φ can only be an equational formula. It is a well-known result of Model Theory that the reducts of any two structures to the language of equational formulas are isomorphic whenever they have the same cardinality⁷. This means that either both \mathcal{M}_1 and \mathcal{M}_2 model φ or neither of them does, which contradicts (1). \square

The above proof immediately suggests an extension of the previous result to the combination of any two consistent theories.⁸

Proposition 3.2 *Let \mathcal{T}_1 and \mathcal{T}_2 be two consistent theories with respective signatures Σ_1 and Σ_2 . Their union is consistent iff there is a model \mathcal{M}_1 of \mathcal{T}_1 and a model \mathcal{M}_2 of \mathcal{T}_2 such that their reducts to $\Sigma_1 \cap \Sigma_2$ are isomorphic.*

Proof: Analogous to that of Prop. 3.1. \square

A well-known general result on the union of theories is Robinson's Consistency Theorem (see [8], for instance). This theorem, however, provides only a sufficient condition for the consistency of the union theory and a somewhat stronger one: $\mathcal{T}_1 \cup \mathcal{T}_2$ is consistent if $\mathcal{T}_1 \cap \mathcal{T}_2$ is a complete theory for the language of $(\Sigma_1 \cap \Sigma_2)$ -sentences.

In our terms, this condition is expressed as follows: for every model \mathcal{M} of one theory there is a model \mathcal{M}' of the other such that the reducts of \mathcal{M} and \mathcal{M}' to Σ are elementarily equivalent⁹. Now, elementary equivalence between structures is a strictly more general relation than isomorphism. In some sense, however, our consistency result can be seen as more general than Robinson's for requiring the existence of just one pair of structures, related through the isomorphism of their $(\Sigma_1 \cap \Sigma_2)$ -reducts, as opposed to an infinite number of them, related through the elementary equivalence of their $(\Sigma_1 \cap \Sigma_2)$ -reducts.

In general, the conditions of Prop. 3.2 on the models of the component theories are not so easy to verify. With signature-disjoint theories, however, some cases are immediate. For instance, a consequence of Prop. 3.1, to which we will appeal later, is that the combination of signature-disjoint theories admitting infinite models is always safe.

Corollary 3.3 *Let \mathcal{T}_1 and \mathcal{T}_2 be as in Prop. 3.1 but such that they both admit an infinite model. Then, their union is consistent.*

⁷In other words, the theory of equational formulas is κ -categorical for any cardinal κ (see, [22]).

⁸This extension was also independently obtained by Franz Baader ([1]) as a generalization of our Prop. 3.1.

⁹Recall that two Σ -structures are *elementarily equivalent* iff they satisfy exactly the same Σ -sentences and that a theory is complete exactly when every two models of it are elementarily equivalent.

Proof: Immediate consequence of the fact that, by the Löwenheim-Skolem-Tarski Theorem (see [8]), if each \mathcal{T}_i admits an infinite model then it admits infinite models of any cardinality (\geq the cardinality of its signature). \square

In the field of equational theories¹⁰ consistency in the general sense is not an issue since all equational theories admit trivial models. A stronger version of consistency is then used, let us call it *E-consistency* here: an equational theory is *E-consistent* iff it admits models of cardinality greater than 1. Franz Baader has shown in [1] that a well-known result in Unification Theory (see [20, 21]) is easily derivable as a consequence of Cor. 3.3.

Corollary 3.4 *Let \mathcal{E}_1 and \mathcal{E}_2 be two signature-disjoint equational theories. If \mathcal{E}_1 and \mathcal{E}_2 are E-consistent, then their union is E-consistent as well.*

Proof: For $i = 1, 2$, let E_i be a presentation of \mathcal{E}_i and consider the set $F_i := E_i \cup \{\bar{\exists}(x \neq y)\}$. Observe that, since \mathcal{E}_i is *E-consistent* by assumption, E_i admits infinite models¹¹ and clearly so does F_i . By Cor. 3.3 then $F_1 \cup F_2$ admits at least one model; moreover, since it entails $\bar{\exists}(x \neq y)$, it only admits non-trivial models. The claim follows by monotonicity observing that $\mathcal{E}_1 \cup \mathcal{E}_2$ is included in $F_1 \cup F_2$. \square

3.2. SOUNDNESS AND COMPLETENESS RESULTS

We are now ready to prove the correctness of the combination procedure. We start with a rigorous definition of the notions that we have been using informally in the previous section.

Definition 3.2 *Consider a Σ -theory \mathcal{T} . We say that a formula φ is satisfiable in \mathcal{T} iff it is satisfiable in some model of \mathcal{T} , that is, iff there exists a model $\mathcal{M} \in \text{Mod}(\mathcal{T})$ such that $\mathcal{M} \models \bar{\exists} \varphi$.*

This definition is the dual of the standard definition of unsatisfiability for formulas (as opposed to sentences) when we follow the convention of considering free variables as implicitly universally quantified.

Recall that we apply the combination procedure to signature-disjoint theories. Prop. 3.1 provides a condition on the component theories, namely that there is a model for one theory and a model for the other which have the same cardinality, that guarantees the consistency of their union and, as a consequence, that the satisfiability problem in it is not trivial. Unfortunately, that conditions alone is not sufficient for the correctness of the

¹⁰Recall that an equational theory is a first order theory admitting an axiomatization, or *presentation*, all of whose axioms are universally quantified equations.

¹¹For instance, the free algebras in countably many generators.

combination procedure. Problems might arise with theories that admit only finite models. We explain this point with the help of an example.

Example. Consider a theory \mathcal{T}_1 admitting models of cardinality at most 2 and a signature-disjoint theory \mathcal{T}_2 admitting models of any cardinality. Assume that f is a functor of \mathcal{T}_1 , g a functor of \mathcal{T}_2 and neither of them is defined as a constant function in its respective theory¹². The union \mathcal{T} of \mathcal{T}_1 and \mathcal{T}_2 is consistent by the Prop. 3.1, so consider the input formula

$$\varphi := fx \neq fy \wedge gx \neq gz \wedge gy \neq gz$$

The procedure splits φ into

$$L_1 := \{fx \neq fy\} \quad \text{and} \quad L_2 := \{gx \neq gz, gy \neq gz\}.$$

Observe that only possible arrangements of the variables shared between L_1 and L_2 are $\{x = y\}$ and $\{x \neq y\}$.

Now, $L_1 \cup \{x = y\}$ is clearly unsatisfiable and so the procedure fails on that arrangement. With the other arrangement however, both L_i 's are satisfiable in their respective theories and so the procedure concludes that φ is satisfiable in \mathcal{T} . Unfortunately,

$$\mathcal{T} \models \varphi \rightarrow (x \neq y \wedge x \neq z \wedge y \neq z)$$

which means that φ is unsatisfiable in \mathcal{T} because, as \mathcal{T}_1 , \mathcal{T} only has models of cardinality less than 3.

It is an easy consequence of the Löwenheim-Skolem-Tarski Theorem that cases like the above do not appear if we only consider formulas that are satisfiable in infinite models. Hence, we will restrict our attention to the class of theories in which formulas are satisfiable if and only if they are satisfiable in an infinite model and prove that for this class the combination procedure is sound and complete.

Definition 3.3 (Stable-infiniteness [18]) *A consistent, quantifier free theory \mathcal{T} with signature Σ is called stably-infinite iff any quantifier-free Σ -formula is satisfiable in \mathcal{T} iff it is satisfiable in an infinite model of \mathcal{T} .*

It follows from the above definition that every stably infinite theory admits at least one infinite model.¹³

¹²That is, $\mathcal{T}_1 \not\models \forall x, y \, fx = fy$ and similarly for g .

¹³We would like to point out that just having infinite models, although necessary, is not sufficient for stable-infiniteness. Consider the theory $\{p(z) \rightarrow x = y\}$; the theory admits infinite models but the quantifier-free formula $p(z)$ is only satisfiable in the trivial models of the theory.

In the following, we will indicate with $\Delta(\tilde{x})$ the quantifier-free formula obtained as the conjunction of all possible disequalities (modulo symmetry) between distinct variables of \tilde{x} . For instance, if \tilde{x} is $\{x_1, x_2, x_3\}$, then $\Delta(\tilde{x})$ is $x_1 \neq x_2 \wedge x_1 \neq x_3 \wedge x_2 \neq x_3$. Recalling the definition of arrangement, it is immediate that $\Delta(\tilde{x})$ is the arrangement generated by the discrete partition of \tilde{x} .

We start with some lemmas involving equational formulas.

Lemma 3.5 *Consider a theory \mathcal{E} over the language of equational formulas. If a closed equational formula ψ is valid in an infinite model of \mathcal{E} , then it is valid in every infinite model of \mathcal{E} .*

Proof: Again, immediate consequence of the Upward Löwenheim-Skolem-Tarski theorem (see [8]), since, as we saw earlier, all the models of \mathcal{E} with the same cardinality are isomorphic. \square

Lemma 3.6 *Consider a theory \mathcal{E} as above. Assume that an equational formula ψ with free variables \tilde{x} is satisfied in a model \mathcal{M} of \mathcal{E} by an assignment ϑ of different individuals to each variable in \tilde{x} . Then ψ is satisfied in \mathcal{M} by any assignment of different individuals to each variable in \tilde{x} .*

Equivalently, for all $\mathcal{M} \in \text{Mod}(\mathcal{E})$,

$$\text{if } \mathcal{M} \models \exists \tilde{x} (\Delta(\tilde{x}) \wedge \psi) \text{ then } \mathcal{M} \models \forall \tilde{x} (\Delta(\tilde{x}) \rightarrow \psi).$$

Proof: It can be shown¹⁴ that ψ is equivalent to a formula of the form

$$\bigvee_{i \in I} (\psi_i \wedge \varphi_i)$$

where each φ_i is a sentence and ψ_i is an arrangement of \tilde{x} . If $\mathcal{M} \models (\Delta(\tilde{x}) \wedge \psi)\vartheta$ for some model $\mathcal{M} \in \text{Mod}(\mathcal{E})$ and assignment ϑ , there exists an $i \in I$ such that $\mathcal{M} \models (\psi_i \wedge \varphi_i)\vartheta$. Since ϑ also satisfies $\Delta(\tilde{x})$, ψ_i cannot contain any equality and so must be equal to $\Delta(\tilde{x})$. The claim follows immediately from the fact that φ_i is closed and so satisfied in \mathcal{M} by any assignment. \square

The proof of the theorem below is based on the following corollary (see [22]) of the above mentioned consistency result by Craig and Robinson.

Lemma 3.7 (Craig Interpolation Lemma) *If $\mathcal{T}_1 \cup \mathcal{T}_2 \models \varphi_1 \rightarrow \varphi_2$, where for $i = 1, 2$, φ_i is a formula in the language of \mathcal{T}_i , there exists a formula ψ , whose free variables are among the free variables shared by φ_1 and φ_2 , such that $\mathcal{T}_1 \models \varphi_1 \rightarrow \psi$ and $\mathcal{T}_2 \models \psi \rightarrow \varphi_2$.*

¹⁴See Lemmas 1.5.6 and 1.5.7 of [8], for instance.

We are now ready for the main result of this section.

Proposition 3.8 *Let \mathcal{T}_1 and \mathcal{T}_2 be two stably-infinite, signature-disjoint theories and let $\varphi_1 \in \text{sCNF}(\mathcal{T}_1)$ and $\varphi_2 \in \text{sCNF}(\mathcal{T}_2)$. Let \tilde{v} be the set of variables shared by φ_1 and φ_2 . If $\varphi_i \wedge \Delta(\tilde{v})$ is satisfiable in \mathcal{T}_i for $i = 1, 2$, then $\varphi_1 \wedge \varphi_2$ is satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$.*

Proof: Ad absurdum, assume that $\varphi_1 \wedge \varphi_2$ is unsatisfiable in $\mathcal{T} := \mathcal{T}_1 \cup \mathcal{T}_2$, then $\mathcal{T} \models \varphi_1 \rightarrow \neg\varphi_2$. By the lemma above, there exists a formula ψ , whose free variables \tilde{x} are in \tilde{v} , such that $\mathcal{T}_1 \models \varphi_1 \rightarrow \psi$ and $\mathcal{T}_2 \models \varphi_2 \rightarrow \neg\psi$. Again, since \mathcal{T}_1 and \mathcal{T}_2 are signature disjoint, ψ must be an equational formula. Now, \mathcal{T}_1 is stably infinite and $\varphi_1 \wedge \Delta(\tilde{x})$ is satisfiable in \mathcal{T}_1 , therefore it is satisfiable in an infinite model \mathcal{M} of \mathcal{T}_1 and so is $\psi \wedge \Delta(\tilde{x})$.

Observe that since \mathcal{T}_1 contains \mathcal{E}_\emptyset , the theory of equality, \mathcal{M} is a model of \mathcal{E}_\emptyset as well. By Lemma 3.6, it follows that $\forall \tilde{x} (\Delta(\tilde{x}) \rightarrow \psi)$ is valid in \mathcal{M} . By Lemma 3.5, it follows that $\forall \tilde{x} (\Delta(\tilde{x}) \rightarrow \psi)$ is valid in every infinite model of \mathcal{E}_\emptyset .

In the same way, we can show that $\forall \tilde{x} (\Delta(\tilde{x}) \rightarrow \neg\psi)$ is valid in every infinite model of \mathcal{E}_\emptyset , which leads to a contradiction. \square

In the above proposition, we assumed that the two φ_i 's are satisfied by an assignment of a different individual to each of their shared variables. The following corollary shows that there is no loss of generality in considering only such assignments since we can always *eliminate*, by identification, those shared variables that would be assigned to the same individual.

Corollary 3.9 *Consider φ_1 and φ_2 as above and let $ar(\tilde{v})$ be an arrangement of their shared variables. If $\varphi_i \wedge ar(\tilde{v})$ is satisfiable in \mathcal{T}_i for $i = 1, 2$, then $\varphi_1 \wedge \varphi_2$ is satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$.*

Proof: Assume that φ_i has the form $\varphi_i(\tilde{v}, \tilde{z}_i)$ where \tilde{z}_i are the non shared variables of φ_i . Given the equivalence relation that generates $ar(\tilde{v})$, we choose an *identification* of the elements of \tilde{v} , that is, a substitution σ from \tilde{v} to \tilde{v} that substitutes each variable in a same equivalence class with a given representative for that class. Now, let $\tilde{u} := \tilde{v}\sigma$

Clearly, $(\varphi_i(\tilde{v}, \tilde{z}_i) \wedge ar(\tilde{v}))\sigma$ is still satisfiable in \mathcal{T}_i . Observe however that $\varphi_i\sigma$ has now the form $\varphi_i(\tilde{u}, \tilde{z}_i)$ and $ar(\tilde{v})\sigma$ is actually (equivalent to) $\Delta(\tilde{u})$. By Prop. 3.8 then, $\varphi_1(\tilde{u}, \tilde{z}_1) \wedge \varphi_2(\tilde{u}, \tilde{z}_2)$ is satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$ and hence $\varphi_1(\tilde{v}, \tilde{z}_1) \wedge \varphi_2(\tilde{v}, \tilde{z}_2) \wedge ar(\tilde{v})$ is satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$ as well. \square

We can now prove the correctness of the combination procedure for the union \mathcal{T} of two signature-disjoint, stably-infinite theories, \mathcal{T}_1 and \mathcal{T}_2 .

Proposition 3.10 (Soundness) *If one of the pairs $\langle \psi_1, \psi_2 \rangle$ output by the decomposition phase of the combination procedure is such that ψ_i is satisfiable in \mathcal{T}_i for $i = 1, 2$, then the input formula is satisfiable in \mathcal{T} .*

Proof: We have already seen that Step 1 of the decomposition phase preserves satisfiability. The claim then is an immediate consequence of Cor. 3.9. \square

Proposition 3.11 (Completeness) *If a formula $\varphi \in \text{sCNF}(\mathcal{T})$ is satisfiable in \mathcal{T} , then there exists an output pair $\langle \psi_1, \psi_2 \rangle$ of the decomposition phase such that ψ_i is satisfiable in \mathcal{T}_i for $i = 1, 2$.*

Proof: Assume φ is satisfiable in a model \mathcal{M} of \mathcal{T} , then $\check{\varphi} := \varphi_1 \wedge \varphi_2$ is satisfiable in \mathcal{M} for some valuation ϑ . This valuation induces an arrangement $ar(\tilde{v})$ on the set \tilde{v} of shared variables between φ_1 and φ_2 , where for any $x, y \in \tilde{v}$, $(x = y) \in ar(\tilde{v})$ if $x\vartheta = y\vartheta$ and $(x \neq y) \in ar(\tilde{v})$ otherwise.

Clearly, we have that $\mathcal{M} \models (ar(\tilde{v}) \wedge \varphi_1 \wedge \varphi_2)\vartheta$, which implies that $\mathcal{M} \models (ar(\tilde{v}) \wedge \varphi_1)\vartheta$ and $\mathcal{M} \models (ar(\tilde{v}) \wedge \varphi_2)\vartheta$. Since, for $i = 1, 2$, the reduct of \mathcal{M} to the signature of \mathcal{T}_i is a model of \mathcal{T}_i , we have that $ar(\tilde{v}) \wedge \varphi_i$ is satisfiable in \mathcal{T}_i .

The claim follows from the fact that $\langle \varphi_1 \wedge ar(\tilde{v}), \varphi_2 \wedge ar(\tilde{v}) \rangle$ is indeed a possible output pair of the decomposition procedure. \square

Notice that for the results above we need neither to postulate that the satisfiability problem is decidable in the component theories nor that the theories are axiomatizable. When satisfiability is in fact decidable, we obtain the stronger correctness result below.

Corollary 3.12 *Assume that for $i = 1, 2$, Sat_i is a decision procedure for the satisfiability in \mathcal{T}_i of formulas of $\text{sCNF}(\mathcal{T}_i)$. Then, a formula $\varphi \in \text{sCNF}(\mathcal{T})$ is satisfiable in \mathcal{T} if and only if the combination procedure succeeds on φ .*

Proof: Immediate consequence of Propositions 3.10 and 3.11 and the easily proved fact that the procedure halts on every input whenever both Sat_1 and Sat_2 do.

4. Conclusions and Further Developments

In this paper we have described a non-deterministic version of the Nelson-Oppen combination procedure and given a novel proof of Nelson and Oppen's original results.

We believe that our proof is relevant for at least two reasons. First, it avoids the problematic concept of *residue* of a formula, introduced in Nelson and Oppen's proofs, which is only defined for infinite interpretations of the formula itself.¹⁵ Second, it shows that equality propagation between the

¹⁵It looks like the authors had not realized this initially. As a matter of fact, the proofs given in [17] were incorrect. In later papers on the methods, [18] and [16], the problem

satisfiability procedures of the component theories, the main idea of the method, can be confined to a restricted set of variables.

Another contribution of the paper was a characterization of the consistency of the union of first-order theories.

In an attempt to extend the combination procedure to more general cases, we are confronted with two issues, among others, that we believe are very significant and should deserve further investigation.

The first issue is the stable-infiniteness requirement on the component theories. There certainly are interesting constraint theories that are not stably-infinite¹⁶. We have seen in Sect. 3.2 what kind of problems can arise if one of the component theories is not stably infinite, in particular if it only admits finite models. Stable-infiniteness does not seem to be a necessary condition for the correctness of the combination procedure although it is the most general sufficient condition identified so far. We conjecture that there might exist weaker requirements on the component theories which are sufficient for the procedure's correctness in the case of signature-disjoint theories and quantifier-free input formulas.

The second issue is the disjointness requirement on the signatures of the component theories. The procedure can be easily extended to the case of component theories with signatures sharing only constant symbols¹⁷; unfortunately, the correctness proof reported here does not lift to this extension because—like Nelson and Oppen's— it is based on the exclusive model-theoretic properties of equational formulas. We have found, however, a simple constructive proof of Prop. 3.2 that leads naturally to a constructive proof of an analogous to Prop. 3.8 for the case of component theories whose signatures share a finite number of constants. The case of shared function symbols is, understandably, much harder because of the infinite number of terms that are then shared by the languages of the component theories. We are currently trying to identify further model-theoretic restrictions on the component theories, beside stable infiniteness, that might lead to some controlled form of term sharing and therefore suggest further extensions of the procedure.

was side-stepped by restricting attention to infinite models only and implicitly claiming that such restriction did not invalidate the generality of the results. This is indeed true but not totally immediate.

¹⁶With theories of finite domains being the most prominent examples, of course.

¹⁷In essence, this can be done by including the shared constants in the computation of arrangements.

5. Acknowledgements

We would like to thank Alan Frisch for initially pointing out Nelson and Oppen's method, Franz Baader for a long and constructive series of discussions on combination methods and their deep implications, Joshua Caplan for a number of illuminating discussions on some model-theoretic issues related to this work, and the anonymous referees for their valuable comments and suggestions.

This work is partially supported by grant DACA88-94-0014 from the US Army Construction Engineering Laboratories.

References

1. Franz Baader, July 1995. Personal communication.
2. Franz Baader and Klaus U. Schulz. Unification in the union of disjoint equational theories: Combining decision procedures. In *Proceedings of the 11th International Conference on Automated Deduction*, volume 607 of *Lecture Notes in Artificial Intelligence*, pages 50–65. Springer-Verlag, 1992.
3. Franz Baader and Klaus U. Schulz. Combination of constraint solving techniques: An algebraic point of view. In *Proceedings of the 6th International Conference on Rewriting Techniques and Applications*, volume 914 of *Lecture Notes in Computer Science*, pages 50–65. Springer-Verlag, 1995.
4. Franz Baader and Klaus U. Schulz. Combination techniques and decision problems for disunification. *Theoretical Computer Science*, 142:229–255, 1995.
5. Franz Baader and Klaus U. Schulz. On the combination of symbolic constraints, solution domains, and constraint solvers. In *Proceedings of the First International Conference on Principles and Practice of Constraint Programming, Cassis (France)*, September 1995.
6. Alexandre Boudet. Unification in a combination of equational theories: An efficient algorithm. In M. E. Stickel, editor, *Proceedings of the 10th International Conference on Automated Deduction*, volume 449 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 1990.
7. Alexandre Boudet. Combining unification algorithms. *Journal of Symbolic Computation*, 16(6):597–626, December 1993.
8. C. C. Chang and H. Jerome Keisler. *Model Theory*, volume 73 of *Studies in logic and the foundations of mathematics*. North-Holland, Amsterdam-New York-Oxford-Tokyo, 1990.
9. E. Domenjoud, F. Klay, and C. Ringeissen. Combination techniques for non-disjoint equational theories. In A. Bundy, editor, *Proceedings 12th International Conference on Automated Deduction, Nancy (France)*, volume 814 of *Lecture Notes in Artificial Intelligence*, pages 267–281. Springer-Verlag, 1994.
10. A. Herold. Combination of unification algorithms. In J. Siekmann, editor, *Proceedings 8th International Conference on Automated Deduction, Oxford (UK)*, volume 230 of *Lecture Notes in Artificial Intelligence*, pages 450–469. Springer-Verlag, 1986.
11. Joxan Jaffar and Jean-Louis Lassez. Constraint Logic Programming. Technical Report 86/74, Monash University, Victoria, Australia, June 1986.
12. Joxan Jaffar and Jean-Louis Lassez. Constraint Logic Programming. In *POPL'87: Proceedings 14th ACM Symposium on Principles of Programming Languages*, pages 111–119, Munich, January 1987. ACM.
13. Joxan Jaffar and Michael Maher. Constraint Logic Programming: A Survey. *Journal of Logic Programming*, 19/20:503–581, 1994.
14. Hélène Kirchner and Christophe Ringeissen. A constraint solver in finite algebras

- and its combination with unification algorithms. In K. Apt, editor, *Proc. Joint International Conference and Symposium on Logic Programming*, pages 225–239. MIT Press, 1992.
15. Hélène Kirchner and Christophe Ringeissen. Constraint solving by narrowing in combined algebraic domains. In P. Van Hentenryck, editor, *Proc. 11th International Conference on Logic Programming*, pages 617–631. The MIT press, 1994.
 16. Greg Nelson. Combining satisfiability procedures by equality-sharing. *Contemporary Mathematics*, 29:201–211, 1984.
 17. Greg Nelson and Derek C. Oppen. Simplification by cooperating decision procedures. *ACM Trans. on Programming Languages and Systems*, 1(2):245–257, October 1979.
 18. Derek C. Oppen. Complexity, convexity and combinations of theories. *Theoretical Computer Science*, 12, 1980.
 19. Christophe Ringeissen. Unification in a combination of equational theories with shared constants and its application to primal algebras. In *Proceedings of the 1st International Conference on Logic Programming and Automated Reasoning*, volume 624 of *Lecture Notes in Artificial Intelligence*, pages 261–272. Springer-Verlag, 1992.
 20. Manfred Schmidt-Schauß. Unification in a combination of disjoint equational theories. In *Proceedings of the 9th International Conference on Automated Deduction*, volume 310 of *Lecture Notes in Computer Science*, pages 378–396. Springer-Verlag, 1988.
 21. Manfred Schmidt-Schauß. Combination of unification algorithms. *Journal of Symbolic Computation*, 8(1–2):51–100, 1989.
 22. Joseph R. Shoenfield. *Mathematical Logic*. Addison-Wesley, Reading, MA, 1967.
 23. Robert E. Shostak. A practical decision procedure for arithmetic with function symbols. *Journal of the ACM*, 26(2):351–360, April 1979.
 24. Robert E. Shostak. Deciding combinations of theories. *Journal of the ACM*, 31:1–12, 1984.
 25. E. Tidén. Unification in combinations of collapse-free theories with disjoint sets of function symbols. In J. Siekmann, editor, *Proceedings 8th International Conference on Automated Deduction, Oxford (UK)*, volume 230 of *Lecture Notes in Artificial Intelligence*, pages 431–449. Springer-Verlag, 1986.
 26. Cesare Tinelli. Extending the CLP scheme to unions of constraint theories. Master’s thesis, Department of Computer Science, University of Illinois, Urbana-Champaign, Illinois, October 1995.
 27. K. Yelik. Combining unification algorithms for confined equational theories. *Journal of Symbolic Computation*, 3(1), April 1987.
 28. K. Yelik. Unification in combinations of collapse-free regular theories. *Journal of Symbolic Computation*, 3(1–2):153–182, April 1987.