Non-Disjoint Unions of Theories and Combinations of Satisfiability Procedures: First Results

Cesare Tinelli, Christophe Ringeissen

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Non-Disjoint Unions of Theories and Combinations of Satisfiability Procedures: First Results

Cesare Tinelli, Christophe Ringeissen

N° 3402

April 1998

THÈME 2

Rapport de recherche

ISSN 0249-6399

Non-Disjoint Unions of Theories and Combinations of Satisfiability Procedures: First Results[*]

Cesare Tinelli[†], Christophe Ringeissen[‡]

Thème 2 — Génie logiciel
et calcul symbolique
Projet PROTHEO

Rapport de recherche n ° 3402 — April 1998 — 63 pages

Abstract: In this paper we outline a theoretical framework for the combination of decision procedures for the satisfiability of constraints with respect to a constraint theory. We describe a general combination method which, given a procedure that decides constraint satisfiability with respect to a constraint theory $\mathcal{T}_1$ and one that decides constraint satisfiability with respect to a constraint theory $\mathcal{T}_2$, is able to produce a procedure that (semi-)decides constraint satisfiability with respect to the union of $\mathcal{T}_1$ and $\mathcal{T}_2$. We also provide some model-theoretic conditions on the constraint language and the component constraint theories for the method to be sound and complete, with special emphasis on the case in which the signatures of $\mathcal{T}_1$ and $\mathcal{T}_2$ are non-disjoint.

Key-words: Combination of Satisfiability Procedures, Decision Problems, Constraint-based Reasoning, Automated Deduction.

(Résumé : tsvp)

[*] Also available as UIUC Technical Report No. UIUCDCS-R-98-2044 UILU-ENG-98-1707.

[†] Department of Computer Science, University of Illinois at Urbana-Champaign, 1304 W. Springfield Ave. Urbana, IL 61801 – USA. Email: `tinelli@cs.uiuc.edu`

[‡] Email: `Christophe.Ringeissen@loria.fr`

Unité de recherche INRIA Lorraine
Technopôle de Nancy-Brabois, Campus scientifique,
615 rue de Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY (France)
Téléphone : (33) 83 59 30 30 – Télécopie : (33) 83 27 83 19

Unions de théories non-disjointes et combinaisons de procédures de satisfaisabilité: premiers résultats

Résumé : Nous présentons dans cet article un cadre formel pour la combinaison des procédures de satisfaisabilité de contraintes par rapport à une théorie contrainte. Nous décrivons une méthode de combinaison générale qui, étant données une procédure de satisfaisabilité pour une théorie $\mathcal{T}_1$ et une autre pour une théorie $\mathcal{T}_2$, est capable de produire une procédure de (semi-)decision pour l'union de $\mathcal{T}_1$ et $\mathcal{T}_2$. On suit une approche fondée sur la théorie des modèles pour fournir également des conditions assurant la correction et la complétude de la méthode, en particulier lorsque les signatures de $\mathcal{T}_1$ et $\mathcal{T}_2$ sont non-disjointes.

Mots-clé : Combinaison de procédures de satisfaisabilité, Problèmes de décision, Raisonnement avec contraintes, Déduction automatique.

Contents

1 Introduction

An established approach to problem solving, common to fields as diverse as Planning and Design, Operation Research, and Image Recognition, represents problem solving in terms of constraint satisfaction: a problem in a particular application domain is described as a set of constraints over the space of the possible solutions. For automated problem solving, a major advantage of constraint-based approaches is efficiency. It is often possible to implement a fast constraint solver for a given constraint domain (or theory) by smartly exploiting some of the features of the domain itself. As a consequence, a major disadvantage of constraint-based approaches is specialization. If a problem requires constraint solving (or, more generally, reasoning) outside the specific constraint domain, a constraint solver alone is not enough.

An emerging deductive paradigm, which comes in many incarnations and flavors (Theory Resolution [Sti85], the Substitutional Framework [Fri91], Constrained Resolution [B9̈4], Constrained Logic [Com93], Constraint Logic Programming [JM94, HS88], Deduction With Constraints [KKR90], T-resolution [FP95], Theory Consolution [BFP92]), tries to combine the advantages of both general purpose and specialized problem solving by providing hybrid frameworks in which a general-purpose reasoner is augmented by number of fast, specialized solvers.

Although this paradigm, which we generically refer to as *Constraint-based Reasoning*, has proven rather successful—especially in Logic Programming—it has been limited so far to frameworks operating essentially over a *single* constraint domain or theory. Many potential applications of Constraint-based Reasoning, however, include input problems over several constraint domains. Syntactically, these are problems expressed in a combination of the constraint languages corresponding to each constraint domain. Operationally, they are problems whose solution would require a main (general-purpose) reasoner to interact with and coordinate several specialized solvers.

In general terms, to deal with such problems the main reasoner must be able to 1) extract from an input problem specification those parts that can be solved by a particular constraint solver, 2) assign each extracted subproblem to the corresponding solver, and 3) compose the various solvers' (local) solutions into solutions of the original problem. In other words, the main reasoner must be able, in effect, to combine the various constraint solvers into a *virtual* solver for a constraint domain or theory which is, in turn, a combination of the solvers' domains or theories. The reason that to date there are almost no frameworks or systems able to do this[1] is probably that the combination of constraint domains or theories and their solvers raises challenging model-theoretic and computational issues, which have only

[1]The Constraint Logic Programming systems PrologIII and CLP($\mathcal{R}$) provide a limited form of domain combination essentially by adopting a multi-sorted constraint domain.

recently begun to be investigated. In general, most model-theoretic properties of single constraint domains and theories and most computational properties of single constraint solvers are not closed under "combination". For every sensible notion of combination, domains and solvers have to satisfy more or less stringent conditions to be meaningfully combinable.

1.1 Related Work

Most of the current work on combination has been restricted to the Unification Problem and derivates (disunification, matching, and so on) [BS92, BS93, Bou90, Bou93, DKR94, Her86, KR94a, KR94b, Rin92, SS88, SS89, Tid86, Yel87b, Yel87a]. In essence, in this context the input language is restricted to quantifier-free formulas over a functional signature (no predicate symbols other than equality), the constraint theories to be combined are equational theories typically with disjoint signatures, theory combination is defined simply as (set-theoretic) union, and constraint satisfiability with respect to a theory is defined as satisfiability in a certain free model of the theory. Very little work exists on the combination of more general constraint languages and theories [NO79, Sho84, CLS96, KR94b] or domains [BS95a, BS95b].

One of the first general methods for combining constraint solvers was proposed by Nelson and Oppen in [NO79]—although not literally in the way we describe it here[2]. Besides in its relative generality, the Nelson-Oppen method's appeal also lies in its easy integrability, in principle, into most Constraint-based Reasoning framework. The integration of the method into the Constraint Logic Programming Scheme is described in [TH98]. The method's main limitations are its restriction to a quantifier-free language of constraints and, more importantly, to component theories that share no function or predicate symbols besides the equality symbol.

1.2 Our Approach

The problem of combining non-disjoint theories is of prime interest for real-life applications but since it cannot be achieved in full generality it is important to identify appropriate sub-classes of non-disjoint theories for which a modular approach is still possible. Modular aspects of non-disjoint first-order theories have extensively investigated in Rewriting Theory (see [Gra96] for instance). Given two rewrite systems satisfying some property P, the problem is to determine whether this property P still holds for the union of rewrite systems. Such kind of problem was first addressed for disjoint union of rewrite systems, and then extended to some non-disjoint cases, where for instance the shared function symbols are in fact *constructor symbols* [MT93, Mid94]. An attempt to extend combination techniques developed for

[2] Strictly speaking, theirs is a method for combining decision procedures for the validity of universal formulas with respect to a given first-order theory. More precisely, if P_i is a procedure that decides validity of universal sentences in a theory $\mathcal{T}_i$ $(i = 1, \ldots, n)$, the Nelson-Oppen method yields a procedure for deciding validity of universal sentences in the theory $\mathcal{T}_1 \cup \cdots \cup \mathcal{T}_n$.

unification to certain classes of non-disjoint union of equational theories was first presented in [Rin92, DKR94]. This extension is still based on an adequate notion of *constructors*. Then, the same ideas have been applied to extend the Nelson-Oppen method [Rin96b].

In this paper, which combines and extends the results in [Rin96b] and [TH96], we fully present a combination method that, although inspired in spirit by Nelson and Oppen's, is more general than theirs in many respects. In particular, we do not restrict ourselves to *free* constructors (syntactic equality between shared terms) but we also allow some relations between constructors.

Our method is built independently of Nelson and Oppen's results, which in fact can be given as corollaries of ours, and so we will not describe their approach here. The reader interested in a description of the original Nelson-Oppen method is referred instead to the original paper [NO79] and its sequels [Nel84, Opp80].

1.3 Paper Outline

The paper is organized as follows. In Section 2, we introduce a number of key concepts and lemmas that we will be using in the paper. While the concepts in this section are essential to understanding the rest of the paper, the various results in Section 2.2 should be probably skipped on a first reading. In Section 3, we present the main issues of the "combination problem." We first describe a way to *meaningfully* combine two constraint domains with different languages into a constraint domain for the union language. Then we show how the satisfiability problem in a combination domain can be solved by solving corresponding satisfiability problems in the component domains. In Section 4, we lift our domain combination results to the combination of theories. There, we introduce the pivotal notion of *N-O-combinability* which defines the model-theoretic requirements that any two theories have to satisfy for the extended Nelson-Oppen method to be applicable to them. In Section 5, we then describe our version of the extended combination method and show its correctness. N-O-combinability is a rather abstract notion and therefore it may be difficult in practice to identify pairs of N-O-combinable theories. In Section 6 then, we elaborate on this problem and discuss more concrete criteria for showing that two theories are N-O-combinable and hence that the extended combination method is applicable to the satisfiability problem in their union. In Section 8, we discuss the currently unresolved issues of our approach and point out the directions we intend to follow to further develop our combination framework.

2 The Constraint Satisfiability Problem

We will only consider constraint satisfiability problems that are expressible in a first-order language. In such a context, a *constraint domain* is formalized by an appropriate first-order structure, a *constraint problem* by a (possibly complex) first-order formula, the problem's *variables* by the free variables of the formula, the problem's *solutions* by a set of maps from the free variables into the universe of the structure.[3]

In the most general case, constraint satisfiability considers satisfiability of constraints with respect to a whole *class* of constraint domains. A constraint is satisfiable if it has solutions in one of the domains in the class. In practice, a class of domains is almost invariably either a singleton class or a (first-order) axiomatizable class. In the first case, the given class is composed by a single structure—the integer/real/complex numbers, the term algebra, the finite domains, the booleans, and so on. In the second case, the class is composed of the structures that model a certain first-order theory.

There is an on-going debate in the field on whether it is better to consider satisfiability with respect to single structures or classes thereof. We will not enter this debate here, but for the sake of clarity we would like to point out that the first option is often really an idealization. In fact, constraint satisfiability which respect to a single domain is often undecidable for rich enough domains. The approach commonly followed to work with these constraint domains then is to identify a (decidable) theory that best axiomatizes the domain[4] and then approximate the constraint satisfiability problem by a corresponding *constraint entailment* problem, where a constraint is considered satisfiable if its existential closure is entailed by the theory, that is, it is satisfiable in *every* model of the theory.

In this paper, we will not be concerned with the entailment problem. We will restrict our attention to constraint satisfiability with respect to an axiomatizable class of constraint domains. In other words, we will fix a constraint theory and say that a constraint is satisfiable if it has solutions in *any* model of the theory.

2.1 Notation and Conventions

We assume that the reader is familiar with the basic concepts of Mathematical Logic [Sho67], Universal Algebra [Wec92], and Model Theory [Hod93b]. A specific description of the notation we use in this paper and the main conventions we adopt is given below.

The letters a, b, x, y, z in general denote elements of a set, u, v, w denote logical variables, r, s, t denote first-order terms, φ, ψ, γ denote first-order formulas, $\top$ denotes the constantly true formula, Σ, Ω, Δ denote signatures, that is, sets of function and predicate symbols each with an associated arity. Some of the above letters may be subscripted and/or have an over-tilde which will represent a finite sequence. For

[3] For a more general approach in the same spirit, see [Smo89].

[4] For instance, the theory of real closed fields for the domain of the real numbers, Pressburger arithmetic for the domain of the integer numbers, and so on.

instance, $\tilde{x}$ stands for an n-sequence of the form $(x_1, x_2, \ldots, x_n)$, for $n \geq 0$.[5] We will call $\tilde{x}$ *discrete* if all of its elements are distinct. We denote by $\tilde{x}; \tilde{y}$ the sequence obtained by concatenating $\tilde{x}$ with $\tilde{y}$. Where R is any symbol standing for a binary relation, $\tilde{x}$ stands for $(x_1, \ldots, x_n)$, and $\tilde{y}$ stands for $(y_1, \ldots, y_n)$, $\tilde{x} R \tilde{y}$ will denote the set $\{x_1 R y_1, \ldots, x_2 R y_n\}$. For uniformity, we will use the tilde notation also for members of a Cartesian product. When convenient, we will also treat $\tilde{x}$ as the set of its elements.

In general, $\mathcal{V}ar(t)$ is the set of t's variables and $\mathcal{V}ar(\varphi)$ is the set of φ's free variables. The notation $\varphi(\tilde{v})$ is used to indicate that the free variables of φ are *exactly* the ones in $\tilde{v}$, that is, $\mathcal{V}ar(\varphi(\tilde{v})) = \tilde{v}$. Analogously, $t(\tilde{v})$ is used to indicate that the variables of t are exactly the ones in $\tilde{v}$.[6] In both cases, it is understood that the elements of $\tilde{v}$ are all distinct. Whenever we write $f(\tilde{v})$, where f is a functor, it is also understood that the length of $f(\tilde{v})$ equals the arity of f. The shorthands $\tilde{\exists}\,\varphi$ and $\tilde{\forall}\,\varphi$ stand for the existential, respectively universal, closure of φ. The symbol $\equiv$ is used in a formula to denote equality; $s \not\equiv t$ is an abbreviation for $\neg(s \equiv t)$. We will systematically identify the union of finite sets of formulas with their logical conjunction and vice versa.

If $h : A \to B$ is a map and $\tilde{a} \in A^n$, the expression $h(\tilde{a})$ denotes the tuple $(h(a_1), \ldots, h(a_n))$. If R is an n-ary relation over A then, the expression $h(R)$ denotes the relation $\{h(\tilde{a}) \mid \tilde{a} \in R\}$. If A is a set, $Card(A)$ denotes the cardinality of A.

Unless otherwise specified, V will always denote a fixed, countably-infinite set of variables. If Σ is a signature, Σ^{P} denotes its subset of predicate symbols and Σ^{F} its subset of functors, that is, function and constant symbols; for any subset U of V, $T(\Sigma, U)$ denotes the set of (first-order) Σ-terms over U; $Card(\Sigma)$ denotes the cardinality of Σ, unless Σ has only finitely many symbols; in that case $Card(\Sigma)$ is defined by convention as (the cardinality of) the first infinite ordinal, that is, $Card(\Sigma) := \omega$.

In general, $\mathcal{L}$ denotes a sub-language of the language of the first-order formulas, that is, a syntactically definable class of first-order formulas (such as, for instance, the class of atomic/ existential/equational/Horn/ ... formulas). The notation $\mathcal{L}^{\Sigma}$ restricts the formulas of $\mathcal{L}$ to a specific signature Σ. Analogously, *Qff* (Qff^{Σ}) denotes the class of all the quantifier-free (Σ)-formulas. For convenience, we will *always* assume that $\top \in \mathcal{L}^{\Sigma}$ for any $\mathcal{L}$ and Σ.

The letters $\mathcal{A}$, $\mathcal{B}$, $\mathcal{C}$, $\mathcal{F}$, possibly with subscripts, denote first-order structures. The corresponding Roman letter denotes the universe of the structure. Unless otherwise specified, the symbol Σ subscripted with the corresponding Roman letter denotes the signature of the structure $(\Sigma_A, \Sigma_{A_1}, \Sigma_B, \ldots)$.

[5] Notice that $\tilde{x}_1$ denotes a sequence of index 1, not the first element of the sequence $\tilde{x}$.

[6] This notation is non-standard, as $\varphi(\tilde{v})$ generally indicates that the free variables of φ are *included* in $\tilde{v}$. We use it here because it greatly simplifies the enunciation of most of our results.

If $\mathcal{A}$ is an Ω-structure, $Card(\mathcal{A})$ denotes the cardinality of A; if s is a symbol of Ω, $s^{\mathcal{A}}$ denotes the interpretation of s given by $\mathcal{A}$; if $\Sigma \subseteq \Omega$, $\mathcal{A}^{\Sigma}$ denotes the reduct of $\mathcal{A}$ to Σ; if $X \subset A$ and X is disjoint with Ω, $\mathcal{A}^{\Omega \cup X}$ denotes the natural expansion of $\mathcal{A}$ to $\Omega \cup X$ in which every element of X is a constant symbol denoting itself; if φ is an Ω-sentence, $\mathcal{A} \models \varphi$ means that $\mathcal{A}$ satisfies φ or, equivalently, that φ is true in $\mathcal{A}$; if $\varphi(x_1, \ldots, x_n)$ is an Ω-formula and α a valuation on $\mathcal{A}$ of $(x_1, \ldots, x_n)$, $(\mathcal{A}, \alpha) \models \varphi$ means that α satisfies φ in $\mathcal{A}$; alternatively, if $\alpha(x_i) = a_i$ for $i \in \{1, \ldots, n\}$, we will sometimes write $\mathcal{A} \models \varphi[a_1, \ldots, a_n]$; in either case, we will say that α is an $\mathcal{A}$-solution of φ; if $t(\tilde{v})$ is an Ω-term and α a valuation of $\tilde{v}$, $t^{(\mathcal{A},\alpha)}$ denotes the interpretation of t given by $(\mathcal{A}, \alpha)$. If $\mathbf{K}$ is a class of Ω-structures and φ an Ω-sentence, we will say that $\mathbf{K}$ entails φ and write $\mathbf{K} \models \varphi$ if $\mathcal{A} \models \varphi$ for all $\mathcal{A} \in \mathbf{K}$.

We will only consider *first-order* theories with equality, which means, in particular, that the equality symbol will be treated as a logical constant. If $\mathcal{T}$ is a theory, that is, a set of sentences, $\Sigma_{\mathcal{T}}$ stands for the smallest signature that includes all the non-logical symbols of $\mathcal{T}$; for any $\Sigma \subseteq \Sigma_{\mathcal{T}}$, $\mathcal{T}^{\Sigma}$ denotes the set of all the Σ-sentences entailed by $\mathcal{T}$; $Mod(\mathcal{T})$ denotes the set of all the $\Sigma_{\mathcal{T}}$ structures that model $\mathcal{T}$. Unless specified otherwise, whenever we say that some structure $\mathcal{A}$ is a model of a theory $\mathcal{T}$, we mean that $\mathcal{A}$ is a $\Sigma_{\mathcal{T}}$-model of $\mathcal{T}$—that is, $\mathcal{A} \in Mod(\mathcal{T})$. We will generally identify every theory $\mathcal{T}$ with its deductive closure with respect to the class of first-order $\Sigma_{\mathcal{T}}$-sentences. Because of this, we will also follow the common abuse of calling *empty theory* (of some signature Σ) the deductive closure (with respect to Σ) of the empty set of formulas.

2.2 Preliminary Notions

2.2.1 Terms and Substitutions

Where Σ is a functional signature and V a countably infinite set of variables, we will also denote by $T(\Sigma, V)$ the term algebra with signature Σ and generators V, that is, the Σ-structure in which every term is interpreted as itself. This algebra is provably absolutely-free over V (see later).

Definition 2.1 (Substitution) *Given a functional signature Σ, a* substitution *is an endomorphism of $T(\Sigma, V)$ whose restriction to V coincides with the identity except for finitely many places.*

The identity of $T(\Sigma, V)$ is called the *empty* substitution and usually denoted by ε. We will follow the established convention of writing substitution applications in postfix form. In a common abuse of terminology, if σ is a substitution, the set

$$\mathcal{D}om(\sigma) \quad := \quad \{v \in V \mid v\sigma \neq v\}$$

is called the *domain* of σ while the set

$$\mathcal{R}an(\sigma) \quad := \quad \{v\sigma \mid v \in \mathcal{D}om(\sigma)\}$$

is called the *range* of σ. We will also often use the set

$$\mathcal{VR}an(\sigma) \quad := \quad \mathcal{V}ar(\mathcal{R}an(\sigma)).$$

We will use the notation $\{v_1 \leftarrow t_1, \ldots, v_n \leftarrow t_n\}$ to denote a substitution σ such that $\mathcal{D}om(\sigma) = \{v_1, \ldots, v_n\}$ and $v_i\sigma = t_i$ for all $i \in \{1, \ldots, n\}$. We are mainly interested in *idempotent* substitutions, that is, substitutions σ such that $\sigma \circ \sigma = \sigma$. For each $U \subseteq V$, we will denote by $\mathrm{SUB}(U)$ the set of idempotent substitutions whose domain (in the sense above) is included in U. Substitutions are extended from terms to arbitrary first-order formulas (and sets thereof) by renaming quantified variables when necessary to avoid capturing of free variables.

We mainly consider two special types of substitutions: identifications of variables and instantiations into non-variable terms.

Definition 2.2 (Identification) *If $\tilde{v}$ is a (finite) set of variables, we define the set of* identifications of $\tilde{v}$ *as follows,*

$$\mathrm{ID}(\tilde{v}) \quad := \quad \{\sigma \in \mathrm{SUB}(\tilde{v}) \mid \mathcal{VR}an(\sigma) \subseteq \tilde{v} \setminus \mathcal{D}om(\sigma)\}.$$

To every identification $\xi \in \mathrm{ID}(\tilde{v})$, we will associate the set

$$\xi_{\neq} := \bigcup_{u,v \in \tilde{v}\xi,\ u \neq v} \{u \not\equiv v\}. \tag{1}$$

Every substitution in $\mathrm{ID}(\tilde{v})$ defines a partition of $\tilde{v}$ and identifies all the variables in the same block with a representative of that block. Observe that the empty substitution belongs to $\mathrm{ID}(\tilde{v})$ for any $\tilde{v}$ and that the associated set, which we will denote by $\varepsilon_{\neq}(\tilde{v})$, is made of all the possible disequations between the elements of $\tilde{v}$. Also observe that $\varepsilon_{\neq}(\tilde{v})$ is satisfied exactly when no two variables in $\tilde{v}$ are assigned with the same individual.

Definition 2.3 (Instantiation) *If $\tilde{v}$ is a (finite) set of variables and Σ a finite signature, we define the set of* Σ-instantiations of $\tilde{v}$ *as follows,*

$$\mathrm{IN}^{\Sigma}(\tilde{v}) \quad := \quad \{\sigma \in \mathrm{SUB}(\tilde{v}) \mid \mathcal{R}an(\sigma) \subseteq T(\Sigma, V) \setminus V\}.$$

In the following, we will only consider Σ-instantiations whose domain is entirely made of fresh variables. To every instantiation $\rho \in \mathrm{IN}^{\Sigma}(\tilde{v})$, we will associate the set

$$\rho^{\Sigma}_{\neq} := \bigcup_{v \in \mathcal{VR}an(\rho),\ f_i \in \Sigma^{\mathrm{F}}} \{\forall \tilde{u}_i\ v \not\equiv f_i(\tilde{u}_i)\}, \tag{2}$$

which we simply denote by $\rho_{\neq}$ when Σ is clear from the context.

Observe that the empty substitution belongs to $\mathrm{IN}^{\Sigma}(\tilde{v})$ for any $\tilde{v}$ and Σ and that the associated set, which we will always denote by $\varepsilon^{\Sigma}_{\neq}(\tilde{v})$, is satisfied if and only if all the variables in $\tilde{v}$ are assigned with individuals that are not in the range of any Σ-function.

The following type of formula will play a crucial role in our combination results.

Definition 2.4 (Σ-Restricted Formula) *Let $\mathcal{L}^{\Omega}$ be a class of Ω-formulas and Σ a finite subset of Ω. We say that a formula $\psi(\tilde{u})$ is* Σ-restricted on $\tilde{v}$ *if it is logically equivalent to a formula ψ' of the form*

$$\varphi(\tilde{u}) \wedge \varepsilon_{\neq}(\tilde{v}) \wedge \varepsilon^{\Sigma}_{\neq}(\tilde{v}),$$

where $\varphi \in \mathcal{L}^{\Omega}$ and $\tilde{v} \subseteq \tilde{u}$. We say that ψ' is a Σ-restricted form of φ on $\tilde{v}$. *We call φ the* body *of ψ' and $\varepsilon_{\neq}(\tilde{v}) \wedge \varepsilon^{\Sigma}_{\neq}(\tilde{v})$ the* Σ-restriction *of ψ'.*

Notice that if a formula is Σ-restricted on some $\tilde{v}$, it is satisfied in a structure $\mathcal{A}$ only by valuations mapping the variables of $\tilde{v}$ to distinct Σ-isolated individuals of $\mathcal{A}$. In general, we will simply say that a formula is Σ-restricted if it is Σ-restricted on a subset of its free variables. We will say that the formula is *totally Σ-restricted* if it is Σ-restricted on *all* of its free variables.

We will denote by $Res(\mathcal{L}^{\Omega}, \Sigma)$ the class of all the formulas with a Σ-restricted form whose body belongs to $\mathcal{L}^{\Omega}$ and by $TRes(\mathcal{L}^{\Omega}, \Sigma)$ the class of all the formulas with a totally Σ-restricted form whose body belongs to $\mathcal{L}^{\Omega}$. Note that $\mathcal{L}^{\Omega}$ and $TRes(\mathcal{L}^{\Omega}, \Sigma)$ are always included in $Res(\mathcal{L}^{\Omega}, \Sigma)$.

2.2.2 Structures and Theories

The following definition formalizes the notion of satisfiability that we will use throughout the paper.

Definition 2.5 *We say that a Σ-formula φ is satisfiable in a Σ-structure $\mathcal{A}$ if its* existential *closure is satisfied by $\mathcal{A}$, that is, if $\mathcal{A} \models \tilde{\exists}\, \varphi$.*[7] *If* **K** *is a class of Σ-structures, we say that φ is* satisfiable in **K** *if it is satisfiable in some member of* **K**.

As mentioned then, given a Σ-theory $\mathcal{T}$, we will say a that a Σ-formula φ is *satisfiable in $\mathcal{T}$* if it is satisfiable in $Mod(\mathcal{T})$. By the above, it follows that φ is satisfiable in $\mathcal{T}$ exactly when the theory $\mathcal{T} \cup \{\tilde{\exists}\, \varphi\}$ is consistent.

We will implicitly appeal to the following lemma almost constantly in the rest of the paper.

[7]Notice that some authors define satisfiability of formulas with respect to satisfiability of their universal closures instead.

Lemma 2.6 Let $\mathcal{A}$ be a Σ-structure, $\varphi(\tilde{v})$ a Σ-formula, and $\tilde{a}$ a sequence in A having $\tilde{v}$'s length. Then,

$$\mathcal{A} \models \varphi[\tilde{a}] \quad \text{iff} \quad \mathcal{A}' \models \varphi[\tilde{a}]$$

for any expansion[8] $\mathcal{A}'$ of $\mathcal{A}$ to a signature $\Sigma' \supseteq \Sigma$.

The following is another general result involving signature reductions to which we will appeal later. Notice that the result is not as trivial as it looks and, in fact, does not hold if the signature Σ below is strictly contained in $\Sigma_1 \cap \Sigma_2$.

Lemma 2.7 *Let $\mathcal{T}_1$ be a Σ_1-theory, $\mathcal{T}_2$ a Σ_2-theory, and $\Sigma := \Sigma_1 \cap \Sigma_2$. Then for all Σ-sentences φ,*

$$(\mathcal{T}_1 \cup \mathcal{T}_2)^\Sigma \models \varphi \quad \textit{iff} \quad \mathcal{T}_1{}^\Sigma \cup \mathcal{T}_2{}^\Sigma \models \varphi$$

Proof. Let φ be a Σ-sentence and recall that for any theory $\mathcal{T}$,

$$\mathcal{T}^\Sigma := \{\varphi \mid \varphi \ \Sigma\text{-sentence}, \mathcal{T} \models \varphi\}.$$

($\Leftarrow$) Immediate consequence of the obvious fact that $\mathcal{T}_1{}^\Sigma \cup \mathcal{T}_2{}^\Sigma \subseteq (\mathcal{T}_1 \cup \mathcal{T}_2)^\Sigma$.

($\Rightarrow$) Assume that $(\mathcal{T}_1 \cup \mathcal{T}_2)^\Sigma \models \varphi$ or, equivalently, that the theory $\mathcal{T}_1 \cup (\mathcal{T}_2 \cup \{\neg\varphi\})$ is inconsistent. By the Craig Interpolation Lemma [Hod93b], there is a Σ-sentence ψ such that $\mathcal{T}_1 \models \neg\psi$ and $\mathcal{T}_2 \cup \{\neg\varphi\} \models \psi$. By logical reasoning, we also have that $\mathcal{T}_2 \models \neg\psi \rightarrow \varphi$. Observing that both φ and $\neg\psi$ are Σ-sentences, we can then conclude that $\neg\psi \in \mathcal{T}_1{}^\Sigma$ and $(\neg\psi \rightarrow \varphi) \in \mathcal{T}_2{}^\Sigma$ from which the claim follows immediately. $\square$

Our combination result will often require that the theories under consideration satisfy the following property.

Definition 2.8 (Collapse-Free) *We say that a class of Σ-structures or a Σ-theory is* collapse free *if it entails no sentences of the form $\tilde{\forall}\ (v \equiv t)$ where v is a variable and t a Σ-term different from v.*[9]

Notice that a theory $\mathcal{T}$ is collapse-free iff the class $\mathit{Mod}(\mathcal{T})$ is collapse-free and that every collapse-free theory admits non-trivial models (otherwise, it would entail $\tilde{\forall}\ (u \equiv v)$).

In Universal Algebra, equational theories are commonly defined as theories axiomatized by a set of (universally quantified) equations. Here, we extend such notion to theories whose signature may include predicate symbols as well.

[8] Recall that a structure $\mathcal{A}'$ is an expansion of a structure $\mathcal{A}$ if $\mathcal{A}$ is a reduct of $\mathcal{A}'$.

[9] Our definition is slightly more restrictive than the standard one in which t is required to be a non-variable term. According to that definition, if Σ has no functors, the trivial Σ-theory is collapse-free. However, the two definitions coincide for non-trivial consistent theories (see later), the theories of interest in this paper.

Definition 2.9 (Σ-Atomic Theory) *We say that a theory is* Σ-atomic *if it is axiomatized by a set of Σ-formulas of the form $\tilde{\forall}\,\varphi$, where φ is atomic.*

Notice that every equational theory of signature Σ is a Σ-atomic theory because, by definition, it is axiomatized by a set of universally quantified equations between Σ-terms. For this reason we will generally use the symbol E to denote a given Σ-atomic theory.

If E is a Σ-atomic theory, we call $Mod(E)$ a *Σ-variety*. By the above observation, this definition is consistent with that of variety in Universal Algebra. A Σ-variety is *non-trivial* if it contains non-trivial structures, that is, structures of cardinality greater than 1. We will say that E is non-trivial if $Mod(E)$ is non-trivial. In Universal Algebra and Unification Theory a trivial equational theory is often called *inconsistent*. In the context of Model Theory in which this paper moves, inconsistent is used instead as a synonym of *unsatisfiable*, that is, admitting no models. We have chosen the terminology *trivial/non-trivial* here in order to avoid this possible source of confusion. As usual, we will denote by $=_E$ the least congruence on $T(\Sigma, V)$ generated by E. Therefore, given Σ-terms s, t, we will write $s =_E t$ to mean that $E \models \tilde{\forall}\, s \equiv t$.

Where $\mathcal{T}$ is an Ω-theory and $\Sigma \subseteq \Omega$, we will denote by $\mathcal{T}^{\Sigma}_{\mathrm{At}}$ the Σ-*atomic theory of* $\mathcal{T}$, that is, the set of all the universally quantified Σ-atoms entailed by $\mathcal{T}$. Notice that, by definition, $\mathcal{T}$ and $\mathcal{T}^{\Sigma}$ have the same Σ-atomic theory and so the above notation is unambiguous. We will also refer to $Mod(\mathcal{T}^{\Sigma}_{\mathrm{At}})$ as the Σ-*variety of* $\mathcal{T}$ and often identify it with $\mathcal{T}^{\Sigma}_{\mathrm{At}}$. In a totally analogous way, we will call Σ-atomic theory of an Ω-structure $\mathcal{A}$ the set of all the universally quantified Σ-atoms valid in $\mathcal{A}$.

We will use the standard types of morphisms from Model Theory. We recall their definition here for completeness.

Definition 2.10 (Homomorphism, Embedding) *Given two Σ-structures $\mathcal{A}$ and $\mathcal{B}$ and a map $h\colon A \to B$, we say that h is a* homomorphism of $\mathcal{A}$ into $\mathcal{B}$ *if*

1. *for any n-ary $f \in \Sigma^{\mathrm{F}}$ and $\tilde{a} \in A^n$, $h(f^{\mathcal{A}}(\tilde{a})) = f^{\mathcal{B}}(h(\tilde{a}))$ and*
2. *for any n-ary $P \in \Sigma^{\mathrm{P}}$ and $\tilde{a} \in A^n$, $\tilde{a} \in P^{\mathcal{A}}$ implies $h(\tilde{a}) \in P^{\mathcal{B}}$.*

We say that h is an embedding of $\mathcal{A}$ into $\mathcal{B}$ *if*

1. *h is an injective homomorphism of $\mathcal{A}$ into $\mathcal{B}$ and*
2. *for any n-ary $P \in \Sigma^{\mathrm{P}}$ and $\tilde{a} \in A^n$, $h(\tilde{a}) \in P^{\mathcal{B}}$ implies $\tilde{a} \in P^{\mathcal{A}}$.*

We say that $\mathcal{A}$ is embeddable in $\mathcal{B}$ *if there exists an embedding of $\mathcal{A}$ into $\mathcal{B}$.*

Definition 2.11 (Isomorphism) *Given two Σ-structures $\mathcal{A}$ and $\mathcal{B}$ and a map h we say that h is an* isomorphism of $\mathcal{A}$ onto $\mathcal{B}$, *and write $h : \mathcal{A} \cong \mathcal{B}$, if h is a surjective embedding of $\mathcal{A}$ into $\mathcal{B}$. We say that h is an* automorphism of $\mathcal{A}$ *if h is an isomorphism of $\mathcal{A}$ onto itself.*

We will say that $\mathcal{A}$ and $\mathcal{B}$ are *isomorphic*, and write $\mathcal{A} \cong \mathcal{B}$, if there exists an isomorphism of $\mathcal{A}$ onto $\mathcal{B}$.

Definition 2.12 (Substructure) *Let $\mathcal{A}$ and $\mathcal{B}$ be two Σ-structures such that $A \subseteq B$. We say that $\mathcal{A}$ is an* substructure *of $\mathcal{B}$ and write $\mathcal{A} \subseteq \mathcal{B}$ if the inclusion map of A into B is an embedding of $\mathcal{A}$ into $\mathcal{B}$.*

A formula is called *universal* if it is in Prenex Normal Form and its quantifier prefix contains only universal quantifiers (if any). We will say that a theory is universal if it is axiomatizable by a set of universal sentences. We will refer to the following two well-known properties of universal theories.

Lemma 2.13 *Let $\mathcal{B}$ be a Σ-structure and $\psi(\tilde{v})$ a universal Σ-formula. Assume that $\mathcal{B} \models \psi[\tilde{a}]$ for some $\tilde{a}$ in $\mathcal{B}$. Then, $\psi(\tilde{v})$ is satisfied by every $\mathcal{A} \subseteq \mathcal{B}$ whose carrier includes $\tilde{a}$.*

Lemma 2.14 *Let $\mathcal{T}$ be an universal Σ-theory and $\mathcal{A}$ a Σ-structure. Then, $\mathcal{A}$ is embeddable in a model of $\mathcal{T}$ iff $\mathcal{A}$ is a model of $\mathcal{T}$. In particular, a Σ-structure is a model of $\mathcal{T}$ iff it is a substructure of a model of $\mathcal{T}$.*

If $\mathcal{A}$ is a Σ-structure and $X \subseteq A$, there is a unique smallest substructure of $\mathcal{A}$ including X. We will indicate such substructure with $\langle X \rangle_{\mathcal{A}}$. If $\mathcal{A} = \langle X \rangle_{\mathcal{A}}$ for some $X \subseteq A$, we usually say that X *generates* $\mathcal{A}$, or that X is *a set of generators for* $\mathcal{A}$.

We say that X is a *non-redundant* set of generators for $\mathcal{A}$ if X generates $\mathcal{A}$ and no proper subset of X generates $\mathcal{A}$. We say that $\mathcal{A}$ is *finitely-generated* if it admits a finite set of generators. While every structure admits a set of generators (its own carrier, for instance), not every structure admits a *non-redundant* set of generators. It is easy to show, however, that all finitely generated structures do so.

Lemma 2.15 *Let Y be a non-redundant set of generators for a structure $\mathcal{A}$. Then, for all $X \subseteq Y$, X is a non-redundant set of generators for $\langle X \rangle_{\mathcal{A}}$.*

Proof. If X is redundant there is an non-empty set $X' \subseteq X$ such that $X \setminus X'$ generates $\langle X \rangle_{\mathcal{A}}$. But then $Y \setminus X'$ generates $\mathcal{A}$, against the assumption that Y is non-redundant for $\mathcal{A}$. □

For brevity, we will often make use of the following definitions.

Definition 2.16 (Σ-Isolated Individual) *Let $\mathcal{A}$ be a structure and Σ a subset of Σ_A. We say that $a \in A$ is a* Σ-isolated individual *of $\mathcal{A}$ if a is not in the range of the interpretation of any function symbol of Σ, that is, if there is no $g \in \Sigma^{\mathrm{F}}$ and tuple $\tilde{x}$ in A such that $a = g^{\mathcal{A}}(\tilde{x})$.*

We say that an individual a is, simply, an *isolated individual* of a structure $\mathcal{A}$ if a is a Σ_A-isolated individual of $\mathcal{A}$. Since the set of $\mathcal{A}$'s Σ-isolated individuals coincides with the set of $\mathcal{A}^\Sigma$'s isolated individuals, for every $\Sigma \subseteq \Sigma_A$, we will use $Is(\mathcal{A}^\Sigma)$ to denote either of them.

Definition 2.17 (Σ-generators) *Let $\mathcal{A}$ be a structure and Σ a subset of Σ_A. We say that $\mathcal{A}$ is* Σ-generated *by a set $X \subseteq A$, or that X is a set of* Σ-generators *of $\mathcal{A}$, if $\mathcal{A}^\Sigma$ is generated by X.*

With a slight abuse we will also say that a *set* $A' \subseteq A$ is Σ-generated by a set $X \subseteq A$, if A' is contained in the carrier of the substructure of $\mathcal{A}^\Sigma$ generated by X ($A' \subseteq B$ where $\mathcal{B} := \langle X \rangle_{\mathcal{A}^\Sigma}$).

It is immediate that when $(\Sigma_A)^{\mathrm{F}} \subseteq \Sigma \subseteq \Sigma_A$, the notions of generators and Σ-generators coincide. It should also be immediate that the set of Σ-isolated individuals of a structure is necessarily included in every set of Σ-generators for that structure.

Definition 2.18 (Σ-Independent Set) *Let $\mathcal{A}$ be a structure, $\Sigma \subseteq \Sigma_A$, and $X \subseteq A$. We say that X is* Σ-independent *in $\mathcal{A}$ if no $\{x\} \subseteq X$ is Σ-generated by $X \setminus \{x\}$.*

Notice that any set of Σ-isolated individuals of a structure $\mathcal{A}$ is Σ-independent in $\mathcal{A}$ and that a set of Σ-generators for $\mathcal{A}$ is non-redundant if and only if it is Σ-independent in $\mathcal{A}$.

Lemma 2.19 *Let $\mathcal{B}$ be an uncountable structure and Σ a countable subset of Σ_B. Then, for all $\tilde{b}$ in B, there is a countably infinite subset of B which is Σ-independent in $\mathcal{B}$ and Σ-generates $\tilde{b}$ in $\mathcal{B}$.*

Proof. Since $\tilde{b}$ is finite, there certainly is a finite subsets of $\tilde{b}$ (possibly the empty set) which is Σ-independent and generates $\tilde{b}$ in $\mathcal{B}$. Let X_0 be such a set. Then, there must be an $x_1 \in B \setminus X_0$ such that $X_1 := X_0 \cup \{x_1\}$ is Σ-independent in $\mathcal{B}$. Otherwise, $\mathcal{B}$ would be Σ-generated by X_0, which is impossible as both X_0 and Σ are countable while $\mathcal{B}$ is not. Iterating the above argument, we can define a family $\{X_n \mid n < \omega\}$ of finite, Σ-independent subsets of B such that $X_n \subset X_{n+1}$ for all $n < \omega$. Let $X := \bigcup_{n<\omega} X_n$. For including X_0, X Σ-generates $\tilde{b}$. We show by contradiction that X, which is countably infinite, is Σ-independent in $\mathcal{B}$.

Assume that there is an $x_i \in X$ which is contained in the carrier of $\langle X \setminus \{x_i\}\rangle_{\mathcal{B}^\Sigma}$. Then, we can show that there is a *finite* subset $\tilde{x}$ of X such that x_i is contained in the carrier of $\langle \tilde{x} \rangle_{\mathcal{B}^\Sigma}$. Clearly, there is an $n < \omega$ such that X_n includes $\tilde{x} \cup \{x_i\}$. But then, by the above, X_n is not Σ-independent in $\mathcal{B}$. □

2.2.3 Free Structures

A fundamental notion of Universal Algebra is that of *free algebra.* This notion can be extended to structures in a natural way (see [Hod93b] for example). Our attention for free structures here is motivated by the crucial role they play in our combination results. We will adopt the following among the several (equivalent) definitions of free structure found in the literature.

Definition 2.20 (Free Structure) *Given a class* $\mathbf{K}$ *of* Σ*-structures and a set* X*, we say that structure* $\mathcal{A} \in \mathbf{K}$ *is* free over X in $\mathbf{K}$ *if* $\mathcal{A}$ *is generated by* X *and every map from* X *into the universe of a structure* $\mathcal{B} \in \mathbf{K}$ *can be extended to a (unique) homomorphism of* $\mathcal{A}$ *into* $\mathcal{B}$*. We call* X *a* basis *of* $\mathcal{A}$*.*

We will say that a Σ-structure is *absolutely free* if it is free in the class of all the Σ-structures.

It is immediate from the definition that a Σ-structure $\mathcal{A}$ is free in some class of Σ-structures if and only if it is free in the singleton class $\{\mathcal{A}\}$. As a consequence, we will simply say that a structure $\mathcal{A}$ is free (over X) if it is free in $\{\mathcal{A}\}$ (over X). Free structures have the following characterization.

Proposition 2.21 ([Hod93b]) *Let* $\mathcal{A} \in \mathbf{K}$ *for some class* $\mathbf{K}$ *of* Σ*-structures, and* $X \subseteq \mathcal{A}$*. Then,* $\mathcal{A}$ *is free over* X *in* $\mathbf{K}$ *iff*

1. X *generates* $\mathcal{A}$ *and*
2. $\mathbf{K} \models \tilde{\forall}\, \varphi$ *for all* Σ*-atoms* $\varphi(\tilde{v})$ *for which there is a discrete tuple* $\tilde{x}$ *in* X *such that* $\mathcal{A} \models \varphi[\tilde{x}]$*.*

It is not difficult to show that every basis of a free structure is non-redundant as a set of generators and that the same structure can be free over more than one basis. Free structures in a collapse-free class, however, have unique bases.

Proposition 2.22 *The basis of a structure free in a collapse-free class is unique and coincides with the set of the structure's isolated individuals.*

Proof. Assume that a Σ-structure $\mathcal{A}$ is free over some set X in a collapse-free class of Σ-structures. For being a set of generators for $\mathcal{A}$, X must contain all of $\mathcal{A}$'s isolated individuals, as we observed earlier. Ad absurdum, assume it also contains a non-isolated individual y. Since y is not isolated, there is a non-variable Σ-term $t(\tilde{v})$ and a sequence $\tilde{x}$ in X such that $y = t^{\mathcal{A}}[\tilde{x}]$.[10]

That means that $\mathcal{A}$ satisfies the atomic formula $(u \equiv t)$ with an assignment of elements of X to the formula's variables. By Prop. 2.21 then, the sentence $\tilde{\forall}\,(u \equiv t)$ is valid in the class, against the assumption that the class is collapse free. $\square$

[10] Incidentally, notice that $y \in \tilde{x}$ otherwise X would be redundant.

It can be shown that, given a signature Σ, every free structure is free in some Σ-variety, and in particular, absolutely free structures are free in the Σ-variety of the empty theory. When a structure is free in an axiomatizable class of Σ-structures, a corresponding Σ-variety is readily identified.

Proposition 2.23 *Let* $\mathbf{K} := Mod(\mathcal{T})$ *for some* Σ*-theory* $\mathcal{T}$*. Then, for all* $\mathcal{A} \in \mathbf{K}$ *and* $X \subseteq A$*,* $\mathcal{A}$ *is free over* X *in* $Mod(\mathcal{T})$ *iff* $\mathcal{A}$ *is free over* X *in* $Mod(\mathcal{T}^{\Sigma}_{\mathrm{At}})$*.*

Proof. The right-to-left implication is trivial as $Mod(\mathcal{T}) \subseteq Mod(\mathcal{T}^{\Sigma}_{\mathrm{At}})$. For the converse, let $\varphi(\tilde{v})$ be a Σ-atom and assume that $\mathcal{A} \models \varphi[\tilde{x}]$ for some discrete $\tilde{x}$ in X. By Prop. 2.21, it is enough to show that $\mathcal{T}^{\Sigma}_{\mathrm{At}} \models \tilde{\forall}\, \varphi$. By assumption and thanks to the same proposition, we know that $\mathcal{T} \models \tilde{\forall}\, \varphi$. Recalling the definition of $\mathcal{T}^{\Sigma}_{\mathrm{At}}$, we can then conclude that $\tilde{\forall}\, \varphi \in \mathcal{T}^{\Sigma}_{\mathrm{At}}$, from which the claim follows immediately. □

The connection between Σ-atomic theories and free Σ-structures is highlighted by the following corollary of Prop. 2.21.

Corollary 2.24 *Let* E *be a* Σ*-atomic theory. For all* Σ*-structures* $\mathcal{A}$ *free in* $Mod(E)$ *over an infinite basis, the* Σ*-atomic theory of* $\mathcal{A}$ *coincides with* E*.*[11]

The above result also entails that a free Σ-structure with an infinite basis is free (over that basis) in its own Σ-variety.

The following is another well-known result which is an immediate consequence of the definition of free structure.

Lemma 2.25 *If two* Σ*-structures* $\mathcal{A}$ *and* $\mathcal{B}$ *are free in the same* Σ*-variety over respective bases* X *and* Y *having the same cardinality, then any bijection of* X *onto* Y *extends to an isomorphism of* $\mathcal{A}$ *onto* $\mathcal{B}$*.*

By this result we can identify all the structures that are free in a Σ-variety $\mathbf{K}$ over a set of a given cardinality κ. We will denote any of these structures as $\mathcal{F}_{\kappa}(\mathbf{K})$. Where E is a Σ-atomic theory, we will also use $\mathcal{F}_{\kappa}(E)$ as an abbreviation of $\mathcal{F}_{\kappa}(Mod(E))$.

A consequence of both the definition of free structure and that of Σ-variety is that, modulo isomorphism, all the free structures of a non-trivial Σ-variety form an embedding chain.

Lemma 2.26 *A structure* $\mathcal{A}$ *is free over* X *in a non-trivial* Σ*-variety* $\mathbf{K}$ *iff there is a structure* $\mathcal{B}$ *free in* $\mathbf{K}$ *over* Y *such that* $\mathcal{A} := \langle X \rangle_{\mathcal{B}}$ *and* $X \subset Y$*.*

[11] Actually, the result is even more general. It is possible to show that the positive theory of $\mathcal{A}$ coincides with the positive theory of E.

The lemma above implies that the property of being free is preserved under reduction of the set of generators. The property, however, is not preserved under reduction of the signature.[12] The problem is that, in general, the reduct of a structure may need a larger set of generators than the original structure does. For example, consider the signature $\Omega := \{prec, succ\}$ and the (equational) theory[13]

$$E_1 \quad := \quad \left\{ \begin{array}{l} \forall x \; x \equiv prec(succ(x)), \\ \forall x \; x \equiv succ(prec(x)) \end{array} \right\} \tag{3}$$

The integers $\mathcal{Z}$ are a free model of E_1 over a basis of cardinality 1 when *prec* and *succ* are interpreted in the obvious way. Any singleton set of integers is a basis for $\mathcal{Z}$. The number zero, for instance, generates all the positive integers with (the function denoted by) *succ* and the negative ones with *prec*. Now, if $\Sigma := \{succ\}$, $\mathcal{Z}^\Sigma$ is definitely not free because it does not even admit a non-redundant set of generators which, as we saw, is a necessary condition for a structure to be free.

We can easily show, however, that the reduct of a free structure is itself free whenever the signature reduction does not alter the set of generators; in other words, whenever the remaining symbols are still enough to generate the whole carrier of the structure from the original set of generators.

Proposition 2.27 *Let $\mathcal{T}$ be an Ω-atomic theory, $\mathcal{A}$ an Ω-structure free in Mod($\mathcal{T}$) over some basis X, and $\Sigma \subseteq \Omega$. If X is also a set of generators for $\mathcal{A}^\Sigma$, then $\mathcal{A}^\Sigma$ is free in Mod($\mathcal{T}^\Sigma_{\mathrm{At}}$) over X.*

Proof. Assume that X generates $\mathcal{A}^\Sigma$. Let $\varphi(\tilde{v})$ be a Σ-atom and assume that $\mathcal{A}^\Sigma \models \varphi[\tilde{x}]$ for some discrete $\tilde{x}$ in X. By Prop. 2.21, it is enough to show that $\mathcal{T}^\Sigma_{\mathrm{At}} \models \tilde{\forall}\, \varphi$. Clearly, $\mathcal{A} \models \varphi[\tilde{x}]$. Since $\tilde{x}$ is in the basis of $\mathcal{A}$, it follows by Prop. 2.21 that $\mathcal{T} \models \tilde{\forall}\, \varphi$. But then, $\mathcal{T}^\Sigma_{\mathrm{At}} \models \tilde{\forall}\, \varphi$ as well by construction of $\mathcal{T}^\Sigma_{\mathrm{At}}$ given that φ is a Σ-atom. $\square$

Verifying that a signature reduction does not alter the set of generators is usually not so easy. A sufficient condition can be found in the easily proven proposition below, which uses the following restriction of the common notion of *explicit definability* [Hod93b].

Definition 2.28 (Σ-definable) *Let $\mathcal{T}$ be an Ω-theory, $\Sigma \subseteq \Omega$, and $f \in (\Omega \setminus \Sigma)^{\mathrm{F}}$. We say that f is* Σ-definable in $\mathcal{T}$ *if there is a $t \in T(\Sigma, V)$ and a discrete $\tilde{u}$ in V such that $\mathcal{T} \models \tilde{\forall}\, f(\tilde{u}) \equiv t$.*

An immediate result of this definition is the following.

[12] The reason we are interested in preserving freeness under signature reduction will become apparent later.

[13] With a slight abuse of notation, we will sometimes use the metasymbols x, y, z also as actual variables in examples of formulas.

Proposition 2.29 *Let $\mathcal{T}$ be an Ω-theory, $\Sigma \subseteq \Omega$, and $\mathcal{A}$ an Ω-structure free in the Ω-variety of $\mathcal{T}$ over some infinite set X. If f is Σ-definable in $\mathcal{T}$ for all $f \in (\Omega \setminus \Sigma)^{\mathrm{F}}$, then X is a set of Σ-generators for $\mathcal{A}$.*

Nonetheless, there are free structures some of whose reducts, although requiring a larger set of generators, are still free. Consider the signature $\Omega := \{0, succ, +\}$ and the non-trivial (equational) theory

$$E \quad := \quad \left\{ \begin{array}{ll} \forall x, y, z & x + (y + z) = (x + y) + z \\ \forall x, y & x + y = y + x \\ \forall x, y & x + succ(y) = succ(x + y) \\ \forall x & x + 0 = x \end{array} \right\} \tag{4}$$

Let $\Sigma := \{0, succ\}$, $\Delta := \{+\}$, and let $\mathcal{A}$ be a free model of E over a non-empty basis X. It is easy to see that $\mathcal{A}$ is not Σ-generated by X, as no valuation into X makes a non-variable Δ-term equal to a non-variable Σ-term.[14] It is also easy to see, however, that each Ω-term is equal in $\mathcal{A}$ to a term of the form $succ^n(t)$ where t is either 0 or a Δ-term. All this entails that the closure Y of X under $+^{\mathcal{A}}$ is a non-redundant set of Σ-generators for $\mathcal{A}$. We leave it to the reader to verify that $\mathcal{A}^{\Sigma}$ is absolutely-free over Y and that E^{Σ}_{At} is empty, which entails that $\mathcal{A}^{\Sigma}$ is free in E^{Σ}_{At} over Y.

The theory above belongs to a class of structures that admit *constructors* in a precise sense. To explain our notion of constructors we first need to introduce some notation.

Let E now be any Ω-atomic theory. For every $\Sigma \subseteq \Omega$, we define the following subsets of $T(\Omega, V)$.

$$\begin{array}{lcllll} G_E(\Sigma, V) & := & \{ \, r & \mid & r \in T(\Omega, V), & \\ & & & & r \neq_E f(\tilde{t}) \text{ for all } f \in \Sigma \text{ and } \tilde{t} \text{ in } T(\Omega, V) \, \} \\ & & & & & \\ ST_E(\Sigma, V) & := & \{ \, s\rho & \mid & s(\tilde{v}) \in T(\Sigma, V), \; \rho \in \mathrm{SUB}(V), & \\ & & & & \tilde{v} \subseteq \mathcal{D}om(\rho), \; \tilde{v}\rho \subseteq G_E(\Sigma, V) \, \} \end{array}$$

In essence, $G_E(\Sigma, V)$ is made, modulo E equivalence, of Ω-terms whose root symbol is not in Σ, while $ST_E(\Sigma, V)$ is made of terms that can be obtained by substituting terms from $G_E(\Sigma, V)$ into the variables of a Σ-term. Notice that the substitution in question is idempotent and that its domain is not necessarily restricted the variables of the Σ-term. Furthermore,

$$T(\Sigma, V) \subseteq ST_E(\Sigma, V) \quad \text{and} \quad G_E(\Sigma, V) \subseteq ST_E(\Sigma, V).$$

[14] In other words, the range of $+^{\mathcal{A}}$ contains Σ-isolated individuals that are not in X.

In fact, for all terms $s \in T(\Sigma, V)$, $s = s'\rho$ where s' is a variable-disjoint variant of s and ρ a renaming of the variables of s' into those of s. In addition, for all terms $r \in G_E(\Sigma, V)$, $r = s\rho$ where s is a variable in $V \setminus \mathcal{V}ar(r)$ and ρ the substitution $\{s \leftarrow r\}$.

Definition 2.30 (Constructors) *Let E be a non-trivial Ω-atomic theory, $\Sigma \subseteq \Omega$. We say that Σ is a* set of constructors modulo for E *if the following holds.*

1. $V \subseteq G_E(\Sigma, V)$.
2. *For all $t \in T(\Omega, V)$, there is an $s\rho \in ST_E(\Sigma, V)$ such that*
$$t =_E s\rho.$$
3. *For all n-ary $P \in \Sigma^{\mathrm{P}} \cup \{\equiv\}$ and $\tilde{s}\rho$ in $ST_E(\Sigma, V)$, where $u\rho \neq_E v\rho$ for every distinct $u, v \in \mathcal{D}om(\rho)$,*[15]
$$E \models \tilde{\forall}\, P(\tilde{s}\rho) \quad \textit{iff} \quad E \models \tilde{\forall}\, P(\tilde{s}).$$

In general, we will say that Σ is a set of constructors for an arbitrary Ω-theory $\mathcal{T}$, if Σ is a set of constructors for $\mathcal{T}^{\Omega}_{\mathrm{At}}$.

When Ω has no predicate symbols, *3* reduces to:

3. *For all $s\rho, s'\rho \in ST_E(\Sigma, V)$ where $u\rho \neq_E v\rho$ for every distinct $u, v \in \mathcal{D}om(\rho)$,*
$$s\rho =_E s'\rho \quad \textit{iff} \quad s =_E s',$$

We leave it to the reader to verify that, $\Sigma := \{0, succ\}$ is a set of constructors for the theory E in the last example above.

Below, we show that free structures in theories with constructors do admit free reducts. For this, we will assume that Σ, Ω, E are defined as in Def. 2.30 and start with the following easy lemmas.

Lemma 2.31 *Let $\mathcal{A}$ a structure free in $Mod(E)$ over the countably infinite set X. Where α is a bijection of V onto X, let*
$$Y \quad = \quad \{r^{(\mathcal{A},\alpha)} \mid r \in G_E(\Sigma, V)\}.$$
If Σ is a set of constructors for E, then

- *Y is a set of Σ-generators for $\mathcal{A}$,*

[15] We have used of the same ρ simply for notational convenience. Equivalently but less compactly, we could have chosen a ρ_i for each s_i in $\tilde{s} = (s_1, \ldots, s_n)$ and added the requirement that for all $i, j \in \{1, \ldots, n\}$ ρ_i and ρ_j agree on the variables shared by s_i and s_j.

- $X \subseteq Y$,
- $Y = Is(\mathcal{A}^\Sigma)$.

Proof. We start by noting that by the definition of generator, under the interpretation $(\mathcal{A}, \alpha)$ each individual of A is denoted by (at least) a term of $T(\Omega, V)$. We first show that for all $y \in A$, $y \in Is(\mathcal{A}^\Sigma)$ iff $y \in Y$.

In fact, let $y \in Is(\mathcal{A}^\Sigma)$ and choose any $r \in T(\Omega, V)$ such that $y = r^{(\mathcal{A},\alpha)}$. We show that $r \in G_E(\Sigma, V)$ which entails that $y \in Y$ by construction of Y. Assume that $r \notin G_E(\Sigma, V)$. Then, there is an $f \in \Sigma$ and $\tilde{t}$ in $T(\Omega, V)$ such that $r =_E f(\tilde{t})$. In particular, since $\mathcal{A}$ is a model of E and α a valuation of $\mathcal{V}ar(r \equiv f(\tilde{t}))$, this means that $y = r^{(\mathcal{A},\alpha)} = f(\tilde{t})^{(\mathcal{A},\alpha)}$. But this contradicts the fact that y is Σ-isolated. Conversely, let $y \in Y$, choose any $r \in G_E(\Sigma, V)$ such that $y = r^{(\mathcal{A},\alpha)}$, and assume that y is not Σ-isolated. Then, there must be an $f \in \Sigma$ and $\tilde{t} \in T(\Omega, V)$ such that $(\mathcal{A}, \alpha) \models r \equiv f(\tilde{t})$. By Prop. 2.21 then, we can conclude that $r =_E f(\tilde{t})$, against the fact that $r \in G_E(\Sigma, V)$.

That $X \subseteq Y$ is an immediate consequence of the assumption $V \subseteq G_E(\Sigma, V)$ at Point 1 of Def. 2.30. To see that $\mathcal{A}$ is Σ-generated by Y, simply notice that, by Point 2 of Def. 2.30, every term $t \in T(\Omega, V)$ is equivalent in E to a term of the form $s(r_1, \ldots, r_m)$ where $s \in T(\Omega, V)$ and each $r_i \in G_E(\Sigma, V)$. By all the above, this means that $t^{(\mathcal{A},\alpha)}$ is Σ-generated by $\{r_1^{(\mathcal{A},\alpha)}, \ldots, r_m^{(\mathcal{A},\alpha)}\}$, a subset of Y. □

Proposition 2.32 *Let $\mathcal{A}$ a structure free in $Mod(E)$ over the countably infinite set X. If Σ is a set of constructors for E then $\mathcal{A}^\Sigma$ is free in $Mod(E^\Sigma)$ over $Y := Is(\mathcal{A}^\Sigma)$.*

Proof. From Lemma 2.31 we know that $\mathcal{A}^\Sigma$ is generated by Y. Observing that Y is infinite for containing X by the same lemma, let φ be a Σ-atom and $\tilde{y}$ a discrete tuple of Y such that $\mathcal{A}^\Sigma \models \varphi[\tilde{y}]$. By Prop. 2.21 then, we only need to show that $E^\Sigma \models \tilde{\forall}\, \varphi$.

We know from Lemma 2.31 again that $Y = \{r^{(\mathcal{A},\alpha)} \mid r \in G_E(\Sigma, V)\}$ for any given bijection α of V onto X. Together with the assumption that $\mathcal{A}^\Sigma \models \varphi[\tilde{y}]$, this entails that we can choose an α and an instantiation ρ of $\tilde{v}$ such that $\mathcal{D}om(\rho) = \tilde{v}$, $\tilde{v}\rho \subseteq G_E(\Sigma, V)$, $\mathcal{D}om(\rho) \cap \mathcal{R}an(\rho) = \emptyset$, $(\mathcal{A}, \alpha) \models u\rho \not\equiv v\rho$ for all distinct $u, v \in \tilde{v}$, and $(\mathcal{A}, \alpha) \models \varphi\rho$. As α is an injection into X and $\mathcal{A}$ is free in $Mod(E)$ over X, we can conclude by Prop. 2.21 that $E \models \tilde{\forall}\, (\varphi\rho)$.

Now observe that $\varphi\rho$ has the form $P(\tilde{s}\rho)$ where $P \in \Sigma^{\mathrm{P}} \cup \{\equiv\}$ and $\tilde{s}$ is in $T(\Sigma, V)$. It is easy to see from the construction of ρ that $\tilde{s}\rho$ is in $ST_E(\Sigma, V)$ and $u\rho \neq_E v\rho$ for every distinct $u, v \in \mathcal{D}om(\rho)$. The claim the follows from Point 3 of Def. 2.30. □

2.2.4 Locally Free Structures

A type of structure strictly related to free structures is the following.

Definition 2.33 (Locally Free Structure) *We say that $\mathcal{A}$ is* locally free *in a class* $\mathbf{K}$ *of* Σ*-structures if every finitely-generated substructure of* $\mathcal{A}$ *is free in* $\mathbf{K}$.

We will say that a Σ-structure is *locally absolutely free* if it is locally free in the class of all the Σ-structures. By definition, any substructure of a locally free structure is itself locally free. A perhaps not so immediate property of locally free structures is the following.

Proposition 2.34 *If a locally free structure in a collapse-free class* $\mathbf{K}$ *admits a non-redundant set* X *of generators, then it is free over* X *in* $\mathbf{K}$.

Proof. Let $\mathcal{B}$ be a Σ-structure with a non-redundant set of generators X and assume that $\mathcal{B}$ is locally free in some class $\mathbf{K}$ of Σ-structures. Let $\varphi(\tilde{v})$ be an atomic Σ-formula and $\tilde{x}$ a sequence of distinct elements of X such that $\mathcal{B} \models \varphi[\tilde{x}]$. By Prop. 2.21, it is enough to show that $\mathbf{K} \models \tilde{\forall}\, \varphi$.

Let $\mathcal{A} := \langle \tilde{x} \rangle_{\mathcal{B}}$ and $I := Is(\mathcal{A}^{\Sigma})$. Notice that $\mathcal{A}$ is free in $\mathbf{K}$, for being a finitely generated substructure of a locally free structure in $\mathbf{K}$, and that, by Prop. 2.22, I is the only basis for $\mathcal{A}$. By construction of $\mathcal{A}$ and Lemma 2.15, $\tilde{x}$ is a non-redundant set of generators for $\mathcal{A}$. From what we observed earlier, I as well is a non-redundant set of generators for $\mathcal{A}$. It follows immediately, as $I \subseteq \tilde{x}$, that $\tilde{x} = I$. Now notice that $\mathcal{A} \models \varphi[\tilde{x}]$ as well because $\mathcal{A} \subseteq \mathcal{B}$ and φ is atomic. Then, by Prop. 2.21 applied to $\mathcal{A}$, we obtain that $\mathbf{K} \models \tilde{\forall}\, \varphi$. □

It is easy to see that every substructure of an absolutely free structure is absolutely free. This immediately entails that absolutely free structures are also *locally* absolutely free. The converse, however, is not true. In fact, consider the Σ-algebra $\mathcal{A}$ with domain $A := Z_1 \cup Z_2$ and signature $\Sigma := \{s\}$ where Z_1 and Z_2 are two disjoint copies of the integers and s is interpreted as the successor function on both Z_1 and Z_2. It can be shown that $\mathcal{A}$ is locally absolutely free. Now consider the subalgebra $\mathcal{B} := \langle Z_1 \cup \{x\} \rangle_{\mathcal{A}}$ where x is any element of Z_2. Observing that $\mathcal{B}$ does not admit non-redundant sets of generators, it is immediate that $\mathcal{B}$ cannot be free. By Lemma 2.34, however, we can claim the following special case.

Corollary 2.35 *For any signature* Σ*, the class of the locally absolutely free* Σ*-structures with a* non-redundant *set of generators coincides with that of the absolutely free* Σ*-structures.*

We are mainly interested in the class of locally absolutely free structures because its most important subclass, the class of locally absolutely free algebras, is ubiquitous in the field of Symbolic Computation. This class was shown to be axiomatizable by

Mal'cev (see [Mal71]). For completeness, we report its axiomatization below. Since its most prominent models are the term algebras, also known as the algebras of the finite trees, we will refer to it as the *theory of the finite trees*.

Finite Trees. The theory of the finite trees (over some signature Σ) is the universal theory $\mathcal{FT}^{\Sigma}$ given by the following axiom schemas.

- For every $f \in \Sigma^{\mathrm{F}}$,

$$\tilde{\forall}\, (f(\tilde{u}) \equiv f(\tilde{v}) \rightarrow \tilde{u} \equiv \tilde{v}).$$

- For every $f, g \in \Sigma^{\mathrm{F}}$, $f \neq g$,

$$\tilde{\forall}\, f(\tilde{u}) \not\equiv g(\tilde{v}).$$

- For every $t(\tilde{v}) \in T(\Sigma, V) \setminus V$ and $v \in \tilde{v}$,

$$\tilde{\forall}\, v \not\equiv t(\tilde{v}).$$

Proofs of the following characterization can be found in [Mal71, Mah88], among others.

Proposition 2.36 *A Σ-algebra $\mathcal{A}$ is locally absolutely free iff $\mathcal{A} \in Mod(\mathcal{FT}^{\Sigma})$.*

3 Combining Constraint Domains

We are mainly concerned with the question of how to solve constraint satisfiability problems with respect to several constraint theories by combining in a modular fashion the satisfiability procedures available for the single theories. We will tackle this question at the domain level first and then extend our approach to the theory level in the next section. To start with, we must be able to recast a given satisfiability problem as a *combined satisfiability problem*. That is, we must be able to, first, describe the solution structure as a proper combination of two or more distinct *component* structures; second, decompose the problem into a number of "pure" subproblems, each solvable over a component structure; third, combine the subproblem solutions, each ranging over one of the component structures, into a solution for the original problem, ranging over the combined structure.

We begin by providing a viable definition of combined structure. For simplicity, here and in the rest of the paper, we will mostly consider combinations of just two components.

Definition 3.1 (Model Fusion) *Given two structures $\mathcal{A}$ and $\mathcal{B}$, we say that the $(\Sigma_A \cup \Sigma_B)$-structure $\mathcal{F}$ is a* fusion[16] *of $\mathcal{A}$ and $\mathcal{B}$ if there exist a map h_{A-F} and a map h_{B-F} such that*

$$h_{A-F} : \mathcal{A} \cong \mathcal{F}^{\Sigma_A} \quad \textit{and} \quad h_{B-F} : \mathcal{B} \cong \mathcal{F}^{\Sigma_B}.$$

We will sometimes use the notation $\langle \mathcal{F}, h_{A-F}, h_{B-F} \rangle$ to indicate the fusion structure and the relative isomorphisms. Essentially, a fusion of two structures $\mathcal{A}$ and $\mathcal{B}$, when it exists, is a structure that, if seen as a Σ_A-structure, is identical to $\mathcal{A}$, and, if seen as a Σ_B-structure, is identical to $\mathcal{B}$.

We will denote by $Fus(\mathcal{A}, \mathcal{B})$ the set of all the fusions of structures $\mathcal{A}$ and $\mathcal{B}$. By the above definition, it is immediate that $Fus(\mathcal{A}, \mathcal{B}) = Fus(\mathcal{B}, \mathcal{A})$ and that $Fus(\mathcal{A}, \mathcal{B})$ is an abstract class, that is, it is closed under isomorphism. Furthermore, it is not difficult to show, although rather tedious, that $Fus(\mathcal{A}, \mathcal{B})$ may contain non-isomorphic structures. Intuitively, however, all the elements of $Fus(\mathcal{A}, \mathcal{B})$ should be pairwise isomorphic at least over the symbols that $\mathcal{A}$ and $\mathcal{B}$ share. Such intuition is indirectly confirmed by the the proposition below, establishing a necessary and sufficient condition for the existence of fusions.

Proposition 3.2 *Given structures $\mathcal{A}$ and $\mathcal{B}$,*

$$Fus(\mathcal{A}, \mathcal{B}) \neq \emptyset \quad \textit{iff} \quad \mathcal{A}^{\Sigma_A \cap \Sigma_B} \cong \mathcal{B}^{\Sigma_A \cap \Sigma_B}.$$

Proof: Let $\Sigma := \Sigma_A \cap \Sigma_B$.

($\Rightarrow$) Let $\mathcal{C} \in Fus(\mathcal{A}, \mathcal{B})$. By definition we have that $\mathcal{A} \cong \mathcal{C}^{\Sigma_A}$ and $\mathcal{B} \cong \mathcal{C}^{\Sigma_B}$. From the fact that $\Sigma \subseteq \Sigma_A$ and $\Sigma \subseteq \Sigma_B$ it follows immediately that

$$\mathcal{A}^{\Sigma} \cong \mathcal{C}^{\Sigma} \quad \text{and} \quad \mathcal{B}^{\Sigma} \cong \mathcal{C}^{\Sigma},$$

which implies that $\mathcal{A}^{\Sigma} \cong \mathcal{B}^{\Sigma}$.

($\Leftarrow$) Let h be a map such that $h \colon \mathcal{A}^{\Sigma} \cong \mathcal{B}^{\Sigma}$. Consider a $(\Sigma_A \cup \Sigma_B)$-structure $\mathcal{C}$ with universe B and such that

for all $p \in (\Sigma_A \cup \Sigma_B)^{\mathrm{P}}$,

$$p^{\mathcal{C}} := \begin{cases} h(p^{\mathcal{A}}) & \text{if} \quad p \in (\Sigma_A \setminus \Sigma_B) \\ p^{\mathcal{B}} & \text{if} \quad p \in \Sigma_B \end{cases}$$

for all $g \in (\Sigma_A \cup \Sigma_B)^{\mathrm{F}}$ of arity n and $\tilde{b} \in B^n$,

[16] We initially chose the term "fusion" to avoid overloading the term "amalgamation", which has a more specific meaning in the Model Theory literature. We have later discovered that [PT97] does use "amalgamation" for the same type of combined structure as ours and that [Hol95] uses "fusion" for a rather different type of combined structure.

$$g^{\mathcal{C}}(\tilde{b}) := \begin{cases} h(g^{\mathcal{A}}(h^{-1}(\tilde{b}))) & \text{if} \quad g \in (\Sigma_A \setminus \Sigma_B) \\ g^{\mathcal{B}}(\tilde{b}) & \text{if} \quad g \in \Sigma_B \end{cases}$$

In practice, $\mathcal{C}$ interprets Σ_B-symbols as $\mathcal{B}$ does and Σ_A-symbols as images, through h, of the corresponding function/relations in $\mathcal{A}$. We prove below that $h\colon \mathcal{A} \cong \mathcal{C}^{\Sigma_A}$.

If P is an n-ary predicate symbol of $\Sigma_A \setminus \Sigma_B$, for each $\tilde{a} \in A^n$,

$$\begin{array}{lll} \tilde{a} \in P^{\mathcal{A}} & \text{iff} \quad h(\tilde{a}) \in h(P^{\mathcal{A}}) & \text{(by def. of } h(P^{\mathcal{A}})) \\ & \text{iff} \quad h(\tilde{a}) \in P^{\mathcal{C}} & \text{(by 3.1);} \end{array}$$

if P is an n-ary predicate symbol of Σ, for each $\tilde{a} \in A^n$,

$$\begin{array}{lll} \tilde{a} \in P^{\mathcal{A}} & \text{iff} \quad h(\tilde{a}) \in P^{\mathcal{B}} & (h\colon \mathcal{A}^{\Sigma} \cong \mathcal{B}^{\Sigma}) \\ & \text{iff} \quad h(\tilde{a}) \in P^{\mathcal{C}} & \text{(by 3.1);} \end{array}$$

if g is an n-ary function symbol of $\Sigma_A \setminus \Sigma_B$, for each $\tilde{a} \in A^n$,

$$\begin{array}{rcll} h(g^{\mathcal{A}}(\tilde{a})) & = & h(g^{\mathcal{A}}(h^{-1}(h(\tilde{a})))) & \text{(by def. of inverse)} \\ & = & g^{\mathcal{C}}(h(\tilde{a})) & \text{(by 3.1);} \end{array}$$

if g is an n-ary function symbol of Σ, for each $\tilde{a} \in A^n$,

$$\begin{array}{rcll} h(g^{\mathcal{A}}(\tilde{a})) & = & g^{\mathcal{B}}(h(\tilde{a})) & (h\colon \mathcal{A}^{\Sigma} \cong \mathcal{B}^{\Sigma}) \\ & = & g^{\mathcal{C}}(h(\tilde{a})) & \text{(by 3.1).} \end{array}$$

By 3.1, it is immediate that $id\colon \mathcal{B} \cong \mathcal{C}^{\Sigma_B}$, where id is the identity of B. It follows from the definition of fusion that $\langle \mathcal{C}, h, id \rangle$ is a fusion of $\mathcal{A}$ and $\mathcal{B}$. □

In essence, two structures admit a fusion exactly when they have the same cardinality and interpret in the same way the symbols shared by their signatures.

We know that for each structure there is at least one set of individuals, the set of generators, which determines the structure univocally. For pairs of structures admitting fusions it is sometimes possible to identify a pair of sets of individuals that, in a sense, determines the possible fusions between the two structures.

Definition 3.3 (Fusible Structures) *Consider two structures $\mathcal{A}$ and $\mathcal{B}$, a set $X \subseteq A$, and a set $Y \subseteq B$ with X's same cardinality. We say that $\mathcal{A}$ is* fusible with $\mathcal{B}$ over $\langle X, Y \rangle$ *if every injection from a finite subset of X into Y can be extended to an isomorphism of $\mathcal{A}^{\Sigma_A \cap \Sigma_B}$ onto $\mathcal{B}^{\Sigma_A \cap \Sigma_B}$.*

Since $\mathcal{A}$ is fusible with $\mathcal{B}$ over $\langle X, Y \rangle$ then $\mathcal{B}$ is fusible with $\mathcal{A}$ over $\langle Y, X \rangle$, in general, we will simply say that $\mathcal{A}$ and $\mathcal{B}$ are *fusible over* $\langle X, Y \rangle$. In analogy with generators, we call *fusors* the elements of X and those of Y.

One may wonder whether there are *practical* sufficient conditions for "fusibility". We will mainly appeal to those suggested by the following result.

Proposition 3.4 *Let $\mathcal{A}$ and $\mathcal{B}$ be two structures and $\Sigma := \Sigma_A \cap \Sigma_B$. Assume that $\mathcal{A}^\Sigma$ is free over X and $\mathcal{B}^\Sigma$ is free over Y in the same class of Σ-structures. If $Card(X) = Card(Y)$, then $\mathcal{A}$ and $\mathcal{B}$ are fusible over $\langle X, Y\rangle$.*

Proof. Given a finite set $X_0 \in X$, consider any injective map $h : X_0 \to Y$. Since X_0 is finite and $Card(X) = Card(Y)$, h can always be extended to a bijection from X onto Y. By Lemma 2.25, h can be extended to an isomorphism of $\mathcal{A}^\Sigma$ onto $\mathcal{B}^\Sigma$. □

We will use this result in Sec. 6 to define a broad class of theories to which our combination method applies.

Having elected fusions as our combination structures, we now show in what sense it is possible to go from satisfiability in a structure $\mathcal{A}$ and in a structure $\mathcal{B}$ to satisfiability in a fusion of theirs. We will start with the simplest type of combined satisfiability problem: given a formula φ satisfiable in a structure $\mathcal{A}$ and a formula ψ satisfiable in a structure $\mathcal{B}$, what can we say about the satisfiability of their conjunction?

The question is perhaps more interesting if $\mathcal{A}$ and $\mathcal{B}$ do not have the same signature. Although we will never make such an assumption, we will refer to formulas like φ and ψ above as *pure* formulas, that is, formulas expressed strictly in the language of the component structures (or later, theories).

Lemma 3.5 *Let $\mathcal{A}$ and $\mathcal{B}$ be two structures of respective signatures Ω and Δ such that $\mathcal{A}$ and $\mathcal{B}$ are fusible over some pair $\langle X, Y\rangle$. Let $\varphi(\tilde{u};\tilde{v})$ an Ω-formula and $\psi(\tilde{w};\tilde{v})$ be a Δ-formula such that $\tilde{u} \cap \tilde{w} = \emptyset$. If φ is satisfiable in $\mathcal{A}$ with $\tilde{v}$ taking distinct values over X and ψ is satisfiable in $\mathcal{B}$ with $\tilde{v}$ taking distinct values over Y, then $\varphi \wedge \psi$ is satisfiable in a fusion of $\mathcal{A}$ and $\mathcal{B}$.*

Proof. Let $\Sigma := \Omega \cap \Delta$ and $\tilde{v} := (v_1, \ldots, v_m)$. Assume that

$$\mathcal{A} \models \varphi[\tilde{a};\tilde{x}] \quad \text{and} \quad \mathcal{B} \models \psi[\tilde{b};\tilde{y}]$$

where $\tilde{x} := (x_1, \ldots, x_m)$ is in X, $\tilde{y} := (y_1, \ldots, y_m)$ is in Y, and neither of them contains repetitions. Consider the map $h\colon \tilde{x} \to Y$ such that,

$$h(x_j) = y_j \quad \text{for all} \quad j \in \{1, \ldots, m\}.$$

By construction of $\tilde{x}$ and $\tilde{y}$, h is injective. Since $\mathcal{A}$ is fusible with $\mathcal{B}$ over $\langle X, Y\rangle$, h can be extended to an isomorphism h_{A-B} of $\mathcal{A}^\Sigma$ onto $\mathcal{B}^\Sigma$. Now, where $K := \{k_1, \ldots, k_m\}$ is a set of constant symbols not appearing in $\Omega \cup \Delta$, we define $\mathcal{A}^{\Omega\cup K}$ as the expansion of $\mathcal{A}$ to $\Omega \cup K$ and $\mathcal{B}^{\Delta\cup K}$ as the expansion of $\mathcal{B}$ to $\Delta \cup K$ such that, for every $j \in \{1, \ldots, m\}$,

$$k_i^{\mathcal{A}^{\Omega\cup K}} = x_i \quad \text{and} \quad k_i^{\mathcal{B}^{\Delta\cup K}} = y_i. \tag{5}$$

It is not difficult to see that h_{A-B} is an isomorphism of $\mathcal{A}^{\Sigma\cup K}$ onto $\mathcal{B}^{\Sigma\cup K}$ as well. By Prop. 3.2, it follows that $Fus(\mathcal{A}^{\Omega\cup K}, \mathcal{B}^{\Delta\cup K})$ is not empty. Consider any $\mathcal{F} \in Fus(\mathcal{A}^{\Omega\cup K}, \mathcal{B}^{\Delta\cup K})$. We show that $\varphi_1 \wedge \varphi_2$ is satisfiable in $\mathcal{F}^{\Omega\cup\Delta}$. The claim will then follow from the easily proven fact that $\mathcal{F}^{\Omega\cup\Delta} \in Fus(\mathcal{A}, \mathcal{B})$.

Consider the instantiation $\sigma := \{v_1 \leftarrow k_1, \ldots, v_m \leftarrow k_m\}$. By assumption, $\mathcal{A} \models \varphi[\tilde{a}; \tilde{x}]$ and so, by construction of $\mathcal{A}^{\Omega\cup K}$ and σ, $\mathcal{A}^{\Omega\cup K} \models \tilde{\exists}\,(\varphi\sigma)$. From the fact that $\mathcal{F}^{\Omega\cup K} \cong \mathcal{A}^{\Omega\cup K}$ it follows that $\mathcal{F} \models \tilde{\exists}\,(\varphi\sigma)$. Similarly, we can show that $\mathcal{F} \models \tilde{\exists}\,(\psi\sigma)$. By elementary logical reasoning and the fact that $\mathcal{V}ar(\varphi\sigma) \cap \mathcal{V}ar(\psi\sigma) = \emptyset$, it follows that $\mathcal{F} \models \tilde{\exists}(\varphi\sigma \wedge \psi\sigma)$ and therefore that $\mathcal{F} \models \tilde{\exists}(\varphi \wedge \psi)$, which implies, by Lemma 2.6, that $\mathcal{F}^{\Omega\cup\Delta} \models \tilde{\exists}\,(\varphi \wedge \psi)$. □

Since component structures with the same cardinality but no symbols in common are always fusable, we immediately obtain the following special case.

Lemma 3.6 *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two signature-disjoint structures with same cardinality and, for $i = 1, 2$, consider the Σ_{A_i}-formula $\varphi_i(\tilde{u}_i; \tilde{v})$, where $\tilde{u}_1 \cap \tilde{u}_2 = \emptyset$. If $\varphi_i \wedge \varepsilon_{\neq}(\tilde{v})$ is satisfiable in $\mathcal{A}_i$, for $i = 1, 2$, then $\varphi_1 \wedge \varphi_2$ is satisfiable in a fusion of $\mathcal{A}_1$ and $\mathcal{A}_2$.*

Proof. For $i = 1, 2$, let α_i be a valuation such that $(\mathcal{A}_i, \alpha_i) \models \varphi_i \wedge \varepsilon_{\neq}(\tilde{v})$. Observe that, because of $\varepsilon_{\neq}(\tilde{v})$, α_i assigns pairwise distinct individuals to the shared variables of φ_i. The result follows then from Lemma. 3.5 noting that two equinumerous structures $\mathcal{A}$ and $\mathcal{B}$ are trivially fusible over $\langle A,\ B\rangle$ when their signatures are disjoint. □

This last result can be interpreted in constraint solving terms as follows. Each φ_i represents a problem in the variables $\tilde{u}_i \cup \tilde{v}$ over the domain modeled by $\mathcal{A}_i$, while $\varphi := \varphi_1 \wedge \varphi_2$ represents a (composite) problem in the variables $\tilde{u}_1 \cup \tilde{u}_2 \cup \tilde{v}$ over the domain modeled by some fusion of $\mathcal{A}_1$ and $\mathcal{A}_2$. In order to *merge* a solution s_1 of φ_1 and a solution s_2 of φ_2 into a solution of φ it is necessary that s_1 and s_2 agree, so to speak, on the values that they assign to the shared variables, if any. The role of $\varepsilon_{\neq}(\tilde{v})$ is exactly that of assuring such merging by requiring that the shared variables take distinct values over the fusors of $\mathcal{A}_1$ and $\mathcal{A}_2$.

Now, what if either φ_i is satisfiable only with valuations that assign the same value to some of the shared variables? For instance, what if $\mathcal{A}_1 \models \varphi_1 \rightarrow (v_i \equiv v_j)$ for some $v_i, v_j \in \tilde{v}$? It should be clear that, if all the $\mathcal{A}_1$-solutions of φ_1 identify some variables in $\tilde{v}$, for $\varphi_1 \wedge \varphi_2$ to be satisfiable in a fusion of $\mathcal{A}_1$ and $\mathcal{A}_2$[17] there must exist an $\mathcal{A}_2$-solution of φ_2 that also identifies those variables. We can then generalize Lemma. 3.6 to encompass the case just illustrated by considering a formula of the form $\varphi_i\xi$, where $\xi \in \mathrm{ID}(\tilde{v})$. More precisely, a formula obtained from φ_i by a syntactical identification of those shared variables that will be (semantically)

[17] That is, for subproblems solutions to be *mergeable* into solutions of the composite problem.

identified by the $\mathcal{A}_i$-solutions. Then, the constraint $\xi_{\neq}$, which is nothing else but $\varepsilon_{\neq}(\tilde{v}\xi)$, can be used in the same way $\varepsilon_{\neq}(\tilde{v})$ was used before.

Proposition 3.7 *For $i = 1, 2$, let $\mathcal{A}_i$ and φ_i be as in Lemma. 3.6. If*

$$\varphi_i\xi \wedge \xi_{\neq}$$

is satisfiable in $\mathcal{A}_i$ for some $\xi \in \mathrm{ID}(\tilde{v})$, then $\varphi_1 \wedge \varphi_2$ is satisfiable in a fusion of $\mathcal{A}_1$ and $\mathcal{A}_2$.

The proposition above is the (generalized) syntactic counterpart of Lemma 3.5 in the case of signature-disjoint structures. The addition of a simple constraint guarantees that the (new) shared variables take distinct values over the fusors of the component structures as Lemma. 3.5 requires. Since equinumerous structures with disjoint signatures are fusible over their whole carriers, the task here was essentially trivial. When two fusible structures are not signature-disjoint, however, they are likely to be fusible only over proper subsets of their carriers. In that case, it is generally impossible to force the shared variables to range over the two sets of fusors by the simple addition of a first-order constraint like $\xi_{\neq}$.[18] One case in which this is possible is when the fusors in question are also Σ-isolated, for some finite set Σ of symbols shared by the two structures signatures.

Lemma 3.8 *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two structures and let Σ be a finite subset of $\Sigma_{A_1} \cap \Sigma_{A_2}$. Assume that for $i = 1, 2$, there is a set X_i such that $Is(\mathcal{A}_i{}^{\Sigma}) \subseteq X_i \subseteq A_i$ and $\mathcal{A}_1$ and $\mathcal{A}_2$ are fusible over $\langle X_1, X_2\rangle$. For $i = 1, 2$, consider the Σ_{A_i}-formula $\varphi_i(\tilde{u}_i; \tilde{v})$, where $\tilde{u}_1 \cap \tilde{u}_2 = \emptyset$. If the Σ-restricted formula*

$$\varphi_i \wedge \varepsilon_{\neq}(\tilde{v}) \wedge \varepsilon_{\neq}^{\Sigma}(\tilde{v})$$

is satisfiable in $\mathcal{A}_i$ for $i = 1, 2$, then $\varphi_1 \wedge \varphi_2$ is satisfiable in a fusion of $\mathcal{A}_1$ and $\mathcal{A}_2$.

Proof. By assumption, for $i = 1, 2$, there is a sequence $\tilde{a}_i$ and a sequence $\tilde{x}_i$ of individuals of A_i such that $\mathcal{A}_i \models \varphi_i[\tilde{a}_i; \tilde{x}_i] \wedge \varepsilon_{\neq}[\tilde{x}_i] \wedge \varepsilon_{\neq}^{\Sigma}[\tilde{x}_i]$. By Lemma. 3.5, all we need to show is that $\tilde{x}_i$ is composed of pairwise distinct elements of X_i.

That $\tilde{x}_i$ does not contain repetitions is entailed by the fact that $\varepsilon_{\neq}[\tilde{x}_i]$ is true in $\mathcal{A}_i$. To see that $\tilde{x}_i$ is included in X_i, just recall that $\varepsilon_{\neq}^{\Sigma}[\tilde{x}_i]$ is true exactly when $\tilde{x}_i$ is a set of Σ-isolated individuals and that all Σ-isolated individuals of A_i are in X_i by assumption. □

In the lemma above, the requirement that both sets of fusors contain the Σ-isolated individuals of their respective structures, allows us to use a first-order formula, $\varepsilon_{\neq}(\tilde{v}) \wedge \varepsilon_{\neq}^{\Sigma}(\tilde{v})$, to force the variables shared by the two pure formulas to take distinct values over the fusors. But now, what if either φ_i is satisfiable only with

[18] Put another way, the property of being a fusor is not first-order definable.

valuations that map some shared variables to individuals that are not Σ-isolated? Well, we can still apply the above result if these individuals are Σ-generated by Σ-isolated elements. We do this by first instantiating each shared variable in question with a suitable Σ-term over fresh variables and then forcing both the new variables and the untouched shared variables to range over the Σ-isolated individuals, as we have done before.

First however, let us introduce for conciseness the following restricted notion of fusibility.

Definition 3.9 (Σ-fusibility) *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two structures and Σ be a* finite *subset of $\Sigma_{A_1} \cap \Sigma_{A_2}$. We say that $\mathcal{A}_1$ and $\mathcal{A}_2$ are* Σ-fusible *if for $i = 1, 2$ there is a set X_i such that $Is({\mathcal{A}_i}^{\Sigma}) \subseteq X_i \subseteq A_i$ and $\mathcal{A}_1$ and $\mathcal{A}_2$ are fusible over $\langle X_1, X_2 \rangle$.*

A little observation on the above definition is in order here. Recalling the definition of fusibility, it is not difficult to see that when two structures $\mathcal{A}_1$ and $\mathcal{A}_2$ as above are fusible over some pair $\langle X_1, X_2 \rangle$, every bijection between two finite subsets of X_i extends to an automorphism of ${\mathcal{A}_i}^{\Sigma}$ $(i = 1, 2)$. This entails, in particular, that all the elements of X_i satisfy exactly the same Σ-formulas in one variable. As a consequence, we obtain that a member of X_i is Σ-isolated in $\mathcal{A}_i$ only if *every* member of X_i is Σ-isolated in $\mathcal{A}_i$. Therefore, unless $Is({\mathcal{A}_1}^{\Sigma})$ and $Is({\mathcal{A}_2}^{\Sigma})$ are empty, if $\mathcal{A}_1$ and $\mathcal{A}_2$ are Σ-fusible, the pair of sets on which they are fusible is univocally determined and coincides with $\langle Is({\mathcal{A}_1}^{\Sigma}), Is({\mathcal{A}_2}^{\Sigma}) \rangle$.

Proposition 3.10 *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two structures Σ-fusible for some finite $\Sigma \subseteq \Sigma_{A_1} \cap \Sigma_{A_2}$. For $i = 1, 2$, consider the Σ_{A_i}-formula $\varphi_i(\tilde{u}_i; \tilde{v})$, where $\tilde{u}_1 \cap \tilde{u}_2 = \emptyset$. If*

$$(\varphi_i \rho \wedge \rho_{\neq})\xi \wedge \xi_{\neq}$$

is satisfiable in $\mathcal{A}_i$ for some $\rho \in \mathrm{IN}^{\Sigma}(\tilde{v})$ and $\xi \in \mathrm{ID}(\mathcal{VR}an(\rho))$, then $\varphi_1 \wedge \varphi_2$ is satisfiable in a fusion of $\mathcal{A}_1$ and $\mathcal{A}_2$.

Proof. For $i = 1, 2$, assume that $(\varphi_i \rho \wedge \rho_{\neq})\xi \wedge \xi_{\neq}$ is satisfiable in $\mathcal{A}_i$, where ρ and ξ are as described above. Let $\varphi'_i := \varphi_i \rho \xi$ and $\tilde{w} := \mathcal{V}ar(\varphi'_1) \cap \mathcal{V}ar(\varphi'_2)$ and observe that $\rho_{\neq}\xi = \varepsilon^{\Sigma}_{\neq}(\tilde{w})$, $\xi_{\neq} = \varepsilon_{\neq}(\tilde{w})$, and hence $(\varphi_i \rho \wedge \rho_{\neq})\xi \wedge \xi_{\neq}$ has the form

$$\varphi'_i(\tilde{u}_i, \tilde{w}) \wedge \varepsilon_{\neq}(\tilde{w}) \wedge \varepsilon^{\Sigma}_{\neq}(\tilde{w}).$$

From the assumptions and Lemma. 3.8 we have that $\varphi'_1 \wedge \varphi'_2$ is satisfiable in a fusion of $\mathcal{A}_1$ and $\mathcal{A}_2$. The claim follows then immediately from the observation that $(\varphi'_1 \wedge \varphi'_2) = (\varphi_1 \wedge \varphi_2)\rho\xi$. $\square$

The result above is not as general as we would like. The satisfiability of $(\varphi_i \rho \wedge \rho_{\neq})\xi \wedge \xi_{\neq}$ in $\mathcal{A}_i$, although sufficient, is typically not necessary for the satisfiability of $\varphi_1 \wedge \varphi_2$ in some fusion of $\mathcal{A}_1$ and $\mathcal{A}_2$. It does become necessary, however, if the fusion in question is Σ-generated by its Σ-isolated individuals alone.

Proposition 3.11 *Assume that two structures $\mathcal{A}_1$ and $\mathcal{A}_2$ admit a fusion $\mathcal{F}$ which is Σ-generated by its Σ-isolated individuals, for some finite $\Sigma \subseteq \Sigma_{A_1} \cap \Sigma_{A_2}$. For $i = 1, 2$, consider the Σ_{A_i}-formula $\varphi_i(\tilde{u}_i; \tilde{v})$, with $\tilde{u}_1 \cap \tilde{u}_2 = \emptyset$. Then, if $\varphi_1 \wedge \varphi_2$ is satisfiable in $\mathcal{F}$, there is a $\rho \in \mathrm{IN}^\Sigma(\tilde{v})$ and a $\xi \in \mathrm{ID}(\mathcal{VR}an(\rho))$ such that $(\varphi_i\rho \wedge \rho_{\neq})\xi \wedge \xi_{\neq}$ is satisfiable in $\mathcal{A}_i$ for $i = 1, 2$.*

Proof. Let X be the set of $\mathcal{F}$'s Σ-isolated individuals. By assumption, there is a valuation α such that $(\mathcal{F}, \alpha) \models \varphi_1 \wedge \varphi_2$. We show that α and X induce an instantiation ρ and identification ξ that satisfy the claim.

For all $v_j \in \tilde{v}$, such that $\alpha(v_j) \notin X$, we choose any non-variable Σ-term $t_j(\tilde{w}_j)$ and sequence $\tilde{x}_j$ in X such that $\alpha(v_j) = t_j^{\mathcal{A}}[\tilde{x}_j]$.[19] We assume, with no loss of generality, that all the variables in each $\tilde{w}_j$ are new and expand α to these variables by mapping each of them to the corresponding element of $\tilde{x}_j$. Then, we choose the instantiation $\rho \in \mathrm{IN}^\Sigma(\tilde{v})$ such that, for all $v_j \in \tilde{v}$,

$$v_j\rho = \begin{cases} v_j & \text{if } \alpha(v_j) \in X \\ t_j(\tilde{w}_j) & \text{otherwise} \end{cases}$$

and the identification $\xi \in \mathrm{ID}(\tilde{v}\rho)$ such that, for all $v, w \in \tilde{v}\rho$,

$$v\xi = w\xi \quad \text{iff} \quad \alpha'(v) = \alpha'(w),$$

where α' is the expansion of α just described. We leave to the reader to verify that $(\mathcal{F}, \alpha') \models (\varphi_i\rho \wedge \rho_{\neq})\xi \wedge \xi_{\neq}$, for $i = 1, 2$. Now, $(\varphi_i\rho \wedge \rho_{\neq})\xi \wedge \xi_{\neq}$ is actually a Σ_{A_i}-formula and so is also satisfied by $\mathcal{F}^{\Sigma_{A_i}}$. The claim then follows from the fact that $\mathcal{F}^{\Sigma_i}$ is isomorphic to $\mathcal{A}_i$ by definition of fusion. $\square$

4 Fusions and Unions of Theories

Our interest in fusions is motivated by their close link to unions of theories, as illustrated in the following proposition. First, if $\mathcal{T}_1$ and $\mathcal{T}_2$ are two theories, let $Fus(\mathcal{T}_1, \mathcal{T}_2)$ denote the following class of structures:

$$\bigcup_{\mathcal{A} \in Mod(\mathcal{T}_1),\, \mathcal{B} \in Mod(\mathcal{T}_2)} Fus(\mathcal{A}, \mathcal{B}).$$

Proposition 4.1 *For any two theories $\mathcal{T}_1$ and $\mathcal{T}_2$,*

$$Fus(\mathcal{T}_1, \mathcal{T}_2) \quad = \quad Mod(\mathcal{T}_1 \cup \mathcal{T}_2).$$

[19] The existence of such a term and sequence is guaranteed by the assumption that X Σ-generates $\mathcal{F}$.

Proof: For $i = 1, 2$, let $\Sigma_i := \Sigma_{\mathcal{T}_i}$.

($\subseteq$) Assume that $\mathcal{F}$ is a fusion of some $\mathcal{A} \in Mod(\mathcal{T}_1)$ and $\mathcal{B} \in Mod(\mathcal{T}_2)$. From the definition of fusion we have that, $\mathcal{A} \cong \mathcal{F}^{\Sigma_1}$ and $\mathcal{B} \cong \mathcal{F}^{\Sigma_2}$. Therefore, $\mathcal{F}$ models every sentence of $\mathcal{T}_1$ and every sentence of $\mathcal{T}_2$. It follows immediately that $\mathcal{F}$ models $\mathcal{T}_1 \cup \mathcal{T}_2$.

($\supseteq$) Immediate consequence of the obvious fact that any $\mathcal{C} \in Mod(\mathcal{T}_1 \cup \mathcal{T}_2)$ is a fusion of $\mathcal{C}^{\Sigma_1}$ and $\mathcal{C}^{\Sigma_2}$ and that $\mathcal{C}^{\Sigma_i}$ models $\mathcal{T}_i$, for $i = 1, 2$. □

Recalling Prop. 3.2 on the existence of fusions we have the following corollary, first proved in [Rin96b] and [TH96].[20]

Corollary 4.2 *The union of a theory $\mathcal{T}_1$ and a theory $\mathcal{T}_2$ is consistent iff there is a model of $\mathcal{T}_1$ and a model of $\mathcal{T}_2$ such that their reducts to $\Sigma_{\mathcal{T}_1} \cap \Sigma_{\mathcal{T}_2}$ are isomorphic.*

We will later see that all the theories we consider for combination will satisfy the right-hand-side condition in the above corollary, therefore it will indeed make sense to work on their union.

Definition 4.3 *Where $\mathcal{L}$ is a class of formulas and Σ_1 and Σ_2 two signatures, we will call* disjoint product *of $\mathcal{L}^{\Sigma_1}$ and $\mathcal{L}^{\Sigma_2}$ and denote with $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ the following subset of $\mathcal{L}^{\Sigma_1} \times \mathcal{L}^{\Sigma_2}$,*

$$\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2} \quad := \quad ((\mathcal{L}^{\Sigma_1} \setminus \mathcal{L}^{\Sigma_2}) \times \mathcal{L}^{\Sigma_2}) \quad \cup \quad (\mathcal{L}^{\Sigma_1} \times (\mathcal{L}^{\Sigma_2} \setminus \mathcal{L}^{\Sigma_1})).$$

In practice, $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ coincides with $\mathcal{L}^{\Sigma_1} \times \mathcal{L}^{\Sigma_2}$ minus all the pairs $\langle \varphi, \psi \rangle$ such that both φ and ψ belong to $\mathcal{L}^{\Sigma_1} \cap \mathcal{L}^{\Sigma_2}$. Observe that if $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$, φ_i contains predicate and function symbols from Σ_i only ($i = 1, 2$). For this reason, we will call φ_i the *i-pure* component of $\langle \varphi_1, \varphi_2 \rangle$. For convenience, we will say that $\langle \varphi_1, \varphi_2 \rangle$ is satisfiable in a structure (theory) if $\varphi_1 \wedge \varphi_2$ is satisfiable in the structure (theory).

We now have all we need to define a class of theories for which our combination method is provably correct, as we will see in the next section.

Definition 4.4 (N-O-combinable Theories) *Let $\mathcal{L}$ be a class of formulas and $\mathcal{T}_1, \mathcal{T}_2$ two theories with respective signatures Σ_1, Σ_2 such that $\Sigma := \Sigma_1 \cap \Sigma_2$ is finite.*

- *We say that $\mathcal{T}_1$ and $\mathcal{T}_2$ are* partially N-O-combinable over $\mathcal{L}$ *if Cond. 4.1 below holds for all $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$.*

- *We say that $\mathcal{T}_1$ and $\mathcal{T}_2$ are* (totally) N-O-combinable over $\mathcal{L}$ *if both Cond. 4.1 and Cond. 4.2 below hold for all $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$.*

[20] See [TH96] for a brief discussion of the significance of such result.

Condition 4.1 *For all* $\rho \in \mathrm{IN}^{\Sigma}(\tilde{v})$ *and* $\xi \in \mathrm{ID}(\mathcal{VR}an(\rho))$ *with* $\tilde{v} := \mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2)$*, if*

$$\psi_i \;:=\; (\varphi_i\rho \wedge \rho_{\neq})\xi \wedge \xi_{\neq}$$

is satisfiable in $\mathcal{T}_i$ *for* $i = 1, 2$*, then* ψ_i *is satisfiable in a model* $\mathcal{A}_i$ *of* $\mathcal{T}_i$ *such that* $\mathcal{A}_1$ *and* $\mathcal{A}_2$ *are* Σ*-fusible.*

Condition 4.2 *If* $\varphi_1 \wedge \varphi_2$ *is satisfiable in* $\mathcal{T}_1 \cup \mathcal{T}_2$*, it is satisfiable in a model of* $\mathcal{T}_1 \cup \mathcal{T}_2$ *that is* Σ*-generated by its* Σ*-isolated individuals.*

The use of $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ in the definition above instead of $\mathcal{L}^{\Sigma_1} \times \mathcal{L}^{\Sigma_2}$ is a necessary technicality to guarantee the existence of pairs of N-O-combinable theories at all. As an example of what can go wrong with $\mathcal{L}^{\Sigma_1} \times \mathcal{L}^{\Sigma_2}$, assume that $\mathcal{L}$ is closed under conjunction and negation and take any two theories $\mathcal{T}_1$ and $\mathcal{T}_2$ of signature Σ_1 and Σ_2, respectively, with $\Sigma := \Sigma_1 \cap \Sigma_2$ non-empty. Then, $\langle \varphi, \neg\varphi \rangle \in \mathcal{L}^{\Sigma_1} \times \mathcal{L}^{\Sigma_2}$ for any $\varphi \in \mathcal{L}^{\Sigma}$, but it is obvious that, against the requirements of Cond. 4.1, no model of $\mathcal{T}_1$ satisfying φ is fusible with a model of $\mathcal{T}_2$ satisfying $\neg\varphi$.[21]

N-O-combinable theories satisfy the properties below which we will use in Sect. 5 to produce a sound and complete combination method. In the following, let $\mathcal{T}_1$, $\mathcal{T}_2$, Σ_1, Σ_2, Σ, and $\mathcal{L}$ be as in Def. 4.4.

Proposition 4.5 *Assume that* $\mathcal{T}_1$ *and* $\mathcal{T}_2$ *are* partially *N-O-combinable over* $\mathcal{L}$*. Then, for all* $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ *and* $\tilde{v} = \mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2)$*,* $\varphi_1 \wedge \varphi_2$ *is satisfiable in* $\mathcal{T}_1 \cup \mathcal{T}_2$ *if there is a* $\rho \in \mathrm{IN}^{\Sigma}(\tilde{v})$ *and* $\xi \in \mathrm{ID}(\mathcal{VR}an(\rho))$ *such that* $(\varphi_i\rho \wedge \rho_{\neq})\xi \wedge \xi_{\neq}$ *is satisfiable in* $\mathcal{T}_i$ *for* $i = 1, 2$*.*

Proof. Immediate consequence of Cond. 4.1, Prop. 3.10, and Prop. 4.1. □

If $\mathcal{T}_1$ and $\mathcal{T}_2$ satisfy Cond. 4.2 as well, the implication in the proposition above becomes a full equivalence.

Theorem 4.6 *When* $\mathcal{T}_1$ *and* $\mathcal{T}_2$ *are totally N-O-combinable over* $\mathcal{L}$ *the following statements are equivalent for all* $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ *and* $\tilde{v} = \mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2)$*.*

1. *There exists a* $\rho \in \mathrm{IN}^{\Sigma}(\tilde{v})$ *and* $\xi \in \mathrm{ID}(\mathcal{VR}an(\rho))$ *such that, for* $i = 1, 2$*,* $(\varphi_i\rho \wedge \rho_{\neq})\xi \wedge \xi_{\neq}$ *is satisfiable in* $\mathcal{T}_i$*.*
2. $\varphi_1 \wedge \varphi_2$ *is satisfiable in* $\mathcal{T}_1 \cup \mathcal{T}_2$*.*

[21] We do not even need $\mathcal{L}$ to be closed under negation and conjunction. It is enough that there is a formula $\varphi \in \mathcal{L}^{\Sigma_1}$, say, and a formula $\psi \in \mathcal{L}^{\Sigma}$ such that $\mathcal{T}_1 \models \neg\tilde{\exists}\,(\varphi \wedge \psi)$. Then, for no theory $\mathcal{T}_2$ will $\langle \varphi, \psi \rangle$ satisfy Cond. 4.1.

- **Instantiation**

  $\dfrac{\langle \varphi_1, \varphi_2 \rangle}{\langle \varphi_1 \rho \wedge \rho_{\neq}, \varphi_2 \rho \wedge \rho_{\neq} \rangle}$ for some $\rho \in \mathrm{IN}^{\Sigma}(\tilde{v})$ with $\tilde{v} := \mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2)$.

- **Identification**

  $\dfrac{\langle \varphi_1, \varphi_2 \rangle}{\langle \varphi_1 \xi \wedge \xi_{\neq}, \varphi_2 \xi \wedge \xi_{\neq} \rangle}$ for some $\xi \in \mathrm{ID}(\tilde{v})$ with $\tilde{v} := \mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2)$.

Figure 1: Derivation rules

Proof. It is enough to show that $(2 \Rightarrow 1)$. But that is an immediate consequence of Cond. 4.2, Prop. 4.1, and Prop. 3.11. $\square$

When combining of two theories it is necessary to make sure that their combination is *meaningful* in the first place, that is, it is not inconsistent (or trivial). This is particularly important when one considers theories as we do that share non logical symbols, because it is much easier for such theories to have constradicting consequences. Now, an implicit consequence of Def. 4.4 is that in all *non degenerate* cases in which two theories are at least partially N-O-combinable, their union is indeed consistent and so it does make sense to combine them. In fact, assume that there is a pair $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ such that each φ_i is satisfiable in $\mathcal{T}_i$. From Cond. 4.1 we can conclude that a model of $\mathcal{T}_1$ and model of $\mathcal{T}_2$ can be fused into a model of $\mathcal{T}_1 \cup \mathcal{T}_2$. Recalling that fusions have the cardinality of their component structures, we can make an even stronger claim: if each φ_i is satisfiable in a non-trivial model of $\mathcal{T}_i$ then $\mathcal{T}_1 \cup \mathcal{T}_2$ has a non-trivial model. In conclusion, we have the following.

Corollary 4.7 *Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be partially N-O-combinable over $\mathcal{L}$. If there is a pair $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ such that $\mathcal{T}_i \cup \{\tilde{\exists}\, \varphi_i\}$ is consistent (non-trivial) for $i = 1, 2$, then $\mathcal{T}_1 \cup \mathcal{T}_2$ is consistent (non-trivial).*

5 Combining Satisfiability Procedures

In this section, we will show that when a certain type of satisfiability problem is decidable for two N-O-combinable theories, it is possible to build a decision procedure for a corresponding satisfiability problem in the union theory, using the very decision procedures for the component theories.

In the following, we will assume that

- $\mathcal{L}$ is some class of formulas closed under identification and instantiation of variables;
- Σ_1 and Σ_2 are two *countable* signatures

1. Input: $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$.
2. Generate the pair $\langle \varphi'_1, \varphi'_2 \rangle := \langle \varphi_1 \rho \wedge \rho_{\neq}, \varphi_2 \rho \wedge \rho_{\neq} \rangle$ by an application of the **Instantiation** rule.
3. Generate the pair $\langle \psi_1, \psi_2 \rangle := \langle \varphi'_1 \xi \wedge \xi_{\neq}, \varphi'_2 \xi \wedge \xi_{\neq} \rangle$ by an application of the **Identification** rule.
4. Succeed if ψ_1 is satisfiable in $\mathcal{T}_1$ and ψ_2 is satisfiable in $\mathcal{T}_2$. Fail otherwise.

Figure 2: The Combination Method.

such that $\Sigma := \Sigma_1 \cap \Sigma_2$ is finite;

- $\mathcal{T}_1$ is a Σ_1-theory and $\mathcal{T}_2$ a Σ_2-theory.

First, we define two non-deterministic derivation rules, shown in Fig. 1, over pairs of formulas. Then, we define a (non-deterministic) combination procedure, shown in Fig. 2, that uses these rules to verify the satisfiability in $\mathcal{T}_1 \cup \mathcal{T}_2$ of pairs of formulas in $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$. Given the input problem $\langle \varphi_1, \varphi_2 \rangle$, the method first guesses an instantiation ρ into Σ-terms of the shared variables in the input pair $\langle \varphi_1, \varphi_2 \rangle$. Then, it guesses an identification ξ of the new set of shared variables. Finally, it checks that each member $\varphi_i \rho \xi$ of the new pair is satisfiable in the corresponding theory under the restriction $\rho_{\neq} \xi \wedge \xi_{\neq}$, succeeding only when both members are satisfiable.

In essence, the procedure is a non-deterministic version of the Nelson and Oppen combination method, but it extends that method in at least three ways: (1) it does not require that the input formulas be quantifier-free; (2) it allows the signatures of the component theories to share up to a finite number of symbols. (3) it considers only identifications over the free variables shared by the two input formulas, whereas Nelson and Oppen's effectively considers identifications over all the variables. The latter improvement is important at least for practical computational concerns (see [TH96] for a discussion) and has also been considered in the combination method for the unification problem introduced in [BS96].

Prop. 4.5 tells us that this extension is sound for component theories that are partially N-O-combinable over the chosen language $\mathcal{L}$. Let $\langle \varphi_1, \varphi_2 \rangle$ be an input pair of the combination method.

Proposition 5.1 (Soundness) *If one of the possible outputs of the identification step (Step 3) is a pair $\langle \psi_1, \psi_2 \rangle$ such that ψ_i is satisfiable in $\mathcal{T}_i$ for $i = 1, 2$, then $\langle \varphi_1, \varphi_2 \rangle$ is satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$.*

```
Comb ⟨φ1, φ2⟩ =
        let ⟨ψ1, ψ2⟩ = Ident Inst ⟨φ1, φ2⟩
        in
            return (Sat1 ψ1 and Sat2 ψ2)
```

Figure 3: The Combination Procedure.

If the component theories are totally N-O-combinable over $\mathcal{L}$, Theor. 4.6 tells us that the method is also complete, in the sense specified below.

Proposition 5.2 (Completeness) *If $\langle \varphi_1, \varphi_2 \rangle$ is satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$, then there is a pair $\langle \psi_1, \psi_2 \rangle$ among the possible outputs of the identification step such that ψ_i is satisfiable in $\mathcal{T}_i$ for $i = 1, 2$.*

Recall that formulas of the form $(\varphi\rho \wedge \rho_{\neq})\xi \wedge \xi_{\neq}$ are Σ-restricted formulas in the sense of Def. 2.4. It follows that, if $\mathcal{T}_1$ and $\mathcal{T}_2$ are N-O-combinable and the satisfiability in $\mathcal{T}_i$ ($i = 1, 2$) of Σ-restricted formulas is decidable, it is possible to use the combination method above to obtain a semi-decision procedure for the satisfiability in $\mathcal{T}_1 \cup \mathcal{T}_2$ of formulas in $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$.

More precisely, assume that for $i = 1, 2$, Sat_i is a boolean procedure that decides the satisfiability in $\mathcal{T}_i$ for the formulas of $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$. Also assume that *Inst* and *Ident* are two (non-deterministic) procedures that implement, respectively, the instantiation and the identification rules of Fig. 1. Then, recalling that non-deterministic procedures are said to succeed if one of their possible runs is successful, it is immediate that the procedure defined in Fig. 3 has the following property.

Proposition 5.3 *The procedure Comb succeeds on an input $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ iff $\langle \varphi_1, \varphi_2 \rangle$ is satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$.*

The reason *Comb* is only a *semi*-decision procedure, even for totally N-O combinable theories, is that in general it may diverge[22] on unsatisfiable inputs. In fact, whenever Σ contains a function symbol of non-zero arity and the set of variables shared by the two formulas in the input is nonempty, there is an infinite number of possible instantiations over that set; if the input pair is unsatisfiable in the union theory, however, none of these instantiations will make both calls to Sat_1 and Sat_2

[22] Strictly speaking, we should say something like: "it may infinitely fail". The reason is that the non-determinism in the instantiation step is generally unbounded. It should be clear that, at the cost of a less elegant definition, we could give an equivalent reformulation of the procedure according to the standard (that is, bounded) notion of non-determinism. (For instance, by considering all instantiations ρ into terms of height n first, then those into terms of height $n + 1$, and so on.) According to that definition, the procedure would diverge in the conventional sense.

succeed.[23] Notice that *Comb* can be easily modified so that it will not diverge on input pairs containing an i-pure formula that is already unsatisfiable in $\mathcal{T}_i$, and hence in $\mathcal{T}_1 \cup \mathcal{T}_2$. The non-termination problem arises only for *genuine* combination questions, that is, input pairs that are unsatisfiable in the union theory even if each of their pure formulas is satisfiable in the corresponding component theory.

We are currently investigating the cases in which *Comb* can be turned into a decision procedure. For now, it is interesting to notice that, although *Comb* is only a semi-decision procedure, it does yield strong decidability results in case of axiomatizable component theories. In fact, as pointed out, *Comb* diverges only on the inputs that are *not* satisfiable in the union theory. This means that, when *Comb* is applicable, the set of pairs satisfiable in the union theory is recursively enumerable. Now, it is an immediate consequence of the completeness of First-Order Predicate Calculus that the set of formulas *un*satisfiable in an axiomatizable theory is recursively enumerable. It follows that if *Comb* is applicable to two theories $\mathcal{T}_1$ and $\mathcal{T}_2$ such that $\mathcal{T}_1 \cup \mathcal{T}_2$ is axiomatizable, then the set of pairs satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$ is recursive.

Observing that a sufficient condition for $\mathcal{T}_1 \cup \mathcal{T}_2$ to be axiomatizable is that both $\mathcal{T}_1$ and $\mathcal{T}_2$ are axiomatizable, we then obtain at once the following decidability result.

Proposition 5.4 *Assume that, for $i = 1, 2$, $\mathcal{T}_i$ is axiomatizable and the satisfiability in $\mathcal{T}_i$ of formulas of $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ is decidable. If $\mathcal{T}_1$ and $\mathcal{T}_2$ are N-O-combinable over $\mathcal{L}$, then the satisfiability in $\mathcal{T}_1 \cup \mathcal{T}_2$ of formulas in $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ is decidable.*

So far, we have used a rather limited language, $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$, to express mixed constraints: we have considered only the mixed constraints that are expressible as the conjunction of two pure constraints which, in addition, do not share non-logical symbols. There is no doubt that our combination results would be much nicer if they could be given in terms of $\mathcal{L}^{\Sigma_1 \cup \Sigma_2}$ instead. This is in fact possible, but at the cost of some closure assumptions on $\mathcal{L}$.[24] We describe such assumptions below and conclude the section by showing, as an example, how they let us improve Prop. 5.4. In the next section, we will use them to get more specific combination results.

Definition 5.5 *Given two signatures Ω_1 and Ω_2, we say that a class $\mathcal{L}$ of formulas is* closed under purification wrt $\langle \Omega_1, \Omega_2 \rangle$ *if for every $\varphi \in \mathcal{L}^{\Omega_1 \cup \Omega_2}$, there is a finite set $\{\langle \varphi_1^i, \varphi_2^i \rangle\}_{i<m} \subseteq \mathcal{L}^{\Omega_1} \otimes \mathcal{L}^{\Omega_2}$ such that*

- *$\varphi_1^i \wedge \varphi_2^i \in \mathcal{L}^{\Omega_1 \cup \Omega_2}$ for all $i < m$ and*
- *φ is satisfiable iff $\bigvee_{i<m} (\varphi_1^i \wedge \varphi_2^i)$ is satisfiable.*

We call $\bigvee_{i<m} (\varphi_1^i \wedge \varphi_2^i)$ a disjunctive pure form of φ *(wrt $\langle \Omega_1, \Omega_2 \rangle$).*

[23] If both calls succeeded, by the procedure's soundness, the input would be satisfiable.

[24] Notice that we barely made any assumptions on $\mathcal{L}$, so far.

NewComb $\varphi =$
let $\langle \psi_1, \psi_2 \rangle =$ *Ident Inst Pure* φ
in
return ($Sat_1\ \psi_1$ **and** $Sat_2\ \psi_2$)

Figure 4: The New Combination Procedure.

We will say that $\mathcal{L}$ is *effectively* closed under purification wrt $\langle \Omega_1, \Omega_2 \rangle$ if, for each formula $\varphi \in \mathcal{L}^{\Omega_1 \cup \Omega_2}$, a disjunctive pure form of φ is effectively computable. The class of quantifier-free formulas is an example of such a class. As a matter of fact, *Qff* is effectively closed under purification with respect to any pair $\langle \Omega_1, \Omega_2 \rangle$ of signatures.

In fact, let $\varphi \in Qff^{\Omega}$. There is a well known algorithm that can first convert φ into a logically equivalent formula $\bigvee_{i<m} \varphi^i$ in disjunctive normal form. Then, we can apply another well-known algorithm (see [Rin96b, TH96] among others) to each disjunct φ^i that, by means of variable abstraction, produces a pair $\langle \varphi_1^i, \varphi_2^i \rangle$ such that $\varphi_1^i \wedge \varphi_2^i \in Qff^{\Omega_1} \otimes Qff^{\Omega_2}$ and $\tilde{\exists}\, (\varphi_1^i \wedge \varphi_2^i)$ and $\tilde{\exists}\, \varphi$ are logically equivalent. It is easy to see then that $\{\langle \varphi_1^i, \varphi_2^i \rangle\}_{i<m}$ satisfies all the requirements in the definition above.

When $\mathcal{L}$ is effectively closed under purification wrt our initial pair of signatures $\langle \Sigma_1, \Sigma_2 \rangle$, we can extend *Comb* to the procedure *NewComb*, defined in Fig. 4, by adding a procedure, *Pure*, that, given an input formula φ from $\mathcal{L}^{\Sigma_1 \cup \Sigma_2}$, returns (non-deterministically) one disjunct of φ's disjunctive pure form.[25] Observing that φ is satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$ if and only if some disjunct of its disjunctive pure form is satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$, it is immediate that the new procedure is correct as well.

Corollary 5.6 *The procedure NewComb succeeds on an input $\varphi \in \mathcal{L}^{\Sigma_1 \cup \Sigma_2}$ iff φ is satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$.*

We can now express the previous decidability result more neatly as follows.

Proposition 5.7 *Assume that, for $i = 1, 2$, $\mathcal{T}_i$ is axiomatizable and the satisfiability in $\mathcal{T}_i$ of formulas of $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ is decidable. If $\mathcal{L}$ is effectively closed under purification and $\mathcal{T}_1$ and $\mathcal{T}_2$ are N-O-combinable over $\mathcal{L}$, then the satisfiability in $\mathcal{T}_1 \cup \mathcal{T}_2$ of formulas of $\mathcal{L}^{\Sigma_1 \cup \Sigma_2}$ is decidable.*

The above proposition seems to suggest that we get a somewhat *weaker* decidability result for the union theory given that we consider satisfiability of restricted formulas in the component theories but only satisfiability of unrestricted formulas in the union theory. However, we do have a really modular result.

[25] We assume with no loss of generality that the disjunctive pure form of φ is unique.

Corollary 5.8 *Assume that $\mathcal{L}$ is effectively closed under purification, $\mathcal{T}_1$ and $\mathcal{T}_2$ are N-O-combinable over $\mathcal{L}$, and $\mathcal{T}_i$ is axiomatizable for $i = 1, 2$. Then, if the satisfiability in $\mathcal{T}_i$ of formulas of $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ is decidable, the satisfiability in $\mathcal{T}_1 \cup \mathcal{T}_2$ of formulas of $Res(\mathcal{L}^{\Sigma_1 \cup \Sigma_2}, \Sigma)$ is also decidable.*

A rigorous proof of this corollary is easy but rather tedious. The following informal argument should suffice. Given a formula φ, *NewComb* first purifies φ into a pair $\langle \varphi_1, \varphi_2 \rangle$, then specializes $\langle \varphi_1, \varphi_2 \rangle$ into a pair $\langle \varphi_1 \rho \xi, \varphi_2 \rho \xi \rangle$, and finally adds to each $\varphi_i \rho \xi$ the Σ-restriction $\rho_{\neq} \xi \wedge \xi_{\neq}$ on some of its shared variables before passing it to Sat_i. It is possible to show that all our combination results lift to the case in which non-shared variables are also considered.[26] Now, if the input φ comes already equipped with some Σ-restriction $\varepsilon_{\neq} \wedge \varepsilon_{\neq}^{\Sigma}$, it is enough for *NewComb* to purify φ into $\langle \varphi_1, \varphi_2 \rangle$ and then generate the formulas $(\varphi_i' \rho \wedge \rho_{\neq}) \xi \wedge \xi_{\neq}$ as before where φ_i' is now $\varphi_i \wedge \varepsilon_{\neq} \wedge \varepsilon_{\neq}^{\Sigma}$. It is a simple exercise to show that $(\varphi_i' \rho \wedge \rho_{\neq}) \xi \wedge \xi_{\neq}$ is an element of $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$.[27]

As a criterion for identifying component theories to which our combination method applies, Def. 4.4 is perhaps too abstract. In the next section, we try to establish more concrete conditions that are sufficient for N-O-combinability.

6 Classes of N-O-combinable Theories

In the simple case of signature-disjoint theories, a sufficient condition for N-O-combinability is the restriction to component theories that are *stably-infinite*. The notion of stable-infiniteness was first introduced in [Opp80] to show the correctness of the Nelson-Oppen combination method. We discuss this case in the next subsection and show how the original combination results by Nelson and Oppen are in fact subsumed by ours. In the subsequent subsection, we introduce our generalization of stable-infiniteness which can be used to identify N-O-combinable theories with non-disjoint signatures.

6.1 Disjoint Signatures

Looking back at Lemma 3.7, one realizes that all we need there for our combination result is that the component structures in which the pure formulas are satisfiable have the same cardinality. One way to enforce this with theories is to restrict attention to those that satisfy the following property.

[26] Having to consider only shared variables is in a sense an optimization of this more general case.

[27] Here, it is important to keep in mind that $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ is closed under logical equivalence.

Definition 6.1 (Stably-Infinite Theory) *We say that a consistent Ω-theory $\mathcal{T}$ is* stably-infinite *if every quantifier-free Ω-formula satisfiable in $\mathcal{T}$ is satisfiable in an infinite model of $\mathcal{T}$.*[28]

It is easy to see that stably-infinite theories admit infinite models and so, by the Upward and Downward Löwenheim-Skolem theorems [Hod93b], admit models of any infinite cardinality[29]. This entails, first, that if a formula is satisfiable in a stably-infinite theory, it is satisfiable in models of the theory of arbitrary, infinite cardinality. Second (by Cor. 4.2), that the union of two stably-infinite, signature-disjoint theories $\mathcal{T}_1$ and $\mathcal{T}_2$ is always consistent. In addition, we have the following.

Proposition 6.2 *Any two stably-infinite, signature-disjoint theories are totally N-O-combinable over Qff.*

Proof. Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be our theories. Let Σ_i be the signature of $\mathcal{T}_i$ for $i = 1, 2$ and let $\Sigma := \Sigma_1 \cap \Sigma_2 (= \emptyset)$. First we show that $\mathcal{T}_1$ and $\mathcal{T}_2$ satisfy Cond. 4.1 for all $\langle \varphi_1, \varphi_2 \rangle \in \mathit{Qff}^{\Sigma_1} \otimes \mathit{Qff}^{\Sigma_2}$.

Observe that $\psi_i := (\varphi_i \rho \wedge \rho_{\neq})\xi \wedge \xi_{\neq}$, where $\tilde{v} := \mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2)$, is equivalent to a formula in Qff^{Σ_i} for $i = 1, 2$, as every $\rho \in \mathrm{IN}^{\Sigma}(\tilde{v})$ coincides with the empty instantiation and $\rho_{\neq}$ with the empty set. It follows by the stable-infiniteness of the theories that if ψ_i is satisfiable in $\mathcal{T}_i$, it is satisfiable in a model $\mathcal{A}_i$ of $\mathcal{T}_i$ whose cardinality is $\mathit{Card}(\Sigma_1 \cup \Sigma_2)$. We have already seen that, then, $\mathcal{A}_1$ and $\mathcal{A}_2$ are trivially Σ-fusible.

To see that $\mathcal{T}_1$ and $\mathcal{T}_2$ satisfy Cond. 4.2 as well, simply notice that since Σ is empty, every individual of any model of $\mathcal{T}_1 \cup \mathcal{T}_2$ is Σ-isolated. □

As a consequence of the above proposition, we obtain the following simplified version of Theor. 4.6.

Theorem 6.3 *Consider two stably infinite-theories $\mathcal{T}_1$ and $\mathcal{T}_2$ with respective, disjoint signatures Σ_1 and Σ_2. For $i = 1, 2$, let $\varphi_i \in \mathit{Qff}^{\Sigma_i}$. Then, where $\tilde{v} := \mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2)$, the following are equivalent:*

1. *$\varphi_i \xi \wedge \xi_{\neq}$ is satisfiable in $\mathcal{T}_i$ for each $i = 1, 2$ and some $\xi \in \mathrm{ID}(\tilde{v})$;*
2. *$\varphi_1 \wedge \varphi_2$ is satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$.*

It is exactly on this result that the soundness and completeness of the original combination method by Nelson and Oppen ultimately rests (see [Rin96b, TH96]).

[28] The original definition in [Opp80] considered only universal theories. We have lifted this restriction because it is inessential; see Theor. 6.3

[29] Greater than, or equal to, the cardinality of their signature, to be precise.

6.2 Non-Disjoint Signatures

We have tried to extend the notion of stable-infiniteness so that it would provide, along with some additional requirements, a sufficient condition for the N-O-combinability of theories with non-disjoint signatures.

Definition 6.4 (Σ-Stable Theory) *Let $\mathcal{T}$ be a consistent theory of signature Ω, Σ a finite subset of Ω, $\mathbf{K}$ a Σ-variety, and $\mathcal{L}$ a class of formulas. We say that $\mathcal{T}$ is* Σ-stable wrt $\mathbf{K}$ over $\mathcal{L}^\Omega$ *if every formula of $\mathcal{L}^\Omega$ satisfiable in $\mathcal{T}$ is satisfiable in a model $\mathcal{A}$ of $\mathcal{T}$ such that $\mathcal{A}^\Sigma$ is free in $\mathbf{K}$ over a basis of cardinality $Card(\Omega)$.*[30]

We will simply say that $\mathcal{T}$ is *Σ-stable over* $\mathcal{L}^\Omega$ if $\mathcal{T}$ is Σ-stable over $\mathcal{L}^\Omega$ wrt its own Σ-variety ($Mod(\mathcal{T}^\Sigma_{\mathrm{At}})$).

That Σ-stability is an extension of stable-infiniteness is given by the fact that the stably-infinite theories are exactly those theories that are Σ-stable over *Qff* where Σ is an empty signature.

Proposition 6.5 *Let $\mathcal{T}$ be a Ω-theory and Σ an empty signature. Then, $\mathcal{T}$ is stably-infinite iff $\mathcal{T}$ is Σ-stable over Qff^Ω.*

Proof. ($\Rightarrow$) Assume that $\mathcal{T}$ is stably-infinite and let $\psi \in Qff^\Omega$ be satisfiable in $\mathcal{T}$. By definition of stable infiniteness, $\mathcal{T} \cup \{\tilde{\exists}\, \psi\}$ has an infinite model and in particular one of cardinality $Card(\Omega)$. Call it $\mathcal{A}$ and observe that $\mathcal{A}^\Sigma$ is trivially absolutely free over its whole carrier and that the Σ-atomic theory of $\mathcal{T}$ is simply the empty theory (over the empty signature). It follows that ψ is satisfiable in a model of $\mathcal{T}$ whose reduct to Σ is free in the Σ-atomic theory of $\mathcal{T}$ over a basis of cardinality $Card(\Omega)$.

($\Leftarrow$) Assume that $\mathcal{T}$ is Σ-stable over Qff^Ω and let $\psi \in Qff^\Omega$ be satisfiable in $\mathcal{T}$. By definition of Σ-stability, ψ is satisfiable in a model of $\mathcal{T}$ containing at least $Card(\Omega)$ individuals and so it is satisfiable in an infinite model of $\mathcal{T}$. □

Thanks to the above result, we know that the class of Σ-stable theories is non-empty. One might wonder, however, whether there are non trivial languages for which this class is effectively larger than that of stably-infinite theories. We will answer this question affirmatively in Sect. 6.3. For now, let us see how a theory admitting infinite models can still fail to be Σ-stable (wrt its own Σ-variety).

Example 6.6 Consider the theory $\mathcal{T} := \{\varphi \to a \equiv b\}$ where φ is a non-tautological ground formula and a and b are constant symbols. Let $\Sigma := \{a, b\}$ and observe that the Σ-atomic theory of $\mathcal{T}$ is empty. Now, a equals b in every model of $\mathcal{T}$ that satisfies φ, therefore the model's reduct to Σ is certainly not free in the Σ-variety of $\mathcal{T}$. It follows that $\mathcal{T}$ is not Σ-stable over the language of φ.

[30] Observe that this basis will be infinite even when Ω is finite.

Example 6.7 As in the previous example we can show that, where a, b and c are constant symbols and $\Sigma := \{a, b, c\}$, the theory $\mathcal{T} := \{(a \equiv b) \vee (a \equiv c)\}$ is Σ-stable for no class of formulas for the simple fact that the Σ-reduct of no model of $\mathcal{T}$ is absolutely free (that is, free in the Σ-atomic theory of $\mathcal{T}$, which is again empty).

The theorem below provides a sufficient condition for the N-O-combinability of two theories with non (necessarily) disjoint signatures. In the following, we will fix two countable[31] signatures Σ_1 and Σ_2 and assume that $\Sigma := \Sigma_1 \cap \Sigma_2$ is finite. First, we need the following maybe not so trivial result.

Lemma 6.8 *Let $\mathcal{L}$ be a class of formulas, $\mathcal{T}_1, \mathcal{T}_2$ two theories of respective signature Σ_1, Σ_2, and E_0 a Σ-atomic theory. If E_0 is the Σ-atomic theory of both $\mathcal{T}_1$ and $\mathcal{T}_2$ and each $\mathcal{T}_i$ is Σ-stable over $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$, then E_0 is also the Σ-atomic theory of $\mathcal{T}_1 \cup \mathcal{T}_2$.*

Proof. It is immediate that $E_0 \subseteq (\mathcal{T}_1 \cup \mathcal{T}_2)^{\Sigma}_{\mathrm{At}}$. We show that $(\mathcal{T}_1 \cup \mathcal{T}_2)^{\Sigma}_{\mathrm{At}} \subseteq E_0$. First recall that we always assume that $\mathcal{L}$ contains the universally true sentence. Together with Def. 6.4, this entails that each $\mathcal{T}_i$ has a model $\mathcal{A}_i$ whose Σ-reduct is free in $Mod(E_0)$ over a countably-infinite set. It follows by Prop. 3.4 and Prop. 4.1 that $\mathcal{A}_1$ and $\mathcal{A}_2$ are fusible in a model $\mathcal{F}$ of $\mathcal{T}_1 \cup \mathcal{T}_2$. Since by definition of fusion, $\mathcal{F}^{\Sigma}$ is isomorphic to ${\mathcal{A}_1}^{\Sigma}$, say, we can conclude that $\mathcal{F}^{\Sigma}$ as well is free in $Mod(E_0)$ over a (countably) infinite set X.

Now, let φ be a Σ-atom such that $\mathcal{T}_1 \cup \mathcal{T}_2 \models \tilde{\forall}\, \varphi$. Then $\mathcal{F}^{\Sigma} \models \tilde{\forall}\, \varphi$, as well since $\mathcal{F}$ is a model of $\mathcal{T}_1 \cup \mathcal{T}_2$ and φ is a Σ-formula. In particular, since $Card(X)$ is greater than $Card(\mathcal{V}ar(\varphi))$, there is a discrete $\tilde{x}$ in X such that $\mathcal{F}^{\Sigma} \models \varphi[\tilde{x}]$. It follows by Prop. 2.21 that $E_0 \models \tilde{\forall}\, \varphi$. Given that E_0 is the Σ-atomic theory of $\mathcal{T}_1$, say, it is easy to show that in fact $\tilde{\forall}\, \varphi \in E_0$. $\square$

Theorem 6.9 *For all classes $\mathcal{L}$ of formulas and theories $\mathcal{T}_1, \mathcal{T}_2$ of respective signature Σ_1, Σ_2, we have the following.*

1. *If $\mathcal{T}_1$ and $\mathcal{T}_2$ have the same Σ-atomic theory E_0 and each $\mathcal{T}_i$ is Σ-stable over $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$, then $\mathcal{T}_1$ and $\mathcal{T}_2$ are partially N-O-combinable over $\mathcal{L}$.*
2. *If, in addition, E_0 is collapse-free and $\mathcal{T}_1 \cup \mathcal{T}_2$ is Σ-stable over $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$, then $\mathcal{T}_1$ and $\mathcal{T}_2$ are totally N-O-combinable over $\mathcal{L}$.*

Proof. Let $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ and $\tilde{v} := \mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2)$.

(1) It suffices to show that $\langle \varphi_1, \varphi_2 \rangle$ satisfies Cond. 4.1. Assume that there is a $\rho \in \mathrm{IN}^{\Sigma}(\tilde{v})$ and a $\xi \in \mathrm{ID}(\mathcal{VR}an(\rho))$ such that $\psi_i := (\varphi_i \rho \wedge \rho_{\neq})\xi \wedge \xi_{\neq}$ is satisfiable in $\mathcal{T}_i$, for $i = 1, 2$. Observe that $\psi_i \in Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ and so, by the Σ-stability of

[31] All we need really is that $Card(\Sigma_1) = Card(\Sigma_2)$. We assume that it is ω only for simplicity.

$\mathcal{T}_i$, it is satisfiable in some $\mathcal{A}_i \in Mod(\mathcal{T}_i)$ such that ${\mathcal{A}_i}^\Sigma$ is free in $Mod(E_0)$ over a countably-infinite set X_i. By Prop. 3.4 then, $\mathcal{A}_1$ and $\mathcal{A}_2$ are fusible over $\langle X_1, X_2\rangle$. To see that they are Σ-fusible observe that, for $i = 1, 2$, X_i is a set of Σ-generators for $\mathcal{A}_i$ and so it necessarily includes $Is({\mathcal{A}_i}^\Sigma)$.

(2) It suffices to show that $\langle \varphi_1, \varphi_2\rangle$ satisfies Cond. 4.2. Let $\mathcal{T} := \mathcal{T}_1 \cup \mathcal{T}_2$ and assume that $\langle \varphi_1, \varphi_2\rangle$ is satisfiable in $\mathcal{T}$. As $\mathcal{T}$ is Σ-stable over $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ by assumption, $\langle \varphi_1, \varphi_2\rangle$ is satisfiable in a model $\mathcal{A}$ of $\mathcal{T}$ whose reduct to Σ is free in the Σ-variety of $\mathcal{T}$. Since the Σ-variety of $\mathcal{T}$ is $Mod(E_0)$ by Lemma 6.8 and E_0 is collapse-free by assumption, we have by Prop. 2.22 that $\mathcal{A}^\Sigma$ is generated by its isolated individuals. In conclusion, we have shown that $\varphi_1 \wedge \varphi_2$ is satisfiable in a model of $\mathcal{T}$ that is Σ-generated by its Σ-isolated individuals. □

Total (as opposed to partial) N-O-combinability of the component theories is important for our combination procedure because it guarantees its completeness, as we have seen in the previous section. An irksome feature of the theorem above is that the theorem explicitly assumes the Σ-stability of $\mathcal{T}_1 \cup \mathcal{T}_2$ over $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ in order to yield the total N-O-combinability of $\mathcal{T}_1$ and $\mathcal{T}_2$. It would be much nicer, however, if the Σ-stability of a union theory were provable from the Σ-stability of its component theories. Unfortunately, that is not the case in general. More information on either the constraint language or the component theories is needed.

For instance, if Σ is empty and $\mathcal{L} = Qff$, $\mathcal{T}_1 \cup \mathcal{T}_2$ is indeed Σ-stable over $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ whenever both $\mathcal{T}_1$ and $\mathcal{T}_2$ are Σ-stable over $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$—which is simply Qff^{Σ_i}. Recalling Prop. 6.5, to prove this it is enough to show the following.

Proposition 6.10 *The union of two signature-disjoint stably-infinite theories is stably-infinite.*

Proof. Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be stably-infinite theories of signature Ω_1 and Ω_2 respectively, and assume that $\Omega_1 \cap \Omega_2 = \emptyset$. Recall that $\mathcal{T}_1 \cup \mathcal{T}_2$ is consistent and consider a formula $\varphi \in Qff^{\Omega_1 \cup \Omega_2}$ satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$. Since Qff is closed under purification wrt $\langle \Omega_1, \Omega_2\rangle$, we can assume with no loss of generality that φ is of the form $\varphi_1 \wedge \varphi_2$ where $\langle \varphi_1, \varphi_2\rangle \in \mathcal{L}^{\Omega_1} \otimes \mathcal{L}^{\Omega_2}$. Let ξ be the identification of $\mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2)$ induced by any valuation satisfying $\varphi_1 \wedge \varphi_2$ in a model $\mathcal{A}$ of $\mathcal{T}_1 \cup \mathcal{T}_2$. Clearly, for $i = 1, 2$, the pure formula $\varphi_i\xi \wedge \xi_{\neq}$ is also satisfiable in $\mathcal{A}$ and hence in $\mathcal{T}_i$, as $\mathcal{A}^{\Omega_i} \in Mod(\mathcal{T}_i)$ for $i = 1, 2$. We leave it to the reader to verify that each $\varphi_i\xi \wedge \xi_{\neq}$ is then satisfiable in a model $\mathcal{A}_i$ of $\mathcal{T}_i$ such that $\mathcal{A}_1$ and $\mathcal{A}_2$ have the same infinite cardinality. By Prop. 3.7 then, $\varphi_1 \wedge \varphi_2$ is satisfiable in some fusion $\mathcal{F}$ of $\mathcal{A}_1$ and $\mathcal{A}_2$, which is an infinite model of $\mathcal{T}_1 \cup \mathcal{T}_2$ by Prop. 4.1. □

The proposition above, not only turns Theor. 6.3 into a mere corollary of Theor. 6.9, but makes the combination method immediately scalable, by iteration, to the combination of more than two signature-disjoint theories. In fact, if $\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3$ are stably-infinite and pairwise signature-disjoint theories, we obtain a combined satisfiability procedure for $\mathcal{T}_1 \cup \mathcal{T}_2 \cup \mathcal{T}_3$ simply by combining, for example, the satisfiability procedure for $\mathcal{T}_1$ with the combined satisfiability procedure for $\mathcal{T}_2 \cup \mathcal{T}_3$.

In more general settings, things are not so smooth. As we have seen, Σ-stability over Σ-restricted formulas and same Σ-atomic theory are enough to prove the extended combination method sound for theories sharing exactly the symbols of Σ. To prove the method complete, however, we must also verify that the Σ-atomic theory is collapse-free and that the union theory too is Σ-stable (albeit on a simpler language). Alternatively, we can always try to verify directly that the union theory satisfies Cond. 4.2, which may be an easier task in some cases, as the examples in Sect. 6.3 will show.

There is no doubt that the notion of Σ-stability would be more useful to our ends if Σ-stability over Σ-restricted formulas were modular with respect to theory union. Unfortunately, what we can show is a much weaker modularity result, in terms of *totally* Σ-restricted formulas.

Proposition 6.11 *Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be as in Theor. 6.9 and assume that each $\mathcal{T}_i$ is Σ-stable over $TRes(\mathcal{L}^{\Sigma_i}, \Sigma)$. If $\mathcal{T}_1$ and $\mathcal{T}_2$ have the same Σ-atomic theory E_0, then $\mathcal{T}_1 \cup \mathcal{T}_2$ is Σ-stable over $TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$.*

Proof. Let $\psi(\tilde{v}) \in TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$ be satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$. We know that ψ has the form $\varphi_1 \wedge \varphi_2 \wedge \varepsilon_{\neq}(\tilde{v}) \wedge \varepsilon_{\neq}^{\Sigma}(\tilde{v})$ with $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$. With no loss of generality, we can assume that $\mathcal{V}ar(\varphi_1) = \mathcal{V}ar(\varphi_2) = \tilde{v}$, otherwise, we can always conjoin each φ_i with a suitable tautological formula in the missing variables.[32] It is easy to see that

$$\psi_i := \varphi_i(\tilde{v}) \wedge \varepsilon_{\neq}(\tilde{v}) \wedge \varepsilon_{\neq}^{\Sigma}(\tilde{v})$$

is satisfiable in $\mathcal{T}_i$ for $i = 1, 2$. In particular, since $\psi_i \in TRes(\mathcal{L}^{\Sigma_i}, \Sigma)$ and $\mathcal{T}_i$ is Σ-stable over $TRes(\mathcal{L}^{\Sigma_i}, \Sigma)$ by assumption, ψ_i is satisfiable in a model $\mathcal{A}_i$ of $\mathcal{T}_i$ such that ${\mathcal{A}_i}^{\Sigma}$ is free in $Mod(E_0)$ over a countably-infinite basis. Again, we can show by Prop. 3.4 that $\mathcal{A}_1$ and $\mathcal{A}_2$ are Σ-fusible and thus conclude, by a corollary of Prop. 3.10[33], that ψ is satisfiable in a fusion $\mathcal{F}$ of $\mathcal{A}_1$ and $\mathcal{A}_2$. We have already seen that $\mathcal{F} \in Mod(\mathcal{T}_1 \cup \mathcal{T}_2)$ and $\mathcal{F}^{\Sigma}$ is free in $Mod(E_0)$ over a countably-infinite basis. To complete the proof then, it is enough to recall that, by Lemma 6.8, the Σ-atomic theory of $\mathcal{T}_1 \cup \mathcal{T}_2$ coincides with E_0. □

The above result is not sufficient for our needs given that, in general, the class of formulas $TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$ is strictly included in $Res(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$. One might argue, however, that if we limit ourselves to totally Σ-restricted formulas, we do get

[32] It is essentially harmless to assume that $\mathcal{L}$ contains tautological formulas over arbitrary sets of variables.

[33] We need the corollary because the result in Prop. 3.10 is in terms of $\varphi_1 \wedge \varphi_2$ not $\varphi_1 \wedge \varphi_2 \wedge \varepsilon_{\neq}(\tilde{v}) \wedge \varepsilon_{\neq}^{\Sigma}(\tilde{v})$.

the nice modularity and completeness results we long for. Unfortunately, that is not quite the case.

In fact, recall that our ultimate goal is to work with formulas in $\mathcal{L}^{\Sigma_1 \cup \Sigma_2}$, whether they have an attached Σ-restriction or not. As we saw, these formulas can be managed by the combination method provided that $\mathcal{L}$ is effectively closed under purification wrt $\langle \Sigma_1, \Sigma_2 \rangle$. What we do then is, first, to put an input formula $\varphi(\tilde{v}) \in \mathcal{L}^{\Sigma_1 \cup \Sigma_2}$ into disjunctive pure form and, then, test the satisfiability of its disjuncts, which are members of $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$. Now, the pure form φ' of φ may have a different (typically larger) set of free variables. Therefore, even if we start with the totally Σ-restricted formula

$$\varphi(\tilde{v}) \wedge \varepsilon_{\neq}(\tilde{v}) \wedge \varepsilon^{\Sigma}_{\neq}(\tilde{v}),$$

after purification we may end up with a partially Σ-restricted formula of the form

$$\varphi'(\tilde{u}) \wedge \varepsilon_{\neq}(\tilde{v}) \wedge \varepsilon^{\Sigma}_{\neq}(\tilde{v}).$$

When $\mathcal{L}$ coincides with *Qff*, it is possible to generate φ' so that

- $\tilde{v} \subseteq \tilde{u}$ and
- $\varphi' \models u_i \equiv t_i$ for all $u_i \in \tilde{u} \setminus \tilde{v}$,

where t_i is a Σ_1- or Σ_2-term whose root symbol is not in Σ. This entails that we can extend the Σ-restriction of φ to the whole $\tilde{u}$ without loss of solutions only if we know that the function symbols in $(\Sigma_2 \cup \Sigma_2) \setminus \Sigma$ only generate Σ-isolated individuals. To be more precise, let us consider the following stronger form of Σ-stability.

Definition 6.12 *Let $\mathcal{T}$ be a consistent theory of signature Ω, Σ a finite subset of Ω, and $\mathcal{L}$ a class of formulas. We say that $\mathcal{T}$ is* strongly *Σ-stable over $\mathcal{L}^{\Omega}$ if every formula of $\mathcal{L}^{\Omega}$ satisfiable in $\mathcal{T}$ is satisfiable in a model $\mathcal{A}$ of $\mathcal{T}$ such that*

- *$\mathcal{A}^{\Sigma}$ is free in the Σ-variety of $\mathcal{T}$ over a set X of cardinality $Card(\Omega)$,*
- *for all $f \in (\Omega \setminus \Sigma)^{\mathrm{F}}$, $f^{\mathcal{A}}(A) \subseteq X \cup \{c^{\mathcal{A}} \mid c \in \Sigma, c$ constant$\}$.*[34]

By our earlier observations, and recalling the proof of Prop. 6.11, it is not difficult to see that strong Σ-stability does yield the kind of modularity we are seeking.

Proposition 6.13 *Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be as in Theor. 6.9. Assume that each $\mathcal{T}_i$ is strongly Σ-stable over $TRes(Qff^{\Sigma_i}, \Sigma)$. If $\mathcal{T}_1$ and $\mathcal{T}_2$ have the same Σ-atomic theory E_0 and E_0 is collapse-free, then $\mathcal{T}_1 \cup \mathcal{T}_2$ is strongly Σ-stable over $TRes(Qff^{\Sigma_1 \cup \Sigma_2}, \Sigma)$.*

Which leads immediately to the following nice corollary of Theor. 6.14.

[34] Where $f^{\mathcal{A}}(A)$ denotes the range of $f^{\mathcal{A}}$.

Corollary 6.14 *Two theories $\mathcal{T}_1$ and $\mathcal{T}_2$ with respective signature Σ_1 and Σ_2 are totally N-O combinable over $TRes(Qff, \Sigma)$ if*

1. *each $\mathcal{T}_i$ is strongly Σ-stable over $TRes(Qff^{\Sigma_i}, \Sigma)$,*
2. *$\mathcal{T}_1$, $\mathcal{T}_2$ have the same Σ-atomic theory E_0,*
3. *E_0 is collapse-free.*

Unfortunately, while the restriction to totally Σ-restricted formulas per se may still be a sensible one and lead to useful applications (ground formulas, for instance, are totally Σ-restricted for any Σ), the same cannot be said for the restriction to strongly Σ-stable theories. The concept of strong Σ-stability is perhaps too stringent to have interesting instances.

6.3 Some Σ-Stable Theories

In this section, we give some examples of classes of Σ-stable theories and show which theories within these classes are N-O-combinable. We believe that more classes can and should be identified in order to better assess the practical significance of the extended combination method. For now, we can see the results below and their proofs as a set of general guidelines on how to apply Theor. 6.9 in practice.

Again, we will consider only *countable* signatures. While some results could be given for greater cardinalities, considering only countable signatures is a sensible restriction given that we are ultimately interested in building decision procedures (which by definition, only consider countable input alphabets).

Example 6.15 Let $\mathbf{T}$ be the class of all stably-infinite theories $\mathcal{T}$ such that, for all constant symbols $k_1, k_2 \in \Sigma_{\mathcal{T}}$, either $\mathcal{T} \models (k_1 \equiv k_2)$ or $\mathcal{T} \models (k_1 \not\equiv k_2)$.

Lemma 6.16 *Consider the class $\mathbf{T}$ defined in Example 6.15. If $\mathcal{T}_i$ is a theory in $\mathbf{T}$ of signature Σ_i and Σ a finite set of constant symbols in Σ_i, then $\mathcal{T}_i$ is Σ-stable over $Res(Qff^{\Sigma_i}, \Sigma)$.*

Proof. Let $\kappa := Card(\Sigma_i)$. It is enough to notice, first, that $\mathcal{A}^{\Sigma}$ is free in the Σ-variety of $\mathcal{T}_i$ over a basis of cardinality κ, for every $\mathcal{A} \in Mod(\mathcal{T}_i)$ of cardinality κ; second, that $Res(Qff^{\Sigma_i}, \Sigma) \subseteq Qff^{\Sigma_i}$ and so, by the stable-infiniteness of $\mathcal{T}_i$, every Σ-restricted formula satisfiable in $\mathcal{T}_i$ is satisfiable in a model of $\mathcal{T}_i$ of cardinality κ. $\square$

Proposition 6.17 *Let $\mathbf{T}$ be the class defined in Example 6.15 and let $\mathcal{T}_1, \mathcal{T}_2 \in \mathbf{T}$ have respective signatures Σ_1 and Σ_2 such that $\Sigma := \Sigma_1 \cap \Sigma_2$ is a finite set of constant symbols. If $\mathcal{T}_1 \models (k_1 \equiv k_2)$ iff $\mathcal{T}_2 \models (k_1 \equiv k_2)$ for all $k_1, k_2 \in \Sigma$, then $\mathcal{T}_1$ and $\mathcal{T}_2$ are totally N-O-combinable over Qff.*

Proof. By Lemma 6.16, $\mathcal{T}_i$ is Σ-stable over $Res(\mathit{Qff}^{\Sigma_i}, \Sigma)$ for $i = 1, 2$. It is immediate that $\mathcal{T}_1$ and $\mathcal{T}_2$ have the same Σ-atomic theory E_0. It follows by Theor. 6.9(1) then that $\mathcal{T}_1$ and $\mathcal{T}_2$ are partially N-O-combinable over Qff. To see that they are totally N-O-combinable, it is enough to notice that, since Σ is a set of constant symbols, *every* model of $\mathcal{T}_1 \cup \mathcal{T}_2$ is Σ-generated by its Σ-individuals, which satisfies Cond. 4.2. $\square$

This result states in essence that the Nelson-Oppen method is trivially extensible to theories sharing constants—provided that the theories are complete "over" these constants and identify them in the same way.

Example 6.18 Given a finite signature Σ containing at least one function symbol of non-zero arity, let $\mathbf{T}_\Sigma$ be the class of all universal theories $\mathcal{T}$ such that $\Sigma \subseteq (\Sigma_\mathcal{T})^{\mathrm{F}}$ and $\mathcal{T}^\Sigma = \mathcal{FT}^\Sigma$ (see Sec. 2.2.4).

In essence, all the automated deduction systems based on the resolution rule of inference with syntactic unification (including pure Prolog, but not Datalog) operate on theories $\mathcal{T}$ of this sort where $\Sigma = (\Sigma_\mathcal{T})^{\mathrm{F}}$.

Lemma 6.19 *Consider the class $\mathbf{T}_\Sigma$ defined in Example 6.18 and the class $\mathcal{L}$ of universal formulas. If $\mathcal{T}_i$ is a theory in $\mathbf{T}_\Sigma$ with signature Σ_i such that $\Sigma = {\Sigma_i}^{\mathrm{F}}$, then $\mathcal{T}_i$ is Σ-stable over $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$.*

Proof. First notice that since $\mathcal{T}_i^\Sigma = \mathcal{FT}^\Sigma$ and Σ contains at least one function symbol of non-zero arity, all models of $\mathcal{T}_i$ are infinite. Moreover, since $\Sigma = {\Sigma_i}^{\mathrm{F}}$, every set of Σ-generators for a model $\mathcal{A}$ of $\mathcal{T}_i$ is also a set of generators for $\mathcal{A}$. Now, suppose that $\varphi(\tilde{v}) \in Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ is satisfied in a model $\mathcal{B}$ of $\mathcal{T}_i$ by some tuple $\tilde{b}$. Since $\mathcal{B}$ is infinite, we can assume without loss of generality that it is uncountable. By Lemma 2.19 then, there is a countably-infinite $X \subseteq B$ that is Σ-independent in $\mathcal{B}$ and Σ-generates $\tilde{b}$. Let $\mathcal{A} := \langle X \rangle_\mathcal{B}$. By construction, X is a non-redundant set of generators for $\mathcal{A}$ and $\tilde{b}$ is in A. Observing that φ is equivalent to a universal formula, we can conclude by Lemma 2.13 and Lemma 2.14 that $\mathcal{A}$ as well is a model of $\mathcal{T}_i$ that satisfies φ. Recalling that $\mathcal{T}_i^\Sigma = \mathcal{FT}^\Sigma$, we know by Prop. 2.36 and Cor. 2.35 that $\mathcal{A}^\Sigma$ is an absolutely free algebra. Since the Σ-variety of $\mathcal{T}_i$ coincides with the Σ-variety of the empty theory, it follows that $\mathcal{A}^\Sigma$ is free over X in the Σ-variety of $\mathcal{T}_i$.

In conclusion, we have shown that an arbitrary formula $\varphi \in Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ satisfiable in $\mathcal{T}_i$, is also satisfiable is a model of $\mathcal{T}_i$ whose Σ-reduct is free in the Σ-variety of $\mathcal{T}_i$ over a basis of cardinality $\omega = Card(\Sigma_i)$, which proves the claim. $\square$

Proposition 6.20 *Let $\mathbf{T}_\Sigma$ be the class defined in Example 6.18 and let $\mathcal{T}_1, \mathcal{T}_2 \in \mathbf{T}_\Sigma$ have respective signatures Σ_1 and Σ_2 such that $\Sigma = \Sigma_1 \cap \Sigma_2 = {\Sigma_1}^{\mathrm{F}} = {\Sigma_2}^{\mathrm{F}}$. If $\mathcal{L}$ is the class of universal formulas, then $\mathcal{T}_1$ and $\mathcal{T}_2$ are totally N-O-combinable over $\mathcal{L}$.*

Proof. Obviously, both $\mathcal{T}_1$ and $\mathcal{T}_2$ have the same Σ-atomic theory, which is empty. Since $\mathcal{T}_i$ is Σ-stable over $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ for $i = 1, 2$, as shown above, we can conclude by Theor. 6.9(1) that $\mathcal{T}_1$ and $\mathcal{T}_2$ are partially N-O-combinable over $\mathcal{L}$.

Now, by Lemma 2.7, it is easy to see that $(\mathcal{T}_1 \cup \mathcal{T}_2)^{\Sigma} = \mathcal{FT}^{\Sigma}$. It follows that $\mathcal{T}_1 \cup \mathcal{T}_2 \in \mathbf{T}_{\Sigma}$ as well, which entails that $\mathcal{T}_1 \cup \mathcal{T}_2$ is Σ-stable over the class $Res(\mathcal{L}^{\Sigma_1 \cup \Sigma_2}, \Sigma)$ and so over its subclass $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$. Observing that the Σ-atomic theory of $\mathcal{T}_1 \cup \mathcal{T}_2$ is also empty and thus definitely collapse-free, we obtain by Theor. 6.9(2) that $\mathcal{T}_1$ and $\mathcal{T}_2$ are totally N-O-combinable over $\mathcal{L}$. $\square$

The above result requires the two component theories to have *exactly* the same set of function symbols. We can achieve a little more generality for free, however, if $\mathcal{T}_1$ and $\mathcal{T}_2$ do not share predicate symbols and are such that $\mathcal{T}_i^{\Delta_i} = \mathcal{FT}^{\Delta_i}$ for $i = 1, 2$ where $\Delta_i := \Sigma_i^{\mathrm{F}}$. In that case in fact, we can consider $\mathcal{T}_i' := \mathcal{T}_i \cup \mathcal{FT}^{\Delta_1 \cup \Delta_2}$ for $i = 1, 2$ as the component theories and $\mathcal{T}_1 \cup \mathcal{T}_2 \cup \mathcal{FT}^{\Delta_1 \cup \Delta_2}$ as the combined theory. As long has the pure parts of the input problem have symbols exclusively from Σ_1 or Σ_2, we can use the satisfiability procedures for $\mathcal{T}_1$ and $\mathcal{T}_2$ to generate a satisfiability procedure for $\mathcal{T}_1 \cup \mathcal{T}_2 \cup \mathcal{FT}^{\Delta_1 \cup \Delta_2}$. We must point out, however, that the union theory we consider here is larger than $\mathcal{T}_1 \cup \mathcal{T}_2$.

Example 6.21 Let $\mathbf{T}_{\Sigma}$ be the class of all collapse-free Horn clause theories $\mathcal{T}$ such that $\Sigma \subseteq \Sigma_{\mathcal{T}}$ and every $f \in (\Sigma_{\mathcal{T}} \setminus \Sigma)^{\mathrm{F}}$ is Σ-definable in $\mathcal{T}$.

Notice that every Ω-atomic theory is a Horn clause theory.

Lemma 6.22 *Consider the class $\mathbf{T}_{\Sigma}$ defined in Example 6.21 and the class $\mathcal{L}$ of the definite goals*[35]*. If $\mathcal{T}_i$ is an Σ_i-theory in $\mathbf{T}_{\Sigma}$, then $\mathcal{T}_i$ is Σ-stable over $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$.*

Proof. Since $\mathcal{T}_i$ is non-trivial for being collapse-free, it can be shown to have a model $\mathcal{A}$ that is free in $Mod(\mathcal{T}_i)$ over a basis X of cardinality $Card(\Sigma_i)$ (see [Hod93a] for instance). In particular, by Lemma 2.23, $\mathcal{A}$ is free over X in the Σ_i-variety of $\mathcal{T}_i$. By Prop. 2.29 and Prop. 2.27 then, $\mathcal{A}^{\Sigma}$ is free over X in the Σ-variety of $\mathcal{T}_i$. We show below that every formula of $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ unsatisfiable in $\mathcal{A}$ is unsatisfiable in $\mathcal{T}_i$. From this, it will follow that $\mathcal{T}_i$ is Σ-stable over $\mathcal{L}$.

Assume that $\psi \in Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ is unsatisfiable in $\mathcal{A}$. Recall that ψ has the form $\varphi(\tilde{u}) \wedge \varepsilon_{\neq}(\tilde{v}) \wedge \varepsilon_{\neq}^{\Sigma}(\tilde{v})$ where φ is a definite goal and $\tilde{v} \subseteq \tilde{u}$. Let α be a valuation assigning each variable of $\tilde{u}$ to a distinct generator of $\mathcal{A}$. Since $\mathcal{T}_i$ is collapse-free, by Prop. 2.22 every such generator is Σ-isolated and so $(\mathcal{A}, \alpha) \models \varepsilon_{\neq}(\tilde{v}) \wedge \varepsilon_{\neq}^{\Sigma}(\tilde{v})$. Since ψ is unsatisfiable in $\mathcal{A}$, we obtain that $(\mathcal{A}, \alpha) \models \neg\varphi(\tilde{u})$. Now, $\neg\varphi$ is equivalent to a conjunction of atoms. It follows from an immediate corollary of Prop. 2.21, that $\mathcal{T}_i \models \tilde{\forall}\, \neg\varphi$, which entails that φ, and so ψ, is unsatisfiable in $\mathcal{T}_i$. $\square$

[35] A definite goal is a disjunction of negated atoms.

Proposition 6.23 *Let $\mathbf{T}_\Sigma$ be the class defined in Example 6.21 and let $\mathcal{T}_1, \mathcal{T}_2 \in \mathbf{T}$ with respective signatures Σ_1 and Σ_2 and same Σ-atomic theory E_0. If $\mathcal{L}$ is the class of the definite goals, then $\mathcal{T}_1$ and $\mathcal{T}_2$ are totally N-O-combinable over $\mathcal{L}$.*

Proof. Since each $\mathcal{T}_i$ is Σ-stable over $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ and $\mathcal{T}_1$ and $\mathcal{T}_2$ have the same Σ-atomic theory, we can conclude by Theor. 6.9(1) that $\mathcal{T}_1$ and $\mathcal{T}_2$ are partially N-O-combinable over $\mathcal{L}$.

By Lemma 6.8, E_0 is also the Σ-atomic theory of $\mathcal{T}_1 \cup \mathcal{T}_2$. This entails that $\mathcal{T}_1 \cup \mathcal{T}_2$ is a collapse-free Horn clause theory and that every function symbol in $(\Sigma_1 \cup \Sigma_2) \setminus \Sigma$ is Σ-definable in $\mathcal{T}_1 \cup \mathcal{T}_2$. It follows that $\mathcal{T}_1 \cup \mathcal{T}_2 \in \mathbf{T}_\Sigma$ as well and so is Σ-stable over $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ for being Σ-stable over $Res(\mathcal{L}^{\Sigma_1 \cup \Sigma_2}, \Sigma)$ by the previous lemma. Since the Σ-variety of $\mathcal{T}_1 \cup \mathcal{T}_2$ is obviously collapse-free, we obtain by Theor. 6.9(2) that $\mathcal{T}_1$ and $\mathcal{T}_2$ are totally N-O-combinable over $\mathcal{L}$. □

Although the above result looks quite interesting, its relevance is perhaps purely academic because of the severe rigidity of the given $\mathcal{L}$. In fact, $\mathcal{L}$ is not even closed under conjunction, a minimal requirement for about every constraint language.

Example 6.24 Let Σ be a finite signature and $\mathcal{A}^\Sigma$ be any free Σ-structure with a countably-infinite basis. Let $\mathbf{T}_{\mathcal{A}^\Sigma}$ be the class of all the theories $\mathcal{T}$ such that $\mathcal{T}$ is the complete theory of an expansion of $\mathcal{A}^\Sigma$ to a larger signature.

Lemma 6.25 *Consider the class $\mathbf{T}_{\mathcal{A}^\Sigma}$ defined in Example 6.24 and the class $\mathcal{L}$ of first-order formulas. If $\mathcal{T}_i$ is a theory in $\mathbf{T}_{\mathcal{A}^\Sigma}$ with signature Σ_i then $\mathcal{T}_i$ is Σ-stable over $\mathcal{L}^{\Sigma_i}$.*

Proof. Let $\mathcal{A}^{\Sigma_i}$ be the structure axiomatized by $\mathcal{T}_i$. By definition of $\mathcal{T}_i$, a Σ_i-sentence is entailed by $\mathcal{T}_i$ exactly when it is valid in $\mathcal{A}^{\Sigma_i}$. This entails, first, that a Σ_i-formula is satisfiable in $\mathcal{T}_i$ iff it is satisfiable in $\mathcal{A}^{\Sigma_i}$; second, that $\mathcal{T}_i$ and $\mathcal{A}^\Sigma$ have the same Σ-variety. Then, the claim follows from the fact that $\mathcal{A}^\Sigma$ is free over some countably-infinite set X and so, by Cor. 2.24, is free over X in its Σ-variety. □

Proposition 6.26 *Let $\mathbf{T}_{\mathcal{A}^\Sigma}$ be the class defined in Example 6.24 and assume that the Σ-atomic theory E_0 of $\mathcal{A}^\Sigma$ is collapse-free. Let $\mathcal{T}_1, \mathcal{T}_2 \in \mathbf{T}_{\mathcal{A}^\Sigma}$ have respective signatures Σ_1 and Σ_2 such that $\Sigma_1 \cap \Sigma_2 = \Sigma$. If $\mathcal{L}$ is the class of first-order formulas, then $\mathcal{T}_1$ and $\mathcal{T}_2$ are partially N-O-combinable over $\mathcal{L}$.*

Proof. For $i = 1, 2$, let $\mathcal{A}^{\Sigma_i}$ be the structure axiomatized by $\mathcal{T}_i$. By construction of $\mathbf{T}_{\mathcal{A}^\Sigma}$, E_0 is the Σ-atomic theory of both $\mathcal{T}_1$ and $\mathcal{T}_2$. By the previous lemma and Theor. 6.9, we can show again that $\mathcal{T}_1$ and $\mathcal{T}_2$ are partially N-O-combinable over $\mathcal{L}$. □

Complete theories are interesting because the satisfiability problem and the entailment problem (of existential closures) coincide for them. A formula is satisfiable in a complete theory iff it is satisfiable in every model of the theory. Now, in essence, the result above says that two complete theories $\mathcal{T}_1$ and $\mathcal{T}_2$ are N-O-combinable if they both have a model whose reduct to their shared symbols is free over countably-many generators in the same collapse-free variety. In that case then, if the entailment problem is decidable for both $\mathcal{T}_1$ and $\mathcal{T}_2$, we can use the combination method to obtain a procedure for the satisfiability problem in $\mathcal{T}_1 \cup \mathcal{T}_2$. Unfortunately, with the above result we can only show the soundness of the combination method, not its completeness. Another thing to notice is that the combination method will yield a procedure for the satisfiability problem in $\mathcal{T}_1 \cup \mathcal{T}_2$ but not one for the entailment problem. The reason is simply that in general the union of two complete theories is not a complete theory.

An interesting generalization of the previous example can be given in terms of the more general notion of Σ-stability. Given an Ω-structure $\mathcal{A}$, a finite subset Σ of Ω, and a class $\mathcal{L}$ of formulas, let us call *Σ-restricted theory of $\mathcal{A}$* the set

$$Th(\mathcal{A})(\mathcal{L},\Sigma) \quad := \quad \{\tilde{\exists}\, \psi \mid \psi \in Res(\mathcal{L}^{\Omega},\Sigma), \mathcal{A} \models \tilde{\exists}\, \psi\}.$$

The theory above is not complete in general but is such that a Σ-restricted formula of $\mathcal{L}$ is satisfiable in it iff it is satisfiable in $\mathcal{A}$.[36] If the Σ-reduct of $\mathcal{A}$ is free in a Σ-variety $\mathbf{K}$, we can easily show that $Th(\mathcal{A})(\mathcal{L},\Sigma)$ is Σ-stable wrt $\mathbf{K}$ over $Res(\mathcal{L}^{\Omega},\Sigma)$. Now, if $\mathcal{B}$ is another structure whose signature includes Σ and whose Σ-reduct is isomorphic to $\mathcal{A}^{\Sigma}$, it should be possible to show, although not directly by Theor. 6.9, that $Th(\mathcal{A})(\mathcal{L},\Sigma)$ and $Th(\mathcal{B})(\mathcal{L},\Sigma)$ are (partially) N-O-combinable.

The following is another variation of Example 6.24 where the whole structure is now free, not just its Σ-reduct.

Example 6.27 Where Σ is a finite signature, let $\mathbf{T}_\Sigma$ be the class of all the theories $\mathcal{T}$ such that $\mathcal{T}$ is the complete theory of some free structure with a countably-infinite basis and Σ is a set of constructors for $\mathcal{T}$ modulo $\mathcal{T}^{\Sigma}_{\mathrm{At}}$.

Lemma 6.28 *Consider the class $\mathbf{T}_\Sigma$ defined in Example 6.27 and the class $\mathcal{L}$ of first-order formulas. If $\mathcal{T}_i$ is a theory in $\mathbf{T}_\Sigma$ with signature Σ_i then $\mathcal{T}_i$ is Σ-stable over $\mathcal{L}^{\Sigma_i}$.*

Proof. Let Σ_i be the signature of $\mathcal{T}_i$ and $\mathcal{A}_i$ the free structure axiomatized by $\mathcal{T}_i$. As before, a Σ_i-formula is satisfiable in $\mathcal{T}_i$ iff it is satisfiable in $\mathcal{A}_i$. All we need to show then is that ${\mathcal{A}_i}^{\Sigma}$ is free in the Σ-variety of $\mathcal{T}_i$ over a countably-infinite basis.

Now, $\mathcal{A}_i$ and $\mathcal{T}_i$ have the same Σ_i-atomic theory E by construction of $\mathcal{T}_i$, and $\mathcal{A}_i$ is free in $Mod(E)$ (over some countably-infinite set) by Cor. 2.24. Since Σ is a set of constructors for E modulo E^{Σ}_{At} by assumption, we can conclude by Prop. 2.32 that ${\mathcal{A}_i}^{\Sigma}$ is free over some countably-infinite set Y in the Σ-variety of E. The claim then follows from the trivial fact that E and $\mathcal{T}_i$ have the same Σ-variety. □

[36] Borrowing the terminology from a similar concept in Constraint Logic Programming [JM94], we could say that $Th(\mathcal{A})(\mathcal{L},\Sigma)$ is *satisfaction complete wrt* $Res(\mathcal{L}^{\Omega},\Sigma)$.

Proposition 6.29 *Let $\mathbf{T}_\Sigma$ be the class defined in Example 6.27. Let $\mathcal{T}_1, \mathcal{T}_2 \in \mathbf{T}_\Sigma$ have respective signatures Σ_1 and Σ_2 such that $\Sigma_1 \cap \Sigma_2 = \Sigma$ and $\mathcal{T}_1$ and $\mathcal{T}_2$ have the same Σ-atomic theory E_0. If $\mathcal{L}$ is the class of first-order formulas, then $\mathcal{T}_1$ and $\mathcal{T}_2$ are partially N-O-combinable over $\mathcal{L}$.*

Proof. Since Σ is a set of constructors for $\mathcal{T}_i$ it is easy to see from Def. 2.30 that E_0 is collapse-free. Given this, the proof is analogous to that of Prop. 6.26. □

Again we only have a partial combinability result here. However, thanks to Prop. 6.11, we can easily show the total N-O-combinability of $\mathcal{T}_1$ and $\mathcal{T}_2$ over the language $TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$.

7 Applications to Theories with Functional Signature

In this section, we show some of the decidability results that can be obtained from our combination method in the case of theories with a functional signature. For this purpose we will use concepts and terminology taken from Unification Theory and Rewrite System. For space constraints, we must forget their definition and instead refer the reader to [Sie89, Wec92].

To start with, recall that equational theories of signature Σ are a special case of Σ-atomic theories in which Σ contains only function symbols.[37] In analogy with the previous terminology then we will sometimes call these theories Σ-*equational theories*. Now, it so happens that most decision problems for a given (non-trivial) equational theory E are in fact decision problems in the corresponding free algebra $\mathcal{F}_\omega(E)$:

- E-*unifiability*, the satisfiability of a conjunction of equations in $\mathcal{F}_\omega(E)$;
- E-*disunifiability*, the satisfiability of a conjunction of equations and disequations in $\mathcal{F}_\omega(E)$;
- E-*equality*, the validity of an equation in $\mathcal{F}_\omega(E)$.

Where E_1 and E_2 are two equational theories, let $\mathcal{T}_{E_i}$ denote the complete theory of $\mathcal{F}_\omega(E_i)$ for $i = 1, 2$. In the following we will show how to use our combination method to obtain decidability results for the satisfiability in $\mathcal{T}_{E_1} \cup \mathcal{T}_{E_2}$ of quantifier-free formulas with a total Σ-restriction where Σ is the intersection of E_1's and E_2's signatures.

First let us point our that, satisfiability in $\mathcal{T}_E$ of quantifier-free formulas is reducible to disunifiability in $\mathcal{F}_\omega(E)$. A non-deterministic decision algorithm for this problem can be constructed if E-equality is decidable and E-unification is finitary. In fact, suppose we are interested in the satisfiability of a conjunction φ of equations and disequations. Then we can do the following.

[37] In such a context, Σ-structures are simply Σ-algebras.

1. Let
   E_φ be the set of φ's equations and
   D_φ be the set of φ's disequations.
2. Succeed if there is a most general E-unifier μ of E_φ
   such that $s \neq_E t$ for each $s \not\equiv t \in D_\varphi\mu$.
   Fail otherwise.

Given the above, the following result is easy to prove.

Proposition 7.1 *Let E be a non-trivial Σ-equational theory. If E-equality is decidable and there exists a finitary E-unification algorithm, then satisfiability in $\mathcal{T}_E$ of formulas in Qff^Σ is decidable.*

In the following, we consider the special case where a Ω-equational theory E is split into a Σ-equational theory, with $\Sigma \subseteq \Omega$, plus an Ω-equational theory E_R generated by a rewrite relation $\rightarrow_R$ on $T(\Omega, V)$ [JK86]:

$$E_R \quad := \quad \{\tilde{\forall}\, s \equiv t \mid s, t \in T(\Omega, V), s \rightarrow_R t\}.$$

Let $\rightarrow_{R!}$ be the normalizing relation of $\rightarrow_R$ defined by:

$$t \rightarrow_{R!} t' \quad \text{iff} \quad t \rightarrow_R^* t' \text{ and } t' \text{ is } R\text{-irreducible}^{38},$$

and assume that it is *weakly normalizing*, that is, for any Ω-term t, there exists at least one t' such that $t \rightarrow_{R!} t'$. The term t' is called a R-normal form of t. Then, consider a collapse-free Σ-equational theory E_0, let $E := E_0 \cup E_R$, and assume that

$$s =_E t,\ s \rightarrow_{R!} s',\ \text{and } t \rightarrow_{R!} t' \quad \text{implies} \quad s' =_{E_0} t' \tag{6}$$

for all $s, t, s', t' \in T(\Omega, V)$ where, by a slight abuse of notation, we let $=_{E_0}$ denote the equivalence relation on $T(\Omega, V)$ (not just $T(\Sigma, V)$) generated by E_0. According to (6), any two normal forms of the same term t are E_0-equivalent. Where $t{\downarrow_R}$ denotes an arbitrary R-normal form of t, also assume that

$$f(t_1, \ldots, t_n){\downarrow_R} \quad =_{E_0} \quad f(t_1{\downarrow_R}, \ldots, t_n{\downarrow_R}) \tag{7}$$

for all $f \in \Sigma$. Under these assumptions we can then prove that Σ is a set of constructors for E.

Proposition 7.2 *Σ is a set of constructors for E.*

[38] That is, $t' \rightarrow_R t''$ for no $t'' \in T(\Omega, V)$.

Proof. We show that all three conditions of Def. 2.30 are satisfied by using well-known facts in the context of combining decision algorithms for the word-problem in a union of (collapse-free) signature-disjoint equational theories [SS89, Nip91, Rin96a, BT97]. Here, the union of interest is E_0 plus the empty collapse-free equational theory E_1 generated by function symbols in $\Omega \setminus \Sigma$.

(*1.*) Suppose that $V \not\subseteq G_E(\Sigma, V)$. Then, there must be a $v \in V$, an $f \in \Sigma$, and a $\tilde{t}$ in $T(\Omega, V)$ such that $v =_E f(\tilde{t})$. By (7) and the fact that $v{\downarrow_R} = v$, we can conclude that $v =_{E_0} f(\tilde{t}{\downarrow_R})$. This is impossible as the union E_0 plus the $(\Omega \setminus \Sigma)$-theory E_1 is still collapse-free.

(*2.*) It is easy to see that $t =_E t{\downarrow_R}$. We show that $t{\downarrow_R} \in ST_E(\Sigma, V)$. If we assume the contrary, there must be an $f \in \Sigma$, a $\tilde{t'}$ in $T(\Omega, V)$, and a subterm r of $t{\downarrow_R}$ with $r(\epsilon) \notin \Sigma$ such that $r =_E f(\tilde{t'})$. Since $t{\downarrow_R}$ is irreducible by construction we know that $r{\downarrow_R} =_{E_0} r$ which entails by (6) and (7) above that $r =_{E_0} f(\tilde{t}'{\downarrow_R})$. This is impossible since two terms with root symbols in two signature-disjoint equational theories E_0 and E_1 cannot be $E_0 \cup E_1$-equal, provided that E_0 and E_1 are collapse-free.

(*3.*) Let $s\rho, s'\rho \in ST_E(\Sigma, V)$ such that $u\rho \neq_E v\rho$ for every distinct $u, v \in \tilde{v}$ where $\tilde{v} := \mathcal{D}om(\rho)$. Recall that ρ is idempotent, which is to say that $\tilde{v} \cap \mathcal{R}an(\rho) = \emptyset$. We show that $s\rho =_E s'\rho$ iff $s =_{E_0} s'$.

If $s =_{E_0} s'$ then $s =_E s'$ by construction of E, which immediately entails $s\rho =_E s'\rho$ for any Ω-instantiation ρ. If $s\rho =_E s'\rho$, we know by (6) above that

$$s\rho{\downarrow_R} =_{E_0} s'\rho{\downarrow_R} .$$

Let ρ' be a substitution such that $v\rho' := v\rho{\downarrow_R}$ for each $v \in \tilde{v}$. By a inductive argument again based on (7) above, it is possible to show that

$$s\rho' =_{E_0} s\rho{\downarrow_R} \quad \text{and} \quad s'\rho{\downarrow_R} =_{E_0} s'\rho',$$

which entails that $s\rho' =_{E_0} s'\rho'$. Now, assuming with no loss of generality that $\mathcal{V}ar(v\rho{\downarrow_R}) \subseteq \mathcal{V}ar(v\rho)$ for each $v \in \tilde{v}$, it is not difficult to see that ρ' satisfies the property $\forall u, v \in \mathcal{D}om(\rho')$, $u = v \Leftrightarrow u\rho' =_{E_0} v\rho'$. Thus we can apply in a straightforward way the combination algorithm for the word-problem in the union of collapse-free signature-disjoint theories E_0 and E_1, and it follows that $s =_{E_0} s'$. □

Consider now an Ω-structure $\mathcal{A}_i$ free in $Mod(E)$ over the countably-infinite set X. We know from Prop. 2.32 that the $\mathcal{A}^\Sigma$ is free in $Mod(E_0)$ over $Is(\mathcal{A}^\Sigma)$. We also know from Prop. 2.31 that

$$Is(\mathcal{A}^\Sigma) \quad = \quad \{r^{(\mathcal{A},\alpha)} \mid r \in G_E(\Sigma, V)\}$$

for every bijective valuation α of V onto X. We show below that $Is({\mathcal{A}_i}^\Sigma)$ can be given a syntactic characterization in terms of the set of Ω-terms whose reduced form under ${\downarrow_R}$ starts with a non-Σ symbol. More precisely, we can show the following.

Proposition 7.3 *Let α be a bijective valuation of V onto X. Then,*

$$Is(\mathcal{A}^{\Sigma}) \quad = \quad \{r^{(\mathcal{A},\alpha)} \mid r \in T(\Omega, V),\ r{\downarrow}_R(\epsilon) \notin \Sigma\}.$$

Proof. By the above, it is enough to show that for every $r \in T(\Omega, V)$, $r{\downarrow}_R(\epsilon) \notin \Sigma$ iff $r \in G_E(\Sigma, V)$.

($\Rightarrow$) Assume that $r{\downarrow}_R(\epsilon) \notin \Sigma$ but $r \notin G_E(\Sigma, V)$. Then, there is an $f \in \Sigma$ and a $\tilde{t}$ in $T(\Omega, V)$ such that $r =_E f(\tilde{t})$. From (6) and (7) above it easy to see then that $r{\downarrow}_R =_{E_0} f(\tilde{t}{\downarrow}_R)$. We have already seen that such an E_0-equality between terms rooted respectively by function symbols in $\Omega\backslash\Sigma$ and Σ contradicts the assumption that E_0 is collapse-free.

($\Leftarrow$) Assume that $r{\downarrow}_R(\epsilon) \in \Sigma$. Then, since $r =_E r{\downarrow}_R$ trivially, we know that there is an $f \in \Sigma$ and a $\tilde{t}$ in $T(\Omega, V)$ such that $r =_E f(\tilde{t})$, which immediately implies that $r \notin G_E(\Sigma, V)$. $\square$

Where $\mathcal{T}_E$ is again defined as the complete theory of $\mathcal{F}_\omega(E)$, also denoted by $\mathcal{A}$ in Prop. 7.3, we can then decide the satisfiability in $\mathcal{T}_E$ of a formula

$$\varphi(\tilde{v}) \wedge \varepsilon_{\neq}(\tilde{v}) \wedge \varepsilon_{\neq}^{\Sigma}(\tilde{v}) \quad \in \quad TRes(Qff^{\Omega}, \Sigma)$$

essentially by adding to the procedure seen for Prop. 7.1 a check on the root symbols of the solutions computed by E-unification.

Proposition 7.4 *The satisfiability in $\mathcal{T}_E$ of formulas in $TRes(Qff^{\Omega}, \Sigma)$ is decidable if*

- *there is a computable function that, for any term t, returns the R-normal form $t{\downarrow}_R$ of t,*
- *E_0-equality is decidable,*
- *there exists a finitary E-unification algorithm.*

Proof. Recall that for a totally Σ-restricted formula to be satisfiable, its variables must be assigned to distinct Σ-isolated individuals. By the above then, it is enough to verify that there is a substitution that satisfies φ and (a) maps no variable of φ to a term whose R-normal form starts with a symbol in Σ, (b) maps no two variables to E-equivalent terms. More precisely, we can do the following.

1. Let

   E_φ be the set of φ's equations and

   D_φ be the set of φ's disequations.

2. Succeed if there is a $\mu \in \mathrm{SUB}(\tilde{v})$ such that

(a) μ is a most general E-unifier of E_φ where $v\mu = v\mu{\downarrow_R}$ for all $v \in \tilde{v}$

(b) $s{\downarrow_R} \neq_{E_0} t{\downarrow_R}$ for all $s \not\equiv t \in D_\varphi\mu$,

(c) $v\mu(\epsilon) \notin \Sigma$ for all $v \in \tilde{v}$,

(d) $u\mu \neq_{E_0} v\mu$ for all distinct $u, v \in \tilde{v}$.

Fail otherwise.

Conditions (a) and (b) in Step 2 above are met if and only if φ is satisfiable in $\mathcal{F}_\omega(E)$ and so in $\mathcal{T}_E$. Conditions (c) and (d) are met if and only if, among φ's possible solutions, there is one that assigns φ's variables to distinct Σ-isolated individuals. When conditions (a) and (b) are met but (c) and (d) are not, then every normalized instance of μ satisfies φ in $\mathcal{F}_\omega(E)$ but does not satisfy the Σ-restriction. In this case, we can conclude that the whole Σ-restricted formula is unsatisfiable in $\mathcal{F}_\omega(E)$ and so in $\mathcal{T}_E$. □

We are now ready to combine two theories $\mathcal{T}_{E_1}$ and $\mathcal{T}_{E_2}$ respectively built over the signatures Σ_1 and Σ_2 such that $\Sigma_1 \cap \Sigma_2 = \Sigma$. For $i = 1, 2$ let $\rightarrow_{R_i}$ be a rewrite relation on $T(\Sigma_i, V)$, E_{R_i} the corresponding Σ_i-equational theory, and $E_i := E_{R_i} \cup E_0$ where E_0 is defined as before. Also assume that for $i = 1, 2$, $\rightarrow_{R_i}$ and E_0 satisfy the same properties satisfied by $\rightarrow_R$ and E_0 above.

Where $\mathcal{T}_{E_i}$ is defined as the complete theory of $\mathcal{F}_\omega(E_i)$ for $i = 1, 2$, we then obtain the following result.

Proposition 7.5 *Assume that Σ is finite. The satisfiability in $\mathcal{T}_{E_1} \cup \mathcal{T}_{E_2}$ of formulas in*
$TRes(Qff^{\Sigma_1} \otimes Qff^{\Sigma_2}, \Sigma)$ is decidable if for $i = 1, 2$

- *there is a computable function that, for any term t, returns the R_i-normal form $t{\downarrow_{R_i}}$ of t,*
- *E_0-equality is decidable,*
- *there exists a finitary E_i-unification algorithm.*

Proof. We know from Prop. 6.29 that $\mathcal{T}_{E_1}$ and $\mathcal{T}_{E_2}$ are totally N-O-combinable over the class $TRes(Qff^{\Sigma_1} \otimes Qff^{\Sigma_2}, \Sigma)$ provided that Σ is finite. From Prop. 7.4 we know that the satisfiability in $\mathcal{T}_{E_i}$ of formulas from $TRes(Qff^{\Sigma_i}, \Sigma)$ is decidable. Therefore we can apply our combination method and obtain a procedure that converges on all satisfiable inputs. Recall that our method does not in general yield a decision procedure because of the infinitary non-determinism of the instantiation step. Now, for problems in $TRes(Qff^{\Sigma_1} \otimes Qff^{\Sigma_2}, \Sigma)$ the only instantiation that preserves the satisfiability of the problem is the empty instantiation Hence, in this case we can skip the instantiation step altogether and obtain an always convergent combination procedure. The claim follows then immediately. □

Here is an example to which the above result applies.

Example 7.6 Consider two endomorphisms h_1 and h_2 over the same commutative symbol $*$. For $i = 1, 2$, let $\rightarrow_{R_i}$ be the rewrite relation generated by the rewrite system

$$R_i \quad := \quad \{h_i(x * y) \rightarrow h_i(x) * h_i(y)\}.$$

Let E_0 be $\{x * y = y * x\}$, E_{R_i} be the equational theory generated by $\rightarrow_{R_i}$, and E_i be $E_0 \cup E_{R_i}$. The relation $\rightarrow_{R_i}$ is obviously normalizing and is also convergent (modulo E_0), which means that $s =_{E_i} t$ iff $s\downarrow_{R_i} =_{E_0} t\downarrow_{R_i}$, for all terms s, t. Observing that

$$(s * t)\downarrow_{R_i} \quad = \quad (s\downarrow_{R_i}) * (t\downarrow_{R_i})$$

for all terms s, t, it is not difficult to see that $*$ is a constructor for E_i. The decidability of E_i-equality follows from the decidability of E_0-equality. Moreover, a finitary E_i-unification algorithm can be easily derived from the one known for E-unification where E is the unitary theory $\{h(x * y) = h(x) * h(y)\}$. Therefore, there exists an algorithm for deciding the satisfiability of totally restricted formulas in $\mathcal{T}_{E_i}$ and so Prop. 7.5 applies.

As a final note, we would like to stress that the above application of our combination techniques does not decide the disunifiability problem in $E_1 \cup E_2$ as one might be induced to believe. That problem is a satisfiability problem in the complete theory $\mathcal{T}_{E_1 \cup E_2}$ of $\mathcal{F}_\omega(E_1 \cup E_2)$, but *does not* corresponds to a satisfiability problem in the union $\mathcal{T}_{E_1} \cup \mathcal{T}_{E_2}$ of complete theories of $\mathcal{F}_\omega(E_1)$ and $\mathcal{F}_\omega(E_2)$, as in our case. One should be aware that $\mathcal{T}_{E_1 \cup E_2}$ and $\mathcal{T}_{E_1} \cup \mathcal{T}_{E_2}$ are generally different.

8 Conclusions

Constraint-based Reasoning is a promising and increasingly popular computational paradigm for computer programming and automated deduction. It combines the versatility of general-purpose reasoners with the high performance of specialized constraint solvers. Full-scale applicability of the paradigm is presently hindered by the difficulty of integrating several constraint solvers *modularly* into one general-purpose reasoner. Like every Engineering field, software production has benefited immensely from design techniques based on the use of modular components. It is now commonplace in many application domains to design and produce software artifacts by assembling pre-existing software modules with well-defined functionalities. This is not yet the case in Automated Deduction where any modularity technique has to be worked out at a model-theoretic level first. Combining different reasoners is a Mathematical Logic problem before being a Software Engineering one. In the case of constraint solvers for instance, where each solver decides constraint satisfiability with respect to a specific theory, to combine solvers first of all means to combine their theories.

Whatever the notion of theory combination, some theories simply cannot be combined because that would lead to an inconsistent theory or to an undecidable constraint satisfiability problem. This is typically the case when a combined theory is obtained as the union of two non-signature disjoint theories—the main focus of this paper. For the theories that can be combined, an ideal combination method is one that, given an input problem over the combined theory, is able to first break that problem into parts solvable separately by the single solvers, and then compose (conceptually) the solutions obtained from the solvers into a global solution for the original problem. General techniques for combining constraint solvers and prove their combination sound and complete are yet to be developed. The few combination methods devised to date have limited applicability because of their stringent requirements on either the input constraint language or the kind of constraint theories to combine.

8.1 Contributions

We have considered one of the existing combination methods, the Nelson-Oppen method, and tried to generalize it by lifting the restriction that the input formula be quantifier-free and that the combined theories share no non-logical symbols. To do this however, we first had to investigate the main model-theoretic issues involved in theory combination.

We have defined the concept fusion of two structures and shown in what sense it is a viable notion of model combination. We have also defined the concept of fusibility and shown how local satisfiability properties of arbitrary first-order constraints with respect to two fusible structures relates to satisfiability of conjunctive constraints in a fusion of the structures. We have then shown that, thanks to the close relation between fusion of structures and union of theories, it is possible to obtain combination results for constraint satisfiability with respect to theories and their unions from the corresponding combination results for constraint satisfiability with respect to the theories' models and their fusions. The model-theoretic conditions on the component theories that make the combination results possible are collected in the concept of N-O-combinability. We have shown that an extension of the Nelson-Oppen method can be applied in a sound a complete way to two N-O-combinable theories to produce a constraint satisfiability procedure for the union of the two theories. Finally, we have singled out some sufficient conditions for N-O-combinability by using the concept of Σ-stability, a natural extension of stable-infiniteness for the case of non disjoint unions of theories.

The results of this investigation have been useful not only in generalizing Nelson and Oppen's results, as shown in this paper, but also in providing a general theoretical framework for the combination of satisfiability problems. In fact, some work is already under way which uses the combination results presented here to propose a novel combination method for a different kind of combination problem: the combination of decision procedures for the word problem (initial results were published in [BT97]).

8.2 Further Research

We consider the research described in this paper a first step in providing theoretical bases for the combination of satisfiability procedures. More work needs to be done to both extend our results further and identify concrete cases from the Constrained Reasoning research to which such results can be applied. In particular, we think it will be necessary to

- refine the notion of Σ-stability introduced in Sect. 6.2 to obtain more modular combination results which can scale up to the combination of more than two theories;
- produce specific examples of Σ-stable theories of practical significance;
- investigate cases in which the combination method in Sect. 5 can yield a decision procedure, as opposed to just a semi-decision procedure;

In addition, we would like to compare our work with some recent research by Baader and Schulz [BS98], who have also developed a model-theoretic framework for the combination problem, but in the context of constraint entailment and in case of disjoint unions of theories. In particular, we would like first to verify whether quasi-free structures, introduced in [BS98] as an extension of free structures, can be used (as it seems) to define more general sufficient conditions for fusibility and hence Σ-stability; then relate fusions, our type of combination structure, with amalgamated product, the type of combination structure used in [BS98], and see if our results on Σ-stability and constructors can contribute to extending the Baader-Schulz method to the non-disjoint case.

Finally, we would like to recast the results presented here in the context of many-sorted (or better order-sorted [GM92]) logic. The case for using many-sorted logic in automated deduction, and computer science in general, is well understood.[39] The language of classical first-order logic is too *permissive* for constraint-based reasoning because it allows constraints one would consider ill-typed in the intended domain of application. For instance, a constraint such as $([1,2] + [1|Y] = 0)$[40] does not make sense if one is interested in the domain of lists of real numbers. Unfortunately, any single-sorted theory meant to axiomatize such domain will admit models in which the constraint above is perfectly meaningful, possibly even satisfiable! One of the merits of a sort-based approach in this case would be to reduce the number

[39] As a matter of fact, most of the automated reasoning/constraint programming systems developed today use sorted languages one way or another.

[40] Where we are using a Prolog-like syntax.

of unintended models of the constraint theory and eliminate unintuitive expressions from the constraint language. The need for sorts becomes possibly more pressing in a theory combination context: even if two theories $\mathcal{T}_1$ and $\mathcal{T}_2$ are adequately described using a single sort, their combination may not be. For instance, we could think of obtaining the theory of lists of real numbers by combining the theory of lists and the theory of real numbers. While each theory has a nice single-sorted axiomatization, their combination gives rise to pointless formulas like $[1,2]+[1|Y]=0$ above.

We believe that, in addition to more closely reflecting the current research in automated deduction and in declarative programming, the payoff of reformulating the work of this paper in a order-sorted context would be the possibility of obtaining completeness results for our combination method more easily. Intuitively, completeness is easier to achieve if we reduce both the size of the input language (by disallowing ill-sorted constraints) and the number of possible models of the combined theory (by disallowing models not conforming to the sort structure of the theory). Another advantage of using sorts would be the reduction of the non-determinism in both the instantiation and the identification step of the method, because variables would only be replaceable by terms or variables with a compatible sort.

The cost of an order-sorted reformulation, in addition to having to deal with a more cluttered notation and more complex proofs, is the necessity of adopting more elaborated forms of theory combination than simple set-theoretic union of signatures and axioms. For instance, at the very least it will be necessary to specify how the sorted signature of one component theory combines with the sorted signature of another to yield a *meaningful* signature for the combined theory.

Consider again the lists and reals example where $\mathcal{T}_1$ is an axiomatizable theory of (flat) lists with a sort list for lists and sort element for list elements, and $\mathcal{T}_2$ is an axiomatizable theory of the real numbers with the single sort real. We may as well be able to axiomatize their combination by taking the union of $\mathcal{T}_1$'s and $\mathcal{T}_2$'s axioms. However, we will also need to say that every real number can be a list element. Specifically, we will need to define a sort structure for the combined theory that not only includes those of the component theories but also expresses the fact that real is a subsort of element. In general terms, since the sort structure of an order-sorted theory is in essence a theory itself (stating which symbols are of what sort and how sorts are *included* in one another), to meaningfully combine two order-sorted theories it will also be necessary to specify how their *sort theories* combine.

8.3 Acknowledgments

We would like to thank Mehdi Harandi for his precious suggestions and Reza Ziaei and Hélène Kirchner for their careful proof-reading of some parts of the paper and some comments on how improve its readability. The first author was partially supported by grant DACA88-94-0014 from the US Army Construction Engineering Research Laboratories.

References

[B9̈4] Hans-Jürgen Bürckert. A resolution principle for constraint logics. *Artificial Intelligence*, 66:235–271, 1994.

[BFP92] Peter Baumgartner, Ulrich Furbach, and Uwe Petermann. A unified approach to theory reasoning. Research Report 15–92, Universität Koblenz-Landau, Koblenz, Germany, 1992. Fachberichte Informatik.

[Bou90] 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.

[Bou93] Alexandre Boudet. Combining unification algorithms. *Journal of Symbolic Computation*, 16(6):597–626, December 1993.

[BS92] Franz Baader and Klaus U. Schulz. Unification in the union of disjoint equational theories: Combining decision procedures. In D. Kapur, editor, *Proceedings of the 11th International Conference on Automated Deduction*, volume 607 of *Lecture Notes in Artificial Intelligence*, pages 50–65. Springer-Verlag, 1992.

[BS93] Franz Baader and Klaus U. Schulz. Combination techniques and decision problems for disunification. In C. Kirchner, editor, *Proceedings of the 5th International Conference on Rewriting Techniques and Applications*, volume 690 of *Lecture Notes in Computer Science*, pages 301–315. Springer-Verlag, 1993.

[BS95a] 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.

[BS95b] 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)*, volume 976 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, September 1995.

[BS96] Franz Baader and Klaus U. Schulz. Unification in the union of disjoint equational theories: Combining decision procedures. *Journal of Symbolic Computation*, 21(2):211–243, February 1996.

[BS98] Franz Baader and Klaus U. Schulz. Combination of constraint solvers for free and quasi-free structures. *Theoretical Computer Science*, 192:107–161, 1998.

[BT97] Franz Baader and Cesare Tinelli. A new approach for combining decision procedures for the word problem, and its connection to the Nelson- Oppen combination method. In W. McCune, editor, *Proceedings of the 14th International Conference on Automated Deduction (Townsville, Australia)*, volume 1249 of *Lecture Notes in Artificial Intelligence*, pages 19–33, 1997.

[CLS96] David Cyrluk, Patrick Lincoln, and Natarajan Shankar. On Shostak's decision procedure for combinations of theories. In M. A. McRobbie and J.K. Slaney, editors, *Proceedings of the 13th International Conference on Automated Deduction, (New Brunswick, NJ)*, volume 1104 of *Lecture Notes in Artificial Intelligence*, pages 463–477. Springer-Verlag, 1996.

[Com93] Hubert Comon. Constraints in term algebras (A short survey). In M. Nivat, C. Rattay, T. Rus, and G. Scollo, editors, *Proceedings of Algebraic Methodology and Software Technology*, pages 145–152, Enschede, Netherlands, June 1993.

[DKR94] 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.

[FP95] Andrea Formisano and Alberto Policriti. T-resolution: Refinements and model elimination. Technical Report UDMI/32/12/95, Dipartimento di Matematica e Informatica, Università di Udine, Udine, Italy, 1995.

[Fri91] Alan M. Frisch. The substitutional framework for sorted deduction: Fundamental results on hybrid reasoning. *Artificial Intelligence*, 49(1-3):161–198, January 1991.

[GM92] Joseph A. Goguen and José Meseguer. Order sorted algebra I. equational deduction for multiple inheritance, overloading, exceptions and partial operations. *Theoretical Computer Science*, 105(2):217–273, November 1992.

[Gra96] Bernhard Gramlich. On termination and confluence properties of disjoint and constructor-sharing conditional rewrite systems. *Theoretical Computer Science*, 165(1):97–131, 1996.

[Her86] 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.

[Hod93a] Wilfrid Hodges. Logical features of Horn clauses. In D.M. Gabbay, C.J. Hogger, and J.A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 1, pages 449–503. Oxford University Press, 1993.

[Hod93b] Wilfrid Hodges. *Model Theory*, volume 42 of *Enclyclopedia of mathematics and its applications*. Cambridge University Press, 1993.

[Hol95] Kitty L. Holland. An introduction to fusions of strongly minimal sets: The geometry of fusions. *Archive for Mathematical Logic*, 34:395–413, 1995.

[HS88] Markus Höhfeld and Gert Smolka. Definite relations over constraint languages. LILOG Report 53, IWBS, IBM Deutschland, Postfach 80 08 80, 7000 Stuttgart 80, Germany, October 1988.

[JK86] J.-P. Jouannaud and Hélène Kirchner. Completion of a set of rules modulo a set of equations. *SIAM*, 15(4):1155–1194, 1986. Preliminary version in Proceedings 11th ACM Symposium on Principles of Programming Languages, Salt Lake City (USA), 1984.

[JM94] Joxan Jaffar and Michael Maher. Constraint Logic Programming: A Survey. *Journal of Logic Programming*, 19/20:503–581, 1994.

[KKR90] C. Kirchner, H. Kirchner, and M. Rusinowitch. Deduction with symbolic constraints. *Revue Française d'Intelligence Artificielle*, 4(3):9–52, 1990. Special issue on Automatic Deduction.

[KR94a] Hélène Kirchner and Christophe Ringeissen. Combining symbolic constraint solvers on algebraic domains. *Journal of Symbolic Computation*, 18(2):113–155, 1994.

[KR94b] Hélène Kirchner and Christophe Ringeissen. Constraint solving by narrowing in combined algebraic domains. In P. Van Hentenryck, editor, *Proceedings of the 11th International Conference on Logic Programming*, pages 617–631. The MIT press, 1994.

[Mah88] M. J. Maher. Complete axiomatizations of finite, rational and infinite trees. In *LICS'88: Proceedings 3rd Symposium on Logic in Computer Science*, pages 348–357, Edinburgh, UK, June 1988.

[Mal71] A. I. Mal'cev. *The metamathematics of algebraic systems*, volume 66 of *Studies in logic and the foundations of mathematics.* North-Holland, Amsterdam-New York-Oxford-Tokyo, 1971.

[Mid94] Aart Middeldorp. Completeness of combinations of conditional constructor systems. *Journal of Symbolic Computation*, 17(1):3–21, January 1994.

[MT93] Aart Middeldorp and Yoshihito Toyama. Completeness of combinations of constructor systems. *Journal of Symbolic Computation*, 15(3):331–348, March 1993.

[Nel84] Greg Nelson. Combining satisfiability procedures by equality-sharing. *Contemporary Mathematics*, 29:201–211, 1984.

[Nip91] Tobias Nipkow. Combining matching algorithms: The regular case. *Journal of Symbolic Computation*, 12:633–653, 1991.

[NO79] Greg Nelson and Derek C. Oppen. Simplification by cooperating decision procedures. *ACM Trans. on Programming Languages and Systems*, 1(2):245–257, October 1979.

[Opp80] Derek C. Oppen. Complexity, convexity and combinations of theories. *Theoretical Computer Science*, 12, 1980.

[PT97] Anand Pillay and Akito Tsuboi. Amalgamations preserving $\aleph_0$-categoricity. *Journal of Symbolic Logic*, 1997. (accepted for publication).

[Rin92] 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.

[Rin96a] Christophe Ringeissen. Combining decision algorithms for matching in the union of disjoint equational theories. *Information and Computation*, 126(2):144–160, May 1996.

[Rin96b] Christophe Ringeissen. Cooperation of decision procedures for the satisfiability problem. In F. Baader and K.U. Schulz, editors, *Frontiers of Combining Systems: Proceedings of the 1st International Workshop, Munich (Germany)*, Applied Logic, pages 121–140. Kluwer, March 1996.

[Sho67] Joseph. R. Shoenfield. *Mathematical Logic.* Addison-Wesley, Reading, MA, 1967.

[Sho84] Robert E. Shostak. Deciding combinations of theories. *Journal of the ACM*, 31:1–12, 1984.

[Sie89] Jörg H. Siekmann. Unification theory. *Journal of Symbolic Computation*, 7(3-4):207–274, March–April 1989.

[Smo89] Gert Smolka. *Logic Programming over Polymorphically Order-Sorted Types*. PhD thesis, Fachbereich Informatik, Universität Kaiserslautern, Germany, 1989.

[SS88] 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.

[SS89] Manfred Schmidt-Schauß. Combination of unification algorithms. *Journal of Symbolic Computation*, 8(1–2):51–100, 1989. Special issue on unification. Part II.

[Sti85] Mark E. Stickel. Automated deduction by theory resolution. *Journal of Automated Reasoning*, 1(4):333–355, 1985.

[TH96] Cesare Tinelli and Mehdi Harandi. A new correctness proof of the Nelson–Oppen combination procedure. In F. Baader and K.U. Schulz, editors, *Frontiers of Combining Systems: Proceedings of the 1st International Workshop, Munich (Germany)*, Applied Logic, pages 103–120. Kluwer, March 1996.

[TH98] Cesare Tinelli and Mehdi Harandi. Constraint logic programming over unions of constraint theories. *Journal of Logic and Functional Programming*, 1998. (to appear).

[Tid86] 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.

[Wec92] Wolfgang Wechler. *Universal Algebra for Computer Scientists*, volume 25 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, Berlin-Heidelberg-New York, 1992.

[Yel87a] Kathy Yelick. Combining unification algorithms for confined equational theories. *Journal of Symbolic Computation*, 3(1), April 1987.

[Yel87b] Kathy Yelick. Unification in combinations of collapse-free regular theories. *Journal of Symbolic Computation*, 3(1–2):153–182, April 1987.

Unit´e de recherche INRIA Lorraine, Technopˆole de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unit´e de recherche INRIA Rennes, Irisa, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unit´e de recherche INRIA Rhˆone-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unit´e de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unit´e de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

Éditeur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
ISSN 0249-6399