Unions of Non-Disjoint Theories and Combinations of Satisfiability Procedures

Cesare Tinelli [1]

Department of Computer Science
University of Iowa
Iowa City, Iowa 52242 – USA
`tinelli@cs.uiowa.edu`

Christophe Ringeissen

INRIA Lorraine and CRIN-CNRS
BP 101
54602 Villers-lès-Nancy Cedex, France
`Christophe.Ringeissen@loria.fr`

Abstract

In this paper we outline a theoretical framework for the combination of decision procedures for constraint satisfiability. We describe a general combination method which, given a procedure that decides constraint satisfiability with respect to a constraint theory T_1 and one that decides constraint satisfiability with respect to a constraint theory T_2, produces a procedure that (semi-)decides constraint satisfiability with respect to the union of T_1 and T_2. We provide a number of 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 the component theories are non-disjoint. We also describe some general classes of theories to which our combination results apply, and relate our approach to some of the existing combination methods in the field.

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

[1] Partially supported by the National Science Foundation under grant no. 9972311.

1 Introduction

An established approach to problem solving is to recast problems in terms of constraint satisfaction. 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 application domain by intelligently exploiting some of the features of the domain itself. A major disadvantage is, of course, specialization. If a problem also requires solving constraints outside the constraint domain, a constraint reasoner alone is not enough. [2]

Now, many potential applications of constraint-based approaches in fields as diverse as software/hardware verification, program synthesis, computational linguistics, expert systems, and so on, are often faced with *heterogeneous problems*, that is, problems spanning over several constraint domains at once. Semantically, these are problems in a domain which is a combination of various constraint domains. Syntactically, they are problems whose constraints are expressed in a combination of the constraint languages of each constraint domain. To deal with heterogeneous problems, one can certainly try to build from scratch a constraint reasoner for the combined domain. However, if constraint reasoners are already available for the various components of the domain, it is sensible to think of obtaining a reasoner for the combined domain by somehow combining the available reasoners. Ideally, such a reasoner must be able to

- extract from the problem specification the constraints that can be handled by a component reasoner, for each such reasoners,
- assign these extracted constraints to the corresponding reasoner, and
- compose, at least in principle, the local solutions from the various reasoners into global solutions for the original problem.

To date, there are very few results on the combination of constraint domains and their reasoners. The fact is that, as desirable as it is from both a knowledge and a software engineering standpoint, this sort of combination raises several challenging model-theoretic and computational issues. Although the computational aspects of combination have been investigated for some time (see (Schulz, 2000) for a recent account), only recently have people started to study the logical and model-theoretic background of general methodologies for combining constraint reasoners. This paper represents our contribution to

[2] We use the term *domain* here in a loose sense. Typically a (constraint) domain, a semantical notion, is represented by a logical (constraint) theory, a syntactical one, which axiomatizes the domain's properties of interest. Also, we speak generically of *constraint reasoners*, as opposed to constraint solvers, to include those cases in which it not necessary to actually produce a solution of the input constraints, but it is enough to discover if the constraints are satisfiable, according to some adopted notion of satisfiability.

this study.

1.1 Previous Work

Most of the current work on the combination of constraints reasoners regards the combination of solvers for equational constraints, in particular, algorithms for E-unification (Herold, 1986; Schmidt-Schauß, 1989; Ringeissen, 1992; Boudet, 1993; Baader and Schulz, 1996) and related problems (Domenjoud et al., 1994; Baader and Schulz, 1995b; Kirchner and Ringeissen, 1994a,b). In this context, the constraint language is restricted to quantifier-free formulae over a functional signature (no predicate symbols other than equality), each component constraint domain is axiomatized by an equational theory and the combined domain is axiomatized by the union of these theories.

The emergence of general constraint-based paradigms, such as constraint logic programming (Jaffar and Maher, 1994), constrained resolution (Bürckert, 1994) and what is generally referred to as *theory-reasoning* (Baumgartner et al., 1992), raises the problem of combining reasoners for *first-order*, but not necessarily equational, constraints. The existing work on the combination of such reasoners is better understood by first realizing that combination problems can be divided into two broad classes, depending on the kind of constraint satisfiability considered by the component reasoners.

The first class comprises constraint reasoners for which satisfiability is defined in terms of *validity* of existential closures in a given constraint theory: a constraint is satisfiable if its existential closure is a logical consequence of the constraint theory. Constraint-based reasoning frameworks using reasoners of this sort are mostly based on the constraint logic programming scheme by J. Jaffar and J.-L. Lassez (Jaffar and Maher, 1994).

The second class comprises constraint reasoners for which satisfiability is defined in terms of *consistency* of existential closures with the constraint theory: a constraint is satisfiable if its existential closure is true in at least one model of the theory. Some constraint-based reasoning frameworks using reasoners of this sort are the constraint logic programming scheme of M. Höhfeld and G. Smolka (Höhfeld and Smolka, 1988), the deduction with constraints framework (Kirchner et al., 1990), constrained resolution (Bürckert, 1994), constraint contextual rewriting (Armando and Ranise, 1998), and—at least at the ground level—all theory-reasoning frameworks (Baumgartner et al., 1992).

Essentially all existing results in the combination of constraint reasoners in the first class come from the work of F. Baader and K. Schulz (Baader and Schulz, 1995a,c; Kepser and Schulz, 1996; Baader and Schulz, 1998), which lifts and extends to a first-order setting earlier combination results in the equational case.

In this paper, we are interested in the combination of constraint reasoners of the second class. Early work on this topic comes from research in automated software verification. The actual problem of interest there was the validity of assertions (expressed as universal formulae) in theories axiomatizing common data types. This problem, however, was conveniently recast as a satisfiability problem since a formula is entailed by a theory exactly when its negation is satisfiable in no models of that theory.

Initial combination results were provided by R. Shostak in (Shostak, 1979) and in (Shostak, 1984). Shostak's approach is limited in scope and not very modular—admitedly on purpose, for efficiency reasons. A rather general and completely modular combination method was proposed by G. Nelson and D. Oppen in (Nelson and Oppen, 1979) and then slightly revised in (Nelson, 1984). Given, for $i = 1, \ldots, n$, a procedure P_i that decides the satisfiability of quantifier-free formulae in a universal theory T_i, their method yields a procedure that decides the satisfiability of quantifier-free formulae in the theory $T_1 \cup \cdots \cup T_n$. A declarative and non-deterministic view of the procedure was suggested by Oppen in (Oppen, 1980). In (Tinelli and Harandi, 1996), C. Tinelli and M. Harandi followed up on this suggestion describing a non-deterministic version of the Nelson-Oppen combination procedure and providing a simpler correctness proof. A similar approach had also been followed by C. Ringeissen in (Ringeissen, 1993), which describes the procedure as a set of derivation rules applied non-deterministically.

All the works mentioned above share one major restriction on the constraint languages of the component reasoners: they must have no function or relation symbols in common. (The only exception is the equality symbol, which is however regarded as a logical constant.) This restriction has proven really hard to lift. A testament of this is that, more than two decades after Nelson and Oppen's original work, their combination results are still state of the art.

Results on *non-disjoint* combination do exist, but they are still quite limited. To start with, some results on the union of non-disjoint equational theories can be obtained as a byproduct of the research on the combination of term rewriting systems. Modular properties of term rewriting systems have been extensively investigated (see the overviews in (Ohlebusch, 1995; Gramlich, 1996) for instance). Using some of these properties it is possible to derive combination results for the word problem in the union of equational theories sharing *constructors*.[3] Outside the work on modular term rewriting, the first

[3] The word problem in an equational theory E is the problem of determining whether a given equation $s \equiv t$ is valid in E—or, equivalently, whether a disequation $\neg(s \equiv t)$ is (un)satisfiable in E. In a term rewriting system, a constructor is a function symbol that does not appear as the top symbol of a rewrite rule's

combination results for the word problem in the union of non-disjoint constraint theories were given in (Domenjoud et al., 1994) as a consequence of some combination techniques based on an adequate notion of (shared) constructors. The second of us used similar ideas later in (Ringeissen, 1996b) to extend the Nelson-Oppen method to theories sharing constructors in a sense close to that of (Domenjoud et al., 1994).

To our knowledge, the only new work since (Ringeissen, 1996b) on the combination of constraint reasoners for constraint theories with symbols in common is the one described in this paper and in a series of related papers by F. Baader and the first of us, the most recent and comprehensive of which is (Baader and Tinelli, 2001). These papers discuss a very general decision procedure for the word problem in the union of equational theories with non-disjoint signatures.[4] The procedure's correctness proof is based on some of the model-theoretic results reported here. Part of the work reported in this paper is also described in (Tinelli, 1999); a preliminary account was given in (Tinelli and Ringeissen, 1998).

1.2 Our Contribution

In this paper we focus on constraint satisfiability problems—in the sense of constrained consistency explained above—which are expressible in the language of first-order logic, or a fragment of it. For these problems, a constraint domain is formalized by a first-order structure (in the sense of Model Theory) and axiomatized by a first-order theory. Problem constraints are represented by sets of first-order formulae, constraint variables by free variables of formulae, constraint solutions by mappings of free variables into the universe of a constraint structure.

In this context, we are specifically concerned with the following combination problem: given two constraint theories T_1 and T_2 and a class $\mathcal{L}$ of constraints, how can a procedure deciding the satisfiability of $\mathcal{L}$-constraints in T_1 and a procedure deciding the satisfiability of $\mathcal{L}$-constraints in T_2 be combined into a procedure deciding the satisfiability of $\mathcal{L}$-constraints in $T_1 \cup T_2$?

This problem is unsolvable in its full generality as there exist union theories $T_1 \cup T_2$ in which constraint satisfiability is undecidable even if it is decidable in their component theories. Our main research effort then has consisted in developing appropriate restrictions on T_1, T_2 and $\mathcal{L}$ that make the above

left-hand side.

[4] An alternative but, as it turns out, equivalent approach to this topic has been very recently proposed by C. Fiorentini and S. Ghilardi in (Fiorentini and Ghilardi, 2001).

combination problem solvable. As mentioned earlier, Nelson and Oppen had already identified some: $\mathcal{L}$ is the class of quantifier-free formulae and T_1 and T_2 are universal with no non-logical symbols in common. This paper relaxes those restrictions to languages that are not necessarily quantifier-free and to theories that are not necessarily universal and have up to a finite number of non-logical symbols in common.

We start to discuss the main issues of the combination problem above in Section 3, after providing some formal preliminaries in Section 2. We first describe what we consider the most basic notion of combined structure, which we call *fusion*, and then provide a necessary and sufficient condition for two structures with arbitrary signatures to be combinable into a fusion: the structures reducts to their common signature must be isomorphic. Then, we show under what conditions the satisfiability of "mixed" constraints in a fusion structure is reducible to the satisfiability of pure [5] constraints in the fusion's components. The main requirement is that the two component structures have a set of elements X and Y, respectively, such that any injection from a finite subset of X into Y extends to an isomorphism of the structures' reducts to the common signature.

In Section 4, we lift the results in the previous section from fusions of structures to unions of theories. This lifting is possible for theories that are *N-O-combinable* over a given class $\mathcal{L}$ of constraints. The essence of N-O-combinability, a rather technical notion, is that the satisfiability in a theory $T_1 \cup T_2$ of the conjunction $\varphi_1 \wedge \varphi_2$ of two pure constraints can be reduced to the *local* satisfiability of φ_1 in T_1 and of φ_2 in T_2 by adding to both formulae an appropriate Σ-*restriction*, a particular kind of first-order restriction on the free variables shared by φ_1 and φ_2. Adding a restriction on the values of the shared variables is in the spirit of the Nelson-Oppen combination procedure [6], but tailored to the case of theories with not necessarily disjoint signatures.

In Section 5, we then describe an extension of the Nelson-Oppen procedure that, by guessing the right Σ-restrictions, is sound and complete for N-O-combinable theories. Our combination procedure is only a semi-decision procedure in general because the set of possible Σ-restrictions is infinite whenever the component theories share function symbols. Nonetheless, it yields the following modular decidability result for the union of two N-O-combinable and axiomatizable theories T_1 and T_2: if the satisfiability in each T_i of pure constraints with Σ-restrictions is decidable then the satisfiability in $T_1 \cup T_2$ of mixed constrains with Σ-restrictions is also decidable. This generalizes both

[5] By *pure* we mean made only of symbols from one of the two theories.

[6] More precisely, of its non-deterministic version, where the added restrictions are simply conjunctions of equations and disequations between shared variables. See, e.g., (Tinelli and Harandi, 1996) for details.

Nelson and Oppen's combination results and Ringeissen's initial results in (Ringeissen, 1996b).

The definition of N-O-combinable theories is rather abstract and imposes conditions on the two theories as a pair, not individually. As a consequence, it is not immediate to tell when two theories are N-O-combinable. We dedicate the rest of the paper to developing more "local" restrictions sufficient for N-O-combinability.

In Section 6, we discuss some criteria for showing that two theories are N-O-combinable. In particular, we define a local property for component theories that, with some additional conditions, makes them N-O-combinable. This property, which we call *stable Σ-freeness*, is an extension of Nelson and Oppen's idea of stable-infiniteness of a theory. In essence, a theory T is stably Σ-free (over a certain constraint language) if every constraint (in the language) satisfiable in T is satisfiable in a model of T whose Σ-reduct is a free structure with infinitely-many generators.

As discovered by previous research on non-disjoint combination, it is easier to combine theories whose shared function symbols are constructors in an appropriate sense. In Section 7, we provide our own definition of constructors, discuss its main properties, and argue that it generalizes previous notions of constructors in the literature. The main idea is that a subsignature of a theory T is a set of constructors for T if every term has a normal form (in T) such that its top part is made only of constructors and the equivalence in T of two normal forms reduces, in a precise sense, to the equivalence of their top parts. This notion of constructors is interesting in its own right, but we use it in this paper mainly to provide an example of a large class of stably Σ-free theories.

In Section 8, we then present some examples of classes of stably Σ-free theories that are N-O-combinable and discuss one of them in detail. In this class the theories will share constructors in the sense of Section 7.

Section 9 concludes the paper with some directions for further research.

2 Formal Preliminaries

We start by introducing some of the basic notions from Model Theory and Universal Algebra that we use in the paper. For the most part we will closely adhere to the notation and terminology of (Hodges, 1993) and (Wechler, 1992).

A *signature* Σ consists of a set Σ^{P} of *relation symbols* and a set Σ^{F} of *function symbols*, each with an associated *arity*, an integer $n \geq 0$. A *constant* symbol is a function symbol of zero arity. A *functional* signature is a signature with no relation symbols. We use the letters Σ, Ω, Δ to denote signatures.

Throughout the paper, we fix a countably-infinite set V of *variables*, disjoint with any signature Σ. For any $X \subseteq V$, $T(\Sigma, X)$ denotes the set of Σ-*terms* over the variables X, i.e., first-order terms of signature Σ^{F} and variables from X. If t is a term, $t(\epsilon)$ denotes the *top symbol* of t, that is, $t(\epsilon) = t$ if t is a variable in V, and $t(\epsilon) = f$ if $t = f(t_1, \ldots, t_n)$ for $n \geq 0$. We generally use u, v, w to denote logical variables, and r, s, t to denote Σ-terms.

We use φ, ψ, γ to denote first-order formulae. The symbols $\top, \bot$ respectively denote the universally true and universally false formula; $\equiv$ denotes equality in formulae; $s \not\equiv t$ is an abbreviation for $\neg(s \equiv t)$. If t is a term and φ a formula, $\mathcal{V}ar(t)$ denotes the set of t's variables while $\mathcal{V}ar(\varphi)$ denotes the set of φ's *free* variables. This notation is extended in the obvious way to sets of terms or formulae. As usual, we call a formula *ground* if it has no variables, and a *sentence* if it has no free variables.

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

Symbols with a tilde on top denote finite sequences. For instance, $\tilde{x}$ stands for an n-sequence of the form $(x_1, x_2, \ldots, x_n)$, for $n \geq 0$.[7] We denote by $\tilde{x}, \tilde{y}$ the sequence obtained by concatenating $\tilde{x}$ with $\tilde{y}$. We use the tilde notation for members of a Cartesian product as well. Whenever convenient, we will also treat $\tilde{x}$ as the set of its elements.

The notation $\varphi(v_1, \ldots, v_n)$ indicates that the free variables of the formula φ are *exactly* the ones in $(v_1, \ldots, v_n)$, i.e., $\mathcal{V}ar(\varphi) = \{v_1, \ldots, v_n\}$.[8] Similarly for $t(v_1, \ldots, v_n)$, where t is a term. In both cases it is understood that the elements of $(v_1, \ldots, v_n)$ are pairwise distinct. We will also use the notation $\varphi(\tilde{v})$ and $t(\tilde{v})$ whenever convenient. When we write $f(\tilde{v})$, where f is a function symbol, it is also understood that the length of $\tilde{v}$ equals the arity of f. For any formula $\varphi(v_1, \ldots, v_n)$, $\tilde{\exists}\varphi$ and $\tilde{\forall}\varphi$ denote respectively the existential and the universal closure of φ. For notational convenience, we will systematically identify finite sets of formulae with the conjunction of their elements (and identify the empty

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

[8] This notation is non-standard, as $\varphi(v_1, \ldots, v_n)$ generally indicates that the free variables of φ are *included* in $\{v_1, \ldots, v_n\}$. We use it here because it simplifies the enunciation of most of our results.

set of formulae with $\top$).

We use the standard notion of substitution, extended from terms to arbitrary first-order formulae (and sets thereof) by renaming quantified variables when necessary to avoid capturing of free variables. We denote the empty substitution by ε and write substitution applications in postfix form. Also, if σ is a substitution we call the sets

$$\mathcal{D}om(\sigma) := \{v \in V \mid v\sigma \neq v\} \text{ and } \mathcal{R}an(\sigma) := \{v\sigma \mid v \in \mathcal{D}om(\sigma)\}$$

respectively the *domain* and the *range* of σ. A substitution σ such that $\mathcal{D}om(\sigma) = \{v_1, \ldots, v_n\}$ and $v_i\sigma = t_i$ for all $i \in \{1, \ldots, n\}$ will be denoted by $\{v_1 \leftarrow t_1, \ldots, v_n \leftarrow t_n\}$. With no loss of generality we only consider *idempotent* substitutions, that is, substitutions σ such that $\sigma \circ \sigma = \sigma$. For each $U \subseteq V$, $\mathrm{SUB}(U)$ denotes the set of idempotent substitutions whose domain (in the sense above) is included in U.

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

Let $\mathcal{A}$ be a structure of signature Σ. If f is a symbol of Σ, $f^{\mathcal{A}}$ denotes the interpretation of f in $\mathcal{A}$. If Ω is a subsignature of Σ, $\mathcal{A}^{\Omega}$ denotes the *reduct* of $\mathcal{A}$ to Ω, that is, the structure obtained from $\mathcal{A}$ by "forgetting" the symbols not in Ω. If U is a set of variables in V, a *valuation* of U is a mapping of U into A. The pair $(\mathcal{A}, \alpha)$ defines an *interpretation*, mapping the terms in $T(\Sigma, U)$ to elements of A, and Σ-formulae φ with free variables in U to true or false. For all $t \in T(\Sigma, U)$, $[\![t]\!]^{\mathcal{A}}_{\alpha}$ denotes the element of A which $(\mathcal{A}, \alpha)$ assigns to t. Using the function $t^{\mathcal{A}}$ induced by t on $\mathcal{A}$, we may also write such an element as $t^{\mathcal{A}}(\tilde{a})$, where $\tilde{a}$ is the tuple of values assigned by α to $\tilde{v}$. We say that $(\mathcal{A}, \alpha)$ *satisfies* a Σ-formula $\varphi(\tilde{v})$, or that α *satisfies* φ *in* $\mathcal{A}$, if $(\mathcal{A}, \alpha)$ maps φ to true. In that case, we write $(\mathcal{A}, \alpha) \models \varphi$. Alternatively, if $\tilde{a}$ is the tuple of values assigned by α to $\tilde{v}$, we may write $\mathcal{A} \models \varphi[\tilde{a}]$. In either case, we will call α an $\mathcal{A}$-*solution* of φ. If φ has no free variables, the choice of α is irrelevant and so we write just $\mathcal{A} \models \varphi$. We say that φ is *satisfiable* in $\mathcal{A}$ if there is a valuation of $\mathcal{V}ar(\varphi)$ that satisfies φ in $\mathcal{A}$ (equivalently, if $\mathcal{A} \models \tilde{\exists}\, \varphi$). We write $\mathcal{A} \models \varphi$ and say that $\mathcal{A}$ *models* φ if *every* valuation of $\mathcal{V}ar(\varphi)$ into A satisfies φ (equivalently, if $\mathcal{A} \models \tilde{\forall}\, \varphi$).

If $\mathbf{K}$ is a class of Σ-structures, we say that φ is *satisfiable in* $\mathbf{K}$ if it is satisfiable in at least one member of $\mathbf{K}$. We 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 say that $\mathbf{K}$ is *non-trivial* if it contains non-trivial structures, that is, structures of cardinality greater than 1.

If $\mathcal{A}$ is a Σ-structure and $X \subseteq A$, $\langle X \rangle_{\mathcal{A}}$ denotes the substructure of $\mathcal{A}$ generated by X. Recall that X is said to *generate* $\mathcal{A}$, or to be a *a set of generators for* $\mathcal{A}$, if $\mathcal{A} = \langle X \rangle_{\mathcal{A}}$. 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 $\mathcal{A}$. While every structure admits a set of generators (its whole universe, for instance), not every structure admits a non-redundant set of generators. Non-redundant sets of generators have the following, easily provable property.

Lemma 1 *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 $\langle X \rangle_{\mathcal{A}}$.*

For brevity, we will often use the definitions below, where $\mathcal{A}$ is any structure and Σ a subsignature of Σ_A.

Definition 2 (Σ-generators) *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.*

It is immediate that when $(\Sigma_A)^{\mathrm{F}} \subseteq \Sigma \subseteq \Sigma_A$, the notions of generators and Σ-generators coincide.

Definition 3 (Σ-Isolated Individual) *An element $a \in A$ is a* Σ-isolated individual *of $\mathcal{A}$ if a is not in the range of the interpretation of any function symbol of Σ, i.e., if there is no $g \in \Sigma^{\mathrm{F}}$ of arity $n \geq 0$ and n-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 $\mathcal{A}$ 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, we will use $Is(\mathcal{A}^{\Sigma})$ to denote either of them. Notice that each Σ-isolated individual of a structure is necessarily included in every set of Σ-generators for that structure. Moreover, any set of Σ-generators consisting only of Σ-isolated individuals is necessarily non-redundant.

A structure $\mathcal{B}$ is an *expansion* of a structure $\mathcal{A}$ if $\mathcal{A}$ is a reduct of $\mathcal{B}$. We will implicitly appeal to the following fact almost constantly in the rest of the paper.

Lemma 4 *Let $\mathcal{A}$ be an Σ-structure, $\varphi(\tilde{v})$ a Σ-formula, and α a valuation of $\tilde{v}$ into A. Then, for any expansion $\mathcal{B}$ of $\mathcal{A}$ to a signature $\Omega \supseteq \Sigma$, $(\mathcal{A}, \alpha) \models \varphi$ iff $(\mathcal{B}, \alpha) \models \varphi$.*

A *first-order theory* is a set of first-order sentences. A Σ-theory is a theory all of whose sentences have signature Σ. All the theories we consider will be first-order theories *with equality*, which means that equality symbol $\equiv$ will always be interpreted as the identity relation.

As usual, a Σ-structure $\mathcal{A}$ is a *model* of a Σ-theory T if $\mathcal{A}$ models every sentence in T. We denote by $Mod^{\Sigma}(T)$, or just $Mod(T)$ when Σ is clear from context, the class of all the Σ-models of T. We say that T is *non-trivial* if $Mod(T)$ is non-trivial. A Σ-formula φ is *satisfiable in* T if it is satisfiable in $Mod(T)$. By the above, a formula φ is satisfiable in T exactly when the theory $T \cup \{\tilde{\exists}\, \varphi\}$ has a model. Two Σ-formulae φ and ψ are *equisatisfiable in* T if for every model $\mathcal{A}$ of T, φ is satisfiable in $\mathcal{A}$ if and only if ψ is satisfiable in $\mathcal{A}$. We say simply that two formulae are *equisatisfiable* if they are equisatisfiable in the empty theory. [9]

The Σ-theory T *entails* φ, written $T \models \varphi$, if $Mod(T) \models \varphi$. If T' is another Σ-theory, we write $T \models T'$ if T entails every sentence in T'. For all Σ-terms s, t, we write $s =_T t$ and say that s and t are *equivalent in* T iff $T \models s \equiv t$. If Ω is a subsignature of Σ we call Ω-restriction of T, or also Ω*-theory of* T, the set T^{Ω} of all the Ω-sentences entailed by T.

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. [10] Notice that a theory T is collapse-free iff the class $Mod(T)$ 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 defined as theories axiomatized by a set of (universally quantified) equations. Here, we extend such a notion to theories whose signature may include predicate symbols as well. We say that a theory is *atomic* if it is axiomatized by a set of sentences of the form $\tilde{\forall}\,\varphi$, where φ is an atomic formula. We use the symbol H to denote a given atomic theory. It can be shown (see, e.g., (Hodges, 1993)) that a class $\mathbf{K}$ of Σ-structures is closed under the formation of substructures, homomorphic images, and direct products exactly when it is axiomatized by some atomic Σ-theory H. In analogy to the equational case then, we call $Mod(H)$ a Σ*-variety.*

If T is a Σ-theory, $\mathrm{At}(T)$ denotes the *atomic theory of* T, the set of all the universally quantified Σ-atoms entailed by T. For any $\Omega \subseteq \Sigma$, we then call $\mathrm{At}(T^{\Omega})$, the set of all universally quantified Ω-atoms entailed by T, *the atomic* Ω*-theory of* T. Similarly, we call *atomic* Ω*-theory of* Σ*-structure* $\mathcal{A}$, and denote by $\mathrm{At}(\mathcal{A}^{\Omega})$, the set of all the universally quantified Ω-atoms modeled by $\mathcal{A}$. We

[9] Note that although logically equivalent formulae are equisatisfiable, the converse is not true. For instance, the formulae $x \equiv a$ and $x \equiv a \wedge y \equiv a$, where x, y are variables and a is a constant symbol, are equisatisfiable but are not logically equivalent.

[10] 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 function symbols the trivial Σ-theory is collapse-free. In any case, the two definitions coincide for non-trivial theories, the theories of interest in this paper.

refer to $Mod(\mathrm{At}(T^{\Omega}))$ as the *Ω-variety of T* and often identify it with $\mathrm{At}(T^{\Omega})$.

3 Combining Constraint Domains

As mentioned in the introduction, 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 proposing a general notion of combined structure, which we call *fusion* [11] . Our primary goal is to identify a *minimal* set of requirements that make a structure a viable combination of a number of given structures. As it turns out, the notion of fusion, which we give below, is general enough to include the type of combined structures found in the literature and, at the same time, provide the basis for all the combination results given in this paper. For simplicity, we will mostly consider combinations of just two component structures.

In the following, and in the rest of the paper, we rely on the standard notions of morphisms of structures from Model Theory (Hodges, 1993). We write $\mathcal{A} \cong \mathcal{B}$ to state that the structures $\mathcal{A}$ and $\mathcal{B}$ are isomorphic, and write $h\colon \mathcal{A} \cong \mathcal{B}$ to state that h is an isomorphism of $\mathcal{A}$ onto $\mathcal{B}$.

Definition 5 (Fusion) *Given two structures $\mathcal{A}$ and $\mathcal{B}$, a $(\Sigma_A \cup \Sigma_B)$-structure $\mathcal{F}$ is a* fusion *of $\mathcal{A}$ and $\mathcal{B}$ iff there exist a map h_{A-F} and a map h_{B-F} such that*

[11] 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 (Pillay and Tsuboi, 1997) does use "amalgamation" for the same type of combined structure as ours while (Holland, 1995) uses "fusion" for a rather different type of combined structure. Our notion of fusion is closely related to the one employed in algebraic approaches to modal logics (see, e.g., (Wolter, 1998)).

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

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 $\mathcal{B}$. Notice that the signatures of the two structures are not necessarily disjoint.

Baader and Schulz's free amalgamated product (Baader and Schulz, 1998) and Kepser and Schulz's rational amalgamation (Kepser and Schulz, 1996) of two quasi-free structures are both readily shown to be a fusion of those structures. Similarly, the amalgamation construction given by Ringeissen in (Ringeissen, 1996b) can also be shown to produce a fusion.

In principle, one could imagine a notion of fusion based on more general morphisms than isomorphisms. For instance, we could say that a structure $\mathcal{F}$ is a fusion of the structures $\mathcal{A}$ and $\mathcal{B}$ in Definition 5 if $\mathcal{A}$ is embeddable in $\mathcal{F}^{\Sigma_A}$ and $\mathcal{B}$ is embeddable in $\mathcal{F}^{\Sigma_B}$. A justification that the definition we give is the right one for our purposes will be provided in Section 4, where we show that all models of a union theory are fusions of models of its component theories.

We denote by $\mathit{Fus}(\mathcal{A}, \mathcal{B})$ the set of all the fusions of two structures $\mathcal{A}$ and $\mathcal{B}$. By Definition 5, it is immediate that $\mathit{Fus}(\mathcal{A}, \mathcal{B}) = \mathit{Fus}(\mathcal{B}, \mathcal{A})$ and that $\mathit{Fus}(\mathcal{A}, \mathcal{B})$ is an abstract class, i.e., it is closed under isomorphism. Note that $\mathit{Fus}(\mathcal{A}, \mathcal{B})$ will usually contain non-isomorphic structures. [12] Intuitively, however, all of its members should be isomorphic over the symbols shared by $\mathcal{A}$ and $\mathcal{B}$. Such an intuition is confirmed by the proposition below, which establishes a necessary and sufficient condition for the existence of fusions.

Proposition 6 *For all structures $\mathcal{A}$ and $\mathcal{B}$,*

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

PROOF. Let $\Sigma := \Sigma_A \cap \Sigma_B$. [13]

($\Rightarrow$) Let $\mathcal{C} \in \mathit{Fus}(\mathcal{A}, \mathcal{B})$. 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 $\mathcal{A}^{\Sigma} \cong \mathcal{C}^{\Sigma}$

[12] For example, assume that the signatures of $\mathcal{A}$ and $\mathcal{B}$ are disjoint and each contains some constant symbol. Then, these two symbols may denote the same individual in one fusion of $\mathcal{A}$ and $\mathcal{B}$ and distinct individuals in another.

[13] To simplify the notation, here and in the rest of the paper we adopt the following notational convention. If $h : C \to D$ is a map and $\tilde{c} \in C^n$, the expression $h(\tilde{c})$ denotes the tuple $(h(c_1), \ldots, h(c_n))$. If R is an n-ary relation over C, the expression $h(R)$ denotes the relation $\{h(\tilde{c}) \mid \tilde{c} \in R\}$.

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

($\Leftarrow$) Let h be a (bijetive) 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 $P \in (\Sigma_A \cup \Sigma_B)^{\mathrm{P}}$,

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

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

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

The structure $\mathcal{C}$ interprets Σ_B-symbols the way $\mathcal{B}$ 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 $\Sigma_A \setminus \Sigma$, for each $\tilde{a} \in A^n$,

$$\begin{aligned} \tilde{a} \in P^{\mathcal{A}} &\text{ iff } h(\tilde{a}) \in h(P^{\mathcal{A}}) && \text{(by def. of } h(P^{\mathcal{A}}) \text{ and injectivity of } h) \\ &\text{ iff } h(\tilde{a}) \in P^{\mathcal{C}} && \text{(by constr. of } \mathcal{C}); \end{aligned}$$

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

$$\begin{aligned} \tilde{a} \in P^{\mathcal{A}} &\text{ iff } h(\tilde{a}) \in P^{\mathcal{B}} && (h\colon \mathcal{A}^\Sigma \cong \mathcal{B}^\Sigma) \\ &\text{ iff } h(\tilde{a}) \in P^{\mathcal{C}} && \text{(by constr. of } \mathcal{C}); \end{aligned}$$

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

$$\begin{aligned} h(g^{\mathcal{A}}(\tilde{a})) &= h(g^{\mathcal{A}}(h^{-1}(h(\tilde{a})))) && \text{(by bijectivity of } h) \\ &= g^{\mathcal{C}}(h(\tilde{a})) && \text{(by constr. of } \mathcal{C}); \end{aligned}$$

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

$$\begin{aligned} h(g^{\mathcal{A}}(\tilde{a})) &= g^{\mathcal{B}}(h(\tilde{a})) && (h\colon \mathcal{A}^\Sigma \cong \mathcal{B}^\Sigma) \\ &= g^{\mathcal{C}}(h(\tilde{a})) && \text{(by constr. of } \mathcal{C}); \end{aligned}$$

By construction of $\mathcal{C}$, 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 $\mathcal{B}$. $\square$

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.

Given an isomorphism h of $\mathcal{A}^{\Sigma}$ and $\mathcal{B}^{\Sigma}$, we will call *canonical fusion of $\mathcal{A}$ and $\mathcal{B}$ induced by h* the fusion of $\mathcal{A}$ and $\mathcal{B}$ constructed like the fusion $\langle \mathcal{C}, h, id \rangle$ in the proof above.

We know that for each structure there is at least one set of individuals, the set of generators, that 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 7 (Fusible Structures) *Consider two structures $\mathcal{A}$ and $\mathcal{B}$, a set $X \subseteq A$, and a set $Y \subseteq B$ with X's 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 $\langle X, Y \rangle$ whenever $\mathcal{B}$ is fusible with $\mathcal{A}$ over $\langle Y, X \rangle$, for brevity 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.

Observe that $\mathcal{A}$ and $\mathcal{B}$ admit a fusion whenever $\mathcal{A}$ and $\mathcal{B}$ are fusible over some $\langle X, Y \rangle$. In that case in fact, according to the definition above, the empty mapping from X to Y extends to an isomorphism of $\mathcal{A}^{\Sigma_A \cap \Sigma_B}$ onto $\mathcal{B}^{\Sigma_A \cap \Sigma_B}$. But then, $Fus(\mathcal{A}, \mathcal{B})$ is non-empty by Proposition 6.

We will provide some sufficient conditions for the fusibility of two structures in Section 6.2. For now, our interest in fusions in general and fusible structures in particular is motivated by the fact that, under the right conditions, satisfiability in a fusion of two fusible structures reduces to satisfiability in each of them. To show this 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?

Lemma 8 *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 $\varphi(\tilde{u}, \tilde{v})$ be an Ω-formula and $\psi(\tilde{w}, \tilde{v})$ 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 $\tilde{a}$, $\tilde{b}$ consist of arbitrary elements of A, B, respectively, $\tilde{x} := (x_1, \ldots, x_m)$ is in X, $\tilde{y} := (y_1, \ldots, y_m)$ is in Y, and neither $\tilde{x}$ nor $\tilde{y}$ 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 $j \in \{1, \ldots, m\}$,

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

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. 6, it follows that $\mathit{Fus}(\mathcal{A}^{\Omega \cup K}, \mathcal{B}^{\Delta \cup K})$ is not empty. Consider any $\mathcal{F} \in \mathit{Fus}(\mathcal{A}^{\Omega \cup K}, \mathcal{B}^{\Delta \cup K})$. 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 $\mathcal{F}^{\Omega \cup \Delta} \in \mathit{Fus}(\mathcal{A}, \mathcal{B})$.

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 4, that $\mathcal{F}^{\Omega \cup \Delta} \models \tilde{\exists}\,(\varphi \wedge \psi)$. $\square$

The lemma above contains the most important model-theoretic result of this paper, in the sense that all the combination results we present here will ultimately rest on it. To be able to use it effectively, however, we need a more syntactical characterization. We will give this characterization in two steps, starting with the simple case of structures with disjoint signature and then moving to the general case.

Consider the structures $\mathcal{A}$ and $\mathcal{B}$, and the sentences φ and ψ given in Lemma 8. When the signatures of $\mathcal{A}$ and $\mathcal{B}$ have no symbols in common, the sufficient condition for the satisfiability of $\varphi \wedge \psi$ can be expressed syntactically by adding to both φ and ψ a simple constraint on the free variables they share. We will define this constraint using the notion of *variable identification*.

Definition 9 (Identification) *Given a finite set U of variables, the set of* identifications of U *is defined as follows,*[14]

$$\mathrm{ID}(U) := \{\xi \in \mathrm{SUB}(U) \mid \mathcal{R}an(\xi) \subseteq U \setminus \mathcal{D}om(\xi)\}.$$

Every substitution in $\mathrm{ID}(U)$ defines a partition of U and identifies all the variables in the same block with a representative of that block. To each $\xi \in \mathrm{ID}(U)$ we will associate the set of constraints

$$dif_{\xi}(U) := \bigcup_{u,v \in U\xi,\ u \neq v} \{u \not\equiv v\}$$

expressing the fact that any two variables not identified by ξ must take distinct values. We will write just dif_{ξ} when the set U is clear from context.

Observe that the empty substitution over the variables U always belongs to $\mathrm{ID}(U)$ and that the associated set of constraints, which we will denote simply by $dif(U)$, is made of all the possible disequations between distinct elements of U. Also observe that $dif(U)$ is satisfied exactly when no two variables in U are assigned to the same individual.

We can now use $dif(U)$ to obtain an immediate special case of Lemma 8.

Lemma 10 *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two signature-disjoint structures with same cardinality and, for $i = 1, 2$, consider the $\Sigma_{\mathcal{A}_i}$-formula $\varphi_i(\tilde{u}_i, \tilde{v})$, where $\tilde{u}_1 \cap \tilde{u}_2 = \emptyset$. If $\varphi_i \wedge dif(\tilde{v})$ 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 $(\mathcal{A}_i, \alpha_i) \models \varphi_i \wedge dif(\tilde{v})$. Observe that, because of $dif(\tilde{v})$, α_i assigns pairwise distinct individuals to the shared variables of φ_i. The result follows then from Lemma 8 noting that two equinumerous structures $\mathcal{A}$ and $\mathcal{B}$ are trivially fusible over $\langle A, B\rangle$ when their signatures are disjoint. $\square$

[14] Recall that $\mathrm{SUB}(U)$ is the set of idempotent substitutions whose domain is included in U.

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 they assign to the shared variables, if any. The role of $\mathit{dif}(\tilde{v})$ is exactly that of assuring such a 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 $\mathcal{A}_2$ [15] there must exist an $\mathcal{A}_2$-solution of φ_2 that also identifies these variables. We can then generalize Lemma 10 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) identified by the $\mathcal{A}_i$-solutions. Then, the constraint dif_ξ, which is nothing but $\mathit{dif}(\tilde{v}\xi)$, can be used in the same way $\mathit{dif}(\tilde{v})$ was used before.

Proposition 11 *For $i = 1, 2$, let $\mathcal{A}_i$ and φ_i be as in Lemma 10. If, for $i = 1, 2$,*

$$\varphi_i \xi \wedge \mathit{dif}_\xi$$

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 above proposition is the syntactic counterpart of Lemma 8 in the case of signature-disjoint structures. The addition of a simple constraint guarantees that the shared variables (after the identification) take distinct values over the fusors of the component structures, as the lemma requires. Since equinumerous structures with disjoint signatures are fusible over their whole carriers, the task here was essentially trivial.

The converse of Proposition 11 holds as well—we will prove a more general version of it in the next subsection for structures with non-necessarily disjoint signature. This already provides a sound and complete combination method to decide the satisfiability in $\mathit{Fus}(\mathcal{A}_1, \mathcal{A}_2)$ of a formula $\varphi_1 \wedge \varphi_2$ like the one in the proposition: consider all possible identifications ξ of the variables shared by φ_1 and φ_2 until one is found that makes $\varphi_i \xi \wedge \mathit{dif}_\xi$ satisfiable in $\mathcal{A}_i$, for $i = 1, 2$. The combination method is also always terminating in this case because there

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

are only finitely-many identifications to consider. Unfortunately, things are not so nice and simple when $\mathcal{A}_1$ and $\mathcal{A}_2$ have symbols in common.

3.2 Non-disjoint Signatures

When two structures are not signature-disjoint, they are likely to be fusible only over sets of fusors that are properly contained in their universes. Now, since the property of being a fusor does not appear to be first-order definable, this means that, in general, it may not be possible to force a variable to range over a set of fusors by the simple addition of a first-order constraint like dif_{ξ}, as we did in the previous subsection. One case in which it is possible is when the fusors in question are also Σ-isolated, where Σ is a finite set of symbols shared by the two structures' signatures. But to see that we will need some more definitions and notation.

Definition 12 (Instantiation) *Given a finite set U of variables and a finite signature Σ, the set of* Σ-instantiations of U *is defined as follows,*

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

Note that a Σ-instantiation of U either fixes an element of U or maps it to a non-variable Σ-term. To avoid name conflicts, given that an instantiation may introduce variables not in its domain, we will only consider Σ-instantiations ρ such that the variables occurring in $\mathcal{R}an(\rho)$ are all fresh. To every instantiation $\rho \in \mathrm{IN}^{\Sigma}(U)$, we will associate the set

$$iso_{\rho}^{\Sigma}(U) := \bigcup_{v \in \mathcal{V}ar(U\rho),\, f_i \in \Sigma^{\mathrm{F}}} \{\forall \tilde{u}_i \; v \not\equiv f_i(\tilde{u}_i)\},$$

which we will denote just by iso_{ρ} when Σ and U are clear from the context.

Observe that the set iso_{ρ}^{Σ} is satisfied by a valuation α if and only if α maps the variables in $U\rho$ to individuals that are not in the range of any Σ-function, i.e., Σ-isolated individuals. Also observe that the empty substitution belongs to $\mathrm{IN}^{\Sigma}(U)$ for any U and Σ. We will denote its associated set simply by $iso^{\Sigma}(U)$.

As we did in the previous subsection, we can use $iso^{\Sigma}(U)$ together with $dif(U)$ to obtain a special case of Lemma 8.

Lemma 13 *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 formula*

$$\varphi_i \wedge iso^{\Sigma}(\tilde{v}) \wedge dif(\tilde{v})$$

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 $\mathcal{A}_i \models \varphi_i[\tilde{a}_i, \tilde{x}_i] \wedge \mathit{iso}^{\Sigma}[\tilde{x}_i] \wedge \mathit{dif}[\tilde{x}_i]$. By Lemma 8, 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 $\mathit{dif}[\tilde{x}_i]$ is true in $\mathcal{A}_i$. To see that $\tilde{x}_i$ is included in X_i, just recall that $\mathit{iso}^{\Sigma}[\tilde{x}_i]$ 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. □

From the proof above and that of Lemma 8 it is clear that we actually have a slightly stronger result: when the conditions of the Lemma 13 hold, the whole formula $\varphi_1 \wedge \varphi_2 \wedge \mathit{iso}^{\Sigma}(\tilde{v}) \wedge \mathit{dif}(\tilde{v})$ is in fact satisfiable in a fusion of $\mathcal{A}_1$ and $\mathcal{A}_2$.

In Lemma 13, the requirement that both sets of fusors contain the Σ-isolated individuals of their respective structures allows us to use a first-order formula, $\mathit{iso}^{\Sigma}(\tilde{v}) \wedge \mathit{dif}(\tilde{v})$, to force the variables shared by the two pure formulae to take distinct values over the fusors. But now, what if either φ_i is satisfiable only with valuations that map some shared variables to individuals that are not Σ-isolated? 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 did before.

To formalize the intuitions above it is convenient to introduce the following restricted notion of fusibility.

Definition 14 (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 *iff for* $i = 1, 2$ *there is a set* X_i *such that* $\mathit{Is}({\mathcal{A}_i}^{\Sigma}) \subseteq X_i \subseteq A_i$ *and* $\mathcal{A}_1$ *and* $\mathcal{A}_2$ *are fusible over* $\langle X_1, X_2 \rangle$*.*

A little clarification 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-formulae 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 $\mathit{Is}({\mathcal{A}_1}^{\Sigma})$ and $\mathit{Is}({\mathcal{A}_2}^{\Sigma})$ 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 15 *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*

$$(\varphi_i \rho \wedge iso_\rho)\xi \wedge dif_\xi$$

is satisfiable in $\mathcal{A}_i$ for some $\rho \in \mathrm{IN}^{\Sigma}(\tilde{v})$ and $\xi \in \mathrm{ID}(\mathcal{V}ar(\tilde{v}\rho))$, 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 $(\varphi_i \rho \wedge iso_\rho)\xi \wedge dif_\xi$ is satisfiable in $\mathcal{A}_i$, where ρ and ξ are as described above. Where $\varphi'_i := \varphi_i \rho \xi$ and $\tilde{w} := \mathcal{V}ar(\tilde{v}\rho)\xi$, it is easy to see that $iso_\rho \xi = iso^{\Sigma}(\tilde{w})$ and $dif_\xi = dif(\tilde{w})$, which means that $(\varphi_i \rho \wedge iso_\rho)\xi \wedge dif_\xi$ has actually the form

$$\varphi'_i(\tilde{u}_i, \tilde{w}) \wedge iso^{\Sigma}(\tilde{w}) \wedge dif(\tilde{w}).$$

From the assumptions and Lemma 13 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 $(\varphi'_1 \wedge \varphi'_2) = (\varphi_1 \wedge \varphi_2)\rho\xi$. $\square$

This proposition is both a syntactic specialization of Lemma 8 and a proper generalization of Proposition 11 to the case of structures with arbitrary signatures. It should already be clear though that any combination method based on it will not in general be terminating, as the number of possible instantiations ρ above becomes infinite once the structures share a function symbol of non-zero arity.

Furthermore, being a specialization of Lemma 8, Proposition 15 provides just a sufficient condition for the joint satisfiability of $\varphi_1 \wedge \varphi_2$. The satisfiability of $(\varphi_i \rho \wedge iso_\rho)\xi \wedge dif_\xi$ in $\mathcal{A}_i$, although sufficient, is typically not necessary for the satisfiability of $\varphi_1 \wedge \varphi_2$ in $Fus(\mathcal{A}_1, \mathcal{A}_2)$. It does become necessary, however, if $\mathcal{A}_1$ and $\mathcal{A}_2$ have a fusion Σ-generated by its Σ-isolated individuals alone.

Proposition 16 *Let $\mathcal{A}_1, \mathcal{A}_2$ be two structures with respective signatures Σ_1, Σ_2 and admitting a fusion $\mathcal{F}$ which is Σ-generated by its Σ-isolated individuals, for some finite $\Sigma \subseteq \Sigma_1 \cap \Sigma_2$. For $i = 1, 2$, consider the Σ_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 $\xi \in \mathrm{ID}(\mathcal{V}ar(\tilde{v}\rho))$ such that $(\varphi_i \rho \wedge iso_\rho)\xi \wedge dif_\xi$ 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 $\alpha(v_j) = t_j^{\mathcal{F}}[\tilde{x}_j]$.[16] 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 $v, w \in \mathcal{V}ar(\tilde{v}\rho)$,

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

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

It should be noted that the requirement that a structure (in the case above, a fusion) be Σ-generated by its Σ-isolated individuals is rather strong. It is easy to find natural examples of structures that are not. For instance, let $\mathcal{A}$ be the integers with zero, successor and predecessor and let Σ consist of the zero and successor symbols. Now, although the set of $\mathcal{A}$'s Σ-isolated individuals is empty—as every integer is the successor of another one—the structure $\mathcal{A}$ is not Σ-generated by the empty set. However, we will see in Section 7 that there is a large and interesting class of structures Σ-generated by their Σ-isolated individuals.

3.3 Σ-Restricted Formulae

We will use formulae with an added constraint of the form $\mathit{iso}^{\Sigma}(\tilde{v}) \wedge \mathit{dif}(\tilde{v})$ often enough to justify the following definition.

Definition 17 (Σ-Restricted Formula) *Given a finite signature Σ and a (possibly empty) tuple of variables $\tilde{v}$ we say that a formula ψ is* Σ-restricted

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

on $\tilde{v}$, *or simply,* Σ*-restricted, if it has the form*

$$\varphi \wedge iso^{\Sigma}(\tilde{v}) \wedge dif(\tilde{v}).$$

We call φ *the* body *of* ψ *and* $iso^{\Sigma}(\tilde{v}) \wedge dif(\tilde{v})$ *the* Σ-restriction *of* ψ.

We will often use the abbreviation $res^{\Sigma}(\tilde{v})$ for the Σ-restriction $iso^{\Sigma}(\tilde{v}) \wedge dif(\tilde{v})$. According to the above definition, a formula of the form $(\varphi\rho \wedge iso_{\rho})\xi \wedge dif_{\xi}$ (like those seen in Proposition 15), where $\rho \in \mathrm{IN}^{\Sigma}(\tilde{u})$ with $\tilde{u} = \mathcal{V}ar(\varphi)$ and $\xi \in \mathrm{ID}(\mathcal{V}ar(\tilde{v}\rho))$, is in fact a Σ-restricted formula with body $\varphi\rho\xi$ and Σ-restriction $iso_{\rho}\xi \wedge dif_{\xi}$.

All combination results in this paper will require Σ-restricted formulae. Many of them will hold only for formulae Σ-restricted on *all* of their free variables. We call such formulae *totally* Σ*-restricted*. More precisely, a Σ-restricted formula $\varphi \wedge res^{\Sigma}(\tilde{v})$ is totally Σ-restricted if $\mathcal{V}ar(\varphi) \subseteq \tilde{v}$. Notice that closed formulae, and ground formulae in particular, are always totally Σ-restricted for any Σ.

Where $\mathcal{L}$ is a class of formulae and Σ a finite subset of a signature Ω, we will denote by $Res(\mathcal{L}^{\Omega}, \Sigma)$ the class of all the Σ-restricted formulae whose body belongs to $\mathcal{L}^{\Omega}$. Similarly, we will denote by $TRes(\mathcal{L}^{\Omega}, \Sigma)$ the class of all the totally Σ-restricted formulae whose body belongs to $\mathcal{L}^{\Omega}$.

By definition, $\mathcal{L}^{\Omega}$ and $TRes(\mathcal{L}^{\Omega}, \Sigma)$ are always included in $Res(\mathcal{L}^{\Omega}, \Sigma)$. For the common case in which $\mathcal{L}$ is Qff, notice that Qff^{Ω} is usually *strictly* included in $Res(Qff^{\Omega}, \Sigma)$. In fact, unless Σ contains at most constant symbols (or $\tilde{v}$ is empty), the $iso^{\Sigma}(\tilde{v})$ component of every Σ-restricted formula contains universal quantifiers. Finally, notice that when Σ is empty, every $\psi \in Res(\mathcal{L}^{\Omega}, \Sigma)$ is simply of the form $\varphi \wedge dif(\tilde{v})$. Then, $\mathcal{L}^{\Omega}$, $TRes(\mathcal{L}^{\Omega}, \Sigma)$ and $Res(\mathcal{L}^{\Omega}, \Sigma)$ all coincide if $\mathcal{L}$ is closed under conjunction with disequations—as is the case with Qff.

Understanding Σ*-restrictions*

The effect of Σ-restrictions is clear by looking at the definition of iso^{Σ} and dif: they constraint some variables to be distinct Σ-isolated individuals. Given that the notion of Σ-isolated individual is quite technical, what may not be clear at this point is whether Σ-restrictions have a place in common constraint solving practice. We show below that there are situations in which Σ-restrictions arise naturally.

In this discussion, we will identify a Σ-restriction with its iso^{Σ} component and ignore the dif component as the latter is essentially unproblematic for satisfiability concerns. In fact, the satisfiability of a formula $\varphi(\tilde{v})$ is reducible to the satisfiability of the formula

$$(\varphi \wedge dif(\tilde{v}))\xi_1 \vee \cdots \vee (\varphi \wedge dif(\tilde{v}))\xi_n$$

where $\xi_1, \ldots, \xi_n$ are all the (finitely many) identifications of $\tilde{v}$. Therefore, by considering a finite number of identifications we can turn any satisfiability problem into one with additional *dif* constraints without changing its set of solutions. That is not the case for iso^Σ constraints because in general we may need to consider infinitely many Σ-instantiations of the constraint φ; and even that will not be enough if φ is satisfied only by values that are not Σ-generated by Σ-isolated individuals.

Now, as in many applications of logics to computer science, Σ-restrictions are better understood in terms of (data) types, or *sorts*, in logic parlance. Even if classical first-order logic—the one used in this paper—has no explicit notion of sort, we do think of elements in a given domain as naturally grouped in sorts, sets of individuals with common features. Correspondingly, we think of functions as mapping tuples of values of certain sorts to values of some fixed sort, and of relations as subsets of the Cartesian products of certain sorts.[17] We show that under the right—and quite reasonable—conditions, a constraint like $iso^\Sigma(v)$ on a variable v amounts to requiring that the value of v does not belong to a certain sort.

In fact, suppose Ω is the signature of interest and Σ collects only function symbols f of Ω that have some fixed sort S as codomain (i.e., the intended type of f is $S_1 \times \cdots \times S_n \to S$). In every Ω-structure including S in its universe and in which all the elements of S are Σ-generated, the only Σ-isolated individuals are those that do not belong to S. For such structures then, a Σ-restriction of the form $iso^\Sigma(v)$ denotes the restriction that $\alpha(v) \notin S$ for every valuation α of v.

Example 18 With $\Omega := \{\mathsf{0}, \mathsf{s}, \mathsf{nil}, \mathsf{cons}, \mathsf{length}\}$, consider the Ω-structure $\mathcal{A}$ whose universe A is made of pairwise disjoint sorts N, L and I where N is the set of the natural numbers, L the set of the LISP lists over A (including non-nil terminated lists), and I a set of *ill-sorted individuals*. The constants $\mathsf{0}$ and nil are interpreted by $\mathcal{A}$ in the obvious way. The interpretation of the other symbols is such that a) $\mathsf{cons}^\mathcal{A}$ is the injective function behaving as the LISP list constructor and mapping values of A into L as expected, b) $\mathsf{s}^\mathcal{A}$ coincides over N with the successor function and injects the elements of $L \cup I$ into I, c) $\mathsf{length}^\mathcal{A}$ coincides over L with the list length function and injects the elements of $N \cup I$ into I. Now let $\Sigma := \{\mathsf{nil}, \mathsf{cons}\}$. The Σ-isolated individuals of $\mathcal{A}$ are

[17] Notoriously, this picture is complicated by the fact that all functions and relations are total in classical first-order logic and so each first-order structure also has to specify how a function or relation behaves over input values that do not have the intended sort.

exactly the elements of $N \cup I$. Therefore, the Σ-restriction $iso^\Sigma(v)$ is equivalent in $\mathcal{A}$ to the requirement that v is not a list.

The above example provides insights on Σ-instantiations as well. In fact, L contains by construction no circular lists[18]: every list in $\mathcal{A}$ is a (possibly nested) list of *atoms*, the elements of $N \cup I$. In our terminology, this is the same as saying that $\mathcal{A}$ is Σ-generated by its Σ-isolated individuals.

Now, let φ be an Ω-formula satisfiable in $\mathcal{A}$ and assume for simplicity that φ has just one free variable, v. If the value of v that satisfies φ is not a list, then this value is Σ-isolated and so it satisfies $\varphi \wedge iso^\Sigma(v)$ as well. If the value of v is a list, then it can be denoted by some Σ-term $t(\tilde{u})$ whose variables are mapped to non-lists values; these values satisfy the formula $\varphi\rho \wedge iso^\Sigma_\rho(\tilde{u})$ where ρ is the Σ-instantiation $\{v \leftarrow t(\tilde{u})\}$. From this it is not hard to see that an Ω-formula $\varphi(\tilde{v})$ is satisfiable in the structure $\mathcal{A}$ above if and only if there is a $\rho \in \mathrm{IN}^\Sigma(\tilde{v})$ and a $\xi \in \mathrm{ID}(\mathcal{V}ar(\tilde{v}\rho))$ such that $(\varphi_i\rho \wedge iso_\rho)\xi \wedge dif_\xi$.

To conclude this section, we show another structure $\mathcal{B}$ that combines in a natural way LISP lists with some other data-type, is Σ-fusible with the structure $\mathcal{A}$ above, and has a fusion with $\mathcal{A}$ that is Σ-generated by its Σ-isolated individuals.

Example 19 Let $\Delta := \{\mathsf{a}, \mathsf{b}, \cdot, \mathsf{nil}, \mathsf{cons}\}$ and consider the Δ-structure $\mathcal{B}$ whose universe B is made of pairwise disjoint sorts W, L and J, where L is again the set of the LISP lists but over B this time, W is the set of strings over the characters a, b, and J is the set of $\mathcal{B}$'s ill-sorted individuals. The symbols in $\Sigma := \{\mathsf{nil}, \mathsf{cons}\}$ are interpreted by $\mathcal{B}$ in a way similar to that of the previous example. The characters are interpreted as distinct elements of S. The binary symbol $\cdot$ is interpreted an associative operator that behaves over $W \times W$ as string concatenation and maps pairs not in $W \times W$ to elements of J. In this structure, the Σ-isolated individuals are exactly the elements of $W \cup J$.

First we show that $\mathcal{A}$ and $\mathcal{B}$ have a fusion. Observing that the sets $N \cup I$ and $W \cup J$ are both countably infinite, let h be any bijection of the former onto the latter. Recalling that $\mathcal{A}$ is Σ-generated by $N \cup I$, let h_{A-B} be the (necessarily) unique Σ-homomorphic extension of h to A mapping $\mathsf{nil}^\mathcal{A}$ to $\mathsf{nil}^\mathcal{B}$ and $\mathsf{cons}^\mathcal{A}(a_1, a_2)$ to $\mathsf{cons}^\mathcal{B}(h_{A-B}(a_1), h_{A-B}(a_2))$ for all $a_1, a_2 \in A$. One can easily show that h_{A-B} is in fact a bijection of A onto B, which entails that $h_{A-B} : \mathcal{A}^\Sigma \cong \mathcal{B}^\Sigma$. It follows from Proposition 6 that $\mathcal{A}$ and $\mathcal{B}$ have a fusion. Now, let $\mathcal{F}$ be the canonical fusion of $\mathcal{A}$ and $\mathcal{B}$ induced by h_{A-B}. Since $\mathcal{F}^\Sigma$ coincides with $\mathcal{B}^\Sigma$ it is immediate that $\mathcal{F}$ is Σ-generated by its Σ-isolated individuals.

[18] Formally, there are no Σ-terms t such that $(\mathcal{A}, \alpha) \models v \equiv t$ for some $v \in \mathcal{V}ar(t)$ and valuation α.

Finally, we claim that $\mathcal{A}$ and $\mathcal{B}$ are Σ-fusible. But instead of proving it directly, we will do it by applying some general results about the fusibility of *free* structures. But for that we must wait until Section 6.

4 Fusions and Unions of Theories

The combined satisfiability results of the previous section can be lifted from structures to theories. What makes this possible is the close link between fusions and unions of theories, as illustrated in the proposition below. If T_1 and T_2 are two theories, let $Fus(T_1, T_2)$ denote the following class of structures:

$$Fus(T_1, T_2) := \bigcup_{\mathcal{A} \in Mod(T_1),\, \mathcal{B} \in Mod(T_2)} Fus(\mathcal{A}, \mathcal{B}).$$

Proposition 20 *For any theories T_1 and T_2, $Fus(T_1, T_2) = Mod(T_1 \cup T_2)$.*

PROOF. For $i = 1, 2$, let Σ_i be the signature of T_i.

($\subseteq$) Assume that $\mathcal{F}$ is a fusion of some $\mathcal{A} \in Mod(T_1)$ and $\mathcal{B} \in Mod(T_2)$. 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 T_1 and every sentence of T_2. It follows immediately that $\mathcal{F}$ models $T_1 \cup T_2$.

($\supseteq$) Immediate consequence of the obvious fact that any $\mathcal{C} \in Mod(T_1 \cup T_2)$ is a fusion of $\mathcal{C}^{\Sigma_1}$ and $\mathcal{C}^{\Sigma_2}$ and that $\mathcal{C}^{\Sigma_i}$ models T_i, for $i = 1, 2$. $\square$

Recalling Proposition 6 on the existence of fusions, we have the following corollary, first proved, independently, in (Ringeissen, 1996b) and (Tinelli and Harandi, 1996).

Corollary 21 *The union of a Σ_1-theory T_1 and a Σ_2-theory T_2 is consistent iff there is a model of T_1 and a model of T_2 such that their reducts to $\Sigma_1 \cap \Sigma_2$ are isomorphic.*

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

In the rest of the paper, we will be mostly interested in pairs of formulae belonging to the Cartesian product $\mathcal{L}^{\Sigma_1} \times \mathcal{L}^{\Sigma_2}$, for a given class $\mathcal{L}$ of formulae and signatures Σ_1 and Σ_2. For technical reasons we explain later, we will only consider pairs in which at most one of the members is or contains a formula made entirely of *shared* symbols, i.e., symbols in $\Sigma_1 \cap \Sigma_2$. We formalize this restriction in the definition below.

Definition 22 *Where $\mathcal{L}$ is a class of formulae and Σ_1 and Σ_2 two signatures, we call* disjoint product *of $\mathcal{L}^{\Sigma_1}$ and $\mathcal{L}^{\Sigma_2}$ and denote by $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ the following subset of $\mathcal{L}^{\Sigma_1} \times \mathcal{L}^{\Sigma_2}$:*

$$\begin{aligned}\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2} := &\ \{\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \times \mathcal{L}^{\Sigma_2} \mid \textit{no subformula of } \varphi_2 \textit{ is in } \mathcal{L}^{\Sigma_1} \setminus \{\top\}\} \\ \cup &\ \{\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \times \mathcal{L}^{\Sigma_2} \mid \textit{no subformula of } \varphi_1 \textit{ is in } \mathcal{L}^{\Sigma_2} \setminus \{\top\}\}\end{aligned}$$

Since $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ is a subset of $\mathcal{L}^{\Sigma_1} \times \mathcal{L}^{\Sigma_2}$, all of its pairs $\langle \varphi_1, \varphi_2 \rangle$ are such that φ_i contains predicate and function symbols from Σ_i only ($i = 1, 2$). For this reason, we call φ_i the *i-pure* component of $\langle \varphi_1, \varphi_2 \rangle$.[19] For convenience, we say that the pair $\langle \varphi_1, \varphi_2 \rangle$ is satisfiable in a structure (theory) iff $\varphi_1 \wedge \varphi_2$ is satisfiable in the structure (theory).

We are now ready to identify a class of theories whose satisfiability procedures can be combined in a modular way to yield a satisfiability procedure for their union, as we will see in Section 5.

Definition 23 (N-O-combinable Theories) *Let $\mathcal{L}$ be a class of formulae and T_1, T_2 two theories with respective signatures Σ_1, Σ_2 such that $\Sigma := \Sigma_1 \cap \Sigma_2$ is finite.*

- *We say that T_1 and T_2 are* partially N-O-combinable over *$\mathcal{L}$ if Condition 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 T_1 and T_2 are* (totally) N-O-combinable over *$\mathcal{L}$ if both Condition 4.1 and Condition 4.2 below hold for all $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$.*

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

$$\psi_i := (\varphi_i \rho \wedge iso_{\rho})\xi \wedge dif_{\xi}$$

is satisfiable in T_i for $i = 1, 2$, then ψ_i is satisfiable in a model $\mathcal{A}_i$ of 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 $T_1 \cup T_2$, it is satisfiable in a model of $T_1 \cup T_2$ that is Σ-generated by its Σ-isolated individuals.*

[19] Observe that $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ also contains pairs of the form $\langle \varphi_1, \top \rangle$ or $\langle \top, \varphi_2 \rangle$—effectively making every i-pure formula a member of $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$.

While Condition 4.2 is straightforward and easy to understand, Condition 4.1 may be hard to grasp at an intuitive level. To get a better idea it is helpful to concentrate on the case in which ρ is the empty instantiation, as the other cases are reducible to this one. For that case, the condition is roughly saying that if each set $T \cup \{\varphi_i\}$ is satisfied by Σ-isolated individuals, the only way for $T \cup \{\varphi_1\}$ and $T \cup \{\varphi_2\}$ to contradict each other is to disagree on which variables of $\tilde{v}$ get the same value and which don't.

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 T_1 and T_2 of signature Σ_1 and Σ_2, respectively, with $\Sigma := \Sigma_1 \cap \Sigma_2$ non-empty. Then, $\langle \varphi_1 \wedge \varphi,\, \varphi_2 \wedge \neg\varphi \rangle \in \mathcal{L}^{\Sigma_1} \times \mathcal{L}^{\Sigma_2}$ for any $\varphi \in \mathcal{L}^{\Sigma}$, $\varphi_1 \in \mathcal{L}^{\Sigma_1}$ and $\varphi_2 \in \mathcal{L}^{\Sigma_2}$; but it is obvious that, against the requirements of Condition 4.1, for no ρ and ξ is a model of T_1 satisfying $((\varphi_1 \wedge \varphi)\rho \wedge \mathit{iso}_\rho)\xi \wedge \mathit{dif}_\xi$ fusible with a model of T_2 satisfying $((\varphi_2 \wedge \neg\varphi)\rho \wedge \mathit{iso}_\rho)\xi \wedge \mathit{dif}_\xi$.[20]

We point out that even the current definition of $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ could be improved, as it still rules out many theories that one would like to be N-O-combinable. A simple example of such theories is any pair of theories of the form $T_1 \cup T_2$ and $T_2 \cup T_3$ where T_1, T_2 and T_3 are pairwise signature-disjoint. Not all of such pairs are N-O-combinable even if they represent a trivial case of non-disjoint combination. To see that, let

$$T_1 := \{\forall x, y.\ P_1(x, y) \Rightarrow x \equiv y\},\ \ T_2 := \{a \equiv a,\ b \equiv b\},\ \text{and}$$
$$T_3 := \{\forall x, y.\ P_3(x, y) \Rightarrow x \not\equiv y\}.$$

Then consider the pair of pure formulae $\langle P_1(x, y),\, P_3(x, y) \rangle$, the instantiation $\rho := \{x \leftarrow a, x \leftarrow b\}$ and the identification $\xi := \{\}$. Again, models of $T_1 \cup T_2$ satisfying $(P_1(x, y)\rho \wedge \mathit{iso}_\rho)\xi \wedge \mathit{dif}_\xi = P_1(a, b)$ and models $T_2 \cup T_3$ satisfying $(P_3(x, y)\rho \wedge \mathit{iso}_\rho)\xi \wedge \mathit{dif}_\xi = P_3(a, b)$ do exist, but they are obviously not fusible.

In any case, we doubt that the current definition of $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ can be drastically improved unless one renounces to a strictly syntactical definition.

When combining two theories one should make sure that their combination is *meaningful* to start with, that is, it is not inconsistent (or trivial). This is particularly important when one considers, as we do, theories that share non-logical symbols, as it is much easier for such theories to have contradicting consequences. Now, a first consequence of Definition 23 is that N-O-combinable

[20] 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 $T_1 \models \neg\tilde{\exists}(\varphi \wedge \psi)$. Then, for no theory T_2 will $\langle \varphi,\, \psi \rangle$ satisfy Condition 4.1.

consistent theories do have a consistent union, and so it does make sense to combine them.

Proposition 24 *Let T_1 and T_2 be partially N-O-combinable over $\mathcal{L}$. If T_1 and T_2 are consistent, then $T_1 \cup T_2$ is consistent.*

PROOF. Let φ_1 and φ_2 both be $\top$. From an earlier observation we know that $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$. If, for $i = 1, 2$, T_i is consistent, then φ_i is trivially satisfiable in a model of T_i. Observing that $\mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2) = \emptyset$, we can conclude from Condition 4.1 (by considering the empty instantiation and identification) that $\varphi_1 \wedge \varphi_2$ is satisfiable in a fusion of a model of T_1 and a model of T_2. By Proposition 20, this fusion is a model of $T_1 \cup T_2$. $\square$

If the class $\mathcal{L}$ contains disequations of variables, we can show in a similar way that $T_1 \cup T_2$ is non-trivial whenever T_1 and T_2 are N-O-combinable and non-trivial.

N-O-combinable theories are suitable candidates for combination methods for satisfiability thanks to the properties below. Let T_1, T_2, Σ_1, Σ_2, Σ, and $\mathcal{L}$ be as in Definition 23.

Proposition 25 *Assume T_1 and 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 $T_1 \cup T_2$ if there is a $\rho \in \mathrm{IN}^{\Sigma}(\tilde{v})$ and $\xi \in \mathrm{ID}(\mathcal{V}ar(\tilde{v}\rho))$ such that $(\varphi_i \rho \wedge iso_\rho)\xi \wedge dif_\xi$ is satisfiable in T_i for $i = 1, 2$.*

PROOF. Immediate consequence of Condition 4.1, Proposition 15 and Proposition 20. $\square$

If T_1 and T_2 satisfy Condition 4.2 as well, the implication in the proposition above becomes a full equivalence.

Theorem 26 *When T_1 and T_2 are totally N-O-combinable over $\mathcal{L}$ the following 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 $\xi \in \mathrm{ID}(\mathcal{V}ar(\tilde{v}\rho))$ such that, for $i = 1, 2$, $(\varphi_i \rho \wedge iso_\rho)\xi \wedge dif_\xi$ is satisfiable in T_i.
(2) $\varphi_1 \wedge \varphi_2$ is satisfiable in $T_1 \cup T_2$.

PROOF. It is enough to show that $(2 \Rightarrow 1)$. But that is an immediate consequence of Condition 4.2, Proposition 20 and Proposition 16. $\square$

Input: $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$

Instantiation Generate the pair $\langle \gamma_1, \gamma_2 \rangle := \langle \varphi_1\rho \wedge iso_\rho, \varphi_2\rho \wedge iso_\rho \rangle$ for some $\rho \in \mathrm{IN}^\Sigma(\tilde{v})$ with $\tilde{v} := \mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2)$.

Identification Generate the pair $\langle \psi_1, \psi_2 \rangle := \langle \gamma_1\xi \wedge dif_\xi, \gamma_2\xi \wedge dif_\xi \rangle$ for some $\xi \in \mathrm{ID}(\mathcal{V}ar(\tilde{v}\rho))$.

Check Succeed if ψ_1 is satisfiable in T_1 and ψ_2 is satisfiable in T_2. Fail otherwise.

Fig. 1. The Combination Procedure.

We exploit the above properties of N-O-combinable theories in the next section where we describe a sound and complete general procedure for combining constraint reasoners for N-O-combinable theories.

5 Combining Satisfiability Procedures

We show in this section 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. We do this by means of a combination procedure whose correctness relies on the combination results of the previous section.

In the following, we will fix

- a class of formulae $\mathcal{L}$ closed under identification and instantiation of free variables;
- two *countable* signatures Σ_1 and Σ_2 such that $\Sigma := \Sigma_1 \cap \Sigma_2$ is finite;
- a Σ_1-theory T_1 and a Σ_2-theory T_2.

Our combination procedure is defined in Figure 1. It considers the satisfiability in $T_1 \cup T_2$ of formulae from $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ by reducing it non-deterministically to the satisfiability in T_1 and in T_2 of pure Σ-restricted formulae. Given the input problem $\langle \varphi_1, \varphi_2 \rangle$, the procedure first applies to $\langle \varphi_1, \varphi_2 \rangle$ an arbitrary instantiation ρ (into Σ-terms) of the variables shared by φ_1 and φ_2. Then, it applies an arbitrary identification ξ of the shared variables in the new pair. Lastly, it checks that each member $\varphi_i\rho\xi$ of the final pair is satisfiable in the corresponding theory under the restriction $iso_\rho\xi \wedge dif_\xi$, succeeding only when both members are satisfiable.

In essence, the procedure is a non-deterministic version of the Nelson-Oppen combination procedure (Nelson and Oppen, 1979), but it extends that procedure in a number of ways: (1) it does not require that the input formulae be quantifier-free, (2) it does not require (correspondingly) that the component theories be universal, (3) it allows the signatures of the component theories to share up to a finite number of symbols, (4) it considers only identifications over the free variables shared by the two input formulae, whereas Nelson and Oppen's considers identifications over all the variables. The latter improvement is significant for practical computational concerns if not theoretical ones because it reduces the number of possibles choices in the identification and instantiation steps. It has also been considered by Baader and Schulz in their own combination methods, starting with the one described in (Baader and Schulz, 1996).

Proposition 25 immediately tells us that for partially N-O-combinable component theories T_1 and T_2 the procedure in Figure 1 is *sound* in this sense: an input pair $\langle\varphi_1,\, \varphi_2\rangle$ is satisfiable in $T_1 \cup T_2$ if one of the possible outputs of the identification step is a pair $\langle\psi_1, \psi_2\rangle$ such that ψ_i is satisfiable in T_i for $i = 1, 2$. For totally N-O-combinable component theories the procedure is also *complete*: an input pair $\langle\varphi_1,\, \varphi_2\rangle$ is satisfiable in $T_1 \cup T_2$ only if one of the pairs $\langle\psi_1, \psi_2\rangle$ above is such that ψ_i is satisfiable in T_i for $i = 1, 2$.

The formula $\psi_i = (\varphi_i\rho \wedge \mathit{iso}_\rho)\xi \wedge \mathit{dif}_\xi$ $(i = 1, 2)$ generated by the procedure's identification step is a Σ-restricted formula in the sense of Definition 17. More precisely, ψ_i is an element of $\mathit{Res}(\mathcal{L}^{\Sigma_i}, \Sigma)$, given that $\varphi_i \in \mathcal{L}^{\Sigma_i}$ and $\mathcal{L}$ is closed under identification and instantiation. For the check step of procedure to be effective then it must be able to resort, for $i = 1, 2$, to a procedure that decides the satisfiability in T_i of formulae in $\mathit{Res}(\mathcal{L}^{\Sigma_i}, \Sigma)$. In that case, recalling that non-deterministic procedures are said to succeed iff one of their possible runs is successful, we can claim by the above the following result.

Proposition 27 *Assume that T_1 and T_2 be totally N-O-combinable over $\mathcal{L}$ and the satisfiability in T_i of formulae in $\mathit{Res}(\mathcal{L}^{\Sigma_i}, \Sigma)$ is decidable, for $i = 1, 2$. Then, the combination procedure 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 $T_1 \cup T_2$.*

We point our that, contrary to what Proposition 27 might seem to imply, the combination procedure is in general only able to *semi*-decide the satisfiability in $T_1 \cup T_2$ of formulae in $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$. The problem lies in the unbounded non-determinism of the identification step. As we have already observed, whenever Σ contains a function symbol of non-zero arity and the set of variables shared by the two formulae in the input is nonempty, there is an infinite number of possible instantiations over that set. In that case, if the input pair is unsatisfiable in the union theory, by the procedure's soundness, none of these instantiations will lead to formulae ψ_1 and ψ_2 in the check step that are both satisfiable in their respective theory. It follows that the procedure will in general diverge[21] on unsatisfiable inputs.

Note that the procedure can be easily reformulated so that it will not diverge on input pairs containing an i-pure formula that is already unsatisfiable in T_i, and hence in $T_1 \cup T_2$. The non-termination problem arises only for *genuine* combination questions, input pairs that are unsatisfiable in the union theory even if each of their pure members is satisfiable in the corresponding component theory. We will illustrate later some special cases in which the combination procedure can be modified so that it always terminates.

Interestingly, even if it is only a semi-decision procedure, the procedure does yield decidability results when the component theories are axiomatizable.[22] In fact, as pointed out, the procedure will diverge only on those inputs that are *not* satisfiable in the union theory. This means that when the procedure is applicable, the set of pairs satisfiable in the union theory is recursively enumerable. Now, by the completeness of first-order predicate calculus, the set of formulae *un*satisfiable in an axiomatizable theory is also recursively enumerable. It follows that if our procedure is applicable to two theories T_1 and T_2 such that $T_1 \cup T_2$ is axiomatizable, the set of pairs satisfiable in $T_1 \cup T_2$ is recursive. Although this observation does not provide us with a practical decision procedure for satisfiability in $T_1 \cup T_2$, it does lead to the following decidability result—together with the observation that $T_1 \cup T_2$ is axiomatizable whenever both T_1 and T_2 are.

Proposition 28 *Assume that, for $i = 1, 2$, T_i is axiomatizable and the satisfiability in T_i of formulae of $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ is decidable. If T_1 and T_2 are N-O-combinable over $\mathcal{L}$, then the satisfiability in $T_1 \cup T_2$ of formulae in $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ is decidable.*

Up to now, we have used a rather weak language for (mixed) constraints, namely $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$. We have considered only constraints expressible as the conjunction of two pure formulae which, in addition, share non-logical symbols is a very limited way. In general however, combined satisfiability problems are not expressed in the nice separated format given by $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$, but rather as mixed constraints in $\mathcal{L}^{\Sigma_1 \cup \Sigma_2}$. Our combination results would certainly be more useful then if they could be given in terms of $\mathcal{L}^{\Sigma_1 \cup \Sigma_2}$ instead. This is in fact

[21] Strictly speaking, we should say something like: "it will infinitely fail". 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.

[22] A theory is *axiomatizable* if its deductive closure coincides with the deductive closure of a recursive set of sentences.

possible, but at the cost of some closure assumptions on $\mathcal{L}$.[23] We describe such assumptions in the following and then show, as an example, how they let us improve on Proposition 28.

Definition 29 *Given two signatures Ω_1 and Ω_2, we say that a class $\mathcal{L}$ of formulae is* purifiable w.r.t. $\langle\Omega_1, \Omega_2\rangle$ *if for every $\varphi \in \mathcal{L}^{\Omega_1\cup\Omega_2}$, there is a finite set $\{\langle\varphi_j^1, \varphi_j^2\rangle\}_{j<m} \subseteq \mathcal{L}^{\Omega_1} \otimes \mathcal{L}^{\Omega_2}$ such that*

(1) $\varphi_j^1 \wedge \varphi_j^2 \in \mathcal{L}^{\Omega_1\cup\Omega_2}$ for all $j < m$,
(2) φ and $\bigvee_{j<m}(\varphi_j^1 \wedge \varphi_j^2)$ are equisatisfiable.

We call $\bigvee_{j<m}(\varphi_j^1 \wedge \varphi_j^2)$ a disjunctive pure form of φ (w.r.t. $\langle\Omega_1, \Omega_2\rangle$). *We say that $\mathcal{L}$ is* effectively *purifiable w.r.t. $\langle\Omega_1, \Omega_2\rangle$ if for each formula $\varphi \in \mathcal{L}^{\Omega_1\cup\Omega_2}$, a disjunctive pure form of φ is effectively computable.*

If the class $\mathcal{L}$ specified at the beginning of this section is effectively purifiable with respect to our initial pair of signatures $\langle\Sigma_1, \Sigma_2\rangle$, we can modify the combination procedure of Figure 1, by adding a "preprocessing" step that, given an input formula φ from $\mathcal{L}^{\Sigma_1\cup\Sigma_2}$, computes a disjunctive pure form ψ of φ and the returns—in a *don't know* non-deterministic way—one of ψ's disjuncts.

Given that φ is satisfiable in $T_1 \cup T_2$ if and only if some disjunct of its disjunctive pure form is satisfiable in $T_1 \cup T_2$, it is immediate that the new procedure is correct as well. With the new procedure we can then conclude by Proposition 28 that when T_1 and T_2 are N-O-combinable over $\mathcal{L}$ and the satisfiability in T_i of formulae of $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ is decidable for $i = 1, 2$, then the satisfiability in $T_1 \cup T_2$ of formulae of $\mathcal{L}^{\Sigma_1\cup\Sigma_2}$ is also decidable. As a matter of fact, we can prove something a little stronger, going from the satisfiability in $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ to the satisfiability in $Res(\mathcal{L}^{\Sigma_1\cup\Sigma_2}, \Sigma)$.

Proposition 30 *Assume that $\mathcal{L}$ is effectively purifiable w.r.t. $\langle\Sigma_1, \Sigma_2\rangle$, T_1 and T_2 are N-O-combinable over $\mathcal{L}$, and T_i is axiomatizable for $i = 1, 2$. If the satisfiability in T_i of formulae of $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ is decidable, then the satisfiability in $T_1 \cup T_2$ of formulae of $Res(\mathcal{L}^{\Sigma_1\cup\Sigma_2}, \Sigma)$ is also decidable.*

The result above is interesting not only because it allow us to work in $\mathcal{L}^{\Sigma_1\cup\Sigma_2}$, as opposed to $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$, but also because it can lead to decidability results for more than two theories by iteration. Suppose in fact that, in addition to the theories in the proposition, there is a third axiomatizable theory T_3 of signature Σ_3 whose common signature with $T_1 \cup T_2$ is also Σ and for which the satisfiability of formulae of $Res(\mathcal{L}^{\Sigma_3}, \Sigma)$ is decidable. Then, if $\mathcal{L}$ is effectively purifiable w.r.t. $\langle\Sigma_1 \cup \Sigma_2, \Sigma_3\rangle$ and $T_1 \cup T_2$ and T_3 are N-O-combinable over $\mathcal{L}$, by the above, the satisfiability in $T_1 \cup T_2 \cup T_3$ of formulae of $Res(\mathcal{L}^{\Sigma_1\cup\Sigma_2\cup\Sigma_3}, \Sigma)$

[23] Notice that we have hardly made any assumptions on $\mathcal{L}$ so far.

is also decidable.

Proving Proposition 30 is easy but tedious. The following informal argument should suffice. Recall that given a formula φ, the new combination procedure first purifies it 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 $\mathit{iso}_\rho\xi \wedge \mathit{dif}_\xi$ before passing the pair 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 for possible instantiation and identification.[24] Now, if the input φ is already of the form $\varphi \wedge \mathit{res}^\Sigma(\tilde{v})$ with $\varphi \in \mathcal{L}^{\Sigma_1\cup\Sigma_2}$, it is enough for the procedure to purify φ into $\langle\varphi_1, \varphi_2\rangle$ and then generate the formulae $(\varphi'_i\rho \wedge \mathit{iso}_\rho)\xi \wedge \mathit{dif}_\xi$ as before with the only differences that φ'_i is now $\varphi_i \wedge \mathit{res}^\Sigma(\tilde{v})$, ρ is chosen so that it does not instantiate any variables in $\tilde{v}$, and ξ is chosen so that it does not identify any two variables in $\tilde{v}$. It is a simple exercise to show that each $(\varphi'_i\rho \wedge \mathit{iso}_\rho)\xi \wedge \mathit{dif}_\xi$ can be effectively reduced[25] to a logically equivalent formula in $\mathit{Res}(\mathcal{L}^{\Sigma_i}, \Sigma)$, which can then be processed by T_i's satisfiability procedure.

5.1 An Effectively Purifiable Class of Formulae

We conclude this section by showing that an important class of formulae, the quantifier-free formulae, is effectively purifiable w.r.t. any pair of signatures. For that we first need to give a precise definition to some concepts we have been using only informally so far.

Let us fix again two arbitrary countable signatures Σ_1 and Σ_2 and let $\Sigma := \Sigma_1 \cap \Sigma_2$. We call *shared symbols* the elements of Σ and *shared terms* the elements of $T(\Sigma, V)$. Observe that when Σ is empty, the only shared terms are the variables. We call (strict) 1-*symbols* the elements of Σ_1 ($\Sigma_1 \setminus \Sigma$) and (strict) 2-*symbols* the elements of Σ_2 ($\Sigma_2 \setminus \Sigma$). Shared symbols are both 1- and 2-symbols, and they are strict for neither signature. A term $t \in T(\Sigma_1 \cup \Sigma_2, V)$ is an i-*term* iff its top symbol $t(\epsilon)$ is an element of $V \cup \Sigma_i$ ($i = 1, 2$). Variables and terms t with top symbol in $\Sigma_1 \cap \Sigma_2$ are both 1- and 2-terms. For $i = 1, 2$, an i-term is *pure* iff it contains only i-symbols and variables.

There is a standard *purification procedure* that when Σ_1 and Σ_2 are disjoint can convert any set S of literals of signature $\Sigma_1 \cup \Sigma_2$ into a set of pure literals (see (Baader and Schulz, 1995a) among others). The purification process is achieved by replacing "alien" subterms by new variables and adding appropriate new equations to S. Intuitively, an alien subterm of an i-term t is a maximal

[24] Considering only shared variables is in a sense an optimization of this more general case.
[25] Exploiting the associativity, commutativity, and idempotency of $\wedge$.

subterm of t that is not itself an i-term. The gist of the procedure then is to abstract by a fresh variable v_s each alien subterm s of an atom in S and add the equation $v_s \equiv s$ to S. The abstraction process is applied repeatedly to S until no more subterms can be abstracted. This procedure always terminates and produces a set of literals that is satisfiable in a $(\Sigma_1 \cup \Sigma_2)$-structure $\mathcal{A}$ iff the original set S is satisfiable in $\mathcal{A}$.

Now, for disjoint Σ_1 and Σ_2 a formal definition of the notion of alien subterm to be used by the purification procedure is straightforward. If one allows Σ_1 and Σ_2 to share symbol, however, things gets tricky because one has to decide how to consider shared function symbols (see (Baader and Tinelli, 2001) for a detailed discussion). We adopt the following definition among a number of possible ones.

Definition 31 (Alien subterms) *Let $t \in T(\Sigma_1 \cup \Sigma_2, V)$. If the top symbol of t is a strict i-symbol, then a subterm s of t is an* alien subterm *of t iff it is not an i-term and it is maximal with this property, i.e., every proper superterm of s in t is an i-term.*

If the top symbol of t is a shared symbol, then for $i = 1, 2$, let S_i be the set of all (proper) maximal subterms of t whose top symbol is a strict i-symbol.

- *If $S_1 \neq \emptyset$, then t is considered to be a 1-term, i.e., a subterm s of t is an* alien subterm *of t iff it is not a 1-term and it is maximal with this property.*
- *If $S_1 = \emptyset$ and $S_2 \neq \emptyset$, then t is considered to be a 2-term, i.e., a subterm s of t is an* alien subterm *of t iff it is not a 2-term and it is maximal with this property.*[26]

We extend the definition of alien subterm from terms to atomic formulae by treating the formula's predicate symbol as if it was a function symbol—with the equality symbol being treated a shared symbol.

With this definition of alien subterm, the purification procedure described earlier can be applied, unchanged and with the same results, to a set of $(\Sigma_1 \cup \Sigma_2)$-literals regardless of whether Σ_1 and Σ_2 are disjoint or not. Relying on this procedure, we can finally show the following.

Proposition 32 *The class Qff of quantifier-free formulae is effectively purifiable w.r.t. $\langle \Sigma_1, \Sigma_2 \rangle$.*

PROOF. Let $\varphi \in Qff^{\Sigma_1 \cup \Sigma_2}$. We first convert φ into its disjunctive normal form, a logically equivalent formula of the form $\bigvee_{j<m} \varphi_j$, where every disjunct φ_j is a conjunction of literals. Then, for each $j < m$, we apply the purification

[26] If $S_1 = \emptyset$ and $S_2 = \emptyset$, then t is pure and so it has no aliens subterms.

procedure to the set of literals in φ_j and produce a set S_j of pure literals. Finally, we collect the Σ_1-literals of S_j into a conjunction φ_j^1 and the Σ_2-literals of S_j into a conjunction φ_j^2, making sure that Σ-literals are either all collected in φ_j^1 or all collected in φ_j^2. This process is clearly effective. Furthermore, it is easy to verify that $\bigvee_{j<m}(\varphi_j^1 \wedge \varphi_j^2)$ is a disjunctive pure form of φ. $\square$

Incidentally, notice that even if the process described in the proof above is non-deterministic (because of the choice of where to collect shared literals), for our purposes this is a *don't-care* kind of non-determinism since all the disjunctive pure forms that can be obtained this way are equisatisfiable with the original formula.

6 Identifying N-O-combinable Theories

The combination method presented in the previous section applies correctly to pairs of N-O-combinable theories. Now, as defined in Definition 23, N-O-combinability is a rather abstract notion, expressing conditions not on the single theories but on both of them as a pair. As a consequence, it is not immediate to see whether two given theories are N-O-combinable.

In this section, we attempt to establish sufficient conditions for N-O-combinability that are less abstract and more "local" to the theories. As we will see, our attempts are only partially successful. More research, and maybe new insights, on this are needed. Once again, it will be beneficial to start with the simple case of theories with disjoint signatures, and then move to the general case.

6.1 Disjoint Signatures

A sufficient, and local, condition for the N-O-combinability of two signature-disjoint theories over the language of quantifier-free formulae has been known for quite some time. It was introduced in (Oppen, 1980) to justify the correctness of the Nelson-Oppen combination method. There, each theory T_i is required to be *stably-infinite*, that is, universal and such that every quantifier-free formula satisfiable in T_i is satisfiable in an infinite model of T_i. In the following, we show that the notion of stable-infiniteness can be extended to arbitrary theories and parameterized by the language of interest. Then, we use this extended and parameterized notion to show how the original combination results by Nelson and Oppen are subsumed by ours.

Looking back at Lemma 10 one realizes that, with disjoint signatures, all is needed for the combination result there is that the component structures that satisfy the pure formulae have the same cardinality. One way to guarantee this with theories is to restrict one's attention to those satisfying the following property.

Definition 33 (Stably-Infinite Theory) *Let $\mathcal{L}$ be a class of formulae and T a consistent theory of signature Ω. We say that T is* stably-infinite over $\mathcal{L}^\Omega$ *iff every formula of $\mathcal{L}^\Omega$ satisfiable in T is satisfiable in an infinite model of T.*

It is immediate that complete theories admitting infinite models are stably-infinite over the whole language of first-order formulae. In (Baader and Tinelli, 1997), it is shown that equational theories augmented with the non-triviality axiom $\exists x \exists y. x \not\equiv y$ are stably infinite over the class of quantifier-free formulas. We prove below that this result can be generalized to any theory axiomatized by Horn sentences. [27]

Proposition 34 *Every consistent Horn theory T of signature Ω such that $T \models \exists x \exists y. x \not\equiv y$ is stably infinite over $\mathcal{L}^\Omega$, where $\mathcal{L}$ is the class of Horn formulae or the class of quantifier-free formulae.*

PROOF. Let $\mathcal{L}$ be the class of Horn formulae first and φ a member of $\mathcal{L}^\Omega$ satisfiable in T. It is enough to show that the theory $T' := T \cup \{\tilde{\exists}\, \varphi\}$ has an infinite model.

Observe that $\tilde{\exists}\, \varphi$ is a Horn sentence, which entails that T' is Horn theory as well. From the consistency of T and the assumption that $T \models \exists x \exists y. x \not\equiv y$, we know that T' admits a non-trivial model $\mathcal{A}$. By a result originally due to Alfred Horn, the class of models of a Horn theory is closed under direct products (see, e.g. (Hodges, 1993)). This means that the direct product $\mathcal{B}$ of $\mathcal{A}$ with itself countably infinitely many times, say, is a model of T'. Now, $\mathcal{B}$ is infinite by definition of direct product and the fact that the set A has at least two elements.

If $\mathcal{L}$ is *Qff*, we can prove the claim by reduction to the previous case, observing that a quantifier-free formula is satisfiable in T iff one of the disjuncts of its disjunctive normal form is, and that conjunctions of literals are Horn formulae. $\square$

Some specific examples of stably-infinite theories, useful in software verification, can be found in (Oppen, 1980).

---

[27] A Horn formula is a first-order formula of the form $Q.\, \varphi_1 \wedge \cdots \wedge \varphi_n$, where Q is an *arbitrary* quantifier prefix and each φ_i is a disjunction of literals other than $\perp$ and $\neg\top$, at most one of which is positive.

One consequence of Definition 33 is that stably-infinite theories admit infinite models and so, by the Upward and Downward Löwenheim-Skolem theorems (Hodges, 1993), admit models of any infinite cardinality[28]. 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 an application of Corollary 21), that the union of two stably-infinite, signature-disjoint theories is always consistent. In addition, for classes of formulae closed under variable identification we have the following.

Proposition 35 *Let $\mathcal{L}$ be a class of formulae closed under variable identification and T_1, T_2 two theories with respective signatures Σ_1, Σ_2 such that $\Sigma := \Sigma_1 \cap \Sigma_2 = \emptyset$. If T_i is stably-infinite over $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ for $i = 1, 2$, then T_1 and T_2 are totally N-O-combinable over $\mathcal{L}$.*

PROOF. First we show that T_1 and T_2 satisfy Condition 4.1. Let $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$, $\tilde{v} := \mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2)$, $\rho \in \mathrm{IN}^{\Sigma}(\tilde{v})$ and $\xi \in \mathrm{ID}(\mathcal{V}ar(\tilde{v}\rho))$. Now, each $(\varphi_i \rho \wedge iso_\rho)\xi \wedge dif_\xi$ is logically equivalent to the formula $\psi_i := \varphi_i \xi \wedge dif_\xi$ since ρ necessarily coincides with the empty instantiation (as $\Sigma = \emptyset$) and iso_ρ with the empty set. Given that $\mathcal{L}$ is closed under variable identification, it is immediate that $\psi_i \in Res(\mathcal{L}^{\Sigma_i}, \Sigma)$. From the stable-infiniteness of T_i it follows that if ψ_i is satisfiable in T_i, it is satisfiable in a model $\mathcal{A}_i$ of T_i of cardinality κ, for any infinite κ greater than or equal to the cardinality of $\Sigma_1 \cup \Sigma_2$. We have already seen that structures like $\mathcal{A}_1$ and $\mathcal{A}_2$ are trivially Σ-fusible.

To see that T_1 and T_2 satisfy Condition 4.2 as well, simply notice that since Σ is empty, every individual of any model of $T_1 \cup T_2$ is Σ-isolated. $\square$

As a consequence of the above proposition, we obtain the following simplified version of Theorem 26.

Theorem 36 *Let $\mathcal{L}$ a class of formulae closed under variable identification and T_1, T_2 two theories with disjoint signatures Σ_1, Σ_2, respectively. For $i = 1, 2$, assume that T_i is stably-infinite over $Res(\mathcal{L}^{\Sigma_i}, \emptyset)$ and let $\varphi_i \in \mathcal{L}^{\Sigma_i}$. Then, where $\tilde{v} := \mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2)$, the following are equivalent:*

(1) $\varphi_i \xi \wedge dif_\xi$ is satisfiable in T_i for each $i = 1, 2$ and some $\xi \in \mathrm{ID}(\tilde{v})$;
(2) $\varphi_1 \wedge \varphi_2$ is satisfiable in $T_1 \cup T_2$.

The soundness and completeness of the Nelson-Oppen combination method (in the case of two component theories) can be proved by an application of the theorem above, observing that the class *Qff* is closed under variable

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

identification and that $Res(Qff^{\Omega}, \emptyset)$ coincides with Qff^{Ω} for any signature Ω. See (Ringeissen, 1996b) or (Tinelli and Harandi, 1996) for more details.

6.2 Non-disjoint Signatures

We now move to the question of finding local sufficient conditions for N-O-combinability for theories that might share function or predicate symbols. We first focus on the problem of showing that two theories are partially N-O-combinable (that is, satisfying Condition 4.1). Then we consider what extra conditions must hold for them to be totally N-O-combinable (that is, to satisfy Condition 4.2 as well).

In the previous subsection, to provide sufficient conditions for N-O-combinability we looked for restrictions that would guarantee the existence of fusible models. There, it was enough to have restrictions that guaranteed the existence of two models with the same cardinality. Now that the theories' signatures may have a non-empty intersection Σ, the two models must be Σ-fusible (cf. Definition 14). The question then is: what structures are Σ-fusible?

A sufficient condition for two structures to be Σ-fusible is that their Σ-reducts are *free* in the same variety over the same set of generators. We will prove this fact in the following and use it to define a general class of N-O-combinable theories. Before that, we present the definition and the properties of free structures that we will need.

6.2.1 Free Structures

The concept of free structure is a natural extension to first-order logic of the concept of free algebra from Universal Algebra. We adopt the following among the many (equivalent) definitions in the literature.

Definition 37 (Free Structure) *Given a class* $\mathbf{K}$ *of* Σ*-structures and a set* X*, a* Σ*-structure* $\mathcal{A}$ *is* free for $\mathbf{K}$ over X *iff*

(1) $\mathcal{A}$ *is generated by* X*;*
(2) every map from X *into the universe of a structure* $\mathcal{B} \in \mathbf{K}$ *extends to a (necessarily unique) homomorphism of* $\mathcal{A}$ *into* $\mathcal{B}$*.*

We say that $\mathcal{A}$ *is free* in $\mathbf{K}$ *over* X *(or free over* X *in* $\mathbf{K}$*) if* $\mathcal{A}$ *is free for* $\mathbf{K}$ *over* X *and* $\mathcal{A} \in \mathbf{K}$*. In either case, we call* X *a* basis *of* $\mathcal{A}$*.*

For convenience, given a Σ-theory T, we will sometimes say that $\mathcal{A}$ *is free over* X *in* T, if $\mathcal{A}$ is free over X in $Mod(T)$. In that case, we will also say that $\mathcal{A}$

is a free model of T.[29]

It is immediate from the above 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). A structure free over an empty basis is called *initial.*[30] A Σ-structure free in the class of all Σ-structures is called *absolutely free.*

We will also use the following characterization of freeness.

Proposition 38 ((Hodges, 1993)) *Let $\mathbf{K}$ be a class of Σ-structures, $\mathcal{A}$ a Σ-structure, and X a subset of A. Then, $\mathcal{A}$ is free for $\mathbf{K}$ over X iff*

(1) X generates $\mathcal{A}$ and
(2) $\mathbf{K} \models \tilde{\forall}\,\varphi$ for all Σ-atoms $\varphi(\tilde{v})$ such that $\mathcal{A} \models \varphi[\tilde{x}]$ for some sequence $\tilde{x}$ of pairwise distinct elements of X.

Free models with infinite bases are *canonical* for atomic formulae, in the sense specified by the following corollary of Proposition 38.

Corollary 39 *Let T be a theory of signature Σ and $\mathcal{A}$ a Σ-structure free in T over an infinite basis. Then, for all atomic Σ-formulae φ,*

$$\mathcal{A} \models \tilde{\forall}\,\varphi \quad \textit{iff} \quad T \models \tilde{\forall}\,\varphi.$$

Equivalently, the atomic theory of $\mathcal{A}$ coincides with the atomic theory of T.

It is possible to show that every basis of a free structure is non-redundant as a set of generators, and that a structure can be free over more than one basis (Hodges, 1993). Free structures in a collapse-free class, however, have unique bases.

Proposition 40 *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. Let $\mathcal{A}$ be a Σ-structure 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 X also contains a non-isolated individual y. Since y is not isolated and X generates $\mathcal{A}$,

[29] To avoid misunderstandings, notice that for $\mathcal{A}$ to be a free model of T it is not enough that $\mathcal{A}$ is a model of T free for some class. It must be free for the class *Mod*(T).

[30] This definition is equivalent to the more common definition of initial structure according to which a structure $\mathcal{A}$ is initial (in a class $\mathbf{K}$) if, for all structures $\mathcal{B} \in \mathbf{K}$, there is a unique homomorphism from $\mathcal{A}$ into $\mathcal{B}$.

there is a non-variable Σ-term $t(\tilde{v})$ and a sequence $\tilde{x}$ in X with no repetitions such that $y = t^{\mathcal{A}}[\tilde{x}]$. [31]

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 Proposition 38 then, the sentence $\tilde{\forall}\,(u \equiv t)$ is entailed by the class, against the assumption that the class is collapse-free. $\square$

Free structures have a close connection with varieties. In fact, every non-trivial Σ-variety contains structures free in it. Furthermore, every free Σ-structure is free in some Σ-variety (Hodges, 1993), and in particular, absolutely free Σ-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 41 *Let $\mathbf{K} := Mod(T)$ for some Σ-theory T. For all $\mathcal{A} \in \mathbf{K}$ and $X \subseteq A$, if $\mathcal{A}$ is free in $Mod(T)$ over X then $\mathcal{A}$ is free in $Mod(\mathrm{At}(T))$ over X.*

PROOF. 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 Proposition 38, it is enough to show that $\mathrm{At}(T) \models \tilde{\forall}\varphi$. By assumption and thanks to the same proposition, we know that $T \models \tilde{\forall}\,\varphi$. Recalling the definition of $\mathrm{At}(T)$, we can then conclude that $\tilde{\forall}\,\varphi \in \mathrm{At}(T)$, from which the claim follows immediately. $\square$

The above result also entails that a free Σ-structure with an infinite basis is free (over that basis) in its own Σ-variety $Mod(H)$, where H is the set of all the Σ-atoms modeled by $\mathcal{A}$.

The free structures of a variety can be identified modulo isomorphism according to the following immediate consequence of Definition 37.

Lemma 42 *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 $\mathcal{B}$.*

We are now ready to prove our earlier claim on the fusibility of structures with a free Σ-reduct.

Proposition 43 *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 Σ-fusible.*

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

PROOF. We start by showing that $\mathcal{A}$ and $\mathcal{B}$ are fusible over $\langle X, Y\rangle$. Given a finite set $X_0 \subseteq X$, consider any injective map $h\colon X_0 \to Y$. Since X_0 is finite and $Card(X) = Card(Y)$, h can always be extended to a bijection from X onto Y. By Lemma 42 then, h can be extended to an isomorphism of $\mathcal{A}^\Sigma$ onto $\mathcal{B}^\Sigma$. To see that $\mathcal{A}$ and $\mathcal{B}$ are Σ-fusible, recall that the isolated individuals of a structure are included in every set that generates that structure. Since X generates $\mathcal{A}^\Sigma$ and Y generates $\mathcal{B}^\Sigma$ by assumption, we have that $Is(\mathcal{A}^\Sigma) \subseteq X$ and $Is(\mathcal{B}^\Sigma) \subseteq Y$, from which the claim follows. □

Notice that in the result above the Σ-reducts of the structures are required to be free, not the whole structures. Also notice that this is indeed a generalization of the signature-disjoint case. In fact, when Σ is empty the Σ-reduct of any structure is (trivially) free over the whole carrier of the structure.

A pair of structures that satisfy the proposition above are the structures seen in Example 18 and Example 19 of Section 3. The structure $\mathcal{A}$ in the first example combined natural numbers and LISP lists, whereas the structure $\mathcal{B}$ in the second example combined strings and LISP lists. Recall that, as data structures, two LISP lists are equal if and only if they are both nil or are both non-nil and have equal head and tail. Mathematically, this means that an equation between two terms in the signature $\Sigma := \{\mathsf{nil}, \mathsf{cons}\}$ is valid in $\mathcal{A}^\Sigma$ (or $\mathcal{B}^\Sigma$) if and only if the two terms are identical. From the fact that, as we have seen in the examples, $\mathcal{A}^\Sigma$ is generated by the set $N \cup I$ and $\mathcal{B}^\Sigma$ is generated by the set $W \cup J$, it easily follows that they are both free in the empty Σ-theory, respectively over $N \cup I$ and $W \cup J$. Since both $N \cup I$ and $W \cup J$ are countably infinite, we can conclude by Proposition 43 that $\mathcal{A}$ and $\mathcal{B}$ are Σ-fusible.

6.2.2 Stably Σ-free Theories

We can use Proposition 43 to extend the notion of stable-infiniteness so that it provides, along with some additional requirements, a sufficient condition for the N-O-combinability of theories with non-disjoint signatures.

Definition 44 (Stably Σ-free Theory) *Let T be a consistent theory of signature Ω, Σ a finite subset of Ω, $\mathcal{L}$ a class of formulae and κ the first infinite cardinal such that $\kappa \geq Card(\Omega)$. The theory T is* stably Σ-free over $\mathcal{L}^\Omega$ *iff every formula of $\mathcal{L}^\Omega$ satisfiable in T is satisfiable in a model $\mathcal{A}$ of T such that $\mathcal{A}^\Sigma$ is free in $Mod(\mathrm{At}(T^\Sigma))$, the Σ-variety of T, over a basis of cardinality κ.*

As said, the notion of stable Σ-freeness is meant to generalize that of stable-infiniteness for pairs of theories whose shared signature is Σ. Indeed, when Σ is empty the two notions coincide.

Proposition 45 *Let $\mathcal{L}$ be a class of formulae, T a consistent theory of signature Ω, and Σ an empty signature. Then, T is stably-infinite over $\mathcal{L}^\Omega$ iff T is stably Σ-free over $\mathcal{L}^\Omega$.*

PROOF. Let κ be the first infinite cardinal such that $\kappa \geq \mathit{Card}(\Omega)$.

($\Rightarrow$) Assume that T is stably-infinite over $\mathcal{L}^\Omega$ and let $\psi \in \mathcal{L}^\Omega$ be satisfiable in T. By definition of stable-infiniteness, $T \cup \{\tilde{\exists}\, \psi\}$ has an infinite model and so, as observed earlier, one of cardinality κ. Call it $\mathcal{A}$ and notice that $\mathcal{A}^\Sigma$ is absolutely free over A. Moreover, the atomic Σ-theory of T is empty. In fact, since Σ has no symbols, the only non-empty atomic Σ-theory is the one axiomatized by $\{\forall x \forall y.\ x \equiv y\}$. However, this cannot be the Σ-theory of T because otherwise all of T's models would be trivial, against the fact that T has an infinite model. It conclusion, we have shown that ψ is satisfiable in a model of T whose reduct to Σ is free in the Σ-variety of T over a basis of cardinality κ.

($\Leftarrow$) Assume that T is stably Σ-free over $\mathcal{L}^\Omega$ and let $\psi \in \mathcal{L}^\Omega$ be satisfiable in T. By Definition 44, ψ is satisfiable in a model of T containing at least κ individuals and so it is satisfiable in an infinite model of T. $\square$

We will see in Section 8 that the class of stably Σ-free theories is non-empty for all signatures Σ. For now, it might be interesting to see how a stably-infinite theory can fail to be stably Σ-free when Σ is non-empty.

Example 46 *Consider the Ω-theory $T := \{a \not\equiv b,\ c \not\equiv d \vee a \equiv d\}$ where a, b, c and d are constant symbols. It is easy to see that T is a consistent Horn theory entailing $\exists x \exists y. x \not\equiv y$. Therefore, it is stably infinite over Qff^Ω by Proposition 34. Now let $\Sigma := \{a, d\}$ and observe that the atomic Σ-theory of T is empty. Then consider any model of T satisfying the quantifier-free formula $c \equiv d$. Such a model exists because $T \cup \{c \equiv d\}$ is consistent, as one can easily see. Moreover, in it a and d are equal. Now, the model's reduct to Σ is certainly not free in the Σ-variety of T, otherwise the atomic Σ-theory of T would contain the equation $a \equiv b$. It follows that T is not stably Σ-free over Qff^Ω.*

We show below that under certain conditions stably Σ-free theories are N-O-combinable. To do that we will fix

- a class $\mathcal{L}$ of formulae closed under identification and instantiation and
- two countable signatures Σ_1 and Σ_2 [32] such that $\Sigma := \Sigma_1 \cap \Sigma_2$ is finite.

[32] All we need really is that Σ_1 and Σ_2 have the same cardinality whenever one of them is not countable. We assume that they are both countable for simplicity.

Lemma 47 *Let T_1, T_2 be two consistent theories of respective signature Σ_1, Σ_2, and H_0 an atomic theory of signature Σ. If H_0 is the atomic Σ-theory of both T_1 and T_2 and each T_i is stably Σ-free over some class of formulae, then H_0 is also the atomic Σ-theory of $T_1 \cup T_2$.*

PROOF. Let $T := T_1 \cup T_2$. It is immediate that $H_0 \subseteq \mathrm{At}(T^\Sigma)$. We show that $\mathrm{At}(T^\Sigma) \subseteq H_0$. First recall that we assume that every class of formulae contains a universally true sentence. Together with Definition 44, this entails that for $i = 1, 2$, T_i has a model $\mathcal{A}_i$ whose Σ-reduct is free in H_0 over a countably-infinite set. It follows by Proposition 43 and Proposition 20 that $\mathcal{A}_1$ and $\mathcal{A}_2$ are fusible in a model $\mathcal{F}$ of T. 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 H_0 (over some countably infinite set).

Now, let $\tilde{\forall}\, \varphi \in \mathrm{At}(T^\Sigma)$, which means that φ is a Σ-atom such that $T \models \tilde{\forall}\, \varphi$. Then, $\mathcal{F}^\Sigma \models \tilde{\forall}\, \varphi$ as well because $\mathcal{F}$ is a model of T and $\tilde{\forall}\, \varphi$ is a Σ-formula. Since $\mathcal{F}^\Sigma$ is a free model of H_0 with an infinite basis, we have by Corollary 39 that $H_0 \models \tilde{\forall}\varphi$. Recalling that H_0 is the atomic Σ-theory of T_1, we can conclude that $\tilde{\forall}\, \varphi \in H_0$. $\square$

Theorem 48 *For all consistent theories T_1, T_2 of respective signature Σ_1, Σ_2, we have the following.*

(1) If T_1 and T_2 have the same atomic Σ-theory H_0 and each T_i is stably Σ-free over $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$, then T_1 and T_2 are partially N-O-combinable over $\mathcal{L}$.

(2) If, in addition, H_0 is collapse-free and $T_1 \cup T_2$ is stably Σ-free over $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$, then T_1 and 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 Condition 4.1. Let $\rho \in \mathrm{IN}^\Sigma(\tilde{v})$ and $\xi \in \mathrm{ID}(\mathcal{V}ar(\tilde{v}\rho))$ such that $\psi_i := (\varphi_i \rho \wedge iso_\rho)\xi \wedge dif_\xi$ is satisfiable in T_i for $i = 1, 2$. We already know that ψ_i belongs to $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$; therefore, by the stable Σ-freness of T_i, it is satisfiable in some $\mathcal{A}_i \in Mod(T_i)$ such that ${\mathcal{A}_i}^\Sigma$ is free in $Mod(H_0)$ over a countably-infinite set X_i. The models $\mathcal{A}_1$ and $\mathcal{A}_2$ are Σ-fusible by Proposition 43.

(2) It suffices to show that $\langle \varphi_1, \varphi_2 \rangle$ satisfies Condition 4.2. Let $T := T_1 \cup T_2$ and assume that $\langle \varphi_1, \varphi_2 \rangle$ is satisfiable in T. As T is stably Σ-free 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 T whose reduct to Σ is free in the Σ-variety of T. Since the Σ-variety of T is $Mod(H_0)$ by Lemma 47, and H_0 is collapse-free by assumption, we have by Proposition 40 that $\mathcal{A}^\Sigma$ is

generated by its isolated individuals. In conclusion, $\varphi_1 \wedge \varphi_2$ is satisfiable in a model of T that is Σ-generated by its Σ-isolated individuals. $\square$

Total (as opposed to partial) N-O-combinability of the component theories is important for our combination method because it guarantees its completeness, as we have seen in Section 5. An irksome feature of the theorem above is that it explicitly assumes that $T_1 \cup T_2$ is stably Σ-free over $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ in order to yield the total N-O-combinability of T_1 and T_2.

It would be better if the stable Σ-freeness of a union theory could be proved from the stable Σ-freeness of its component theories. Unfortunately, we have not been able to do that. In fact, we believe that it is unlikely to be the case in general. More constraints on either the language or the component theories are needed. For instance, it is possible to show that if Σ is empty, then $T_1 \cup T_2$ is indeed stably Σ-free over $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ whenever both T_1 and T_2 are stably Σ-free over $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$.

Although we are not able to show in general that stable Σ-freeness over Σ-restricted formulae is modular with respect to the union of theories, we can show a weaker result in terms of *totally* Σ-restricted formulae.

Proposition 49 *Let T_1, T_2 be two consistent theories of respective signature Σ_1, Σ_2, such that T_i is stably Σ-free over $TRes(\mathcal{L}^{\Sigma_i}, \Sigma)$ for $i = 1, 2$. If T_1 and T_2 have the same atomic Σ-theory H_0, then $T_1 \cup T_2$ is stably Σ-free over $TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$* [33] .

PROOF. Let $\varphi_1 \wedge \varphi_2 \wedge res^{\Sigma}(\tilde{u})$ be an element of $TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$ satisfiable in $T_1 \cup T_2$, where $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ and $\mathcal{V}ar(\varphi_1 \wedge \varphi_2) \subseteq \tilde{u}$. We show that the formula is satisfiable in a model of $T_1 \cup T_2$ whose Σ-reduct is free in the atomic Σ-theory of $T_1 \cup T_2$ over a countably infinite base.

Clearly, the sentence $\psi_i := \varphi_i \wedge res^{\Sigma}(\tilde{u})$ is satisfiable in T_i for $i = 1, 2$. In particular, since $\psi_i \in TRes(\mathcal{L}^{\Sigma_i}, \Sigma)$ and T_i is stably Σ-free over $TRes(\mathcal{L}^{\Sigma_i}, \Sigma)$ by assumption, ψ_i is satisfiable in a model $\mathcal{A}_i$ of T_i such that $\mathcal{A}_i{}^{\Sigma}$ is free in H_0 over a countably-infinite basis. By Proposition 43, $\mathcal{A}_1$ and $\mathcal{A}_2$ are Σ-fusible.

Since the shared variables of φ_1 and φ_2 are included in the restriction $res^{\Sigma}(\tilde{u}) = iso^{\Sigma}(\tilde{u}) \wedge dif(\tilde{u})$, we can already conclude by Lemma 13 that $\varphi_1 \wedge \varphi_2$ is satisfiable in a fusion $\mathcal{F}$ of $\mathcal{A}_1$ and $\mathcal{A}_2$. By an argument similar to the observation after Lemma 13, we can actually show that the whole $\varphi_1 \wedge \varphi_2 \wedge res^{\Sigma}(\tilde{u})$ is satisfiable in $\mathcal{F}$.

[33] By a small abuse of notation, we treat here each pair in $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$ as the conjunction of its components.

We have already seen that $\mathcal{F} \in \mathit{Mod}(T_1 \cup T_2)$ and $\mathcal{F}^{\Sigma}$ is free in H_0 over a countably-infinite basis. To complete the proof then, it is enough to recall that, by Lemma 47, the atomic Σ-theory of $T_1 \cup T_2$ coincides with H_0. $\square$

The above result is not sufficient for our needs given that, in general, the class $\mathit{TRes}(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$ is strictly included in $\mathit{Res}(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$. One might argue, however, that if we limit ourselves to totally Σ-restricted formulae, we do get the kind of modularity and completeness results we desire. As a matter of fact, we can show that our combination procedure is sound and complete for all partially Σ-restricted formulae of the form $\varphi_1 \wedge \varphi_2 \wedge \mathit{res}^{\Sigma}(\tilde{u})$ in which $\tilde{u}$ includes the variables shared by φ_1 and φ_2. Unfortunately, even this is not enough.

In fact, our ultimate goal is to work with formulae in $\mathcal{L}^{\Sigma_1 \cup \Sigma_2}$, whether they have an attached Σ-restriction or not. As we saw, these formulae can be dealt with by our combination method provided that $\mathcal{L}$ is effectively purifiable w.r.t. $\langle \Sigma_1, \Sigma_2 \rangle$. What we do then is first convert 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, these disjuncts may have a different (typically larger) set of free variables. Therefore, even if we start with the totally Σ-restricted formula $\varphi(\tilde{v}) \wedge \mathit{res}^{\Sigma}(\tilde{v})$, after purification we may end up with partially Σ-restricted formulae of the form $\varphi_1 \wedge \varphi_2 \wedge \mathit{res}^{\Sigma}(\tilde{v})$ where not all the shared variables of φ_1 and φ_2 are included in $\tilde{v}$.

When $\mathcal{L}$ coincides with *Qff*, it is possible to generate the disjuncts $\varphi_1 \wedge \varphi_2$ so that

- $U := \mathcal{V}ar(\varphi_1 \wedge \varphi_2) \setminus \tilde{v}$ consists only of shared variables and
- $\varphi_1 \wedge \varphi_2 \models u_i \equiv t_i$ for all $u_i \in U$, where t_i is a pure term.

This entails that we can extend the Σ-restriction of φ to $\mathcal{V}ar(\varphi_1 \wedge \varphi_2)$ without loss of solutions only if we are guaranteed that the terms t_i above denote only Σ-isolated individuals.

We show in Section 8 that a situation like this can in fact be achieved for certain pairs of component theories. A crucial feature of some of these theories will be that their shared symbols are *constructors* in the sense formally defined in the next section.

7 Theories with Constructors

There are several definitions of constructors in Computer Science, but they are all based on the same fundamental idea. In essence, a set of constructors is a set of constants and functions that can be used to construct a computable data type. For instance, zero and the successor function are the constructors of the positive integer data type, the empty stack and the push function are the constructors of the stack data type, and so on.

In symbolic computation, constructors are the *symbols* that denote constructor functions. As such, they can be given syntactical definitions such as the one used in term rewriting (see later). The algebraic approaches to abstract data types, however, provide insights for formally understanding constructor symbols at a semantic level. In the algebraic ADT literature (see, e.g., (Ehrig and Mahr, 1985, 1990)), abstract data types are typically defined by initial algebras.[34] In that context, the constructors of an initial algebra $\mathcal{A}$ of signature Ω are those function symbols of Ω that can be used to incrementally generate the universe of $\mathcal{A}$ out of an initially empty set. Non-constructors then are function symbols that, while also denoting maps from A to A, are not necessary to *build* A. More formally, we could say that a signature $\Sigma \subseteq \Omega$ is a set of constructors for $\mathcal{A}$ if the empty set, which is a set of (Ω-)generators for $\mathcal{A}$, is also a set of (Σ-)generators for $\mathcal{A}^{\Sigma}$.

We could think of extending this notion to non-initial free algebras by saying that a signature $\Sigma \subseteq \Omega$ is a set of constructors for a free algebra $\mathcal{A}$ with signature Ω and basis X, if X, which is a set of generators for $\mathcal{A}$, is also a set of generators for $\mathcal{A}^{\Sigma}$. As it turns out, this straightforward generalization is more restrictive than it needs be. To see that, consider the equational theory E of signature $\Omega := \{\mathsf{0}, \mathsf{s}, +\}$ axiomatized by the sentences:

$$\boxed{\begin{array}{ll} \forall x, y, z. & x + (y + z) \equiv (x + y) + z \\ \forall x, y. & x + y \equiv y + x \\ \forall x, y. & x + \mathsf{s}(y) \equiv \mathsf{s}(x + y) \\ \forall x. & x + \mathsf{0} \equiv x \end{array}}$$

The algebra of the natural numbers with addition is an initial model of this theory (where s denotes the successor function). Now, the reduct of this algebra to the signature $\Sigma := \{\mathsf{0}, \mathsf{s}\}$ is also initial, which means that Σ is a set of constructors for the algebra. We would like to say then that Σ is also a set of constructors for all the free models of E, but this is not the case. In fact, if $\mathcal{A}$ is an Ω-algebra free in E over a *nonempty* set X, the individual $x +^{\mathcal{A}} x$ of $\mathcal{A}$,

[34] Recall that an initial algebra is a free algebra with an empty basis.

for any $x \in X$, cannot be generated from X by $0^{\mathcal{A}}$ and $s^{\mathcal{A}}$ alone. Therefore, X is not a set of generators for $\mathcal{A}^{\Sigma}$. The interesting point of this example is that $\mathcal{A}^{\Sigma}$ is in fact a free algebra. And while not free over X, it is free over an easily definable superset of X which includes all the individuals that, like $x +^{\mathcal{A}} x$, are not generated from X by $0^{\mathcal{A}}$ and $s^{\mathcal{A}}$ alone. Moreover, $\mathcal{A}^{\Sigma}$ is free precisely in the Σ-variety $Mod(E^{\Sigma})$.

We have developed our notion of constructors around the observation above and have found it very useful in the combination results described later in the paper. The key facts about constructors used for those results are that free structures with a set Σ of constructors are Σ-generated by their Σ-isolated individuals and are Σ-fusible.

This idea of constructors was introduced in (Tinelli and Ringeissen, 1998) after a similar one in (Domenjoud et al., 1994), and further refined with Franz Baader in (Baader and Tinelli, 1998) in the context of equational theories. In the following, we provide a unified treatment of the results in (Tinelli and Ringeissen, 1998) and (Baader and Tinelli, 1998) for the case of arbitrary first-order theories. Our definition of constructors is rather general. As initially shown in (Baader and Tinelli, 1999), it includes the constructors in (Domenjoud et al., 1994). We will show in the appendix that, under quite reasonable assumptions, it also includes the constructors used in term rewriting. [35]

7.1 Constructors

For the rest of the section let us fix a signature Ω and a subsignature Σ of Ω. Given a subset G of $T(\Omega, V)$, we denote by $T(\Sigma, G)$ the set of terms over the "variables" G. More precisely, every member of $T(\Sigma, G)$ is obtained from a term $s \in T(\Sigma, V)$ by replacing the variables of s with terms from G. To express this construction we will denote any such term by $s(\tilde{r})$ $s(\tilde{r})$ where $\tilde{r}$ is a tuple with no repetitions collecting the terms of G that replace the variables of s. Note that this notation is consistent with the fact that $G \subseteq T(\Sigma, G)$. In fact, every $r \in G$ can be represented as $s(r)$ where s is a variable of V. Also notice that $T(\Sigma, V) \subseteq T(\Sigma, G)$ whenever $V \subseteq G$. In this case, every $s \in T(\Sigma, V)$ can be trivially represented as $s(\tilde{v})$ where $\tilde{v}$ are the variables of s.

For every theory T with signature Ω and every subset Σ of Ω, we define the following subset of $T(\Omega, V)$:

[35] After the submission of this paper, the definition has been generalized even further. See (Baader and Tinelli, 2001) for more details.

$$G_T(\Sigma, V) := \{r \in T(\Omega, V) \mid r \neq_T t \text{ for all } t \in T(\Omega, V) \text{ with } t(\epsilon) \in \Sigma\}.$$

In essence, $G_T(\Sigma, V)$ is made, modulo equivalence in T, of Ω-terms whose top symbol is not in Σ.

We start with a syntactical definition of our notion of constructors for a theory. We then show that for theories admitting free models with an infinite basis, this definition has a simple model-theoretic characterization. We will use both the syntactical definition and the semantical characterization of constructors in the next sections, as convenient.

Definition 50 (Constructors) *Let T be a non-trivial theory of signature Ω, $\Sigma \subseteq \Omega$, and $G := G_T(\Sigma, V)$. The signature Σ is a* set of constructors for T *iff the following holds:*

(1) $V \subseteq G$.
(2) For all $t \in T(\Omega, V)$, there is an $s(\tilde{r}) \in T(\Sigma, G)$ such that

$$t =_T s(\tilde{r}).$$

(3) For all n-ary $P \in \Sigma^{\mathrm{P}} \cup \{\equiv\}$ and $s_1(\tilde{r}_1), \ldots, s_n(\tilde{r}_n) \in T(\Sigma, G)$,

$$T \models \tilde{\forall}\, P(s_1(\tilde{r}_1), \ldots, s_n(\tilde{r}_n)) \quad \textit{iff} \quad T \models \tilde{\forall}\, P(s_1(\tilde{v}_1), \ldots, s_n(\tilde{v}_n))$$

where $\tilde{v}_1, \ldots, \tilde{v}_n$ are fresh variables abstracting $\tilde{r}_1, \ldots, \tilde{r}_n$ so that two terms are abstracted by the same variable iff they are equivalent in T.

Notice that when Σ has no predicate symbols, condition *(3)* reduces to:

(3) For all $s_1(\tilde{r}_1), s_2(\tilde{r}_2) \in T(\Sigma, G)$,

$$s_1(\tilde{r}_1) =_T s_2(\tilde{r}_2) \quad \textit{iff} \quad s_1(\tilde{v}_1) =_T s_2(\tilde{v}_2)$$

where $\tilde{v}_1, \tilde{v}_2$ are fresh variables abstracting $\tilde{r}_1, \tilde{r}_2$ so that two terms are abstracted by the same variable iff they are equivalent in T.

It is easy to see that any set of constant symbols of Ω is a set of constructors for any Ω-theory T. It is also easy to show that the whole Ω is a set of constructors for T if and only if T is collapse-free.

The following is an another immediate consequence of the definition of constructors.

Proposition 51 *For every theory T and signature Σ, Σ is a set of constructors for T iff Σ is a set of constructors for* $\mathrm{At}(T)$.

We show below that when Σ is a set of constructors for an Ω-theory T admitting a free model $\mathcal{A}$ with an infinite basis [36], the Σ-reduct of $\mathcal{A}$ is free in T^{Σ},

[36] A large class of theories admitting free models with infinite bases is the class of

the Σ-theory of T, with over a set determined by $G_T(\Sigma, V)$. For this purpose, we will use the following properties of $G_T(\Sigma, V)$.

Lemma 52 *For all non-trivial theories T of signature Ω,*

(1) $G_T(\Sigma, V)$ is closed under equivalence in T;
(2) $G_T(\Sigma, V)$ is nonempty iff $V \subseteq G_T(\Sigma, V)$;
(3) If $V \subseteq G_T(\Sigma, V)$, then T^Σ is collapse-free.

PROOF. Let $G := G_T(\Sigma, V)$. We prove only points 2 and 3, as 1 is trivial.

(2) Since V is assumed to be countably infinite, $V \subseteq G$ obviously implies that G is nonempty. We prove the other direction by proving its contrapositive. Assume that there exists a variable $v \in V \setminus G$. By definition of G then, there exists an $f \in \Sigma$ and a tuple $\tilde{t}$ of Ω-terms such that $v =_T f(\tilde{t})$. Now consider any $r \in T(\Omega, V)$. By applying the substitution $\{v \leftarrow r\}$ to the equation $v \equiv f(\tilde{t})$, we obtain a tuple of Ω-terms $\tilde{t'}$ such that $r =_T f(\tilde{t'})$, which means that $r \notin G$. From the generality of r it follows that G is empty.

(3) Again, we prove the contrapositive. Assume that T^Σ is not collapse-free. Since T is non-trivial by assumption, there must exist a *non-variable* Σ-term s and a variable $v \in V$ such that $v =_{T^\Sigma} s$. By definition of G this implies that $v \notin G$, and thus $V \not\subseteq G$. $\square$

Proposition 53 *Let T a Ω-theory admitting a free model $\mathcal{A}$ with a countably infinite basis X and let α be a bijective valuation of V onto X.*[37] *If Σ is a set of constructors for T then $\mathcal{A}^\Sigma$ is free in T^Σ over the superset Y of X defined as follows:*

$$Y := \{[\![r]\!]^{\mathcal{A}}_\alpha \mid r \in G_T(\Sigma, V)\}.$$

PROOF. Let $G := G_T(\Sigma, V)$ and assume that Σ is a set of constructors for T. First notice that $X \subseteq Y$ because $V \subseteq G$. Then observe that since $\mathcal{A}$ is a model of T, its reduct $\mathcal{A}^\Sigma$ is a model of T^Σ. We show that $\mathcal{A}^\Sigma$ is Σ-generated by Y. In fact, let a be an element of A—which is also the carrier of $\mathcal{A}^\Sigma$. We know that as an Ω-structure $\mathcal{A}$ is generated by X; thus there exists a term $t \in T(\Omega, V)$ such that $a = [\![t]\!]^{\mathcal{A}}_\alpha$. By Definition 50(2), the term $t \in T(\Omega, V)$ is equivalent in T to a term $s(\tilde{r}) \in T(\Sigma, G)$. Since $\mathcal{A}$ is a model of T, this implies that $a = [\![t]\!]^{\mathcal{A}}_\alpha = [\![s(\tilde{r})]\!]^{\mathcal{A}}_\alpha$, from which it easily follows by definition of Y that a is Σ-generated by Y.

non-trivial universal Horn theories.

[37] Such a valuation α exists since V is assumed to be countably infinite.

The above entails that $\mathcal{A}^\Sigma$ satisfies the first condition of Proposition 38. To show that it is free in T^Σ then it is enough to show that it also satisfies the second condition of the same proposition.

Thus, consider any terms $s_1(\tilde{v}_1), \dots, s_n(\tilde{v}_n) \in T(\Sigma, V)$, relation symbol $P \in \Sigma^{\mathrm{P}} \cup \{\equiv\}$, and injection β of $V_0 := \mathcal{V}ar(P(s_1(\tilde{v}_1), \dots, s_n(\tilde{v}_n)))$ into Y such that

$$(\mathcal{A}^\Sigma, \beta) \models P(s_1(\tilde{v}_1), \dots, s_n(\tilde{v}_n)).$$

By definition of Y we know that for all $v \in V_0$, there is a term $r_v \in G$ such that $\beta(v) = [\![r_v]\!]^{\mathcal{A}}_\alpha$. Using these terms we can construct two tuples $\tilde{r}_1$ and $\tilde{r}_2$ of terms in G such that, for $i = 1, 2$, the term $s_i(\tilde{r}_i)$ is obtained from $s_i(\tilde{v}_i)$ by replacing each variable v in $\tilde{v}_i$ by the term r_v, and $(\mathcal{A}, \alpha) \models P(s_1(\tilde{r}_1), \dots, s_2(\tilde{r}_n))$. Since $\mathcal{A}$ is free in T over X and α is injective as well we can conclude by Proposition 38(2) that $T \models \tilde{\forall}\, P(s_1(\tilde{r}_1), \dots, s_2(\tilde{r}_n))$.

Now, by the injectivity of β we know that $r_u \neq_T r_v$ for distinct variables $u, v \in V_0$. Therefore, considered the other way round, the atom $P(s_1(\tilde{v}_1), \dots, s_n(\tilde{v}_n))$ can be obtained from $P(s_1(\tilde{r}_1), \dots, s_2(\tilde{r}_n))$ by abstracting the terms $\tilde{r}_1, \dots, \tilde{r}_n$ so that two terms are abstracted by the same variable iff they are equivalent in T. But then, by Point 3 of Definition 50 we can conclude that $T \models \tilde{\forall}\, P(s_1(\tilde{v}_1), \dots, s_n(\tilde{v}_n))$. Since $\tilde{\forall}\, P(s_1(\tilde{v}_1), \dots, s_n(\tilde{v}_n))$ is a Σ-sentence, this is the same as saying that $T^\Sigma \models \tilde{\forall}\, P(s_1(\tilde{v}_1), \dots, s_n(\tilde{v}_n))$. □

The freeness of the structure $\mathcal{A}^\Sigma$ above is therefore necessary for Σ to be a set of constructors for T. It becomes also sufficient when T^Σ is collapse-free, as the following theorem shows.

Theorem 54 *Let T a Ω-theory admitting a free model $\mathcal{A}$ over a countably infinite set. Then, Σ is a set of constructors for T iff*

- *the Σ-reduct of $\mathcal{A}$ is free in T^Σ and*
- *T^Σ is collapse-free.*

PROOF. As before, let X be a countably infinite basis of $\mathcal{A}$, α a bijective valuation of V onto X, $G := G_T(\Sigma, V)$, and $Y := \{[\![r]\!]^{\mathcal{A}}_\alpha \mid r \in G\}$.

($\Rightarrow$) By Proposition 53, $\mathcal{A}^\Sigma$ is free in T^Σ. By Lemma 52(3), the fact that $V \subseteq G$ (cf. Condition (1) of Definition 50) implies that T^Σ is collapse-free.

($\Leftarrow$) Assume that T^Σ is collapse-free and $\mathcal{A}^\Sigma$ is free in T^Σ over some set Z. First, notice that Z cannot be the empty set. Otherwise, $\mathcal{A}$ would also be generated by the empty set, making X a redundant set of generators, which is impossible because $\mathcal{A}$ is free over X by assumption.

We prove $Y = Z$ by first proving that $Y \subseteq Z$ and then that $Z \subseteq Y$. Ad absurdum, assume that $Y \not\subseteq Z$ and let $y \in Y \setminus Z$. Since $\mathcal{A}$ is Ω-generated by X and Σ-generated by Z, we know that there exist a *non-variable* Σ-term s and a tuple $\tilde{t}$ of Ω-terms such that $[\![t_i]\!]^{\mathcal{A}}_{\alpha} \in Z$ for all elements t_i of $\tilde{t}$, and $y = [\![s(\tilde{t})]\!]^{\mathcal{A}}_{\alpha}$. By definition of Y we know that there is a term $r \in G$ such that $y = [\![r]\!]^{\mathcal{A}}_{\alpha}$. As $\mathcal{A}$ is free in T and α is injective, we can then conclude by Proposition 38(2) that $r =_T s(\tilde{t})$, but then r cannot be in G. It follows that $Y \subseteq Z$.

To show that $Z \subseteq Y$, let $z \in Z$. Since $\mathcal{A}$ is Ω-generated by X, there exists an Ω-term r such that $z = [\![r]\!]^{\mathcal{A}}_{\alpha}$. We prove by contradiction that r is an element of G, which will then entail by construction of Y that $z \in Y$. Therefore, assume that $r \notin G$. Then, there must be a function symbol $f \in \Sigma$ and a tuple of Ω-terms $\tilde{t}$ such that $r =_T f(\tilde{t})$. Since the elements of $\tilde{t}$ are all Σ-generated by Z, there is a variable v, a *non-variable* Σ-term s, and an injective mapping β of $\mathcal{V}ar(s) \cup \{v\}$ into Z such that $\beta(v) = z = [\![s]\!]^{\mathcal{A}^{\Sigma}}_{\beta}$.[38] As $\mathcal{A}^{\Sigma}$ is free in T^{Σ} over Z, we obtain that $v =_{T^{\Sigma}} s$. But this contradicts the fact that T^{Σ} is collapse-free. It follows that $r \in G$ and so $z \in Y$.

In conclusion, we have shown that Z is nonempty and coincides with $Y = \{[\![r]\!]^{\mathcal{A}}_{\alpha} \mid r \in G\}$. In particular, this means that G is nonempty either. The first condition in Definition 50 follows then directly from Lemma 52(2). The second condition follows by Proposition 38(2) and Corollary 39, given that $\mathcal{A}$ is free in T and Σ-generated by $Y = Z$. Similarly, the third condition follows from Proposition 38(2). $\square$

We can now give an alternative formulation of Theorem 54 by means of the following corollary.

Corollary 55 *Let T a Ω-theory admitting a free model $\mathcal{A}$ over a countably infinite set. Then, the following are equivalent.*

(1) Σ is a set of constructors for T.
(2) $\mathcal{A}^{\Sigma}$ is free in T^{Σ} over $Is(\mathcal{A}^{\Sigma})$.[39]

PROOF. ($1 \Rightarrow 2$) By Theorem 54, $\mathcal{A}^{\Sigma}$ is free in the collapse-free theory T^{Σ}. By Proposition 40, the unique basis of $\mathcal{A}^{\Sigma}$ coincides with $Is(\mathcal{A}^{\Sigma})$.

($2 \Rightarrow 1$) Let $\mathcal{A}^{\Sigma}$ be free in T^{Σ} over $Is(\mathcal{A}^{\Sigma})$. By Theorem 54, it is enough to show that T^{Σ} is collapse-free. Assume the contrary. Then, since T^{Σ} is non-trivial for admitting the infinite model $\mathcal{A}^{\Sigma}$, there must be a variable v and a non-variable

[38] Note that v may be an element of $\mathcal{V}ar(s)$.
[39] Recall the $Is(\mathcal{A}^{\Sigma})$ is the set of all the isolated individuals of $\mathcal{A}^{\Sigma}$ (cf. Definition 3).

Σ-term s such that $v =_{T^\Sigma} s$. From the fact then that variables are equivalent in T^Σ, and so in $\mathcal{A}^\Sigma$, to a term starting with a Σ-symbol, it easily follows that no individual of $\mathcal{A}^\Sigma$ is Σ-isolated. Therefore, $\mathit{Is}(\mathcal{A}^\Sigma)$ is empty. But then, we can argue as in the proof of Theorem 54 that $\mathcal{A}$ is generated by the empty set, which is impossible as $\mathcal{A}$ is free over an infinite set by assumption. $\square$

Later in the paper we will consider theories T for which $G_T(\Sigma, V)$ is closed under instantiation into itself, by which we mean that replacing the variables of a term in $G_T(\Sigma, V)$ by terms in $G_T(\Sigma, V)$ yields a term also in $G_T(\Sigma, V)$.

Definition 56 *Let T be a of signature Ω and $\Sigma \subseteq \Omega$. We say that $G_T(\Sigma, V)$ is* closed under instantiation into itself *iff $r\sigma \in G_T(\Sigma, V)$ for all terms $r \in G_T(\Sigma, V)$ and substitutions $\sigma \in \mathrm{SUB}(V)$ such that $\mathcal{R}an(\sigma) \subseteq G_T(\Sigma, V)$.*

When $G_T(\Sigma, V)$ is closed under instantiation into itself, the set $\mathit{Is}(\mathcal{A}^\Sigma)$ exhibits in turn the following closure property.

Lemma 57 *Let T a Ω-theory admitting a free model $\mathcal{A}$ over a countably infinite set X and assume that Σ is a set of constructors for T. If $G_T(\Sigma, V)$ is closed under instantiations into itself, then*

$$[\![r]\!]^{\mathcal{A}}_{\beta} \in \mathit{Is}(\mathcal{A}^\Sigma)$$

for all terms $r \in G_T(\Sigma, V)$ and valuations β of $\mathcal{V}ar(r)$ into $\mathit{Is}(\mathcal{A}^\Sigma)$.

PROOF. Let $r(\tilde{v}) \in G := G_T(\Sigma, V)$ and β a valuation of $\tilde{v}$ into $\mathit{Is}(\mathcal{A}^\Sigma)$. We have seen that $X \subseteq \mathit{Is}(\mathcal{A}^\Sigma) = \{[\![r]\!]^{\mathcal{A}}_{\alpha} \mid r \in G\}$ for any bijective valuation α of V onto X. This means that for each $v \in \tilde{v}$ there is a term $r_v \in G$ such that $\beta(v) = [\![r_v]\!]^{\mathcal{A}}_{\alpha}$. It follows that there is a substitution σ into G such that $[\![r]\!]^{\mathcal{A}}_{\beta} = [\![r\sigma]\!]^{\mathcal{A}}_{\alpha}$. The claim then follows immediately from the assumption that G is closed under instantiation into itself. $\square$

7.2 Normal Forms

Condition 2 of Definition 50 says that when Σ is a set of constructors for an Ω-theory T, every term $t \in T(\Omega, V)$ is equivalent in T to a term $s(\tilde{r}) \in T(\Sigma, G)$, where $G := G_T(\Sigma, V)$. We call $s(\tilde{r})$ a *normal form of t in T*.[40] We say that a term t is *in normal form* if it is a member of $T(\Sigma, G)$. Because $V \subseteq G$, it is immediate that Σ-terms are in normal form, as are terms in G.

[40] Notice that in general, a term may have more than one normal form.

We point out that, according to our definition, it is not necessarily the case that all the variables occurring in the normal form of a term also occur in the term itself. However, it is possible to make this assumption with no loss of generality whenever Σ contains a constant symbol. A proof of this fact can be found in (Tinelli and Ringeissen, 2001), an extended version of this paper. A similar result is also shown in (Fiorentini and Ghilardi, 2001).

We will be interested in normal forms that are computable in the following sense.

Definition 58 (Computable Normal Forms) *Let Σ be a set of constructors for a theory T of signature Ω and consider a map*

$$\mathrm{NF}_T^\Sigma\colon T(\Omega, V) \to T(\Sigma, G_T(\Sigma, V)).$$

We say that normal forms are computable for Σ and T by NF_T^Σ *iff* NF_T^Σ *is computable and* $\mathrm{NF}_T^\Sigma(t)$ *is a normal form of* t*, i.e.,* $\mathrm{NF}_T^\Sigma(t) =_T t$*.*

We will simply say that *normal forms are computable for* Σ *and* T if there is a function NF_T^Σ such that normal forms are computable for Σ and T by NF_T^Σ.

Although we will not needed it here, we point out an important consequence of Definition 58: if normal forms are computable for Σ and T, it is always possible to tell whether a term is in normal form or not. Again, a proof of this can be found in (Tinelli and Ringeissen, 2001).

7.3 Examples

We provide below some examples of theories admitting constructors for situations other than the trivial ones already mentioned. But first, let us consider some counter-examples.

- The signature $\Sigma := \{f\}$ is not a set of constructors for the theory $T := \{\forall x.\ x \equiv f(g(x))\}$ because it does not satisfy Definition 50(1), as one can easily show.
- The signature $\Sigma := \{f\}$ is not a set of constructors for the theory $T := \{\forall x.\ g(x) \equiv f(g(x))\}$ because it does not satisfy Definition 50(2). In fact, the term $g(x)$ does not have a normal form.
- The subsignature $\Sigma := \{f\}$ of $\Omega := \{f, g\}$ is not a set of constructors for the theory $T := \{\forall x.\ f(g(x)) \equiv f(f(g(x)))\}$. It is easy to show that $G_T(\Sigma, V) = V \cup \{g(t) \mid t \in T(\Omega, V)\}$ and that conditions (1) and (2) of Definition 50 hold. However, condition (3) does not hold since $f(g(x)) =_T f(f(g(x)))$ even if $f(y) \neq_T f(f(y))$.

- By a similar argument, one can show that the subsignature $\Sigma := \{P\}$ of $\Omega := \{P, g\}$ is not a set of constructors for the theory $T := \{\forall x.\ P(g(x))\}$.

The theory of the natural numbers with addition considered earlier is indeed an example of a theory with constructors.

Example 59 *Consider the signature $\Sigma_{59} := \{\mathsf{0}, \mathsf{s}, +\}$ and the theory E_{59} axiomatized by the sentences:*

$$\boxed{\begin{array}{ll} \forall x, y, z. & x + (y + z) \equiv (x + y) + z \\ \forall x, y. & x + y \equiv y + x \\ \forall x, y. & x + \mathsf{s}(y) \equiv \mathsf{s}(x + y) \\ \forall x. & x + \mathsf{0} \equiv x \end{array}}$$

The signature $\Sigma := \{\mathsf{0}, \mathsf{s}\}$ is a set of constructors for E_{59} in the sense of Definition 50 (see (Baader and Tinelli, 1998) for a proof). In particular, $G_T(\Sigma, V)$ is the set of all terms that either a variable or a (possibly nested) addition of variables. Furthermore, every normal form looks like $\mathsf{s}^n(r)$ where $n \geq 0$ and r is either $\mathsf{0}$ or a term in $G_T(\Sigma, V)$. It is interesting to notice that $G_T(\Sigma, V)$ is closed under instantiation into itself.

The following is another simple, but this time non-equational, example of a theory with constructors.

Example 60 *Consider the signature $\Sigma_{60} := \{\mathsf{0}, \mathsf{s}, +, \mathsf{Even}\}$ and the theory T_{60} axiomatized by E_{59} above plus the sentences:*

$$\boxed{\begin{array}{l} \quad\ \mathsf{Even}(\mathsf{0}) \\ \forall x.\ \mathsf{Even}(x) \Rightarrow \mathsf{Even}(\mathsf{s}(\mathsf{s}(x))) \end{array}}$$

It is not difficult to show that the signature $\Sigma := \{\mathsf{0}, \mathsf{s}, \mathsf{Even}\}$ is a set of constructors for T_{60}. Interestingly, Σ is not a set of constructors if we also add the axiom $\forall x.\ \mathsf{Even}(x + x)$. The reason is that then, since $x + x$ is in $G_T(\Sigma, V)$, the sentence $\forall y.\ \mathsf{Even}(y)$ should also be entailed by the theory according to Definition 50(3), but it is not.

The next examples differ from the previous ones in that their equational Σ-theory is no longer empty.

Example 61 *Consider the signature $\Sigma_{61} := \{\mathsf{0}, \mathsf{1}, \mathsf{rev}, \cdot\}$ and the theory E_{61}*

axiomatized by the sentences:

$$\boxed{\begin{array}{ll} \forall x, y, z. & x \cdot (y \cdot z) \equiv (x \cdot y) \cdot z \\ \forall x, y. & \mathsf{rev}(x \cdot y) \equiv \mathsf{rev}(y) \cdot \mathsf{rev}(x) \\ \forall x. & \mathsf{rev}(\mathsf{rev}(x)) \equiv x \\ & \mathsf{rev}(0) \equiv 0 \\ & \mathsf{rev}(1) \equiv 1 \end{array}}$$

We show in the appendix that the signature $\Sigma := \{0, 1, \cdot\}$ *is a set of constructors for* E_{61}*. The set* $G_T(\Sigma, V)$ *is the equivalence closure in* E_{61} *of the set* $V \cup \{\mathsf{rev}(v) \mid v \in V\}$*. Moreover, every normal form is a concatenation (with* $\cdot$*) of terms in* $\{0, 1\} \cup G_T(\Sigma, V)$*. In this case too* $G_T(\Sigma, V)$ *is closed under instantiation into itself.*

Example 62 *Consider the signature* $\Sigma_{62} := \{0, 1, \mathsf{rev}, \cdot, \mathsf{Prefix}\}$ *and the theory* T_{62} *axiomatized by* E_{61} *plus the sentences:*

$$\boxed{\begin{array}{ll} \forall x. & \mathsf{Prefix}(x, x) \\ \forall x, y. & \mathsf{Prefix}(x, x \cdot y) \end{array}}$$

Again, it is not difficult to see that the signature $\Sigma := \{0, 1, \cdot, \mathsf{Prefix}\}$ *is a set of constructors for* T_{62}*.*

8 A Class of N-O-combinable Theories

Our main goal in the last section was to identify sufficient conditions for N-O-combinability, which lead us to the idea of stable Σ-freeness. In (Tinelli and Ringeissen, 2001) we describe some simple cases of stably Σ-free theories with N-O-combinable members. For instance, we show that theories sharing at most finitely-many constant symbols and entailing that these symbols are distinct are N-O-combinable over the quantifier-free formulae. We also show that universal theories sharing all of their function symbols are N-O-combinable over the universal formulae, provided that each theory's restriction to the function symbols coincides with the theory of finite trees (see (Tinelli and Ringeissen, 2001) for more details and further examples).

For space constraints we discuss here only one, major, case of stably Σ-free theories with N-O-combinable members: the class of complete theories with constructors. As usual, let us fix two countable signatures Σ_1, Σ_2 with finite intersection Σ and two theories T_1, T_2 of respective signature Σ_1, Σ_2.

We will assume that for $i = 1, 2$,

- T_i is the (complete) theory of some free Σ_i-structure $\mathcal{A}_i$ with a countably infinite basis;
- Σ is a finite set of constructors for T_i.

We also assume that the two theories agree on their constructors in the sense that $\mathrm{At}(\mathcal{A}_1^{\Sigma}) = \mathrm{At}(\mathcal{A}_2^{\Sigma})$.

Our goal is to show that T_1 and T_2 are N-O-combinable over some effectively purifiable language $\mathcal{L}$ by using the fact that each T_i is stably Σ-free over any $\mathcal{L}^{\Sigma_i}$. Recall that if we can show this, then we know we can use our combination procedure in a sound and complete way to (semi)-decide the satisfiability in $T_1 \cup T_2$ of formulae in $Res(\mathcal{L}^{\Sigma_1 \cup \Sigma_2}, \Sigma)$, once we have for $i = 1, 2$ a decision procedure for the satisfiability in T_i of formulae in $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$.

We can easily show that T_1 and T_2 are partially N-O-combinable over an arbitrary $\mathcal{L}$, which makes our procedure sound. Our current results are not strong enough to show that T_1 and T_2 are *totally* N-O-combinable over $\mathcal{L}$—which would make the combination procedure also complete. But they suffice to show that the procedure is complete for input formulae which are already totally Σ-restricted.

Although this is a strong restriction in general, it has a remarkable side-effect. As we prove in the following, with some additional assumptions on the computability of normal forms in T_1 and in T_2, we can turn our combination procedure into a decision procedure for the satisfiability in $T_1 \cup T_2$ of totally restricted quantifier-free formulae, even when T_1 and T_2 share infinitely-many terms.

We start by showing that the component theories are stably Σ-free over any class of formulae and (totally) N-O-combinable over totally Σ-restricted pairs of pure formulae.

Lemma 63 *For every class $\mathcal{L}$ of formulae, T_i is stably Σ-free over $\mathcal{L}^{\Sigma_i}$ for $i = 1, 2$.*

PROOF. Let $i \in \{1, 2\}$. Since T is the theory of $\mathcal{A}_i$, a Σ_i-formula is satisfiable in 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 $\mathrm{At}(T_i^{\Sigma})$ over a countably-infinite set. Now, since Σ is a set of constructors for T_i and $\mathcal{A}_i$ is obviously a free model of T_i, we know from Theorem 54 that $\mathcal{A}^{\Sigma}$ is free in T_i^{Σ} over some countably infinite set Y. From this and Proposition 41, it is easy to see that $\mathcal{A}^{\Sigma}$ is also free in $\mathrm{At}(T_i^{\Sigma})$ over Y. $\square$

Proposition 64 *For any class $\mathcal{L}$ of first-order formulae, T_1 and T_2 are totally N-O-combinable over $TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$.*

PROOF. Let H_0 be the atomic Σ-theory of T_1. By the construction of T_1 and T_2 and the assumption that $\mathrm{At}(\mathcal{A}_1^\Sigma) = \mathrm{At}(\mathcal{A}_2^\Sigma)$, it is immediate that H_0 is also the atomic Σ-theory of T_2. By Lemma 63, for $i = 1, 2$, T_i is stably Σ-free over any class of formulae, in particular over $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$. We can then conclude by Theorem 48(1), that T_1 and T_2 are partially N-O-combinable over $\mathcal{L}$.

From Lemma 63 again and Proposition 49, we also have that $T_1 \cup T_2$ is Σ-stable over $TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$. Since H_0 is collapse-free by Theorem 54, we can show exactly as in the proof of Theorem 48(2) that T_1 and T_2 are totally N-O-combinable over $TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$. $\square$

By virtue of the above result we can use our combination method to yield, trivially, a decision procedure for the satisfiability in $T := T_1 \cup T_2$ of formulae in $TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$ whenever the satisfiability in T_i of formulae in $TRes(\mathcal{L}^{\Sigma_i}, \Sigma)$ is decidable for $i = 1, 2$. In fact, we can modify the combination procedure so that, given a formula

$$\varphi_1 \wedge \varphi_2 \wedge iso^\Sigma(\tilde{v}) \wedge dif(\tilde{v}) \in TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma),$$

it considers it as the input pair $\langle \varphi_1, \varphi_2 \rangle$. However, since all the shared variables of φ_1 and φ_2 are Σ-restricted, the procedure this time chooses, deterministically, only the empty substitution in both the instantiation and the identification step. At this point, our decidability claim follows immediately.

Now, the decidability of the satisfiability of formulae in $TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$ is not terribly exciting because, as already observed, if one is interested in totally Σ-restricted formulae, he is more likely to be interested in the satisfiability of formulae in $TRes(\mathcal{L}^{\Sigma_1 \cup \Sigma_2}, \Sigma)$, not of those in $TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$.

We show below, however, that under some more assumption of T_1 and T_2, the result provided by Proposition 64 is enough for deciding the satisfiability in $T := T_1 \cup T_2$ of a specific instance of $TRes(\mathcal{L}^{\Sigma_1 \cup \Sigma_2}, \Sigma)$, namely $TRes(\mathit{Qff}^{\Sigma_1 \cup \Sigma_2}, \Sigma)$, the class of totally restricted *quantifier-free* formulae of signature $\Sigma_1 \cup \Sigma_2$. The reason is that the satisfiability in T of such formulae becomes effectively reducible to the satisfiability in T of formulae in $TRes(\mathit{Qff}^{\Sigma_1} \otimes \mathit{Qff}^{\Sigma_2}, \Sigma)$.

Here are the additional assumptions, which we will make from now on: for $i = 1, 2$,

- $G_{T_i}(\Sigma, V)$ is closed under instantiation into itself (cf. Definition 56);

- the word problem for T_i is decidable,
- normal forms are computable for Σ and T_i.

We start with some useful lemmas about $T := T_1 \cup T_2$.

Lemma 65 *Every model of T has Σ-isolated individuals.*

PROOF. Assume by contradiction that there is a model $\mathcal{B}$ of T with no Σ-isolated individuals. Then the Σ-sentence $\varphi := \neg\exists v.\ iso^{\Sigma}(v)$ is true in $\mathcal{B}$ and hence in $\mathcal{B}^{\Sigma_1}$, say. Since $\mathcal{B}^{\Sigma_1}$ is a model of T_1 and T_1 is the complete theory of $\mathcal{A}_1$, we can conclude that φ is true in $\mathcal{A}_1$ as well. But this is impossible because $\mathcal{A}_1$ has infinitely many Σ-isolated individuals by Proposition 53 and Corollary 55. $\square$

The following lemma states that in *every* model of T the terms of $G_{T_i}(\Sigma, V)$ $(i = 1, 2)$ map Σ-isolated individuals to Σ-isolated individuals.

Lemma 66 *For all $i = 1, 2$, $v \in V$, and $r(\tilde{v}) \in G_{T_i}(\Sigma, V)$,*

$$T \models v \equiv r(\tilde{v}) \wedge iso^{\Sigma}(\tilde{v}) \Rightarrow iso^{\Sigma}(v) \tag{1}$$

PROOF. Let $i \in \{1, 2\}$. As T includes T_i, the complete theory of $\mathcal{A}_i$, it is enough to show that the Σ_i-sentence in (1) above holds in $\mathcal{A}_i$.

Let β be any valuation of V such that $(\mathcal{A}_i, \beta) \models v \equiv r(\tilde{v}) \wedge iso^{\Sigma}(\tilde{v})$. To satisfy $iso^{\Sigma}(\tilde{v})$ in $\mathcal{A}_i$, β must map every variable in $\tilde{v}$ to an element of $Is({\mathcal{A}_i}^{\Sigma})$. Since $G_{T_i}(\Sigma, V)$ is closed under instantiation into itself, we obtain by Lemma 57 that $\beta(v) = [\![r]\!]^{\mathcal{A}_i}_{\beta} \in Is({\mathcal{A}_i}^{\Sigma})$, which means that $(\mathcal{A}_i, \beta) \models iso^{\Sigma}(v)$. The claim then follows from the generality of β. $\square$

As we have seen in Section 7, for $i = 1, 2$, Σ_i-terms have a normal form in T_i that is a Σ-term over the "variables" $G_{T_i}(\Sigma, V)$. Something analogous holds for $(\Sigma_1 \cup \Sigma_2)$-terms in T, where a set of "variables" can be built incrementally out of $G_{T_1}(\Sigma, V)$ and $G_{T_2}(\Sigma, V)$.

Definition 67 *The set $G^*_T(\Sigma, V)$ is inductively defined as follows:*

(1) Every variable is an element of $G^*_T(\Sigma, V)$, that is, $V \subseteq G^*_T(\Sigma, V)$.
(2) Assume that $r(\tilde{v}) \in G_{T_i}(\Sigma, V)$ for $i = 1$ or $i = 2$ and $\tilde{r}$ is a tuple of elements of $G^*_T(\Sigma, V)$ such that the following holds:
   (a) $r(\tilde{v}) \neq_T v$ for all variables $v \in V$;
   (b) $r_j(\epsilon) \notin \Sigma_i$ for all components r_j of $\tilde{r}$;

(c) *the tuples* $\tilde{v}$ *and* $\tilde{r}$ *have the same length;*
(d) $r_j \neq_T r_k$ *if* r_j, r_k *occur at different positions in the tuple* $\tilde{r}$.
Then $r(\tilde{r}) \in G^*_T(\Sigma, V)$.

Notice that for $i = 1, 2$ every non-collapsing element of G_i is in $G^*_T(\Sigma, V)$ for $i = 1, 2$ because the components of $\tilde{r}$ above can also be variables. Also notice that an element r of $G^*_T(\Sigma, V)$ cannot have a shared symbol (i.e., a symbol in Σ) as top symbol since r is a variable or it "starts" with an element of G_i.

In (Tinelli, 1999), it is shown that under the given assumptions on T_1 and T_2, Σ is also a set of constructors for T, normal forms are computable for Σ and T, and every normal form can be assumed to be in $T(\Sigma, G^*_T(\Sigma, V))$.[41] We will appeal to these facts in Proposition 69.

Lemma 68 *Let* φ *be a conjunction of* $(\Sigma_1 \cup \Sigma_2)$*-literals all of whose arguments are terms in* $T(\Sigma_1 \cup \Sigma_2, G^*_T(\Sigma, V))$. *Then,* φ *can be effectively converted into a finite set* S *which is equisatisfiable with* φ *in* T *and is partitioned into the sets*

$$L_1, \quad L_2, \quad F_1 := \{v^1_j \equiv r^1_j\}_{j \in J_1}, \quad F_2 := \{v^2_j \equiv r^2_j\}_{j \in J_2},$$

where

(1) L_1 *is made of literals of signature* Σ_1 *and* L_2 *is made of literals of signature* $\Sigma_2 \setminus \Sigma_1$;
(2) $\mathcal{V}ar(S) \setminus \mathcal{V}ar(\varphi) = \{v^i_j\}_{i,j}$;
(3) *for all* $i = 1, 2$ *and* $j \in J_i$,
  (a) v^i_j *does not occur in* L_i *and occurs only once in* F_i;
  (b) $r^i_j \in G_{T_i}(\Sigma, V) \setminus V$;
(4) *for all* $j \in J_1$, $v^1_j \in \mathcal{V}ar(L_2)$ *or* $v^1_j \in \mathcal{V}ar(r^2_k)$ *for some* $k \in J_2$;
  for all $j \in J_2$, $v^2_j \in \mathcal{V}ar(L_1)$ *or* $v^2_j \in \mathcal{V}ar(r^1_k)$ *for some* $k \in J_1$;

Furthermore, let $\tilde{v} := \mathcal{V}ar(\varphi)$, $\tilde{u} := \mathcal{V}ar(S)$, $\mathcal{A}$ *a model of* T *and* α *a valuation of* V *into* A. *If* $(\mathcal{A}, \alpha) \models S \cup iso^\Sigma(\tilde{v})$ *then* $(\mathcal{A}, \alpha) \models S \cup iso^\Sigma(\tilde{u})$.

PROOF. We simply apply to φ the purification procedure seen in Section 5 and collect in F_i $(i = 1, 2)$ the Σ_i-equations added by the purification process, in L_1 the purified literals of signature Σ_1, and in L_2 the remaining literals.

Then, Point 1 and point 2 are trivial. Point 3a is a consequence of the fact that each alien subterm is abstracted by a fresh variable. Point 3b follows from the definition of $G^*_T(\Sigma, V)$. Point 4 follows from the fact that each v^i_j is an abstraction variable.

[41] A proof of this for the equational case can also be found in (Baader and Tinelli, 2001).

Now let $\mathcal{A} \in Mod(T)$ and α a valuation such that $(\mathcal{A}, \alpha) \models S \cup iso^{\Sigma}(\tilde{v})$. Then define the binary relation $\succ$ on $F := F_1 \cup F_2$ as follows: for all $(v \equiv r), (v' \equiv r') \in F$,

$$(v \equiv r) \succ (v' \equiv r') \text{ iff } v' \in \mathcal{V}ar(r).$$

From the properties in the previous points and the fact that F consists only of equations added by purification it is not hard to show that $\succ$ is an acyclic relation. Then, by a simple well-founded induction argument based on $\succ$ one can show using Lemma 66 that $(\mathcal{A}, \alpha) \models iso^{\Sigma}(v^i_j)$ for all $i = 1, 2$ and $j \in J_i$. It follows by point 2 above and the definition of iso^{Σ} that $(\mathcal{A}, \alpha) \models S \cup iso^{\Sigma}(\tilde{u})$. $\square$

We are now ready to prove our reducibility claim.

Proposition 69 *The satisfiability in T of formulae in $TRes(Qff^{\Sigma_1 \cup \Sigma_2}, \Sigma)$ is effectively reducible to the satisfiability in T of formulae in the subclass $TRes(Qff^{\Sigma_1} \otimes Qff^{\Sigma_2}, \Sigma)$.*

PROOF. Let $\psi(\tilde{v}) := \varphi \wedge res^{\Sigma}(\tilde{v})$ be a formula of $TRes(Qff^{\Sigma_1 \cup \Sigma_2}, \Sigma)$ and assume for simplicity that $\tilde{v}$ is non-empty. This assumption is with no loss of generality because $\tilde{v}$ can be empty only when φ is a ground formula. But then, where v is an arbitrary variable, φ is trivially equisatisfiable in T by Lemma 65 with the totally Σ-restricted formula $\varphi \wedge res^{\Sigma}(v)$, which is effectively computable from φ.

Clearly, $\psi(\tilde{v})$ can be effectively converted into the logically equivalent formula

$$\psi_1 \wedge res^{\Sigma}(\tilde{v}) \vee \cdots \vee \psi_n \wedge res^{\Sigma}(\tilde{v})$$

where $\psi_1 \vee \cdots \vee \psi_n$ is φ's disjunctive normal form. Each ψ_i above is a conjunction of literals and $\psi(\tilde{v}) = \varphi(\tilde{v}) \wedge res^{\Sigma}(\tilde{v})$ is satisfiable in a model $\mathcal{A}$ of T if and only if for some $i \in \{1, \ldots, n\}$ the totally restricted formula $\psi_i \wedge res^{\Sigma}(\tilde{v})$ is satisfiable in $\mathcal{A}$. With no loss of generality then assume that φ is just a conjunction of literals and consider the following procedure with input $\varphi \wedge res^{\Sigma}(\tilde{v})$.

(1) Replace each argument t in each atom of φ by its computable normal form, which we know is an element of $T(\Sigma, G^*_T(\Sigma, V))$.
(2) Convert φ into the set $S := L_1 \cup L_2 \cup F_1 \cup F_2$ as in Lemma 68.
(3) For $i = 1, 2$, let φ_i be the conjunction of all the literals in $L_i \cup F_i$ and output the formula $\varphi_1 \wedge \varphi_2 \wedge res^{\Sigma}(\tilde{v})$.

From our assumptions and the procedure's construction it is clear that $\varphi_1 \wedge \varphi_2 \wedge res^{\Sigma}(\tilde{v})$ is computable from the initial formula $\varphi \wedge res^{\Sigma}(\tilde{v})$ and equisatisfiable with it in T. Now, in general, $\varphi_1 \wedge \varphi_2 \wedge res^{\Sigma}(\tilde{v})$ will be only partially Σ-restricted. In fact, step 1 above may introduce some new variables $\tilde{v}_1$ because the computed normal forms may have variables not occurring in the original terms, and step 2 will introduce further new variables $\tilde{v}_2$ whenever φ has non-pure literals.

The variables in $\tilde{v}_1$ are just a technical nuisance and can be identified with any variable of $\tilde{v}$ without loss of generality. The following brief argument should suffice in proving that. Suppose the computed normal form t' of a term t in the original φ has "extra variables", that is, variables not occurring in t. Recalling that $t =_T t'$, it is not hard to see that the denotation of t' in any model of T will not depend on the value assigned to the extra variables. Therefore, these variables can all be identified with an arbitrary variable; for instance one in $\tilde{v}$—which is non-empty by assumption. In the following then, we will assume that $\tilde{v}_1$ is enclosed in $\tilde{v}$, and concentrate on $\tilde{v}_2$ instead.

We show below that the partially Σ-restricted formula $\varphi_1 \wedge \varphi_2 \wedge \mathit{res}^\Sigma(\tilde{v})$ is satisfiable in T if and only if there is an identification ξ of $\tilde{u} := \tilde{v} \cup \tilde{v}_2$ that identifies no variables in $\tilde{v}$ and makes the totally Σ-restricted formula $(\varphi_1 \wedge \varphi_2)\xi \wedge \mathit{res}^\Sigma(\tilde{u}\xi)$ satisfiable in T. From this, the proposition's claim will then easily follow.

Assume there is a $\xi \in \mathrm{ID}(\tilde{u})$ such that ξ identifies no two variables in $\tilde{v}$ and $(\varphi_1 \wedge \varphi_2)\xi \wedge \mathit{res}^\Sigma(\tilde{u}\xi)$ is satisfiable in T. Observing that $\tilde{v}$ is contained in $\tilde{u}\xi$, we can conclude by the definition of res^Σ that $(\varphi_1 \wedge \varphi_2)\xi \wedge \mathit{res}^\Sigma(\tilde{v})$ is satisfiable in T. But then, $\varphi_1 \wedge \varphi_2 \wedge \mathit{res}^\Sigma(\tilde{v})$ is also satisfiable in T.

Now assume that $\varphi_1 \wedge \varphi_2 \wedge \mathit{res}^\Sigma(\tilde{v})$ is satisfiable in T. By construction of φ_i and definition of res^Σ, we can conclude that $S \cup \mathit{iso}^\Sigma(\tilde{v}) \cup \mathit{dif}(\tilde{v})$ is satisfiable in T, where S is the set generated at step 2 of the procedure above. By Lemma 68 then $S' := S \cup \mathit{iso}^\Sigma(\tilde{u}) \cup \mathit{dif}(\tilde{v})$ is satisfiable in T. Notice that every valuation satisfying S' in a model of T will assign distinct individuals to the variables in $\tilde{v}$. Let α be any such valuation and let ξ be the identification of $\tilde{u}$ induced by α.[42] It is immediate that ξ identifies no two variables in $\tilde{v}$ and that the set

$$(S \ \cup \ \mathit{iso}^\Sigma(\tilde{u}) \ \cup \ \mathit{dif}(\tilde{v}))\xi$$

is satisfiable in T. But this is equivalent to saying that $S\xi \cup \mathit{iso}^\Sigma(\tilde{u}\xi) \cup \mathit{dif}(\tilde{u}\xi)$ is satisfiable in T. It follows from the construction of φ_i and the definition of res^Σ that $(\varphi_1 \wedge \varphi_2)\xi \wedge \mathit{res}^\Sigma(\tilde{u}\xi)$ is satisfiable in T. $\square$

Finally, we obtain the following decidability result.

Theorem 70 *Let T_1, T_2 be such that for $i = 1, 2$,*

[42] That is, the substitution that identifies two variables in $\tilde{u}$ iff α maps them to the same individual.

- *T_i is the (complete) theory of some free Σ_i-structure $\mathcal{A}_i$ with a countably infinite basis;*
- $\mathrm{At}(\mathcal{A}_1^\Sigma) = \mathrm{At}(\mathcal{A}_2^\Sigma)$*;*
- *$G_{T_i}(\Sigma, V)$ is closed under instantiation into itself;*
- *Σ is a finite set of constructors for T_i;*
- *normal forms are computable for Σ and T_i;*
- *the word problem for T_i is decidable.*

If the satisfiability in T_i of formulae in $TRes(Qff^{\Sigma_i}, \Sigma)$ is decidable for $i = 1, 2$, then the satisfiability in $T := T_1 \cup T_2$ of formulae in $TRes(Qff^{\Sigma_1 \cup \Sigma_2}, \Sigma)$ is also decidable.

PROOF. By Proposition 64, Proposition 69, and our earlier observation on how to use our combination procedure deterministically with totally restricted formulae. □

An interesting and immediate corollary of the theorem above is that, under the same assumptions on T_1 and T_2, if the satisfiability of totally Σ-restricted quantifier-free formulae is decidable in each theory, then the satisfiability of *ground* $(\Sigma_1 \cup \Sigma_2)$-formulae is decidable in their union.

In their full generality, the conditions on T_1 and T_2 for the combination result above might appear somewhat arcane. The reader might be wondering what kinds of theories are there that satisfy them all. A more specific class of theories that does so is presented in (Tinelli and Ringeissen, 2001) for the case complete theories of free algebras. There, we reformulate the above conditions in terms of more familiar properties of equational theories, and provide some specific examples as well.

9 Conclusions and Further Research

In this paper we have described some general conditions for the combination of satisfiability procedures for constraint theories and languages that may have symbols in common. Building on the main ideas behind the combination method by Nelson and Oppen, we have developed a general non-deterministic procedure for reducing constraint satisfiability in a combined theory to constraint satisfiability in its component theories. To achieve this, we have started by investigating the main model-theoretic issues involved in theory combination.

We have defined the concept of 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 the local satisfiability of arbitrary first-order constraints with respect to two fusible structures relates to the 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 also possible to obtain combination results for constraint satisfiability with respect to theories and their unions.

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 our generalization of the Nelson-Oppen procedure can be applied in a sound and complete way to N-O-combinable theories and produce a constraint satisfiability procedure for the union of the theories.

Then, we have provided some sufficient conditions for N-O-combinability by using the concept of stable Σ-freeness, a natural extension of Nelson and Oppen's stable-infiniteness requirement for theories with non-disjoint signatures. Finally, we have illustrated an applications of our combination results to the case of theories sharing constructors.

We believe that the work described here provides a better understanding of the principles of combining constraint reasoners in the case of non-disjoint signatures. Undoubtedly, more work needs to be done to improve the scope of our theoretical results as well as identify concrete cases from the constraint-based reasoning practice to which such results can be applied.

In particular, we think that an improved definition of N-O-combinability is needed. The current one basically states that two theories are N-O-combinable if whenever a constraint φ_1 is satisfiable in one of them and a constraint φ_2 is satisfiable in the other, the only way for φ_1 and φ_2 to be inconsistent in the union theory is to entail "incompatible" Σ-restrictions for their shared variables. On the one hand, it appears that this condition is strong enough to rule out many examples of constraint theories used in constraint-based reasoning. On the other hand, it seems that a less restrictive definition of N-O-combinability would correspondingly require a more general definition of Σ-restriction; and at the moment—other than making every Σ-formula a possible Σ-restriction—it is not clear just what this definition could be.

If the definition of N-O-combinability cannot be reasonably improved, the problem remains of finding good sufficient conditions for it. The stable Σ-freeness property, which we have identified for this purpose, is not completely satisfactory for the reasons we have explained in Subsection 6.2. More work in this direction is also needed. For practical purposes, an alternative to finding general sufficient conditions for N-O-combinability may be to look at concrete cases of theories one would be interested in combining and try to show directly that they are N-O-combinable. For some of these theories it might even be possible to show that there is a finite bound on the number of Σ-restrictions that need to be considered for completeness sake. In that case, the combination procedure might be turned into one that converges on all inputs.

Finally, we think it might be beneficial to recast our results in terms of many-sorted (or better order-sorted (Goguen and Meseguer, 1992)) logic. In a sense, 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. The case for a sorted logic is possibly even more pressing in a combination context: even if two theories T_1 and T_2 are adequately described with no sorts, their combination may not be.[43] Reformulating our model-theoretic results and definitions into many-sorted logic might make it easier for two given theories to be N-O-combinable. The intuition is that N-O-combinability is easier to achieve if one reduces both the constraint 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).

Adopting a sorted framework would also have the practical advantage of reducing the non-determinism of the procedure's instantiation and identification steps because shared variables would only be replaceable by terms or variables of a compatible sort. Furthermore, it would make Σ-restrictions more natural. In fact, similarly to what we have seen in Example 18, under reasonable assumptions on Σ and the sort structure, including the assumption that Σ consists of the constructors of a certain sort S, declaring a free variable to be of a sort other than S would make it automatically Σ-restricted.

A Our Constructors vs. Constructors in Term Rewriting

In (Baader and Tinelli, 2001) it is shown that our notion of constructors subsumes the one given in (Domenjoud et al., 1994). In this appendix we show that it is also a natural generalization of the notion of constructors used in Term Rewriting.

Specifically, we prove that the set of constructors of any confluent and (weakly) normalizing term rewriting system (TRS) R is also a set of constructors in the sense of Definition 50 for the equational theory induced by R. We will not

[43] For instance, one could think of obtaining the theory of lists of real numbers as the union of the theory of lists and the theory of real numbers. Now, while each theory has an adequate unsorted axiomatization, their combination gives rise to pointless formulae such as $[1,2] + [1] \equiv 0$.

provide a direct proof of such a claim. Instead, we will show that the claim is a corollary of a more general result about TRSs *modulo* an equational theory, as defined in (Jouannaud and Kirchner, 1986).

We will assume that the reader is familiar with Term Rewriting and so we will introduce only the terminology and the notation needed to prove our claims. Comprehensive introductions to the field can be found in (Baader and Nipkow, 1998; Dershowitz and Jouannaud, 1990; Wechler, 1992), among others. Since all the signatures in question will be functional and all the theories of interest equational, we will speak of *algebras* rather than structures. Similarly, since the only atomic formulae will be equations, we will speak of the *equational* theory of a theory/algebra instead of the atomic theory.

We will first consider the equational Ω-theory E generated by a term rewriting system R modulo a set of collapse-free Σ-equations, for some $\Sigma \subseteq \Omega$. We will see that, under reasonable conditions, Σ is a set of constructors for E.

Constructors in term rewriting, which we call TRS-*constructors* here, are defined as follows.

Definition 71 (TRS-constructors) *Let Ω be a functional signature and R a TRS over $T(\Omega, V)$. We say that a signature $\Sigma \subseteq \Omega$ is a set of* TRS-constructors *for R if no symbol of Σ occurs at the top of the left-hand side of a rule in R.*

For the rest of the subsection, let

- Ω be a functional signature, and Σ a subset of Ω,
- E an equational theory of signature Ω,
- E_0 a collapse-free equational theory of signature Σ and
- R a set of rewrite rules built over $T(\Omega, V)$.

We will often need to consider the equivalence in E_0 of terms from $T(\Omega, V)$, not just $T(\Sigma, V)$. Formally, this is done by considering the Ω-theory E_0^Ω defined as the union of E_0 and the empty $(\Omega \setminus \Sigma)$-theory. To simplify the notation, we will often write $s =_{E_0} t$ instead of $s =_{E_0^\Omega} t$, for Ω-terms s, t that are equivalent in E_0^Ω.

Definition 72 *We denote by $S = (R, E_0)$ the TRS R* modulo E_0*, that is, the TRS whose rewrite relation $\rightarrow_S$ over $T(\Omega, V)$ is defined as follows. For all $s, t \in T(\Omega, V)$, $s \rightarrow_S t$ if there exists a position p, a substitution σ, and a rule $l \rightarrow r \in R$ such that $s_{|p} =_{E_0} l\sigma$ and $t = s[p \hookleftarrow r\sigma]$.*

We say that a term t' is a normal form (w.r.t. $\rightarrow_S$) *of an Ω-term t iff t' is irreducible by $\rightarrow_S$ and $t \xrightarrow{*}_S t'$. We say that two Ω-terms t_1, t_2 are* joinable modulo E_0 *iff there are two Ω-terms t_1', t_2' such that $t_1 \xrightarrow{*}_S t_1'$, $t_2 \xrightarrow{*}_S t_2'$, and $t_1' =_{E_0} t_2'$.*

As customary, the notation $s_{|p}$ above denotes the subterm of s at position p, $s[p \hookleftarrow r\sigma]$ denotes the term obtained by replacing $s_{|p}$ in s by $r\sigma$, and $\xrightarrow{*}_S$ denotes the reflexive transitive closure of $\rightarrow_S$. Note that, when the theory E_0 is empty, $\rightarrow_S$ is a term rewriting relation in the usual sense. Correspondingly, the definitions of normal form and of joinable modulo E_0 reduce to the usual ones.

An example of a TRS R modulo E_0 is the following.

Example 73 *E_0 is the theory presented by the axiom:*

$$\forall x, y, z.\ x \cdot (y \cdot z) \equiv (x \cdot y) \cdot z$$

and R is the TRS consisting of the rules:

$$\begin{array}{llll} \mathsf{rev}(x \cdot y) & \rightarrow \mathsf{rev}(y) \cdot \mathsf{rev}(x), & \qquad \mathsf{rev}(0) \rightarrow 0, \\ \mathsf{rev}(\mathsf{rev}(x)) & \rightarrow x, & \qquad \mathsf{rev}(1) \rightarrow 1. \end{array}$$

Observe that $\Sigma := \{\cdot, 0, 1\}$ is a set of TRS-constructors for R.

Definition 74 (Completeness) *The TRS $S = (R, E_0)$ is* semi-complete for *the theory E iff all of the following hold:*

(1) the relation $=_E$ coincides with $(=_{E_0} \cup \leftrightarrow_S)^*$ on $T(\Omega, V)$—or, equivalently, E is axiomatized by $E_0 \cup \{\tilde{\forall}\, l \equiv r \mid l \rightarrow r \in R\}$;

(2) the relation $\rightarrow_S$ is normalizing*, i.e., every Ω-term t has a normal form w.r.t $\rightarrow_S$;*

(3) the relation $\rightarrow_S$ is confluent modulo E_0*, i.e., for all Ω-terms t, t_1, t_2 such that $t_1 \xleftarrow{*}_S t \xrightarrow{*}_S t_2$, t_1 and t_2 are joinable modulo E_0.*

We say that S is complete for *E iff it is semi-complete for E and $\rightarrow_S$ is* terminating*, i.e., there is no infinite sequence $(t_0, t_1, t_2, \ldots)$ such that $t_0 \rightarrow_S t_1 \rightarrow_S t_2 \rightarrow_S \cdots$.*

It is not difficult to show that when the TRS $S = (R, E_0)$ is semi-complete for E, E is non-trivial, every Ω-term is equivalent in E to its normal forms w.r.t. $\rightarrow_S$, and for all $s, t \in T(\Omega, V)$ and respective normal forms s', t',

$$s =_E t \text{ iff } s' =_{E_0} t'.$$

From this it follows that any two normal forms of the same term t are equivalent in E_0. For this reason, we will identify them all and denote them by $t\downarrow_S$.

(Semi-)Complete TRSs form a natural class of rewrite systems. The reason is that if a TRS $S = (R, E_0)$ is complete for some theory E, and the matching problem [44] and word problems in E_0 are decidable, then the normal form $t{\downarrow_S}$ of every term t is computable; as a consequence, the word problem in E is also decidable. [45]

To prove that TRS-constructors are constructors for E in the sense of Definition 50, we will appeal to well-known results from the research on the combination of decision procedures for the word-problem in a union of collapse-free, signature-disjoint equational theories (Schmidt-Schauß, 1989; Nipkow, 1991; Ringeissen, 1996a; Baader and Tinelli, 1997). Here, the union of interest will be E_0^Ω, the union of the (collapse-free) equational Σ-theory E_0 with the (collapse-free) empty $(\Omega \setminus \Sigma)$-theory.

Lemma 75 *Let E_1 and E_2 be two collapse-free equational theories of respective signature Σ_1 and Σ_2, with $\Sigma_1 \cap \Sigma_2 = \emptyset$. Then, the following holds.*

(1) The theory $E_1 \cup E_2$ is collapse-free.
(2) For all $t_1, t_2 \in T(\Sigma_1 \cup \Sigma_2, V)$ such that $t_i(\epsilon) \in \Sigma_i$ for $i = 1, 2$,

$$t_1 \neq_{E_1 \cup E_2} t_2.$$

(3) Let $\sigma \in \mathrm{SUB}(V)$, $i \in \{1, 2\}$, and let s, t be two i-pure non-variable terms such that
- *$(v\sigma)(\epsilon) \notin \Sigma_i$ for all $v \in \mathcal{V}ar(s \equiv t)$ and*
- *$u\sigma \neq_{E_1 \cup E_2} v\sigma$ for all distinct $u, v \in \mathcal{V}ar(s \equiv t)$.*

Then, $s\sigma =_{E_1 \cup E_2} t\sigma$ iff $s =_{E_i} t$.

A property of S that follows from the lemma above is the following.

Proposition 76 *If $S = (R, E_0)$ is semi-complete for E and Σ is a set of TRS-constructors for R, then*

$$f(t_1, \ldots, t_n){\downarrow_S} =_{E_0} f(t_1{\downarrow_S}, \ldots, t_n{\downarrow_S})$$

for all n-ary $f \in \Sigma$ and $t_1, \ldots, t_n \in T(\Omega, V)$.

Another property of S is that every Σ-term is in normal form w.r.t. $\to_S$.

Lemma 77 *If $S = (R, E_0)$ is semi-complete for E and Σ is a set of TRS-constructors for R, then $t{\downarrow_S} = t$ for all $t \in T(\Sigma, V)$.*

A proof of the two results above is given in (Tinelli and Ringeissen, 2001).

[44] Recall that the problem of *matching a term t_1 against a term t_2* in E_0 is the problem of determining whether there is a substitution σ such that $t_1\sigma =_{E_0} t_2$.
[45] Actually, by standard results in term rewriting, it can be shown that the word problem in E is decidable already when S is semi-complete for E.

An easily provable consequence of Lemma 77 is that, under its assumptions, two Σ-terms are equivalent in E exactly when they are equivalent in E_0. In other words, E_0 axiomatizes the equational Σ-theory of E.

We now show that when Σ is a set of TRS-constructors for R, the set $G_E(\Sigma, V)$ defined at the beginning of Subsection 7.1 coincides with the set of terms whose normal forms w.r.t. $\rightarrow_S$ do not start with a Σ-symbol.[46]

Lemma 78 *Assume that $S = (R, E_0)$ is semi-complete for E and Σ is a set of* TRS*-constructors for R. Then,*

$$G_E(\Sigma, V) = \{r \in T(\Omega, V) \mid r{\downarrow_S}(\epsilon) \notin \Sigma\}.$$

PROOF. Let $r \in T(\Omega, V)$.

($\subseteq$) Recalling the definition of $G_E(\Sigma, V)$, it is obvious that $r \notin G_E(\Sigma, V)$ whenever $r{\downarrow_S}(\epsilon) \in \Sigma$, given that $r =_E r{\downarrow_S}$.

($\supseteq$) Assume ad absurdum that $r{\downarrow_S}(\epsilon) \notin \Sigma$ but $r \notin G_E(\Sigma, V)$. Then, there is an $f \in \Sigma$ and a $\tilde{t}$ in $T(\Omega, V)$ such that $r =_E f(\tilde{t})$. By Definition 74 and Proposition 76, we can then conclude that $r{\downarrow_S} =_{E_0^\Omega} f(\tilde{t}{\downarrow_S})$. Now, if $r{\downarrow_S}(\epsilon)$ is in $\Omega \setminus \Sigma$, the above equivalence contradicts point 2 of Lemma 75. If $r{\downarrow_S}(\epsilon)$ is a variable, the equivalence contradicts the fact that E_0^Ω is collapse free by Lemma 75(1). $\square$

Together with Proposition 76, Lemma 78 has the following consequence.

Lemma 79 *Let $G := G_E(\Sigma, V)$. Assume that $S = (R, E_0)$ is semi-complete for E and Σ is a set of* TRS*-constructors for R. Then,*

$$t{\downarrow_S} \in T(\Sigma, G)$$

for all $t \in T(\Omega, V)$.

PROOF. Let $t \in T(\Omega, V)$ and assume that $t{\downarrow_S} \notin T(\Sigma, G)$. Then, it is not difficult to show by the results above that there must be a subterm r of $t{\downarrow_S}$ with $r(\epsilon) \notin \Sigma$, a function symbol $f \in \Sigma$, and a tuple $\tilde{t}$ in $T(\Omega, V)$, such that $r =_E f(\tilde{t})$. By Definition 74(3) then we have that $r{\downarrow_S} =_{E_0} f(\tilde{t}){\downarrow_S}$. Now, $r{\downarrow_S} = r$ as r is the subterm of the irreducible term $t{\downarrow_S}$, and $f(\tilde{t}){\downarrow_S} =_{E_0} f(\tilde{t}{\downarrow_S})$

[46] Notice that when $S = (R, E_0)$ is semi-complete for E, a term has a normal form with top symbol in Σ iff all its normal forms have their top symbol in Σ, as one can easily show.

by Proposition 76. But this entails that $r =_{E_0^\Omega} f(\tilde{t}\downarrow_S)$, which is impossible by Lemma 75(2). $\square$

We are now ready to prove the main result of this subsection.

Proposition 80 *If $S = (R, E_0)$ is semi-complete for E and Σ is a set of* TRS-*constructors for R, then Σ is a set of constructors for E.*

PROOF. We prove the claim by showing that the three conditions of Definition 50 are satisfied. Let $G := G_E(\Sigma, V)$.

(1) Let $v \in V$. Since $v = v\downarrow_S$ by Lemma 77, we can immediately conclude by Lemma 78 that $v \in G$. It follows that $V \subseteq G$.

(2) Let $t \in T(\Omega, V)$. We have already observed that $t =_E t\downarrow_S$. From Lemma 79 we also know that $t\downarrow_S \in T(\Sigma, G)$.

(3) Let $s_1(\tilde{r}_1), s_2(\tilde{r}_2) \in T(\Sigma, G)$ and $s_1(\tilde{v}_1), s_2(\tilde{v}_2)$ be the corresponding terms obtained by abstracting $\tilde{r}_1, \tilde{r}_2$ with fresh variables so that terms equivalent in E are abstracted by the same variable. We show that $s_1(\tilde{r}_1) =_E s_2(\tilde{r}_2)$ iff $s_1(\tilde{v}_1) =_E s_2(\tilde{v}_2)$.

The right-to-left implication is immediate, hence assume that $s_1(\tilde{r}_1) =_E s_2(\tilde{r}_2)$. From the hypothesis that (R, E_0) is semi-complete for E we can conclude that

$$s_1(\tilde{r}_1)\downarrow_S =_{E_0^\Omega} s_2(\tilde{r}_2)\downarrow_S.$$

Recalling that s_1 and s_2 are Σ-terms, we can show by a simple inductive argument based on Proposition 76 that

$$s_1(\tilde{r}_1\downarrow_S) =_{E_0^\Omega} s_2(\tilde{r}_2\downarrow_S).$$

Assuming that E-equivalent terms in $\tilde{r}_1, \tilde{r}_2$ have the same normal w.r.t. $\rightarrow_S$,[47] it is easy to see that each $s_i(\tilde{r}_i\downarrow_S)$ is the result of applying to $s_i(\tilde{v}_i)$ a substitution σ satisfying Point 3 of Lemma 75. By that lemma, it then follows that $s_1(\tilde{v}_1) =_{E_0} s_2(\tilde{v}_2)$ and so $s_1(\tilde{v}_1) =_E s_2(\tilde{v}_2)$. $\square$

We would like to stress that, although the preconditions in Proposition 80 entail that Σ is a set of constructors for E, they do not entail that normal

[47] Such an assumption is with no loss of generality because normal forms of E-equivalent terms are E_0-equivalent and so can be identified in $\tilde{r}_1\downarrow_S, \tilde{r}_2\downarrow_S$.

forms in the sense of Definition 58 are computable. A sufficient condition for the computability of normal forms, under the assumptions of Proposition 80, is that E_0-matching with free constants is decidable. A proof of this can be found in (Tinelli and Ringeissen, 2001).

Finally, we can produce a result like the above for conventional TRSs again by observing that such systems are TRSs modulo the empty equational theory.

Corollary 81 *Let R be a TRS over $T(\Omega, V)$. If $\rightarrow_R$ is semi-complete and Σ is a set of* TRS*-constructors for R, then Σ is a set of constructors for the equational theory induced by R.*

To summarize, for semi-complete term rewriting systems, our notion of constructors is a generalization of the notion of TRS-constructors. In addition, it is a strict generalization, given that the equational theory over TRS-constructors is always empty (as one can easily see), which need not be the case for our constructors.

We conclude this section by sketching how the above results can be used to prove that the signature Σ in Example 61 of Section 7.3 is indeed a set of constructors. Consider the TRS $S := (R, E_0)$ where E_0 and R are defined as in Example 73. Clearly, E_0 is collapse-free, $\rightarrow_R$ is terminating (therefore, normalizing) and $\Sigma := \{\mathsf{0}, \mathsf{1}, \cdot\}$ is a set of TRS-constructors for R. It is not difficult to show that $\rightarrow_R$ is confluent modulo E_0. It follows by Proposition 80 that $\Sigma := \{\mathsf{0}, \mathsf{1}, \cdot\}$ is a set of constructors for E_{61}.

Acknowledgements

We thank the anonymous referees for their extremely thorough reviews and for their numerous and invaluable suggestions for improving the presentation of the paper.

References

Armando, A., Ranise, S., Nov. 1998. Constraint contextual rewriting. In: Caferra, R., Salzer, G. (Eds.), Proceedings of the 2nd International Workshop on First Order Theorem Proving, FTP'98, Vienna (Austria). pp. 65–75.

Baader, F., Nipkow, T., 1998. Term Rewriting and All That. Cambridge University Press, United Kingdom.

Baader, F., Schulz, K. U., 1995a. Combination of constraint solving techniques: An algebraic point of view. In: Proceedings of the 6th International

Conference on Rewriting Techniques and Applications, RTA'95. Vol. 914 of Lecture Notes in Computer Science. Springer-Verlag, pp. 50–65.

Baader, F., Schulz, K. U., 1995b. Combination techniques and decision problems for disunification. Theoretical Computer Science 142, 229–255.

Baader, F., Schulz, K. U., Sep. 1995c. 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). Vol. 976 of Lecture Notes in Artificial Intelligence. Springer-Verlag.

Baader, F., Schulz, K. U., Feb. 1996. Unification in the union of disjoint equational theories: Combining decision procedures. Journal of Symbolic Computation 21 (2), 211–243.

Baader, F., Schulz, K. U., 1998. Combination of constraint solvers for free and quasi-free structures. Theoretical Computer Science 192, 107–161.

Baader, F., Tinelli, C., 1997. A new approach for combining decision procedures for the word problem, and its connection to the Nelson-Oppen combination method. In: McCune, W. (Ed.), Proceedings of the 14th International Conference on Automated Deduction (Townsville, Australia). Vol. 1249 of Lecture Notes in Artificial Intelligence. Springer-Verlag, pp. 19–33.

Baader, F., Tinelli, C., Oct. 1998. Deciding the word problem in the union of equational theories. Tech. Rep. UIUCDCS-R-98-2073, Department of Computer Science, University of Illinois at Urbana-Champaign.

Baader, F., Tinelli, C., 1999. Deciding the word problem in the union of equational theories sharing constructors. In: Narendran, P., Rusinowitch, M. (Eds.), Proceedings of the 10th International Conference on Rewriting Techniques and Applications (Trento, Italy). Vol. 1631 of Lecture Notes in Computer Science. Springer-Verlag, pp. 175–189.

Baader, F., Tinelli, C., 2001. Deciding the word problem in the union of equational theories. Information and Computation (to appear).

Baumgartner, P., Furbach, U., Petermann, U., 1992. A unified approach to theory reasoning. Research Report 15–92, Universität Koblenz-Landau, Koblenz, Germany, fachberichte Informatik.

Boudet, A., Dec. 1993. Combining unification algorithms. Journal of Symbolic Computation 16 (6), 597–626.

Bürckert, H.-J., 1994. A resolution principle for constraint logics. Artificial Intelligence 66, 235–271.

Dershowitz, N., Jouannaud, J.-P., 1990. Rewriting systems. In: van Leeuwen, J. (Ed.), Handbook of Theoretical Computer Science. Elsevier Publishers, Amsterdam, pp. 243–320.

Domenjoud, E., Klay, F., Ringeissen, C., 1994. Combination techniques for non-disjoint equational theories. In: Bundy, A. (Ed.), Proceedings of the 12th International Conference on Automated Deduction, Nancy (France). Vol. 814 of Lecture Notes in Artificial Intelligence. Springer-Verlag, pp. 267–281.

Ehrig, H., Mahr, B., 1985. Fundamentals of Algebraic Specification 1: Equations and Initial Semantics. Vol. 6 of EATCS Monographs on Theoretical Computer Science. Springer-Verlag, New York, N.Y.

Ehrig, H., Mahr, B., 1990. Fundamentals of Algebraic Specification 2: Module Specifications and Constraints. Vol. 21 of EATCS Monographs on Theoretical Computer Science. Springer-Verlag, New York, N.Y.

Fiorentini, C., Ghilardi, S., 2001. Combining word problems through rewriting in categories with products. Theoretical Computer Science (to appear).

Goguen, J. A., Meseguer, J., Nov. 1992. Order sorted algebra I. Equational deduction for multiple inheritance, overloading, exceptions and partial operations. Theoretical Computer Science 105 (2), 217–273.

Gramlich, B., 1996. On termination and confluence properties of disjoint and constructor-sharing conditional rewrite systems. Theoretical Computer Science 165 (1), 97–131.

Herold, A., 1986. Combination of unification algorithms. In: Siekmann, J. (Ed.), Proceedings 8th International Conference on Automated Deduction, Oxford (UK). Vol. 230 of Lecture Notes in Artificial Intelligence. Springer-Verlag, pp. 450–469.

Hodges, W., 1993. Model Theory. Vol. 42 of Enclyclopedia of mathematics and its applications. Cambridge University Press.

Höhfeld, M., Smolka, G., Oct. 1988. Definite relations over constraint languages. LILOG Report 53, IWBS, IBM Deutschland, Postfach 80 08 80, 7000 Stuttgart 80, Germany.

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

Jaffar, J., Maher, M., 1994. Constraint Logic Programming: A Survey. Journal of Logic Programming 19/20, 503–581.

Jouannaud, J.-P., Kirchner, H., 1986. Completion of a set of rules modulo a set of equations. SIAM Journal on Computing 15 (4), 1155–1194, preliminary version in Proceedings 11th ACM Symposium on Principles of Programming Languages, Salt Lake City (USA), 1984.

Kepser, S., Schulz, K. U., Aug. 1996. Combination of constraint systems II: Rational amalgamation. In: Freuder, E. C. (Ed.), Proceedings of the 2nd International Conference on Principles and Practice of Constraint Programming, Cambridge, MA, USA. Vol. 1118 of Lecture Notes in Computer Science. Springer-Verlag, pp. 282–296.

Kirchner, C., Kirchner, H., Rusinowitch, M., 1990. Deduction with symbolic constraints. Revue Française d'Intelligence Artificielle 4 (3), 9–52, special issue on Automatic Deduction.

Kirchner, H., Ringeissen, C., 1994a. Combining symbolic constraint solvers on algebraic domains. Journal of Symbolic Computation 18 (2), 113–155.

Kirchner, H., Ringeissen, C., 1994b. Constraint solving by narrowing in combined algebraic domains. In: Van Hentenryck, P. (Ed.), Proceedings of the 11th International Conference on Logic Programming. The MIT press, pp. 617–631.

Nelson, G., 1984. Combining satisfiability procedures by equality-sharing. In:

Bledsoe, W. W., Loveland, D. W. (Eds.), Automated Theorem Proving: After 25 Years. Vol. 29 of Contemporary Mathematics. American Mathematical Society, Providence, RI, pp. 201–211.

Nelson, G., Oppen, D. C., Oct. 1979. Simplification by cooperating decision procedures. ACM Trans. on Programming Languages and Systems 1 (2), 245–257.

Nipkow, T., 1991. Combining matching algorithms: The regular case. Journal of Symbolic Computation 12, 633–653.

Ohlebusch, E., 1995. Modular properties of composable term rewriting systems. Journal of Symbolic Computation 20 (1), 1–41.

Oppen, D. C., 1980. Complexity, convexity and combinations of theories. Theoretical Computer Science 12.

Pillay, A., Tsuboi, A., Dec. 1997. Amalgamations preserving $\aleph_0$-categoricity. The Journal of Symbolic Logic 62 (4), 1070–1074.

Ringeissen, C., 1992. Unification in a combination of equational theories with shared constants and its application to primal algebras. In: Voronkov, A. (Ed.), Proceedings of the 1st International Conference on Logic Programming and Automated Reasoning. Vol. 624 of Lecture Notes in Artificial Intelligence. Springer-Verlag, pp. 261–272.

Ringeissen, C., Dec. 1993. Combinaison de résolutions de contraines. Thèse de Doctorat d' Université, Université de Nancy 1, Nancy, France.

Ringeissen, C., May 1996a. Combining decision algorithms for matching in the union of disjoint equational theories. Information and Computation 126 (2), 144–160.

Ringeissen, C., Mar. 1996b. Cooperation of decision procedures for the satisfiability problem. In: Baader, F., Schulz, K. (Eds.), Frontiers of Combining Systems: Proceedings of the 1st International Workshop, Munich (Germany). Applied Logic. Kluwer Academic Publishers, pp. 121–140.

Schmidt-Schauß, M., July/August 1989. Unification in a combination of arbitrary disjoint equational theories. Journal of Symbolic Computation 8 (1–2), 51–100, special issue on unification. Part II.

Schulz, K. U., Mar. 2000. Why combined decision problems are often intractable. In: Kirchner, H., Ringeissen, C. (Eds.), Proceedings of the 3rd International Workshop on Frontiers of Combining Systems, FroCoS'2000, Nancy (France). Vol. 1794 of Lecture Notes in Artificial Intelligence. Springer-Verlag, pp. 217–244.

Shostak, R. E., Apr. 1979. A practical decision procedure for arithmetic with function symbols. Journal of the ACM 26 (2), 351–360.

Shostak, R. E., 1984. Deciding combinations of theories. Journal of the ACM 31, 1–12.

Tinelli, C., May 1999. Combination of decidability procedures for automated deduction and constraint-based reasoning. PhD dissertation, Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana-Champaign, Illinois.

Tinelli, C., Harandi, M. T., Mar. 1996. A new correctness proof of the Nelson–

Oppen combination procedure. In: Baader, F., Schulz, K. (Eds.), Frontiers of Combining Systems: Proceedings of the 1st International Workshop (Munich, Germany). Applied Logic. Kluwer Academic Publishers, pp. 103–120.
Tinelli, C., Ringeissen, C., Apr. 1998. Non-disjoint unions of theories and combinations of satisfiability procedures: First results. Tech. Rep. UIUCDCS-R-98-2044, Department of Computer Science, University of Illinois at Urbana-Champaign, (also available as INRIA research report no. RR-3402).
Tinelli, C., Ringeissen, C., Apr. 2001. Unions of non-disjoint theories and combinations of satisfiability procedures. Tech. Rep. 01-02, Department of Computer Science, University of Iowa, (extended version of this paper).
Wechler, W., 1992. Universal Algebra for Computer Scientists. Vol. 25 of EATCS Monographs on Theoretical Computer Science. Springer-Verlag, Berlin-Heidelberg-New York.
Wolter, F., 1998. Fusions of modal logics revisited. In: Kracht, M., de Rijke, M., Wansing, H., Zakharyaschev, M. (Eds.), Advances in Modal Logic. CSLI, Stanford, CA.