Unions of Non-Disjoint Theories and Combinations of Satisfiability Procedures

Cesare Tinelli
tinelli@cs.uiowa.edu

Christophe Ringeissen[1]
Christophe.Ringeissen@loria.fr

Department of Computer Science
University of Iowa

Report No. 01-02

April 2001

[1]Address: INRIA Lorraine and CRIN-CNRS, BP 101, 54602 Villers-lès-Nancy Cedex, France.

Unions of Non-Disjoint Theories and Combinations of Satisfiability Procedures

Cesare Tinelli
Department of Computer Science
University of Iowa
14 McLean Hall, Iowa City, IA 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

August 1, 2001

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.

Keywords: combination of satisfiability procedures, decision problems, constraint-based reasoning, automated deduction.

Contents

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.[1]

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 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 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 [Sch00] 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 this study.

[1] 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.

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 and related problems [BS95b, Bou93, DKR94, Her86, KR94a, KR94b, Rin92, SS89]. 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 [JM94], constrained resolution [B$\ddot{9}$4] and what is generally referred to as *theory-reasoning* [BFP92], 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 [JM94].

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 [HS88], the deduction with constraints framework [KKR90], constrained resolution [B$\ddot{9}$4] constraint contextual rewriting [AR98], and—at least at the ground level—all theory-reasoning frameworks [BFP92].

Essentially all existing results in the combination of constraint reasoners in the first class come from the work of F. Baader and K. Schulz [BS95a, BS95c, KS96, BS98], 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 [Sho79] and in [Sho84]. Shostak's approach is limited in scope and not very modular—admittedly on purpose, for efficiency reasons. A rather general and completely modular combination method was proposed by G. Nelson and D. Oppen in [NO79] and then slightly revised in

[Nel84]. 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 [Opp80]. In [TH96], C. Tinelli (the first of us) 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 (the second of us) in [Rin93], which describes the procedure as a set of derivation rules applied non-deterministically.

All the work mentioned above shares 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 main 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 [Ohl95, Gra96] 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*.[2] Outside the work on modular term rewriting, the first combination results for the word problem in the union of non-disjoint constraint theories were given in [DKR94] as a consequence of some combination techniques based on an adequate notion of (shared) constructors. The second of us used similar ideas later in [Rin96b] to extend the Nelson-Oppen method to theories sharing constructors in a sense close to that of [DKR94].

To our knowledge, the only new work since [Rin96b] 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 [BT01]. These papers discuss a very general decision procedure for the word problem in the union of equational theories with non-disjoint signatures.[3] The procedure's correctness proof is based on some of the model-theoretic results reported here. Part of the work reported here is also described in [Tin99]; a preliminary account was given in [TR98].

[2]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 left-hand side.

[3]An alternative and, as it turns out, equivalent approach to this topic has been very recently proposed by C. Fiorentini and S. Ghilardi in [FG01].

1.2 Our Contribution

In this paper we focus on constraint satisfiability problems 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 components. Our main research effort then has consisted in developing appropriate restrictions on T_1 and T_2 and $\mathcal{L}$ that make the above combination problem solvable. As mentioned earlier, Nelson and Oppen's 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 a *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 basic "mixed" constraints in a fusion structure is reducible to the satisfiability of pure constraints[4] in the fusion 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,[5] but tailored to the case of theories with not necessarily

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

[5]More precisely, of its non-deterministic version, where the added restrictions are simply con-

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 Nelson and Oppen's combination results and Ringeissen's initial results in [Rin96b].

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 a stably-infinite 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 describe some examples of classes of stably Σ-free theories that are N-O-combinable. In the most important of these examples, the theories share constructors in the sense of Section 7.

In Section 9, we show that, even if designed to combine satisfiability procedures, our combination method can also be used to combine certain decision procedures for the validity of existential equational constraints. Then, we show that in general our method can be seen as an approximation of a combination method for equational constraints due to Baader and Schulz [BS95b].

In Section 10, we provide a further application of our method, this time to the combination of constraint solvers for certain *initial* structures.

junctions of equations and disequations between shared variables. See, e.g., [TH96] for details.

Section 11 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 [Hod93] and [Wec92].

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*, i.e., first-order terms of signature Σ^{F}. 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 is *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 $n \geq 0$.[6] 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., $\mathcal{V}ar(\varphi) = \{v_1, \ldots, v_n\}$.[7] 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

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

[7] 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.

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 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. As common, we denote the empty substitution by ε and write substitution applications in postfix form. Also, if σ is a substitution we call the sets

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

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 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 of Σ-isolated individuals only 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 set 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.[8]

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.[9] 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 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., [Hod93]) 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 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

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

[9]Our definition is slightly more restrictive than the standard one, in which t is required to be a non-variable term. According to that definition, if Σ has no 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.

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*[10]. 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 will rely on the standard notions of morphisms of structures from Model Theory [Hod93]. We will 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*

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

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 [BS98] and Kepser and Schulz's rational amalgamation [KS96] of two quasi-free structures are both readily shown to be a fusion of those structures. Similarly, the amalgamation construction given by Ringeissen in [Rin96b] can also be shown to produce a fusion.

In principle, one could imagine a notion of a 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

[10]We initially chose the term "fusion" to avoid overloading the term "amalgamation", which has a more specific meaning in the Model Theory literature. We have later discovered that [PT97] does use "amalgamation" for the same type of combined structure as ours while [Hol95] 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., [Wol98]).

contain non-isomorphic structures.[11] 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, establishing 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 $\Sigma := \Sigma_A \cap \Sigma_B$. To simplify the notation, in this proof and in the rest of the paper we adopt the following notational convention. If $h\colon C \to D$ 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 $\{h(\tilde{c}) \mid \tilde{c} \in R\}$.

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

($\Leftarrow$) Let h be a (bijective) 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 $h :\mathcal{A} \cong \mathcal{C}^{\Sigma_A}$.

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

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

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

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

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

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

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

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

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

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 $\mathcal{B}$. □

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* freely 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 freely fusible with $\mathcal{B}$ over $\langle X, Y \rangle$ whenever $\mathcal{B}$ is freely 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 $\tilde{v} := (v_1, \dots, v_m)$. 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, $\tilde{x} := (x_1, \dots, x_m)$ is in X, $\tilde{y} := (y_1, \dots, y_m)$ is in Y, and neither $\tilde{x}$ nor $\tilde{y}$ contains repetitions. Consider the map $h \colon \tilde{x} \to Y$ such that,

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

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 $K := \{k_1, \dots, k_m\}$ 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 $j \in \{1, \dots, m\}$,

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

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 $\sigma := \{v_1 \leftarrow k_1, \dots, v_m \leftarrow k_m\}$. 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 $\mathcal{F}^{\Omega \cup \Delta} \models \tilde{\exists}\,(\varphi \wedge \psi)$. □

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 will need a more syntactic 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.

3.1 Disjoint Signatures

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,*[12]

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

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

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 Σ_{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. □

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 $dif(\tilde{v})$ 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 $\mathcal{A}_2$[13] 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 $dif(\tilde{v}\xi)$, can be used in the same way $dif(\tilde{v})$ was used before.

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

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 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

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

which we will denote just by iso_ρ when Σ and U are clear from the context.

Observe that the set iso^Σ_ρ 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., to 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 $\mathcal{A}_i \models \varphi_i[\tilde{a}_i, \tilde{x}_i] \wedge iso^\Sigma[\tilde{x}_i] \wedge dif[\tilde{x}_i]$. 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 $dif[\tilde{x}_i]$ is true in $\mathcal{A}_i$. To see that $\tilde{x}_i$ is included in X_i, just recall that $iso^\Sigma[\tilde{x}_i]$ 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 is clear that we actually have a slightly stronger result: when the conditions of the Lemma 13 hold, the whole formula $\varphi_1 \wedge \varphi_2 \wedge iso^\Sigma(\tilde{v}) \wedge dif(\tilde{v})$ 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, $iso^\Sigma(\tilde{v}) \wedge dif(\tilde{v})$, 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 $Is(\mathcal{A}_i{}^\Sigma) \subseteq X_i \subseteq A_i$ 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 $Is({\mathcal{A}_1}^{\Sigma})$ and $Is({\mathcal{A}_2}^{\Sigma})$ are empty, if $\mathcal{A}_1$ and $\mathcal{A}_2$ are Σ-fusible, the pair of sets on which they are fusible is univocally determined and coincides with $\langle Is({\mathcal{A}_1}^{\Sigma}), Is({\mathcal{A}_2}^{\Sigma})\rangle$.

Proposition 15 *Let $\mathcal{A}_1$ and $\mathcal{A}_2$ be two structures Σ-fusible for some finite $\Sigma \subseteq \Sigma_{\mathcal{A}_1} \cap \Sigma_{\mathcal{A}_2}$. 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*

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

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 $(\varphi_1' \wedge \varphi_2') = (\varphi_1 \wedge \varphi_2)\rho\xi$. □

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 $\alpha(v_j) = t_j^{\mathcal{F}}[\tilde{x}_j]$.[14] 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. □

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

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

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

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

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 $(\varphi\rho \wedge iso_{\rho})\xi \wedge dif_{\xi}$ (such as 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} will be 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 will contain 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. Since the notion of Σ-isolated individual is quite technical, what may not be clear of 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 consider just the iso^{Σ} component as the Σ-restriction and ignore the dif component, which is essentially unproblematic. 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 only satisfied 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—which we use 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.[15] 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 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[16]: every list in $\mathcal{A}$ is a (possibly nested) list of *atoms*, the elements of $N \cup I$. This is what it means in our terminology for $\mathcal{A}$ to be 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 $\{v \leftarrow t(\tilde{u})\}$. It should be now easy to see that, in general, a 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.

[15] 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.

[16] 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.

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. The Σ-isolated individuals of $\mathcal{B}$ are exactly the elements of $W \cup J$.

First we show that $\mathcal{A}$ and $\mathcal{B}$ have a fusion. Observing that $N \cup I$ 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 $a_1, a_2 \in A$. It should be easy to see that h_{A-B} is in fact a bijection of A onto B, which entails that $h_{A-B} \colon \mathcal{A}^{\Sigma} \cong \mathcal{B}^{\Sigma}$. 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.

Although it is possible to show directly that $\mathcal{A}$ and $\mathcal{B}$ are Σ-fusible, we will do that by using some general results about the fusibility of *free* structures. But for that we will have to 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 $\mathit{Fus}(T_1, T_2)$ denote the following class of structures:

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

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

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

($\subseteq$) Assume that $\mathcal{F}$ is a fusion of some $\mathcal{A} \in \mathit{Mod}(T_1)$ and $\mathcal{B} \in \mathit{Mod}(T_2)$. 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 $\mathcal{C} \in \mathit{Mod}(T_1 \cup T_2)$ is a fusion of $\mathcal{C}^{\Sigma_1}$ and $\mathcal{C}^{\Sigma_2}$ and that $\mathcal{C}^{\Sigma_i}$ models T_i, for $i = 1, 2$. □

Recalling Proposition 6 on the existence of fusions, we have the following corollary, first proved in [Rin96b] and [TH96].

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 on 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 in the following, we will only consider pairs in which at most one of the two formulae is, or has subformulae, 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 $\mathcal{L}^{\Sigma_1} \times \mathcal{L}^{\Sigma_2}$:*

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

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 $\langle \varphi_1, \varphi_2 \rangle$.[17] 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 24 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 24 and Condition 25 below hold for all $\langle \varphi_1, \varphi_2 \rangle \in \mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$.*

Condition 24 *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.

[17] 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 $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$.

Condition 25 *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.*

While Condition 25 is straightforward and easy to understand, it may be hard to grasp Condition 24 at an intuitive level. To do that it is helpful to concentrate on the case in which ρ is the empty instantiation (and iso_ρ is then empty), 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 24, for no ρ and ξ is a model of T_1 satisfying $((\varphi_1 \wedge \varphi)\rho \wedge iso_\rho)\xi \wedge dif_\xi$ fusible with a model of T_2 satisfying $((\varphi_2 \wedge \neg\varphi)\rho \wedge iso_\rho)\xi \wedge dif_\xi$.[18]

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.[19] However, we doubt that much improvement can be achieved without abandoning a strictly syntactical definition of $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$.

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

Proposition 26 *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 24

[18] 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 24.

[19] A case in point are pairs 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 $T_1 := \{\forall x, y.\, P_1(x, y) \Rightarrow x \equiv y\}$, $T_2 := \{a \equiv a,\, b \equiv b\}$ 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 $(P_1(x, y)\rho \wedge iso_\rho)\xi \wedge dif_\xi = P_1(a, b)$ and models $T_2 \cup T_3$ satisfying $(P_3(x, y)\rho \wedge iso_\rho)\xi \wedge dif_\xi = P_3(a, b)$ do exist, but they are obviously not fusible.

(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 $T_1 \cup T_2$. □

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 make viable 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 27 *Let T_1 and T_2 be 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 24, Proposition 15 and Proposition 20. □

If T_1 and T_2 satisfy Condition 25 as well, the implication in the proposition above becomes a full equivalence.

Theorem 28 *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 25, Proposition 20 and Proposition 16. □

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

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

1. 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 $\tilde{v} := \mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2)$.

2. 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 $\xi \in \mathrm{ID}(\mathcal{V}ar(\tilde{v}\rho))$.

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

Figure 1: The Combination Procedure.

- 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 [NO79], 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 (steps 1 and 2). It has also been considered by Baader and Schulz in their own combination methods, starting with the one described in [BS96].

Proposition 27 immediately tells us that the procedure in Figure 1 is sound for component theories that are partially N-O-combinable over the given language $\mathcal{L}$.

Proposition 29 (Soundness) *Let T_1 and T_2 be partially N-O-combinable over $\mathcal{L}$. 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 $i = 1, 2$, then the input pair $\langle\varphi_1, \varphi_2\rangle$ is satisfiable in $T_1 \cup T_2$.*

If the component theories are totally N-O-combinable over $\mathcal{L}$, Theorem 28 tells us that the procedure is also complete, in the sense specified below.

Proposition 30 (Completeness) *Let T_1 and T_2 be totally N-O-combinable over $\mathcal{L}$. If the input pair $\langle\varphi_1, \varphi_2\rangle$ is satisfiable in $T_1 \cup T_2$, then there is a pair $\langle\psi_1, \psi_2\rangle$ among the possible outputs of the identification step such that ψ_i is satisfiable in T_i for $i = 1, 2$.*

The formula ψ_i $(i = 1, 2)$ in the two results above, which has the form $(\varphi_i\rho \wedge iso_\rho)\xi \wedge dif_\xi$, is a Σ-restricted formula in the sense of Definition 17. More precisely, ψ_i is an element of $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ as $\varphi_i \in \mathcal{L}^{\Sigma_i}$ and $\mathcal{L}$ is closed under identification and instantiation. For for Step 3 of the combination 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 $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$. 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 31 *Assume that T_1 and T_2 be totally N-O-combinable over $\mathcal{L}$ and the satisfiability in T_i of formulae in $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ 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 31 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 Step 1. 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 make both formulae ψ_1 and ψ_2 in Step 3 satisfiable in their respective theory. It follows that the procedure will in general diverge[20] on unsatisfiable inputs.

Notice 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

[20]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.

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, though, even if it is only a semi-decision procedure, the procedure does yield decidability results when the theories considered are axiomatizable.[21] 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—once we notice that $T_1 \cup T_2$ is axiomatizable whenever both T_1 and T_2 are.

Proposition 32 *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 always expressible 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 possible, but at the cost of some closure assumptions on $\mathcal{L}$.[22] We describe such assumptions in the following and then show, as an example, how they let us improve on Proposition 32.

Definition 33 *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.*

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

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

If the class $\mathcal{L}$ specified 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. We can now express the previous decidability result more neatly as follows.

Proposition 34 *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 $\mathcal{L}$ is effectively purifiable w.r.t. $\langle\Sigma_1,\, \Sigma_2\rangle$ and T_1 and T_2 are N-O-combinable over $\mathcal{L}$, then the satisfiability in $T_1 \cup T_2$ of formulae of $\mathcal{L}^{\Sigma_1\cup\Sigma_2}$ is decidable.*

The above proposition seems to suggest that we get a somewhat *weaker* decidability result for the union theory, since we start with restricted formulae in the component theories and end up with unrestricted formulae in the union theory. This is not true, as the corollary below shows.

Corollary 35 *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 $i = 1, 2$. Then, if the satisfiability in T_i of formulae of $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ is decidable, 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 because it can lead by iteration to decidability results for more than two theories. Suppose in fact that, in addition to the theories in the corollary, 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 $Res(\mathcal{L}^{\Sigma_1\cup\Sigma_2\cup\Sigma_3}, \Sigma)$ is also decidable.

Proving Corollary 35 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 $iso_\rho\xi \wedge dif_\xi$ 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.[23] Now, if the input φ is already of the form $\varphi \wedge res^\Sigma(\tilde{v})$ 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 $(\varphi_i'\rho \wedge iso_\rho)\xi \wedge dif_\xi$ as before with the only differences that φ_i' is now $\varphi_i \wedge res^\Sigma(\tilde{v})$, 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 $(\varphi_i'\rho \wedge iso_\rho)\xi \wedge dif_\xi$

[23] Considering only shared variables is in a sense an optimization of this more general case.

can be effectively reduced[24] to a logically equivalent formula in $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$, 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 [BS95a] 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 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 [BT01] for a detailed discussion). We adopt the following definition among a number of possible ones.

Definition 36 (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.*

[24] Exploiting the associativity, commutativity, and idempotency of $\wedge$.

- *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.*[25]

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 37 *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 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. □

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 try 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.

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

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 [Opp80] 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 38 (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 [BT97], 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.[26]

Proposition 39 *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 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. [Hod93]). 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. □

[26] 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 $\bot$ and $\neg\top$, at most one of which is positive.

Some specific examples of stably-infinite theories interesting in program verification can be found in [Opp80].

One consequence of Definition 38 is that stably-infinite theories admit infinite models and so, by the Upward and Downward Löwenheim-Skolem theorems [Hod93], admit models of any infinite cardinality[27]. 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 40 *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 24. 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 25 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 28.

Theorem 41 *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 identification and that $Res(Qff^{\Omega}, \emptyset)$ coincides with Qff^{Ω} for any signature Ω. See [Rin96b] or [TH96] for more details.

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

6.2 Non-disjoint Signatures

Let us now consider 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 24). Then, we consider what extra conditions must be true for them to be totally N-O-combinable (that is, satisfying Condition 25 as well).

In the previous subsection, to provide sufficient conditions for the N-O-combinability of two theories with disjoint signatures we looked for restrictions that would guarantee the existence of fusible models. In that case, it was enough to guarantee the existence of two models with the same cardinality. When the theories' signatures have a non-empty common part 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 later and use it to define a general class of N-O-combinable theories. But first, we will go over the definition and the properties of free structures that we will need for that.

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 42 (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.[28]

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).

[28] 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 $Mod(T)$.

A structure free over an empty basis is called *initial*.[29] A structure of signature Σ free in the class of all Σ-structures is called *absolutely free*.

We will often use the following characterization of freeness.

Proposition 43 (Characterization of Free Structure [Hod93]) *Let $\mathbf{K}$ be any 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 43.

Corollary 44 *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 [Hod93]. Free structures in a collapse-free class, however, have unique bases.

Proposition 45 *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 $\mathcal{A}$, there is a non-variable Σ-term $t(\tilde{v})$ and a sequence $\tilde{x}$ in X with no repetitions such that $y = t^{\mathcal{A}}[\tilde{x}]$.[30]

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

Free structures have a close connection to varieties. In fact, every non-trivial Σ-variety contains structures free in it. Furthermore, every free Σ-structure is free in some Σ-variety [Hod93], 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.

[29] This definition is equivalent to a 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 $\mathcal{B}$.

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

Proposition 46 *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 43, 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. □

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 42.

Lemma 47 *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 48 *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.*

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 47 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 48 that $\mathcal{A}$ and $\mathcal{B}$ are Σ-fusible.

6.2.2 Stably Σ-free Theories

We can use Proposition 48 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 49 (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 50 *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 $\kappa \geq Card(\Omega)$.

($\Rightarrow$) 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, $\forall x \forall y.\ x \equiv y$ is clearly not a consequence of T given the assumption that T is stably-infinite. It follows 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 49, 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 51 *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 39.*

Now let $\Sigma := \{a, d\}$ and observe that the atomic Σ-theory of T is empty. Since a equals d in every model of T that satisfies the quantifier-free formula $c \equiv d$, the model's reduct to Σ is certainly not free in the Σ-variety of T. 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[31] such that $\Sigma := \Sigma_1 \cap \Sigma_2$ is finite.

Lemma 52 *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 49, 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 48 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 44 that $H_0 \models \tilde{\forall}\, \varphi$. Recalling that H_0 is the atomic Σ-theory of T_1, we can conclude that $\tilde{\forall}\, \varphi \in H_0$. □

Theorem 53 *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 24. 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 $i = 1, 2$.

[31] 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.

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 48.

(2) It suffices to show that $\langle \varphi_1, \varphi_2 \rangle$ satisfies Condition 25. 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 52, and H_0 is collapse-free by assumption, we have by Proposition 45 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. □

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 much nicer instead, 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 54 *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 $TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$*[32].

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 48, $\mathcal{A}_1$ 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

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

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 $\mathcal{F}$.

We have already seen that $\mathcal{F} \in Mod(T_1 \cup T_2)$ 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 52, the atomic Σ-theory of $T_1 \cup T_2$ coincides with H_0. □

The above result is not sufficient for our needs given that, in general, the class $TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$ is strictly included in $Res(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$. One might argue, however, that if we limit ourselves to totally Σ-restricted 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 $\varphi_1 \wedge \varphi_2 \wedge res^{\Sigma}(\tilde{u})$ 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, to 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 $\varphi(\tilde{v}) \wedge res^{\Sigma}(\tilde{v})$, after purification we may end up with partially Σ-restricted formulae of the form $\varphi_1 \wedge \varphi_2 \wedge res^{\Sigma}(\tilde{v})$ 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

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

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

We show in Section 8.4 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

In the rest of the paper, we show the range of applications of our combination method by providing ways to identify N-O-combinable theories with decidable satisfiability problem. A major class of N-O-combinable theories will involve theories sharing *constructor* symbols. We define in this section what we mean by constructors, and prove some of their properties which we use later in the paper. The notion of constructors presented here was introduced in [TR98] after a similar one in [DKR94], and further refined with Franz Baader in [BT98] in the context of equational theories. In the following, we provide a unified treatment of the results in [TR98] and [BT98] for the case of arbitrary first-order theories.[33]

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., [EM85, EM90]), 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:

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

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 $\mathcal{A}$, for any $x \in X$, cannot be generated by $\mathsf{0}^{\mathcal{A}}$ and $\mathsf{s}^{\mathcal{A}}$ alone. Therefore, X is not a set of generators for $\mathcal{A}^\Sigma$. The interesting

[33] Recently, the notion has been extended even further, again in the context of equational theories. See [BT01] for more details.

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

thing about this example is that $\mathcal{A}^\Sigma$ is indeed a free algebra. And while it is 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 by $\mathsf{0}^\mathcal{A}$ and $\mathsf{s}^\mathcal{A}$ 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.

We will start with a general syntactic definition of constructors and then provide a semantical characterization of it in terms of free structures with an infinite basis. Then, we will show how our definition extends a similar one from term rewriting.

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 $s(\tilde{r})$ where $\tilde{r}$ is a discrete tuple collecting the terms of G that replace the variables of s. Notice 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 $T(\Omega, V)$:

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

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 55 (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 56 *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,[35] the Σ-reduct of $\mathcal{A}$ is free in T^{Σ} with a basis determined by $G_T(\Sigma, V)$. For this purpose, we will use the following properties of $G_T(\Sigma, V)$.

Lemma 57 *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

[35] An important class of theories admitting free models with infinite bases is the class of non-trivial universal Horn theories.

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 $V \not\subseteq G$. □

Proposition 58 *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.*[36] *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:*

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

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 55(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.

The above entails that $\mathcal{A}^\Sigma$ satisfies the first condition of Proposition 43. 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 $s_1(\tilde{v}_1), \ldots, s_n(\tilde{v}_n) \in T(\Sigma, V)$, relation symbol $P \in \Sigma^{\mathrm{P}} \cup \{\equiv\}$, and injection β of $V_0 := \mathcal{V}ar(P(s_1(\tilde{v}_1), \ldots, s_n(\tilde{v}_n)))$ into Y such that

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

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 $(\mathcal{A}, \alpha) \models P(s_1(\tilde{r}_1), \ldots, s_2(\tilde{r}_n))$. Since $\mathcal{A}$ is free in T over X and α is injective as well we can conclude by Proposition 43(2) that $T \models \tilde{\forall}\, P(s_1(\tilde{r}_1), \ldots, s_2(\tilde{r}_n))$.

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 $P(s_1(\tilde{v}_1), \ldots, s_n(\tilde{v}_n))$ can be obtained from $P(s_1(\tilde{r}_1), \ldots, s_2(\tilde{r}_n))$ by abstracting the terms in $\tilde{r}_1, \ldots, \tilde{r}_n$ so that two terms are abstracted by the same variable iff they are equivalent in T. But then, by Definition 55(3) we can conclude that $T \models \tilde{\forall}\, P(s_1(\tilde{v}_1), \ldots, s_n(\tilde{v}_n))$. Since $\tilde{\forall}\, P(s_1(\tilde{v}_1), \ldots, s_n(\tilde{v}_n))$ is a Σ-sentence, this is the same as saying that $T^\Sigma \models \tilde{\forall}\, P(s_1(\tilde{v}_1), \ldots, s_n(\tilde{v}_n))$. □

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

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 59 *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 58, $\mathcal{A}^\Sigma$ is free in T^Σ. By Lemma 57(3), the fact that $V \subseteq G$ (cf. Condition (1) of Definition 55) 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 43(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 $\beta(v) = z = [\![s]\!]^{\mathcal{A}^\Sigma}_\beta$.[37] 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 55 follows then directly from Lemma 57(2). The second condition follows by Proposition 43(2) and Corollary 44, given that $\mathcal{A}$ is free in T and Σ-generated by $Y = Z$. Similarly, the third condition follows from Proposition 43(2). □

We can now give an alternative formulation of Theorem 59 by means of the following corollary.

[37]Note that v may be an element of $\mathcal{V}ar(s)$.

Corollary 60 *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 $Is(\mathcal{A}^\Sigma)$.*[38]

Proof. ($1 \Rightarrow 2$) By Theorem 59, $\mathcal{A}^\Sigma$ is free in the collapse-free theory T^Σ. By Proposition 45, 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 59, 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 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, $Is(\mathcal{A}^\Sigma)$ is empty. But then, we can argue as in the proof of Theorem 59 that $\mathcal{A}$ is generated by the empty set, which is impossible as $\mathcal{A}$ is free over an infinite set by assumption. □

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 61 *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 $Is(\mathcal{A}^\Sigma)$ exhibits in turn the following closure property.

Lemma 62 *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*

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

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

Proof. Let $r(\tilde{v}) \in G := G_T(\Sigma, V)$ and β a valuation of $\tilde{v}$ into $Is(\mathcal{A}^\Sigma)$. We have seen that $X \subseteq Is(\mathcal{A}^\Sigma) = \{[\![r]\!]^{\mathcal{A}}_\alpha \mid r \in G\}$ 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. □

[38] Recall the $Is(\mathcal{A}^\Sigma)$ is the set of all the isolated individuals of $\mathcal{A}^\Sigma$ (cf. Definition 3).

7.2 Normal Forms

Condition 2 of Definition 55 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*.[39] 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.

We point out that, according to our definition it is not necessarily the case that all the variables occurring the normal form of a term also occur in the term itself. We can guarantee that under the additional assumption that Σ contains a constant symbol.[40]

Proposition 63 *Let Σ be a set of constructors for a theory T of signature Ω. If Σ contains a constant symbol, every $t \in T(\Omega, V)$ has a normal form t' with $\mathcal{V}ar(t') \subseteq \mathcal{V}ar(t)$.*

Proof. First, let us say that a variable v occurring in a term s is *extra for s* (in T) if replacing v in s by a fresh variable produces a term that is equivalent to s in T. Now let $G := G_T(\Sigma, V)$ and $t \in T(\Omega, V)$. Since Σ is a set of constructors for T, there is a term $s(v_1, \dots, v_n) \in T(\Sigma, V)$ and a term $r_1, \dots, r_n \in G$ such that $t =_T s(r_1, \dots, r_n)$.

With no loss of generality we can assume that $r_1, \dots, r_n$ are all inequivalent in T—otherwise we can identify the equivalent terms of $\{r_1, \dots, r_n\}$ and the corresponding variables of $\{v_1, \dots, v_n\}$. Also with no loss of generality we can assume that none of the elements of $v_1, \dots, v_n$ is extra for s. Otherwise we consider the term $s(v_1, \dots, v_n)\sigma\theta$ where σ is the substitution mapping each extra variable of s to a constant symbol of Σ and $\theta := \{v_1 \leftarrow r_1, \dots, v_n \leftarrow r_n\}$. It is easy to see that this term is in $T(\Sigma, G)$ and is equivalent to t in T; therefore it too is a normal form of t.

We now show that each variable of $s(r_1, \dots, r_n)$ not occurring in t, if any, can be "removed" without loss of generality. Suppose that the term r_i for some $i \in \{1, \dots, n\}$ contains (or is) a variable v not in t, and let $\sigma := \{v \leftarrow v'\}$ where v' is a fresh variable. From the equivalence $s(r_1, \dots, r_n) =_T t$, we can conclude that v is extra in $s(r_1, \dots, r_n)$, from which it follows that $s(r_1, \dots, r_i, \dots r_n) =_T s(r_1, \dots, r_i\sigma, \dots r_n)$.

We claim that $r_i\sigma$ is equivalent in T to a term in $(r_1, \dots r_n)$. Assume the contrary. Then, since $r_i\sigma$ is an element of G and $r_1, \dots, r_n$ are all inequivalent in T, we can conclude from Definition 55(3) that $s(v_1, \dots, v_i, \dots v_n) =_T s(v_1, \dots, u, \dots v_n)$ where u is distinct from all $v_1, \dots v_n$. But then v_i is extra for s, against the assumption.

Now let r_j be the term of $(r_1, \dots, r_n)$ equivalent to $r_i\sigma$. If $i \neq j$, consider the term $s(v_1, \dots, v_n)\{v_i \leftarrow v_j\}\{v_1 \leftarrow r_1, \dots, v_n \leftarrow r_n\}$. This term is a normal form of t having less occurrences than $s(r_1, \dots, r_n)$ of the variable v which did not occur in t. If $i = j$, then from $r_i =_T r_i\sigma$ we can conclude that v is extra for r_i and so $r_i =_T r_i\theta$

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

[40] A similar result is also shown in [FG01].

with $\theta := \{v \leftarrow c\}$ for some constant c in Σ. Since G is closed under equivalence in T, this means that $s(r_1, \ldots, r_i\theta, \ldots r_n)$ is also a normal form of t. This term too has less occurrences of v than $s(r_1, \ldots, r_n)$. Repeating the whole process on the new normal form eventually produces a normal form of t all of whose variables are in $\mathcal{V}ar(t)$. □

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

Definition 64 (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.

From the proof of Proposition 63 it is not hard to see that if Σ has a constant symbol, normal forms are computable for Σ and T and term equivalence in T is decidable, then normal forms are computable for Σ and T by a function NF_T^Σ such that $\mathcal{V}ar(\mathrm{NF}_T^\Sigma(t)) \subseteq \mathcal{V}ar(t)$ for all terms t.

Although we will not needed it here, we point out an important consequence of Definition 64: if normal forms are computable for Σ and T, it is always possible to tell whether a term is in normal form or not.

Proposition 65 *Let Σ be a set of constructors for a theory T of signature Ω. If normal forms are computable for Σ and T, the property of being T-reducible is decidable for the terms in $T(\Omega, V)$.*

Proof. Every $t \in T(\Omega, V)$ can be seen as having the form $s(\tilde{r})$ where s is a Σ-term and $\tilde{r}$ are terms with top symbols not in Σ. From the definition of normal form it is immediate that $s(\tilde{r})$ is in normal form exactly when every component of $\tilde{r}$ is in $G := G_T(\Sigma, V)$. But being a member of G is a decidable property of Ω-terms: to test whether any $r \in T(\Omega, V)$ is in G, it is enough to compute $\mathrm{NF}_T^\Sigma(r)$ and look at its top symbol. In fact,

$$r \in G \quad \text{iff} \quad \mathrm{NF}_T^\Sigma(r)(\epsilon) \notin \Sigma.$$

To see that first notice that, by the definition of G, if $\mathrm{NF}_T^\Sigma(r)$ starts with a Σ-symbol then $r \notin G$. Now, if $\mathrm{NF}_T^\Sigma(r)$ does not start with a Σ-symbol, since it is a term in $T(\Sigma, G)$ it must be an element of G, r' say. But then, by definition of NF_T^Σ, r and r' are equivalent in T, which entails that $r \in G$ by Lemma 57(1). □

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.[41]

- The signature $\Sigma := \{f\}$ is not a set of constructors for the theory T axiomatized by $\{\forall x.\ x \equiv f(g(x))\}$ because it does not satisfy Definition 55(1), as one can easily show.

- The signature $\{f\}$ is not a set of constructors for the theory T axiomatized by $\{\forall x.\ g(x) \equiv f(g(x))\}$ because it does not satisfy Definition 55(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 axiomatized by $\{\forall x.\ f(g(x)) \equiv f(f(g(x)))\}$. 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 55 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 $\{P\}$ of $\Omega := \{P, g\}$ is not a set of constructors for the theory axiomatized by $\{\forall x.\ P(g(x))\}$.

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

Example 66 *Consider the signature* $\Sigma_{66} := \{\mathsf{0}, \mathsf{s}, +\}$ *and the theory* E_{66} *axiomatized by the sentences:*

$\forall x, y, z.$	$x + (y + z)$	$\equiv$	$(x + y) + z$
$\forall x, y.$	$x + y$	$\equiv$	$y + x$
$\forall x, y.$	$x + \mathsf{s}(y)$	$\equiv$	$\mathsf{s}(x + y)$
$\forall x.$	$x + \mathsf{0}$	$\equiv$	x

The signature $\Sigma := \{\mathsf{0}, \mathsf{s}\}$ *is a set of constructors for* E_{66} *in the sense of Definition 55 (see [BT98] for a proof). In particular,* $G_T(\Sigma, V)$ *is the set of all terms made of zero or more additions of variables, and each normal forms 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.

[41] In the (counter-)examples below, x, y, z are variables, numbers and identifiers starting with a capital letter are function symbols, identifiers starting with a capital letter are relation symbols.

Example 67 *Consider the signature* $\Sigma_{67} := \{\mathsf{0}, \mathsf{s}, +, \mathsf{Even}\}$ *and the theory* T_{67} *axiomatized by* E_{66} *above plus the sentences:*

$$\begin{array}{|lll|}\hline & \mathsf{Even}(\mathsf{0}) & \\ \forall x. & \mathsf{Even}(x) \;\Rightarrow & \mathsf{Even}(\mathsf{s}(\mathsf{s}(x))) \\ \hline\end{array}$$

It is not difficult to show that the signature $\Sigma := \{\mathsf{0}, \mathsf{s}, \mathsf{Even}\}$ *is a set of constructors for* T_{67}. *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, according to Definition 55(3), since* $x + x$ *is in* $G_T(\Sigma, V)$, $\forall y.\ \mathsf{Even}(y)$ *should also be entailed by the theory, which is not the case.*

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

Example 68 *Consider the signature* $\Sigma_{68} := \{\mathsf{0}, \mathsf{1}, \mathsf{rev}, \cdot\}$ *and the theory* E_{68} *axiomatized by the sentences:*

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

We show in Section 7.4 that the signature $\Sigma := \{\mathsf{0}, \mathsf{1}, \cdot\}$ *is a set of constructors for* E_{68}. *The set* $G_T(\Sigma, V)$ *is the equivalence closure in* E_{68} *of the set* $V \cup \{\mathsf{rev}(v) \mid v \in V\}$. *Moreover, every normal form is a concatenation (with* $\cdot$*) of terms in* $\{\mathsf{0}, \mathsf{1}\} \cup G_T(\Sigma, V)$. *In this case too* $G_T(\Sigma, V)$ *is closed under instantiation into itself.*

Example 69 *Consider the signature* $\Sigma_{69} := \{\mathsf{0}, \mathsf{1}, \mathsf{rev}, \cdot, \mathsf{Prefix}\}$ *and the theory* T_{69} *axiomatized by* E_{68} *plus the sentences:*

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

Again, it is not difficult to see that the signature $\Sigma := \{\mathsf{0}, \mathsf{1}, \cdot, \mathsf{Prefix}\}$ *is a set of constructors for* T_{69}.

7.4 Constructors in Term Rewriting

In Term Rewriting, a function symbol in the signature of a given term rewriting system (TRS for short) R is a constructor for R if it does not occur at the top of the left-hand side of any rule in R. Our constructors are a natural generalization of this notion. We show in the following that the set of constructors of any confluent and (weakly) normalizing TRS R is also a set of constructors in the sense of Definition 55 for the equational theory induced by R. We will not 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 [JK86].

We will assume that the reader is familiar with Term Rewriting and so we will introduce only the terminology and the notation we need to prove our claims. Comprehensive introductions to the field can be found in [BN98, DJ90, Wec92], 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 rather than 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 will call TRS-*constructors* here, are defined as follows.

Definition 70 (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 in 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 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 71 *We denote by $S = (R, E_0)$ the TRS R* modulo *E_0, that is, the TRS whose rewrite relation $\to_S$ over $T(\Omega, V)$ is defined as follows. For all $s, t \in T(\Omega, V)$, $s \to_S t$ if there exists a position p, a substitution σ, and a rule $l \to r \in R$ 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. $\to_S$) *of an Ω-term t iff t' is irreducible by $\to_S$ 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 $\to_S$. Note that, when the theory E_0 is empty, $\to_S$ 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 72 *E_0 is the theory presented by the axiom:*

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

and R is the TRS consisting of the rules:

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

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

Definition 73 (Completeness) *The TRS $S = (R, E_0)$ is* semi-complete for 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 $E_0 \cup \{\tilde{\forall}\, l \equiv r \mid l \to r \in R\}$;*
2. *the relation $\to_S$ is* normalizing, *i.e., every Ω-term t has a normal form w.r.t $\to_S$;*
3. *the relation $\to_S$ is* confluent modulo E_0, *i.e., for all Ω-terms t, t_1, t_2 such that $t_1 \overset{*}{\leftarrow}_S t \overset{*}{\to}_S t_2$, 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 $\to_S$ is* terminating, *i.e., there is no infinite sequence $(t_0, t_1, t_2, \dots)$ such that $t_0 \to_S t_1 \to_S t_2 \to_S \cdots$.*

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. $\to_S$, and for all $s, t \in T(\Omega, V)$ and respective normal forms s', t',

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

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 and denote any of 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 and word problems in E_0 are decidable,[42] then the normal form $t{\downarrow}_S$ of every term t is computable; as a consequence, the word problem in E is also decidable. As a matter of fact, using standard results in term rewriting, it can be shown that the word problem in E is decidable even if S is just semi-complete for E.

To prove that TRS-constructors are constructors for E in the sense of Definition 55, 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 [SS89, Nip91, Rin96a, BT97]. 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.

[42] 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 $t_1\sigma =_{E_0} t_2$.

Lemma 74 *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. *For all $\sigma \in \mathrm{SUB}(V)$, $i \in \{1, 2\}$, and i-pure non-variable terms s, t such that*
   - *$(v\sigma)(\epsilon) \notin \Sigma_i$ for all $v \in \mathcal{V}ar(s \equiv t)$,*
   - *$u\sigma \neq_{E_1 \cup E_2} v\sigma$ for all distinct $u, v \in \mathcal{V}ar(s \equiv t)$,*

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

The first property of S that we can show with the lemma above is the following.

Proposition 75 *If $S = (R, E_0)$ is semi-complete for E and Σ is a set of* TRS-*constructors for R, then*
$$f(t_1, \ldots, t_n)\downarrow_S \quad =_{E_0} \quad f(t_1\downarrow_S, \ldots, t_n\downarrow_S)$$
for all n-ary $f \in \Sigma$ and $t_1, \ldots, t_n \in T(\Omega, V)$.

Proof. Consider the term $f(t_1, \ldots, t_n)$ as above. Since $f(t_1, \ldots, t_n)$ and $f(t_1\downarrow_S, \ldots, t_n\downarrow_S)$ are obviously equivalent in E, we have by the observation after Definition 73 that
$$f(t_1, \ldots, t_n)\downarrow_S \quad =_{E_0} \quad f(t_1\downarrow_S, \ldots, t_n\downarrow_S)\downarrow_S.$$
To prove the claim then it is enough to show that the term $f(t_1\downarrow_S, \ldots, t_n\downarrow_S)$ is irreducible by $\rightarrow_S$. Assume the contrary. Then any rule $l \rightarrow r \in R$ that applies to $t := f(t_1\downarrow_S, \ldots, t_n\downarrow_S)$ must apply at the top of t, which means that $t =_{E_0} l\sigma$ for some substitution σ or, more precisely, that
$$f(t_1\downarrow_S, \ldots, t_n\downarrow_S) \quad =_{E_0^\Omega} \quad l\sigma.$$
It follows by point 2 of Lemma 74 that $(l\sigma)(\epsilon)$ is in Σ as well. But this is impossible because, since Σ is a set of TRS-constructors for R, $l(\epsilon) \notin \Sigma$. □

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

Lemma 76 *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 $t \in T(\Sigma, V)$.*

Proof. We prove the claim by term induction.

(Base case) Let $t = v \in V$ and assume ad absurdum that $v{\downarrow_S} \neq v$. Then, by definition of $v{\downarrow_S}$ and $\rightarrow_S$, there must be an Ω-term t other that v such that $v =_{E_0^\Omega} t$. But this contradicts the fact that E_0^Ω is collapse-free by Lemma 74(1) for being the union of two collapse-free theories.

(Induction Step) Let $t \in T(\Sigma, V) \setminus V$. Then t has the form $f(t_1, \ldots, t_n)$ where, by induction hypothesis, each t_i is irreducible by $\rightarrow_S$. Exactly as in Lemma 75 we can then show that $f(t_1, \ldots, t_n)$ as well is irreducible by $\rightarrow_S$. □

An easily provable consequence of the lemma above is that, under the same assumptions of the lemma, 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.[43]

Lemma 77 *Assume that $S = (R, E_0)$ is semi-complete for E and Σ is a set of* TRS-*constructors for R. Then,*

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

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 73 and Proposition 75, 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 74. If $r{\downarrow_S}(\epsilon)$ is a variable, the equivalence contradicts the fact that E_0^Ω is collapse free by Lemma 74(1). □

Together with Proposition 75, Lemma 77 has the following consequence.

Lemma 78 *Let $G := G_E(\Sigma, V)$ and 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 73(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 $f(\tilde{t}){\downarrow_S} =_{E_0} f(\tilde{t}{\downarrow_S})$ by Proposition 75. But this entails that $r =_{E_0^\Omega} f(\tilde{t}{\downarrow_S})$, which is impossible by Lemma 74(2). □

[43]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.

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

Proposition 79 *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 55 are satisfied. Let $G := G_E(\Sigma, V)$.

(1) Let $v \in V$. Since $v = v{\downarrow_S}$ by Lemma 76, we can immediately conclude by Lemma 77 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 78 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

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

Recalling that s_1 and s_2 are Σ-terms, we can show by a simple inductive argument based on Proposition 75 that

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

Assuming that E-equivalent terms in $\tilde{r}_1, \tilde{r}_2$ have the same normal w.r.t. $\to_S$,[44] 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 74. By that lemma, it then follows that $s_1(\tilde{v}_1) =_{E_0} s_2(\tilde{v}_2)$ and so $s_1(\tilde{v}_1) =_E s_2(\tilde{v}_2)$. □

We would like to stress that, although the preconditions in Proposition 79 entail that Σ is a set of constructors for E, they do not entail that normal forms in the sense of Definition 64 are computable. A sufficient condition for the computability of normal forms, under the assumptions of Proposition 79, is that E_0-matching with free constants is decidable. We will prove that in Subsection 9.2.

Finally, we can produce a result like the above for "conventional" TRSs, i.e. TRSs not modulo some equational theory, again by observing that such systems are TRSs modulo the empty equational theory.

Corollary 80 *Let R be a TRS over $T(\Omega, V)$. If $\to_R$ 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.*

[44]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 $\tilde{r}_1{\downarrow_S}, \tilde{r}_2{\downarrow_S}$.

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 68 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 72. Clearly, E_0 is collapse-free, $\rightarrow_R$ is terminating (therefore, normalizing) and $\Sigma := \{0, 1, \cdot\}$ 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 79 that $\Sigma := \{0, 1, \cdot\}$ is a set of constructors for E_{68}.

8 Identifying Σ-stable Theories

In this section, we give some examples of classes of stably Σ-free theories and show which theories within these classes are N-O-combinable. We believe more classes can and should be identified in order to better assess the practical significance of our combination method in the case of component theories with non-disjoint signatures. For now, we can look at the results below and their proofs as a set of general guidelines on how to apply Theorem 53 in practice.

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

For the rest of this section, we fix two countable signatures Σ_1, Σ_2 such that $\Sigma := \Sigma_1 \cap \Sigma_2$ is finite, and two theories T_1, T_2 of respective signature Σ_1, Σ_2.

8.1 Theories Sharing Constants

We start with the simple case of theories sharing just constant symbols. Assume that for $i = 1, 2$,

- T_i is stably-infinite over Qff^{Σ_i};
- Σ contains only constant symbols;
- for all $k_1, k_2 \in \Sigma$, either $T_i \models (k_1 \equiv k_2)$ or $T_i \models (k_1 \not\equiv k_2)$.

Then, we can show the following.

Lemma 81 *T_i is stably Σ-free over $\mathit{Res}(\mathit{Qff}^{\Sigma_i}, \Sigma)$ for $i = 1, 2$.*

Proof. It is enough to notice the following: first, since Σ is just a set of constants, $Res(Qff^{\Sigma_i}, \Sigma) \subseteq Qff^{\Sigma_i}$ and so, by the stable-infiniteness of T_i, every Σ-restricted formula satisfiable in T_i is satisfiable in a countably infinite model of T_i; second, for any countably infinite model $\mathcal{A}_i$ of T_i, $\mathcal{A}_i{}^{\Sigma}$ is free in the Σ-variety of T_i over a countably infinite basis (because any Σ-equation valid in $\mathcal{A}_i$ is also valid in T_i). □

Proposition 82 *If $T_1 \models (k_1 \equiv k_2)$ iff $T_2 \models (k_1 \equiv k_2)$ for all $k_1, k_2 \in \Sigma$, then T_1 and T_2 are totally N-O-combinable over Qff.*

Proof. By Lemma 81, T_i is stably Σ-free over $Res(Qff^{\Sigma_i}, \Sigma)$ for $i = 1, 2$. It is immediate that T_1 and T_2 have the same atomic Σ-theory. It follows by Theorem 53(1) then that T_1 and T_2 are partially N-O-combinable over *Qff*. To see that they are totally N-O-combinable, it is enough to notice that, since Σ is a set of constant symbols, *every* model of $T_1 \cup T_2$ is Σ-generated by its Σ-individuals, which satisfies Condition 25 directly. □

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

Let us simplify our initial assumptions by requiring that no shared constants are actually equivalent in T_i for $i = 1, 2$. In practice, such a requirement causes no loss of generality, as we can always identify two equivalent constants and remove one of them from the signature. Then, we obtain the following decidability result.

Proposition 83 *Let T_1, T_2 be such that for $i = 1, 2$,*

- *T_i is stably-infinite over Qff^{Σ_i};*
- *Σ contains only constant symbols;*
- *for all $k_1, k_2 \in \Sigma$, $T_i \models (k_1 \not\equiv k_2)$.*

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

Proof. Since T_1 and T_2 are N-O-combinable over *Qff* by Proposition 82, our combination method is applicable in a sound and complete way. This means that the method yields a semi-decision procedure for the satisfiability of formulae in $Qff^{\Sigma_1 \cup \Sigma_2}$. To see that it actually yields a decision procedure, simply observe that the non-determinism in the instantiation step is bounded in this case because the set of shared terms is finite. □

8.2 Theories Sharing the Finite Trees

In this subsection we show that theories obtained as an extension of the *theory of finite trees* are N-O-combinable under certain conditions. Finite trees are a major data structure in Computer Science and Symbolic Computation. The domain of finite trees, which is essentially a term algebra, was first axiomatized by Mal'cev (see [Mal71]). We present this axiomatization below and call it the theory of finite trees.

Definition 84 (Finite Trees) *The theory of the finite trees, over some signature* Σ, *is the universal theory presented by the axioms below.*

- $\tilde{\forall}\,(f(\tilde{u}) \equiv f(\tilde{v}) \Rightarrow \tilde{u} \equiv \tilde{v})$ *for every* $f \in \Sigma^{\mathrm{F}}$
- $\tilde{\forall}\, f(\tilde{u}) \not\equiv g(\tilde{v})$ *for every* $f, g \in \Sigma^{\mathrm{F}}$, $f \neq g$
- $\tilde{\forall}\, v \not\equiv t(\tilde{v})$ *for every* $t(\tilde{v}) \in T(\Sigma, V) \setminus V$ *and* $v \in \tilde{v}$

To facilitate the exposition, in the following we will identify the theory above with its deductive closure and denote it by FT^{Σ}—which is consistent with the way we denote the Σ-restriction of a theory (cf. Section 2).

The models of FT^{Σ} can be give an algebraic characterization; they are all what Mal'cev called *locally absolutely free algebras*. We will use some of the properties of such algebras later. We introduce and prove these properties below, but in the context of a more general class: the class of locally free structures.

Definition 85 (Locally Free Structure) *A structure* $\mathcal{A}$ *of signature* Σ *is* locally free *in a class* $\mathbf{K}$ *of* Σ*-structures if every finitely-generated substructure of* $\mathcal{A}$ *is free in* $\mathbf{K}$ *(over some finite set).*

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

Proposition 86 *If a locally free structure in a collapse-free class* $\mathbf{K}$ *admits a non-redundant set* X *of generators, then it is free over* X *in* $\mathbf{K}$.

Proof. Let $\mathcal{B}$ be a Σ-structure with a non-redundant set of generators X and assume that $\mathcal{B}$ is locally free in some class $\mathbf{K}$ of Σ-structures. Let $\varphi(\tilde{v})$ be an atomic Σ-formula and $\tilde{x}$ a sequence of distinct elements of X such that $\mathcal{B} \models \varphi[\tilde{x}]$. By Proposition 43, it is enough to show that $\mathbf{K} \models \tilde{\forall}\, \varphi$.

Let $\mathcal{A} := \langle \tilde{x} \rangle_{\mathcal{B}}$ and $I := \mathit{Is}(\mathcal{A}^{\Sigma})$. Notice that $\mathcal{A}$ is free in $\mathbf{K}$, for being a finitely generated substructure of a locally free structure in $\mathbf{K}$, and that, by Proposition 45, I is the only basis for $\mathcal{A}$. By construction of $\mathcal{A}$ and Lemma 1, $\tilde{x}$ is a non-redundant set of generators for $\mathcal{A}$. From what we observed earlier, I as well is a non-redundant set of generators for $\mathcal{A}$. It follows immediately, as $I \subseteq \tilde{x}$, that $\tilde{x} = I$. Now notice that $\mathcal{A} \models \varphi[\tilde{x}]$ as well because $\mathcal{A} \subseteq \mathcal{B}$ and φ is atomic. Then, by Proposition 43 applied to $\mathcal{A}$, we obtain that $\mathbf{K} \models \tilde{\forall}\, \varphi$. □

It possible to show that every substructure of an absolutely free structure is absolutely free. This immediately entails that absolutely free structures are also *locally* absolutely free. The converse, however, is not true. In fact, consider the Σ-structure $\mathcal{Z}$ of the integer numbers with signature $\Sigma := \{s\}$ for the successor function. The structure $\mathcal{Z}$ cannot be free because it does not admit a non-redundant set of generators. However, it is easy to see that every finitely generated substructure of $\mathcal{Z}$ is isomorphic to the Σ-structure of the natural numbers, which is absolutely free. Nonetheless, by Lemma 86, we claim the following special case.

Corollary 87 *For any signature Σ, the class of locally absolutely free Σ-structures with a* non-redundant *set of generators coincides with the class of absolutely free Σ-structures.*

For each signatures Σ not containing predicate symbols, the class of locally absolutely free Σ-structures coincides with $Mod(\mathrm{FT}^{\Sigma})$.

Proposition 88 *A Σ-algebra $\mathcal{A}$ is locally absolutely free iff $\mathcal{A} \in Mod(\mathrm{FT}^{\Sigma})$.*

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

We can now move to the the combination of theories extending the theory of finite trees, and show under what conditions they are N-O-combinable.

The extended theories will be *universal*, that is, axiomatized by a set of closed universal formulae, where a universal formula is a formula in Prenex Normal Form whose (possibly empty) quantifier prefix contains only universal quantifiers. We will appeal to the following two properties of universal theories (see [Hod93] among others).

Lemma 89 *Let $\mathcal{B}$ be a Σ-structure and $\varphi(\tilde{v})$ a universal Σ-formula such that $\mathcal{B} \models \varphi[\tilde{a}]$ for some $\tilde{a}$ in $\mathcal{B}$. Then, $\varphi(\tilde{v})$ is satisfiable in every substructure $\mathcal{A}$ of $\mathcal{B}$ whose universe includes $\tilde{a}$.*

Lemma 90 *For every universal Σ-theory T, the class $Mod^{\Sigma}(T)$ is closed under the formation of substructures.*

In addition, we will use some general properties of what we call *Σ-independent sets.*

Definition 91 *Let $\mathcal{B}$ be a structure and $\Sigma \subseteq \Sigma_B$. A set $X \subseteq B$ is* Σ-independent *in* $\mathcal{B}$ *iff X is a non-redundant set of generators for the substructure $\langle X\rangle_{\mathcal{B}^{\Sigma}}$ of $\mathcal{B}^{\Sigma}$ generated by X.*

To simplify the enunciation of the next result let us say that a *set* A_X is Σ-generated by a set X included in the universe of a structure $\mathcal{A}$, with $\Sigma \subseteq \Sigma_A$, if A_X is contained in the universe of $\langle X\rangle_{\mathcal{A}^{\Sigma}}$.

Lemma 92 *Let $\mathcal{B}$ be an uncountable structure, Σ a countable subsignature of Σ_B, and A a finite subset of B. Then, there is a countably infinite subset of B which is Σ-independent in $\mathcal{B}$ and Σ-generates A in $\mathcal{B}$.*

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

Assume that there is an $x \in X$ such that $\{x\}$ is Σ-generated by $X \setminus \{x\}$. Then, we can show that $\{x\}$ is Σ-generated by some finite subset Y of $X \setminus \{x\}$. By construction of X, there is an $n < \omega$ such that X_n includes $Y \cup \{x\}$. But then, by the above, X_n is not Σ-independent in $\mathcal{B}$, against the assumption. □

For the rest of the subsection we will assume that our theories T_1 and T_2 are such that, for $i = 1, 2$,

- T_i is universal,
- Σ_i contains at least one function symbol of non-zero arity,
- $\Sigma = {\Sigma_i}^{\mathrm{F}}$, and
- $T_i^{\Sigma} = \mathrm{FT}^{\Sigma}$.

The *Clark completion* of a Prolog program, which provides the logical semantics of the program (see, e.g., [Llo87]), is an example of a theory of this sort. Each of the theories above is stably Σ-free over Σ-restricted universal formulae.

Lemma 93 *For $i = 1, 2$, T_i is stably Σ-free over $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ where $\mathcal{L}$ is the class of universal formulae.*

Proof. Because of the assumptions that $T_i^{\Sigma} = \mathrm{FT}^{\Sigma}$ for $i = 1, 2$ and Σ contains at least one function symbol of non-zero arity, all models of T_i are infinite. Moreover, as $\Sigma = {\Sigma_i}^{\mathrm{F}}$, every set of Σ-generators for a model $\mathcal{A}$ of T_i is also a set of generators for $\mathcal{A}$. Now, suppose that $\varphi(\tilde{v}) \in Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ is satisfied in a model $\mathcal{B}$ of T_i by some tuple $\tilde{b}$. Since $\mathcal{B}$ is infinite, we can assume without loss of generality that it is uncountable. By Lemma 92 then, there is a countably-infinite $X \subseteq B$ that is Σ-independent in $\mathcal{B}$ and Σ-generates $\tilde{b}$. Let $\mathcal{A} := \langle X \rangle_{\mathcal{B}}$. By construction, X is a non-redundant set of generators for $\mathcal{A}$ and $\tilde{b}$ is in A. Observing that φ is equivalent to a universal formula, we can conclude by Lemma 89 and Lemma 90 that $\mathcal{A}$ as well is a model of T_i that satisfies φ. Now, $\mathcal{A}^{\Sigma}$ is clearly a model of FT^{Σ}, therefore by Proposition 88 and Corollary 87, $\mathcal{A}^{\Sigma}$ is an absolutely free algebra. Since the Σ-variety of T_i coincides with the Σ-variety of the empty theory, it follows that $\mathcal{A}^{\Sigma}$ is free over X in the Σ-variety of T_i.

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

We are almost ready to show that T_1 and T_2 are N-O-combinable over universal formulae. We need one more general result about theory restrictions. This result, which is not limited to universal theories, is not as trivial as it looks and, in fact, does not hold if the signature Σ below is strictly contained in $\Sigma_1 \cap \Sigma_2$.

Lemma 94 *Let Γ_1 be an Ω_1-theory, Γ_2 an Ω_2-theory, and $\Omega := \Omega_1 \cap \Omega_2$. Then, for all Ω-sentences φ,*

$$(\Gamma_1 \cup \Gamma_2)^\Omega \models \varphi \quad \textit{iff} \quad \Gamma_1^\Omega \cup \Gamma_2^\Omega \models \varphi.$$

Proof. ($\Leftarrow$) Immediate consequence of the obvious fact that $\Gamma_1^\Omega \cup \Gamma_2^\Omega \subseteq (\Gamma_1 \cup \Gamma_2)^\Omega$.

($\Rightarrow$) Let φ be a Ω-sentence and assume that $(\Gamma_1 \cup T_2)^\Omega \models \varphi$ or, equivalently, that the theory $\Gamma_1 \cup (\Gamma_2 \cup \{\neg\varphi\})$ is inconsistent. By the Craig Interpolation Lemma [Hod93], there is a Ω-sentence ψ such that $\Gamma_1 \models \neg\psi$ and $\Gamma_2 \cup \{\neg\varphi\} \models \psi$. By logical reasoning, we also have that $\Gamma_2 \models \neg\psi \Rightarrow \varphi$. Observing that both φ and $\neg\psi$ are Ω-sentences, we can then conclude that $\neg\psi \in \Gamma_1^\Omega$ and $(\neg\psi \Rightarrow \varphi) \in \Gamma_2^\Omega$ from which the claim follows immediately. □

Proposition 95 *T_1 and T_2 are totally N-O-combinable over the class of universal formulae.*

Proof. Let $\mathcal{L}$ be the class of universal formulae. Obviously, both T_1 and T_2 have the same atomic Σ-theory, the empty theory. Since, by the previous lemma, T_i is stably Σ-free over $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$ for $i = 1, 2$, we can conclude by Theorem 53(1) that T_1 and T_2 are partially N-O-combinable over $\mathcal{L}$.

Now, by Lemma 94, it is easy to see that $(T_1 \cup T_2)^\Sigma = \mathrm{FT}^\Sigma$. It follows that $T_1 \cup T_2$ satisfies the same preconditions we have on T_1 and T_2, which means that Lemma 93 applies to $T_1 \cup T_2$ as well. In other words, $T_1 \cup T_2$ is stably Σ-free over the class $Res(\mathcal{L}^{\Sigma_1 \cup \Sigma_2}, \Sigma)$. In particular, it is stably Σ-free over the subclass $\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}$. Observing that the Σ-atomic theory of $T_1 \cup T_2$ is also empty and thus definitely collapse-free, we obtain by Theorem 53(2) that T_1 and T_2 are totally N-O-combinable over $\mathcal{L}$. □

Finally, we obtain the following decidability result.

Proposition 96 *Let T_1, T_2 two theories such that for $i = 1, 2$,*

- *T_i is an axiomatizable universal theory of signature Σ_i,*
- *Σ_i contains at least one function symbol of non-zero arity,*
- *$\Sigma = {\Sigma_i}^{\mathrm{F}}$, and*
- *$T_i^\Sigma = \mathrm{FT}^\Sigma$.*

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

Proof. By Corollary 35, observing that *Qff* is effectively purifiable, T_1 and T_2 are totally N-O-combinable over universal formulae, and every formula in *Qff* is universal. □

8.3 Theories Sharing Decomposition Axioms

In the previous subsection T_1 and T_2 shared *all* their function symbols. We can allow them to have non-shared function symbols if each of them is the complete theory of a free structure in which the extra symbols are "disjoint" from the shared ones. In the following, we will assume that $\Sigma = \Sigma_1 \cap \Sigma_2$ contains only function symbols.

Definition 97 *Let Ω be a signature and $\Sigma \subseteq \Omega$ be such that $\Sigma = \Sigma^{\mathrm{F}}$. The Ω-theory* $\mathrm{DEC}(\Omega, \Sigma)$ *is the universal theory presented by the axioms of* FT^{Σ} *plus the axioms*

$$\tilde{\forall}\, f(\tilde{u}) \not\equiv g(\tilde{v}) \quad \text{for every } f \in \Sigma \text{ and } g \in \Omega \setminus \Sigma.$$

An Ω-theory T is Σ-decomposable *if* $T \models \mathrm{DEC}(\Omega, \Sigma)$. *A Ω-structure $\mathcal{A}$ is* Σ-decomposable *if* $\mathcal{A} \in \mathit{Mod}(\mathrm{DEC}(\Omega, \Sigma))$.

Example 98 *Consider the following Σ_i-theories:*

$$\begin{aligned} E_1 &:= \{\forall x.\ (-x) * (-x) \equiv x * x\} \\ E_2 &:= \{\forall x.\ x + (-x) \equiv 0\} \end{aligned}$$

where $\Sigma_1 := \{-, *\}$, $\Sigma_2 := \{-, 0, +\}$, and $\Sigma := \{-\}$. It can be easily shown that the free models of E_1 and of E_2 with an infinite basis are Σ-decomposable structures.

Lemma 99 *If T_1 and T_2 are Σ-decomposable, then $T_1 \cup T_2$ is Σ-decomposable.*

Proof. For $i = 1, 2$, if $T_i \models \mathrm{DEC}(\Sigma_i, \Sigma)$, then $T_1 \cup T_2 \models \mathrm{DEC}(\Sigma_i, \Sigma)$ since $T_i \subseteq T_1 \cup T_2$. Therefore, $T_1 \cup T_2 \models \mathrm{DEC}(\Sigma_1, \Sigma) \cup \mathrm{DEC}(\Sigma_2, \Sigma)$. The claim then follows from the fact that $\mathrm{DEC}(\Sigma_1, \Sigma) \cup \mathrm{DEC}(\Sigma_2, \Sigma) = \mathrm{DEC}(\Sigma_1 \cup \Sigma_2, \Sigma)$. □

Proposition 100 *Let T_1 and T_2 be both Σ-decomposable. Then, the satisfiability in $T_1 \cup T_2$ of formulae in $\mathit{TRes}(\mathit{Qff}^{\Sigma_1 \cup \Sigma_2}, \Sigma)$ is effectively reducible to the satisfiability in $T_1 \cup T_2$ of formulae in the subclass $\mathit{TRes}(\mathit{Qff}^{\Sigma_1} \otimes \mathit{Qff}^{\Sigma_2}, \Sigma)$.*

Proof. Let $\varphi \wedge \mathit{res}^{\Sigma}(\tilde{u})$ be an element of $\mathit{TRes}(\mathit{Qff}^{\Sigma_1 \cup \Sigma_2}, \Sigma)$. As we have seen before, we can assume without loss of generality that φ is a conjunction of literals. Now let $\varphi_1 \wedge \varphi_2$ be the conjunction of pure formulae obtained by purifying φ as described in Subsection 5.1. Again, the formula $\psi := \varphi_1 \wedge \varphi_2 \wedge \mathit{res}^{\Sigma}(\tilde{u})$ need not be totally Σ-restricted because $\varphi_1 \wedge \varphi_2$ may contain a set $\tilde{v}$ of new variables introduced by the purification process. Notice, however, that each member v of $\tilde{v}$ occurs in $\varphi_1 \wedge \varphi_2$ exclusively as the left-hand side of an equation $v \equiv t_v$ where t_v is a term with top-symbol in $\Sigma_i \setminus \Sigma$. Now assume that ψ is satisfiable in a model $\mathcal{A}$ of $T_1 \cup T_2$ and let g be the top symbol of t_v. Since $g \notin \Sigma$ and $T \models \forall \tilde{x}, \tilde{y}\ g(\tilde{x}) \not\equiv f(\tilde{y})$ for all $f \in \Sigma$, every $\mathcal{A}$-solution of ψ must map v to a Σ-isolated individual. It follows that there is an identification ξ of $\tilde{w} := \tilde{u}, \tilde{v}$ such that $\varphi_1\xi \wedge \varphi_2\xi \wedge \mathit{res}^{\Sigma}(\tilde{w}\xi)$ is satisfiable in T. The proposition's claim then is an easy consequence of the fact that $\varphi_1\xi \wedge \varphi_2\xi \wedge \mathit{res}^{\Sigma}(\tilde{w}\xi) \in \mathit{TRes}(\mathit{Qff}^{\Sigma_1} \otimes \mathit{Qff}^{\Sigma_2}, \Sigma)$. □

Lemma 101 *If T_i is Σ-decomposable, then* $\mathrm{At}(T_i^\Sigma)$ *is axiomatized by the empty equational Σ-theory.*

Proof. Since $T_i \models \mathrm{DEC}(\Sigma_i, \Sigma)$, it is easy to prove by structural induction on Σ-terms that for all Σ-equations $s \equiv s'$, $T_i \models s \equiv s'$ iff $s = s'$. Therefore, $\mathrm{At}(T_i^\Sigma)$ corresponds to the empty Σ-theory. □

The following proposition is analogous to Proposition 83 presented for the case of theories sharing only constant symbols.

Proposition 102 *Assume that for $i = 1, 2$,*

- *T_i is stably Σ-free over $TRes(Qff^{\Sigma_i}, \Sigma)$;*
- *T_i is Σ-decomposable.*

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. Moreover,

- *T is stably Σ-free over $TRes(Qff^{\Sigma_1 \cup \Sigma_2}, \Sigma)$;*
- *T is Σ-decomposable.*

Proof. According to Lemma 101, T_1 and T_2 have the same atomic Σ-theory, and so we can apply Proposition 54, Theorem 53, and Theorem 28, where $\mathcal{L} = Qff$. Thanks to Proposition 100, we can substitute $TRes(Qff^{\Sigma_1} \otimes Qff^{\Sigma_2}, \Sigma)$ by $TRes(Qff^{\Sigma_1 \cup \Sigma_2}, \Sigma)$. □

It is rather easy to find examples of theories satisfying the assumptions of the previous proposition.

Proposition 103 *If T_i is the (complete) theory of a Σ-decomposable Σ_i-structure $\mathcal{A}_i$ free over a countably infinite basis X, then*

- *T_i is stably Σ-free over $TRes(Qff^{\Sigma_i}, \Sigma)$;*
- *T_i is Σ-decomposable.*

Proof. It is easy to see that the set $Y := Is(\mathcal{A}^\Sigma)$, which includes X as T_i^Σ is collapse-free, generates $\mathcal{A}_i{}^\Sigma$. We can show by term induction that $s = s'$ for all Σ-terms s, s' for which there is a discrete tuple $\tilde{y}$ in Y such that $\mathcal{A}_i{}^\Sigma \models (s \equiv s')[\tilde{y}]$. By Proposition 43 then, $\mathcal{A}_i^\Sigma$ is free over Y in $Mod(\mathrm{At}(T_i^\Sigma))$, which is the class of all Σ-structures according to Lemma 101. Then, the first point of the proposition follows directly from the definition of stable Σ-freeness, the second point from the definition of T_i. □

From the above results it is then easy to prove the following.

Proposition 104 *Assume that for $i = 1, 2$, T_i is the (complete) theory of some Σ-decomposable Σ_i-structure free over a countably infinite basis. 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.*

When the word problem for each T_i above is decidable, the result in Proposition 104 is actually a special case of a more general combination result for theories sharing constructors. We will prove that result in the next subsection with Theorem 113. Then, to see that Proposition 104 follows from Theorem 113, it will be enough to consider the following properties of theories of Σ-decomposable structures.

Proposition 105 *If T_i is the (complete) theory of a Σ-decomposable structure free over a countably infinite basis X, then*

- *Σ is a finite set of constructors for T_i;*
- *$G_{T_i}(\Sigma, V)$ is closed under instantiation into itself;*
- *normal forms are computable for Σ and T_i.*

Proof. We already know that the Σ-reduct of $\mathcal{A}_i$ is free over X in $\mathrm{At}(T_i^{\Sigma})$, which is empty. It follows that it is free in T_i^{Σ}, where T_i^{Σ} is collapse-free. By Theorem 59, we then have that Σ is a set of constructors for T_i. Moreover, it is easy to see that $G_{T_i}(\Sigma, V)$ is the set of terms in $T(\Sigma_i, V)$ with top-symbol in $V \cup \Sigma_i \setminus \Sigma$. From this it obviously follows that $G_{T_i}(\Sigma, V)$ is closed under instantiation into itself. It also follows that every Σ_i term is in normal form, and so we can choose $\mathrm{NF}^{\Sigma}_{T_i}$ as the identity on $T(\Sigma_i, V)$, which is obviously computable. □

8.4 Theories Sharing Constructors

In this subsection, we consider the combination of complete theories sharing constructors. Specifically, 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;
- $\mathrm{At}(\mathcal{A}_1^{\Sigma}) = \mathrm{At}(\mathcal{A}_2^{\Sigma})$;
- Σ is a finite set of constructors for T_i.

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. However, our current result 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. What we can show is that the procedure is complete for input formulae that are already totally Σ-restricted.

Although this may be a strong restriction in general, it has a remarkable side-effect. As we will show 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 106 *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 complete theory of $\mathcal{A}_i$, we know that 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 59 that $\mathcal{A}^{\Sigma}$ is free in T_i^{Σ} over some countably infinite set Y. From this and Proposition 46, it is easy to see that $\mathcal{A}^{\Sigma}$ is also free in $\mathrm{At}(T_i^{\Sigma})$ over Y. □

Proposition 107 *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 106, for $i = 1, 2$, T_i is is stably Σ-free over any class of formulae, in particular over $Res(\mathcal{L}^{\Sigma_i}, \Sigma)$. We can then conclude by Theorem 53(1), that T_1 and T_2 are partially N-O-combinable over $\mathcal{L}$.

From Lemma 106 again and Proposition 54, 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 59, we can show exactly as in the proof of Theorem 53(2) that T_1 and T_2 are totally N-O-combinable over $TRes(\mathcal{L}^{\Sigma_1} \otimes \mathcal{L}^{\Sigma_2}, \Sigma)$. □

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 chooses, deterministically, 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, one is more likely to be interested in the satisfiability of formulae in $TRes(\mathcal{L}^{\Sigma_1 \cup \Sigma_2}, \Sigma)$, not just 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 107 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 61);
- the word problem for T_i is decidable,
- normal forms are computable for Σ and T_i.

Lemma 108 *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 58 and Corollary 60. □

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

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

Proof. Let $i \in \{1, 2\}$. Since T includes T_i and T_i is 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 62 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. □

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 110 *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 [Tin99], 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 $T(\Sigma, G^*_T(\Sigma, V))$.[45] We will appeal to these facts in Proposition 112.

Lemma 111 *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 $k \in J_1$;

[45] A proof of these facts can also be found in [BT01] for the case of equational theories.

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 term 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_j^i is an abstraction variable.

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 $(v \equiv r), (v' \equiv r') \in F$,

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

From the properties in the previous points and the fact that F consists only of equations added by purification it 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 109 that $(\mathcal{A}, \alpha) \models iso^{\Sigma}(v_j^i)$ for all $i = 1, 2$ and $j \in J_i$. It follows by point 2 above and the definition of iso^{Σ} that $(\mathcal{A}, \alpha) \models S \cup iso^{\Sigma}(\tilde{u})$. □

We are now ready to prove our reducibility claim.

Proposition 112 *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 108 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 111.
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 $\varphi_1 \wedge \varphi_2 \wedge res^\Sigma(\tilde{v})$ 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 $(\varphi_1 \wedge \varphi_2)\xi \wedge res^\Sigma(\tilde{u}\xi)$ 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 $(\varphi_1 \wedge \varphi_2)\xi \wedge res^\Sigma(\tilde{u}\xi)$ is satisfiable in T. Observing that $\tilde{v}$ is contained in $\tilde{u}\xi$, we can conclude by the definition of res^Σ that $(\varphi_1 \wedge \varphi_2)\xi \wedge res^\Sigma(\tilde{v})$ is satisfiable in T. But then, $\varphi_1 \wedge \varphi_2 \wedge res^\Sigma(\tilde{v})$ is also satisfiable in T.

Now assume that $\varphi_1 \wedge \varphi_2 \wedge res^\Sigma(\tilde{v})$ is satisfiable in T. By construction of φ_i and definition of res^Σ, we can conclude that $S \cup iso^\Sigma(\tilde{v}) \cup dif(\tilde{v})$ is satisfiable in T, where S is the set generated at step 2 of the procedure above. By Lemma 111 then $S' := S \cup iso^\Sigma(\tilde{u}) \cup dif(\tilde{v})$ 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.[46] It is immediate that ξ identifies no two variables in $\tilde{v}$ and that the set

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

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

Finally, we obtain the following decidability result.

Theorem 113 *Let T_1, T_2 be such 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;*

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

- $\mathrm{At}(\mathcal{A}_1^\Sigma) = \mathrm{At}(\mathcal{A}_2^\Sigma)$;
- $G_{T_i}(\Sigma, V)$ *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 107, Proposition 112, 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 the following section, where we concentrate on complete theories of free algebras. There, we will reformulate the above conditions in terms of more familiar properties of equational theories, and provide some specific examples as well.

9 Complete Theories of Free Algebras

In this section we establish a connection between our combination method and constraint reasoners for the entailment of equational constraints. Recall that our method combines satisfiability procedures, not entailment procedures. As mentioned in the introduction, some satisfiability procedures check that their input formulae are satisfiable in *at least one* model of the constraint theory, whereas other check that their input formulae are satisfiable in *every* model of the constraint theory. Now, with complete theories this distinction disappears because a formula is satisfiable in a complete theory if and only if it is satisfiable in every model of the theory.

Now, it so happens that most satisfiability problems of the second sort in a given (non-trivial) equational theory E can be recast as satisfiability problems of the second sort in the complete theory of its free algebra over a countably infinite basis. In this section, we exploit this fact to show how our combination results can be specialized in terms of familiar notions from equational reasoning.

A similar approach was already presented in Chapter 4 of [Rin93]. It is argued there that the Nelson-Oppen method can be seen as a basic combination method, which can be turned into a combination method for constraint entailment by the addition of more steps. We follow up on this argument in Subsection 9.3 where we compare our method to a combination method by Baader and Schulz.

We start by briefly recalling the definition of the most important entailment problems for equational theories (see [Sie89, JK91, BS94] for comprehensive surveys), and then recasting them as satisfiability problems in an appropriate complete theory. In the following, E will be a non-trivial equational theory of signature Ω. By Lemma 47, all the free models of E over a countably-infinite basis are isomorphic. We will identify them all and denote any of them $\mathrm{Fr}_\omega(E)$.

E-unification. An *E-unification problem* is a conjunction of Ω-equations. An E-unification problem φ is *solvable* if there is a substitution σ into Ω-terms such that $s\sigma =_E t\sigma$ for every equation $s \equiv t$ of φ. Equivalently, φ is solvable if φ is satisfiable in $\mathrm{Fr}_\omega(E)$.

The substitution σ above is called an *E-unifier* of φ. Sometimes, the set of all E-unifiers of an E-unification problem φ can be denoted by a *minimal complete* set $\mu U_E(\varphi)$ of E-unifiers. In that case, every E-unifier of φ is an *E-instance* of a unifier in $\mu U_E(\varphi)$ w.r.t. an appropriate substitution ordering.[47] E-unification is said to be *finitary* if $\mu U_E(\varphi)$ is finite for all E-unification problems φ. In that case, it may be possible to devise a *E-unification algorithm*, that is, a procedure that returns $\mu U_E(\varphi)$ for every input unification problem φ.

E-unification is decidable if there is an effective procedure for determining whether an arbitrary E-unification problem is solvable. Notice that whereas a decision procedure for E-unification can be trivially derived from an E-unification algorithm, if one exists, a decision procedure for E-unification may exist even if there are no E-unification algorithms.

E-disunification. E-disunification is an extension of E-unification to include negated equations. An *E-disunification problem* is a conjunction of Ω-equations and disequations. An E-disunification problem φ is *solvable* if there is a substitution σ into Ω-terms such that $s\sigma =_E t\sigma$ for every equation $s \equiv t$ of φ, and $s\sigma \neq_E t\sigma$ for every disequation $s \not\equiv t$ of φ. Equivalently, φ is solvable if φ is satisfiable in $\mathrm{Fr}_\omega(E)$.

The notions of E-disunification algorithm and decision procedure can be defined as in E-unification.

E-matching. E-matching can be seen as a restricted kind of E-unification. An *E-matching problem* is a directed Ω-equation, usually represented as $s \leq t$. An E-matching problem $s \leq t$ is *solvable* if there is a substitution σ into Ω-terms such that $s\sigma =_E t$.

The notions of E-matching algorithm and decision procedure can be defined as in E-unification.

E-validity. E-validity[48], commonly referred to as the *word problem in E*, can be seen in turn as a restricted kind of E-matching. An *E-validity problem* is just a

[47] See [Sie89] or [BS94] for more details.

[48] The term "E-validity" is non-standard, we adopt it here for uniformity.

Ω-equation. An E-validity problem $s \equiv t$ is *solvable* if $s =_E t$. In other words, if the E-matching problem $s \leq t$ is solvable with the empty substitution. Equivalently, $s \equiv t$ is solvable if it is satisfied in $\mathrm{Fr}_\omega(E)$ by a valuation that assigns its variables to distinct individuals in the basis of $\mathrm{Fr}_\omega(E)$.

E-unification as defined above is often called *elementary E-unification* because all the unification problems considered are over the signature Ω of E. One usually speaks of *E-unification with (free) constants* if unification problems may contain constant symbols not in Ω, and of *general E-unification* if they may contain arbitrary function symbols not in Ω. In either case, the theory of interest is not really E, but the union of E and the empty theory over the extra symbols. A similar terminology is used for E-disunification and E-matching.

Later, we will use the following.

Fact 114 *The word problem in E is decidable whenever E-matching with constants is decidable.*

To see that it is enough to notice that for any terms s, t in the signature of E, $s =_E t$ iff there is an instantiation σ of the variables of s and t into distinct free constants such that the (ground) E-matching problem $s\sigma \leq t\sigma$ is solvable.

In the following, we fix a non-trivial equational theory E of (functional) signature Ω and denote by T_E the complete Ω-theory of $\mathrm{Fr}_\omega(E)$. Then, we investigate the conditions under which the satisfiability in T_E of totally restricted formulae (cf. Definition 17) is decidable. There, we will often implicitly appeal to the fact that, by construction of T_E, a formula is satisfiable (valid) in T_E iff it is satisfiable (valid) in $\mathrm{Fr}_\omega(E)$.

9.1 Totally Restricted Formulae

First, notice that the satisfiability in T_E of equational (i.e. quantifier-free) formulae is reducible to E-disunifiability. In fact, an equational formula φ is satisfiable in T_E exactly when at least one disjunct φ_i of φ's disjunctive normal form is satisfiable in T_E. But each of these disjuncts is satisfiable in T_E exactly when it is solvable as a E-disunification problem.

Now, a non-deterministic decision algorithm for E-disunifiability can be constructed if the word problem in E is decidable and E-unification is finitary. In fact, suppose we are interested in the satisfiability of a conjunction φ of equations and disequations. Then we can do the following.

1. Let
   E_φ be the set of φ's equations and
   D_φ be the set of φ's disequations.

2. Compute the minimal complete set U of E-unifiers of E_φ.

3. Succeed if there is a $\mu \in U$ such that $s\mu \neq_E t\mu$ for all $(s \not\equiv t) \in D_\varphi$.
   Fail otherwise.

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

Proposition 115 *If the word problem in E is decidable and there exists a finitary E-unification algorithm, then the satisfiability in T_E of formulae in Qff^Ω is decidable.*

Now suppose that some $\Sigma \subseteq \Omega$ is a finite set of constructors for E. When normal forms are computable for Σ and E in the sense of Definition 64, it is also possible to decide the satisfiability in T_E of totally Σ-restricted equational formulae, i.e. formulae

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

with $\mathit{res}^\Sigma(\tilde{v}) := \mathit{dif}(\tilde{v}) \wedge \mathit{iso}^\Sigma(\tilde{v})$. In essence, it is enough to add to the procedure seen above a check on the top symbols of the solutions computed by E-unification to obtain the following result.

Proposition 116 *Whenever*

- *Σ is a finite set of constructors for E,*
- *normal forms are computable for Σ and E,*
- *the word problem in E is decidable,*
- *there exists a finitary E-unification algorithm,*

the satisfiability in T_E of formulae in $\mathit{TRes}(\mathit{Qff}^\Omega, \Sigma)$ is decidable.

Proof. Let NF_E^Σ be the function that computes the normal form for E and Σ for terms in $T(\Omega, V)$. Recall that for a totally Σ-restricted formula to be satisfiable, the free variables of its body must be assigned to distinct Σ-isolated individuals. By the previous result then, it is enough to verify that there is a substitution that satisfies φ and (a) maps no variable of φ to a term whose normal form starts with a symbol in Σ, (b) maps no two variables to E-equivalent terms. More precisely, we can do the following.

1. Let E_φ be the set of φ's equations and D_φ be the set of φ's disequations.

2. Compute the minimal complete set U of E-unifiers of E_φ.

3. Normalize the E-unifiers in U by replacing each $\mu \in U$ by a substitution μ' such that $v\mu' = \mathrm{NF}_E^\Sigma(v\mu)$ for all $v \in \mathcal{D}om(\mu)$.

4. Succeed if there is a μ in (the new) U such that

   (a) $s\mu \neq_E t\mu$ for all $(s \not\equiv t) \in D_\varphi$,

(b) $v\mu(\epsilon) \notin \Sigma$ for all $v \in \mathcal{V}ar(\varphi)$,

(c) $u\mu \neq_E v\mu$ for all distinct $u, v \in \mathcal{V}ar(\varphi)$.

Fail otherwise.

By the various computability assumptions, it is easy to see that the procedure above is effective and terminates on all inputs $\varphi(\tilde{v}) \in \mathit{Qff}^{\Omega}$. We leave it to the reader to show that $\varphi \wedge \mathit{res}^{\Sigma}(\tilde{v})$ is satisfiable in $\mathrm{Fr}_{\omega}(E)$ iff the procedure succeeds on input φ, from which the claim easily follows. □

From the above, we can produce a first specialization of the combination result in Theorem 113. In the following, we will let E_1, E_2 be two equational theories with countable signatures Σ_1, Σ_2 and $\Sigma := \Sigma_1 \cap \Sigma_2$.

Proposition 117 *Assume that for* $i = 1, 2$

- $\mathrm{At}(E_1^{\Sigma}) = \mathrm{At}(E_2^{\Sigma})$;
- $G_{E_i}(\Sigma, V)$ *is closed under instantiation into itself;*
- Σ *is a finite set of constructors for* E_i*;*
- *normal forms are computable for* Σ *and* E_i*;*
- *the word problem for* E_i *is decidable.*
- *there exists a finitary* E_i*-unification algorithm.*

Then, the satisfiability in $T_{E_1} \cup T_{E_2}$ *of formulae in* $\mathit{TRes}(\mathit{Qff}^{\Sigma_1 \cup \Sigma_2}, \Sigma)$ *is decidable.*

Proof. Assume that all the assumptions above hold. To prove the claim then it is enough to show that Theorem 113 is applicable. Now, T_{E_1} and T_{E_2} satisfy by construction the first itemized condition in Theorem 113, while the assumptions above correspond to the remaining itemized conditions in the theorem. Finally, by Proposition 116, the satisfiability in T_{E_i} of formulae in $\mathit{TRes}(\mathit{Qff}^{\Sigma_i}, \Sigma)$ is decidable for $i = 1, 2$. □

9.2 Theories Generated by TRSs

.

It may be still not immediate to show that theories such as E_1 and E_2 above satisfy all the requirements in Proposition 117. One possibility is to try to show that each of them is generated by a term rewriting system $S_i = (R_i, E_0)$ where E_0 is a collapse-free Σ-theory (cf. Section 7.4). The idea would be to assume that S_i is a semi-complete TRS and $\rightarrow_{S_i}$ is effectively computable. The effective computability of $\rightarrow_{S_i}$ may be provided by a general E_0-matching algorithm. If S_i is not just semi-complete but complete, then E_0 is not only collapse-free but also regular[49], and a general E_0-matching algorithm can be constructed from a E_0-matching algorithm with free constants.

[49] An equational theory E is *regular* if $\mathcal{V}ar(s) = \mathcal{V}ar(t)$ whenever $s =_E t$.

Lemma 118 *Let $i \in \{1,2\}$. If $S_i = (R_i, E_0)$ is a complete TRS for E_i, Σ is a finite set of TRS-constructors for R_i, and E_0 is a collapse-free equational Σ-theory such that E_0-matching with free constants is decidable, then*

1. *Σ is a finite set of constructors for E_i;*
2. *normal forms are computable for Σ and E_i;*
3. *the word problem in E_i is decidable.*

Proof. 1) Immediately by Proposition 79.

2) Given that S_i is terminating for being complete by assumption, it can be shown using standard results in Term Rewriting that E_0 is necessarily regular. Together with the assumption that E_0 is collapse-free and E_0-matching with free constants is decidable this entails that general E_0-matching is decidable (see [Rin96a]). Now, consider any decision procedure for general E_0-matching. It can be shown (see [Rin96a] again) that this procedure can be turned into a general E_0-matching procedure. It follows that $\rightarrow_{S_i}$, which is defined in terms of general E_0-matching (specifically E_0-matching with free symbols in $\Omega \setminus \Sigma$), is effectively computable. By Lemma 78 then, we can simply define a normal form function $\mathrm{NF}_{E_i}^{\Sigma}$ so that $\mathrm{NF}_{E_i}^{\Sigma}(t) = t{\downarrow}_{S_i}$ for all $t \in T(\Sigma_i, V)$.

3) Simply notice that, by the computability of $\rightarrow_{S_i}$, the word problem in E_i is effectively reducible to the word problem in E_0, which is decidable by Fact 114. □

From this lemma we can then easily obtain the following corollary of Proposition 116.

Corollary 119 *Let $i \in \{1,2\}$. If there is a TRS R_i modulo and equational Σ-theory E_0 such that*

- *$S_i = (R_i, E_0)$ is complete for E_i,*
- *E_0 is collapse-free,*
- *Σ is a finite set of* TRS*-constructors,*
- *E_0-matching with free constants is decidable,*
- *there exists a finitary E_i-unification algorithm,*

then the satisfiability in T_{E_i} of formulae in $TRes(Qff^{\Sigma_i}, \Sigma)$ is decidable.

This corollary produces in turn the following combination result as a corollary of Proposition 117, provided that each $G_{E_i}(\Sigma, V)$ is closed under instantiation into itself.

Corollary 120 *For $i = 1,2$, let R_i be a TRS over $T(\Sigma_i, V)$ and E_0 an equational Σ-theory such that*

- *$S_i = (R_i, E_0)$ is a complete TRS for E_i,*

- *E_0 is collapse-free,*
- *Σ is a finite set of* TRS*-constructors for R_i,*
- *E_0-matching with free constants is decidable,*
- *there exists a finitary E_i-unification algorithm,*
- *$G_{E_i}(\Sigma, V)$ is closed under instantiation into itself.*

Then, the satisfiability in $T_{E_1} \cup T_{E_2}$ of formulae in $TRes(Qff^{\Sigma_1 \cup \Sigma_2}, \Sigma)$ is decidable.

Here is a simple example of pairs of theories that can be easily shown to satisfy all the conditions in the above corollary.

Example 121 *The equational Σ_i-theories generated by the TRSs*

$$\begin{array}{lcl}(R_1, E_0) & := & (\{h_1(x \cdot y) \rightarrow h_1(x) \cdot h_1(y)\}, \{x \cdot y = y \cdot x\}) \\ (R_2, E_0) & := & (\{h_2(x \cdot y) \rightarrow h_2(x) \cdot h_2(y)\}, \{x \cdot y = y \cdot x\})\end{array}$$

where $\Sigma_i := \{\cdot, h_i\}$ for $i = 1, 2$, and $\Sigma := \{\cdot\}$.

We believe that more interesting examples can be found in the literature. In fact, a lot of work has been already done to obtain unification algorithms for complete TRSs modulo an equational theory by exploiting narrowing techniques [JK86, Han94, MH94]. Therefore, it should be possible to use the existing narrowing-based algorithms to decide the satisfiability of totally restricted equational formulae.

9.3 A Comparison with the Baader-Schulz Procedure

In Subsection 9.1, we have seen how to build decision procedures for the satisfiability of totally restricted equational formulae in the complete theory T_{E_i}, where E_i is a non-trivial equational theory of signature Σ_i. Then, we have applied our combination techniques to obtain a decision procedure for the satisfiability of totally Σ-restricted equational formulae in theories of the form $T_{E_1} \cup T_{E_2}$, where $\Sigma = \Sigma_1 \cap \Sigma_2$.

In [BS95b], Baader and Schulz present a combination method to decide the solvability of disunification problems in the union of signature-disjoint equational theories. Observing that when the shared signature Σ is empty, the class of Σ-restricted disunification problems coincides with the class of disunification problems, one can then be induced to conclude that the combination results in this section are a generalization of Baader and Schulz's to the combination of non-signature-disjoint equational theories.

This is not quite the case. The solvability of disunification problems in E_i coincides with the satisfiability of equational formulae in T_{E_i} and so, by the same token, the solvability of disunification problems in $E_1 \cup E_2$ coincides with the satisfiability of equational formulae in $T_{E_1 \cup E_2}$. Our combination approach, however, yields a procedure for deciding the satisfiability of equational formulae in the theory $T_{E_1} \cup T_{E_2}$, which is not equivalent to the theory $T_{E_1 \cup E_2}$.

Nevertheless, the two theories are not unrelated: unsatisfiability answers in $T_{E_1} \cup T_{E_2}$ are meaningful for $T_{E_1 \cup E_2}$ as well. To show this, we will assume again that $\Sigma = \Sigma_1 \cap \Sigma_2$ is a set of constructors for both E_1 and E_2, and E_1 and E_2 entail the same Σ-equations, keeping in mind that these assumptions are always satisfied whenever Σ is empty. In [Tin99], it is shown that $\mathrm{Fr}_\omega(E_1 \cup E_2)$ is embedded in a fusion of $\mathrm{Fr}_\omega(E_1)$ and $\mathrm{Fr}_\omega(E_2)$. This fact entails the following.

Proposition 122 *Every equational $(\Sigma_1 \cup \Sigma_2)$-formula satisfiable in $T_{E_1 \cup E_2}$ is satisfiable in $T_{E_1} \cup T_{E_2}$.*

Proof. Let φ be an equational $(\Sigma_1 \cup \Sigma_2)$-formula satisfiable in $T_{E_1 \cup E_2}$. By construction of $T_{E_1 \cup E_2}$, φ is satisfiable in $\mathrm{Fr}_\omega(E_1 \cup E_2)$. Since $\mathrm{Fr}_\omega(E_1 \cup E_2)$ is embedded in a fusion $\mathcal{F}$ of $\mathrm{Fr}_\omega(E_1)$ and $\mathrm{Fr}_\omega(E_2)$, and φ is quantifier-free, φ is satisfiable in $\mathcal{F}$. The claim then follows from the fact that $\mathcal{F}$ is a model of $T_{E_1} \cup T_{E_2}$. □

The converse of this proposition does not hold, as the following example shows.

Example 123 *Assume that $\Sigma_1 \cap \Sigma_2 = \emptyset$ and for $i = 1, 2$, Σ_i contains a constant a_i. Then, the formula $x = a_1 \wedge x = a_2$ is satisfiable in any model of $T_{E_1} \cup T_{E_2}$ interpreting a_1 and a_2 identically, but it is unsatisfiable in $T_{E_1 \cup E_2}$, unless the sentence $\forall x\ x = a_i$ holds in $\mathrm{Fr}_\omega(E_1 \cup E_2)$ contradicting the assumption that E_i is non-trivial.*

By the contrapositive of Proposition 122, no equational formula unsatisfiable in $T_{E_1} \cup T_{E_2}$ is satisfiable in $T_{E_1 \cup E_2}$. Therefore, in the disjoint case, we can still relate our combination procedure for the satisfiability in $T_{E_1} \cup T_{E_2}$ to that of Baader and Schulz for the satisfiability in $T_{E_1 \cup E_2}$.[50]

In fact, our approach may be viewed as a simplified form of the Baader-Schulz approach, one that implements only the first two steps of their procedure: the deterministic step transforming an impure formula φ into a conjunction of pure formulae $\varphi_1 \wedge \varphi_2$, and the non-deterministic step producing an identification ξ of $\mathcal{V}ar(\varphi_1) \cap \mathcal{V}ar(\varphi_2)$.

If there are no identifications ξ such that $\varphi_1\xi \wedge \xi_{\neq}$ and $\varphi_2\xi \wedge \xi_{\neq}$ are both satisfiable in T_{E_1} and T_{E_2}, respectively, then our procedure can already conclude that $\varphi_1 \wedge \varphi_2$, and so φ, is unsatisfiable in $T_{E_1 \cup E_2}$, because it is already unsatisfiable in $T_{E_1} \cup T_{E_2}$. Otherwise, φ will have a solution $T_{E_1} \cup T_{E_2}$, but may still have no solutions in $T_{E_1 \cup E_2}$. This explains why the Baader-Schulz procedure requires additional steps (unfortunately, non-deterministic ones) to be able to decide the satisfiability of φ in $T_{E_1 \cup E_2}$.

To conclude this section, we point out another connection between $T_{E_1} \cup T_{E_2}$ and $T_{E_1 \cup E_2}$: they have the same universal theory (whether they share function symbols or not).

[50] A similar argument, relating the Nelson-Oppen and the Baader-Schulz procedures, was also presented in [Kep98].

Proposition 124 *For every universal sentence φ of signature $\Sigma_1 \cup \Sigma_2$,*

$$T_{E_1} \cup T_{E_2} \models \varphi \quad \textit{iff} \quad T_{E_1 \cup E_2} \models \varphi \ .$$

Proof. ($\Rightarrow$) If $T_{E_1} \cup T_{E_2} \models \varphi$, then φ is true in every fusion of $\mathrm{Fr}_\omega(E_1)$ and $\mathrm{Fr}_\omega(E_2)$. Since $\mathrm{Fr}_\omega(E_1 \cup E_2)$ is embedded in one of such fusions and φ is universal, we can conclude that φ is true in $\mathrm{Fr}_\omega(E_1 \cup E_2)$.

($\Leftarrow$) If $\mathrm{Fr}_\omega(E_1 \cup E_2) \models \varphi$, then $E_1 \cup E_2 \models \varphi$ (a consequence of Theorem 3.7 in [BS98]). Since $T_{E_1} \cup T_{E_2} \models E_1 \cup E_2$, we obtain that $T_{E_1} \cup T_{E_2} \models \varphi$ by transitivity of logical entailment. □

If, as we conjecture, $\mathrm{Fr}_\omega(E_1 \cup E_2)$ was in fact a fusion of $\mathrm{Fr}_\omega(E_1)$ and $\mathrm{Fr}_\omega(E_2)$, not just a substructure of one, it is possible to show by a slight variation of the proof above that $T_{E_1} \cup T_{E_2}$ and $T_{E_1 \cup E_2}$ have the same *positive* theory. For now, we can only show is that the positive theory of $T_{E_1} \cup T_{E_2}$ includes that of $T_{E_1 \cup E_2}$.

10 Fusions of Initial Models

As we have already seen, some combination methods consider the combination of procedures that decide satisfiability with respect to fixed structures. This is typical of constraint solvers, which not only say whether an input constraint is satisfiable in a given structure, but also return a set of "solutions" for the constraint, if any.[51] For instance, a unification algorithm for a certain equational theory E decides the satisfiability of unification problems in the free model $\mathrm{Fr}_\omega(E)$. The algorithm described in [BS96] takes a unification algorithm for $\mathrm{Fr}_\omega(E_1)$ and one for a $\mathrm{Fr}_\omega(E_2)$, where E_1 and E_2 are two disjoint equational theories, and produces a unification algorithm for $\mathrm{Fr}_\omega(E_1 \cup E_2)$.

An interesting question is whether our combination results can be used to combine constraint solvers as well. In this section, we show one example in which this is indeed the case. We will consider *initial* structures, that is, structures free over an empty set of generators. Initial structures have a number of properties that make them ideal as solution structures for certain domains of computation. The most important use of initial structures is perhaps in algebraic specification, where they provide a semantics for abstract data types.

Now, in that field since initial structures are often (initial) models of theories obtained as the union of theories themselves admitting initial models. It would be very useful then to be able to build constraint solvers for initial structures in a modular way. As we show below, this is possible in some cases. Under the right conditions, we can use our results to combine a constraint solver for the initial model of a theory T_1 and one for the initial model of a theory T_2, into a constraint solver for the initial model of $T_1 \cup T_2$.

[51]Strictly speaking, a constraint solver returns a simplified version of the input constraint from which the solutions can be easily elicited.

For the rest of the section, let us fix a finite signature Σ and an initial Σ-structure $\mathcal{A}^\Sigma$. We will denote by $T(\Sigma)$ to denote $T(\Sigma, \emptyset)$, the set of *ground* (i.e. variable-free) Σ-terms.

Definition 125 *Where $\mathcal{A}^\Sigma$ is an initial Σ-structure, we denote by* $\mathrm{IT}(\mathcal{A}^\Sigma)$ *the class of theories T such that*

- *$\Sigma \subseteq \Sigma_T$;*
- *T admits an initial model whose Σ-reduct is isomorphic to $\mathcal{A}^\Sigma$.*

The theories presented in the examples of Section 7, for instance, are all of the type defined above. The theories in $\mathrm{IT}(\mathcal{A}^\Sigma)$ are all sufficiently complete (cf. [DJ90]).

Definition 126 (Sufficiently Complete Theory) *Let T be a theory of signature Ω with $\Sigma \subseteq \Omega$. T* is sufficiently complete w.r.t Σ *if for all terms $t \in T(\Omega)$ there is a term $s \in T(\Sigma)$ such that $t =_T s$.*

Proposition 127 *For all $T_i \in \mathrm{IT}(\mathcal{A}^\Sigma)$, T_i is sufficiently complete w.r.t. Σ.*

Proof. Let Σ_i be T_i's signature and let $\mathcal{A}_i$ be an initial model of T_i. Since $\mathcal{A}^\Sigma$ is initial, we know that it is generated by $\emptyset$. From the assumption that ${\mathcal{A}_i}^\Sigma$ and $\mathcal{A}^\Sigma$ are isomorphic it follows that $\mathcal{A}_i$ is actually Σ-generated by $\emptyset$. This entails that for every ground Σ_i-term t, there is a ground Σ-term s such that $\mathcal{A}_i \models t \equiv s$. Given that $\mathcal{A}_i$ is free in T_i, the claim then follows by Proposition 43. □

Since any ground term t of signature Σ_i is equivalent in T_i to a ground Σ-term s, we will call s a Σ-*normal form* of t. Again, a term may have several normal forms in T_i, but they are all equivalent in T_i. We will say that Σ-*normal forms are computable for T_i*, if there is a computable function that maps every term in $T(\Sigma_i)$ to one of its normal forms.

By a fairly standard induction argument, we can show that the union of two theories in the same $\mathrm{IT}(\mathcal{A}^\Sigma)$ class admits Σ-normal forms as well.

Proposition 128 *For $i = 1, 2$, let T_i be a Σ_i-theory in* $\mathrm{IT}(\mathcal{A}^\Sigma)$*. Then, for all terms $t \in T(\Sigma_1 \cup \Sigma_2)$ there is a term $s \in T(\Sigma)$ such that $t =_{T_1 \cup T_2} s$.*

In the situation of the above proposition, we can use the same inductive argument to show that Σ-normal forms are computable for $T_1 \cup T_2$ whenever they are computable for T_1 and for T_2.

Proposition 128 is actually a consequence of the stronger result below.

Proposition 129 *For $i = 1, 2$, let T_i be a Σ_i-theory in* $\mathrm{IT}(\mathcal{A}^\Sigma)$ *with initial model $\mathcal{A}_i$. If $\Sigma_1 \cap \Sigma_2 = \Sigma$, then $\mathcal{A}_1$ and $\mathcal{A}_2$ are Σ-fusible. Moreover, for all isomorphisms h of ${\mathcal{A}_1}^\Sigma$ onto ${\mathcal{A}_2}^\Sigma$, the canonical fusion of $\mathcal{A}_1$ and $\mathcal{A}_2$ w.r.t. h is an initial model of $T_1 \cup T_2$.*

Proof. It is immediate from the definition of $\mathrm{IT}(\mathcal{A}^\Sigma)$ that ${\mathcal{A}_1}^\Sigma$ and ${\mathcal{A}_2}^\Sigma$ are initial in the same class of Σ-structures. By Proposition 48, this trivially implies that $\mathcal{A}_1$ and $\mathcal{A}_2$ are Σ-fusible.

Let h be an isomorphisms of ${\mathcal{A}_1}^\Sigma$ onto ${\mathcal{A}_2}^\Sigma$ and $\mathcal{F}$ the canonical fusion of $\mathcal{A}_1$ and $\mathcal{A}_2$ w.r.t. h. Observing that $\mathcal{F}^{\Sigma_2}$ coincides with $\mathcal{A}_2$ by construction, it is immediate that $\mathcal{F}$ is Σ-generated by the empty set. Since we already know that $\mathcal{F}$ is a model of $T_1 \cup T_2$, by Proposition 43, all we need to show is that every ground atomic formula φ of signature $\Sigma_1 \cup \Sigma_2$ true in $\mathcal{F}$ is entailed by $T_1 \cup T_2$.

Hence, let φ be such a formula. By Proposition 128, we can assume with no loss of generality that all the arguments of φ are ground Σ-terms. It follows that φ is a Σ_i-formula for $i = 1$ or $i = 2$. But then, since $\mathcal{F}^{\Sigma_i}$ is isomorphic to $\mathcal{A}_i$ which is free in T_i, we can conclude, using again Proposition 43, that $T_i \models \varphi$. It follows immediately that $T_1 \cup T_2 \models \varphi$. □

Corollary 130 *The class of theories* $\mathrm{IT}(\mathcal{A}^\Sigma)$ *is closed under union.*

We will consider component theories T_i and constraint languages such that the satisfiability problem in the initial model $\mathcal{A}_i$ of T_i is *finitary*, that is, each constraint φ has at most a finite number of $\mathcal{A}_i$-solutions, denoted by $Sol_{\mathcal{A}_i}(\varphi)$. Since $\mathcal{A}_i$ will be Σ-generated by the empty set, this is the same as saying that for each constraint φ, there are at most a finite number of instantiations ρ of $\mathcal{V}ar(\varphi)$ into ground Σ-terms such that $\varphi\rho$ is true in $\mathcal{A}_i$.

In this setting it is possible to compute in a modular way the solutions of mixed formulae in the initial model of $T_1 \cup T_2$ by using a slight modification of the combination algorithm described in Section 5.

Here is the main idea. Once the mixed input formula has been converted into a pair of the form $\langle\varphi_1, \varphi_2\rangle$, the combination algorithm first asks the constraint solver of T_1, say, for all the solutions for φ_1. Since the returned set of solutions will be a (finite) set S_1 of mappings of φ_1's variables to ground Σ-terms, instead of guessing as in the original formulation, the algorithm can then "deduce" from S_1 an instantiation ρ into ground Σ-terms of *all* the variables shared by φ_1 and φ_2. At this point, no identification step is necessary; the algorithm can simply pass the formula $\varphi_2\rho$ to the other contraint solver and ask for the set S_2 of all its solutions.

It is clear that this combination algorithm will converge for all inputs, as S_1 and S_2 are both guaranteed to be finite. It is also clear how to modify the algorithm so that it becomes a constraint solver for the initial model of $T_1 \cup T_2$, not just a satifiability procedure: it is enough for the algorithm to pair every solution in S_1 with the corresponding solutions in S_2 and return them.

In the following, we assume that Σ-normal forms are computable for component theories. As a direct consequence, the equality of ground $\Sigma_1 \cup \Sigma_2$-terms is decidable provided that the equality of ground Σ-terms is decidable.

Definition 131 *Given a Σ_i-theory $T_i \in \mathrm{IT}(\mathcal{A}^\Sigma)$ with initial model $\mathcal{A}_i$, let C_{T_i} be a language made of quantifier-free Σ_i-formulae of the form*

$$\psi^i \quad := \quad \bigwedge_{j \in J} \varphi_j \wedge \bigwedge_{k \in K} x_k \equiv t_k \tag{2}$$

where

- *for all $j \in J$, φ_j is an atomic Σ_i-formula such that $Sol_{\mathcal{A}_i}(\varphi_j)$ is finite;*
- *for all $k \in K$, x_k occurs only once in ψ^i and t_k is a Σ_i-term such that $Sol_{\mathcal{A}_i}(t_k \equiv s)$ is finite for all $s \in T(\Sigma)$.*

An algorithm that computes all $\mathcal{A}_i$-solutions for each φ_j occurring in a formula $\psi_i \in C_{T_i}$, and all $\mathcal{A}_i$-solutions for each equation $t_k \equiv s$ with $s \in T(\Sigma)$ is called a finitary solver over C_{T_i} in the initial model of T_i.

Where T_1, T_2 are two theories as in the definition above, we will denote by $C_{T_1} \odot C_{T_2}$ the class of quantifier-free $\Sigma_1 \cup \Sigma_2$-formulae that can be reduced via the purification procedure described in Subsection 5.1 to a conjunction $\psi_1 \wedge \psi_2$ with $\psi_1 \in C_{T_1}$ and $\psi_2 \in C_{T_2}$.

According to this definition, we have that $(C_{T_1} \odot C_{T_2})^{\Sigma_i} = C_{T_i}$, but notice that $C_{T_1} \odot C_{T_2}$ is only included in $C_{T_1 \cup T_2}$ In fact, although every formula in $C_{T_1 \cup T_2}$ is purifiable into a formula $\psi_1 \wedge \psi_2$ with each ψ_i having the form given in (2), the solutions of ψ_i need not satisfy the restrictions stated in Definition 131.

From our previous observations on how to use our combination algorithm and the results above it is easy to prove the following.

Theorem 132 *Let $T_i \in \mathrm{IT}(\mathcal{A}^\Sigma)$ for $i = 1, 2$. If there exists a finitary solver over C_{T_i} in the initial model of T_i for $i = 1, 2$, then there exists a finitary solver over $C_{T_1} \odot C_{T_2}$ in the initial model of $T_1 \cup T_2$.*

As a specific application of the results above, we consider the case in which each atomic formula φ_j in (2) in the definition of C_{T_i} is *a Σ_i-match-equation*, that is, a Σ_i-equation of the form $s \equiv t$ where t is a ground Σ_i-term. When Σ-normal forms are computable, we can assume without loss of generality that the right-hand side of each Σ_i-match-equation is in fact a ground Σ-term. Let us call a conjunction of Σ_i-match-equations, a *Σ_i-matching problem*. If the T_i-matching problem is finitary for $i = 1, 2$, each $(\Sigma_1 \cup \Sigma_2)$-matching problem is reducible by purification to a conjunction $\psi_1 \wedge \psi_2$ of pure formulae where each ψ_i verifies the requirements of Definition 131. Thus, we get another modular result for the matching problem case.

Corollary 133 *For $i = 1, 2$, let $T_i \in \mathrm{IT}(\mathcal{A}^\Sigma)$ and assume that Σ-normal-forms are computable for T_i. If there exists a finitary matching algorithm over Σ_i-matching problems in the initial model of T_i for $i = 1, 2$, then there exists a finitary matching algorithm over $\Sigma_1 \cup \Sigma_2$-matching problems in the initial model of $T_1 \cup T_2$.*

We show in the example below how to obtain in practice a finitary matching algorithm for $T_1 \cup T_2$ by combining those for T_1 and for T_2. The combined solving process is very simple, but it shows that our combination methods leads to concrete applications not only in constraint satisfiability but also in constraint solving.

Example 134 *Let T_1 be the theory axiomatized by the sentences:*[52]

$$\begin{array}{lrcl}
\forall x. & \mathsf{0} \times x & \equiv & \mathsf{0}, \\
\forall x. & x \times \mathsf{0} & \equiv & x, \\
\forall x. & \mathsf{0} + x & \equiv & x, \\
\forall x, y. & \mathsf{s}(x) + y & \equiv & \mathsf{s}(x + y), \\
\forall x, y. & \mathsf{s}(x) \times \mathsf{s}(y) & \equiv & \mathsf{s}(y + x \times \mathsf{s}(y)), \\
 & \mathsf{square}(\mathsf{0}) & \equiv & \mathsf{0} \\
\forall x. & \mathsf{square}(\mathsf{s}(x)) & \equiv & \mathsf{s}(\mathsf{square}(x) + 2 \times x).
\end{array}$$

and T_2 the theory axiomatized by the sentences:

$$\begin{array}{lrcl}
 & \mathsf{double}(\mathsf{0}) & \equiv & \mathsf{0}, \\
\forall x. & \mathsf{double}(\mathsf{s}(x)) & \equiv & \mathsf{s}(\mathsf{s}(\mathsf{double}(x))).
\end{array}$$

Where $\Sigma = \{\mathsf{0}, \mathsf{s}\}$ is the set of symbols shared by the two theories, it is easy to see that the respective initial models of T_1 and T_2 have the same Σ-reduct, which is initial in the empty Σ-theory. Also, it is possible to show, using a general method described in [DM99], that T_i-matching is finitary in the initial model of T_i for $i = 1, 2$. Now, consider for instance the $(\Sigma_1 \cup \Sigma_2)$-matching-problem

$$\mathsf{square}(x) + \mathsf{double}(y) = 5$$

in the initial model of $T_1 \cup T_2$. We start by purifying the equation into the equisatisfiable conjunction of two pure equations:

$$\mathsf{square}(x) + z = 5 \wedge z = \mathsf{double}(y)$$

The first equation is a Σ_1-matching-problem in the initial model of T_1 and has the following (Σ-)solutions:

$$x = \mathsf{0} \wedge z = 5, \quad x = 1 \wedge z = 4, \quad x = 2 \wedge z = 1.$$

After computing these solutions, we can successively propagate the shared instantiations of z in the other equation $z = \mathsf{double}(y)$, turning the equation into a Σ_2-matching-problem in the initial model of T_2. For the instantiations $\{z \leftarrow 5\}$ and $\{z \leftarrow 1\}$, the problem has no solutions in that model. For $\{z \leftarrow 4\}$, we get the solution $y = 2$.

We can conclude that the original equation has only one solution in the initial model of $T_1 \cup T_2$, namely $x = 1 \wedge y = 2$.

[52] For notational convenience, we will use the numeral n as an abbreviation for the term $\underbrace{\mathsf{s} \ldots \mathsf{s}}_{n \text{ times}} \mathsf{0}$.

11 Conclusions and Further Research

In this paper we have proposed 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 to theories with non-disjoint signatures. Finally, we have illustrated a number of applications of our combination results and related those results to some of the previous work in the combination literature.

We believe the work described in this paper 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 also appears 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 would be.

If the definition of N-O-combinability cannot be reasonably modified, the problem of finding good sufficient conditions for it still remains. 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 [GM92]) 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.[53] 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.

Acknowledgments

This work was partially supported by the National Science Foundation under grant no. 9972311.

References

[AR98] Alessandro Armando and Silvio Ranise. Constraint contextual rewriting. In R. Caferra and G. Salzer, editors, *Proceedings of the 2nd International*

[53]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 $[1, 2] + [1] \equiv 0$.

Workshop on First Order Theorem Proving, FTP'98, Vienna (Austria), pages 65–75, November 1998.

[B$\ddot{9}$4] Hans-Jürgen Bürckert. A resolution principle for constraint logics. *Artificial Intelligence*, 66:235–271, 1994.

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

[BN98] Franz Baader and Tobias Nipkow. *Term Rewriting and All That.* Cambridge University Press, United Kingdom, 1998.

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

[BS94] F. Baader and J.H. Siekmann. Unification theory. In D.M. Gabbay, C.J. Hogger, and J.A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, pages 41–125. Oxford University Press, Oxford, UK, 1994.

[BS95a] Franz Baader and Klaus U. Schulz. Combination of constraint solving techniques: An algebraic point of view. In *Proceedings of the 6th International Conference on Rewriting Techniques and Applications, RTA'95*, volume 914 of *Lecture Notes in Computer Science*, pages 50–65. Springer-Verlag, 1995.

[BS95b] Franz Baader and Klaus U. Schulz. Combination techniques and decision problems for disunification. *Theoretical Computer Science*, 142:229–255, 1995.

[BS95c] Franz Baader and Klaus U. Schulz. On the combination of symbolic constraints, solution domains, and constraint solvers. In *Proceedings of the First International Conference on Principles and Practice of Constraint Programming, Cassis (France)*, volume 976 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, September 1995.

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

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

[BT97] Franz Baader and Cesare Tinelli. A new approach for combining decision procedures for the word problem, and its connection to the Nelson-Oppen combination method. In W. McCune, editor, *Proceedings of the*

14th International Conference on Automated Deduction (Townsville, Australia), volume 1249 of *Lecture Notes in Artificial Intelligence*, pages 19–33. Springer-Verlag, 1997.

[BT98] Franz Baader and Cesare Tinelli. Deciding the word problem in the union of equational theories. Technical Report UIUCDCS-R-98-2073, Department of Computer Science, University of Illinois at Urbana-Champaign, October 1998.

[BT01] Franz Baader and Cesare Tinelli. Deciding the word problem in the union of equational theories. *Information and Computation*, 2001. (to appear).

[DJ90] Nachum Dershowitz and Jean-Pierre Jouannaud. Rewriting systems. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, pages 243–320. Elsevier Publishers, Amsterdam, 1990.

[DKR94] Eric Domenjoud, Francis Klay, and Christophe Ringeissen. Combination techniques for non-disjoint equational theories. In A. Bundy, editor, *Proceedings of the 12th International Conference on Automated Deduction, Nancy (France)*, volume 814 of *Lecture Notes in Artificial Intelligence*, pages 267–281. Springer-Verlag, 1994.

[DM99] Nachum Dershowitz and Subrata Mitra. Jeopardy. In P. Narendran and M. Rusinowitch, editors, *Proceedings of the 10th International Conference on Rewriting Techniques and Applications*, volume 1631 of *Lecture Notes in Computer Science*. Springer-Verlag, 1999.

[EM85] Hartmut Ehrig and Bernd Mahr. *Fundamentals of Algebraic Specification 1: Equations and Initial Semantics*, volume 6 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, New York, N.Y., 1985.

[EM90] Hartmut Ehrig and Bernd Mahr. *Fundamentals of Algebraic Specification 2: Module Specifications and Constraints*, volume 21 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, New York, N.Y., 1990.

[FG01] Camillo Fiorentini and Silvio Ghilardi. Combining word problems through rewriting in categories with products. *Theoretical Computer Science*, 2001. (to appear).

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

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

[Han94] Michael Hanus. The integration of functions into logic programming: From theory to practice. *Journal of Logic Programming*, 19,20:583–628, 1994.

[Her86] Alexander Herold. Combination of unification algorithms. In J. Siekmann, editor, *Proceedings 8th International Conference on Automated Deduction, Oxford (UK)*, volume 230 of *Lecture Notes in Artificial Intelligence*, pages 450–469. Springer-Verlag, 1986.

[Hod93] Wilfrid Hodges. *Model Theory*, volume 42 of *Enclyclopedia of mathematics and its applications.* Cambridge University Press, 1993.

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

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

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

[JK91] Jean-Pierre Jouannaud and Claude Kirchner. Solving equations in abstract algebras: a rule-based survey of unification. In J.-L. Lassez and G. Plotkin, editors, *Computational Logic. Essays in honor of Alan Robinson*, chapter 8, pages 257–321. MIT Press, Cambridge, MA (USA), 1991.

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

[Kep98] Stephan Kepser. *Combination of Constraint Systems.* PhD dissertation, Centre for Information and Language Processing, University of Munich, Munich, Germany, 1998.

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

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

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

[KS96] Stephan Kepser and Klaus U. Schulz. Combination of constraint systems II: Rational amalgamation. In E. C. Freuder, editor, *Proceedings of the 2nd International Conference on Principles and Practice of Constraint Programming, Cambridge, MA, USA*, volume 1118 of *Lecture Notes in Computer Science*, pages 282–296. Springer-Verlag, August 1996.

[Llo87] John W. Lloyd. *Foundations of Logic Programming.* Spinger-Verlag, Berlin, second edition, 1987.

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

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

[MH94] Aart Middeldorp and Erik Hamoen. Completeness results for basic narrowing. *Applicable Algebra in Engineering, Communication, and Computing*, 5:213–253, 1994.

[Nel84] Greg Nelson. Combining satisfiability procedures by equality-sharing. In W. W. Bledsoe and D. W. Loveland, editors, *Automated Theorem Proving: After 25 Years*, volume 29 of *Contemporary Mathematics*, pages 201–211. American Mathematical Society, Providence, RI, 1984.

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

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

[Ohl95] Enno Ohlebusch. Modular properties of composable term rewriting systems. *Journal of Symbolic Computation*, 20(1):1–41, 1995.

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

[PT97] Anand Pillay and Akito Tsuboi. Amalgamations preserving $\aleph_0$-categoricity. *The Journal of Symbolic Logic*, 62(4):1070–1074, December 1997.

[Rin92] Christophe Ringeissen. Unification in a combination of equational theories with shared constants and its application to primal algebras. In A. Voronkov, editor, *Proceedings of the 1st International Conference on Logic Programming and Automated Reasoning*, volume 624 of *Lecture Notes in Artificial Intelligence*, pages 261–272. Springer-Verlag, 1992.

[Rin93] Christophe Ringeissen. *Combinaison de Résolutions de Contraines.* Thèse de Doctorat d' Université, Université de Nancy 1, Nancy, France, December 1993.

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

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

[Sch00] Klaus U. Schulz. Why combined decision problems are often intractable. In H. Kirchner and Ch. Ringeissen, editors, *Proceedings of the 3rd International Workshop on Frontiers of Combining Systems, FroCoS'2000, Nancy (France)*, volume 1794 of *Lecture Notes in Artificial Intelligence*, pages 217–244. Springer-Verlag, March 2000.

[Sho79] Robert E. Shostak. A practical decision procedure for arithmetic with function symbols. *Journal of the ACM*, 26(2):351–360, April 1979.

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

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

[SS89] Manfred Schmidt-Schauß. Unification in a combination of arbitrary disjoint equational theories. *Journal of Symbolic Computation*, 8(1–2):51–100, July/August 1989. Special issue on unification. Part II.

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

[Tin99] Cesare Tinelli. *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, May 1999.

[TR98] Cesare Tinelli and Christophe Ringeissen. Non-disjoint unions of theories and combinations of satisfiability procedures: First results. Technical Report UIUCDCS-R-98-2044, Department of Computer Science, University of Illinois at Urbana-Champaign, April 1998. (also available as INRIA research report no. RR-3402).

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

[Wol98] Frank Wolter. Fusions of modal logics revisited. In M. Kracht, M. de Rijke, H. Wansing, and M. Zakharyaschev, editors, *Advances in Modal Logic*. CSLI, Stanford, CA, 1998.