Combination of Constraint Solving Techniques: An Algebraic Point of View

Franz Baader
Lehr- und Forschungsgebiet Theoretische Informatik, RWTH Aachen
Ahornstraße 55, 52074 Aachen, Germany
e-mail: baader@informatik.rwth-aachen.de

Klaus U. Schulz*
CIS, Universität München
80538 München, Germany
Wagmüllerstraße 23
e-mail: schulz@cis.uni-muenchen.de

Abstract

In a previous paper we have introduced a method that allows one to combine decision procedures for unifiability in disjoint equational theories. Lately, it has turned out that the prerequisite for this method to apply—namely that unification with so-called linear constant restrictions is decidable in the single theories—is equivalent to requiring decidability of the positive fragment of the first order theory of the equational theories. Thus, the combination method can also be seen as a tool for combining decision procedures for positive theories of free algebras defined by equational theories.

The present paper uses this observation as the starting point of a more abstract, algebraic approach to formulating and solving the combination problem. Its contributions are twofold. As a new result, we describe an (optimized) extension of our combination method to the case of constraint solvers that also take relational constraints (such as ordering constraints) into account. The second contribution is a new proof method, which depends on abstract notions and results from universal algebra, as opposed to technical manipulations of terms (such as ordered rewriting, abstraction functions, etc.)

*This author was supported by the EC Working Group CCL, EP6028.

Contents

1 Introduction

The integration of special inference methods (for restricted classes of problems) into general purpose deductive systems aims at a combination of the efficiency of the special method with the universality of the general method. For example, G. Plotkin [Plo72] proposed to build equational axioms into the unification algorithm used by a resolution theorem prover. This was motivated by the observation that certain axioms such as associativity and commutativity, if left unconstrained in the data base of a deductive system, may force it to go astray. Another example is the integration of special unification and matching algorithms into Knuth-Bendix completion procedures to avoid non-termination of the rewrite systems under consideration [JK86, Bac91]. In both cases, the term "unification algorithm" refers to an algorithm that computes a complete set of unifiers. With the more recent development of constraint approaches to theorem proving [Bür91] and term rewriting [KK89], the role of algorithms that compute complete sets of unifiers is more and more taken on by algorithms that decide solvability of the unification problems. In this setting, more general constraints than the equational constraints $s = t$ of unification problems become important as well. For example, one might be interested in ordering constraints of the form $s \leq t$ on terms [CT94], where the predicate $\leq$ could be interpreted as the subterm ordering or as a reduction ordering (such as the recursive path ordering).

If more than one special inference method is to be integrated into a general purpose deductive system then one must combine these special methods. In order to illustrate this, let us, for the moment, concentrate on unification problems. For example, building-in associativity and commutativity of a symbol f *and* associativity of a symbol g into a resolution theorem prover requires a unification algorithm that can handle mixed terms containing both f and g. More generally, one is thus faced with the following *combination problem:* Let E and F be equational theories over disjoint signatures, and assume that unification algorithms for E and for F are given. How can we combine these algorithms to obtain a unification algorithm for $E \cup F$.

Most of the research on combining unification algorithms was concerned with the combination of algorithms that compute complete sets of unifiers (see [SS89, Bou90] for the most recent results). In [BS92] the problem of combining decision procedures has been solved in a rather general way. The main tool of this combination method is a *decomposition algorithm*, which separates a given unification problem Γ of the joined theory (i.e., an $(E \cup F)$-

unification problem) into pure unification subproblems Γ_E and Γ_F of the single theories. Solutions of these pure problems must satisfy additional conditions, called *linear constant restrictions* in [BS92], to yield a solution of Γ. The main result of [BS92] is that solvability of unification problems in the combined theory $E \cup F$ is decidable, provided that solvability of unification problems with linear constant restrictions is decidable in E and F. It should be noted that this result can easily be lifted to solvability of $(E \cup F)$-unification problems with linear constant restrictions. This combination result has been generalized to disunification [BS93a] and to unification in the union of theories with shared constant symbols [Rin92]. The proof method used in [BS92, BS93a, Rin92]—which depends on an infinite ordered rewrite system obtained by unfailing completion, term abstraction functions, etc.—seems not to facilitate further generalizations, though (see, e.g., the rather technical "shared constructor" condition in [DKR94]). To overcome this problem, we are interested in more abstract formulations of the combination problem, which should yield a better understanding, easier proofs, and thus be a better basis for generalizations.

At first sight, the notion of "unification with linear constant restrictions" seems not to support such an abstract view: it is a technical notion that makes our combination machinery work, but seems to have little further significance. This impression is wrong, however. In [BS93] it is shown that E-unification with linear constant restrictions is decidable iff the positive fragment of the first-order theory of E is decidable. Since the positive theory of E coincides with the positive theory of the E-free Σ-algebra $\mathcal{T}(\Sigma, X)/{=_E}$ over infinitely many generators X, the combination result of [BS92] can be reformulated as follows: Let E and F be equational theories over disjoint signatures Σ and Δ, and let X be a countably infinite set of generators. The positive theory of $\mathcal{T}(\Sigma \cup \Delta, X)/{=_{E\cup F}}$ is decidable, provided that the positive theories of $\mathcal{T}(\Sigma, X)/{=_E}$ and $\mathcal{T}(\Delta, X)/{=_F}$ are decidable.

This observation can be used as the starting point of a more abstract, algebraic approach to formulating and solving the combination problem. Starting with two algebras over disjoint signatures, the goal is to construct a "combined" algebra such that validity of positive formulae in this algebra can be decided by using a decomposition algorithm and decision procedures for the positive theories of the original algebras. Obviously, this can only be achieved if the algebras satisfy some additional properties. We will call an algebra $\mathcal{A}$ *combinable* iff it is generated by a countably infinite set X such that any mapping from a finite subset of X to $\mathcal{A}$ can be extended to a surjective endomorphism of $\mathcal{A}$. For combinable algebras $\mathcal{A}$ and $\mathcal{B}$ over disjoint signatures Σ and Δ, we can define the so-called *free amalgamated product*

$\mathcal{A} \odot \mathcal{B}$, which is a $(\Sigma \cup \Delta)$-algebra.[1] Now a simple modification of the decomposition algorithm of [BS92] can be used to show that the positive theory of $\mathcal{A} \odot \mathcal{B}$ is decidable iff the positive theories of $\mathcal{A}$ and of $\mathcal{B}$ are decidable.

Obviously, the free algebras $\mathcal{T}(\Sigma, X)/{=_E}$ and $\mathcal{T}(\Delta, X)/{=_F}$ over a countably infinite set of generators X are combinable. In this case, the free amalgamated product yields an algebra that is isomorphic to the combined free algebra $\mathcal{T}(\Sigma \cup \Delta, X)/{=_{E \cup F}}$. Thus, the combination result of [BS92] is obtained as a corollary. As described until now, the amalgamation of combinable algebras does not yield a real generalization of this result. Indeed, one can use well-know results from universal algebra to show that an algebra is combinable (as defined above) iff it is a free algebra over countably many generators for an equational theory. What is new, though, is the proof method, which—in contrast to the original proof—only depends on elementary notions from universal algebra (homomorphisms, generators). This new proof can be seen as an adaptation of the proof ideas in [SS89] to the combination of decision procedures. Unlike in [SS89], however, everything is done on the abstract algebraic level instead of on the term level. Interestingly, on this level it is also very easy to prove completeness of an optimized version of the decomposition algorithm of [BS92], which significantly reduces the number of nondeterministic choices.

In addition, the abstract algebraic approach allows for an easier generalization of the results. In fact, instead of algebras we will consider algebraic structures in the following. This means that the signatures may contain both function symbols and predicate symbols, and these additional predicate symbols may occur in the constraint problems to be solved. With the usual notion of homomorphism for structures, most of the results from universal algebra carry over to structures. The combination result for combinable algebras sketched above thus holds for free structures as well. Consequently, we obtain a combination method for constraint solvers of more general constraints than just equational constraints.

The next section recalls some results from universal algebra for free structures. In Section 3 we define the free amalgamated product of free structures, and show that it again yields a free structure. Section 4 describes the decomposition algorithm and proves that it is sound and complete for existential positive input formulae. In Section 5, this result is extended to positive formulae with arbitrary quantifier prefix.

[1] This construction is similar to the one made in [SS89] for free algebras.

2 Free Structures

Let Σ be a signature consisting of a finite set Σ_F of function symbols and a finite set Σ_P of predicate symbols, where each symbol has a fixed arity. We assume that equality $=$ is an additional predicate symbol that does not occur in Σ_P. An *atomic Σ-formula* is an *equation* $s = t$ between Σ_F-terms s, t, or a *relational atomic formula* of the form $p[s_1, \ldots, s_m]$ where p is a predicate symbol in Σ_P of arity m and $s_1, \ldots, s_m$ are Σ_F-terms. A positive Σ-matrix is any Σ-formula obtained from atomic Σ-formulae using conjunction and disjunction only. A positive Σ-formula is obtained from a positive Σ-matrix by adding an arbitrary quantifier prefix, and an existential positive Σ-formula is a positive formula where the prefix consists of existential quantifier only. As usual, we shall sometimes write $t(v_1, \ldots, v_n)$ (resp. $\varphi(v_1, \ldots, v_n)$) to express that t (resp. φ) is a term (resp. formula) whose (free) variables are a subset of $\{v_1, \ldots, v_n\}$. Sentences are formulae without free variables.

A Σ-structure $\mathcal{A}$ has a non-empty carrier set A, and it interprets each $f \in \Sigma_F$ of arity n as an n-ary function $f_\mathcal{A}$ on A, and each $p \in \Sigma_P$ of arity m as an m-ary relation $p_\mathcal{A}$. The interpretation function is extended to terms and formulae as usual. If $t = t(v_1, \ldots, v_n)$ is a Σ-term, then $t_\mathcal{A}$ denotes the n-ary function on A that maps $(a_1, \ldots, a_n)$ to the value of t under the evaluation $\{v_1 \mapsto a_1, \ldots, v_n \mapsto a_n\}$. For a formula $\varphi = \varphi(v_1, \ldots, v_n)$, we write $\mathcal{A} \models \varphi(a_1, \ldots, a_n)$ to express that the formula φ is true in $\mathcal{A}$ under the evaluation $\{v_1 \mapsto a_1, \ldots, v_n \mapsto a_n\}$.

Usually, Σ-constraints are formulae of the form $\varphi(v_1, \ldots, v_n)$ with free variables. A solution of such a constraint (in a fixed Σ-structure $\mathcal{A}$) is an evaluation $\{v_1 \mapsto a_1, \ldots, v_n \mapsto a_n\}$ such that $\mathcal{A} \models \varphi(a_1, \ldots, a_n)$. Obviously, the constraint $\varphi(v_1, \ldots, v_n)$ has a solution in $\mathcal{A}$ iff the formula $\exists v_1 \ldots \exists v_n\ \varphi(v_1, \ldots, v_n)$ is valid in $\mathcal{A}$. In the present paper, we are only interested in solvability of constraints, and will thus usually take this logical point of view. In the following, tuples of variables will often be abbreviated by $\vec{v}, \vec{u}, \vec{w}$, and tuples of elements of a structure by $\vec{a}, \vec{b}$, etc.

Substructures and direct products of structures are defined in the usual way. If the Σ-substructure $\mathcal{B}$ of $\mathcal{A}$ is generated by $X \subseteq A$, we write $\mathcal{B} = \langle X \rangle_\Sigma$. Note that the carrier B of $\mathcal{B}$ consists of all elements of A that are of the form $t_\mathcal{A}(x_1, \ldots, x_n)$ where $t(v_1, \ldots, v_n)$ is a Σ-term and $x_1, \ldots, x_n$ are generators in X. In particular, any element b of B is generated by a finite subset of X.

Later on, we will consider several signatures simultaneously. If Δ is a subset of the signature Σ, then any Σ-structure $\mathcal{A}$ can be considered as a

Δ-structure (called the Δ-reduct of $\mathcal{A}$) by just forgetting about the interpretation of the additional symbols. To make clear with respect to which signature a given Σ-structure $\mathcal{A}$ is currently considered we will sometimes write $\mathcal{A}^\Sigma$ for the full Σ-structure and $\mathcal{A}^\Delta$ for its Δ-reduct.

A Σ-homomorphism is a mapping h between two Σ-structures $\mathcal{A}$ and $\mathcal{B}$ such that

$$\begin{array}{rcl} h(f_\mathcal{A}(a_1,\ldots,a_n)) & = & f_\mathcal{B}(h(a_1),\ldots,h(a_n)) \\ p_\mathcal{A}[a_1,\ldots,a_n] & \Rightarrow & p_\mathcal{B}[h(a_1),\ldots,h(a_n)] \end{array}$$

for all $f \in \Sigma_F$, $p \in \Sigma_P$, $a_1,\ldots,a_n \in A$. A Σ-isomorphism is a bijective Σ-homomorphism whose inverse is also a Σ-homomorphism.

If two homomorphisms agree on the generators of a substructure then they agree on the whole substructure. To be more precise, assume that $h_1, h_2 : \mathcal{A} \to \mathcal{B}$ are Σ-homomorphisms, and that $b \in \langle Y \rangle_\Sigma$. If h_1 and h_2 agree on Y then $h_1(b) = h_2(b)$. In particular, let $h : \mathcal{A} \to \mathcal{A}$ be a Σ-endomorphism, and assume that b is generated by Y. If h is the identity on Y then $h(b) = b$. This and the next property concerning homomorphisms and generators will be used later on. Let $h : \mathcal{A} \to \mathcal{B}$ be a Σ-homomorphism, let $b \in \langle X \rangle_\Sigma$, and let $Z \subseteq B$ be such that each element of $h(X)$ is generated by Z. Then Z is a set of generators for $h(b)$.

There is an interesting connection between surjective homomorphisms and positive formulae, which will frequently be used in proofs.

Lemma 2.1 *Let $h : \mathcal{A} \to \mathcal{B}$ be a surjective homomorphism between the Σ-structures $\mathcal{A}$ and $\mathcal{B}$, $\varphi(v_1,\ldots,v_m)$ be a positive Σ-formula, and $a_1,\ldots,a_m$ be elements of A. Then $\mathcal{A} \models \varphi(a_1,\ldots,a_m)$ implies $\mathcal{B} \models \varphi(h(a_1),\ldots,h(a_m))$.*

A proof of this lemma can be found in [Mal73], pp. 143, 144. As for the case of algebras, Σ-varieties are defined as classes of Σ-structures that are closed under direct products, substructures, and homomorphic images. The well-known Birkhoff Theorem says that a class of Σ_F-algebras is a variety iff it is an equational class, i.e., the class of models of a set of equations [Bir35]. For structures, a similar characterization is possible [Mal71].

Theorem 2.2 *A class $\mathcal{V}$ of Σ-structures is a Σ-variety if, and only if, there exists a set E of atomic Σ-formulae[2] such that $\mathcal{V}$ is the class of models of E.*

[2]As usual, open formulae are here considered as implicitly universally quantified.

In this situation, we say that $\mathcal{V}$ is the Σ-variety defined by E, and we write $\mathcal{V} = \mathcal{V}(E)$. As in the case of varieties of algebras, varieties of structures always have free objects.

Definition 2.3 *Let $\mathcal{K}$ be a class of Σ-structures, and let $\mathcal{A} \in \mathcal{K}$ be a Σ-structure that is generated by the set $X \subseteq A$. Then $\mathcal{A}$ is called* free for $\mathcal{K}$ over X *iff every mapping from X into the carrier of a Σ-structure $\mathcal{B} \in \mathcal{K}$ can be extended to a Σ-homomorphism of $\mathcal{A}$ into $\mathcal{B}$.*[3]

If $\mathcal{A}$ and $\mathcal{B}$ are free Σ-structures for the same class $\mathcal{K}$, and if their sets of generators have the same cardinality then these structures are isomorphic. Every non-trivial variety contains free structures with sets of generators of arbitrary cardinality [Mal71]. Conversely, free structures are always free for some variety.

Theorem 2.4 *Let $\mathcal{A}$ be a Σ-structure that is generated by X. Then the following conditions are equivalent:*

1. *$\mathcal{A}$ is free over X for $\{\mathcal{A}\}$.*
2. *$\mathcal{A}$ is free over X for some class $\mathcal{K}$ of Σ-structures.*
3. *$\mathcal{A}$ is free over X for some Σ-variety.*

Proof. "$1 \rightarrow 2$" and "$3 \rightarrow 1$" are trivial. In order to show "$2 \rightarrow 3$", one considers the variety generated by $\mathcal{K}$, i.e., the closure of $\mathcal{K}$ under building direct products, substructures and homomorphic images. It is easy to see that $\mathcal{A}$ is also free over X for this variety (see [Mal71, Coh65] for details). $\square$

In the following, a Σ-structure $\mathcal{A}$ will be called free (over X) iff it is free (over X) for $\{\mathcal{A}\}$. Let us now analyze how free Σ-structures look like (see [Mal71, Wea93] for more information). Obviously, the Σ_F-reduct of such a structure is a free Σ_F-algebra, and thus it is (isomorphic to) an E-free Σ_F-algebra $\mathcal{T}(\Sigma_F, X)/{=_E}$ for an equational theory E. In particular, the $=_E$-equivalence classes $[s]$ of Σ_F-terms constitute the carrier of $\mathcal{A}$. It remains to be shown how the predicate symbols are interpreted on this carrier. Since $\mathcal{A}$ is free over X, any mapping from X into $T(\Sigma_F, X)/{=_E}$ can be extended to a Σ-endomorphism of $\mathcal{A}$. This, together with the definition of homomorphisms of structures, shows that the interpretation of the predicates must

[3] Since $\mathcal{A}$ is generated by X, this homomorphism is unique.

be closed under substitution, i.e., for all $p \in \Sigma_P$, all substitutions σ, and all terms $s_1, \ldots, s_m$, if $p[[s_1], \ldots, [s_m]]$ holds in $\mathcal{A}$ then $p[[s_1\sigma], \ldots, [s_m\sigma]]$ must also hold in $\mathcal{A}$. Conversely, it is easy to see that any extension of the Σ_F-algebra $\mathcal{T}(\Sigma_F, X)/{=_E}$ to a Σ-structure that satisfies this property is a free Σ-structure over X.

Example 2.5 Let Σ_F be an arbitrary set of function symbols, and assume that Σ_P consists of a single binary predicate symbol $\leq$. Consider the (absolutely free) term algebra $\mathcal{T}(\Sigma_F, X)$. We can extend this algebra to a Σ-structure by interpreting $\leq$ as subterm ordering. Another possibility would be to take a reduction ordering [Der87] such as the lexicographic path ordering. In both cases, we have closure under substitution, which means that we obtain a free Σ-structure.

Free structures over countably infinite sets of generators are canonical for the positive theory of their variety in the following sense:

Theorem 2.6 *Let $\mathcal{A}$ be free over the countably infinite set X for a Σ-variety $\mathcal{V}(E)$, and let ϕ be a positive Σ-formula. Then the following are equivalent:*

1. *ϕ is valid in all elements of $\mathcal{V}(E)$, i.e., ϕ is a logical consequence of the set of atomic formulae E.*
2. *ϕ is valid in $\mathcal{A}$.*

Proof. Without loss of generality we may assume that ϕ is closed. "$1 \rightarrow 2$" is trivial. Now assume that $\mathcal{A} \models \phi$. Suppose that there exists an element of $\mathcal{V}(E)$ where ϕ does not hold. By the theorem of Löwenheim-Skolem there exists a countable member $\mathcal{B}$ of $\mathcal{V}(E)$ such that ϕ does not hold in $\mathcal{B}$. Obviously there exists a surjective mapping from the generators of $\mathcal{A}$ onto $\mathcal{B}$, which can be extended to a surjective homomorphism from $\mathcal{A}$ onto $\mathcal{B}$. By Lemma 2.1, we obtain a contradiction. $\square$

For the purpose of this paper, the following characterization of free structures is useful.

Lemma 2.7 *Let $\mathcal{A}$ be a Σ-structure that is generated by the countably infinite set X. Then the following conditions are equivalent:*

1. *$\mathcal{A}$ is free over X.*

2. *For every finite subset X_0 of X, every mapping $h_0 : X_0 \to A$ can be extended to a* surjective *endomorphism of $\mathcal{A}$.*

Proof. To show "$1 \to 2$," assume that $h_0 : X_0 \to A$ is given. Let $h_1 : X \setminus X_0 \to X$ be a bijection. Let h be an extension of $h_0 \dot\cup h_1$ to an endomorphism of $\mathcal{A}$. By (1), such an endomorphism exists. Since A is generated by X, h is surjective.

To show "$2 \to 1$," assume that $h_0 : X \to A$ is given. Let $X_1 \subseteq X_2 \subseteq X_3 \ldots$ be an increasing chain of subsets of X such that $X = \bigcup_{i=1}^{\infty} X_i$. For $i \geq 1$, let h_i be the restriction of h_0 to X_i. Because of (2) we know that the mappings h_i can be extended to surjective endomorphisms H_i of $\mathcal{A}$.

Let $\mathcal{A}_i$ denote the substructure of $\mathcal{A}$ generated by X_i. It is easy to see that $i < j$ implies that $\mathcal{A}_i$ is a substructure of $\mathcal{A}_j$, and that H_i and H_j coincide on $\mathcal{A}_i$. In addition, any element a of A is generated by finitely many generators, and thus there exists a least index $i(a)$ such that $a \in \mathcal{A}_{i(a)}$.

We define the mapping H_0 from A to A as the "limit" of the homomorphisms H_i; more precisely: $H_0(a) := H_{i(a)}(a)$. It remains to be shown that H_0 is a homomorphism. Thus, let f be an n-ary function symbol, and let $a_1, \ldots, a_n$ be elements of A. For $i := \max\{i(a_1), \ldots, i(a_n)\}$ we have $H_0(a_j) = H_i(a_j)$ for all $j, 1 \leq j \leq n$. In addition, since $f_{\mathcal{A}}(a_1, \ldots, a_n)$ is also in A_i we have $H_0(f_{\mathcal{A}}(a_1, \ldots, a_n)) = H_i(f_{\mathcal{A}}(a_1, \ldots, a_n))$. Since H_i is a homomorphism, we obtain $H_0(f_{\mathcal{A}}(a_1, \ldots, a_n)) = H_i(f_{\mathcal{A}}(a_1, \ldots, a_n)) = f_{\mathcal{A}}(H_i(a_1), \ldots, H_i(a_n)) = f_{\mathcal{A}}(H_0(a_1), \ldots, H_0(a_n))$. The homomorphism condition for predicates can be proved in the same way. □

Note that the second condition in the lemma is the combinability condition that was mentioned in the introduction. The lemma together with Theorem 2.4 shows that this condition holds iff the structure is a free structure for some variety.

The next lemma will be important for the proof of correctness of our combination method. It is a consequence of Lemma 2.1 and the fact that free structures satisfy the combinability condition.

Lemma 2.8 *Let $\mathcal{A}$ be a free Σ-structure over the countably infinite set of generators X, and let*

$$\gamma = \forall \vec{u}_1 \exists \vec{v}_1 \ldots \forall \vec{u}_k \exists \vec{v}_k \ \varphi(\vec{u}_1, \vec{v}_1, \ldots, \vec{u}_k, \vec{v}_k)$$

be a positive Σ-sentence. Then the following conditions are equivalent:

1. $\mathcal{A} \models \forall \vec{u}_1 \exists \vec{v}_1 \ldots \forall \vec{u}_k \exists \vec{v}_k \ \varphi(\vec{u}_1, \vec{v}_1, \ldots, \vec{u}_k, \vec{v}_k)$,

2. *There exist tuples* $\vec{x}_1 \in \vec{X}, \vec{e}_1 \in \vec{A}, \ldots, \vec{x}_k \in \vec{X}, \vec{e}_k \in \vec{A}$ *and finite subsets* $Z_1, \ldots, Z_k$ *of* X *such that*

   (a) $\mathcal{A} \models \varphi(\vec{x}_1, \vec{e}_1, \ldots, \vec{x}_k, \vec{e}_k)$,

   (b) *all generators occurring in the tuples* $\vec{x}_1, \ldots, \vec{x}_k$ *are distinct,*

   (c) *for all* $j, 1 \leq j \leq k$, *the components of* $\vec{e}_j$ *are generated by* Z_j, *i.e., are elements of* $\langle Z_j \rangle_\Sigma$, *and*

   (d) *for all* $j, 1 < j \leq k$, *no component of* $\vec{x}_j$ *occurs in* $Z_1 \cup \ldots \cup Z_{j-1}$.

Proof. "$1 \Rightarrow 2$". First, select an arbitrary tuple $\vec{x}_1$ of distinct generators from X such that this tuple has the same length as $\vec{u}_1$. Since $\mathcal{A}$ satisfies γ, there exists a tuple $\vec{e}_1 \in \vec{A}$ such that

$$(*) \quad \mathcal{A} \models \forall \vec{u}_2 \exists \vec{v}_2 \ldots \forall \vec{u}_k \exists \vec{v}_k \ \varphi(\vec{x}_1, \vec{e}_1, \vec{u}_2, \vec{v}_2, \ldots, \vec{u}_k, \vec{v}_k).$$

Let Z_1 be a finite subset of X such that each component of $\vec{e}_1$ is contained in $\langle Z_1 \rangle_\Sigma$. Obviously, such a finite set exists. Now, we may choose a finite tuple $\vec{x}_2$ of distinct generators from X such that this tuple has the same length as $\vec{u}_2$ and none of its components occurs in Z_1 or $\vec{x}_1$. This is possible because X is infinite by assumption, and Z_1 is finite.

Because of $(*)$, there exist a tuple $\vec{e}_2 \in \vec{A}$ such that

$$\mathcal{A} \models \forall \vec{u}_3 \exists \vec{v}_3 \ldots \forall \vec{u}_k \exists \vec{v}_k \ \varphi(\vec{x}_1, \vec{e}_1, \vec{x}_2, \vec{e}_2, \vec{u}_3, \vec{v}_3, \ldots, \vec{u}_k, \vec{v}_k).$$

Obviously, this argument can be iterated until Condition 2 of the lemma is proved.

"$2 \Rightarrow 1$". Let $\vec{x}_1 \in \vec{X}, \vec{e}_1 \in \vec{A}, \ldots, \vec{x}_k \in \vec{X}, \vec{e}_k \in \vec{A}$ and finite subsets $Z_1, \ldots, Z_k$ of X as in Condition 2 be given. We claim that this implies, for all $i, 0 \leq i \leq k$, the following condition C_i:

C_i: For all $\vec{a}_1 \in \vec{A}$ there exists $\vec{b}_1 \in \vec{A}$, ..., for all $\vec{a}_i \in \vec{A}$ there exists $\vec{b}_i \in \vec{A}$, and there exist $\vec{y}_{i+1}, \ldots, \vec{y}_k \in \vec{X}$, $\vec{b}_{i+1}, \ldots, \vec{b}_k \in \vec{A}$ and finite subsets $U_1, \ldots, U_i, V_1, \ldots, V_k$ of X such that

   (a') $\mathcal{A} \models \varphi(\vec{a}_1, \vec{b}_1, \ldots, \vec{a}_i, \vec{b}_i, \vec{y}_{i+1}, \vec{b}_{i+1}, \ldots, \vec{y}_k, \vec{b}_k)$,

   (b') all generators occurring in the tuples $\vec{y}_{i+1}, \ldots, \vec{y}_k$ are distinct,

   (c') for all $j, 1 \leq j \leq i$, the components of $\vec{a}_j$ are generated by U_j, and for all $j, 1 \leq j \leq k$, the components of $\vec{b}_j$ are generated by V_j, (d') for all $j, i < j \leq k$, no component of $\vec{y}_j$ occurs in $\bigcup_{\nu=1}^{j-1} V_\nu \cup \bigcup_{\mu=1}^{i} U_\mu$.

Obviously, the condition C_k is just Condition 1 of the lemma. We show that condition C_i holds for all $i, 0 \leq i \leq k$, by induction on i. For $i = 0$, validity of C_0 follows from Condition 2.

Now, assume that C_i holds for some $i, 0 \leq i < k$. To show C_{i+1}, assume that an arbitrary tuple $\vec{a}_{i+1} \in \vec{A}$ is given. Let U_{i+1} be a finite subset of X such that each component of $\vec{a}_{i+1}$ is contained in $\langle U_{i+1}\rangle_\Sigma$. For $j = i+1, ..., k$, we define a mapping h_j from a finite set of generators X_j to A by induction on j.

For $j = i + 1$, the set X_{i+1} consists of $V_{i+1} \cup \bigcup_{\nu=1}^{i}(U_\nu \cup V_\nu)$ and the components of $\vec{y}_{i+1}$. The mapping h_{i+1} leaves all elements of $\bigcup_{\nu=1}^{i}(U_\nu \cup V_\nu)$ invariant. It maps (each component of) $\vec{y}_{i+1}$ to (the corresponding component of) $\vec{a}_{i+1}$. The elements of V_{i+1} that have not yet obtained an image this way are mapped in an arbitrary way. Note that this definition of h_{i+1} is consistent because of (b') and (d') of C_i.

Now assume that X_j, h_j are already defined (for some $i + 1 \leq j < k$). The set X_{j+1} is obtained as the union of X_j with V_{j+1} and the components of $\vec{y}_{j+1}$. The mapping h_{j+1} is obtained as follows:

1. Its restriction to X_j coincides with h_j.
2. Let Y_j be a finite subset of X such that any element of $h_j(X_j)$ is contained in $\langle Y_j\rangle_\Sigma$, and let $\vec{z}_j$ be a tuple of distinct generators such that no component of $\vec{z}_j$ occurs in Y_j. (Such a tuple exists since the set of generators was assumed to be infinite, and Y_j is finite.) The mapping h_{j+1} maps (each component of) $\vec{y}_{j+1}$ to (the corresponding component of) $\vec{z}_{j+1}$.
3. The elements of V_{i+1} that have not yet obtained an image this way are mapped in an arbitrary way.

Note that Condition 1 does not conflict with Condition 2 since (b') and (d') of C_i imply that none of the components of $\vec{y}_{j+1}$ occurs in X_j.

Since $\mathcal{A}$ is free over X, and X_k is a finite subset of X, Lemma 2.7 implies that there exists a surjective endomorphism H of $\mathcal{A}$ that extends h_k. By definition of h_k, we have $H(\vec{a}_1) = \vec{a}_1$, $H(\vec{b}_1) = \vec{b}_1$, ..., $H(\vec{a}_i) = \vec{a}_i$, $H(\vec{b}_i) = \vec{b}_i$, $H(\vec{y}_{i+1}) = \vec{a}_{i+1}$, and for $i+1 < j \leq k$, $H(\vec{y}_j) = \vec{z}_j$. Thus, Lemma 2.1 implies

$$\mathcal{A} \models \varphi(\vec{a}_1, \vec{b}_1, \ldots, \vec{a}_i, \vec{b}_i, \vec{a}_{i+1}, H(\vec{b}_{i+1}), \vec{z}_{i+2}, H(\vec{b}_{i+2}), \ldots, \vec{z}_k, H(\vec{b}_k)).$$

This yields (a') of C_{i+1}. For all $j, i+1 \leq j \leq k$, the set V_j of generators of $\vec{b}_j$ is contained in X_j. In addition, any element of $H(X_j) = h_j(X_j)$ is generated by Y_j. Consequently, there exists a subset V_j' of Y_j such that all components of $H(\vec{b}_j)$ are generated by V_j'. Thus, (c') holds for $H(\vec{b}_j)$ and V_j'. It is easy to see that the mapping h_k was constructed such that (b') and (d') hold as well. □

3 Amalgamation of Free Structures

Let Σ and Δ be disjoint signatures, and let X be a countably infinite set (of generators). Let $\mathcal{A}$ be a free Σ-structure over X and and let $\mathcal{B}$ be a free Δ-structure over X. Equivalently, $\mathcal{A}$ is free over X for some Σ-variety $\mathcal{V}(E)$ and $\mathcal{B}$ is free over X for some Δ-variety $\mathcal{V}(F)$ (by Theorem 2.4). The following construction yields a $(\Sigma \cup \Delta)$-structure $\mathcal{A} \odot \mathcal{B}$ that is free over X for the $(\Sigma \cup \Delta)$-variety $\mathcal{V}(E \cup F)$.

We consider two countably infinite supersets X_∞ and Y_∞ of $X_0 := Y_0 := X$ such that $X_\infty \cap Y_\infty = X$ and $X_\infty \setminus X_0$ and $Y_\infty \setminus Y_0$ are infinite. Let $\mathcal{A}_\infty$ be free for $\mathcal{V}(E)$ over X_∞, and let $\mathcal{B}$ be free for $\mathcal{V}(F)$ over Y_∞. Obviously, $\mathcal{A}$ is the substructure of $\mathcal{A}_\infty$ that is generated by $X_0 \subseteq X_\infty$. Since both structures are free for the same variety, and since their generating sets X_0 and X_∞ have the same cardinality, $\mathcal{A}$ and $\mathcal{A}_\infty$ are isomorphic. The same holds for $\mathcal{B}$ and $\mathcal{B}_\infty$.

We shall make a zig-zag construction that defines an ascending tower of Σ-structures $\mathcal{A}_n$, and similarly an ascending tower of Δ-structures $\mathcal{B}_n$. These structures are connected by bijective mappings h_n and g_n. The free amalgamated product $\mathcal{A} \odot \mathcal{B}$ will be obtained as the limit structure, which obtains its functional and relational structure from both towers by means of the limits of the mappings h_n and g_n.

$n = 0$: Let $\mathcal{A}_0 := \mathcal{A} = \langle X_0 \rangle_\Sigma$. We interpret the "new" elements in $A_0 \setminus X_0$ as generators in $\mathcal{B}_\infty$. For this purpose, select a subset $Y_1 \subseteq Y_\infty$ such that $Y_1 \cap Y_0 = \emptyset$, $|Y_1| = |A_0 \setminus X_0|$, and the remaining complement $Y_\infty \setminus (Y_0 \cup Y_1)$ is countably infinite. Choose any bijection $h_0 : Y_0 \cup Y_1 \to A_0$ where $h_0|_{Y_0} = id_{Y_0}$.

Let $\mathcal{B}_0 := \langle Y_0 \rangle_\Delta$. As for $\mathcal{A}_0$, we interpret the "new" elements in $B_0 \setminus Y_0$ as generators in $\mathcal{A}_\infty$. Select a subset $X_1 \subseteq X_\infty$ such that $X_1 \cap X_0 = \emptyset$, $|X_1| = |B_0 \setminus Y_0|$ and the remaining complement $X_\infty \setminus (X_0 \cup X_1)$ is countably infinite. Choose any bijection $g_0 : X_0 \cup X_1 \to B_0$ where $g_0|_{X_0} = id_{X_0}$.

$n \to n+1$: Suppose that $\mathcal{A}_n = \langle \bigcup_{i=0}^n X_i \rangle_\Sigma$ and $\mathcal{B}_n = \langle \bigcup_{i=0}^n Y_i \rangle_\Delta$ are already defined, and that subsets X_{n+1} of X_∞ and Y_{n+1} of Y_∞ are already given. We assume that the complements $X_\infty \setminus \bigcup_{i=0}^{n+1} X_i$ and $Y_\infty \setminus \bigcup_{i=0}^{n+1} Y_i$ are infinite, and that the sets X_i (resp. Y_i) are pairwise disjoint. In addition, we assume that bijections

$$\begin{array}{rll} h_n: & B_{n-1} \cup Y_n \cup Y_{n+1} & \to A_n \\ g_n: & A_{n-1} \cup X_n \cup X_{n+1} & \to B_n \end{array}$$

are defined such that

$$\begin{array}{lrcl} (*) & g_n(h_n(b)) & = & b \textit{ for } b \in B_{n-1} \cup Y_n \\ & h_n(g_n(a)) & = & a \textit{ for } a \in A_{n-1} \cup X_n \\ (**) & h_n(Y_{n+1}) & = & A_n \setminus (A_{n-1} \cup X_n) \\ & g_n(X_{n+1}) & = & B_n \setminus (B_{n-1} \cup Y_n). \end{array}$$

Note that $(**)$ implies that $h_n(B_{n-1} \cup Y_n) = A_{n-1} \cup X_n$ and $g_n(A_{n-1} \cup X_n) = B_{n-1} \cup Y_n$.

We define $\mathcal{A}_{n+1} = \langle \bigcup_{i=0}^{n+1} X_i \rangle_\Sigma$ and $\mathcal{B}_{n+1} = \langle \bigcup_{i=0}^{n+1} Y_i \rangle_\Delta$, and select subsets $Y_{n+2} \subseteq Y_\infty$ and $X_{n+2} \subseteq X_\infty$ such that $Y_{n+2} \cap \bigcup_{i=0}^{n+1} Y_i = \emptyset = X_{n+2} \cap \bigcup_{i=0}^{n+1} X_i$. In addition, the cardinalities must satisfy $|Y_{n+2}| = |A_{n+1} \setminus (A_n \cup X_{n+1})|$ and $|X_{n+2}| = |B_{n+1} \setminus (B_n \cup Y_{n+1})|$, and the remaining complements $Y_\infty \setminus \bigcup_{i=0}^{n+2} Y_i$ and $X_\infty \setminus \bigcup_{i=0}^{n+2} X_i$ must be countably infinite. Let

$$\begin{array}{rcl} \upsilon_{n+1} : Y_{n+2} & \to & A_{n+1} \setminus (A_n \cup X_{n+1}), \\ \xi_{n+1} : X_{n+2} & \to & B_{n+1} \setminus (B_n \cup Y_{n+1}) \end{array}$$

be arbitrary bijections. We define $h_{n+1} := \upsilon_{n+1} \cup g_n^{-1} \cup h_n$ and $g_{n+1} := \xi_{n+1} \cup h_n^{-1} \cup g_n$. In more detail:

$$h_{n+1}(b) = \begin{cases} \upsilon_{n+1}(b) & \text{for } b \in Y_{n+2} \\ h_n(b) & \text{for } b \in B_{n-1} \cup Y_n \cup Y_{n+1} \\ g_n^{-1}(b) & \text{for } b \in B_n \setminus (B_{n-1} \cup Y_n) \end{cases}$$

and

$$g_{n+1}(a) = \begin{cases} \xi_{n+1}(a) & \text{for } a \in X_{n+2} \\ g_n(a) & \text{for } a \in A_{n-1} \cup X_n \cup X_{n+1} \\ h_n^{-1}(a) & \text{for } a \in A_n \setminus (A_{n-1} \cup X_n). \end{cases}$$

Without loss of generality we may assume (for notational convenience) that the construction eventually covers all generators in X_∞ and Y_∞; in other words, we assume that $\bigcup_{i=0}^{\infty} X_i = X_\infty$ and $\bigcup_{i=0}^{\infty} Y_i = Y_\infty$, and thus $\bigcup_{i=0}^{\infty} A_i = A_\infty$ and $\bigcup_{i=0}^{\infty} B_i = B_\infty$. We define the limit mappings

$$\begin{aligned} h_\infty &:= \bigcup_{i=0}^{\infty} h_i : B_\infty \to A_\infty, \\ g_\infty &:= \bigcup_{i=0}^{\infty} g_i : A_\infty \to B_\infty. \end{aligned}$$

It is easy to see that h_∞ and g_∞ are bijections that are inverse to each other: in fact, given $b \in B_\infty$ there is a minimal n such that $b \in B_{n-1}$. By $(*)$ it follows that $g_n(h_n(b)) = b$ and thus $g_\infty(h_\infty(b)) = b$. Accordingly, we obtain $h_\infty(g_\infty(a)) = a$ for all $a \in A_\infty$.

The bijections h_∞ and g_∞ may be used to carry the Δ-structure of $\mathcal{B}_\infty$ to $\mathcal{A}_\infty$ and to carry the Σ-structure of $\mathcal{A}_\infty$ to $\mathcal{B}_\infty$: let f (f') be an n-ary function symbol of Δ (Σ) and $a_1, \ldots, a_n \in A_\infty$ $(b_1, \ldots, b_n \in B_\infty)$. We define

$$\begin{aligned} f_{\mathcal{A}_\infty}(a_1, \ldots, a_n) &:= h_\infty(f_{\mathcal{B}_\infty}(g_\infty(a_1), \ldots, g_\infty(a_n))), \\ f'_{\mathcal{B}_\infty}(b_1, \ldots, b_n) &:= g_\infty(f'_{\mathcal{A}_\infty}(h_\infty(b_1), \ldots, h_\infty(b_n))). \end{aligned}$$

Let p (q) be an n-ary predicate symbol of Δ (Σ) and $a_1, \ldots, a_n \in A_\infty$ $(b_1, \ldots, b_n \in B_\infty)$. We define

$$\begin{aligned} p_{\mathcal{A}_\infty}[a_1, \ldots, a_n] &:\Longleftrightarrow p_{\mathcal{B}_\infty}[g_\infty(a_1), \ldots, g_\infty(a_n)], \\ q_{\mathcal{B}_\infty}[b_1, \ldots, b_n] &:\Longleftrightarrow q_{\mathcal{A}_\infty}[h_\infty(b_1), \ldots, h_\infty(b_n)]. \end{aligned}$$

With this definition, the mappings h_∞ and g_∞ are inverse isomorphisms between the $(\Sigma \cup \Delta)$-structures $\mathcal{A}_\infty$ and $\mathcal{B}_\infty$. Identifying isomorphic structures, we call $\mathcal{A}_\infty^{\Sigma\cup\Delta} \simeq \mathcal{B}_\infty^{\Sigma\cup\Delta}$ the *free amalgamated product* $\mathcal{A} \odot \mathcal{B}$ *of* $\mathcal{A}$ *and* $\mathcal{B}$. As a Σ-structure, $\mathcal{A} \odot \mathcal{B}$ is isomorphic to $\mathcal{A}$, which is free over X for $\mathcal{V}(E)$, and as a Δ-structure it is isomorphic to $\mathcal{B}$, which is free over X for $\mathcal{V}(F)$. Now, we show that as a $(\Sigma \cup \Delta)$-structure it is free over X for $\mathcal{V}(E \cup F)$. First, let us show that it is generated by X.

Lemma 3.1 *As a $(\Sigma \cup \Delta)$-structure, the free amalgamated product $\mathcal{A} \odot \mathcal{B}$ is generated by X, the set of generators of both $\mathcal{A}$ and $\mathcal{B}$.*

Proof. Obviously, $\langle X \rangle_{\Sigma\cup\Delta}$ is a $(\Sigma \cup \Delta)$-substructure of $\mathcal{A}_\infty$. To prove the other direction, assume that a is an element of A_∞, i.e, $a \in A_n$ for some $n \geq 0$. We show $a \in \langle X \rangle_{\Sigma\cup\Delta}$ by induction on n. For $n = 0$ we have $a \in A_0 = \langle X \rangle_\Sigma \subseteq \langle X \rangle_{\Sigma\cup\Delta}$.

Now, assume that $n > 0$. We distinguish three cases: First, let $a \in A_{n-1}$. We obtain $a \in \langle X\rangle_{\Sigma\cup\Delta}$ by the induction hypothesis. Second, assume that $a \in X_n$. Then we have $g_{n-1}(a) \in B_{n-1} \setminus (B_{n-2} \cup Y_{n-1})$. This means that there are a Δ-term t (with m different variables), and elements $b_1, \ldots, b_m$ of $B_{n-2} \cup Y_{n-1}$ such that $g_\infty(a) = g_{n-1}(a) = t_{\mathcal{B}_\infty}(b_1, \ldots, b_m)$. Thus, we have

$$a = h_\infty(g_\infty(a)) = h_\infty(t_{\mathcal{B}_\infty}(b_1, \ldots, b_m)) = t_{\mathcal{A}_\infty}(h_\infty(b_1), \ldots, h_\infty(b_m)),$$

since h_∞ is a $(\Sigma \cup \Delta)$-homomorphism. Since $h_\infty(b_1), \ldots, h_\infty(b_m) \in A_{n-1}$, the induction hypothesis yields $h_\infty(b_1), \ldots, h_\infty(b_m) \in \langle X\rangle_{\Sigma\cup\Delta}$, which shows $a = t_{\mathcal{A}_\infty}(h_\infty(b_1), \ldots, h_\infty(b_m)) \in \langle X\rangle_{\Sigma\cup\Delta}$.

Finally, assume that $a \in A_n \setminus (A_{n-1} \cup X_n)$. Thus, there are a Σ-term s (with m different variables), and elements $a_1, \ldots, a_m$ of $A_{n-1} \cup X_n$ such that $a = s_{\mathcal{A}_\infty}(a_1, \ldots, a_m)$. By induction and by what we have shown above, we know that $a_1, \ldots, a_m \in \langle X\rangle_{\Sigma\cup\Delta}$, which yields $a = s_{\mathcal{A}_\infty}(a_1, \ldots, a_m) \in \langle X\rangle_{\Sigma\cup\Delta}$.
□

Theorem 3.2 *$A \odot B$ is free over X for the $(\Sigma \cup \Delta)$-variety $\mathcal{V}(E \cup F)$.*

Proof. Let $\mathcal{C}$ be a $(\Sigma\cup\Delta)$-structure that satisfies $E\cup F$, i.e., $\mathcal{C} \in \mathcal{V}(E\cup F)$, and let $\varphi : X \to C$ be a mapping.

Because of Lemma 3.1, it remains to be shown that φ can always be extended to a $(\Sigma \cup \Delta)$-homomorphism Φ of $\mathcal{A}_\infty$ to $\mathcal{C}$. For this purpose, we define for all $n \geq 0$ mappings

$$\begin{array}{rcl} \Phi_{\Sigma,n} : A_n & \to & C \\ \Phi_{\Delta,n} : B_n & \to & C \end{array}$$

that satisfy the following properties:

1. $\Phi_{\Sigma,n}$ is a Σ-homomorphism and $\Phi_{\Delta,n}$ is a Δ-homomorphism.
2. If $n > 0$ then, for all $x \in \bigcup_{i=1}^{n} X_i$,

   $$\Phi_{\Sigma,n}(x) = \Phi_{\Delta,n-1}(g_\infty(x)),$$

   and, for all $y \in \bigcup_{i=1}^{n} Y_i$,

   $$\Phi_{\Delta,n}(y) = \Phi_{\Sigma,n-1}(h_\infty(y)).$$

3. If $n > 0$ then the restriction of $\Phi_{\Sigma,n}$ to A_{n-1} yields $\Phi_{\Sigma,n-1}$ and the restriction of $\Phi_{\Delta,n}$ to B_{n-1} yields $\Phi_{\Delta,n-1}$.

4. For all $x \in X$, $\Phi_{\Sigma,n}(x) = \varphi(x) = \Phi_{\Delta,n}(x)$.

$n = 0$: Obviously, $\mathcal{C}$ can be considered as a Σ-structure, and this Σ-structure belongs to the Σ-variety $\mathcal{V}(E)$. Since $\mathcal{A} = \mathcal{A}_0$ is free over $X = X_0$ for $\mathcal{V}(E)$, the mapping $\varphi : X \to C$ can be extended to a Σ-homomorphism from $\mathcal{A}_0$ to $\mathcal{C}$. We call this homomorphism $\Phi_{\Sigma,0}$. Condition 1 is thus trivially satisfied. For Conditions 2 and 3 there is nothing to show since $n = 0$. Condition 4 is satisfied by definition of $\Phi_{\Sigma,0}$. The Δ-homomorphism $\Phi_{\Delta,0} : \mathcal{B}_0 \to \mathcal{C}$ is defined analogously.

$n \to n+1$: Assume that mappings $\Phi_{\Sigma,n}$ and $\Phi_{\Delta,n}$ satisfying Conditions 1–4 are given. We define mappings $\varphi_{\Sigma,n+1} : \bigcup_{i=0}^{n+1} X_i \to C$ and $\varphi_{\Delta,n+1} : \bigcup_{i=0}^{n+1} Y_i \to C$ by

$$\begin{aligned}
\varphi_{\Sigma,n+1}(x) &= \begin{cases} \Phi_{\Delta,n}(g_\infty(x)) & \text{if } x \in X_{n+1} \\ \Phi_{\Sigma,n}(x) & \text{else,} \end{cases} \\
\varphi_{\Delta,n+1}(y) &= \begin{cases} \Phi_{\Sigma,n}(h_\infty(y)) & \text{if } y \in Y_{n+1} \\ \Phi_{\Delta,n}(y) & \text{else.} \end{cases}
\end{aligned}$$

Let $\Phi_{\Sigma,n+1}$ be the extension of $\varphi_{\Sigma,n+1}$ to a Σ-homomorphism of $\mathcal{A}_{n+1}$ to $\mathcal{C}$, and let $\Phi_{\Delta,n+1}$ be the extension of $\varphi_{\Delta,n+1}$ to a Δ-homomorphism of $\mathcal{B}_{n+1}$ to $\mathcal{C}$. We must show that the four conditions are again satisfied.

1. Condition 1 is trivially satisfied.

2. We proof Condition 2 for $\Phi_{\Sigma,n+1}$. For $x \in X_{n+1}$, the condition is satisfied by definition of $\varphi_{\Sigma,n+1}(x)$. For $x \in \bigcup_{i=0}^{n} X_i$ we have $\Phi_{\Sigma,n+1}(x) = \varphi_{\Sigma,n+1}(x) = \Phi_{\Sigma,n}(x)$. We know $\Phi_{\Sigma,n}(x) = \Phi_{\Delta,n-1}(g_\infty(x))$ by assumption. Looking back at the definition of the free amalgamated product, we see that $g_\infty(x)$ is an element of B_{n-1}. By assumption, we know that $\Phi_{\Delta,n-1}$ and $\Phi_{\Delta,n}$ agree on B_{n-1}.

3. By definition, $\Phi_{\Sigma,n}$ and $\Phi_{\Sigma,n+1}$ agree on the generators $\bigcup_{i=1}^{n} X_i$ of $\mathcal{A}_n$.

4. For $x \in X = X_0$, we have $\Phi_{\Sigma,n+1}(x) = \Phi_{\Sigma,n}(x) = \varphi(x)$ by assumption.

This completes the construction of the mappings $\Phi_{\Sigma,n}$ and $\Phi_{\Delta,n}$ ($n \geq 0$). The mappings $\Phi_\Sigma : A_\infty \to C$ and $\Phi_\Delta : B_\infty \to C$ are the limits of these ascending chains of mappings. More precisely, we define mappings $\varphi_\Sigma : \bigcup_{i=0}^{\infty} X_i \to C$ and $\varphi_\Delta : \bigcup_{i=0}^{\infty} Y_i \to C$ by

$$\begin{aligned}
\varphi_\Sigma(x) &:= \Phi_{\Sigma,n}(x) \text{ for } x \in X_n \\
\varphi_\Delta(y) &:= \Phi_{\Delta,n}(y) \text{ for } y \in Y_n
\end{aligned}$$

Now Φ_Σ is the extension of φ_Σ to a Σ-homomorphism of $\mathcal{A}_\infty$ to $\mathcal{C}$, and Φ_Δ is the extension of φ_Δ to a Δ-homomorphism of $\mathcal{B}_\infty$ to $\mathcal{C}$. Since $\mathcal{A}_n$ (resp. $\mathcal{B}_n$) is generated by $\bigcup_{i=0}^n X_i$ (resp. $\bigcup_{i=0}^n Y_i$), the restriction of Φ_Σ to A_n coincides with $\Phi_{\Sigma,n}$ (resp. the restriction of Φ_Δ to B_n coincides with $\Phi_{\Delta,n}$).

By construction Φ_Σ is a Σ-homomorphism and Φ_Δ is a Δ-homomorphism. It remains to be shown that they are $(\Sigma \cup \Delta)$-homomorphisms. In order to show this we prove the following claim:

$$(*) \quad \Phi_\Sigma \circ h_\infty = \Phi_\Delta \quad \text{and} \quad \Phi_\Delta \circ g_\infty = \Phi_\Sigma.[4]$$

From the second identity of $(*)$ we can easily deduce that Φ_Σ is a $(\Sigma \cup \Delta)$-homomorphism. In fact, we already know that it is a Σ-homomorphism. In addition, Φ_Δ is a Δ-homomorphism and g_∞ is a $(\Sigma \cup \Delta)$-homomorphism. Thus the composition $\Phi_\Delta \circ g_\infty$ is a Δ-homomorphism. Accordingly, the first identity of $(*)$ implies that Φ_Δ is a $(\Sigma \cup \Delta)$-homomorphism.

To complete the proof, we show the first identity of $(*)$. (The second follows by symmetry.) Let b be an element of B_∞. Thus there is an $n \geq 0$ such that $b \in B_n \setminus B_{n-1}$. First, assume that $b \in Y_n$. By construction of the free amalgamated product, this implies $h_\infty(b) \in A_{n-1}$, and thus we have

$$\Phi_\Sigma(h_\infty(b)) = \Phi_{\Sigma,n-1}(h_\infty(b)) = \Phi_{\Delta,n}(b) = \Phi_\Delta(b).$$

The second identity holds by Condition 2 in the construction of the mappings $\Phi_{\Delta,n}$ and $\Phi_{\Sigma,n}$ since B_n is generated by $\bigcup_{i=0}^n Y_i$. The third identity follows from the definition of Φ_Δ.

Second, assume that $b \in B_n \setminus (B_{n-1} \cup Y_n)$. In this case we have $h_\infty(b) = g_\infty^{-1}(b) \in X_{n+1}$, and thus $\Phi_\Sigma(h_\infty(b)) = \Phi_{\Sigma,n+1}(g_\infty^{-1}(b)) = \Phi_{\Delta,n}(g_\infty(g_\infty^{-1}(b))) = \Phi_{\Delta,n}(b) = \Phi_\Delta(b)$. □

Free structures for the same variety and over generator sets of equal cardinality are isomorphic. This, together with the above theorem yields:

Corollary 3.3 *Let $\mathcal{A}_1$ be free over X in the Σ_1-variety $\mathcal{V}(E_1)$, $\mathcal{A}_2$ be free over X in the Σ_2-variety $\mathcal{V}(E_2)$, and $\mathcal{A}_3$ be free over X in the Σ_3-variety $\mathcal{V}(E_3)$, where X is countably infinite and Σ_1, Σ_2, and Σ_3 are pairwise disjoint. Then $(\mathcal{A}_1 \odot \mathcal{A}_2) \odot \mathcal{A}_3 \simeq \mathcal{A}_1 \odot (\mathcal{A}_2 \odot \mathcal{A}_3)$.*

Proof. By the theorem, both structures are free over X for the $(\Sigma_1 \cup \Sigma_2 \cup \Sigma_3)$-variety $\mathcal{V}(E_1 \cup E_2 \cup E_3)$. □

[4] Here, composition should be read from right to left, i.e., apply first h_∞ and then Φ_Σ.

4 The Decomposition Algorithm

As in the previous section, let $\mathcal{V}(E)$ be a Σ-variety and $\mathcal{V}(F)$ be a Δ-variety, where Σ and Δ are disjoint signatures. For a countably infinite set of generators X, let $\mathcal{A}$ be free for $\mathcal{V}(E)$ over X, and let $\mathcal{B}$ be free for $\mathcal{V}(F)$ over X. We know that the positive theories of $\mathcal{V}(E)$ and $\mathcal{A}$ (resp. $\mathcal{V}(F)$ and $\mathcal{B}$) coincide (by Theorem 2.6), and that the free amalgamated product $\mathcal{A} \odot \mathcal{B}$ is free for $\mathcal{V}(E \cup F)$ over X (by Theorem 3.2).

In this section, we consider only existential positive $(\Sigma \cup \Delta)$-sentences. The decomposition algorithm described below can be used to reduce validity of such sentences in $\mathcal{A} \odot \mathcal{B}$ (or, equivalently, in $\mathcal{V}(E \cup F)$) to validity of positive sentences in $\mathcal{A}$ and in $\mathcal{B}$.

Before we can describe the algorithm, we must introduce some notation. In the following, V denotes an infinite set of variables used by the first order languages under consideration. Let t be a $(\Sigma \cup \Delta)$-term. This term is called *pure* iff it is either a Σ-term or a Δ-term. An equation is pure iff it is an equation between pure terms of the same signature. A relational formula $p[s_1, \ldots, s_m]$ is pure iff $s_1, \ldots, s_m$ are pure terms of the signature of p. Now assume that t is a non-pure term whose topmost function symbol is in Σ. A subterm s of t is called *alien subterm* of t iff its topmost function symbol belongs to Δ and every proper superterm of s in t has its top symbol in Σ. Alien subterms of terms with top symbol in Δ are defined analogously. For a relational formula $p[s_1, \ldots, s_m]$, alien subterms are defined as follows: if s_i has a top symbol whose signature is different from the signature of p then s_i itself is an alien subterm; otherwise, any alien subterm of s_i is an alien subterm of $p[s_1, \ldots, s_m]$.

Algorithm 1

Let φ_0 be a positive existential $(\Sigma \cup \Delta)$-sentence. Without loss of generality, we may assume that φ_0 has the form $\exists \vec{u}_0\ \gamma_0$, where γ_0 is a conjunction of atomic formulae. Indeed, since existential quantifiers distribute over disjunction, a sentence $\exists \vec{u}_0\ (\gamma_1 \vee \gamma_2)$ is valid iff $\exists \vec{u}_0\ \gamma_1$ or $\exists \vec{u}_0\ \gamma_2$ is valid.

Step 1: Transform non-pure atomic formulae.

(1) Equations $s = t$ of γ_0 where s and t have topmost function symbols belonging to different signatures are replaced by (the conjunction of) two new equations $u = s, u = t$, where u is a new variable. The

quantifier prefix is extended by adding an existential quantification for u.

(2) As a result, we may assign a unique label Σ or Δ to each atomic formula that is not an equation between variables. The label of an equation $s = t$ is the signature of the topmost function symbols of s and/or t. The label of a relational formula $p[s_1, \ldots, s_m]$ is the signature of p.

(3) Now alien subterms occurring in atomic formulae are successively replaced by new variables. For example, assume that $s = t$ is an equation in the current formula, and that s contains the alien subterm s_1. Let u be a variable not occurring in the current formula, and let s' be the term obtained from s by replacing s_1 by u. Then the original equation is replaced by (the conjunction of) the two equations $s' = t$ and $u = s_1$. The quantifier prefix is extended by adding an existential quantification for u. The equation $s' = t$ keeps the label of $s = t$, and the label of $u = s_1$ is the signature of the top symbol of s_1. Relational atomic formulae with alien subterms are treated analogously. This process is iterated until all atomic formulae occurring in the conjunctive matrix are pure. It is easy to see that this is achieved after finitely many iterations.

Step 2: Remove atomic formulae without label.

Equations between variables occurring in the conjunctive matrix are removed as follows: If $u = v$ is such an equation then one removes $\exists u$ from the quantifier prefix and $u = v$ from the matrix. In addition, every occurrence of u in the remaining matrix is replaced by v. This step is iterated until the matrix contains no equations between variables.

Let φ_1 be the new sentence obtained this way. The matrix of φ_1 can be written as a conjunction $\gamma_{1,\Sigma} \wedge \gamma_{1,\Delta}$, where $\gamma_{1,\Sigma}$ is a conjunction of all atomic formulae from φ_1 with label Σ, and $\gamma_{1,\Delta}$ is a conjunction of all atomic formulae from φ_1 with label Δ. There are three different types of variables occurring in φ_1: shared variables occur both in $\gamma_{1,\Sigma}$ and in $\gamma_{1,\Delta}$; Σ-variables occur only in $\gamma_{1,\Sigma}$; and Δ-variables occur only in $\gamma_{1,\Delta}$. Let $\vec{u}_{1,\Sigma}$ be the tuple of all Σ-variables, $\vec{u}_{1,\Delta}$ be the tuple of all Δ-variables, and $\vec{u}_1$ be the tuple of all shared variables.[5] Obviously, φ_1 is equivalent to the sentence

$$\exists \vec{u}_1 \, (\exists \vec{u}_{1,\Sigma} \; \gamma_{1,\Sigma} \wedge \exists \vec{u}_{1,\Delta} \; \gamma_{1,\Delta}) \, .$$

The next two steps of the algorithm are nondeterministic, i.e., a given sentence is transformed into finitely many new sentences. Here the idea is

[5] The order in these tuples can be chosen arbitrarily.

that the original sentence is valid iff at least one of the new sentences is valid.

Step 3: Variable identification.
Consider all possible partitions of the set of all shared variables. Each of these partitions yields one of the new sentences as follows. The variables in each class of the partition are "identified" with each other by choosing an element of the class as representative, and replacing in the sentence all occurrences of variables of the class by this representative. Quantifiers for replaced variables are removed.

Let $\exists \vec{u}_2\, (\exists \vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma} \wedge \exists \vec{u}_{1,\Delta}\ \gamma_{2,\Delta})$ denote one of the sentences obtained by Step 3.

Step 4: Choose signature labels and ordering.
We choose a label Σ or Δ for every (shared) variable in $\vec{u}_2$, and a linear ordering $<$ on these variables.

For each of the choices made in Step 3 and 4, the algorithm yields a pair (α, β) of sentences as output.

Step 5: Generate output sentences.
The sentence $\exists \vec{u}_2(\exists \vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma} \wedge \exists \vec{u}_{1,\Delta}\ \gamma_{2,\Delta})$ is split into two sentences

$$\alpha = \forall \vec{v}_1 \exists \vec{w}_1 \dots \forall \vec{v}_k \exists \vec{w}_k \exists \vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma}$$

and

$$\beta = \exists \vec{v}_1 \forall \vec{w}_1 \dots \exists \vec{v}_k \forall \vec{w}_k \exists \vec{u}_{1,\Delta}\ \gamma_{2,\Delta}.$$

Here $\vec{v}_1\vec{w}_1 \dots \vec{v}_k\vec{w}_k$ is the unique re-ordering of $\vec{u}_2$ along $<$. The variables $\vec{v}_i$ ($\vec{w}_i$) are the variables with label Δ (label Σ).

Thus, the overall output of the algorithm is a finite set of pairs of sentences. Note that the sentences α and β are positive formulae, but they need no longer be existential positive formulae.

Algorithm 1 is an optimization of the decomposition algorithm described in [BS92] in that the nondeterministic steps—which are responsible for the NP-complexity of the algorithm—are applied only to shared variables and not to all variables occurring in the system. This may drastically decrease the number of nondeterministic choices. For example, if there are no shared variables then the algorithm is completely deterministic, and the output formulae are existential positive formulae.

Correctness of Algorithm 1

First, we show soundness of the algorithm, i.e., if one of the output pairs is valid then the original sentence was valid.

Lemma 4.1 *$\mathcal{A} \odot \mathcal{B} \models \varphi_0$ if $\mathcal{A} \models \alpha$ and $\mathcal{B} \models \beta$ for some output pair (α, β).*

Proof. Since $\mathcal{A}^\Sigma$ and $\mathcal{A}^\Sigma_\infty$ are isomorphic Σ-structures, we know that $\mathcal{A}^\Sigma_\infty \models \alpha$. Accordingly, we also have $\mathcal{B}^\Delta_\infty \models \beta$. More precisely, this means

$$\begin{array}{rlcl} (*) & \mathcal{A}^\Sigma_\infty & \models & \forall \vec{v}_1 \exists \vec{w}_1 \ldots \forall \vec{v}_k \exists \vec{w}_k \exists \vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma}(\vec{v}_1, \vec{w}_1, \ldots, \vec{v}_k, \vec{w}_k, \vec{u}_{1,\Sigma}) \\ (**) & \mathcal{B}^\Delta_\infty & \models & \exists \vec{v}_1 \forall \vec{w}_1 \ldots \exists \vec{v}_k \forall \vec{w}_k \exists \vec{u}_{1,\Delta}\ \gamma_{2,\Delta}(\vec{v}_1, \vec{w}_1, \ldots, \vec{v}_k, \vec{w}_k, \vec{u}_{1,\Delta}). \end{array}$$

Because of the existential quantification over $\vec{v}_1$ in $(**)$, there exist elements $\vec{b}_1 \in \vec{B}_\infty$ such that

$$(***)\quad \mathcal{B}^\Delta_\infty \models \forall \vec{w}_1 \ldots \exists \vec{v}_k \forall \vec{w}_k \exists \vec{u}_{1,\Delta}\ \gamma_{2,\Delta}(\vec{b}_1, \vec{w}_1, \ldots, \vec{v}_k, \vec{w}_k, \vec{u}_{1,\Delta}).$$

We consider $\vec{a}_1 := h_\infty(\vec{b}_1)$. Because of the universal quantification over $\vec{v}_1$ in $(*)$ we have

$$\mathcal{A}^\Sigma_\infty \models \exists \vec{w}_1 \ldots \forall \vec{v}_k \exists \vec{w}_k \exists \vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma}(\vec{a}_1, \vec{w}_1, \ldots, \vec{v}_k, \vec{w}_k, \vec{u}_{1,\Sigma}).$$

Because of the existential quantification over $\vec{w}_1$ in this formula there exist elements $\vec{c}_1 \in \vec{A}_\infty$ such that

$$\mathcal{A}^\Sigma_\infty \models \forall \vec{v}_2 \exists \vec{w}_2 \ldots \forall \vec{v}_k \exists \vec{w}_k \exists \vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma}(\vec{a}_1, \vec{c}_1, \vec{v}_2, \vec{w}_2, \ldots, \vec{v}_k, \vec{w}_k, \vec{u}_{1,\Sigma}).$$

We consider $\vec{d}_1 := g_\infty(\vec{c}_1)$. Because of the universal quantification over $\vec{w}_1$ in $(***)$ we have

$$\mathcal{B}^\Delta_\infty \models \exists \vec{v}_2 \forall \vec{w}_2 \ldots \exists \vec{v}_k \forall \vec{w}_k \exists \vec{u}_{1,\Delta}\ \gamma_{2,\Delta}(\vec{b}_1, \vec{d}_1, \vec{v}_2, \vec{w}_2, \ldots, \vec{v}_k, \vec{w}_k, \vec{u}_{1,\Delta}).$$

Iterating this argument, we thus obtain

$$\begin{array}{rcl} \mathcal{A}^\Sigma_\infty & \models & \exists \vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma}(\vec{a}_1, \vec{c}_1, \ldots, \vec{a}_k, \vec{c}_k, \vec{u}_{1,\Sigma}), \\ \mathcal{B}^\Delta_\infty & \models & \exists \vec{u}_{1,\Delta}\ \gamma_{2,\Delta}(\vec{b}_1, \vec{d}_1, \ldots, \vec{b}_k, \vec{d}_k, \vec{u}_{1,\Delta}), \end{array}$$

where $\vec{a}_i = h_\infty(\vec{b}_i)$ and $\vec{d}_i = g_\infty(\vec{c}_i)$ (for $1 \leq i \leq k$). Since h_∞ is a $(\Sigma \cup \Delta)$-isomorphism that is the inverse of g_∞, we also know that

$$\mathcal{A}^\Delta_\infty \models \exists \vec{u}_{1,\Delta}\ \gamma_{2,\Delta}(\vec{a}_1, \vec{c}_1, \ldots, \vec{a}_k, \vec{c}_k, \vec{u}_{1,\Delta}).$$

It follows that

$$\mathcal{A}_\infty^{\Sigma\cup\Delta} \models \exists\vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma}(\vec{a}_1,\vec{c}_1,\ldots,\vec{a}_k,\vec{c}_k,\vec{u}_{1,\Sigma}) \wedge \exists\vec{u}_{1,\Delta}\ \gamma_{2,\Delta}(\vec{a}_1,\vec{c}_1,\ldots,\vec{a}_k,\vec{c}_k,\vec{u}_{1,\Delta}).$$

Obviously, this implies that

$$\mathcal{A}\odot\mathcal{B} \simeq \mathcal{A}_\infty^{\Sigma\cup\Delta} \models \exists\vec{u}_2\,(\exists\vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma} \wedge \exists\vec{u}_{1,\Delta}\ \gamma_{2,\Delta}),$$

i.e., one of the sentences obtained after Step 3 of the algorithm holds in $\mathcal{A}\odot\mathcal{B}$. It is easy to see that this implies that $\mathcal{A}\odot\mathcal{B}\models\varphi_0$. □

Next, we show completeness of the decomposition algorithm, i.e., if the input sentence was valid then there exists a valid output pair.

Lemma 4.2 *If $\mathcal{A}\odot\mathcal{B}\models\varphi_0$ then $\mathcal{A}\models\alpha$ and $\mathcal{B}\models\beta$ for some output pair (α,β).*

Proof. Assume that $\mathcal{A}\odot\mathcal{B}\simeq\mathcal{B}_\infty^{\Sigma\cup\Delta}\models\exists\vec{u}_0\gamma_0$. Obviously, this implies that $\mathcal{B}_\infty^{\Sigma\cup\Delta}\models\exists\vec{u}_1\,(\exists\vec{u}_{1,\Sigma}\ \gamma_{1,\Sigma}(\vec{u}_1,\vec{u}_{1,\Sigma})\wedge\exists\vec{u}_{1,\Delta}\ \gamma_{1,\Delta}(\vec{u}_1,\vec{u}_{1,\Delta}))$, i.e., $\mathcal{B}_\infty^{\Sigma\cup\Delta}$ satisfies the sentence that is obtained after Step 2 of the decomposition algorithm. Thus there exists an assignment $\nu: V\rightarrow B_\infty$ such that $\mathcal{B}_\infty^{\Sigma\cup\Delta}\models\exists\vec{u}_{1,\Sigma}\ \gamma_{1,\Sigma}(\nu(\vec{u}_1),\vec{u}_{1,\Sigma})\wedge\exists\vec{u}_{1,\Delta}\ \gamma_{1,\Delta}(\nu(\vec{u}_1),\vec{u}_{1,\Delta})$.

In Step 3 of the decomposition algorithm we identify two shared variables u and u' of $\vec{u}_1$ if, and only if, $\nu(u)=\nu(u')$. With this choice, $\mathcal{B}_\infty^{\Sigma\cup\Delta}\models\exists\vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma}(\nu(\vec{u}_2),\vec{u}_{1,\Sigma})\wedge\exists\vec{u}_{1,\Delta}\ \gamma_{2,\Delta}(\nu(\vec{u}_2),\vec{u}_{1,\Delta})$, and all components of $\nu(\vec{u}_2)$ are distinct.

In Step 4, a shared variable u in $\vec{u}_2$ is labeled with Δ if $\nu(u)\in B_\infty\setminus(\bigcup_{i=1}^\infty Y_i)$, and with Σ otherwise. In order to choose the linear ordering on the shared variables, we partition the range B_∞ of ν as follows:

$$B_0,\ \ Y_1,\ \ B_1\setminus(B_0\cup Y_1),\ \ Y_2,\ \ B_2\setminus(B_1\cup Y_2),\ \ Y_3,\ \ B_3\setminus(B_2\cup Y_3),\ \ldots$$

Now, let $\vec{v}_1,\vec{w}_1,\ldots,\vec{v}_k,\vec{w}_k$ be a re-ordering of the tuple $\vec{u}_2$ such that the following holds:

1. The tuple $\vec{v}_1$ contains exactly the shared variables whose ν-images are in B_0.
2. For all $i, 1\leq i\leq k$, the tuple $\vec{w}_i$ contains exactly the shared variables whose ν-images are in Y_i.

3. For all $i, 1 < i \leq k$, the tuple $\vec{v}_i$ contains exactly the shared variables whose ν-images are in $B_{i-1} \setminus (B_{i-2} \cup Y_{i-1})$.

Obviously, this implies that the variables in the tuples $\vec{w}_i$ have label Σ, whereas the variables in the tuples $\vec{v}_i$ have label Δ. Note that some of these tuples may be of dimension 0. The re-ordering determines the linear ordering we choose in Step 4. Let

$$\begin{aligned}\alpha &= \forall \vec{v}_1 \exists \vec{w}_1 \ldots \forall \vec{v}_k \exists \vec{w}_k \exists \vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma} \\ \beta &= \exists \vec{v}_1 \forall \vec{w}_1 \ldots \exists \vec{v}_k \forall \vec{w}_k \exists \vec{u}_{1,\Delta}\ \gamma_{2,\Delta}\end{aligned}$$

be the output pair that is obtained by these choices. Let $\vec{y}_i := \nu(\vec{w}_i) \in \vec{Y}$ and $\vec{b}_i := \nu(\vec{v}_i) \in \vec{B}_\infty$. We claim that the sequence $\vec{b}_1, \vec{y}_1, \ldots, \vec{b}_k, \vec{y}_k$ satisfies Condition 2 of Lemma 2.8 for $\varphi = \exists \vec{u}_{1,\Delta}\ \gamma_{2,\Delta}$, the structure $\mathcal{B}_\infty^\Delta$, and appropriate sets $Z_1, \ldots, Z_k$.[6]

Part (a) of this condition is satisfied since $\mathcal{B}_\infty^{\Sigma\cup\Delta} \models \exists \vec{u}_{1,\Delta}\ \gamma_{2,\Delta}(\nu(\vec{u}_2), \vec{u}_{1,\Delta})$, and thus

$$\mathcal{B}_\infty^\Delta \models \exists \vec{u}_{1,\Delta}\ \gamma_{2,\Delta}(\vec{b}_1, \vec{y}_1, \ldots, \vec{b}_k, \vec{y}_k, \vec{u}_{1,\Delta}).$$

Part (b) of the condition is satisfied since the ν-images of all shared variables in $\vec{u}_2$ are distinct according to our choice in the variable identification step. Finally, part (c) and (d) can be satisfied because of our choice of the linear ordering. In fact, any component b of $\vec{b}_j$ belongs to B_{j-1}, and is thus generated by $\bigcup_{i=0}^{j-1} Y_i$, whereas the components of $\vec{y}_j$ are in Y_j. Thus, there exists a (finite) subset Z_j of $\bigcup_{i=0}^{j-1} Y_i$ such that part (c) of Condition 2 of the lemma is satisfied. Part (d) is satisfied as well since Y_j and $\bigcup_{i=0}^{j-1} Y_i$ are disjoint by definition.

Thus, we can apply Lemma 2.8, which yields $\mathcal{B} \simeq \mathcal{B}_\infty^\Delta \models \beta$. In order to show $\mathcal{A} \models \alpha$, we use the fact that $h_\infty : \mathcal{B}_\infty \to \mathcal{A}_\infty$ is a $(\Sigma \cup \Delta)$-isomorphism. Thus, $\mathcal{B}_\infty^{\Sigma\cup\Delta} \models \exists \vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma}(\nu(\vec{u}_2), \vec{u}_{1,\Sigma})$ implies that $\mathcal{A}_\infty^\Sigma \models \exists \vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma}(h_\infty(\nu(\vec{u}_2)), \vec{u}_{1,\Sigma})$.

Let $\vec{x}_i := h_\infty(\vec{b}_i) = h_\infty(\nu(\vec{v}_i))$ and $\vec{a}_i := h_\infty(\vec{y}_i) = h_\infty(\nu(\vec{w}_i))$ (for $i = 1, \ldots, k$). We claim that the sequence $\vec{x}_1, \vec{a}_1, \ldots, \vec{x}_k, \vec{a}_k$ satisfies Condition 2 of Lemma 2.8 for $\varphi = \exists \vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma}$, the structure $\mathcal{A}_\infty^\Sigma$, and appropriate sets $Z'_1, \ldots, Z'_k$.

Obviously, $\mathcal{A}_\infty^\Sigma \models \exists \vec{u}_{1,\Sigma}\ \gamma_{2,\Sigma}(h_\infty(\nu(\vec{u}_2)), \vec{u}_{1,\Sigma})$ implies that part (a) of the condition is satisfied. To see that part (b) is satisfied, recall that, by our

[6]Note that, in contrast to the formulation of the lemma, our sequence starts with a tuple of structure elements instead of generators. The lemma applies nevertheless since in its formulation we did not assume that all tuples have a non-zero dimension.

choice in the variable identification step, the ν-images of different shared variables in $\vec{u}_2$ are distinct. Since h_∞ is a bijection, this holds for their $(h_\infty \circ \nu)$-images as well.

Now part (c) and (d) are an easy consequence of the following properties, which in turn are consequences of the definition of the bijection h_∞ and and its inverse g_∞:

1. Since the components of $\vec{b}_1$ are in B_0, we know that the components of $\vec{x}_1$ are in $X_0 \cup X_1$.
2. For $1 < i \leq k$, the components of $\vec{b}_i$ are in $B_{i-1} \setminus (B_{i-2} \cup Y_{i-1})$. Thus, the components of $\vec{x}_i$ are in X_i.
3. For $1 \leq i \leq k$, the components of $\vec{y}_i$ are in Y_i. Thus, the components of $\vec{a}_i$ are in $A_{i-1} \setminus (A_{i-2} \cup Y_{i-1})$.

The third property shows that there exist (finite) subsets Z'_j of $\bigcup_{i=0}^{j-1} X_i$ satisfying part (c). Obviously, the properties 1 and 2 imply that part (d) is satisfied as well. Thus, we can apply Lemma 2.8, and obtain $\mathcal{A} \simeq \mathcal{A}_\infty^\Sigma \models \alpha$. $\square$

The two lemmas obviously imply the next theorem.

Theorem 4.3 *Let $\mathcal{V}(E)$ be a Σ-variety and $\mathcal{V}(F)$ be a Δ-variety for disjoint signatures Σ and Δ. The positive existential theory of the $(\Sigma \cup \Delta)$-variety $\mathcal{V}(E \cup F)$ is decidable, provided that the positive theories of $\mathcal{V}(E)$ and of $\mathcal{V}(F)$ are decidable.*

If the signatures contain no predicate symbols, this theorem is a reformulation of Theorem 2.1 of [BS92]. What is new here is the algebraic proof method and the fact that relational constraints can be treated as well.

5 Decision Procedures for Positive Theories

A disadvantage of Theorem 4.3 is that it does not show modularity of decidability of the positive theory of varieties of structures. Indeed, the prerequisites of the theorem (decidability of the *full* positive theories of $\mathcal{V}(E)$ and $\mathcal{V}(F)$) are stronger than its consequence (decidability of the *existential* positive theory of $\mathcal{V}(E \cup F)$). This problem will be overcome by the algorithm described in this section, which can be used to reduce decidability of the *full* positive theory of $\mathcal{V}(E \cup F)$ to decision procedures for the positive theories of $\mathcal{V}(E)$ and $\mathcal{V}(F)$. This shows that Theorem 4.3 can be applied iteratively.

Algorithm 2

The input is a positive sentence φ_1 in the mixed signature $\Sigma \cup \Delta$. We assume that φ_1 is in prenex normalform, and that the matrix of φ_1 is in disjunctive normalform. The algorithm proceeds in two phases.

Phase 1

Via Skolemization of universally quantified variables,[7] φ_1 is transformed into an existential sentence φ_1' over the signature $\Sigma \cup \Delta \cup \Gamma_1$. Here Γ_1 is the signature consisting of all the new Skolem function symbols that have been introduced.

Suppose that φ_1' is of the form $\exists \vec{u}_1(\bigvee \gamma_{1,i})$, where the $\gamma_{1,i}$ are conjunctions of atomic formulae. Obviously, φ_1' is equivalent to $\bigvee(\exists \vec{u}_1\ \gamma_{1,i})$, and thus it is sufficient to decide validity of the sentences $\exists \vec{u}_1\ \gamma_{1,i}$. Each of these sentences is used as input for Algorithm 1. Subsequently, $\exists \vec{u}_1 \gamma_1$ denotes one of these sentences.

The atomic formulae in γ_1 may contain symbols from the two (disjoint) signatures Σ and $\Delta \cup \Gamma_1$. In Phase 1 we treat the sentences $\exists \vec{u}_1 \gamma_{1,i}$ by means of Steps 1–4 of Algorithm 1, finally splitting them into positive Σ-sentences α and positive $(\Delta \cup \Gamma_1)$-sentences φ_2. Thus, the output of Phase 1 is a finite set of pairs (α, φ_2).

Phase 2

In the second phase, φ_2 is treated exactly as φ_1 was treated before, applying Skolemization to universally quantified variables and Steps 1–4 of Algorithm 1 a second time. Now we consider the two (disjoint) signatures Δ and $\Gamma = \Gamma_1 \cup \Gamma_2$, where Γ_2 contains the Skolem functions that are introduced by the Skolemization step of Phase 2. We obtain output pairs of the form (β, ρ), where β is a positive sentence over the signature Δ and ρ is a positive

[7] We are Skolemizing *universally* quantified variables since we are interested in validity of the sentence and not in satisfiability.

sentence over the signature Γ. Together with the corresponding sentence α (over the signature Σ) we thus obtain triples (α, β, ρ) as output.

For each of these triple, the sentence α is now tested for validity in $\mathcal{A}$, β is tested for validity in $\mathcal{B}$, and ρ is tested for validity in the absolutely free term algebra $\mathcal{T}(\Gamma, X)$ with countably many generators X, i.e., the free algebra over X for the class of all Γ-algebras.[8]

Correctness of Algorithm 2

We want to show that the original sentence φ_1 is valid iff for one of the output triples, all three components are valid in the respective structures. The proof depends on the following lemma, which exhibits an interesting connection between Skolemization and amalgamation with an absolutely free algebra.

Lemma 5.1 *Let $\mathcal{A}$ be a Σ-structure that is free in $\mathcal{V}(E)$ over the countably infinite set of generators X, and let γ be a positive Σ-sentence. Suppose that the (positive) existential sentence γ' is obtained from γ via Skolemization of the universally quantified variables in γ, introducing the set of Skolem function symbols Γ. Then $\mathcal{A} \models \gamma$ if, and only if, $\mathcal{A} \odot \mathcal{T}(\Gamma, X) \models \gamma'$.*

Proof. Let $\mathcal{B} = \mathcal{T}(\Gamma, X)$. In order to avoid notational overhead, we assume without loss of generality that existential and universal quantifiers alternate in γ,[9] i.e., $\gamma = \forall u_1 \exists v_1 \ldots \forall u_k \exists v_k\ \varphi(u_1, v_1, \ldots, u_k, v_k)$. Skolemization yields the existential formula $\gamma' = \exists v_1 \ldots \exists v_k\ \varphi(f_1, v_1, f_2(v_1), v_2, \ldots, f_k(v_1, \ldots, v_{k-1}), v_k)$. Thus, Γ consists of k distinct new Skolem functions $f_1, f_2, \ldots, f_k$ having the arities $0, 1, \ldots, k-1$, respectively.

First, assume that $\mathcal{A} \models \gamma$. As Σ-structures, $\mathcal{A}$ and $\mathcal{A} \odot \mathcal{B}$ are isomorphic, and thus

$$(*) \quad \mathcal{A} \odot \mathcal{B} \models \forall u_1 \exists v_1 \ldots \forall u_k \exists v_k\ \varphi(u_1, v_1, \ldots, u_k, v_k).$$

Since $\mathcal{A} \odot \mathcal{B}$ can also be considered as a $(\Sigma \cup \Gamma)$-structure, the Skolem symbols $f_1, f_2, \ldots, f_k$ are interpreted by functions $g_1, g_2, \ldots, g_k$ on the carrier $A \odot B$ of $\mathcal{A} \odot \mathcal{B}$. Because of $(*)$, for $g_1 \in A \odot B$ there exists $a_1 \in A \odot B$ such that $\mathcal{A} \odot \mathcal{B} \models \forall u_2 \exists v_2 \ldots \forall u_k \exists v_k\ \varphi(g_1, a_1, u_2, v_2, \ldots, u_k, v_k)$. Iterating this argument, we obtain $a_1, \ldots, a_k \in A \odot B$ such that

$$\mathcal{A} \odot \mathcal{B} \models \varphi(g_1, a_1, g_2(a_1), a_2, \ldots, g_k(a_1, \ldots, a_{k-1}), a_k).$$

[8] Note that Γ contains no predicate symbols.

[9] Obviously one can introduce additional quantifiers over variables not occurring in γ to generate an equivalent formula of this form.

Since the functions g_i are the interpretation of the symbols f_i in $\mathcal{A} \odot \mathcal{B}$, this yields

$$\mathcal{A} \odot \mathcal{B} \models \exists v_1 \ldots \exists v_k \ \varphi(f_1, v_1, f_2(v_1), v_2, \ldots, f_k(v_1, \ldots, v_{k-1}), v_k),$$

i.e., $\mathcal{A} \odot \mathcal{B} \models \gamma'$.

For the converse direction, assume that

$$\mathcal{A} \odot \mathcal{B} \models \exists v_1 \ldots \exists v_k \ \varphi(f_1, v_1, f_2(v_1), v_2, \ldots, f_k(v_1, \ldots, v_{k-1}), v_k).$$

Since $\mathcal{A} \odot \mathcal{B} \simeq \mathcal{A}_\infty^{\Sigma\cup\Gamma}$, there exist $a_1, \ldots, a_k \in A_\infty$ such that

$$(**) \quad \mathcal{A}_\infty^{\Sigma\cup\Gamma} \models \varphi(f_1^{\mathcal{A}}, a_1, f_2^{\mathcal{A}}(a_1), a_2, \ldots, f_k^{\mathcal{A}}(a_1, \ldots, a_{k-1}), a_k),$$

where $f_1^{\mathcal{A}}, \ldots, f_k^{\mathcal{A}}$ denote the functions on A_∞ that interpret the symbols $f_1, \ldots, f_k$.

Our goal is to apply Lemma 2.8. Obviously, $(**)$ shows that the sequence $f_1^{\mathcal{A}}, a_1, f_2^{\mathcal{A}}(a_1), a_2, \ldots, f_k^{\mathcal{A}}(a_1, \ldots, a_{k-1}), a_k$ satisfies part (a) of Condition 2 of Lemma 2.8. It remains to be shown that part (b), (c) and (d) are valid as well (for an appropriate choice of the sets $Z_1, \ldots, Z_k$). The proof will depend on the following four properties, which are an easy consequence of the fact that $\mathcal{B}_\infty^\Gamma$ is an absolutely free Γ-algebra. Note that the carrier of $\mathcal{B}_\infty^\Gamma$ consists of the Γ-terms over the set (of variables) Y_∞, i.e., the symbols f_i interpret themselves.

- (p1) Elements of B_∞ of the form $f_i(b_1, \ldots, b_{i-1})$ and $f_j(b'_1, \ldots, b'_{j-1})$ are distinct if $i \neq j$.
- (p2) Elements of B_∞ of the form $f_i(b_1, \ldots, b_{i-1})$ are elements of $B_\infty \setminus Y_\infty$.
- (p3) If $b \in B_{m+1} \setminus B_m$, then $f_j(\ldots, b, \ldots) \notin B_m \cup Y_{m+1}$.
- (p4) Terms $f_j(b_1, \ldots, b_{j-1})$ are distinct from all their arguments b_ν.

Now, (p1) and (p2) can be used to show part (b) of Condition 2 of Lemma 2.8. By definition of the bijections h_∞ and g_∞, the h_∞-image of $B_\infty \setminus Y_\infty$ is in X_∞, and thus $f_i^{\mathcal{A}}(a_1, \ldots, a_{i-1}) = h_\infty(f_i(g_\infty(a_1), \ldots, g_\infty(a_{i-1}))) \in X_\infty$ by (p2). This shows that the elements $f_i^{\mathcal{A}}(a_1, \ldots, a_{i-1})$ of the sequence are in fact generators, i.e., elements of X_∞. All these generators are different because of (p1). Indeed, since h_∞ is a bijection, (p1) implies

$$f_i^{\mathcal{A}}(a_1, \ldots, a_{i-1}) = h_\infty(f_i(g_\infty(a_1), \ldots, g_\infty(a_{i-1}))) \neq \\ h_\infty(f_j(g_\infty(a_1), \ldots, g_\infty(a_{j-1}))) = f_j^{\mathcal{A}}(a_1, \ldots, a_{j-1})$$

for all $i \neq j$.

To establish (c) and (d), we must find finite sets of generators $Z_1, \ldots, Z_k$ such that *(c)* a_i is generated by Z_i (for all $i, 1 \leq i \leq k$), and *(d)* the generator $f_i^{\mathcal{A}}(a_1, \ldots, a_{i-1})$ is not an element of $Z_1 \cup \ldots \cup Z_{i-1}$ (for all $i, 1 < i \leq k$).

Let Z_k be an arbitrary finite set of generators of a_k. To define Z_{i-1} for all $i, 1 < i \leq k$, let $b_1, \ldots, b_{i-1}$ be the images of $a_1, \ldots, a_{i-1}$ under the bijection g_∞, and let m be the minimal number such that $\{a_1, \ldots, a_{i-1}\} \subseteq A_m$. Obviously, this implies that there exist (finite) sets $Z_1, \ldots, Z_{i-1}$ satisfying (c) such that $Z_1 \cup \ldots \cup Z_{i-1} \subseteq \bigcup_{j=0}^{m} X_j$. This information, however, is not sufficient to infer part (d).

First, we consider the case where the sequence $a_1, \ldots, a_{i-1}$ contains an element $a_j \in A_m \setminus (A_{m-1} \cup X_m)$. Then $b_j = g_\infty(a_j)$ is an element of Y_{m+1}. Property (p3) yields $f_i(b_1, \ldots, b_{i-1}) \notin B_m \cup Y_{m+1}$, and thus $f_i^{\mathcal{A}}(a_1, \ldots, a_{i-1}) = h_\infty(f_i(b_1, \ldots, b_{i-1})) \notin A_m \cup X_{m+1}$. Hence $f_i^{\mathcal{A}}(a_1, \ldots, a_{i-1}) \notin \bigcup_{j=0}^{m} X_j \subseteq A_m \cup X_{m+1}$, and we can take an arbitrary finite subset Z_{i-1} of $\bigcup_{j=0}^{m} X_j$ that generates a_{i-1} (without violating (d)).

Otherwise, the sequence $a_1, \ldots, a_{j-1}$ contains a non-zero number of elements of X_m (these will be called generators of type 1), and possibly some elements of A_{m-1}. The latter elements have (finitely many) generators in $\bigcup_{j=0}^{m-1} X_j$ (which will be called generators of type 2). Recall that $g_\infty(X_m) = B_{m-1} \setminus (B_{m-2} \cup Y_{m-1})$. By (p3), $f_i(b_1, \ldots, b_{i-1}) \notin B_{m-2} \cup Y_{m-1}$, and thus $f_i^{\mathcal{A}}(a_1, \ldots, a_{i-1}) = h_\infty(f_i(b_1, \ldots, b_{i-1})) \notin A_{m-2} \cup X_{m-1}$. This implies that $f_i^{\mathcal{A}}(a_1, \ldots, a_{i-1})$ is different from all generators of type 2. In addition, (p4) says that $f_i(b_1, \ldots, b_{i-1})$ is different from all its arguments $b_1, \ldots, b_{i-1}$. Consequently, $f_i^{\mathcal{A}}(a_1, \ldots, a_{i-1})$ is distinct from all its arguments $a_1, \ldots, a_{i-1}$, and thus from all generators of type 1. Thus, we can define Z_{i-1} to be the set of all generators of type 1 or 2. This completes the proof that Condition 2 of Lemma 2.8 can be satisfied by an appropriate choice of sets of generators $Z_1, \ldots, Z_k$.

Applying the lemma, we obtain

$$\mathcal{A}_\infty^{\Sigma \cup \Gamma} \models \forall u_1 \exists v_1 \ldots \forall u_k \exists v_k \ \varphi(u_1, v_1, \ldots, u_k, v_k).$$

Since $\gamma = \forall u_1 \exists v_1 \ldots \forall u_k \exists v_k \ \varphi(u_1, v_1, \ldots, u_k, v_k)$ is a pure Σ-formula, and since, as Σ-structures, $\mathcal{A}$ and $\mathcal{A}_\infty$ are isomorphic, this shows $\mathcal{A} \models \gamma$. $\square$

Correctness of Algorithm 2 is an easy consequence of this lemma.

Proposition 5.2 *$\mathcal{A} \odot \mathcal{B} \models \varphi_1$ if, and only if, there exists an output triple (α, β, ρ) such that $\mathcal{A} \models \alpha$, $\mathcal{B} \models \beta$, and $\mathcal{T}(\Gamma, X) \models \rho$, where Γ consists of the*

Skolem functions introduced in Phase 1 and 2 of the algorithm and X is a countably infinite set (of variables).

Proof. Assume that $\mathcal{A} \odot \mathcal{B} \models \varphi_1$. By Lemma 5.1, this implies that $(\mathcal{A} \odot \mathcal{B}) \odot \mathcal{T}(\Gamma_1, X) \simeq \mathcal{A} \odot (\mathcal{B} \odot \mathcal{T}(\Gamma_1, X)) \models \varphi_1'$, i.e., the formula obtained from φ_1 by Skolemization. Let $\exists \vec{u}_1 \gamma_1$ be one of the disjuncts in φ_1' satisfied by $\mathcal{A} \odot (\mathcal{B} \odot \mathcal{T}(\Gamma_1, X))$. Since Algorithm 1 is correct, one of the output pairs (α, φ_2) generated by applying Algorithm 1 to $\exists \vec{u}_1 \gamma_1$ satisfies $\mathcal{A} \models \alpha$ and $\mathcal{B} \odot \mathcal{T}(\Gamma_1, X) \models \varphi_2$.

Applying Lemma 5.1 a second time, we obtain $(\mathcal{B} \odot \mathcal{T}(\Gamma_1, X)) \odot \mathcal{T}(\Gamma_2, X) \simeq \mathcal{B} \odot \mathcal{T}(\Gamma_1 \cup \Gamma_2, X) \models \varphi_2'$, where φ_2' is the positive existential sentence that is obtained from φ_2 via Skolemization. Algorithm 1, applied to φ_2', thus yields an output pair (β, ρ) at the end of Phase 2 such that $\mathcal{B} \models \beta$ and $\mathcal{T}(\Gamma_1 \cup \Gamma_2, X) \models \rho$.

It is easy to see that all arguments used during this proof also apply in the other direction. $\square$

The proposition shows that decidability of the positive theory of $\mathcal{A} \odot \mathcal{B}$ can be reduced to decidability of the positive theories of $\mathcal{A}$, $\mathcal{B}$, and of an absolutely free term algebra $\mathcal{T}(\Gamma, X)$. It is well-known that the whole first-order theory of absolutely free term algebras is decidable [Mal71, Mah88, CL89]. Thus, we obtain the desired modularity result:

Theorem 5.3 *Let $\mathcal{V}(E)$ be a Σ-variety and $\mathcal{V}(F)$ be a Δ-variety for disjoint signatures Σ and Δ. The positive theory of the $(\Sigma \cup \Delta)$-variety $\mathcal{V}(E \cup F)$ is decidable, provided that the positive theories of $\mathcal{V}(E)$ and of $\mathcal{V}(F)$ are decidable.*

6 Conclusion and Outlook

We have presented an abstract algebraic approach to the problem of combining constraint solvers for constraint languages over disjoint signatures. The constraints that can be handled this way are built from atomic equational *and relational* constraints with the help of conjunction, disjunction, and both universal and existential quantifiers. Solvability means validity of such (closed) constraint formulae in a free structure, or equivalently in a variety of structures.

Simple examples of free structures with a non-trivial relational part are (absolutely free) term algebras that are equipped with an ordering that is invariant under substitution, such as the lexicographic path ordering or the subterm ordering. For our combination result to apply, however, the positive theory of these structures must be decidable. For a total lexicographic path ordering, this is not the case. For the subterm ordering and for partial lexicographic path orderings, the existential theory is decidable, but the full first-order theory is undecidable [CT94]. Decidability of the positive theory is still an open problem.

Combination of constraint solving techniques in the presence of predicate symbols other than equality have independently been considered by H. Kirchner and Ch. Ringeissen [KR94]. However, their approach is based on the rewriting and abstraction techniques mentioned in the introduction (see, e.g., [BS92, Bou90]). Consequently, the interpretation of the predicate symbols in the combined structure is defined in a rather technical way, and it is not a priori clear what this definition means in an intuitive algebraic sense. It would be interesting to find out whether, for free structures, the combined structure of [KR94] coincides with our free amalgamated product.

We are currently working on a generalization of the notion of "combinable structure" that considerably extends the notion of a "free structure." An example of a structure that is not a free structure, but nevertheless satisfies the generalized combinability condition, is the algebra of rational trees.

References

[Bac91] L. Bachmair. *Canonical Equational Proofs.* Birkhäuser, Boston, Basel, Berlin, 1991.

[Bir35] G. Birkhoff. On the structure of abstract algebras. *Proceedings Cambridge Phil. Soc.*, 31:433–454, 1935.

[Bou90] A. Boudet. Unification in a combination of equational theories: An efficient algorithm. In M.E. Stickel, editor, *Proceedings of the 10th International Conference on Automated Deduction*, volume 449 of *Lecture Notes in Artificial Intelligence*, pages 292–307, Kaiserslautern, Germany, 1990. Springer-Verlag.

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

[BS93] F. Baader and K.U. Schulz. Unification in the union of disjoint equational theories: Combining decision procedures, 1993. Extended version, submitted for publication.

[BS93a] F. Baader and K.U. Schulz. Combination techniques and decision problems for disunification. In *Proceedings of the International Conference on Rewriting Techniques and Applications, RTA 93*, volume 690 of *Lecture Notes in Computer Science*, pages 301–315, Montreal, Canada, 1993. Springer–Verlag.

[Bür91] H.-J. Bürckert. *A Resolution Principle for a Logic with Restricted Quantifiers*, volume 568 of *Lecture Notes in Artificial Intelligence.* Springer-Verlag, 1991.

[Coh65] P.M. Cohn. *Universal Algebra.* Harper & Row, New York, 1965.

[CL89] H. Comon and P. Lescanne. Equational problems and disunification. *J. Symbolic Computation*, 7:371–425, 1989.

[CT94] H. Comon and R. Treinen. Ordering constraints on trees. In *Colloquium on trees in Algebra and Programming (CAAP)*, Lecture Notes in Computer Science, Edinburgh, Scotland, 1994. Springer Verlag.

[Der87] N. Dershowitz. Termination of rewriting. *J. Symbolic Computation*, 3:69–116, 1987.

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

[JK86] J.-P. Jouannaud and H. Kirchner. Completion of a set of rules modulo a set of equations. *SIAM J. Computing*, 15, 1986.

[KK89] C. Kirchner and H. Kirchner. Constrained equational reasoning. In *Proceedings of SIGSAM 1989 International Symposium on Symbolic and Algebraic Computation.* ACM Press, 1989.

[KR94] H. Kirchner and Ch. Ringeissen. Combining symbolic constraint solvers on algebraic domains. Unpublished manuscript, 1994.

[Mah88] M.J. Maher. Complete axiomatizations of the algebras of finite, rational and infinite trees. In *Proceedings of Third Annual Symposium on Logic in Computer Science, LICS'88*, pages 348–357, Edinburgh, Scotland, 1988. IEEE Computer Society.

[Mal71] A.I. Mal'cev. *The Metamathematics of Algebraic Systems*, volume 66 of *Studies in Logic and the Foundation of Mathematics*. North Holland, Amsterdam, London, 1971.

[Mal73] A.I. Mal'cev. *Algebraic Systems*, volume 192 of *Die Grundlehren der mathematischen Wissenschaften in Einzeldarstellungen*. Springer–Verlag, Berlin, 1973.

[Plo72] G. Plotkin. Building in equational theories. *Machine Intelligence*, 7:73–90, 1972.

[Rin92] Ch. Ringeissen. Unification in a combination of equational theories with shared constants and its application to primal algebras. In *Proceedings of LPAR'92*, volume 624 of *Lecture Notes in Artificial Intelligence*, pages 261–272, St. Petersburg, Russia, 1992. Springer-Verlag.

[SS89] M. Schmidt-Schauß. Unification in a combination of arbitrary disjoint equational theories. *J. Symbolic Computation*, 8(1,2):51–99, 1989.

[Wea93] N. Weaver. Generalized varieties. *Algebra Universalis*, 30:27–52, 1993.