Connecting many-sorted theories

Franz Baader[1] and Silvio Ghilardi[2]

[1] Institut für Theoretische Informatik, TU Dresden
[2] Dipartimento di Scienze dell'Informazione, Università degli Studi di Milano

Abstract. Basically, the connection of two many-sorted theories is obtained by taking their disjoint union, and then connecting the two parts through connection functions that must behave like homomorphisms on the shared signature. We determine conditions under which decidability of the validity of universal formulae in the component theories transfers to their connection. In addition, we consider variants of the basic connection scheme.

1 Introduction

The combination of decision procedures for logical theories arises in many areas of logic in computer science, such as constraint solving, automated deduction, term rewriting, modal logics, and description logics. In general, one has two first-order theories T_1 and T_2 over signatures Σ_1 and Σ_2, for which validity of a certain type of formulae (e.g., universal, existential positive, etc.) is decidable. These theories are then combined into a new theory T over a combination Σ of the signatures Σ_1 and Σ_2. The question is whether decidability transfers from T_1, T_2 to their combination T.

One way of combining the theories T_1, T_2 is to build their union $T_1 \cup T_2$. Both the Nelson-Oppen combination procedure [16, 15] and combination procedures for the word problem [19, 17, 5] address this type of combination, but for different types of formulae to be decided. Whereas the original combination procedures were restricted to the case of theories over disjoint signatures, there are now also solutions for the non-disjoint case [8, 22, 6, 9, 11, 3], but they always require some additional restrictions since it is easy to see that in the unrestricted case decidability does not transfer. Similar combination problems have also been investigated in modal logic, where one asks whether decidability of (relativized) validity transfers from two modal logics to their fusion [12, 20, 23, 4]. The approaches in [11, 3] actually generalize these results from equational theories induced by modal logics to more general first-order theories satisfying certain model-theoretic restrictions: the theories T_1, T_2 must be *compatible* with their shared theory T_0, and this shared theory must be *locally finite* (i.e., its finitely generated models are finite). The theory T_i is compatible with the shared theory T_0 iff (i) $T_0 \subseteq T_i$; (ii) T_0 has a model completion T_0^*; and (iii) every model of T_i embeds into a model of $T_i \cup T_0^*$.

In [13], a new combination scheme for modal logics, called $\mathcal{E}$-connection, was introduced, for which decidability transfer is much simpler to show than in the case of the fusion. Intuitively, the difference between fusion and $\mathcal{E}$-connection can be explained as follows. A model of the fusion is obtained from two models of the component logics by identifying their domains. In contrast, a model of the $\mathcal{E}$-connection consists of two separate models of the component logics together with certain connecting relations between their domains. There are also differences in the syntax of the combined logic. In the case of the fusion, the Boolean operators are shared, and all operators can be applied to each other without restrictions. In the case of the $\mathcal{E}$-connection, there are two copies of the Boolean operators, and operators of the different logics cannot be mixed; the only connection between the two logics are new (diamond) modal operators that are induced by the connecting relations.

If we want to adapt this approach to the more general setting of combining first-order theories, then we must consider many-sorted theories since only the sorts allow us to keep the domains separate and to restrict the way function symbols can be applied to each other. Let T_1, T_2 be two many-sorted theories that may share some sorts as well as function and relation symbols. We first build the disjoint union $T_1 \uplus T_2$ of these two theories (by using disjoint copies of the shared parts), and then connect them by introducing *connection functions* between the shared sorts. These connection functions must behave like homomorphisms for the shared function and predicate symbols, i.e., the axioms stating this are added to $T_1 \uplus T_2$. This corresponds to the fact that the new diamond operators in the $\mathcal{E}$-connection approach distribute over disjunction and do not change the false formula $\perp$. We call the combined theory obtained this way the *connection* of T_1 and T_2.

This kind of connection between theories has already been considered in automated deduction (see, e.g., [1, 24]), but only in very restricted cases where both T_1 and T_2 are fixed theories (e.g., the theory of sets and the theory of integers in [24]) and the connection functions have a fixed meaning (like yielding the length of a list). In categorical logic, this type of connection can be seen as an instance of a general co-comma construction in bicategories associated with theories and syntactic interpretations (see, e.g., [25]). However, in this general setting, computational properties of the combined theories have not been considered yet.

This paper is a first step towards providing general results on the transfer of decidability from component theories to their connection. We start by considering the simplest case where there is just one connection function, and show that decidability transfers whenever certain model-theoretic conditions are satisfied. These conditions are weaker than the ones required in [3] for the case of the union of theories.[1] In addition, both the combination procedure and its proof of correctness are much simpler than the ones in [11, 3]. The approach easily extends to the case of several connection functions. We will also consider variants of the general combination scheme where the connection function must satisfy additional properties (like being surjective, an embedding, or an isomorphism), or where a theory is connected with itself. The first variant is, for example, in-

[1] Our conditions are in general not weaker than the ones in [11], although this is the case for all the theories we have considered until now.

teresting since the combination result for the union of theories shown in [11] can be obtained from the variant where one has an isomorphism as connection function. The second case is interesting since it can be used to reduce the global consequence problem in the modal logic **K** to propositional satisfiability, which is a surprising result.

2 Notation and definitions

In this section, we fix the notation and give some important definitions, in particular a formal definition of the connection of two theories.

We use standard *many-sorted first-order logic* (see, e.g., [10]), but try to avoid the notational overhead caused by the presence of sorts as much as possible. Thus, a *signature* Ω consists of a non-empty set of sorts $\mathcal{S}$ together with a set of function symbols $\mathcal{F}$ and a set of predicate symbols $\mathcal{P}$. The function and predicate symbols are equipped with arities from $\mathcal{S}^*$ in the usual way. For example, if the arity of $f \in \mathcal{F}$ is $S_1S_2S_3$, then this means that the function f takes tuples consisting of an element of sort S_1 and an element of sort S_2 as input, and produces an element of sort S_3. We consider logic with equality, i.e., the set of predicate symbols contains a symbol $\approx_S$ for equality in every sort S. Usually, we will just use $\approx$ without explicitly specifying the sort. In this paper we usually assume that signatures are at most countable.

Terms and first-order formulae over Ω are defined in the usual way, i.e., they must respect the arities of function and predicate symbols, and the variables occurring in them are also equipped with sorts. An *Ω-atom* is a predicate symbol applied to (sort-conforming) terms, and an *Ω-literal* is an atom or a negated atom. A *ground* literal is a literal that does not contain variables. We use the notation $\phi(\underline{x})$ to express that ϕ is a formula whose free variables are among the ones in the tuple of variables $\underline{x}$. An *Ω-sentence* is a formula over Ω without free variables. An *Ω-theory* T is a set of Ω-sentences (called the axioms of T). If T, T' are Ω-theories, then we write (by a slight abuse of notation) $T \subseteq T'$ to express that all the axioms of T are logical consequences of the axioms of T'.

From the semantic side, we have the standard notion of an *Ω-structure* $\mathcal{A}$, which consists of non-empty and pairwise disjoint domains A_S for every sort S, and interprets function symbols f and predicate symbols P by functions $f^{\mathcal{A}}$ and predicates $P^{\mathcal{A}}$ according to their arities. By A we denote the union of all domains A_S. Validity of a formula ϕ in an Ω-structure $\mathcal{A}$ ($\mathcal{A} \models \phi$), satisfiability, and logical consequence are defined in the usual way. The Ω-structure $\mathcal{A}$ is a *model* of the Ω-theory T iff all axioms of T are valid in $\mathcal{A}$. If $\phi(\underline{x})$ is a formula with free variables $\underline{x} = x_1, \ldots, x_n$ and $\underline{a} = a_1, \ldots, a_n$ is a (sort-conforming) tuple of elements of A, then we write $\mathcal{A} \models \phi(\underline{a})$ to express that $\phi(\underline{x})$ is valid in $\mathcal{A}$ under the assignment $\{x_1 \mapsto a_1, \ldots, x_n \mapsto a_n\}$. Note that $\phi(\underline{x})$ is valid in $\mathcal{A}$ iff it is valid under all assignments iff its universal closure is valid in $\mathcal{A}$.

An *Ω-homomorphism* between two Ω-structures $\mathcal{A}$ and $\mathcal{B}$ is a mapping $\mu : A \to B$ that is sort-conforming (i.e., maps elements of sort S in $\mathcal{A}$ to elements of sort S in $\mathcal{B}$), and satisfies the condition

$$(*) \qquad \mathcal{A} \models \alpha(a_1, \ldots, a_n) \quad \text{implies} \quad \mathcal{B} \models \alpha(\mu(a_1), \ldots, \mu(a_n))$$

for all Ω-atoms $\alpha(x_1, \ldots, x_n)$ and (sort-conforming) elements $a_1, \ldots, a_n$ of A. In case the converse of $(*)$ holds too, μ is called an *Ω-embedding.* Note that an embedding is something more than just an injective homomorphism since the stronger condition must hold not only for the equality predicate, but for all predicate symbols. If the embedding μ is the identity on A, then we say that $\mathcal{A}$ is a *Ω-substructure* of $\mathcal{B}$.

We say that Σ is a subsignature of Ω (written $\Sigma \subseteq \Omega$) iff Σ is a signature that can be obtained from Ω by removing some of its sorts and function and predicate symbols. If $\Sigma \subseteq \Omega$ and $\mathcal{A}$ is an Ω-structure, then the *Σ-reduct* of $\mathcal{A}$ is the Σ-structure $\mathcal{A}_{|\Sigma}$ obtained from $\mathcal{A}$ by forgetting the interpretations of sorts, function and predicate symbols from Ω that do not belong to Σ. Conversely, $\mathcal{A}$ is called an *expansion* of the Σ-structure $\mathcal{A}_{|\Sigma}$ to the larger signature Ω. If $\mu : \mathcal{A} \to \mathcal{B}$ is an Ω-homomorphism, then the *Σ-reduct* of μ is the Σ-homomorphism $\mu_{|\Sigma} : \mathcal{A}_{|\Sigma} \to \mathcal{B}_{|\Sigma}$ obtained by restricting μ to the sorts that belong to Σ, i.e., by restricting the mapping to the domain of $\mathcal{A}_{|\Sigma}$.

Given a set X of constant symbols not belonging to the signature Ω, but each equipped with a sort from Ω, we denote by Ω^X the extension of Ω by these new constants. If $\mathcal{A}$ is an Ω-structure, then we can view the elements of A as a set of new constants, where $a \in A_S$ has sort S. By interpreting each $a \in A$ by itself, $\mathcal{A}$ can also be viewed as an Ω^A-structure. The *positive diagram* $\Delta^+_\Omega(\mathcal{A})$ of $\mathcal{A}$ is the set of all ground Ω^A-atoms that are true in $\mathcal{A}$, and the *diagram* $\Delta_\Omega(\mathcal{A})$ of $\mathcal{A}$ is the set of all ground Ω^A-literals that are true in $\mathcal{A}$. *Robinson's diagram theorems* [7] say that there is a homomorphism (embedding) between the Ω-structures $\mathcal{A}$ and $\mathcal{B}$ iff it is possible to expand $\mathcal{B}$ to an Ω^A-structure in such a way that it becomes a model of the positive diagram (diagram) of $\mathcal{A}$.

Basic Connections

In the remainder of this section, we introduce our basic scheme for connecting many-sorted theories, and illustrate it with the example of $\mathcal{E}$-connections of modal logics. Let T_1, T_2 be theories over the respective signatures Ω_1, Ω_2, and let Ω_0 be a common subsignature of Ω_1 and Ω_2. We call Ω_0 the *connecting* signature. In addition, let T_0 be an Ω_0-theory[2] that is contained in both T_1 and T_2. We define the new theory $T_1 >_{T_0} T_2$ (called the *connection of* T_1 *and* T_2 *over* T_0) as follows.

The *signature* Ω of $T_1 >_{T_0} T_2$ contains the disjoint union $\Omega_1 \uplus \Omega_2$ of the signatures Ω_1 and Ω_2, where the shared sorts and the shared function and predicate symbols are appropriately renamed, e.g., by attaching labels 1 and 2. Thus, if S $(f,\ P)$ is a sort (function symbol, predicate symbol) contained in both Ω_1

[2] When *defining* the connection of T_1, T_2, the theory T_0 is actually irrelevant; all we need is its signature Ω_0. However, for our decidability transfer results to hold, T_0 and the T_i must satisfy certain model-theoretic properties.

and Ω_2, then S^i (f^i, P^i) for $i = 1, 2$ are its renamed variants in the disjoint union, where the arities are accordingly renamed. In addition, Ω contains a *new function symbol* h_S of arity $S^1 S^2$ for every sort S of Ω_0.

The *axioms* of $T_1 >_{T_0} T_2$ are obtained as follows. Given an Ω_i-formula ϕ, its renamed variant ϕ^i is obtained by replacing all shared symbols by their renamed variants with label i. The axioms of $T_1 >_{T_0} T_2$ consist of

$$\{\phi^1 \mid \phi \in T_1\} \cup \{\phi^2 \mid \phi \in T_2\},$$

together with the universal closures of the formulae

$$\begin{array}{l} h_S(f^1(x_1, \ldots, x_n)) \approx f^2(h_{S_1}(x_1), \ldots, h_{S_n}(x_n)), \\ P^1(x_1, \ldots, x_n) \rightarrow P^2(h_{S_1}(x_1), \ldots, h_{S_n}(x_n)), \end{array}$$

for every function (predicate) symbol f (P) in Ω_0 of arity $S_1 \ldots S_n S$ ($S_1 \ldots S_n$).

Since the signatures Ω_1 and Ω_2 have been made disjoint, and since the additional axioms state that the family of mappings h_S behaves like an Ω_0-homomorphism, it is easy to see that the *models of* $T_1 >_{T_0} T_2$ are formed by triples of the form $(\mathcal{M}^1, \mathcal{M}^2, h^{\mathcal{M}})$, where $\mathcal{M}^1$ is a model of T_1, $\mathcal{M}^2$ is a model of T_2, and $h^{\mathcal{M}}$ is an Ω_0-homomorphism $h^{\mathcal{M}} : \mathcal{M}^1_{|\Omega_0} \rightarrow \mathcal{M}^2_{|\Omega_0}$ between the respective Ω_0-reducts.

Example 1. The most basic variant of the $\mathcal{E}$-connection scheme introduced in [13] is an instance of our approach if one translates it into the algebraic setting. The abstract description systems considered in [13], which cover all the usual modal and description logics, are closely related to to Boolean-based equational theories (see [2] for details). The theory E is called *Boolean-based equational theory* [3] iff its signature Σ has just one sort, equality is the only predicate symbol, the set of function symbols contains the Boolean operators $\sqcap, \sqcup, \neg, \top, \bot$, and its set of axioms consists of identities (i.e., the universal closures of atoms $s \approx t$) and contains the Boolean algebra axioms.

For example, consider the basic modal logic $\mathbf{K}$, where we use only the modal operator $\Diamond$ (since $\Box$ can then be defined). The Boolean-based equational theory $E_{\mathbf{K}}$ corresponding to $\mathbf{K}$ is obtained from the theory of Boolean algebras by adding the identities $\Diamond(x \sqcup y) \approx \Diamond(x) \sqcup \Diamond(y)$ and $\Diamond(\bot) \approx \bot$.

Let us illustrate the notion of an $\mathcal{E}$-connection also on this simple example. To build the $\mathcal{E}$-connection of $\mathbf{K}$ with itself, one takes two disjoint copies of $\mathbf{K}$, obtained by renaming the Boolean operators and the diamonds, e.g., into $\sqcap_i, \sqcup_i, \neg_i, \top_i, \bot_i, \Diamond_i$ for $i = 1, 2$. The signature of the $\mathcal{E}$-connection contains all these renamed symbols together with a new symbol $\Diamond$. However, it is now a two-sorted signature, where symbols with index i are applied to elements of sort S_i and yield as results an element of this sort. The new symbol has arity $S_1 S_2$.[3] The semantics of this $\mathcal{E}$-connection can be given in terms of Kripke structures.

[3] In the $\mathcal{E}$-connection scheme introduced in [13], there is also an inverse diamond operator $\Diamond^-$ with arity $S_2 S_1$, but the algebraic approach introduced in the present paper cannot treat this case (see the conclusion for a discussion).

A Kripke structure for the $\mathcal{E}$-connection consists of two Kripke structures $\mathcal{K}_1, \mathcal{K}_2$ for $\mathbf{K}$ over disjoint domains W_1 and W_2, together with an additional connecting relation $E \subseteq W_2 \times W_1$. The symbols with index i are interpreted in $\mathcal{K}_i$, and the new symbol $\Diamond$ is interpreted as the diamond operator induced by E, i.e., for every $X \subseteq W_1$ we have

$$\Diamond(X) := \{x \in W_2 \mid \exists y \in W_1.\ (x,y) \in E \wedge y \in X\}.$$

This interpretation of the new operator implies that it satisfies the usual identities of a diamond operator, i.e., $\Diamond(x \sqcup_1 y) \approx \Diamond(x) \sqcup_2 \Diamond(y)$ and $\Diamond(\bot_1) \approx \bot_2$, and that these identities are sufficient to characterize its semantics. Thus, the equational theory corresponding to the $\mathcal{E}$-connection of $\mathbf{K}$ with itself consists of these two axioms, together with the axioms of $E_{\mathbf{K}_1}$ and $E_{\mathbf{K}_2}$.

Obviously, this theory is also obtained as the connection of the theory $E_{\mathbf{K}}$ with itself, if the connecting signature Ω_0 consists of the single sort of $E_{\mathbf{K}}$, the predicate symbol $\approx$, and the function symbols $\sqcup, \bot$. As theory T_0 we can take the theory of semilattices, i.e., the axioms that say that $\sqcup$ is associative, commutative, and idempotent, and that $\bot$ is a unit for $\sqcup$.

Example 2. The previous example can be varied by including $\sqcap$ in the connecting signature, and taking as theory T_0 the theory of distributive lattices with a least element $\bot$. It is easy to see that this corresponds to the case of an $\mathcal{E}$-connection where the connecting relation E is required to be a partial function.

3 Positive algebraic completions and compatibility

In order to transfer decidability results from the component theories T_1, T_2 to their connection $T_1 >_{T_0} T_2$ over T_0, the theories T_0, T_1, T_2 must satisfy certain model-theoretic conditions, which we introduce below. The most important one is that T_0 has a positive algebraic completion. Before we can define this concept, we must introduce some notions from model theory [7].

The formula ϕ is called *open* iff it does not contain quantifiers; it is called *universal* iff it is obtained from an open formula by adding a prefix of universal quantifiers; and it is called *geometric* iff it is built from atoms by using conjunction, disjunction, true, false, and existential quantifiers.[4]

The main property of geometric formulae is that they are preserved under homomorphisms in the following sense: if $\mu : \mathcal{A} \to \mathcal{B}$ is a homomorphism between Ω-structures and $\phi(x_1, \ldots, x_n)$ is a geometric formula over Ω, then $\mathcal{A} \models \phi(a_1, \ldots, a_n)$ implies $\mathcal{B} \models \phi(\mu(a_1), \ldots, \mu(a_n))$ for all (sort-conforming) $a_1, \ldots, a_n \in A$. Open formulae are related to embeddings in various ways. First, they are preserved under building sub- and superstructures, i.e., if $\mathcal{A}$ is a substructure of $\mathcal{B}$, $\phi(x_1, \ldots, x_n)$ is an open formula, and $a_1, \ldots, a_n \in A$ are sort-conforming, then $\mathcal{A} \models \phi(a_1, \ldots, a_n)$ iff $\mathcal{B} \models \phi(a_1, \ldots, a_n)$. Moreover, two Ω-theories T, T' entail the same set of open formulae iff every model of T can be embedded into a model of T' and vice versa (see [7] for these and related results).

[4] The latter formulae are called "geometric" in categorical logic [14] since they are preserved under inverse image geometric morphisms.

The theory T is a *universal theory* iff its axioms are universal sentences; it is a *geometric theory* iff it can be axiomatized by using universal closures of geometric sequents, where a geometric sequent is an implication between two geometric formulae. Note that any universal theory is geometric since open formulae are conjunctions of clauses and clauses can be rewritten as geometric sequents.

Definition 1. *Let T be a universal and T^* a geometric theory over Ω. We say that T^* is a* positive algebraic completion *of T iff the following properties hold:*

1. $T \subseteq T^*$;
2. *every model of T embeds into a model of T^*;*[5]
3. *for every geometric formula $\phi(\underline{x})$ there is an* open *geometric formula $\phi^*(\underline{x})$ such that $T^* \models \phi \leftrightarrow \phi^*$.*

It can be shown that the models of T^* are exactly the algebraically closed models of T (see [2]). In particular, this means that the positive algebraic completion of T is unique, provided that it exists.

When trying to show that Property 3 of Definition 1 holds for given theories T, T^*, it is sufficient to consider *simple existential formulae* $\phi(\underline{x})$, i.e., formulae that are obtained from conjunctions of atoms by adding an existential quantifier prefix. In fact, any geometric formula ϕ can be normalized to a disjunction $\phi_1 \vee \ldots \vee \phi_n$ of simple existential formulae ϕ_i by using distributivity of conjunction and existential quantification over disjunction. In addition, if $T^* \models \phi_i \leftrightarrow \phi_i^*$ for geometric open formulae ϕ_i^* $(i = 1, \ldots, n)$, then $\phi_1^* \vee \ldots \vee \phi_n^*$ is also a geometric open formula and $T^* \models (\phi_1 \vee \ldots \vee \phi_n) \leftrightarrow (\phi_1^* \vee \ldots \vee \phi_n^*)$.

The following lemma will turn out to be useful later on.

Lemma 1. *Assume that T, T^* satisfy Property 2 of Definition 1. If $\phi(\underline{x})$ is a simple existential formula and $\phi^*(\underline{x})$ is an open formula, then $T^* \models \phi \to \phi^*$ implies $T \models \phi \to \phi^*$.*

This is an immediate consequence of the facts that $\phi \to \phi^*$ is then equivalent to an open formula, and open formulae are preserved under building substructures.

The first ingredient of our combinability condition is the following notion of compatibility, which is a variant of analogous compatibility conditions introduced in [11, 3] for the case of the union of theories.

Definition 2. *Let $T_0 \subseteq T$ be theories over the respective signatures $\Omega_0 \subseteq \Omega_1$. We say that T is* T_0-algebraically compatible *iff T_0 is universal, has a positive algebraic completion T_0^*, and every model of T embeds into a model of $T \cup T_0^*$.*

The second ingredient is that T_0 must be locally finite, i.e., all finitely generated models of T_0 are finite. To be more precise, we need the following effective variant of local finiteness defined in [11, 3]. Let T_0 be a universal theory over the finite signature Ω_0. Then T_0 is called *effectively locally finite* iff for every tuple of variables $\underline{x}$, one can effectively determine terms $t_1(\underline{x}), \ldots, t_k(\underline{x})$ such that, for every further term $u(\underline{x})$, we have that $T_0 \models u \approx t_i$ for some $i = 1, \ldots, k$.

[5] Equivalently, T and T^* entail the same universal sentences.

4 The main combination result

We are interested in deciding the universal fragments of our theories, i.e., validity of universal formulae (or, equivalently open formulae) in a theory T. This is the decision problem also treated by the Nelson-Oppen combination method (albeit for the union of theories). It is well known that this problem is equivalent to the problem of deciding whether a set of literals is satisfiable in some model of T. We call such a set of literals a *constraint*.

By introducing new free constants (i.e., constants not occurring in the axioms of the theory), we can assume without loss of generality that such constraints contain no variables. In addition, we can transform any ground constraint into an equisatisfiable set of *ground flat literals*, i.e., literals of the form

$$a \approx f(a_1, \ldots, a_n), \quad P(a_1, \ldots, a_n), \quad \text{or} \quad \neg P(a_1, \ldots, a_n),$$

where $a, a_1, \ldots, a_n$ are (sort-conforming) free constants, f is a function symbol, and P is a predicate symbol (possibly also equality).

Theorem 1. *Let T_0, T_1, T_2 be theories over the respective signatures $\Omega_0, \Omega_1, \Omega_2$, where Ω_0 is a common subsignature of Ω_1 and Ω_2. Assume that $T_0 \subseteq T_1$ and $T_0 \subseteq T_2$, that T_0 is universal and effectively locally finite, and that T_2 is T_0-algebraically compatible. Then the decidability of the universal fragments of T_1 and T_2 entails the decidability of the universal fragment of $T_1 >_{T_0} T_2$.*

To prove the theorem, we consider a finite set Γ of ground flat literals over the signature Ω of $T_1 >_{T_0} T_2$ (with additional free constants), and show how it can be tested for satisfiability in $T_1 >_{T_0} T_2$. Since all literals in Γ are flat, we can divide Γ into three disjoint sets $\Gamma = \Gamma_0 \cup \Gamma_1 \cup \Gamma_2$, where Γ_i $(i = 1, 2)$ is a set of literals in the signature Ω_i (expanded with free constants), and Γ_0 is of the form

$$\Gamma_0 = \{h(a_1) \approx b_1, \ldots, h(a_n) \approx b_n\}$$

for free constants $a_1, b_1, \ldots, a_n, b_n$. Here and in the following we omit the sort index when writing the connection functions h_S.

Proposition 1. *The constraint $\Gamma = \Gamma_0 \cup \Gamma_1 \cup \Gamma_2$ is satisfiable in $T_1 >_{T_0} T_2$ iff there exists a triple $(\mathcal{A}, \mathcal{B}, \nu)$ such that*

1. *$\mathcal{A}$ is an Ω_0-model of T_0, which is generated by $\{a_1^{\mathcal{A}}, \ldots, a_n^{\mathcal{A}}\}$;*
2. *$\mathcal{B}$ is an Ω_0-model of T_0, which is generated by $\{b_1^{\mathcal{B}}, \ldots, b_n^{\mathcal{B}}\}$;*
3. *$\nu : \mathcal{A} \to \mathcal{B}$ is an Ω_0-homomorphism such that $\nu(a_j^{\mathcal{A}}) = b_j^{\mathcal{B}}$ for $j = 1, \ldots, n$;*
4. *$\Gamma_1 \cup \Delta_{\Omega_0}(\mathcal{A})$ is satisfiable in T_1;*
5. *$\Gamma_2 \cup \Delta_{\Omega_0}(\mathcal{B})$ is satisfiable in T_2.*

Proof. The only-if direction is simple. In fact, as noted in Section 2, a model $\mathcal{M}$ of $T_1 >_{T_0} T_2$ is given by a triple $(\mathcal{M}^1, \mathcal{M}^2, h^{\mathcal{M}})$, where $\mathcal{M}^1$ is a model of T_1, $\mathcal{M}^2$ is a model of T_2, and $h^{\mathcal{M}} : \mathcal{M}^1_{|\Omega_0} \to \mathcal{M}^2_{|\Omega_0}$ is an Ω_0-homomorphism between the respective Ω_0-reducts. Assume that this model $\mathcal{M}$ satisfies Γ. We can take as $\mathcal{A}$

the substructure of $\mathcal{M}^1_{|\Omega_0}$ generated by (the interpretations of) $a_1, \ldots, a_n$, as $\mathcal{B}$ the substructure of $\mathcal{M}^2_{|\Omega_0}$ generated by (the interpretations of) $b_1, \ldots, b_n$, and as homomorphism ν the restriction of $h^{\mathcal{M}}$ to $\mathcal{A}$. It is easy to see that the triple $(\mathcal{A}, \mathcal{B}, \nu)$ obtained this way satisfies 1.–5. of the proposition.

Conversely, assume that $(\mathcal{A}, \mathcal{B}, \nu)$ is a triple satisfying 1.–5. of the proposition. Because of 4. and 5., there is an Ω_1-model $\mathcal{N}'$ of T_1 satisfying $\Gamma_1 \cup \Delta_{\Omega_0}(\mathcal{A})$ and an Ω_2-model $\mathcal{N}''$ of T_2 satisfying $\Gamma_2 \cup \Delta_{\Omega_0}(\mathcal{B})$. By Robinson's diagram theorem, $\mathcal{N}'$ has $\mathcal{A}$ as an Ω_0-substructure and $\mathcal{N}''$ has $\mathcal{B}$ as an Ω_0-substructure. We assume without loss of generality that $\mathcal{N}'$ is at most countable and that $\mathcal{N}''$ is a model of $T_2 \cup T_0^*$. The latter assumption is by T_0-algebraic compatibility of T_2, and the former assumption is by the Löwenheim-Skolem theorem since our signatures are at most countable. Let us enumerate the elements of $\mathcal{N}'$ as

$$c_1, c_2, \ldots, c_n, c_{n+1}, \ldots,$$

where we assume that $c_i = a_i^{\mathcal{A}}$ $(i = 1, \ldots, n)$, i.e., $c_1, \ldots, c_n$ are generators of $\mathcal{A}$. We define an increasing sequence of sort-conforming functions $\nu_k : \{c_1, \ldots c_k\} \to N''$ (for $k \geq n$) such that, for every ground $\Omega_0^{\{c_1,\ldots,c_k\}}$-atom α we have

$$\mathcal{N}'_{|\Omega_0} \models \alpha(c_1, \ldots, c_k) \quad \text{implies} \quad \mathcal{N}''_{|\Omega_0} \models \alpha(\nu_k(c_1), \ldots, \nu_k(c_k)).$$

We first take ν_n to be ν. To define ν_{k+1} (for $k \geq n$), let us consider the conjunction $\psi(c_1, \ldots, c_k, c_{k+1})$ of the $\Omega_0^{\{c_1,\ldots,c_{k+1}\}}$-atoms that are true in $\mathcal{N}'_{|\Omega_0}$: this conjunction is finite (modulo taking representative terms, thanks to local finiteness of T_0). Let $\phi(x_1, \ldots, x_k)$ be $\exists x_{k+1}.\psi(x_1, \ldots, x_k, x_{k+1})$ and let $\phi^*(x_1, \ldots, x_k)$ be a geometric open formula such that $T_0^* \models \phi \leftrightarrow \phi^*$.

By Lemma 1, $T_0 \models \phi \to \phi^*$, and thus we have $\mathcal{N}'_{|\Omega_0} \models \phi^*(c_1, \ldots, c_k)$ and also $\mathcal{N}''_{|\Omega_0} \models \phi^*(\nu_k(c_1), \ldots, \nu_k(c_k))$ by the induction hypothesis. Since $\mathcal{N}''_{|\Omega_0}$ is a model of T_0^*, there is a b such that $\mathcal{N}''_{|\Omega_0} \models \psi(\nu_k(c_1), \ldots, \nu_k(c_k), b)$ for some b. We now obtain the desired extension ν_{k+1} of ν_k by setting $\nu_{k+1}(c_{k+1}) := b$. Taking $\nu_\infty = \bigcup_{k \geq n} \nu_k$, we finally obtain a homomorphism $\nu_\infty : \mathcal{N}'_{|\Omega_0} \to \mathcal{N}''_{|\Omega_0}$ such that the triple $(\mathcal{N}', \mathcal{N}'', \nu_\infty)$ is a model of $T_1 >_{T_0} T_2$ that satisfies $\Gamma_0 \cup \Gamma_1 \cup \Gamma_2$. $\square$

The above proof uses the assumption that T_0 is locally finite. By using heavier model-theoretic machinery, one can also prove the proposition without using local finiteness of T_0 (see [2]). However, since the proof of Theorem 1 needs this assumption anyway (see below), we gave the above proof since it is simpler.

To conclude the proof of Theorem 1, we describe a *non-deterministic decision procedure* that effectively guesses an appropriate triple $(\mathcal{A}, \mathcal{B}, \nu)$ and then checks whether it satisfies 1.–5. of Proposition 1. To guess an Ω_0-model of T_0 that is generated by a finite set X, one uses effective local finiteness of T_0 to obtain an effective bound on the size of such a model, and then guesses an Ω_0-structure that satisfies this size bound. Once the structures $\mathcal{A}, \mathcal{B}$ are given, one can build their diagrams, and use the decision procedures for T_1 and T_2 to check whether 4. and 5. of Proposition 1 are satisfied. If the answer is yes, then $\mathcal{A}, \mathcal{B}$ are also models of T_0: in fact, if for instance $\Gamma_1 \cup \Delta_{\Omega_0}(\mathcal{A})$ is satisfiable in the model $\mathcal{M}$

of T_1, then $\mathcal{M}$ has $\mathcal{A}$ as a substructure, and this implies $\mathcal{A} \models T_0$ because T_0 is universal and $T_0 \subseteq T_1$. Finally, one can guess a mapping $\nu : A \to B$ that satisfies $\nu(a_j^{\mathcal{A}}) = b_j^{\mathcal{B}}$, and then use the diagrams of $\mathcal{A}, \mathcal{B}$ to check whether ν satisfies the homomorphism condition $(*)$.

The proof of Proposition 1 shows that our decidability transfer result can easily be extended to the case of *several connection functions*, possibly going in both directions. In fact, one simply considers several Ω_0-homomorphisms between $\mathcal{A}$ and $\mathcal{B}$ in 3. of the proposition, and extends them separately to homomorphisms between $\mathcal{N}'$ and $\mathcal{N}''$ (see [2] for more details). If there are also connection functions in the other direction (and thus homomorphisms from $\mathcal{B}$ to $\mathcal{A}$), then T_1 must also be T_0-algebraically compatible.

Examples

When trying to axiomatize the positive algebraic completion T_0^* of a given universal theory T_0, it is sufficient to produce for every simple existential formula $\phi(\underline{x})$ an appropriate geometric and open formula $\phi^*(\underline{x})$. Take as theory T_0^* the one axiomatized by T_0 together with the formulae $\phi \leftrightarrow \phi^*$ for every simple existential formula ϕ. In order to complete the job, it is sufficient to show that every model of T_0 embeds into a model of T_0^*. It should also be noted that one can without loss of generality restrict the attention to simple existential formulae with just one existential quantifier since more than one quantifier can then be treated by iterated elimination of single quantifiers.

In the next example we encounter a special case where the formulae $\phi \leftrightarrow \phi^*$ are already valid in T_0. In this case, we have $T_0 = T_0^*$, and thus the model-embedding condition is trivially satisfied. In addition, any theory T with $T_0 \subseteq T$ is T_0-algebraically compatible.

Example 3. Recall from [3] the definition of a Gaussian theory. Let us call a conjunction of atoms an *e-formula.* The universal theory T_0 is *Gaussian* iff for every e-formula $\phi(\underline{x}, y)$ it is possible to compute an e-formula $\psi(\underline{x})$ and a term $s(\underline{x}, \underline{z})$ with fresh variables $\underline{z}$ such that

$$T_0 \models \phi(\underline{x}, y) \leftrightarrow (\psi(\underline{x}) \wedge \exists \underline{z}.(y \approx s(\underline{x}, \underline{z}))). \tag{1}$$

Any Gaussian theory T_0 is its own positive algebraic completion. In fact, it is easy to see that (1) implies $T_0 \models (\exists y.\phi(\underline{x}, y)) \leftrightarrow \psi(\underline{x})$, and thus the comment given above this example applies.

As a consequence, our combination result applies to all the examples of effectively locally finite Gaussian theories given in [3] (e.g., Boolean algebras, vector spaces over a finite field, empty theory over a signature whose sets of predicates consists of $\approx$ and whose set of function symbols is empty): if the universal theory T_0 is effectively locally finite and Gaussian, and T_1, T_2 are arbitrary theories containing T_0 and with decidable universal fragment, then the universal fragment of $T_1 >_{T_0} T_2$ is also decidable.

Example 4. Let T_0 be the theory of semilattices (see Example 1). This theory is obviously effectively locally finite. In the following, we use the disequation $s \sqsubseteq t$ as an abbreviation for the equation $s \sqcup t \approx t$. Obviously, any equation $s \approx t$ can be expressed by the disequations $s \sqsubseteq t \wedge t \sqsubseteq s$.

The theory T_0 has a positive algebraic completion, which can be axiomatized as follows. Let $\phi(\underline{x})$ be a simple existential formula with just one existential quantifier. Using the fact that $z_1 \sqcup \ldots \sqcup z_n \sqsubseteq z$ is equivalent to $z_1 \sqsubseteq z \wedge \ldots \wedge z_n \sqsubseteq z$, it is easy to see that $\phi(\underline{x})$ is T_0-equivalent to a formula of the form

$$\psi(\underline{x}) \wedge \exists y.((y \sqsubseteq t_1) \wedge \cdots \wedge (y \sqsubseteq t_n) \wedge (u_1 \sqsubseteq s_1 \sqcup y) \wedge \cdots \wedge (u_m \sqsubseteq s_m \sqcup y)), \quad (2)$$

where $\psi(\underline{x}), t_i, s_j, u_k$ do not contain y. Let $\phi^*(\underline{x})$ be the formula

$$\psi(\underline{x}) \quad \wedge \quad \bigwedge_{i=1}^{n} \bigwedge_{j=1}^{m} (u_j \sqsubseteq s_j \sqcup t_i), \quad (3)$$

and let T_0^* be obtained from T_0 by adding to it the universal closures of all formulae $\phi \leftrightarrow \phi^*$.

We prove that T_0^* is contained in the theory of Boolean algebras. In fact, the system of disequations (2) is equivalent, in the theory of Boolean algebras, to

$$\psi(\underline{x}) \wedge \exists y.((y \sqsubseteq t_1) \wedge \cdots \wedge (y \sqsubseteq t_n) \wedge (u_1 \sqcap \neg s_1 \sqsubseteq y) \wedge \cdots \wedge (u_m \sqcap \neg s_m \sqsubseteq y), \quad (4)$$

and hence to

$$\psi(\underline{x}) \wedge (u_1 \sqcap \neg s_1 \sqsubseteq t_1 \sqcap \ldots \sqcap t_n) \wedge \cdots \wedge (u_m \sqcap \neg s_m \sqsubseteq t_1 \sqcap \ldots \sqcap t_n). \quad (5)$$

Finally, it is easy to see that (5) and (3) are equivalent.

Since every semilattice embeds into a Boolean algebra [2], this shows that T_0^* is the positive algebraic completion of T_0. In addition, this implies that any Boolean-based equational theory T is T_0-algebraically compatible since T_0^* is contained in T. Consequently, Theorem 1 covers the case of a basic $\mathcal{E}$-connection (see Example 1) for arbitrary classical modal logics as components.

In [2] we show a similar result for the case where the theory T_0 is the theory of distributive lattices with $\perp$. Thus, our result also covers the case of connecting relations that are partial functions (see Example 2).

Complexity considerations

The complexity of the combined decision procedure described in the proof of Theorem 1 is usually higher than the complexity of the decision procedures for the components. There are two main reasons for this complexity increase. First, one must guess the Ω_0-structures $\mathcal{A}, \mathcal{B}$ as the well as the mapping $\nu : A \to B$. This can be done by a *non-deterministic* procedure whose complexity depends on the bound on the size of Ω_0-models of T_0 with n generators given by the effective local finiteness of T_0. Second, the decision procedures for T_1 and T_2 are respectively applied to $\Gamma_1 \cup \Delta_{\Omega_0}(\mathcal{A})$ and $\Gamma_2 \cup \Delta_{\Omega_0}(\mathcal{B})$. The size of the diagrams again depends on the bound on the size of finitely generated Ω_0-models of T_0.

Let us consider the case where T_0 is the theory of semilattices (see Examples 1 and 4) in more detail. Given generators $a_1, \ldots, a_n$, there are 2^n representative terms, namely all terms of the form $a_{i_1} \sqcup \cdots \sqcup a_{i_k}$ for $\{i_1, \ldots, i_k\} \subseteq \{1, \ldots, n\}$ (where the empty disjunction corresponds to $\bot$). Atoms are of the form $t_1 \approx t_2$ where t_1, t_2 are such representative terms, and thus there are $2^n \cdot 2^n = 2^{2n}$ atoms. One can now guess a possible diagram of an Ω_0-structure by guessing (in non-deterministic exponential time) a subset S of the set of atoms. Given such a subset, the potential diagram is $\Delta_S := \{\alpha \mid \alpha \in S\} \cup \{\neg\alpha \mid \alpha \notin S\}$. Of course, not every such set Δ_S is indeed the diagram of an Ω_0-structure, but the ones that are not will lead to unsatisfiability when satisfiability in T_i of $\Gamma_i \cup \Delta_S$ is tested. Since the size of Δ_S is $O(n \cdot 2^{2n})$, the complexity of this satisfiability test is one exponential higher than the complexity of the satisfiability problem in T_i.

Assume that we have guessed sets S_1, S_2 determining the diagrams of semilattices $\mathcal{A}, \mathcal{B}$ generated by $a_1, \ldots, a_n$ and $b_1, \ldots, b_n$, respectively. Guessing an Ω_0-homomorphism $\nu : \mathcal{A} \to \mathcal{B}$ is not really necessary. In fact, if it exists, such a homomorphism ν is uniquely determined by the requirement that $\nu(a_i) = b_i$ $(i = 1, \ldots, n)$ since the semilattice $\mathcal{A}$ is generated by the $a_1, \ldots, a_n$. Obviously, an Ω_0-homomorphism $\nu : \mathcal{A} \to \mathcal{B}$ with $\nu(a_i) = b_i$ exists iff $\alpha(a_1, \ldots, a_n) \in S_1$ implies $\alpha(b_1, \ldots, b_n) \in S_2$ for all Ω_0-atoms $\alpha(x_1, \ldots, x_n)$. Thus, if one first guesses S_1, then one can start with $S_1' := \{\alpha(b_1, \ldots, b_n) \mid \alpha(a_1, \ldots, a_n) \in S_1\}$ and add some additional atoms when guessing S_2.

To sum up, in the case of T_0 being the theory of semilattices, our combined decision procedure has the following complexity. Its starts with a non-deterministic exponential step that guesses potential diagrams Δ_{S_1} and Δ_{S_2} such that the homomorphism condition $(*)$ is satisfied. Then it tests $\Gamma_i \cup \Delta_{S_i}$ $(i = 1, 2)$ for satisfiability in T_i. Since the size of Δ_{S_i} is exponential, the complexity of this step is one exponential higher than the complexity of deciding the universal fragment of T_i. This shows that our combination procedure has the same complexity as the one for $\mathcal{E}$-connections described in [13].

Let us consider the complexity increase caused by the combination procedure in more detail for the complexity class ExpTime, which is often encountered when considering the global satisfiability problem in modal logic. Thus, assume that the decision procedures for the universal fragments of T_1 and T_2 are in ExpTime, and that T_0 is the theory of semilattices. The combined decision procedure then generates doubly-exponentially many decision problems of exponential size for the component procedures. Each of these component decision problems can be decided in double-exponential time. Thus, the overall complexity of the combined decision procedure is 2ExpTime.

5 A variant of the connection scheme

Here we consider a slightly different combination scheme where a theory T is connected with itself rather than with a copy of itself. Let $T_0 \subseteq T$ be theories over the respective signatures $\Omega_0 \subseteq \Omega$. We use $T_{>T_0}$ to denote the theory whose models are models $\mathcal{M}$ of T endowed with a homomorphism $h : \mathcal{M}_{|\Omega_0} \to \mathcal{M}_{|\Omega_0}$. Thus, the *signature* Ω' of $T_{>T_0}$ is obtained from the signature Ω of T by adding a new function symbol h_S of arity SS for every sort S of Ω_0. The axioms of $T_{>T_0}$ are obtained from the axioms of T by adding

$$h_S(f(x_1, \ldots, x_n)) \approx f(h_{S_1}(x_1), \ldots, h_{S_n}(x_n)),$$
$$P(x_1, \ldots, x_n) \to P(h_{S_1}(x_1), \ldots, h_{S_n}(x_n)),$$

for every function (predicate) symbol f (P) in Ω_0 of arity $S_1 \ldots S_n S$ ($S_1 \ldots S_n$).

Example 5. An interesting example of a theory obtained as such a connection is the theory $E_{\mathbf{K}}$ corresponding to the basic modal logic $\mathbf{K}$ (see Example 1). In fact, let T be the theory of Boolean algebras, and T_0 the theory of semilattices over the signature Ω_0 as defined in Example 1. If we use the symbol $\Diamond$ for the connection function, then $T_{>T_0}$ is exactly the theory $E_{\mathbf{K}}$.

Theorem 2. *Let T_0, T be theories over the respective signatures Ω_0, Ω, where Ω_0 is a subsignature of Ω. Assume that $T_0 \subseteq T$, that T_0 is universal and effectively locally finite, and that T is T_0-algebraically compatible. Then the decidability of the universal fragment of T entails the decidability of the universal fragment of $T_{>T_0}$.*

To prove the theorem, we consider a finite set $\Gamma \cup \Gamma_0$ of ground flat literals over the signature Ω' of $T_{>T_0}$, where Γ is a set of literals in the signature Ω of T (expanded with free constants), and Γ_0 is of the form

$$\Gamma_0 = \{h(a_1) \approx b_1, \ldots, h(a_n) \approx b_n\}.$$

The theorem is an easy consequence of the following proposition, whose proof is similar to the one of Proposition 1.

Proposition 2. *The constraint $\Gamma \cup \Gamma_0$ is satisfiable in $T_{>T_0}$ iff there exists a triple $(\mathcal{A}, \mathcal{B}, \nu)$ such that*

1. *$\mathcal{A}$ is an Ω_0-model of T_0, which is generated by $\{a_1^{\mathcal{A}}, \ldots, a_n^{\mathcal{A}}\}$;*
2. *$\mathcal{B}$ is an Ω_0-model of T_0, which is generated by $\{b_1^{\mathcal{B}}, \ldots, b_n^{\mathcal{B}}\}$;*
3. *$\nu : \mathcal{A} \to \mathcal{B}$ is an Ω_0-homomorphism such that $\nu(a_j^{\mathcal{A}}) = b_j^{\mathcal{B}}$ for $j = 1, \ldots, n$;*
4. *$\Gamma \cup \Delta_{\Omega_0}(\mathcal{A}) \cup \Delta_{\Omega_0}(\mathcal{B})$ is satisfiable in T.*

Applied to the connection of BA with itself w.r.t. the theory of semilattices considered in Example 5, the theorem shows that deciding the universal theory of $E_{\mathbf{K}}$ can be reduced to deciding the universal theory of BA. It is well-known that deciding the universal theory of $E_{\mathbf{K}}$ is equivalent to deciding global consequence in $\mathbf{K}$, and that deciding the universal theory of BA is equivalent to propositional reasoning. Thus, we have shown the (rather surprising) result that the global consequence problem in $\mathbf{K}$ can be reduced to purely propositional reasoning. However, if we directly apply the non-deterministic combination algorithm suggested by Proposition 2, then the complexity of the obtained decision procedure is worse then the known ExpTime-complexity [20] of the problem. The deterministic combination procedure described below overcomes this problem.

A deterministic combination procedure

As pointed out in [18], Nelson-Oppen style combination procedures can be made deterministic in the presence of a certain convexity condition. Let T be a theory over the signature Ω, and let Ω_0 be a subsignature of Ω. Following [21], we say that T is Ω_0*-convex* iff every finite set of ground Ω^X-literals (using additional free constants from X) T-entailing a disjunction of $n > 1$ ground Ω_0^X-atoms, already T-entails one of the disjuncts. Note that universal Horn Ω-theories are always Ω-convex. In particular, this means that equational theories (like BA) are convex w.r.t. any subsignature.

Let $T_0 \subseteq T$ be theories over the respective signatures Ω_0, Ω, where Ω_0 is a subsignature of Ω. If T is Ω_0-convex, then Theorem 2 can be shown with the help of a deterministic combination procedure. (The same is actually also true for Theorem 1, but will not explicitly be shown here.)

Let $\Gamma \cup \Gamma_0$ be a finite set of ground flat literals (with free constants) in the signature of $T_{>T_0}$; suppose also that Γ does not contain the symbol h and that $\Gamma_0 = \{h(a_1) \approx b_1, \dots, h(a_n) \approx b_n\}$. We say that Γ is Γ_0*-saturated* iff for every Ω_0-atom $\alpha(x_1, \dots, x_n)$, $T \cup \Gamma \models \alpha(a_1, \dots, a_n)$ implies $\alpha(b_1, \dots, b_n) \in \Gamma$.

Theorem 3. *Let T_0, T be theories over the respective signatures Ω_0, Ω, where Ω_0 is a subsignature of Ω. Assume that $T_0 \subseteq T$, that T_0 is universal and effectively locally finite, and that T is Ω_0-convex and T_0-algebraically compatible. Then the following deterministic procedure decides whether $\Gamma \cup \Gamma_0$ is satisfiable in $T_{>T_0}$ (where Γ, Γ_0 are as above):*

1. *Γ_0-saturate Γ;*
2. *check whether the Γ_0-saturated set $\widehat{\Gamma}$ obtained this way is satisfiable in T.*

The saturation process (and thus the procedure) terminates because T_0 is locally finite. In addition, if $\Gamma \cup \Gamma_0$ is satisfied in a model $\mathcal{M}$ of $T_{>T_0}$, then the reduct of $\mathcal{M}$ to the signature Ω obviously satisfies $\widehat{\Gamma}$. Conversely, if the Γ_0-saturated set $\widehat{\Gamma}$ is satisfiable in T, then one can use $\widehat{\Gamma}$ to construct a triple $(\mathcal{A}, \mathcal{B}, \nu)$ satisfying 1.–4 of Proposition 2 (see [2] for details).

Example 5 (continued). Let us come back to the connection of $T := BA$ with itself w.r.t. the theory T_0 of semilattices, which yields as combined theory the equational theory $E_{\mathbf{K}}$ corresponding to the basic modal logic $\mathbf{K}$. In this case, checking during the saturation process whether $T \cup \Gamma \models \alpha(\underline{a})$ amounts to checking whether a propositional formula ϕ_Γ (whose size is linear in the size of Γ) implies a propositional formula of the form $\psi_1 \Leftrightarrow \psi_2$, where ψ_1, ψ_2 are disjunctions of the propositional variables from $\underline{a}$. Since there are only exponentially many different formulae of the form $\psi_1 \Leftrightarrow \psi_2$, the saturation process needs at most exponentially many such propositional tests, and the size of the intermediate sets Γ and of the Γ_0-saturated set $\widehat{\Gamma}$ is at most exponential. However, all these sets contain only the free constants $\underline{a}$. Since propositional reasoning can be done in time exponential in the number of propositional variables, this shows that both the saturation process and the final satisfiability test of $\widehat{\Gamma}$ in T can be done in time exponential in the number of free constants $\underline{a}$.

Consequently, we have shown that Theorem 3 yields an ExpTime decision procedure for the global consequence relation in $\mathbf{K}$, which thus matches the known worst-case complexity of the problem.

6 Conditions on the connection functions

Until now, we have considered connection functions that are arbitrary homomorphisms. In this section we impose the additional conditions that the connection functions be surjective, embeddings, or isomorphisms: in this way, we obtain new combined theories, which we denote by $T_1 >^{s}_{T_0} T_2, T_1 >^{em}_{T_0} T_2, T_1 >^{iso}_{T_0} T_2$, respectively. For these combined theories one can show combination results that are analogous to Theorem 1: one just needs different compatibility conditions.

To treat embeddings and isomorphisms, we use the compatibility condition introduced in [11, 3] for the case of unions of theories (see also the introduction of this paper). Following [11, 3], we call this condition T_0-*compatibility* in the following.

Theorem 4. *Let T_0, T_1, T_2 be theories over the respective signatures $\Omega_0, \Omega_1, \Omega_2$, where Ω_0 is a common subsignature of Ω_1 and Ω_2. Assume that $T_0 \subseteq T_1$ and $T_0 \subseteq T_2$, and that T_0 is universal and effectively locally finite.*

1. *If T_2 is T_0-compatible, then the decidability of the universal fragments of T_1 and T_2 entails the decidability of the universal fragment of $T_1 >^{em}_{T_0} T_2$.*
2. *If T_1 and T_2 are T_0-compatible, then the decidability of the universal fragments of T_1 and T_2 entails the decidability of the universal fragment of $T_1 >^{iso}_{T_0} T_2$.*

A proof of this theorem, which is similar to the proof of Theorem 1, can be found in [2]. It is easy to see that the problem of deciding the universal fragment of $T_1 >^{iso}_{T_0} T_2$ is interreducable in polynomial time with the problem of deciding the universal fragment of $T_1 \cup T_2$. Consequently, the proof of part 2. of Theorem 4 yields an alternative proof of the combination result in [11].

To treat $T_1 >^{s}_{T_0} T_2$, we must dualize the notions "algebraic completion" and "algebraic compatibility" (see [2] for the definitions of these dual notions, and the formulation and proof of the corresponding combination result).

7 Conclusion

We have introduced a new scheme for combining many-sorted theories, and have shown under which conditions decidability of the universal fragment transfers from the component theories to their combination. Though this kind of combination has been considered before in restricted cases [13, 1, 24], it has not been investigated in the general algebraic setting considered here.

In this paper, we mainly concentrated on the simplest case of connecting many-sorted theories where there is just one connection function. The approach was then extended to the case of several independent connection functions, and to variants of the general combination scheme where the connection function must satisfy additional properties or where a theory is connected with itself.

On the one hand, our results are more general than the combination results for $\mathcal{E}$-connections of abstract description systems shown in [13] since they are not restricted to Boolean-based equational theories, which are closely related to abstract description systems (see Example 1). For instance, we have shown in Example 3 that any pair of theories T_1, T_2 extending a universal theory T_0 that is effectively locally finite and Gaussian satisfies the prerequisites of our transfer theorem. Examples of such theories having nothing to do with Boolean-based equational theories can be found in [3].

On the other hand, in the $\mathcal{E}$-connection approach introduced in [13], one usually considers not only the modal operator induced by a connecting relation E (see Example 1), but also the modal operator induced by its inverse E^{-1}. It is not adequate to express these two modal operators by independent connection functions going in different directions since this does not capture the relationships that must hold between them. For example, if $\Diamond$ is the diamond operator induced by the connecting relation E, and $\Box^-$ is the box operator induced by its inverse E^-, then the formulae $x \to \Box^- \Diamond x$ and $\Diamond \Box^- y \to y$ are valid in the $\mathcal{E}$-connection. In order to express these relationships in the algebraic setting without assuming the presence of the Boolean operators in the shared theory, one can replace the logical implication $\to$ by a partial order $\leq$, and require that $x \leq r(\ell(x))$ and $\ell(r(y)) \leq y$ hold for the connection functions r, ℓ generalizing the diamond and the inverse box operator. If ℓ, r are also order preserving, then this mean that ℓ, r is a pair of *adjoint functions* for the partial order $\leq$. This suggests a new way of connecting theories through pairs of adjoint functions. Again, we can show transfer of decidability provided that certain algebraic conditions are satisfied.

References

1. Farid Ajili and Claude Kirchner. A modular framework for the combination of symbolic and built-in constraints. In *Proceedings of Fourteenth International Conference on Logic Programming*, pages 331–345, Leuven, Belgium, 1997. The MIT Press.
2. Franz Baader and Silvio Ghilardi. Connecting many-sorted theories. LTCS-Report LTCS-05-04, TU Dresden, Germany, 2005.
See http://lat.inf.tu-dresden.de/research/reports.html.
3. Franz Baader, Silvio Ghilardi, and Cesare Tinelli. A new combination procedure for the word problem that generalizes fusion decidability results in modal logics. In *Proceedings of the Second International Joint Conference on Automated Reasoning (IJCAR'04)*, volume 3097 of *Lecture Notes in Artificial Intelligence*, pages 183–197, Cork (Ireland), 2004. Springer-Verlag.
4. Franz Baader, Carsten Lutz, Holger Sturm, and Frank Wolter. Fusions of description logics and abstract description systems. *Journal of Artificial Intelligence Research*, 16:1–58, 2002.
5. Franz Baader and Cesare Tinelli. A new approach for combining decision procedures for the word problem, and its connection to the Nelson-Oppen combina-

tion method. In *Proceedings of the 14th International Conference on Automated Deduction*, volume 1249 of *Lecture Notes in Artificial Intelligence*, pages 19–33, Townsville (Australia), 1997. Springer-Verlag.
6. Franz Baader and Cesare Tinelli. Deciding the word problem in the union of equational theories. *Information and Computation*, 178(2):346–390, 2002.
7. Chen-Chung Chang and H. Jerome Keisler. *Model Theory*. North-Holland, Amsterdam-London, IIIrd edition, 1990.
8. Eric Domenjoud, Francis Klay, and Christophe Ringeissen. Combination techniques for non-disjoint equational theories. In *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.
9. Camillo Fiorentini and Silvio Ghilardi. Combining word problems through rewriting in categories with products. *Theoretical Computer Science*, 294:103–149, 2003.
10. Jean H. Gallier. *Logic for Computer Science: Foundations of Automatic Theorem Proving*. Harper & Row, 1986.
11. Silvio Ghilardi. Model-theoretic methods in combined constraint satisfiability. *Journal of Automated Reasoning*, 33(3–4):221–249, 2004.
12. Marcus Kracht and Frank Wolter. Properties of independently axiomatizable bimodal logics. *The Journal of Symbolic Logic*, 56(4):1469–1485, 1991.
13. Oliver Kutz, Carsten Lutz, Frank Wolter, and Michael Zakharyaschev. $\mathcal{E}$-connections of abstract description systems. *Artificial Intelligence*, 156:1–73, 2004.
14. Michael Makkai and Gonzalo E. Reyes. *First-Order Categorical Logic*, volume 611 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1977.
15. 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.
16. Greg Nelson and Derek C. Oppen. Simplification by cooperating decision procedures. *ACM Trans. on Programming Languages and Systems*, 1(2):245–257, October 1979.
17. Tobias Nipkow. Combining matching algorithms: The regular case. *Journal of Symbolic Computation*, 12:633–653, 1991.
18. Derek C. Oppen. Complexity, convexity and combinations of theories. *Theoretical Computer Science*, 12:291–302, 1980.
19. Don Pigozzi. The join of equational theories. *Colloquium Mathematicum*, 30(1):15–25, 1974.
20. Edith Spaan. *Complexity of Modal Logics*. PhD thesis, Department of Mathematics and Computer Science, University of Amsterdam, The Netherlands, 1993.
21. Cesare Tinelli. Cooperation of background reasoners in theory reasoning by residue sharing. *Journal of Automated Reasoning*, 30(1):1–31, January 2003.
22. Cesare Tinelli and Christophe Ringeissen. Unions of non-disjoint theories and combinations of satisfiability procedures. *Theoretical Computer Science*, 290(1):291–353, January 2003.
23. Frank Wolter. Fusions of modal logics revisited. In *Advances in Modal Logic*. CSLI, Stanford, CA, 1998.
24. Calogero Zarba. Combining multisets with integers. In *Proc. of the 18th International Conference on Automated Deduction (CADE'18)*, volume 2392 of *Lecture Notes in Artificial Intelligence*, pages 363–376, Copenhagen (Denmark), 2002. Springer-Verlag.
25. Marek W. Zawadowski. Descent and duality. *Ann. Pure Appl. Logic*, 71(2):131–188, 1995.