Model-Theoretic Methods in Combined Constraint Satisfiability

Silvio Ghilardi
Dipartimento di Scienze dell'Informazione
Università degli Studi di Milano
Italy

March, 2004

Abstract

We extend Nelson-Oppen combination procedure to the case of theories which are compatible with respect to a common subtheory in the shared signature. The notion of compatibility relies on model completions and related concepts from classical model theory.

Keywords: Combination, Nelson-Oppen Procedure, Fusion, Superposition Calculus, Quantifier Elimination, Model Completions.

1 Introduction

Constraint solving problems, arising for instance in software verification, can be often formalized as the problem of proving *the unsatisfiability of a set of literals modulo a background theory* [1]. Given the heterogeneity of the involved datatypes, the background theory is usually the union of multiple theories each of them dealing for instance with lists, arrays, booleans, integer or rational numbers, etc. (see e.g. [24] for useful examples).Thus constraint solving problems naturally become *combination problems* and, as such, should hopefully be sensitive to a *modular* approach.

In more formal terms, this raises the following question: is it possible to reduce unsatisfiability of a set of literals modulo the union of two or more theories to instances of the same problem which are relative to only one of the component theories? Nelson-Oppen combination method [21], [24], [27] provides a rather simple and positive answer for the case in which the involved theories are *stably infinite* and operate on *disjoint* signatures. In the literature, there have been some attempts to drop such limitations: for instance, in [29] an asymmetric procedure is introduced which does not require stable infiniteness for all the component theories. In this paper we investigate the case in which the shared signature is not empty and we propose a new approach based on classical model theory (in Robinson's

style). We leave comparison with related work to Section 8 and we shall explain here the basic features of our approach, which relies on a thorough investigation about the general reasons why Nelson-Oppen method works.

In original Nelson-Oppen combination procedure we are given theories T_1, T_2 over disjoint signatures (with equality) Σ_1, Σ_2 and we are dealing with the problem of deciding the consistency of sets of sentences like

$$(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2) \tag{1}$$

where Γ_i $(i = 1, 2)$ are finite sets of ground literals in the signature Σ_i augmented with new *shared* free constants, say $\underline{a} = a_1, \ldots, a_n$. Let us call Σ_0 the common signature $\Sigma_1 \cap \Sigma_2$ (which is reduced in Nelson-Oppen case to the only equality predicate). If (1) has a model $\mathcal{M}$, then we can take the Robinson diagram $\Delta(\mathcal{A})$ of the Σ_0-substructure $\mathcal{A}$ of $\mathcal{M}$ generated by $\underline{a}$ and realize that the two set of sentences

$$T_1 \cup \Gamma_1 \cup \Delta(\mathcal{A}) \qquad \text{and} \qquad T_2 \cup \Gamma_2 \cup \Delta(\mathcal{A}) \tag{2}$$

are consistent as well, because $\mathcal{M}$ is a model of both of them. $\Delta(\mathcal{A})$ simply specifies exhaustively whether for any $i, j = 1, \ldots, n$, we have $\mathcal{M} \models a_i = a_j$ or $\mathcal{M} \models a_i \neq a_j$. Since there are only finitely many such 'arrangements' $\Delta(\mathcal{A})$, it is clear that if the consistency of the two sets (2) is not only a necessary but also a sufficient condition for the consistency of (1), then we actually have a combined decision procedure.

The consistency of both the sets in (2) is not precisely sufficient for the consistency of (1), however it is sufficient whenever T_1, T_2 are *stably infinite*: this means that every model of T_i $(i = 1, 2)$ embeds into an infinite model of T_i.[1] In other words, we have the following fact (here a new ingredient, namely *quantifier elimination*, comes into the picture): the shared universal theory T_0 (which is nothing but the pure equality theory) can be extended to a theory T_0^* in the same signature Σ_0 (namely the theory of an infinite set) which is conservative over the universal fragment of T_0 and which, in addition, enjoys quantifier elimination. In the standard model-theoretic terminology, T_0^* is a *model completion* of T_0 and stable infiniteness means that every model of T_i $(i = 1, 2)$ embeds into a model of $T_i \cup T_0^*$.

This extra hypothesis is indeed sufficient for the consistency of the two sets of sentences in (2) to imply the consistency of the set of sentences in (1), by a simple proof based on Robinson's Joint Consistency Theorem (see the Appendix below for the details).

What does this argument need to work? It needs that: (i) there is a universal theory T_0 in the shared signature Σ_0 which is contained in both T_1 and T_2; (ii) T_0 admits a model-completion T_0^*; (iii) every model of T_i embeds

[1] The usual definition of stable infiniteness required by the Nelson-Oppen procedure is not just this one, but it is equivalent to it by a simple compactness argument, see Proposition 4.2 below.

into a model of $T_i \cup T_0^*$; (iv) T_0 is effectively locally finite (see Section 6 for the definition of this extra condition which is needed for having finitely many arrangements to try).

This analysis makes clear that our argument is quite general and extends Nelson-Oppen procedure to non-disjoint cases: we shall use it in section 6 to obtain, as special cases, results concerning e.g. fusion decidability in propositional modal logic [33].

Whenever only hypothesis (iv) fails, we no longer have a combined decision procedure. In this case, the method presented in this paper might still have some independent interest: for instance, it can be used in order to limit partial residue exchange of cooperating reasoners to disjunctions of positive literals in the shared signature (see section 5 below) and it can be applied in saturation-based theorem-proving in order to block inferences with mixed languages premises, while still retaining refutational completeness (see section 7 below).

We shall begin the paper with a couple of sections recalling the basic model-theoretic ingredients in more detail and providing also examples from natural software verification theories.

2 Preliminaries

A *signature* Σ is a set of functions and predicate symbols (each of them endowed with the corresponding arity). We assume the binary equality predicate symbol '=' to be always present in any signature Σ. In order to avoid confusion, we use in the metalanguage the symbol $\equiv$ (instead of $=$) to mean identity of syntactic expressions.

The signature obtained from Σ by the addition of a set of new constants (that is, 0-ary function symbols) X is denoted by $\Sigma \cup X$ or by Σ^X. We have the usual notions of Σ-*term*, (full first order) -*formula*, -*atom*, -*literal*, -*clause*, -*positive clause*, etc.: e.g. atoms are just atomic formulas, literals are atoms and their negations, clauses are multisets of literals, positive clauses are multisets of atoms, etc. We usually improperly write clauses/positive clauses as disjunctions of the corresponding literals/atoms (in particular, the empty clause is always identified with the empty disjunction $\bot$ expressing syntactic falsity). Letters $\phi, \psi, \ldots$ are used for formulas, whereas letters $A, B, \ldots$ are used for literals and letters $C, D, \ldots$ are used for clauses. Terms, literals, clauses and formulas are called *ground* whenever variables do not appear in them. Formulas without free variables are called *sentences*. The universal (resp. existential) closure of a formula ϕ is the sentence obtained from ϕ by adding it a prefix of universal (resp. existential) quantifiers binding all variables which happen to have a free occurrence in ϕ. If E is a term or a literal or a clause, we occasionally use the standard automatic deduction notations $E_{|p}, E[s]_p$ to denote the subterm at position p in E and the syntactic

expression (term, literal or clause) resulting from E by the replacement of the subterm at position p by the term s.

A Σ-*theory* T is a set of sentences (called the axioms of T) in the signature Σ; however when we write $T \subseteq T'$ for theories, we may mean not just set-theoretic inclusion but the fact that all the axioms for T are logical consequences of the axioms for T'.

From the semantic side, we have the standard notion of a Σ-*structure* $\mathcal{A}$: this is nothing but a support set endowed with an arity-matching interpretation of the function and predicate symbols from Σ. We use $f^{\mathcal{A}}$ (resp. $P^{\mathcal{A}}$) to denote the interpretation of the function symbol f (resp. predicate symbol P) in the structure $\mathcal{A}$. The support set of a structure $\mathcal{A}$ is indicated by the notation $|\mathcal{A}|$. *Truth* of a Σ-formula in $\mathcal{A}$ is defined in any one of the standard ways (so that truth of a formula is equivalent to truth of its *universal* closure). A formula ϕ is *satisfiable* in $\mathcal{A}$ iff its *existential* closure is true in $\mathcal{A}$.[2]

If $\Sigma_0 \subseteq \Sigma$ is a sub-signature of Σ and if $\mathcal{A}$ is a Σ-structure, the Σ_0-*reduct* of $\mathcal{A}$ is the Σ_0-structure obtained from $\mathcal{A}$ by forgetting the interpretation of function and predicate symbols from $\Sigma \setminus \Sigma_0$.

A Σ-structure $\mathcal{A}$ is a *model* of a Σ-theory T (in symbols $\mathcal{A} \models T$) iff all axioms of T are true in $\mathcal{A}$; for models of a Σ-theory T we shall preferably use the letters $\mathcal{M}, \mathcal{N}, \ldots$ to distinguish them from arbitrary Σ-structures. If ϕ is a formula, $T \models \phi$ ('ϕ *is a logical consequence of* T') means that ϕ is true in any model of T (notice that $T \models \phi$ turns out to be equivalent to $T \models \forall \underline{x}\, \phi$, where $\forall \underline{x}\, \phi$ is the universal closure of ϕ). A Σ-theory T is *complete* iff for every Σ-sentence ϕ, either ϕ or $\neg\phi$ is a logical consequence of T; T is *consistent* iff it has a model (i.e. iff $T \not\models \bot$).

The problems we deal with are *word problems*, more precisely, given a Σ-theory T:

- the *word problem* for T is that of deciding whether $T \models A$ holds for a Σ-atom A;

- the *conditional word problem* for T is that of deciding whether $T \models C$ holds for a Σ-clause C containing exactly one positive literal;

- the *clausal word problem* for T is that of deciding whether $T \models C$ holds for a Σ-clause C;

- the *elementary word problem* for T is that of deciding whether $T \models \phi$ holds for a first order Σ-formula ϕ.

[2] In the literature, there is also another definition of satisfiability for formulas which are not sentences (one says that ϕ is satisfiable iff its universal closure is true in some structure). The definition we adopted is the natural one for papers whose main focus is on constraint satisfiability.

Finally, the *constraint satisfiability problem* for the constraint theory T is the problem of deciding whether (the conjunction of) a finite set of Σ-literals is satisfiable in a model of T. The complementary *constraint unsatisfiability problem* (i.e. the problem of deciding whether a finite set of Σ-literals is unsatisfiable in all the models of T) is easily reduced to the clausal word problem: notice in fact that T-unsatisfiability of $A_1 \wedge \cdots \wedge A_n$ is the same as the relation $T \models \neg\exists \underline{x}\,(A_1 \wedge \cdots \wedge A_n)$ (for the appropriate existential closure prefix $\exists \underline{x}$), i.e. as the relation $T \models \forall \underline{x}\,(\neg A_1 \vee \cdots \vee \neg A_n)$. Vice versa, $T \models C$ (where C is the clause $B_1 \vee \cdots \vee B_m$) is equivalent to $T \models \forall \underline{x}\, C$ and hence to the T-unsatisfiability of $\neg B_1 \wedge \cdots \wedge \neg B_m$. In the following, we shall prefer to use free constants instead of variables in constraint satisfiability problems, so that we (equivalently) redefine a constraint satisfiability problem for the theory T as the problem of *establishing the consistency of* $T \cup \Gamma$ *for a finite set* Γ *of ground* $\Sigma^{\underline{a}}$*-literals* (where $\underline{a}$ is a finite set of new constants).

An Σ-*embedding* (or, simply, an embedding) between two Σ-structures $\mathcal{A}$ and $\mathcal{B}$ is any mapping $\mu : |\mathcal{A}| \longrightarrow |\mathcal{B}|$ among the corresponding support sets satisfying the condition

$$(*) \qquad \mathcal{A} \models A \quad \text{iff} \quad \mathcal{B} \models A$$

for all $\Sigma^{|\mathcal{A}|}$-atoms A (here $\mathcal{A}$ is regarded as a $\Sigma^{|\mathcal{A}|}$-structure[3] by interpreting each additional constant $a \in |\mathcal{A}|$ into itself and $\mathcal{B}$ is regarded as a $\Sigma^{|\mathcal{A}|}$-structure by interpreting each additional constant $a \in |\mathcal{A}|$ into $\mu(a)$). Notice the following facts: (a) as we have equality in the language, an embedding is an injective function; (b) an embedding $\mu : \mathcal{A} \longrightarrow \mathcal{B}$ must be an algebraic homomorphism, that is for every n-ary function symbol f and for every $a_1, \ldots, a_n \in |\mathcal{A}|$, we must have $f^{\mathcal{B}}(\mu(a_1), \ldots, \mu(a_n)) = \mu(f^{\mathcal{A}}(a_1, \ldots, a_n))$;[4] (c) for an n-ary predicate symbol P we must have $(a_1, \ldots, a_n) \in P^{\mathcal{A}}$ iff $(\mu(a_1), \ldots, \mu(a_n)) \in P^{\mathcal{B}}$. It is easily seen that an embedding $\mu : \mathcal{A} \longrightarrow \mathcal{B}$ can be equivalently defined as a mapping $\mu : |\mathcal{A}| \longrightarrow |\mathcal{B}|$ satisfying (a)-(b)-(c) above.

If the embedding $\mu : \mathcal{A} \longrightarrow \mathcal{B}$ is an inclusion map, we say that $\mathcal{A}$ is a *substructure* of $\mathcal{B}$ or that $\mathcal{B}$ is an *extension* of $\mathcal{A}$.

In case $(*)$ holds for all first order formulas, the embedding μ is said to be an *elementary* embedding. Correspondingly, in case μ is also an inclusion, we say that $\mathcal{A}$ is an elementary substructure of $\mathcal{B}$ or that $\mathcal{B}$ is an elementary extension of $\mathcal{A}$.

The *diagram* $\Delta(\mathcal{A})$ of a Σ-structure $\mathcal{A}$ is the set of ground $\Sigma^{|\mathcal{A}|}$-literals which are true in $\mathcal{A}$; the *elementary diagram* $\Delta^e(\mathcal{A})$ of a Σ-structure $\mathcal{A}$ is the

[3] We recall that we use $|\mathcal{A}|$ to denote the support of $\mathcal{A}$ and that we use Σ^X (or sometimes also $\Sigma \cup X$) to denote the signature obtained from Σ by adding it the set X as new free constants. Hence, the signature $\Sigma^{|\mathcal{A}|}$ (or $\Sigma \cup |\mathcal{A}|$) is the signature obtained from Σ by adding it new constants naming the elements of the support of $\mathcal{A}$.

[4] To see this, apply $(*)$ to the $\Sigma^{|\mathcal{A}|}$-atom $f(a_1, \ldots, a_n) = a$, where $a \in |\mathcal{A}|$ is just $f^{\mathcal{A}}(a_1, \ldots, a_n)$.

set of $\Sigma^{|\mathcal{A}|}$-sentences which are true in $\mathcal{A}$. Robinson (elementary) diagram theorem [11] says that there is an (elementary) embedding between the Σ-structures $\mathcal{A}$ and $\mathcal{B}$ iff it is possible to expand $\mathcal{B}$ to a $\Sigma^{|\mathcal{A}|}$-structure in such a way that it becomes a model of the (elementary) diagram of $\mathcal{A}$. This theorem (in combination with the compactness of the logical consequence relation) will be repeatedly used without explicit mention in the paper; a typical standard use is the following: suppose that we want to embed $\mathcal{A}$ into a model of a theory T, then it is sufficient to check that $T \cup \Delta(\mathcal{A})$ is consistent.

A formula is *quantifier-free* (or open) iff it does not contain quantifiers. A Σ-theory T is said to *eliminate quantifiers* iff for every formula $\phi(\underline{x})$[5] there is a quantifier-free formula $\phi'(\underline{x})$ such that

$$T \models \phi(\underline{x}) \leftrightarrow \phi'(\underline{x}).$$

There are many well-known theories eliminating quantifiers [11], we give here some examples which can be of interest for software verification.

Example 1. Linear integer arithmetic (i.e. the theory of the structure of integer numbers in the signature $+, 0, 1, \leq, \equiv_n$) eliminates quantifiers [12], [24]; so does rational linear arithmetic [32](i.e. the theory of rational numbers in the signature $+, 0, \leq$). Another well-known classical example from Tarski is real arithmetic (i.e. the theory of real numbers in the signature $+, 0, \cdot, 1, \leq$).

Example 2. The theory of acyclic lists L [24], [25] eliminates quantifiers;[6] this is the theory in the signature

$$\Sigma_L = \{car, cdr, cons\}$$

consisting on the universal closures of the following axioms:

$$\begin{array}{l} cons(car(x), cdr(x)) = x \\ car(cons(x, y)) = x \\ cdr(cons(x, y)) = y \\ x \neq t(x) \end{array}$$

where t is a term built up from x by using finitely many applications of the unary function symbols car, cdr. Notice that all models of L are infinite: in fact, by the first three axioms, the interpretation of the symbol *cons* realizes, in any model $\mathcal{M}$ of L, a bijection $\mathcal{M} \times \mathcal{M} \longrightarrow \mathcal{M}$ and such a bijection cannot exist if the support of $\mathcal{M}$ is a finite set having more than one element (it cannot have just one element because $L \models \forall x(car(x) \neq x)$).

[5] By this notation, we mean that ϕ contains free variables only among the finite set $\underline{x}$.

[6] Quantifier elimination for this theory seems to be well-known (it is included in old Mal'cev work [20]). See in any case the Appendix A in [15] for a full proof.

3 Model Completions

The main ingredient of this paper is the well-known notion of a *model completion* of a theory. There are good chapters on that in all textbooks from Model Theory. We shall recall here just the essential definitions (readers may consult e.g. [19],[11], [30] for further information).

Definition 3.1 *Let T be a Σ-theory and let $T^* \supseteq T$ be a further Σ-theory; we say that T^* is a model completion of T iff (i) every model of T can be embedded into a model of T^* and (ii) for every model $\mathcal{M}$ of T, we have that $T^* \cup \Delta(\mathcal{M})$ is a complete theory (in the signature $\Sigma^{|\mathcal{M}|}$).*

It can be shown that a model completion T^* of a theory T is unique, in case it exists, see [11]. Model completions are related with quantifier elimination as shown in the following easy proposition:

Proposition 3.2 *Let T be a Σ-theory and let $T^* \supseteq T$ a further Σ-theory; we have that T^* is a model completion of T in case (a) every model of T can be embedded into a model of T^* and (b) T^* eliminates quantifiers.*

Proof. Suppose that T^* satisfies conditions (a)-(b) from the statement of the proposition and let us prove that for every model $\mathcal{M}$ of T, we have that $T^* \cup \Delta(\mathcal{M})$ is a complete theory. Consider two models $\mathcal{N}_1, \mathcal{N}_2$ of such a theory, a Σ-formula $\phi(\underline{x})$ and a tuple $\underline{a}$ of elements from $|\mathcal{M}|$: we show that the $\Sigma^{|\mathcal{M}|}$-sentence $\phi(\underline{a})$ (resulting from $\phi(\underline{x})$ by replacing the variables $\underline{x}$ by the names of the $\underline{a}$'s) is true in $\mathcal{N}_1$ iff it is true in $\mathcal{N}_2$: as $\phi(\underline{a})$ is an arbitrary $\Sigma^{|\mathcal{M}|}$-sentence, this shows that the theory $T^* \cup \Delta(\mathcal{M})$ is indeed complete.

Now (by the Robinson Diagram Theorem) $\mathcal{M}$ is a common substructure of $\mathcal{N}_1$ and $\mathcal{N}_2$; moreover $\phi(\underline{x})$ is T^*-equivalent to a quantifier free formula $\phi'(\underline{x})$, hence if $\phi(\underline{a})$ is true in $\mathcal{N}_1$, $\phi'(\underline{a})$ is true in it, $\phi'(\underline{a})$ is true in $\mathcal{M}$ and in $\mathcal{N}_2$, as well,[7] thus establishing that $\phi(\underline{a})$ is true in $\mathcal{N}_2$.

It can be shown that in the case in which T is a *universal* theory (namely, a theory having as axioms only universal closures of quantifier-free formulas), conditions (a)-(b) from Proposition 3.2 are actually equivalent to conditions (i)-(ii) from Definition 3.1 (see for instance [15], [17] or also [11] for a proof). We shall not need such an equivalence in the following, although in this paper we shall consider model completions of universal theories only.

Example 3. The theory of an infinite set is the model completion of the pure theory of equality in the minimum signature containing only the equality predicate; the theory of dense total orders without endpoints is the model completion of the theory of total orders.

[7] Recall that ground formulas are preserved by both substructures and extensions: in fact, such formulas are Boolean combinations of ground atoms and $(*)$ from Section 2 consequently applies.

Example 4. There are many classical examples from algebra: the theory of algebraically closed fields is the model completion of the theory of fields, the theory of divisible torsion free abelian groups is the model completion of the theory of torsion free abelian groups, etc.

Example 5. The theory of atomless Boolean algebras is the model completion of the theory of Boolean algebras (for model completions arising in the algebra of logic, see the book [17]). We recall that an atom in a Boolean algebra $\mathcal{B}$ is a non-zero element which is minimal (among non-zero elements) with respect to the partial order induced by the lattice structure. A Boolean algebra $\mathcal{B}$ is atomless iff it has no atoms.

Example 6. An old result in [30] says, in particular, that universal Horn theories T in finite signatures always have a model completion, provided the following two conditions are satisfied: (a) finitely generated models of T are all finite; (b) amalgamation property holds for models of T.[8] This fact can be used in order to prove the existence of a model completion for theories axiomatizing many interesting discrete structures (like graphs, posets, etc.).

Example 7. It follows from the quantifier elimination result for the theory L of acyclic lists, that this theory is the model completion of itself.

Example 8. If a theory T^* has elimination of quantifiers, then it is the model completion of the theory T axiomatized by the set of universal sentences which are logical consequences of T^*, see [11].

4 T_0-compatibility

The key notion for our combination procedure is the following:

Definition 4.1 *Let T be a theory in the signature Σ and let T_0 be a universal theory in a subsignature $\Sigma_0 \subseteq \Sigma$ admitting a model completion T_0^*. We say that T is* T_0-compatible *iff*

(i) $T_0 \subseteq T$;

(ii) *every model of T can be embedded into a model of $T \cup T_0^*$.*

Condition (ii) can be equivalently stated in a slightly different form:

[8] Amalgamation property for models of T says that for every couple of embeddings $\mu_1 : \mathcal{M}_0 \longrightarrow \mathcal{M}_1$, $\mu_2 : \mathcal{M}_0 \longrightarrow \mathcal{M}_2$ among models of T, there are a model $\mathcal{M}$ of T and further embeddings $\nu_1 : \mathcal{M}_1 \longrightarrow \mathcal{M}$, $\nu_2 : \mathcal{M}_2 \longrightarrow \mathcal{M}$ forming a commutative square. Notice that the existence of a model completion for T always implies amalgamation property for models of T (this is easy to prove, see once again [11]).

Proposition 4.2 *Let T be a theory in the signature Σ and let $T_0 \subseteq T$ be a universal theory in a subsignature $\Sigma_0 \subseteq \Sigma$. If T_0 has a model-completion T_0^*, then T is T_0-compatible iff every quantifier-free Σ-formula which is satisfiable in a model of T is satisfiable also in a model of $T \cup T_0^*$.*

Proof. This is just an immediate application of the following well-known lemma (take T' to be $T \cup T_0^*$).

Lemma 4.3 *Let T and T' be both Σ-theories. The following two conditions are equivalent:*

(i) *every model of T can be embedded into a model of T';*

(ii) *every quantifier-free Σ-formula which is satisfiable in a model of T is satisfiable also in a model of T'.*

Proof. It is evident that if every model of T embeds into a model of T', then every quantifier-free Σ-formula which is satisfiable in a model of T is satisfiable in a model of T' too. The converse is by a compactness argument: let $\mathcal{M}$ be a model of T, we show that $T' \cup \Delta(\mathcal{M})$ is consistent. If not, we have $T' \models \neg\phi(\underline{a})$, where $\phi(\underline{a})$ is a finite conjunction of formulas from $\Delta(\mathcal{M})$. In more detail, this means that there are a quantifier-free Σ-formula $\phi(\underline{x})$, a finite set of elements $\underline{a}$ from the support of $\mathcal{M}$ such that $\phi(\underline{a})$ is true in $\mathcal{M}$ (here $\phi(\underline{a})$ is obviously the ground $\Sigma^{\underline{a}}$–formula resulting from $\phi(\underline{x})$ after the replacements $\underline{x} \mapsto \underline{a}$). As the constants $\underline{a}$ do not belong to the signature Σ, we have that $T' \models \forall\underline{x}\neg\phi(\underline{x})$. But $\phi(\underline{x})$ is a quantifier-free formula satisfiable in a model of T and, by hypothesis, there is a model of T' in which $\phi(\underline{x})$ is satisfiable too, contrary to the fact that $T' \models \neg\exists\underline{x}\phi(\underline{x})$.

Example 9. According to Proposition 4.2, it is evident that T_0-compatibility reduces to the standard notion of stable infiniteness (used in the disjoint Nelson-Oppen combination procedure) in case T_0 is the pure theory of equality (recall in fact that in this case, T_0^* is the theory of an infinite set).

Example 10. Every theory including the theory L of acyclic lists is L-compatible, because L is universal and $L = L^*$.

Example 11. If T_0 is universal, has a model completion T_0^* and if $T \supseteq T_0^*$, then T is certainly T_0-compatible: this trivial case is often interesting (we may take e.g. T_0 to be the theory of linear orders and T to be real arithmetic or linear rational arithmetic).

Example 12. Let T_0 be a universal theory having a model completion T_0^* and let T be any extension of T_0 with free function symbols only. In this case, T is T_0-compatible: to see it, take any model $\mathcal{M}$ of T, embed Σ_0-reduct of $\mathcal{M}$ into a model $\mathcal{M}'$ of T_0^* (see Definition 3.1(i)) and expand in any arbitrary way the interpretation of the free function symbols to the tuples of $\mathcal{M}'$ not entirely belonging to $\mathcal{M}$.

More examples will be supplied in Section 6. An interesting feature of T_0-compatibility is that it is a *modular* property; we formally state this fact now and prove it in the Appendix:

Proposition 4.4 *Let T_1 be a Σ_1-theory and let T_2 be a Σ_2-theory; suppose they are both compatible with respect to a Σ_0-theory T_0 (where $\Sigma_0 := \Sigma_1 \cap \Sigma_2$). Then $T_1 \cup T_2$ is T_0-compatible too.*

5 Combining compatible theories

Let us progressively fix our main data for the whole paper.

Assumption (I). *T_1 is a theory in the signature Σ_1 and T_2 is a theory in the signature Σ_2; Σ_0 is the signature $\Sigma_1 \cap \Sigma_2$.*

Our main aim is that of (semi)deciding the constraint unsatisfiability problem for $T_1 \cup T_2$, given that the corresponding constraint unsatisfiability problems for T_1 and T_2 are (semi)decidable. This amounts to (semi)decide the inconsistency of

$$T_1 \cup T_2 \cup \Gamma, \tag{1}$$

where Γ is a finite set of ground literals in the signature $\Sigma_1 \cup \Sigma_2$, expanded with a finite set of new free constants.

Γ can be *purified*. In fact, for any literal $A \in \Gamma$, we have that (1) is equi-satisfiable with

$$T_1 \cup T_2 \cup (\Gamma \setminus \{A\}) \cup \{A[c]_p,\ c = A_{|p}\}, \tag{2}$$

where p is a term position in A and c is a further new constant. After finitely many transformations like that from (1) to (2), we can reduce our problem to that of establishing the consistency of a set of sentences like

$$(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2), \tag{3}$$

where Γ_1, Γ_2 are as explained in the following:[9]

Assumption (II). *For finitely many new free constants $\underline{a}$, Γ_1 is a finite set of ground literals in the signature $\Sigma_1^{\underline{a}}$ and Γ_2 is a finite set of ground literals in the signature $\Sigma_2^{\underline{a}}$.*

[9] We leave the reader the details of this standard step. Of course, one can apply the above transformation until in each literal there is at most one symbol which is not the identity predicate or a free constant, but this method (although correct) would be inefficient. A better solution, which is commonly adopted, is that of abstracting only 'alien' subterms; however, there are many different possibilities for defining what an alien subterm is in case the signatures are not disjoint (see [5] for a detailed discussion). Notice that at the end of this purification process, Σ_0-literals could be inserted either in Γ_1 or in Γ_2 or in both of them, indifferently.

Clearly the consistency of (3) cannot follow from the mere separate consistency of $T_1 \cup \Gamma_1$ and of $T_2 \cup \Gamma_2$ (for trivial reasons, take e.g. $T_1 = T_2 = \emptyset$, $\Gamma_1 = \{a_1 = a_2\}$ and $\Gamma_2 = \{a_2 = a_3,\ a_1 \neq a_3\}$). To reason modularly, we need some *information exchange* between a reasoner dealing with $T_1 \cup \Gamma_1$ and a reasoner dealing with $T_2 \cup \Gamma_2$.

Craig's interpolation theorem for first order logic ensures that the inconsistency of (3) can be detected by the information exchange of a single Σ_0^a-sentence ϕ such that $T_1 \cup \Gamma_1 \models \phi$ and $T_2 \cup \Gamma_2 \cup \{\phi\} \models \bot$. However, as pointed out in [26], this observation is not very useful, as ϕ might be any first-order formula, whereas we would like - at least - ϕ to be quantifier-free: recall that most existing provers detect inconsistency just by skolemization and saturation, so they are not designed to directly find such a ϕ (if it is not quantifier-free), even in case they can efficiently handle with both $T_1 \cup \Gamma_1$ and $T_2 \cup \Gamma_2$.

Unfortunately, information exchange of Σ_0^a-quantifier free formulas alone is not sufficient, even for syntactically simple T_1 and T_2, to establish the inconsistency of (3) (a counterexample will be supplied in Section 7 below). We consequently need a further assumption in order to get limited information exchange without affecting refutational completeness (this is the only relevant assumption we make, the other two being mere notational conventions):

Assumption (III). *There is a universal Σ_0-theory T_0 such that both T_1 and T_2 are T_0-compatible.*

A finite list

$$C_1, \ldots, C_n$$

of positive ground Σ_0^a-clauses such that for every $k \in \{1, \ldots, n\}$, there is $i \in \{1, 2\}$ such that

$$T_i \cup \Gamma_i \cup \{C_1, \ldots, C_{k-1}\} \models C_k.$$

is called a *positive residue chain.*[10] We can now formulate our main combination results, whose proofs can be found in the Appendix 9:

Theorem 5.1 *With the above assumptions, $(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2)$ is inconsistent iff there is a positive residue chain $C_1, \ldots, C_n$ such that C_n is the empty clause.*

Thus inconsistency can be detected by repeated exchanges of positive ground clauses only; if we allow information exchange consisting of ground sentences, a single exchange step is sufficient:

[10] In the disjoint signatures case, the concept of a positive residue chain is a declarative formalization of the back-and-forth equality propagation mechanism of the original Nelson-Oppen procedure. For the use of positive ground clauses as residues in reasoners' cooperation (within tableaux methods), see [34].

Theorem 5.2 *With the above assumptions,* $(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2)$ *is inconsistent iff there is a ground* $\Sigma_0^{\underline{a}}$*-sentence* ϕ *such that*

$$T_1 \cup \Gamma_1 \models \phi \qquad \textit{and} \qquad T_2 \cup \Gamma_2 \cup \{\phi\} \models \bot.$$

Following [26], we say that our T_i's are Σ_0-convex iff whenever $T_i \cup \Gamma_i \models A_1 \vee \cdots \vee A_n$ (for $n \geq 1$ and for ground $\Sigma_0^{\underline{a}}$-atoms $A_1, \ldots, A_n$), then there is $k \in \{1, \ldots, n\}$ such that $T_i \cup \Gamma_i \models A_k$.[11] For Σ_0-convex theories, an immediate subsumption argument refines Theorem 5.1 in the following way:

Corollary 5.3 *In addition to the above assumptions, suppose also that* T_1, T_2 *are both* Σ_0*-convex. Then* $(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2)$ *is inconsistent iff there is a positive residue chain* $C_1, \ldots, C_n$ *in which* $C_1, \ldots, C_{n-1}$ *are all ground* Σ_0*-atoms and* C_n *is* $\bot$*.*

6 Extensions of Nelson-Oppen combination procedure

We say that a universal Σ_0-theory T_0 is *locally finite* iff Σ_0 is finite and for every finite set $\underline{a}$ of new free constants, there are finitely many $\Sigma_0^{\underline{a}}$-ground terms $t_1, \ldots, t_{k_{\underline{a}}}$ such that for every further $\Sigma_0^{\underline{a}}$-ground term u, we have $T_0 \models u = t_i$ (for some $i \in \{1, \ldots, k_{\underline{a}}\}$). If such $t_1, \ldots, t_{k_{\underline{a}}}$ are effectively computable from $\underline{a}$, then T_0 is said to be *effectively locally finite*. Examples of effectively locally finite theories are the theory of graphs, of partial orders (more generally, any theory whose signature does not contain function symbols), of commutative idempotent monoids, of Boolean algebras, etc.

In a locally finite theory T_0, there are restricted *finite* classes which are *representatives*, up to T_0-equivalence, of the whole classes of $\Sigma_0^{\underline{a}}$-ground literals, clauses, quantifier free sentences, etc. (they are just the ground literals, clauses, quantifier free sentences, etc. containing only the above mentioned terms $t_1, \ldots, t_{k_{\underline{a}}}$). As it is evident that we can limit information exchange to ground positive clauses and quantifier-free sentences in that restricted representative class, both Theorems 5.1 and 5.2 yield *combined decision procedures for the constraint satisfiability problem* in $T_1 \cup T_2$ (in case T_0 is effectively locally finite, both T_1 and T_2 are T_0-compatible and in case the corresponding constraint satisfiability problems for T_1 and T_2 are separately decidable).

The procedure suggested by Theorem 5.1 is just a fair information exchange of positive ground $\Sigma_0^{\underline{a}}$-clauses, to be continued until the situation gets stable or until an inconsistency is detected. Notice that in case T_1, T_2 are Σ_0-convex theories, information exchange can be further limited to ground

[11] Among Σ_0-convex theories we have the important class of universal Horn theories, see [26] again.

$\Sigma_0^{\underline{a}}$-atoms by Corollary 5.3.[12] More formally, the algorithm suggested by Theorem 5.1 can be described as follows:

Algorithm 6.1

Step 1: Purify the finite **input** *set of ground literals* Γ*, thus producing, for some finite set* $\underline{a}$ *of free constants, a finite set* Γ_1 *of ground* $\Sigma_1^{\underline{a}}$*-literals and a finite set* Γ_2 *of ground* $\Sigma_2^{\underline{a}}$*-literals (then* $\Gamma_1 \cup \Gamma_2$ *is* $T_1 \cup T_2$*-equisatisfiable with* Γ*). In the next loop, positive ground* $\Sigma_0^{\underline{a}}$*-clauses are added to* Γ_1, Γ_2*.*

Step 2: Using the decision procedures for T_1, T_2*, check whether* $T_1 \cup \Gamma_1$ *and* $T_2 \cup \Gamma_2$ *are consistent or not (if one of them is not,* **return** *'unsatisfiable').*

Step 3: If $T_i \cup \Gamma_i$ *entails some representative positive ground* $\Sigma_0^{\underline{a}}$*-clause (atom in the* Σ_0*-convex case) not entailed by* $T_j \cup \Gamma_j$ *(*$j \neq i$*) add this positive ground clause (atom) to* Γ_j *and go back to Step 2.*

Step 4: If this step is reached, **return** *'satisfiable'.*

By contrast, the procedure suggested by Theorem 5.2 (which is nothing but an interpolation theorem) identifies all ground $\Sigma_0^{\underline{a}}$-clauses which are logical consequences of $T_1 \cup \Gamma_1$ and check whether their conjunction is consistent with $T_2 \cup \Gamma_2$. There is a third possible (non-deterministic) procedure, which is suggested in the introduction of the paper and which is justified directly by Lemma 9.4 from Appendix 9: as there are only finitely many Σ_0-structures generated by $\underline{a}$ which are models of T_0 (recall that such structures cannot have more than $k_{\underline{a}}$-elements), one simply guesses one of them and check whether its diagram is consistent with both $T_1 \cup \Gamma_1$ and $T_2 \cup \Gamma_2$.

We give a first example to which the above outlined combined procedures apply.

Example 13. Let T_1 be linear rational arithmetic and let T_2 be the theory of total orders endowed with a strict monotonic function f.[13] We take as T_0 the theory of total orders (recall that its model completion T_0^* is the theory of dense total orders without endpoints). T_1 is known to be decidable [10]. We leave the reader the little exercise to prove that the constraint satisfiability problem for T_2 is decidable too: the relevant lemma to be proved shows that any finite total order endowed with a *partial* strict monotonic function embeds into a model of T_2 (this is shown by successively inserting new points and by taking union in the limit). Once this lemma is proved, the satisfiability of a set Γ_2 of $\Sigma_2^{\underline{a}}$-ground literals can be decided by a non-deterministic guessing of such a finite total order endowed with a partial strict monotonic function. As $T_1 \supseteq T_0^*$, T_1 is certainly T_0-compatible. We only need to show that T_2 is T_0-compatible, by embedding each model

[12] This observation, as shown in [24] for the disjoint signatures case, may improve complexity bounds in certain significant situations.

[13] This means that f is constrained by the axiom $\forall x \forall y\, (x < y \rightarrow f(x) < f(y))$.

$\mathcal{M}$ of T_2 into a model $\mathcal{M}'$ of $T_0^* \cup T_2$. It is sufficient to take as $\mathcal{M}'$ the lexicographic product of $\mathcal{M}$ with e.g. the poset of rational numbers[14] (the symbol f is interpreted by putting $f^{\mathcal{M}'}(b,m) = (f^{\mathcal{M}}(b),m)$, the embedding $\mathcal{M} \longrightarrow \mathcal{M}'$ is defined by associating with $b \in \mathcal{M}$ the pair $(b,0)$). Thus our combination results apply and we obtain the decidability of the constraint satisfiability problem for rational linear arithmetic endowed with a strict monotonic function. It is not difficult to see that the complexity of this combined decision algorithm lies in the NP-class (just adapt the arguments in [24]).

In order to explain our applications to fusion decidability in modal logic, we need to fix some terminology. We shall not directly introduce modal logic, rather we insist on the algebraic counterpart of modal logic (this choice makes exposition simpler for the purposes of this paper). A *modal algebra* is just a Boolean algebra $\mathcal{B} = \langle B, \cap, 1, \cup, 0, (-)' \rangle$ endowed with an hemimorphism $\Box$ (a hemimorphism is a function preserving only meets and the top element). Hemimorphisms are also called modal (necessity) operators. The modal operator $\Box$ is said to be *transitive* iff the identity $\Box a \leq \Box\Box a$ holds for every $a \in B$.[15]

Let now Σ_1 be the signature of Boolean algebras augmented with a unary function symbol $\Box_1$ and let Σ_2 be the signature of Boolean algebras augmented with a unary function symbol $\Box_2$. $\mathcal{V}_1$ is the equational theory of a variety V_1 of modal algebras and $\mathcal{V}_2$ is the equational theory of another variety V_2 of modal algebras. For $i \in \{1,2\}$, $\mathcal{V}_i$ is a universal Horn theory, hence it is Σ_i-convex: this means in particular that the decidability of the conditional word problem for $\mathcal{V}_i$ is equivalent to the decidability of the clausal word problem for $\mathcal{V}_i$ (and consequently also to the decidability of the constraint satisfiability problem for $\mathcal{V}_i$).

Proposition 6.2 *Let $\mathcal{V}_1, \mathcal{V}_2$ be as above; then the decidability of the conditional word problem for $\mathcal{V}_1$ and $\mathcal{V}_2$ implies the decidability of the conditional word problem for $\mathcal{V}_1 \cup \mathcal{V}_2$.*

Proof. As already mentioned, we have for free the decidability of the constraint satisfiability problem in $\mathcal{V}_1$ and $\mathcal{V}_2$; we take as T_0 the theory of Boolean algebras (which is locally finite and admits as a model completion the theory of atomless Boolean algebras): in order to apply our results, we simply need to show that $\mathcal{V}_1, \mathcal{V}_2$ are T_0-compatible. We do it for $\mathcal{V}_1$. Let $\mathcal{M}$ be a model of $\mathcal{V}_1$, we show how to embed it into a model $\mathcal{M}'$ of $\mathcal{V}_1$ which is based on an atomless Boolean algebra: this is a well-known and rather trivial fact, which is used also in [33] as a side preliminary lemma. Instead

[14] For density, observe that $(b,m) <_{lex} (b',m')$ implies $(b,m) <_{lex} (b,n) <_{lex} (b',m')$, where $n = \frac{m+m'}{2}$ in case $m < m'$ and $n = m+1$ otherwise (notice that in the latter case, we must have $b < b'$).

[15] Recall that, in a Boolean algebra, $y \leq z$ is a shorthand for $y \cap z = y$.

of reporting the argument used in [33] and in [3] too, we give a more direct one.[16] Define a sequence of models of $\mathcal{V}_1$ by: $\mathcal{M}_0 := \mathcal{M}$, $\mathcal{M}_{k+1} := \mathcal{M}_k \times \mathcal{M}_k$; define also embeddings $\delta_k : \mathcal{M}_k \longrightarrow \mathcal{M}_{k+1}$ by $\delta_k(a) := \langle a, a \rangle$. Now take as $\mathcal{M}'$ the union (more precisely, the inductive limit) of this chain: clearly $\mathcal{M}'$ is atomless as a Boolean algebra (no non-zero element is minimal in it, as any $a \in \mathcal{M}_k$ gets identified with $\langle a, a \rangle = \langle a, 0 \rangle \cup \langle 0, a \rangle$ in $\mathcal{M}_{k+1}$).

Corollary 6.3 *Let $\mathcal{V}_1, \mathcal{V}_2$ be as above; if the modal operators $\Box_1$, $\Box_2$ are transitive, then the decidability of the word problem for $\mathcal{V}_1$ and $\mathcal{V}_2$ implies the decidability of the word problem for $\mathcal{V}_1 \cup \mathcal{V}_2$.*

Proof. This is simply due to the fact that, because of the transitivity of the modal operators, the decidability of the word problem in $\mathcal{V}_i$ implies the decidability of the conditional word problem in $\mathcal{V}_i$ ($i = 1, 2$): in fact the $\mathcal{V}_i$-validity of the conditional equation

$$t_1 \neq u_1 \vee \cdots \vee t_n \neq u_n \vee t = u$$

is equivalent to the $\mathcal{V}_i$-validity of the single equation

$$\Box_i^+(t_1 \leftrightarrow u_1) \cap \cdots \cap \Box_i^+(t_n \leftrightarrow u_n) \leq (t \leftrightarrow u)$$

(where we used the abbreviations $\Box_i^+ x$ for $x \cap \Box_i x$ and $x \leftrightarrow y$ for $(x' \cup y) \cap (x \cup y')$), because of elementary facts concerning congruences in modal algebras.

The last corollary, once read in terms of logics, means exactly fusion decidability for normal extensions of $K4$. Although Proposition 6.2 and Corollary 6.3 do not entirely cover the fusion decidability results of [33], they put some substantial part of them into the appropriate general combination context. Quite recently, the full results of [33] have been considerably strengthened in [3], by joining the techniques presented in this paper with those explained in [5]. For further results (based on a refinement of the combination schema explained in this section) concerning fusion of modal logics sharing a universal modality and nominals, see [16].

7 Pure deductions in a saturation-based framework

Nowadays saturation-based theorem provers are powerful and efficient when treating some instances of specific constraint satisfiability problems: in [1], it is shown that superposition calculus always terminates and gives consequently a decision procedure for such problems, when the involved constraint

[16] We thank L. Santocanale for suggesting this elegant argument.

theory is a theory axiomatizing certain common datatypes such as lists or arrays. In this section we give some further suggestions about a possible use of the ideas explained in Section 5 within saturation-based theorem proving. We show that whenever T_0-compatibility holds, it is possible to cut in a deduction the inferences which are not pure, while still retaining refutational completeness.

An inference among $(\Sigma_1 \cup \Sigma_2)^{\underline{a}}$-clauses

$$\frac{C_1, \ldots, C_n}{C}$$

is *pure* iff there is $i = 1, 2$ such that all the clauses $C_1, \ldots, C_n, C$ are $\Sigma_i^{\underline{a}}$-clauses. Similarly, a deduction is pure iff all inferences in it are pure. Usually pure deductions are not able to detect inconsistency of (the skolemization of) sets of sentences like $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$, however we shall see that this may happen when the T_0-compatibility conditions are satisfied.

In order to realize this program, we first need to skolemize the theories T_1, T_2, thus passing to theories T_1^{sk}, T_2^{sk} in extended signatures $\Sigma_1^{sk}, \Sigma_2^{sk}$; Skolem functions will *not* be considered shared symbols, hence we still have that $\Sigma_0 = \Sigma_1^{sk} \cap \Sigma_2^{sk}$. The first problem we meet is the following: if T_i is T_0-compatible, is T_i^{sk} still T_0-compatible? We do not have a general answer for that, however there is a relevant case in which the answer is affirmative:

Lemma 7.1 *Let T be a Σ-theory which is T_0-compatible for some Σ_0-theory T_0 (here Σ_0 is a subsignature of Σ). If the axioms of T are all $\forall\exists$-sentences, then T^{sk} is T_0-compatible too.*

Proof. Let $\mathcal{M}$ be a model of T^{sk}; we want to embed it into a model of $T^{sk} \cup T_0^*$. As T is T_0-compatible, clearly the Σ-reduct of $\mathcal{M}$ is embeddable into a model $\mathcal{M}'$ of $T \cup T_0^*$. We only have extend the interpretation of the additional Skolem functions from $\mathcal{M}$ to $\mathcal{M}'$. Let $f_1, \ldots, f_m$ be the n-ary Skolem functions coming from the T-axiom[17]

$$\forall x_1 \cdots \forall x_n \exists y_1 \cdots \exists y_m \, \phi(x_1, \ldots, x_n, y_1, \ldots, y_m).$$

Pick $a_1, \ldots, a_n$ from $|\mathcal{M}'|$; as $\mathcal{M}' \models T$, there are $b_1, \ldots, b_m$ such that the quantifier-free ground formula $\phi(a_1, \ldots, a_n, b_1, \ldots, b_m)$ is true in $\mathcal{M}'$. In case $a_1, \ldots, a_n$ all belong to the support of $\mathcal{M}$, we choose b_i to be precisely $f_i^{\mathcal{M}}(a_1, \ldots, a_n)$ (this is possible because ϕ is quantifier-free, hence $\mathcal{M} \models \phi(a_1, \ldots, a_n, b_1, \ldots, b_m)$ entails $\mathcal{M}' \models \phi(a_1, \ldots, a_n, b_1, \ldots, b_m)$ in case the a_j, b_i's are all from the support of $\mathcal{M}$). We put $f_i^{\mathcal{M}'}(a_1, \ldots, a_n)$ equal to b_i and we are done.

[17] Recall that this axiom is replaced in T^{sk} by the axiom

$$\forall x_1 \cdots \forall x_n \, \phi(x_1, \ldots, x_n, f_1(x_1, \ldots, x_n), \ldots, f_m(x_1, \ldots, x_n))$$

containing the new Skolem function $f_1, \ldots, f_m$.

The case covered by Lemma 7.1 is important: recall from [11] that the model completion T_0^* of a universal theory T_0 must have a set of $\forall\exists$-axioms, hence the cases in which T_0-compatibility is trivially ensured by the fact that the axioms for T_0^* are included in the axioms of T are not compromized by the skolemization process (provided the remaining axioms for T are also $\forall\exists$-sentences).

Lemma 7.1 motivates the following extra assumption (in addition to the three assumptions from section 5):

Assumption (IV). *T_1, T_2 are axiomatized by $\forall\exists$-sentences; T_1^{sk}, T_2^{sk} are the skolemizations of T_1, T_2, respectively.*

We take into consideration here the *Superposition Calculus* $\mathcal{I}$ (see [7], [8] and the surveys [9], [22]). As commonly done within paramodulation-based theorem proving, we assume that identity is the only predicate[18] and we treat it symmetrically (this means that the equation $s = t$ is identified with the multiset $\{s, t\}$). Clauses[19] like

$$t_1 = u_1, \ldots, t_n = u_n \Rightarrow t'_1 = u'_1, \ldots, t'_m = u'_m$$

are often identified (e.g. for questions concerning orderings) with the multiset of multisets

$$\{\{t_1, t_1, u_1, u_1\}, \ldots, \{t_n, t_n, u_n, u_n\}, \{t'_1, u'_1\}, \ldots, \{t'_m, u'_m\}\}.$$

We fix a *lexicographic path ordering*[20] induced by a total precedence on the symbols of $\Sigma_1^{sk} \cup \Sigma_2^{sk} \cup \{\underline{a}\}$; assuming for simplicity that our signatures are finite, this induces a reduction ordering $>$ which is total on ground terms.[21] We give to symbols in $\Sigma_0^{\underline{a}}$ *lower precedence* than to symbols in $\Sigma_1^{sk} \setminus \Sigma_0$ and in $\Sigma_2^{sk} \setminus \Sigma_0$. This is essential: as a consequence, ground $\Sigma_0^{\underline{a}}$-clauses will be smaller in the twofold multiset extension of $>$ than all ground clauses containing a proper Σ_1- or Σ_2-symbol.

We recall below the inference rules of the Superposition Calculus $\mathcal{I}$: the application of each rule is subject to the satisfiability of certain *ordering constraints*, which are essential for the efficiency of the calculus; we do not report such constraints here for the sake of simplicity, we just remind that, roughly speaking, they *restrict the application of the rules to maximal*

[18] Atomic formulas $P(t_1, \ldots, t_n)$ are seen as equations $p(t_1, \ldots, t_n) = \top$. For this and other similar details, see [22].

[19] In this section, we use sequent notation for clauses.

[20] It is not clear whether the results explained in this section hold in case a Knuth-Bendix ordering is adopted.

[21] For term-rewriting terminology, see e.g. [2].

terms/literals in the respective clauses (see again [22] for details):

$$\frac{\Gamma' \Rightarrow \Delta', l = r \qquad \Gamma \Rightarrow s = t, \Delta}{\Gamma\mu, \Gamma'\mu \Rightarrow s\mu[r\mu]_p = t\mu, \Delta\mu, \Delta'\mu}$$

$$\frac{\Gamma' \Rightarrow \Delta', l = r \qquad \Gamma, s = t \Rightarrow \Delta}{\Gamma\mu, \Gamma'\mu, s\mu[r\mu]_p = t\mu \Rightarrow \Delta\mu, \Delta'\mu}$$

$$\frac{s = t, \Gamma \Rightarrow \Delta}{\Gamma\mu \Rightarrow \Delta\mu}$$

$$\frac{\Gamma \Rightarrow s = s', t = t', \Delta}{\Gamma\mu, s'\mu = t'\mu \Rightarrow s\mu = t'\mu, \Delta\mu}$$

where, in the first two rules, the premises do not share any variable and $s_{|p}$ is not a variable; the substitution μ is supposed to be a most general unifier of $\{s_{|p}, l\}$ in the first two rules and of $\{s, t\}$ in the last two rules.

Theorem 7.2 *With the above assumptions (I)-(IV), the set of sentences $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$ is inconsistent iff there is a pure $\mathcal{I}$-derivation of the empty clause from $T_1^{sk} \cup \Gamma_1 \cup T_2^{sk} \cup \Gamma_2$.*

Proof. See the Appendix.

Before concluding this section, we shall provide an example in which the assumptions of Theorem 7.2 are satisfied and an example in which such assumptions fail.

Example 14. Let T_1, T_2 be both the theory of Boolean algebras; we assume that the symbols of the bounded distributive lattice language (namely $\cap, \cup, 0, 1$) are shared but that the two complements n_1, n_2 are not. We want to prove that

$$T_1 \cup T_2 \models \forall x(n_1(x) = n_2(x)).$$

If we take T_0 to be the theory of bounded distributive lattices (i.e. of distributive lattices with 0 and 1), we see that T_1, T_2 are T_0-compatible.[22] Negation, skolemization and purification give for instance the two sets of literals $\Gamma_1 = \{a = n_1(c), a \neq b\}$ and $\Gamma_2 = \{b = n_2(c), a \neq b\}$. A pure $\mathcal{I}$-refutation exists: the prover SPASS [31] produces a pure $\mathcal{I}$-refutation

[22]The model completion of T_0 is the theory of atomless Boolean algebras (formulated without the complement in the language), because the theory of atomless Boolean algebras can eliminate quantifiers even without having the complement in the language: in fact, an atomic Boolean formula can be rewritten as the conjunction of equations of the kind $x'_1 \cup \cdots \cup x'_n \cup y_1 \cup \cdots \cup y_m = 1$ (for $n, m \geq 0$) and these equations can be written without complement as $x_1 \cap \cdots \cap x_n \leq y_1 \cup \cdots \cup y_m$.

consisting on 28 steps. However, the system is not programmed in order to avoid impure inferences, so that, during saturation, it impurely derives also (useless) 'mixed' clauses containing both n_1 and n_2. One of them, namely the atom $b \cap n_1(n_2(a)) = b$, is also selected as a given clause.[23]

Example 15. Let T_1 be the theory of Boolean algebras and let T_2' be the theory of pseudocomplemented distributive lattices; these are bounded distributive lattices endowed with a unary operator $(-)^*$ satisfying the condition

$$\forall x \forall y \, (x \cap y = 0 \leftrightarrow y \leq x^*)$$

$(-)^*$ expresses the properties of intuitionistic negation, but in the union theory $T_1 \cup T_2'$ such operator collapses into the classical complement. This is evident, because the axiom for $(-)^*$ implies that x^* is the supremum of the set of the elements y such that $x \cap y = 0$: as the Boolean complement of x enjoys the same property, clearly it must coincide with x^*. This means that $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$ is inconsistent, where Γ_1, Γ_2 are both empty and T_2 is $T_2' \cup \{(a^*)^* \neq a\}$ (here a is a new constant added to the signature of T_2'). A SPASS refutation takes 43 lines and it is highly impure. In fact a pure refutation cannot exist: the Σ_0-clauses (i.e. the bounded distributive lattice clauses) that can be exchanged by $T_1 \cup \Gamma_1$ and $T_2 \cup \Gamma_2$ are insufficient to detect inconsistency, because they are all logical consequences of both the consistent theories $T_1 \cup \Gamma_1 \cup \{0 \neq 1\}$ and $T_2 \cup \Gamma_2$. To show this, it is sufficient to observe that any non-degenerated bounded distributive lattice embeds both into a non-degenerated Boolean algebra and into a pseudocomplemented distributive lattice endowed with a non-complemented element a. This means that the universal sentences in the bounded distributive lattice language which are deducible from either $T_1 \cup \Gamma_1 \cup \{0 \neq 1\}$ or $T_2 \cup \Gamma_2$ are already deducible from the common subtheory of the non-degenerated bounded distributive lattices: thus, the only possible clause information exchange between $T_1 \cup \Gamma_1$ and $T_2 \cup \Gamma_2$ is limited the unit clause $0 \neq 1$. Notice that T_2 is not compatible with the Σ_0-theory T_0 of bounded distributive lattices.

8 Conclusions and related work

In this paper we have extended Nelson-Oppen combination procedure to the case of theories T_1, T_2 over non-disjoint signatures, in presence of compatibility conditions over a common universal subtheory T_0. The extension we proposed applies to examples of real interest giving, as shown in Section 6, combined decidability in case T_0 is effectively locally finite. Whenever T_0 is

[23] This execution is obtained by a default configuration of the prover. By contrast, a non-default RPO-configuration giving n_1, n_2 bigger precedence yields an impure refutational proof.

not locally finite, our method can be used in order to limit residue exchange (see Section 5) or in order to forbid impure inferences in saturation-based theorem proving, thus yielding restrictions on the search space during refutation derivations (see Section 7).

Quantifier elimination has been considered, since the early times of modern symbolic logic, a powerful technique for decision procedures. In actual approaches to combination problems (see e.g. [18]), specific quantifier elimination algorithms are often invoked as specialized reasoners to be integrated within a flexible general setting dealing with multiple theories. This happens, in particular, whenever there is the need of addressing numerical constraints problems [10]: examples of such specialized reasoners are for instance the Fourier-Motzkin quantifier elimination procedure (and its refinements [32]) for linear rational arithmetic or Cooper's [12], [23] quantifier elimination procedure for integer Presburger arithmetic. In contrast to this *local call* for quantifier elimination algorithms, we addressed in this paper quantifier elimination as a *global design* opportunity for integrated provers in the Nelson-Oppen style.

It should be noticed however that quantifier-elimination plays only an indirect role in our approach: in this sense, the existence of a model completion for a universal theory T_0 guarantees a certain behavior in combination problems by itself, independently on how quantifier elimination in the model completion is established (this can be established also by semantic non constructive arguments, as largely exemplified in the model-theoretic literature). In principle, the quantifier elimination complexity/decidability has nothing to do with the complexity/decidability of our combination methods, simply because quantifier elimination algorithms are not used in them. This is crucial, because most quantifier elimination algorithms are subject to heavy complexity lower bounds, which are often structural lower bounds for the decision of the elementary word problem in the corresponding theories (see the book [13]). In fact, there are instructive examples showing how complexity can grow from constraint satisfiability problems to decision problems in full first-order language: the most striking one is the case of acyclic lists, where complexity grows from linear [25] to non-elementary [13], [25].

One may wonder how severe is the crucial condition of T_0-compatibility used in the paper. T_0-compatibility involves two aspects, namely the existence of a model completion T_0^* for T_0 and the embeddability of models of T_i into models of $T_i \cup T_0^*$. As we have shown in the examples, the existence of a model completion seems to be frequent for theories commonly used in software verification. On the one hand, numeric constraint theories often enjoy this property, in the sense that they eliminate quantifiers (thus being model completions of the theories axiomatized by their respective universal consequences). On the other hand, acyclic lists might probably be the paradigm of situations arising in theories axiomatizing natural datatypes. Finally, notice that quantifier elimination strictly depends on the choice of

the language: every theory trivially has quantifier elimination in an extended language with infinitely many definitional axioms, hence the problem of obtaining quantifier elimination seems to be mostly a problem of choosing a sufficiently rich but still natural and manageable language.

The question concerning embeddability of models of T_i into models of $T_i \cup T_0^*$ looks more problematic, for instance because it can fail in significant situations. Further research is necessary on this point, however we stress that there is a relevant case in which the problem disappears. This is the case in which T_i is an extension of T_0^*: we have seen an example in Section 6 where T_i is rational linear arithmetic and T_0 is the theory of linear orders. Another example is the theory of acyclic lists L (which coincides with L^*): any extension of the theory of acyclic lists with extra structures matches our requirements and the advantages of our method (limited residue exchange, elimination of impure inferences, etc) apply to all combinations of theories obtained in this way.

An interesting point for future developments relies on the following research direction: how to relax local finiteness while still guaranteeing the convergence of the modular combination procedure? Here one may for instance take advantage from algebraic notions like 'noetherianity' for congruences. Notice however that local finiteness is already a much weaker requirement than other analogous notions known from the literature. For instance, in [6], a Σ_0-theory T_0 is said to be *finitary modulo a renaming* iff there is a finite set of Σ_0-terms S such that for every further Σ_0-term u there are $t \in S$ and a renaming σ such that $T_0 \models u = t\sigma$. This is a stronger condition than local finiteness: notice in fact that locally finite theories (like Boolean algebras), in which the number $k_{\underline{a}}$ of non-equivalent ground $\Sigma_0^{\underline{a}}$-terms grows more than polynomially in the cardinality of $\underline{a}$, cannot be finitary modulo a renaming.

There have been other efforts in the literature trying to extend Nelson-Oppen combination method to theories sharing function and predicate symbols (different from equality). We leave aside interesting recent results [5], [14] on combined word problems because (as pointed out in [4]) they cannot be appropriately seen as generalizations of Nelson-Oppen combination procedure; rather, we concentrate here on [28] which is directly related to the subject.

The starting point of any attempt to generalize Nelson-Oppen procedure to the non-disjoint case should preliminarily answer the following question: what is the specific feature of the stable infiniteness requirement that we want to generalize? In the present paper we answered the question by saying that infinite models are just *existentially closed*[24] models of the pure theory of equality and based our further investigations on this observation. On the contrary, in [28] the authors notice that infinite models are just *free*

[24]See [11] for this notion and its relationship to model completeness.

models of the pure theory of equality with infinitely many generators. This leads to a completely different research direction, whose results can only very roughly be reported here. Let T be a theory in the signature Σ and let Σ_0 be a finite subsignature of Σ; say that 'T is stably Σ_0-free (over a certain constraint language) iff every constraint (in the language) satisfiable in T is satisfiable in a model of T whose Σ_0-reduct is a free structure with infinitely many generators' (loc. cit. p.296). The main notion in [28] is that of N-O combinability of two theories T_1, T_2 over signatures Σ_1, Σ_2 (admitting Σ_0 as a common subsignature): N-O combinability guarantees, in essence, that 'the satisfiability in a theory $T_1 \cup T_2$ of the conjunction of two pure constraints $\phi_1 \wedge \phi_2$ can be reduced to the local satisfiability of ϕ_1 in T_1 and ϕ_2 in T_2 by adding to both formulas an appropriate Σ_0-restriction, a particular kind of first-order restriction on the free variables shared by ϕ_1 and ϕ_2' (loc. cit. p.295). One of the main results in [28] (Theorem 48, p.328) says that T_1, T_2 are N-O combinable provided they are both stably Σ_0-free and provided some additional conditions are satisfied.

Apart from technical details, we want to underline here that the notion of infinitely generated free and of existentially closed structure are quite different and their coincidence for the pure theory of equality should be considered a rather exceptional fact. To see how the two notions can diverge consider the case of torsion-free abelian groups: free abelian groups are torsion-free, but they are never divisible, whereas divisibility is just the necessary and sufficient condition for being existentially closed. Notice also that free structures may not exist for a theory, whereas every model of a universal theory T_0 embeds into a model which is existentially closed for T_0 (see [11]). These remarks should make evident that the generalization of the Nelson-Oppen procedure presented in [28] is quite different from that presented in this paper.

Before closing, we would like to remark that the idea (suggested in [26]) of using interpolation theorems in order to limit residue exchange in partial theory reasoning (whenever the background reasoner has to deal with combined theories) inspired some of the material presented in Section 5 above. The main problem addressed in [26] is general reasoners' interaction (more precisely, reasoners' interaction in a variable free tableaux framework) and the results explained there do not assume almost anything about the component theories. In this paper, on the contrary, we focus on contexts in which a very *specific* requirement (namely T_0-compatibility) is satisfied. Thanks to this requirement, we can, contrary to the methods of [26], restrict the shape of the residues to be exchanged and get also decidability in the locally finite case.

Acknowledgements The author wishes to thank Cesare Tinelli, Silvio Ranise and Harald Ganzinger for helpful discussions on earlier versions of this paper.

References

[1] Armando A., Ranise S., Rusinowitch M., *Uniform Derivation of Superposition Based Decision Procedures*, in Fribourg L. (ed.) "Proc. of the Annual Conf. on Computer Science Logic (CSL-01)", Paris, France, pp.513-527, (2001).

[2] Baader F., Nipkow T., *Term Rewriting and All That*, Cambridge University Press, (1998).

[3] Baader F., Ghilardi S., Tinelli C., *A New Combination Procedure for the Word Problem that generalizes Fusion Decidability in Modal Logic*, "Proc. of the Second Int. Joint Conf. on Automated Reasoning (IJCAR-04)", Springer LNAI, to appear, (2004). Extended version available at http://www.cs.uiowa.edu/ tinelli/html/papers.html as Technical Report No.03-03, Department of Computer Science, The University of Iowa.

[4] Baader F., Tinelli C., *A new approach for combining decision procedures for the word problem, and its connection to the Nelson-Oppen combination method*, in Mc Cune W. (ed.) "Proc. of the 14th Int. Conf. on Automated Deduction, (CADE-14)", Springer LNCS 1249, pp.19-33, (1997).

[5] Baader F., Tinelli C., *Deciding the Word Problem in the Union of Equational Theories*, Information and Computation, 178, 2, pp.346-390, (2002a).

[6] Baader F., Tinelli C., *Combining Decision Procedures for Positive Theories Sharing Constructors*, in Tison S. (ed.) "Proc. of the 13th Int. Conf. on Rewriting Techniques and Applications", (RTA-02)", Springer LNCS 2378, pp.352-366, (2002b).

[7] Bachmair L., Ganzinger H. *On Restrictions of Ordered Paramodulation with Simplification*, in Stickel M. (ed.) "Proceedings of the 10th Int. Conf. on Automated Deduction (CADE-10)", Springer LNCS 449, pp.427-441, (1990).

[8] Bachmair L., Ganzinger H. *Rewrite-based equational theorem proving with selection and simplification*, Journal of Logic and Computation, vol.4, n.3, pp.217-247, (1994).

[9] Bachmair L., Ganzinger H. *Equational Reasoning in Saturation-Based Theorem Proving*, in Bibel L., Schmitt P.H. (eds.) "Automated Deduction - A Basis for Applications", vol. I, pp.353-397, Kluwer, (1998).

[10] Bockmayr A., Weispfenning V., *Solving Numerical Constraints*, in Robinson A., Voronkov A., (eds.) "Handbook of Automated Reasoning", vol. I, Elsevier/MIT, pp.751-842, (2001).

[11] Chang C.C., Keisler H.J., *Model Theory*, IIIrd edition, North Holland, (1990).

[12] Cooper D. C., *Theorem Proving in Arithmetic without Multiplication*, in Meltzer B., Milchie D., (eds.) "Machine Intelligence 7", pp.91-99, (1972).

[13] Ferrante J., Rackoff C.W., *The Computational Complexity of Logical Theories*, Springer Lecture Notes in Mathematics 718, (1979).

[14] Fiorentini C., Ghilardi S., *Combining Word Problems through Rewriting in Categories with Products*, Theoretical Computer Science, 294, pp.103-149, (2003).

[15] Ghilardi S., *Reasoners' Cooperation and Quantifier Elimination*, Technical Report No.288-03, Dipartimento di Scienze dell'Informazione, Università degli Studi di Milano, (2003). Available at http://homes.dsi.unimi.it/∼ghilardi/

[16] Ghilardi S., Santocanale L., *Algebraic and Model Theoretic Techniques for Fusion Decidability in Modal Logic*, Vardi M., Voronkov A. "Logic for Programming, Artificial Intelligence and Reasoning (LPAR-03)", Springer LNAI 2850, pp.152-166, (2003).

[17] Ghilardi S., Zawadowski M., *Sheaves, Games and Model Completions*, Trends in Logic Series, Kluwer, (2002).

[18] Janicìc P., Bundy A., *A General Setting for Flexibly Combining and Augmenting Decision Procedures*, Journal of Automated Reasoning, 28, pp.257-305, (2002).

[19] MacIntyre A., *Model Completeness*, in Barwise J. (ed.), "Handbook of Mathematical Logic", North Holland, pp.139-180, (1977).

[20] Mal'cev A. I., *Axiomatizable Classes of Locally Free Algebras of Certain Types*, Sibirsk. Mat. Ž., 3, pp.729-743, (1962).

[21] Nelson G., Oppen D., *Simplification by Cooperating Decision Procedures*, ACM Transactions on Programming Languages and Systems, 1(2), pp.245-257, (1979).

[22] Nieuwhenhuis R., Rubio A., *Paramodulation-Based Theorem Proving*, in Robinson A., Voronkov A., (eds.) "Handbook of Automated Reasoning", vol. I, Elsevier/MIT, pp.371-533, (2001).

[23] Oppen D., *A* $2^{2^{2^{pn}}}$ *-Upper Bound on the Complexity of Presburger Arithmetic*, Journal of Computer and Systems Sciences, 16(3), pp.323-332, (1978).

[24] Oppen D., *Complexity, Convexity and Combination of Theories*, Theoretical Computer Science, 12, pp.291-302, (1980a).

[25] Oppen D., *Reasoning about Recursively Defined Data Structures*, Journal of the ACM, 27, 3, pp.403-411, (1980b).

[26] Tinelli C., *Cooperation of Background Reasoners in Theory Reasoning by Residue Sharing*, Journal of Automated Reasoning, 30(1), pp.1-31, (2003).

[27] Tinelli C., Harandi M., *A New Correctness Proof of the Nelson-Oppen Combination Procedure*, in Baader F., Schulz K. (eds.) "1st International Workshop on Frontiers of Combining Systems (FroCos-96)", Applied Logic Series, vol. 3, Kluwer Academic Publishers, pp.103-120, (1996).

[28] Tinelli C., Ringeissen C., *Unions of Non-Disjoint Theories and Combination of Satisfiability Procedures*, Theoretical Computer Science 290(1), pp.291-353, (2003).

[29] Tinelli C., Zarba C., *Combining Non-Stably Infinite Theories*, "Proc. of the Int. Workshop on First Order Theorem Proving (FTP-03)", Electronic Notes in Theoretical Computer Science, vol. 86, n.1, (2003).

[30] Wheeler W. H., *Model-Companions and Definability in Existentially Complete Structures*, Israel Journal of Mathematics, 25, pp.305-330, (1976).

[31] Weidenbach C., Afshordel B., Brahm U., Cohrs C., Engel T., Keen E., Theobalt C., Topic D., *System Description: Spass Version 1.0.0*, in Ganzinger H. (ed.) "Proc. of the 16-th Int. Conf. on Automated Deduction (CADE-16)", Springer LNAI 1632, pp.314-318, (1999). See also the SPASS Home Page http://spass.mpi-sb.mpg.de/

[32] Weispfenning V., *The Complexity of Linear Problems in Fields*, Journal of Symbolic Computation, 5 (1-2), pp.3-27, (1988).

[33] Wolter F., *Fusions of Modal Logics Revisited*, in Kracht M., De Rijke M., Wansing H., Zakharyaschev M. (eds.) "Advances in Modal Logic", CSLI, Stanford, (1998).

[34] Zarba, C. G., *A Tableau Calculus for Combining Non-Disjoint Theories*, in Egly U., Fermüller C. (eds.), "Automated Reasoning with Analytical Tableaux and Related Methods", Springer LNCS 2381, pp.315-329, (2002).

9 Appendix: proofs of the combination results

This Appendix is devoted to the proofs of Theorems 5.1, 5.2, 7.2 and of Proposition 4.4. To begin with, we need to introduce little further model-theoretic background.

Theorem 9.1 *(Robinson's Joint Consistency Theorem) Let H_1, H_2 be, respectively, consistent Θ_1, Θ_2-theories and let Θ_0 be the signature $\Theta_1 \cap \Theta_2$. Suppose that there is a complete Θ_0-theory H_0 such that $H_0 \subseteq H_1$ and $H_0 \subseteq H_2$; then $H_1 \cup H_2$ is a consistent $\Theta_1 \cup \Theta_2$-theory.*

The proof of this theorem can be easily deduced from Craig's Interpolation Theorem (alternatively, a direct proof using a double chain argument is possible, see [11], pp. 141-142).

Next, we need a lemma which is a little variation on Robinson Diagram Theorem. Let $\mathcal{A}$ be a Σ-structure and let $\underline{a}$ be a finite set of elements from $|\mathcal{A}|$;[25] we say that $\underline{a}$ *generates* $\mathcal{A}$ iff for every $b \in |\mathcal{A}|$ there is a Σ-term $t(\underline{x})$ such that $\mathcal{A} \models t(\underline{a}) = b$.[26] More generally, for a finite subset $\underline{a} \subseteq |\mathcal{A}|$, the Σ-*substructure generated by* $\underline{a}$ is the smallest substructure of $\mathcal{A}$ containing $\underline{a}$: it is easily seen that the support of such a substructure is formed by the elements $b \in |\mathcal{A}|$ such that we have $\mathcal{A} \models b = t(\underline{a})$, for some Σ-term $t(\underline{x})$.

[25] By abuse, we sometimes confuse a finite set with a tuple of distinct elements.

[26] Obviously, $t(\underline{a})$ is the $\Sigma^{\underline{a}}$-ground term resulting from the replacements $\underline{x} \mapsto \underline{a}$. Notice that, here and in similar situations, $\mathcal{A}$ is expanded to a $\Sigma^{\underline{a},b}$-structure by interpreting the names of the $\underline{a}, b$'s into themselves.

Lemma 9.2 *Suppose that the Σ-structures $\mathcal{A}$, $\mathcal{A}'$ are generated by the finite subsets $\underline{a} = (a_1, \ldots, a_n)$ and $\underline{a}' = (a'_1, \ldots, a'_n)$, respectively. Let $\underline{x}$ be the tuple of variables $(x_1, \ldots, x_n)$ and suppose also that for every Σ-atom $A(\underline{x})$,[27] we have $\mathcal{A} \models A(\underline{a})$ iff $\mathcal{A} \models A(\underline{a}')$. Then the bijection $a_i \mapsto a'_i$ extends to an isomorphism between $\mathcal{A}$ and $\mathcal{A}'$.*

Proof. Define the isomorphism μ as follows: as for every element $b \in |\mathcal{A}|$, we have $\mathcal{A} \models b = t(\underline{a})$ for some Σ-term $t(\underline{x})$, let us put $\mu(b) = b'$ iff $\mathcal{A}' \models b' = t(\underline{a}')$. This is well-defined and does not depend on the choice of t, because for every $u(\underline{x})$, we have $\mathcal{A} \models t(\underline{a}) = u(\underline{a})$ iff $\mathcal{A}' \models t(\underline{a}') = u(\underline{a}')$. By considering the term x_i, it is immediate to see that for every $i \in \{1, \ldots, n\}$, we have $\mu(a_i) = a'_i$.

Moreover, for every atom $B(y_1, \ldots, y_m)$ and for every m-tuple of elements $b_1, \ldots, b_m \in |\mathcal{A}|$, we have

$$\mathcal{A} \models B(b_1, \ldots, b_m) \quad \text{iff} \quad \mathcal{A}' \models B(\mu(b_1), \ldots, \mu(b_m)).$$

In fact, if $\mathcal{A} \models b_j = t_j(\underline{a})$ for $j \in \{1, \ldots, m\}$, the condition $\mathcal{A} \models B(b_1, \ldots, b_m)$ is equivalent to $\mathcal{A} \models B(t_1(\underline{a}), \ldots, t_m(\underline{a}))$, the condition $\mathcal{A}' \models B(\mu(b_1), \ldots, \mu(b_m))$ is equivalent to $\mathcal{A}' \models B(t_1(\underline{a}'), \ldots, t_m(\underline{a}'))$ and $B(t_1(\underline{x}), \ldots, t_m(\underline{x}))$ is a Σ-atom fitting the hypothesis of the lemma. Hence μ is an embedding; it is also surjective (that is, an isomorphism), because for every $b' \in |\mathcal{A}'|$ we have $\mathcal{A}' \models b' = t(\underline{a}')$ (for a suitable term $t(\underline{x})$) and consequently $b' = \mu(b)$, where b is such that $\mathcal{A} \models b = t(\underline{a})$.

Let us now work *under the Assumptions* (I)-(II)-(III) from Section 5: that is, we fix a Σ_1-theory T_1, a Σ_2-theory T_2 (both compatible with respect to a $\Sigma_0 = \Sigma_1 \cap \Sigma_2$-universal theory T_0); we also fix a finite set of free constants $\underline{a}$, a finite set Γ_1 of ground $\Sigma_1^{\underline{a}}$-literals and a finite set Γ_2 of ground $\Sigma_2^{\underline{a}}$-literals.

Say that a set Γ_0 of *positive* ground $\Sigma_0^{\underline{a}}$-clauses is *saturated* iff it is closed under the two rules

$$T_1 \cup \Gamma_1 \cup \Gamma_0 \models C \qquad \Rightarrow \qquad C \in \Gamma_0$$

$$T_2 \cup \Gamma_2 \cup \Gamma_0 \models C \qquad \Rightarrow \qquad C \in \Gamma_0.$$

Lemma 9.3 *Suppose that Γ_0 is saturated and does not contain the empty clause. Then there are $\Sigma_i^{\underline{a}}$-models $\mathcal{M}_i$ $(i = 1, 2)$ such that*

$$\mathcal{M}_1 \models T_1 \cup \Gamma_1 \cup \Gamma_0 \quad \textit{and} \quad \mathcal{M}_2 \models T_2 \cup \Gamma_2 \cup \Gamma_0;$$

moreover $\mathcal{M}_1$ and $\mathcal{M}_2$ share the same Σ_0-substructure generated by the elements (denoted by) $\underline{a}$.

[27] Recall that we use the notation $A(\underline{x})$ to mean that A contains at most the free variables $\underline{x}$.

Proof. A set Δ of of ground Σ_0^a-literals is said to be *exhaustive* iff it contains for every ground Σ_0^a-atom A, either A itself or its negation. The statement of the lemma is proved if we are able to find an exhaustive set Δ of ground Σ_0^a-literals which is consistent with both $T_1 \cup \Gamma_1 \cup \Gamma_0$ and $T_2 \cup \Gamma_2 \cup \Gamma_0$. In this case, in fact, given any two models $\mathcal{M}_1 \models T_1 \cup \Gamma_1 \cup \Gamma_0 \cup \Delta$ and $\mathcal{M}_2 \models T_2 \cup \Gamma_2 \cup \Gamma_0 \cup \Delta$, we have that their Σ_0-substructures generated by $\underline{a}$ are Σ_0-isomorphic by Lemma 9.2 (hence we may assume that they are just the same substructure, by renaming some elements in one of the supports, if needed).

We shall adapt the notion of productive clause used today in proofs of refutational completeness for various variants of resolution and paramodulation calculi (see [7], [8]). Consider any terminating strict total order on ground Σ_0^a-atoms and extend it to a terminating strict total order $>$ for positive ground Σ_0^a-clauses by taking its standard multiset extension. We shall define increasing sets Δ_C^+ (varying $C \in \Gamma_0$) of ground Σ_0^a-atoms as follows. Recall that, as the empty clause is not in Γ_0, all positive clauses in Γ_0 are of the kind $A \vee A_1 \vee \cdots \vee A_n$ $(n \geq 0)$.

The definition is by transfinite induction on $>$. Say that the clause $C \equiv A \vee A_1 \vee \cdots \vee A_n$ from Γ_0 is *productive* (and produces the atom A) iff i) $\{A\} > \{A_1, \ldots, A_n\}$ and ii) $A_1, \ldots, A_n \notin \Delta_{<C}^+$ (where $\Delta_{<C}^+$ is $\bigcup_{D<C} \Delta_D^+$). Now, if C is productive and produces the atom A, we let Δ_C^+ to be $\Delta_{<C}^+ \cup \{A\}$, otherwise (if Δ is not productive) Δ_C^+ is simply $\Delta_{<C}^+$.

Let Δ^+ be $\bigcup_{C \in \Gamma_0} \Delta_C^+$ and Δ be $\Delta^+ \cup \{\neg A \,|\, A$ is a ground Σ_0^a-atom not belonging to $\Delta^+\}$. By construction, $\Delta \models \Gamma_0$, so we simply need to show that $T_1 \cup \Gamma_1 \cup \Delta$ and $T_2 \cup \Gamma_2 \cup \Delta$ are consistent. We need a preliminary claim.

Claim: if the clause $A \vee A_1 \vee \cdots \vee A_n$ is productive and A is the maximum atom in it, then $A_1, \ldots, A_n \notin \Delta^+$: this is evident, as the A_i's could only be produced by clauses smaller than $A \vee A_1 \vee \cdots \vee A_n$.

Suppose now that $T_1 \cup \Gamma_1 \cup \Delta$ is not consistent (the case $i = 2$ is analogous). Then there are ground atoms $B_1, \ldots, B_m \notin \Delta^+$ and productive clauses

$$\begin{array}{rcl} C_1 & \equiv & A_1 \vee A_{11} \vee \cdots \vee A_{1k_1} \\ & & \cdots \\ C_n & \equiv & A_n \vee A_{n1} \vee \cdots \vee A_{nk_n} \end{array}$$

(with maximum atoms $A_1, \ldots, A_n$, respectively), such that

$$T_1 \cup \Gamma_1 \cup \{A_1, \ldots, A_n, \neg B_1, \ldots, \neg B_m\} \models \bot,$$

i.e. such that

$$T_1 \cup \Gamma_1 \cup \{A_1, \ldots, A_n\} \models B_1 \vee \cdots \vee B_m.$$

By trivial logical manipulations, it follows that

$$T_1 \cup \Gamma_1 \cup \{C_1, \ldots, C_n\} \models \bigvee_{i,j} A_{ij} \vee B_1 \vee \cdots \vee B_m.$$

As $C_1, \ldots, C_n$ are clauses in Γ_0 and as Γ_0 is saturated, the clause

$$D \equiv \bigvee_{i,j} A_{ij} \vee B_1 \vee \cdots \vee B_m$$

is also in Γ_0. By construction (whether D is productive or not)[28] some of the atoms $\{A_{11}, \ldots, A_{nk_n}, B_1, \ldots, B_m\}$ are in Δ^+. By the claim, $A_{11}, \ldots, A_{nk_n}$ cannot be there, so one of the B_j's is in Δ^+, contradiction.

Recall that if Σ is a signature, $\mathcal{M}$ is a Σ-structure and $X \subseteq |\mathcal{M}|$, we may tacitly consider $\mathcal{M}$ as a Σ^X-structure by interpreting the name of each $b \in X$ into b. In the next lemma, the assumption on T_0-compatibility is crucial:

Lemma 9.4 *Let $\mathcal{M}_1$ be a Σ_1-model of T_1 and let $\mathcal{M}_2$ be a Σ_2-model of T_2; suppose also that $\mathcal{M}_1$ and $\mathcal{M}_2$ share a common Σ_0-substructure $\mathcal{A}$ (hence, in particular, they can be regarded as $\Sigma_1^{|\mathcal{A}|}$- and as $\Sigma_2^{|\mathcal{A}|}$-structures, respectively). Then there are a $(\Sigma_1 \cup \Sigma_2)^{|\mathcal{A}|}$-model $\mathcal{M}$ of $T_1 \cup T_2$ and two $\Sigma_i^{|\mathcal{A}|}$-embeddings $\mu_i : \mathcal{M}_i \longrightarrow \mathcal{M}$ ($i = 1, 2$).*[29]

Proof. By using suitable embeddings (supplied by the definition T_0-compatibility), we can embed the $\mathcal{M}_i$ into models $\mathcal{M}_i'$ of $T_i \cup T_0^*$. By renaming some elements in the supports if needed, we can also freely suppose that $\mathcal{A}$ is still a common substructure of $\mathcal{M}_1'$ and $\mathcal{M}_2'$ and that the sets $|\mathcal{M}_1'| \setminus |\mathcal{A}|$ and $|\mathcal{M}_2'| \setminus |\mathcal{A}|$ are disjoint. Consider the elementary diagrams $\Delta^e(\mathcal{M}_1')$, $\Delta^e(\mathcal{M}_2')$ of $\mathcal{M}_1'$, $\mathcal{M}_2'$; we first show that $\Delta^e(\mathcal{M}_1') \cup \Delta^e(\mathcal{M}_2')$ is consistent as a $\Sigma_1 \cup \Sigma_2 \cup |\mathcal{M}_1'| \cup |\mathcal{M}_2'|$-theory.

First notice that $T_0^* \cup \Delta(\mathcal{A})$ is a complete $\Sigma_0^{|\mathcal{A}|}$-theory: we have $\mathcal{A} \models T_0$ (recall that T_0 is universal, that $\mathcal{M}_1' \models T_0$, that $\mathcal{A}$ is a substructure of $\mathcal{M}_1'$ and that truth of universal formulas is preserved in substructures), hence the definition of model completion applies. Also, the theories $\Delta^e(\mathcal{M}_1')$ and $\Delta^e(\mathcal{M}_2')$ are both extensions of $T_0^* \cup \Delta(\mathcal{A})$. Now we can simply invoke Robinson's Joint Consistency Theorem 9.1: in our case, the signature Θ_1 of the statement of Therem 9.1 is $\Sigma_1 \cup |\mathcal{M}_1'|$ and the signature Θ_2 is $\Sigma_2 \cup |\mathcal{M}_2'|$, so that the signature Θ_0 is $\Sigma_0 \cup |\mathcal{A}|$, because the sets $|\mathcal{M}_1'| \setminus |\mathcal{A}|$ and $|\mathcal{M}_2'| \setminus |\mathcal{A}|$ are disjoint.

[28] If the maximal atom has more than one occurrence in D, we can erase all but one of such occurrences, thus getting a positive ground clause D' which is in Γ_0 too, because Γ_0 is saturated and $D \models D'$.

[29] This implies, in particular, that $\mu_1(a) = \mu_2(a)$ holds for all $a \in |\mathcal{A}|$: in fact, $\mathcal{M}$ is a $(\Sigma_1 \cup \Sigma_2)^{|\mathcal{A}|}$-model and the μ_i are $\Sigma_i^{|\mathcal{A}|}$-embeddings, hence we have $\mu_1(a) = a^{\mathcal{M}} = \mu_2(a)$.

Having established that $\Delta^e(\mathcal{M}'_1) \cup \Delta^e(\mathcal{M}'_2)$ is consistent, any model of its fits our $\mathcal{M}$: in fact, notice that such an $\mathcal{M}$ is a model of T_1 and of T_2, because it is a model of the elementary diagrams of $\mathcal{M}'_1$ and $\mathcal{M}'_2$. Moreover, we have $\Sigma_i^{|\mathcal{A}|}$-embeddings $\mathcal{M}_i \longrightarrow \mathcal{M}'_i \longrightarrow \mathcal{M}$, as required.[30]

Prof of Theorem 5.1 Suppose that there is no positive residue chain ending up with the empty clause. We build a saturated set Γ_0 of positive ground $\Sigma_0^{\underline{a}}$-clauses in ω steps. Let Θ_0 be the empty set; if Θ_k has already been defined, let Θ_{k+1} be the set of positive ground $\Sigma_0^{\underline{a}}$-clauses C such that $T_i \cup \Gamma_i \cup \Theta_k \models C$ holds for $i = 1$ or $i = 2$. Clearly $\Theta_k \subseteq \Theta_{k+1}$; moreover, by the compactness theorem for first order logic, it is clear that $\Gamma_0 = \bigcup_k \Theta_k$ is saturated. Notice also that a clause C belongs to Θ_{k+1} ($k \geq 0$) iff there is a positive residue chain $C_1, \ldots, C_n, C$ such that $C_1, \ldots, C_n$ all belong to Θ_k (this is easily proved by induction on k and by compactness again).[31] Consequently, Γ_0 does not contain the empty clause and Lemma 9.3 applies. This means that there are a $\Sigma_1^{\underline{a}}$-model $\mathcal{M}_1 \models T_1 \cup \Gamma_1$ and a $\Sigma_2^{\underline{a}}$-model $\mathcal{M}_2 \models T_2 \cup \Gamma_2$, whose Σ_0-substructures generated by $\underline{a}$ are the same. We can now apply Lemma 9.4 to $\mathcal{M}_1, \mathcal{M}_2, \mathcal{A}$ (where $\mathcal{A}$ is this Σ_0-substructure generated by $\underline{a}$). By that Lemma, there are a $(\Sigma_1 \cup \Sigma_2)^{|\mathcal{A}|}$-model $\mathcal{M} \models T_1 \cup T_2$ and $\Sigma_i^{|\mathcal{A}|}$-embeddings $\mathcal{M}_i \longrightarrow \mathcal{M}$. As $\mathcal{M}_i \models \Gamma_i$, we have also $\mathcal{M} \models \Gamma_i$ ($i = 1, 2$) (recall that the Γ_i's are sets of ground $\Sigma_0^{\underline{a}}$-literals, so their truth is preserved by $\Sigma_i^{|\mathcal{A}|}$-embeddings). Thus $\mathcal{M} \models T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$, so the latter set is indeed consistent.

Proof of Theorem 5.2 We reduce this Theorem to the previous one. If $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$ is inconsistent, there is a positive residue chain $C_1, \ldots, C_n$ ending up with the empty clause. Say that C_k is an i-residue ($i = 1, 2$) iff $T_i \cup \Gamma_i \cup \{C_1, \ldots, C_{k-1}\} \models C_k$. Let ψ_k (for $k = 1, \ldots, n$) be the quantifier-free ground $\Sigma_0^{\underline{a}}$-formula $\neg C_1 \vee \cdots \vee \neg C_{k-1} \vee C_k$ and let ϕ be the conjunction of the ψ_k such that C_k is a 1-residue. Clearly, $T_1 \cup \Gamma_1 \models \phi$. Moreover, by induction, it is easy to see that $T_2 \cup \Gamma_2 \cup \{\phi\} \models C_j$ for all $j = 1, \ldots, n$: in fact, if C_j is a 2-residue, then $T_2 \cup \Gamma_2 \cup \{C_1, \ldots, C_{j-1}\} \models C_j$ (hence $T_2 \cup \Gamma_2 \cup \{\phi\} \models C_j$ by induction hypothesis) and if C_j is a 1-residue, then $\{\phi\} \models \neg C_1 \vee \cdots \vee \neg C_{j-1} \vee C_j$ (hence $T_2 \cup \Gamma_2 \cup \{\phi\} \models C_j$ again by induction hypothesis). As $T_2 \cup \Gamma_2 \cup \{\phi\} \models C_j$ holds for all j, we have in particular that $T_2 \cup \Gamma_2 \cup \{\phi\} \models \bot$ for $j = n$.

The same argument used for the proof of Lemma 9.4 gives also the

Proof of Proposition 4.4 Take a model $\mathcal{M}$ of $T_1 \cup T_2$ and embeds its Σ_i-reducts into models $\mathcal{M}_i$ of $T_i \cup T_0^*$ ($i = 1, 2$). We can freely suppose that the

[30] Notice that the latter embedding is elementary, whereas the former needs not be such.

[31] The induction step is as follows: if $C \in \Theta_{k+1}$, then there are $C_1, \ldots, C_n \in \Theta_k$ such that $T_i \cup \Gamma_i \cup \{C_1, \ldots, C_n\} \models C$ holds for $i = 1$ or $i = 2$. Now it is sufficient to append C to any juxtaposition of positive residue chains ending up in $C_1, \ldots, C_n$.

embeddings are inclusions and that we have $|\mathcal{M}| = |\mathcal{M}_1| \cap |\mathcal{M}_2|$ for supports. Now $T_0^* \cup \Delta(\mathcal{M})$ is a complete theory (here $\Delta(\mathcal{M})$ is the diagram of $\mathcal{M}$ as a Σ_0-structure), hence by Robinson Joint Consistency Theorem 9.1 there is a model $\mathcal{N}$ of $\Delta^e(\mathcal{M}_1) \cup \Delta^e(\mathcal{M}_2)$. It follows that $\mathcal{N}$ is a $\Sigma_1 \cup \Sigma_2 \cup |\mathcal{M}_1| \cup |\mathcal{M}_2|$-model of $T_1 \cup T_2 \cup T_0^*$ and that there are $\Sigma_i^{|\mathcal{M}|}$-embeddings $\mu_i : \mathcal{M}_i \longrightarrow \mathcal{N}$. In particular, for $b \in |\mathcal{M}|$, we have $\mu_1(b) = b^{\mathcal{N}} = \mu_2(b)$; let us call μ the common restriction of μ_1 and μ_2 to $|\mathcal{M}|$. We show that μ is a $\Sigma_1 \cup \Sigma_2$-embedding of $\mathcal{M}$ into $\mathcal{N}$. Observe in fact that for every n-ary Σ_i-function symbol f and for every n-tuple $\underline{b}$ of elements from the support of $\mathcal{M}$, we have[32]

$$\mu(f^{\mathcal{M}}(\underline{b})) = \mu_i(f^{\mathcal{M}_i}(\underline{b})) = f^{\mathcal{N}}(\mu_i(\underline{b})) = f^{\mathcal{N}}(\mu(\underline{b}));$$

analogously, for every n-ary Σ_i-predicate symbol P, we have

$$\mathcal{M} \models P(\underline{b}) \text{ iff } \mathcal{M}_i \models P(\underline{b}) \text{ iff } \mathcal{N} \models P(\mu_i(\underline{b})) \text{ iff } \mathcal{N} \models P(\mu(\underline{b})).$$

This proves that $\mu : \mathcal{M} \longrightarrow \mathcal{N}$ is a $\Sigma_1 \cup \Sigma_2$-embedding.

It remains to prove Theorem 7.2: here *we need also Assmption* (IV) from Section 7, that is we suppose that T_1, T_2 are axiomatized by $\forall\exists$-sentences and that T_1^{sk}, T_2^{sk} are the skolemizations of T_1, T_2, respectively. Recall also that we run Superposition Calculus $\mathcal{I}$ with respect to a lexicographic path ordering giving lower precedence to shared Σ_0^a-symbols.

Our main task is that of reaching the model-theoretic configuration of Lemma 9.4 within the model generation construction of [7], [8]. As this is a quite laborious construction, we cannot make the exposition completely self-contained. However, we shall closely follow notations and definitions from [22].

Proof of Theorem 7.2 We know that $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$ is consistent iff so is $T_1^{sk} \cup \Gamma_1 \cup T_2^{sk} \cup \Gamma_2$. Suppose that there is no pure $\mathcal{I}$-derivation of the empty clause from $T_1^{sk} \cup \Gamma_1 \cup T_2^{sk} \cup \Gamma_2$. The saturation process with respect to pure $\mathcal{I}$-inferences produces a set of clauses S that can be represented as the union of two sets S_1, S_2 of $\Sigma_1^{sk} \cup \{\underline{a}\}$- and of $\Sigma_2^{sk} \cup \{\underline{a}\}$-clauses, respectively. Such sets can be taken so that they contain *the same* Σ_0^a-clauses (otherwise said, Σ_0^a-clauses from S are put in both of them); moreover, each of them is saturated with respect to $\mathcal{I}$-inferences from $T_i^{sk} \cup \Gamma_i$. Let S_1^{gr}, S_2^{gr} be the sets of $\Sigma_1^{sk} \cup \{\underline{a}\}$- and of $\Sigma_2^{sk} \cup \{\underline{a}\}$-ground instances of clauses in S_1 and S_2. Following the model generation procedure, we can define, by transfinite induction on ground clauses ordering, two convergent ground rewrite sytems R^1, R^2 whose associated normal Herbrand models $\mathcal{H}_1, \mathcal{H}_2$ are models of S_1, S_2 (hence also of $T_1^{sk} \cup \Gamma_1$ and of $T_2^{sk} \cup \Gamma_2$), respectively. The definition of R^i $(i = 1, 2)$ is as follows. Say that a $\Sigma_i^{sk} \cup \{\underline{a}\}$-ground clause $C \in S_i^{gr}$ of the kind

$$\Gamma \Rightarrow l = r, \Delta$$

[32]Here, if $\underline{b} = (b_1, \ldots, b_n)$, we write e.g. $\mu(\underline{b})$ for the tuple $(\mu(b_1), \ldots, \mu(b_n))$.

is *productive* and that that $Gen^i(C) = \{l \to r\}$ iff the following conditions are satisfied:

- $R^i_C \not\models C$;
- $l > r$, $l > u$ (for all u occurring in Γ), $\{l, r\} > \{u, v\}$ (for every equation $u = v$ occurring in Δ);
- l is not reducible by R^i_C;
- $R^i_C \not\models r = t'$, for every equation of the kind $l = t'$ occurring in Δ

(here R^i_C is the rewrite system $\bigcup_{D \in S_i^{gr}, D<C} Gen^i(D)$). If C is not productive, let us put $Gen^i(C) = \emptyset$. Finally take R^i to be $\bigcup_{C \in S_i^{gr}} Gen^i(C)$.

From the construction, it is clear that $\mathcal{H}_1, \mathcal{H}_2$ share the same Σ_0-substructure generated by $\underline{a}$ (up to a Σ_0-isomorphism): in fact, since ground $\Sigma_0^{\underline{a}}$-clauses/terms are smaller than clauses/terms containing a proper Σ_1- or Σ_2-symbol, it is immediate to prove by transfinite induction that a $\Sigma_0^{\underline{a}}$-term can only be reduced by a rule produced by a $\Sigma_0^{\underline{a}}$-clause and that, if C is such a clause, then $Gen^1(C) = Gen^2(C)$. Thus, R^1 and R^2 reduce a ground $\Sigma_0^{\underline{a}}$-term to the same normal form and this implies that $\mathcal{H}_1, \mathcal{H}_2$ share the same Σ_0-substructure $\mathcal{A}$ generated by $\underline{a}$.

By Lemma 7.1, T_1^{sk} and T_2^{sk} are both T_0-compatible, hence we can apply Lemma 9.4 to them and to the models $\mathcal{H}_1, \mathcal{H}_2$, thus getting a model $\mathcal{M}$ of $T_1^{sk} \cup T_2^{sk}$ and two $\Sigma_i^{sk} \cup |\mathcal{A}|$-embeddings $\mathcal{H}_i \longrightarrow \mathcal{M}$. As ground literals are preserved by embeddings, $\mathcal{M}$ is also a model of $T_1^{sk} \cup \Gamma_1 \cup T_2^{sk} \cup \Gamma_2$, as desired.