Combination Results for Many Sorted Theories with Overlapping Signatures

Vijay Ganesh[1], Sergey Berezin[1], Cesare Tinelli[2], David L. Dill[1]
[1]Stanford University
[2]University of Iowa

21st April 2004

Abstract

We present a combination result for many-sorted first-order theories whose signatures may share common symbols (i.e. *overlapping* or *non-disjoint* signatures), extending the recent results by Ghilardi for the unsorted case. Furthermore, we give practical conditions under which the combination method becomes a semi-decision procedure, and additional sufficient conditions which turn it into a decision procedure.

Several theories which are practically useful in formal verification have overlapping signatures (e.g. linear arithmetic and bit-vectors). We demonstrate how their decision procedures can be combined using our results. In addition, we obtain a many-sorted version of the Nelson-Oppen method as a special case of our combination result.

Contents

1 Introduction

Decision procedures are becoming increasingly important in formal verification and related areas. As a result, many efficient decision procedures have been developed for various theories like linear arithmetic, theory of equality, uninterpreted functions etc. However, practical verification problems often yield formulas which span over several theories, and it is very desirable to have a decision procedure for the combination of these theories.

A naive approach is to build a monolithic decision procedure for the union of a chosen set of theories. A more systematic and modular approach is to combine existing decision procedures for the individual theories into a decision procedure for the combined theory. Modular combinations allow one to take existing off-the-shelf decision procedures (either as an implementation or as an algorithm) and add it to the combination framework as a component, without having to re-implement the rest of the decision procedure for the combined theory.

The most well-known combination approaches to date for *unsorted first-order logic* (FOL) are the Nelson-Oppen [NO79] and Shostak [Sho84, RS01] methods and their variations [Bar03] which enable one to combine decision procedures for quantifier-free first-order theories with disjoint signatures. Various extensions have been studied in [BT97, Ghi03].

Many tools have been built based on Nelson-Oppen and Shostak combination methods for FOL. Examples include ICS [FORS01], Simplify [DNS], Verifun [FJOS03], etc.

However, it is our belief that many verification problems are naturally expressed in *many-sorted first-order logic* (MSL) [Man96], since hardware and software systems are usually written in typed languages. Consequently, we believe that it is natural for the users to expect support for many-sorted input languages in verification tools. Moreover, individual decision procedures are usually written for theories with specific models in mind, for instance, finite strings over $\{0, 1\}$ and integers. The FOL combination of the individual theories over these models will have a model whose elements behave both as strings and integers, which may be unnatural and confusing both to the developers and the users. For instance, FOL formulas in the union theory may have perfectly valid but ill-sorted behaviors. This could confuse the users, thus lowering the utility of the combination. These considerations were the primary motivation for us in deciding upon MSL as the input language of CVC [SBD02] and its successor CVC Lite [BB03].

Since any MSL formula can be translated into an equivalent FOL formula using *relativization* (i.e. introduction of unary *sort predicates*), one possible implementation of such an input language is to translate MSL theories into FOL theories and use an existing FOL combination algorithm. However, Such a translation of MSL theories into FOL theories results in theories whose signatures share the unary predicates corresponding to the sorts of MSL, i.e. we have overlapping signatures (Note: we assume sharing of sort symbols between the signatures of the individual MSL theories. This is a necessary assumption, for otherwise the combination is uninteresting). The FOL Nelson-Oppen method is no longer applicable and one has to use combination results for FOL theories with overlapping signatures (such as [Ghi03]), which are much more involved and do not guarantee decidability. Therefore, the existing combination results for FOL become inadequate. This provides us with the motivation to consider the MSL Nelson-Oppen combination result. There are added benefits to considering MSL Nelson-Oppen Combinations. It usually is easier to prove that a sorted theory is stably infinite over a certain set of sorts, than it is to prove that its unsorted version is stably infinite as a whole. Also, one can now combine theories with sorts admitting only finite interpretations, as long as these sorts are not shared between the theories.

There are practical settings in which MSL theories may have overlapping signatures (not just sorts but constant, function and relation symbols), thus motivating us go beyond MSL Nelson-Oppen Combinations, and to consider combination results for MSL theories with overlapping signatures. For example, it is natural to implement the theory of linear arithmetic and the theory of bit-vectors as two separate theories within the combination framework (here the sorts for the theory of linear arithmetic are $\mathcal{R}$, $\mathcal{Z}$ and the sorts for the theory of fixed bitvectors are $\mathcal{Z}$ and bit-vectors of length n for each $n \in \mathcal{N}$). However, these two theories share integer constants and the '+' operator, which implies that the MSL Nelson-Oppen method is inapplicable. We have to establish a combination result for theories with overlapping signatures to cover this case. In fact, in this paper we first establish such a combination result, and then subsequently derive the MSL Nelson-Oppen combination result from it.

A cursory knowledge of MSL might prompt one to ask "don't the combination results for MSL follow automatically from the results for FOL?". Logicians have noticed long ago that MSL is a quite different logic compared to FOL. Although MSL can be translated into FOL, some properties of MSL do not directly follow from their FOL counterparts [Fef68, Fef74, Man96].

In particular, MSL combination results cannot be derived from the corresponding results for FOL. For instance, extending the FOL Nelson-Oppen method to MSL yields a method in which only well-sorted arrangements are considered. Thus, only a subset of all possible arrangements (relative to the FOL case) are taken into account, and consequently, the completeness for the MSL Nelson-Oppen combination method does not follow from the completeness of the FOL Nelson-Oppen method. Similarly, the combination results for MSL theories with overlapping signatures do not follow directly from the corresponding FOL results by Ghilardi [Ghi03]. In particular, the theorems regarding the combination method in Section 7 differ from their FOL counterparts in important ways.

There has been almost no work done in the area of combination results for MSL. The only work we are aware are the unpublished results by Tinelli & Zarba. In their work they consider only theories with disjoint signatures, but have some results for order sorted logic (MSL with subsorts). We have not considered order sorted logic in this work, but we believe that our methods are more general, amenable to extensions and consequently can be extended to theories in order sorted logic with overlapping signatures. Our work is a non-trivial extension of Ghilardi's combination results for FOL [Ghi03] to many-sorted logic.

We now provide an overview of the paper. In Section 2 we present a syntax and semantics for MSL. The logic in the rest of the paper is MSL unless we specify otherwise. In Section 3 we provide a bird's eye-view of how the combination results for MSL have been derived from first principles and give a high-level structure of the proofs of important results. In Section 4 we provide some basic model theoretic notions for MSL and prove many useful theorems that will be used in subsequent sections. In Section 5 we state Feferman's Interpolation Theorem [Fef74, Fef68] for MSL, which is one of key ingredients for establishing the MSL combination results just as Craig's Interpolation Lemma [CK98] is for the FOL case and derive the many-sorted version of Robinson's Joint Consistency Theorem. In Section 6 we define the all important notion of model-completion and the notion of T_0-compatible theories. In Subsection 6.1 we show that a theory which admits elimination of quantifiers is also submodel-complete. In Subsection 6.2 we use the idea of submodel-completeness to establish that if two theories T_1 and T_2 are T_0-compatible then so is their union $T_1 \cup T_2$, and furthermore, we show that if both T_1 and T_2 are individually consistent then the union $T_1 \cup T_2$ is consistent provided T_1 and T_2 have models $\mathcal{M}_1$ and $\mathcal{M}_2$ with a common submodel.

In Section 7 we finally state and prove the combination result for MSL theories with overlapping signatures. In Section 8 we derive the MSL Nelson-Oppen method from the results in the previous sections. In Section 9 we provide a list of decidability conditions under which the combination is rendered decidable. In Section 10 we present some instances where the results are directly applicable to CVC Lite. Finally, we conclude in Section 11.

2 Preliminaries

This section describes the syntax and semantics of a first-order many-sorted logic (MSL) and gives various basic model-theoretic definitions. For convenience and clarity of definitions, we use a notion of *decorated symbols*, that is, symbols which carry a sort declaration explicitly in them. While decorated symbols are cumbersome to write in practice, at the theoretical level they dramatically simplify or eliminate a number of problems that vex more standard definitions of sorted logics. With decorated symbols sort inference is trivial, terms have a unique sort, set operations on signatures are straightforward and *ad-hoc* overloading is a non-issue.

2.1 Syntax

We assume that there exist fixed, infinite and pairwise disjoint sets of sorts $\mathcal{S}$, function symbols $\mathcal{F}$, constant symbols $\mathcal{C}$, predicate symbols $\mathcal{P}$, and variables $\mathcal{V}$. A *decorated function symbol* $f_{\bar{s},\, s_r}$ is a tuple $\langle f,\, \bar{s},\, s_r \rangle \in \mathcal{F} \times \mathcal{S}^+ \times \mathcal{S}$, and intuitively, it means that f expects arguments of the sorts $\bar{s} = \langle s_1,\, \ldots,\, s_n \rangle$, and returns the result of the sort s_r. Similarly, a *decorated predicate symbol* $p_{\bar{s}}$ is a tuple $\langle p,\, \bar{s} \rangle \in \mathcal{P} \times \mathcal{S}^*$ (i.e. the predicate p takes n arguments of sorts $\bar{s} = \langle s_1,\, \ldots,\, s_n \rangle$, and is true or false when $n = 0$). Finally, *decorated constant symbols and variables* c_s and x_s are pairs $\langle c,\, s \rangle \in \mathcal{C} \times \mathcal{S}$ and $\langle x,\, s \rangle \in \mathcal{V} \times \mathcal{S}$, respectively. Here $\mathcal{S}^*$ (similarly, $\mathcal{S}^+$) denotes a set of tuples (non-empty tuples) of elements of $\mathcal{S}$.

For the rest of the paper, we assume that each symbol and variable is uniquely decorated. Since every symbol is decorated uniquely, we will often omit the decorations (subindices) when it is convenient, and will also omit the word "decorated" when referring to variables and function, constant and predicate symbols.

A *first-order many-sorted signature* is a tuple $\Sigma = (P,\, F,\, C,\, S)$, where $S \subseteq \mathcal{S}$ is a set of sort symbols, $P \subseteq \mathcal{P} \times S^*$ is a set of decorated predicate symbols, $F \subseteq \mathcal{F} \times S^+ \times S$ is a set of decorated function symbols, and $C \subseteq \mathcal{C} \times S$ is a set of decorated constant symbols. For two signatures $\Sigma_1 = (P_1,\, F_1,\, C_1,\, S_1)$ and $\Sigma_2 = (P_2,\, F_2,\, C_2,\, S_2)$, we define:

$$
\begin{aligned}
\Sigma_1 \cup \Sigma_2 &\;=\; (P_1 \cup P_2,\, F_1 \cup F_2,\, C_1 \cup C_2,\, S_1 \cup S_2) \\
\Sigma_1 \cap \Sigma_2 &\;=\; (P_1 \cap P_2,\, F_1 \cap F_2,\, C_1 \cap C_2,\, S_1 \cap S_2) \\
\Sigma_1 \subseteq \Sigma_2 &\;\text{ iff }\; P_1 \subseteq P_2,\, F_1 \subseteq F_2,\, C_1 \subseteq C_2,\, S_1 \subseteq S_2.
\end{aligned}
$$

Definition. For a signature $\Sigma = (P,\ F,\ C,\ S)$ we define Σ-terms, Σ-atoms, and Σ-formulas.

Σ-term t:

$$t \ ::= x \mid c \mid f(t_1, \ldots, t_n)$$

where $x \in \mathcal{V}$ is a variable and $t_1, \ldots, t_n$ are Σ-terms, $c \in C$, $f \in F$.

Σ-atom a:

$$a \ ::= p(t_1, \ldots, t_n) \mid t_1 \approx t_2 \mid \mathsf{false} \mid \mathsf{true}$$

where $t_1, .., t_n$ are Σ-terms, and $p \in P$. false is the universally false atom and true is the universally true atom. In our logic, the equality predicate symbol $\approx$ is a logical symbol, and is not a part of any signature.

Σ-formula φ:

$$\varphi \ ::= a \mid \neg\varphi_1 \mid \varphi_1 \wedge \varphi_2 \mid (\exists x_s)\, \varphi_1$$

where φ_1, φ_2 are Σ-formulas, a is a Σ-atom, x_s is a variable whose sort is s, and $\exists$ is the existential quantifier. We will also use logical connectives $\rightarrow, \leftarrow, \leftrightarrow, \vee$ and the universal quantifier $\forall$ as the usual shorthands built out of $\neg, \wedge$ and $\exists$.

Definition. We define the notions of the sort of Σ-terms and well-formedness of Σ-terms and Σ-formulas.

Sort of Σ-term: A term t is *well-formed* and *has a sort s* (denoted by $t : s$), if $t : s$ can be derived by the following rules:

$$\frac{}{x_s : s} \qquad \frac{}{c_s : s} \qquad \frac{t_1 : s_1 \ \cdots \ t_n : s_n}{f_{s_1 \ldots s_n, s}(t_1,\ \ldots,\ t_n) : s}$$

Well-formed Σ-formula: A formula φ is *well-formed* (denoted by $\varphi : \mathsf{wff}$ for *well-formed formula*) if $\varphi : \mathsf{wff}$ can be derived by the following rules:

$$\mathsf{true} \qquad \mathsf{false} \qquad \frac{t_1 : s_1 \ \cdots \ t_n : s_n}{p_{s_1 \ldots s_n}(t_1,\ \ldots,\ t_n) : \mathsf{wff}} \qquad \frac{t_1 : s \quad t_2 : s}{t_1 \approx t_2 : \mathsf{wff}} \qquad \frac{\varphi : \mathsf{wff}}{\neg\varphi : \mathsf{wff}}$$

$$\frac{\varphi_1 : \mathsf{wff} \quad \varphi_2 : \mathsf{wff}}{\varphi_1 \wedge \varphi_2 : \mathsf{wff}} \qquad \frac{\varphi : \mathsf{wff}}{(\exists x_s)\varphi : \mathsf{wff}}$$

The set of all such well-formed formulas is referred to as the *first-order many-sorted language L.*

We have the usual notion of free and bound variables. A Σ-literal is a Σ-atom or its negation. A Σ-clause is a disjunction of Σ-literals. A Σ-term or a Σ-literal is called *ground* if it does not have any variables. A Σ-formula is called *closed* if it does not contain any free variables. Closed Σ-formulas are also called Σ-*sentences.* A sentence is called *universal* (*existential*) if its prenex normal form has only universal (existential) quantifiers.

Definition. Theory: A Σ-*theory* T is a non-empty set of Σ-sentences.

Notation 1. A formula with free variables $x_1, \ldots, x_n$ is typically denoted as $\varphi(x_1, \ldots, x_n)$. Formulas with free variables are also called *open formulas.* Henceforth, we will drop the "Σ" from Σ-formula, Σ-atom etc., if it is clear from context.

2.2 Semantics

Following are definitions of a model, model of theory etc. These definitions have some differences from their first-order counterpart. Most of these definitions are many-sorted versions of the definitions given in the Chang & Keisler book on model theory [CK98].

Definition. Model $\mathcal{A}$, varable interpretation α, term and formula evaluation $\varphi^{\mathcal{A},\alpha}$, model of a theory $\mathcal{A} \models T$, consistent theory, complete theory.
Model $\mathcal{A}$: For a signature $\Sigma = (P, F, C, S)$, a Σ-*model* is a pair:

$$\mathcal{A} = \langle A, I \rangle,$$

where $A = \{A_s \mid s \in S\}$ is an S-indexed family of non-empty sets, called *sort-domains*, and I is a mapping of symbols from Σ to the corresponding constants, functions and predicates over the sort-domains. Namely, a function symbol $f_{s_1 \ldots s_n, s} \in F$ is interpreted as $I(f_{s_1 \ldots s_n, s}) = f^{\mathcal{A}}$, where $f^{\mathcal{A}}$ is a total function from $A_{s_1} \times \cdots \times A_{s_n}$ to A_s; a constant symbol $c_s \in C$ is interpreted as $I(c_s) = c^{\mathcal{A}} \in A_s$, and a predicate symbol $p_{s_1 \ldots s_n} \in P$ is interpreted as a relation $I(p_{s_1 \ldots s_n}) = p^{\mathcal{A}} \in A_{s_1} \times \cdots \times A_{s_n}$.

For any Σ-model, we also say that symbols of Σ are *interpreted* in $\mathcal{A}$, or $\mathcal{A}$ interprets the symbols.

Variable Interpretation: Let $V \subseteq \mathcal{V}$ be a set of variables. Let $V_s \subseteq V$ denote the set of all variables whose sort is s. A *variable interpretation* α of V over a Σ-model $\mathcal{A}$ is an S-indexed family of functions $\alpha = \{\alpha_s : V_s \rightarrow A_s \mid s \in S\}$, where S is the set of sorts of Σ. For any $x_s \in V_s$, $\alpha(x_s)$ denotes $\alpha_s(x_s)$. For any $a \in A_s$, $x_s \in V_s$ and $s \in S$, we denote by $\alpha[x_s \mapsto a]$ a new variable interpretation over $\mathcal{A}$ that maps x_s to a and is otherwise identical to α. We call the pair $(\mathcal{A}, \alpha)$ a Σ-*interpretation* over V. A Σ-interpretation $(\mathcal{A}, \alpha)$ over V induces a mapping $(t)^{\mathcal{A},\alpha}$ over Σ-terms into elements of $\mathcal{A}$ (also referred to as *evaluation of terms* in $\mathcal{A}$). This mapping can be further extended to evaluate Σ-formulas to $\{\mathsf{true}, \mathsf{false}\}$ in the model $\mathcal{A}$, as defined below.
Evaluation of terms and formulas (Satisfaction of formulas): For a Σ-model $\mathcal{A}$ and a variable interpretation α, we denote the *evauation of a term* in $\mathcal{A}$ as $t^{\mathcal{A},\alpha} \in A_s$, where t is of sort s, and denote the *evaluation of formulas* in the model $\mathcal{A}$ as $\varphi^{\mathcal{A},\alpha} \in \{\mathsf{true}, \mathsf{false}\}$. If $\varphi^{\mathcal{A},\alpha} = \mathsf{true}$, then we say that $\mathcal{A}$ *satisfies* φ *under the variable interpretation* α or $(\mathcal{A}, \alpha)$ *satisfies* φ (also denoted as $\mathcal{A}, \alpha \models \varphi$). We define the evaluation of terms and formulas inductively in Figure 1.

For ground terms t (similarily, closed formulas, or sentences φ) it is easy to see that the variable interpretation α is irrelevant in determining the value of $t^{\mathcal{A},\alpha}$ (similarily for $\varphi^{\mathcal{A},\alpha}$), and hence, we just write $t^{\mathcal{A}}(\varphi^{\mathcal{A}})$ to denote their evaluation in the model $\mathcal{A}$. We say that a Σ-sentence φ is *true in a model* $\mathcal{A}$ (alternatively $\mathcal{A}$ *satisfies* φ or $\mathcal{A}$ is a *model of* φ, written as $\mathcal{A} \models \varphi$) iff every Σ-interpretation $(\mathcal{A}, \alpha)$ satisfies φ.
Model of a theory: If all sentences of a theory T are true in a model $\mathcal{A}$ then we say $\mathcal{A}$ is a *model of* T and write $\mathcal{A} \models T$.
Entailment: We say that a theory *entails* a sentence φ (or φ is a *consequence* of T), written as $T \models \varphi$, iff every model of T is also a model of φ. In particular, $\psi \models \varphi$ denotes that the sentence ψ entails φ. Similarly, $T \models \Gamma$ for a set of sentences Γ denotes that $T \models \varphi$ for every $\varphi \in \Gamma$.
Consistent theory: A theory is *inconsistent* if for every Σ-sentence φ we have $T \models \varphi$, and in particular, $T \models \mathsf{false}$. A theory is *consistent* otherwise. It can be easily shown that a consistent theory always has a model.
Complete Theory: A theory is *complete* if for all sentences φ either $T \models \varphi$ or $T \models \neg\varphi$ but not both.
Axioms of a theory: We say that a subset Ax_T of sentences of T are the *axioms* of T if for every $\varphi \in T$ we have $Ax_T \models \varphi$.

$$
\begin{array}{rcl}
x^{\mathcal{A},\alpha} & = & \alpha(x) \\
c^{\mathcal{A},\alpha} & = & c^{\mathcal{A}} \\
f(t_1,\,\ldots,\,t_n)^{\mathcal{A},\alpha} & = & f^{\mathcal{A}}(t_1^{\mathcal{A},\alpha},\,\ldots,\,t_n^{\mathcal{A},\alpha}) \\
(t_1 = t_2)^{\mathcal{A},\alpha} & = & \left\{\begin{array}{ll} \mathsf{true} & \text{if } t_1^{\mathcal{A},\alpha} \approx t_2^{\mathcal{A},\alpha} \\ \mathsf{false} & \text{otherwise} \end{array}\right. \\
p(t_1,\,\ldots,\,t_n)^{\mathcal{A},\alpha} & = & \left\{\begin{array}{ll} \mathsf{true} & \text{if } t_1^{\mathcal{A},\alpha},\ldots,t_n^{\mathcal{A},\alpha} \in p^{\mathcal{A}} \\ \mathsf{false} & \text{otherwise} \end{array}\right. \\
\mathsf{false}^{\mathcal{A},\alpha} & = & \mathsf{false} \\
\mathsf{true}^{\mathcal{A},\alpha} & = & \mathsf{true} \\
(\neg\varphi)^{\mathcal{A},\alpha} & = & \left\{\begin{array}{ll} \mathsf{true} & \text{if } \varphi^{\mathcal{A},\alpha} \in \mathsf{false} \\ \mathsf{false} & \text{otherwise} \end{array}\right. \\
(\varphi_1 \wedge \varphi_2)^{\mathcal{A},\alpha} & = & \left\{\begin{array}{ll} \mathsf{true} & \text{if } \varphi_1^{\mathcal{A},\alpha} \in \mathsf{true} \text{ and } \varphi_2^{\mathcal{A},\alpha} \in \mathsf{true} \\ \mathsf{false} & \text{otherwise} \end{array}\right. \\
((\exists x_s)\,\varphi)^{\mathcal{A},\alpha} & = & \left\{\begin{array}{ll} \mathsf{true} & \text{if there exists } a \in A_s \\ & \text{such that } \varphi^{\mathcal{A},\alpha[x\mapsto a]} = \mathsf{true} \\ \mathsf{false} & \text{otherwise} \end{array}\right.
\end{array}
$$

Figure 1: Evaluation of a formula

Subtheory: We say that T' is a *subtheory* of T, written as $T' \subseteq T$, if every sentence T' is entailed by T.
Closure of a theory: Let T be a Σ-theory and let Ax_T be its axioms. The set of all Σ-sentences which are entailed by Ax_T is called the *closure of the theory* T or the *clsosure of the set* Ax_T.
Union Theory (Combination theory, Combination of theories): Given two theory T_1 and T_2, we say that the $\Sigma_1 \cup \Sigma_2$-theory obtained by the closure of the set $Ax_{T_1} \cup Ax_{T_2}$ (the usual set union of the axioms of T_1 and T_2), denoted by $T_1 \cup T_2$, is called the *union theory*, or the *union of* T_1 and T_2. The $\Sigma_1 \cup \Sigma_2$-theory $T_1 \cup T_2$ is also sometimes referred to as the *combination of theories* T_1 and T_2.
Equivalence: We say that two Σ-sentences φ and ψ are *equivalent* if the evaluation of φ and ψ is the same in all models.
Equisatisfiability: We say that two Σ-sentences φ and ψ are *equisatisfiable* if φ is satisfiable iff so is ψ.
Decidable Theories: We say that a Σ-theory T is (semi)-decidable if there is a (semi)-decision procedure which determines whether a Σ-sentence φ is entailed by T. We also say that the *problem* $T \models \varphi$ *is (semi)-decidable.* The set of all universal Σ-sentences entailed by a Σ-theory T is called the *universal fragment of* T. We say that a Σ-theory T is *universally (semi)-decidable* if there is a (semi)-decision procedure D which determines whether a universal Σ-sentence φ is entailed by T. We also call D a (semi)-decision procedure for the universal fragment of T.

Notation 2. Unless we explicitly specify, the terms *models, theories, formulas* etc. refer to Σ*-models,* Σ*-theories,* Σ*-formulas* etc.

3 Overview of the Results

The problem that we are trying to solve in this paper is the following:

Suppose there is a Σ_1-theory T_1 and a Σ_2-theory T_2 with overlapping signatures (i.e. $\Sigma_0 = \Sigma_1 \cap \Sigma_2$ is non-empty and there is common subtheory $T_1 \supseteq T_0 \subseteq T_2$), and an arbitrary $\Sigma_1 \cup \Sigma_2$-sentence φ. The problem we want to solve is "can we determine if $T_1 \cup T_2 \models \varphi$, provided that for any Σ_i-sentence φ_i we can determine whether $T_i \models \varphi_i$ where $i \in \{1, 2\}$". We call this the entailment problem for the theory $T_1 \cup T_2$.

The entailment problem for $T_1 \cup T_2$ can be further divided into two subproblems:

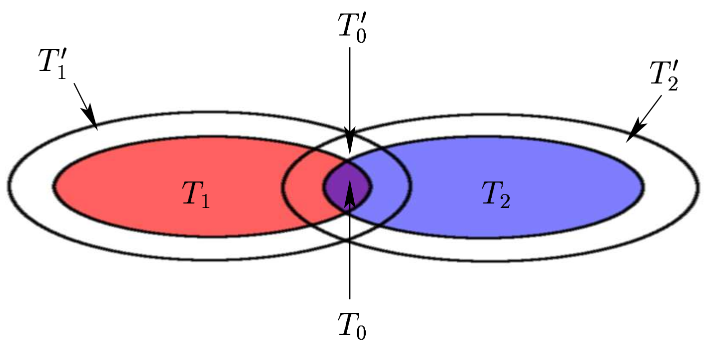

Figure 2: Use of Robinson's Consistency Theorem in Ghilardi's approach.

1. If T_1 and T_2 are both individually consistent, then is their union $T_1 \cup T_2$ consistent as well? Clearly if either of them is inconsistent, then so is the union.

2. Is there a method to determine whether $T_1 \cup T_2 \models \varphi$, provided we have methods to solve the entailment problem for each T_i.

The naivest approach to solving the first subproblem is to check if T_1 and T_2 have models $\mathcal{M}_1$ and $\mathcal{M}_2$ respectively, such that the Σ_0-reduct of $\mathcal{M}_1$ is elementarily equivalent to the Σ_0-reduct of $\mathcal{M}_2$. Clearly, such an approach is not very effective and infeasible in general. On the other hand, if certain restrictions are placed on the common subtheory T_0 such that the Σ_0-reducts of the models of T_1 and T_2 (and hence the models of T_0) are elementarily equivalent, then the consistency of the union is easily established.

Abraham Robinson proposed one such restriction in his now famous theorem called the *Robinson's Joint Consistency Theorem* [Rob56, CK98], regarding the consistency of the union of two consistent first-order (unsorted) theories. In order for the union theory $T_1 \cup T_2$ to be consistent for two individually consistent theories T_1 and T_2, Robinson's Theorem requires the existence of a theory $T_1 \supseteq T_0 \subseteq T_2$, such that T_0 is complete. Theorem 10 (in Section 5) is our many-sorted version of Robinson's Theorem, which relies on *Feferman's Interpolation Lemma* [Fef68, Fef74] for many-sorted logics (Lemma 9 in Section 5 as well). However, the completeness condition in Theorem 10 is the same as in the FOL case and is too strong for most theories of interest.

Following Ghilardi's work [Ghi03], we introduce a different and more practical set of conditions on the consistent many-sorted theories T_1 and T_2 so as to render the union $T_1 \cup T_2$ consistent. Namely, both T_1 and T_2 must be T_0-compatible for some theory T_0 (see Definition 13 in Section 6), where $T_1 \supseteq T_0 \subseteq T_2$, and there must exist two models $\mathcal{M}_1 \models T_1$ and $\mathcal{M}_2 \models T_2$ that share a common substructure for the shared signature.

The idea here is that, if T_1 and T_2 are T_0-compatible theories then one can extend them suitably to T_1' and T_2' respectively, such that there is a complete common subtheory $T_0' \supseteq T_0$ (Theorem 15 in Section 6, illustrated in Figure 2). Hence, by our many-sorted version of Robinson's Theorem their union $T_1' \cup T_2'$ is consistent, which in turn implies the consistency of $T_1 \cup T_2$.

Furthermore, we provide a mathematical construction (based on Ghilardi's approach [Ghi03]) for reducing the problem $T_1 \cup T_2 \models \varphi$ to the entailment problems $T_1 \models \varphi_1$ and $T_2 \models \varphi_2$ for the individual theories T_1 and T_2 (Theorem 19 in Section 7), where φ is a universal $\Sigma_1 \cup \Sigma_2$-sentence, φ_i is a universal Σ_i-sentence for $i \in \{1, 2\}$. The completeness of this construction is guaranteed by the Finite Residue Chain Theorem (Theorem 20 in Section 7).

Following the approach pioneered by Nelson and Oppen [NO79], this method can be described in two high-level steps:

1. *Purify* $\neg\varphi$ into equisatisfiable sets of pure ground formulas Γ_1 and Γ_2 over the signatures $\Sigma_1^{\bar{a}}$ and $\Sigma_2^{\bar{a}}$ respectively, where $\bar{a}$ is a finite set of fresh constant symbols. The idea is that once the formulas are purified then we can easily check if $T_i \cup \Gamma_i$ are individually consistent, $i \in \{1, 2\}$. Note that not all formulas can be purified, but universal sentences can always be purified.

2. Check if $T_i \cup \Gamma_i$ are individually consistent by exchanging information between the decision procedures for the respective theories. If false is exchanged at any point, then conclude that $T_1 \cup T_2 \models \varphi$ (by the Finite Residue Chain Theorem).

An immediate application of Theorems 15 and 20 (Sections 6 and 7) is the extension of the Nelson-Oppen combination result to many-sorted logic, when the shared signature is empty (i.e. the only shared predicate symbol is equality (since it is infact a logical symbol), but there may be shared sorts). The common subtheory in this case is the *theory of pure equality* $T_=$ over the shared sorts.

In the unsorted case, the Nelson-Oppen method requires that two quantifier-free theories T_1 and T_2 be *stably-infinite* [NO79] for the union theory $T_1 \cup T_2$ to be consistent, provided that T_1 and T_2 are individually consistent.

In our many-sorted extension of the Nelson-Oppen method (Theorem **29** in Subsection 8.3), we require **(a)** stably-infiniteness of T_1 and T_2 over the shared sorts (see Definition 22 in section 8) and **(b)** the existence of models $\mathcal{M}_1 \models T_1$ and $\mathcal{M}_2 \models T_2$ such that the domains for the shared sorts are the same in both models.

To establish the consistency of the union theory $T_1 \cup T_2$ in the many-sorted Nelson-Oppen method we use Theorem 15. This theorem requires that the individual theories T_1 and T_2 satisfy two conditions, namely the $T_=$-compatibility condition and the common submodel property. To satisfy the $T_=$-compatibility, we show that staby-infinteness (above condition (a)) implies $T_=$-compatibility (Lemma 27 in Subsection 8.2). Condition (b) provides the common submodel property. Hence, the consistency of $T_1 \cup T_2$ follows from the above conditions (a) and (b) by Theorem 15.

The proof of Lemma 27 relies on the fact that the theory of infinite sorts T_S is the model completion of $T_=$ (see Definition 22 in Section 8 and Definition 13 in Section 6), which in turn follows from the fact that T_S admits elimination of quantifiers (Theorem 24 in Subsection 8.1).

A non-deterministic decision procedure for the consistency of $T_1 \cup T_2 \cup \{\phi\}$, for any universal $\Sigma_1 \cup \Sigma_2$-sentence ϕ, consists of the following steps: purification of terms in ϕ by introducing a finite set of fresh shared constants $\bar{c} = \{c_1, \ldots, c_n\}$, translating ϕ into an equisatisfiable set of pure ground formulas $\Gamma_1 \cup \Gamma_2$ (that is, Γ_1 and Γ_2 are over the signatures $\Sigma_1^{\bar{c}}$ and $\Sigma_2^{\bar{c}}$, respectively), picking an *arrangement* A (see Definition 22) over these constants, and checking that $T_1 \cup \Gamma_1 \cup \{A\}$ and $T_2 \cup \Gamma_2 \cup \{A\}$ are individually consistent. Since the number of different arrangements A is finite, the termination of this procedure is obvious.

Notice, that we only consider the well-sorted arrangements, and therefore, the completeness of this algorithm does not immediately follow from the completeness of the unsorted Nelson-Oppen procedure, since the latter requires all the arrangements to be considered, even those that would be ill-sorted in our case.

We show completeness of our many-sorted version of the Nelson-Oppen method in Theorem 29 in Section 8.3. Namely, we show that if $(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2)$ is consistent, then there exists an arrangement A such that $T_1 \cup \Gamma_1 \cup \{A\}$ and $T_2 \cup \Gamma_2 \cup \{A\}$ are individually consistent. This fact follows from the existence of the common substructure guaranteed by Theorems 19 and 20, and this substructure effectively provides the arrangement A.

In Section 9 we list several decidability conditions and in Section 10 we provide some concrete applications of our results.

4 Some Basic Notions from Model Theory

Let $\Sigma = (P, F, C, S)$ and $\Sigma' = (P', F', C', S')$ be two first-order many-sorted signatures such that $\Sigma \subseteq \Sigma'$. We describe important notions such as isomorphism between models, elementarily equivalent models etc. below.

Definition. Domain Mapping, Isomorphism, Elementarily equivalent, Submodel, Elementary submodel.

Domain mapping: Given two sets of S-indexed families of sort-domains $A = \{A_s \mid s \in S\}$ and $B = \{B_s \mid s \in S\}$, an S-indexed family of functions $H = \{h_s : A_s \to B_s \mid s \in S\}$is called a *domain mapping.* We write $H(a)$ for $a \in A_s$ to denote $h_s(a)$. Similarly, for a variable interpretation α over $\mathcal{A}$, $\beta = H \circ \alpha)$ is a variable interpretation over $\mathcal{B}$ such that $\beta(x) = H \circ \alpha(x)$. Notice that we allow for overlapping sort-domains, although we do not limit the universe of models to those with specific relationships over the sort-domains (that would be order-sorted logic or MSL with subsorts). Even in the presence of overlapping sort-domains the following definitions should work fine provided the functions in the domain mappings behave identically over the elements in the intersection of the sort-domains.

Isomorphism: A Σ-*isomorphism* between Σ-model $\mathcal{A}$ and and Σ'-model $\mathcal{B}$ is a domain mapping $H = \{h_s : A_s \to B_s \mid s \in S\}$ where each h_s is a bijection from the sort-domain A_s into the sort-domain B_s, and the following conditions are satisfied:

1. For every predicate symbol $p_{s_1 \ldots s_n} \in P$ we have

   $$p^{\mathcal{A}}(a_1, \ldots, a_n) \text{ iff } p^{\mathcal{B}}(H(a_1), \ldots, H(a_n))$$

   for all tuples $\langle a_1, \ldots, a_n \rangle \in A_{s_1} \times \cdots \times A_{s_n}$.

2. For every function symbol $f_{s_1 \ldots s_n, s} \in F$ we have

   $$H(f^{\mathcal{A}}(a_1, \ldots, a_n)) = f^{\mathcal{B}}(H(a_1), \ldots, H(a_n))$$

   for all tuples $\langle a_1, \ldots, a_n \rangle \in A_{s_1} \times \cdots \times A_{s_n}$.

3. For each constant $c^{\mathcal{A}} \in A_s$, the corresponding constant $c^{\mathcal{B}} \in B_s$ is such that $H(c^{\mathcal{A}}) = c^{\mathcal{B}}$.

Isomorphic models: A Σ-model $\mathcal{A}$ and a Σ'-model $\mathcal{B}$ are Σ-*isomorphic* (written as $\mathcal{A} \cong \mathcal{B}$) if there exists a Σ-isomorphism between $\mathcal{A}$ and $\mathcal{B}$.
Equivalence of models: Σ-model $\mathcal{A}$ and Σ'-model $\mathcal{A}'$ are said to be Σ-*equivalent* (written as $\mathcal{A} \sim \mathcal{A}'$) whenever for any quantifier-free Σ-sentence σ (i.e. quantifier-free ground Σ-formulas) we have

$$\mathcal{A} \models \sigma \text{ iff } \mathcal{A}' \models \sigma$$

Elementary equivalence of models: Σ-model $\mathcal{A}$ and Σ'-model $\mathcal{A}'$ are said to be *elementarily* Σ-*equivalent* (written as $\mathcal{A} \equiv \mathcal{A}'$) whenever for any Σ-sentence σ we have

$$\mathcal{A} \models \sigma \text{ iff } \mathcal{A}' \models \sigma$$

Submodel: A Σ-model $\mathcal{A}$ is called a Σ-*submodel* of a Σ'-model $\mathcal{B}$, or $\mathcal{B}$ a Σ'-*extension* of $\mathcal{A}$ (written as $\mathcal{A} \subseteq \mathcal{B}$), iff $A_s \subseteq B_s$ for each sort $s \in S$, and the following conditions hold:

- For each predicate symbol $p_{s_1 \ldots s_n} \in P$, we have $p^{\mathcal{A}} = p^{\mathcal{B}} \cap (A_{s_1} \times \cdots \times A_{s_n})$.
- For each function symbol $f_{s_1 \ldots s_n, s} \in F$ and every tuple $\langle a_1, \ldots, a_n \rangle \in A_{s_1} \times \cdots \times A_{s_n}$ we have $f^{\mathcal{A}}(a_1, \ldots, a_n) = f^{\mathcal{B}}(a_1, \ldots, a_n)$.
- For each constant symbol $c \in C$, we have $c^{\mathcal{A}} = c^{\mathcal{B}}$.

Elementary Submodel: We say that a Σ-model $\mathcal{A}$ is an *elementary* Σ-*submodel* of Σ'-model $\mathcal{B}$ (written as $\mathcal{A} \preceq \mathcal{B}$), if $\mathcal{A} \subseteq \mathcal{B}$ and for all Σ-formulas $\varphi(x_1, \ldots, x_n)$ and all variable interpretations α over $\mathcal{A}$, we have $\varphi^{\mathcal{A},\alpha} \iff \varphi^{\mathcal{B},\alpha}$. We also say that $\mathcal{B}$ is an *elementary* Σ'-*extension* of $\mathcal{A}$.
Embedding: An (elementary) Σ-*embedding* $H : \mathcal{A} \to \mathcal{B}$ from Σ-model $\mathcal{A}$ into Σ'-model $\mathcal{B}$ is an isomorphism between $\mathcal{A}$ and $\mathcal{C}$, where $\mathcal{C}$ is some (elementary) Σ-submodel of $\mathcal{B}$.

We say that $\mathcal{A}$ is (elementarily) Σ-*embedded* into $\mathcal{B}$ (written as $\mathcal{A} \to \mathcal{B}$) if there exists an (elementary) Σ-embedding from $\mathcal{A}$ to $\mathcal{B}$. In other words, there exists a Σ-model $\mathcal{C}$ such that $\mathcal{A} \cong \mathcal{C} \subseteq \mathcal{B}$ (or $\mathcal{A} \cong \mathcal{C} \preceq \mathcal{B}$ for the elementary case); that is, there is a Σ-model $\mathcal{C}$ isomorphic to $\mathcal{A}$, such that $\mathcal{C}$ is an (elementary) submodel of $\mathcal{B}$.

Proposition 3. *A domain mapping $H = \{h_s : A_s \to B_s \,|\, s \in S\}$ is an (elementary) Σ-embedding from a Σ-model $\mathcal{A}$ into a Σ'-model $\mathcal{B}$ iff for any (Σ-formula) quantifier free Σ-formula $\varphi(x_1, \ldots, x_n)$ and every variable interpretation α over $\mathcal{A}$, we have*

$$\mathcal{A}, \alpha \models \varphi \textit{ iff } \mathcal{B}, H \circ \alpha \models \varphi.$$

Proof. An easy induction over the (Σ-formulas) quantifier free Σ-formulas, similar to the proof of Theorem 2.2.16 in [CK98] for FOL. □

5 Robinson's Joint Consistency Theorem

Many of the well known theorems from model theory for FOL like the Compactness Theorem and Substitution Lemma also hold for many-sorted first-order logic [Man96]. We shall use these theorems in the subsequent sections without proof. In this section, we state the Feferman's Interpolation Lemma and prove the many-sorted analog of the Robinson's Joint Consistency Theorem (also known as Robinson's Consistency Theorem). These theorems form the basis for our combination result.

Definition. Expansion (reduction) of a model: Let $\Sigma = (P, F, C, S)$ and $\Sigma' = (P', F', C', S')$ be two first-order many-sorted signatures such that $\Sigma \subseteq \Sigma'$. Let $\mathcal{M}'$ be a Σ'-model. The Σ-*reduct* of $\mathcal{M}'$ (written as $\mathcal{M}' \mid_\Sigma$) is the Σ-model whose sort-domains are the same as the sort-domains of $\mathcal{M}'$ over the sorts in S, for any $p \in P$ we have $p^{\mathcal{M}'|_\Sigma} = p^{\mathcal{M}'}$, for any $f \in F$ we have $f^{\mathcal{M}'|_\Sigma} = f^{\mathcal{M}'}$, for any $c \in C$ we have $c^{\mathcal{M}'|_\Sigma} = c^{\mathcal{M}'}$. We also call the Σ'-model $\mathcal{M}'$ a Σ'-*expansion* of the Σ-model $\mathcal{M}' \mid_\Sigma$.

Notation 4. We denote by $\varphi(a_1, \ldots, a_n)$ the Σ-formula obtained by substituting a_1 for x_1,...,a_n for x_n in the open Σ-formula $\varphi(x_1, \ldots, x_n)$, where $a_1, \ldots, a_n$ are decorated constant symbols of the appropriate sort in Σ. Note that Substitution Lemma holds for MSL [Man96].

Definition. $\mathcal{A}_\mathcal{X}$, the canonical $\Sigma_\mathcal{X}$-expansion of $\mathcal{A}$.

Signature $\Sigma_\mathcal{X}$: For a signature $\Sigma = (P, F, C, S)$, let $\mathcal{A} = \langle A, I \rangle$ be a Σ-model. Let $\mathcal{X} = \{X_s \subseteq A_s \,|\, s \in S\}$ be an S-indexed family of non-empty subsets of sort-domains of $\mathcal{A}$. We define $\Sigma_\mathcal{X} = (P, F, C_\mathcal{X}, S)$ to be the signature such that $C_\mathcal{X} = C \cup \{a_s \in X_s \,|\, X_s \in \mathcal{X}\}$ (that is, we expand the set

of constants C with the (appropriately decorated) elements of all $X_s \in \mathcal{X}$). For reasons of convenience, we may denote, $\Sigma_{\mathcal{X}}$, *the expansion of* Σ, *as* $(\Sigma)_{\mathcal{X}}$ (this notation is convenient when the signatures are already subscripted).
Canonical $\Sigma_{\mathcal{X}}$-expansion of $\mathcal{A}$: The $\Sigma_{\mathcal{X}}$-model $\mathcal{A}_{\mathcal{X}} = \langle A, I_{\mathcal{X}} \rangle$ is the *canonical* $\Sigma_{\mathcal{X}}$*-expansion of* $\mathcal{A}$ such that $I_{\mathcal{X}}(w) = I(w)$ for all symbols $w \in C \cup F \cup P$, and $I_{\mathcal{X}}(a) = a$ for all $a \in X_s$, $X_s \in \mathcal{X}$. We often have the case that $\mathcal{X} = A$, in which case the expanded signature is denoted by $\Sigma_{\mathcal{A}}$, and the corresponding canonical $\Sigma_{\mathcal{A}}$-expansion is denoted by $\mathcal{A}_{\mathcal{A}}$. Note that there may be other $\Sigma_{\mathcal{X}}$-models which extend $\mathcal{A}$ by interpreting the symbols from $\mathcal{X}$ as something other than themselves. Such models are termed *non-canonical* $\Sigma_{\mathcal{X}}$*-expansions* of $\mathcal{A}$.

Definition. Diagram $\Delta(\mathcal{A})$, elementary diagram $E\Delta(\mathcal{A})$, theory of a model $Th(\mathcal{A})$.

(Elementary) Diagram: The *diagram* $\Delta(\mathcal{A})$ of a Σ-model $\mathcal{A}$ is the set of all quantifier-free ground $\Sigma_{\mathcal{A}}$-formulas true in $\mathcal{A}_{\mathcal{A}}$. The *elementary diagram* $E\Delta(\mathcal{A})$ of a Σ-model $\mathcal{A}$ is the set of all $\Sigma_{\mathcal{A}}$-sentences true in $\mathcal{A}_{\mathcal{A}}$.
Theory of a model: The theory $Th(\mathcal{A})$ is the set of all Σ-sentences true in $\mathcal{A}$.

Theorem 5. *Let $\mathcal{A}$ and $\mathcal{B}$ be Σ-models, where $\Sigma = (P, F, C, S)$. For each Σ-formula $\varphi(x_1, \ldots, x_n)$ and every variable interpretation α over $\mathcal{A}$ the following statements hold:*

1. $\mathcal{A}, \alpha \models \varphi \quad \text{iff} \quad \mathcal{A}_{\mathcal{A}} \models \varphi(\alpha(x_1), , \ldots, \alpha(x_n))$

2. Let $M = \{m_s : A_s \to B_s \mid s \in S\}$ be a domain mapping from the sort-domains of $\mathcal{A}$ to the corresponding sort-domains of $\mathcal{B}$. Then

$$\mathcal{B}, M \circ \alpha \models \varphi \quad \text{iff} \quad \mathcal{B}' \models \varphi(\alpha(x_1), \ldots, \alpha(x_n)),$$

where $\mathcal{B}'$ is a $\Sigma_{\mathcal{A}}$-expansion of $\mathcal{B}$ such that the constant symbols $a \in C_{\mathcal{A}}$ in $\Sigma_{\mathcal{A}}$ corresponding to the elements of $\mathcal{A}$ are interpreted as $M(a)$ in $\mathcal{B}'$.

Proof. The proof is straightforward by a structural induction over φ. □

5.1 Robinson's Diagram Lemmas

Theorem 6. *(Robinson's Diagram Lemma). Let* $\mathcal{A} = \langle A, I_{\mathcal{A}} \rangle$ *and* $\mathcal{B} = \langle B, I_{\mathcal{B}} \rangle$ *be* Σ*-models. Then the following statements hold:*

1. $\mathcal{A}$ is embedded in $\mathcal{B}$ iff the Σ-model $\mathcal{B}$ can be expanded to a $\Sigma_{\mathcal{A}}$-model of $\Delta(\mathcal{A})$, the diagram of $\mathcal{A}$.

2. $\mathcal{A}$ is elementarily embedded in $\mathcal{B}$ iff $\mathcal{B}$ can be expanded to a $\Sigma_{\mathcal{A}}$-model of $Th(\mathcal{A}_{\mathcal{A}})$, the theory of the $\Sigma_{\mathcal{A}}$-model $\mathcal{A}_{\mathcal{A}}$ (Note: $Th(\mathcal{A}_{\mathcal{A}})$ is the same as the $E\Delta(\mathcal{A})$).

Proof. We first give the proof of statement 1.

($\Rightarrow$) Let $M : \mathcal{A} \to \mathcal{B}$ be an embedding from $\mathcal{A}$ into $\mathcal{B}$. We expand $\mathcal{B}$ to a $\Sigma_{\mathcal{A}}$-model $\mathcal{B}'$ by interpreting all constant symbols from $\Sigma_{\mathcal{A}}$ that correspond to $a \in A_s$ as $M(a)$. For every quantifier-free $\Sigma_{\mathcal{A}}$-formula $\varphi(x_1, \ldots, x_n)$ and a variable interpretation α over $\mathcal{A}$ we have:

$$\begin{array}{llll} \mathcal{A}_{\mathcal{A}} \models \varphi(\alpha(x_1), \ldots, \alpha(x_n)) & \text{iff} & \mathcal{A}, \alpha \models \varphi & \text{by thm 5(1)} \\ & \text{iff} & \mathcal{B}, M(\alpha) \models \varphi & \text{by prop 3} \\ & \text{iff} & \mathcal{B}' \models \varphi(\alpha(x_1), \ldots, \alpha(x_n)). & \text{by thm 5(2)} \end{array}$$

This immediately implies that $\mathcal{B}' \models \Delta(\mathcal{A})$.

($\Leftarrow$) Suppose that the Σ-model $\mathcal{B}$ can be expanded to a $\Sigma_{\mathcal{A}}$-model $\mathcal{B}'$ of $\Delta(\mathcal{A})$. Then we define a domain mapping M to be $M(a) = a^{\mathcal{B}'}$ for all $a \in A_s$, where $A_s \in A$, and show that M is an embedding from $\mathcal{A}$ into $\mathcal{B}$. For any quantifier-free $\Sigma_{\mathcal{A}}$-formula $\varphi(x_1, \ldots, x_n)$ and a variable interpretation α over $\mathcal{A}$ we have:

$$\begin{array}{llll} \mathcal{B}, M(\alpha) \models \varphi & \text{iff} & \mathcal{B}' \models \varphi(\alpha(x_1), \ldots, \alpha(x_n)) & \text{by thm 5(2)} \\ & \text{iff} & \mathcal{A}_{\mathcal{A}} \models \varphi(\alpha(x_1), \ldots, \alpha(x_n)) & \text{by assump. } \mathcal{B}' \models \Delta(\mathcal{A}) \\ & \text{iff} & \mathcal{A}, \alpha \models \varphi. & \text{by thm 5(1)} \end{array}$$

Thus, $\varphi^{\mathcal{B}, M(\alpha)} \iff \varphi^{\mathcal{A}, \alpha}$. By proposition 3, M is an embedding from $\mathcal{A}$ to $\mathcal{B}$.

The proof of statement 2 is identical to the proof of statement 1, except that $\Delta(\mathcal{A})$ is replaced by $Th(\mathcal{A}_{\mathcal{A}})$. □

Theorem 7. *A* Σ*-theory* T *is complete iff any two models of* T *are elementarily equivalent.*

Proof. Follows directly from the definitions of a complete theory and elementary equivalence of models. ☐

Theorem 8. *Let $\mathcal{A}, \mathcal{B}$ be two Σ-models. Then the following claims hold:*

- If $\mathcal{A} \subseteq \mathcal{B}$ then $\mathcal{A} \sim \mathcal{B}$ (a model is equivalent to its submodel)
- If $\mathcal{A} \preceq \mathcal{B}$ then $\mathcal{A} \equiv \mathcal{B}$ (a model is elementary equivalent to its elementary submodel)
- If $\mathcal{A} \cong \mathcal{B}$ then $\mathcal{A} \equiv \mathcal{B}$ (isomorphic models are elementarily equivalent)
- If $\mathcal{A} \equiv \mathcal{B}$ and $\mathcal{B} \equiv \mathcal{C}$, then $\mathcal{A} \equiv \mathcal{C}$ (transitivity of elementary equivalence).
- If $\mathcal{A} \cong \mathcal{B}$ and $\mathcal{B} \cong \mathcal{C}$ then $\mathcal{A} \cong \mathcal{C}$ (transitivity of isomorphism).

Proof. The first two claims follow directly from the definition of submodel and elementary submodel. The third claim is proven by a straightforward induction over the structure of formulas, and the fourth claim follows from the definition of elementary equivalence. The fifth claim follows easily from the definition of isomorphism. ☐

5.2 Feferman's Interpolation Lemma and Robinson's Consistency Theorem

We state the Feferman's interpolation Lemma [Fef68, Fef74] and we use it prove the many-sorted version of the Robinson's Consistency Theorem.

Definition. Let $\tilde{\varphi}$ be the negation normal form of a Σ-sentence φ (i.e. all the negations in $\tilde{\varphi}$ are applied only to the atomic formulas). Define $Un(\varphi) \subseteq S$ and $Ex(\varphi) \subseteq S$ to be sets of sorts such that $s \in Un(\varphi)$ ($s \in Ex(\varphi)$) iff $\tilde{\varphi}$ contains a universal (existential) quantifier over a variable of the sort s. Let $\mathsf{Sort}(\varphi)$ be the set of sorts of all the terms in φ. We write $\mathsf{Const}(\varphi)$, $\mathsf{Fun}(\varphi)$ and $\mathsf{Pred}(\varphi)$ respectively for the sets of constant, function and predicate symbols in φ.

Given functions $F_1, \ldots, F_m$ over arbitrary Σ-sentences to sets, we say that θ is an interpolant for $(\varphi \models \psi)$ w.r.t $F_1, \ldots, F_m$ if (i) θ is a Σ-sentence, (ii) $\varphi \models \theta$ and $\theta \models \psi$ and (iii) $F_i(\theta) \subseteq F_i(\varphi) \cap F_i(\psi)$ for each of $i = 1, \ldots, m$.

Lemma 9. *Feferman's Interpolation Lemma: Suppose $\varphi \models \psi$ for Σ-sentences φ and ψ. Then there is an interpolant θ w.r.t the functions* Const*,* Fun*,* Sort *and* Pred*, and in addition, $Un(\tilde{\theta}) \subseteq Un(\tilde{\varphi})$ and $Ex(\tilde{\theta}) \subseteq Ex(\tilde{\psi})$.*

Next, we use Feferman's interpolation Lemma and compactness to prove the many-sorted version of the Robinson Consistency Theorem.

Theorem 10. *Robinson Joint Consistency Theorem (many-sorted version): Let Σ_1 and $\Sigma_2$be two signatures and let $\Sigma = \Sigma_1 \cap \Sigma_2$. Suppose T is a complete Σ-theory such that $T_1 \supseteq T \subseteq T_2$, where T_1 and T_2 are consistent theories in $\Sigma_1 \Sigma_2$, respectively. Then $T_1 \cup T_2$ is a consistent $\Sigma_1 \cup \Sigma_2$-theory.*

Proof. Suppose $T_1 \cup T_2$ is inconsistent. Then by compactness (also refer [Man96]), there exist finite subtheories $L_1 \subseteq T_1$ and $L_2 \subseteq T_2$ such that $L_1 \cup L_2$ is inconsistent. Let σ_1 be the conjunction of the sentences in L_1, and σ_2 be the conjunction of sentences in L_2. It follows that $\sigma_1 \models \neg\sigma_2$. By the Feferman Interpolation Lemma, we have an interpolant such that $\sigma_1 \models \theta$ and $\theta \models \neg\sigma_2$ where θ is a Σ-sentence. Since $T_1 \models \sigma_1$ we have that $T_1 \models \theta$. Since T_1 is consistent, $T_1 \not\models \neg\theta$, and hence $T \not\models \neg\theta$. Moreover, $T_2 \models \neg\theta$ and by consistency of T_2 we have $T_2 \not\models \theta$, hence $T \not\models \theta$. But this contradicts the hypothesis that T is a complete Σ-theory. $\square$

6 Compatibility and Consistency of the Union Theory

Although Robinson's Joint Consistency Theorem is a classic result which allows one to establish the consistency of the union theory $T_1 \cup T_2$ of the theories T_1 and T_2, the conditions it requires of T_1 and T_2 are too strong. In order for the union theory $T_1 \cup T_2$ to be consistent for two individually consistent theories T_1 and T_2, Robinson's Theorem requires the existence of a theory $T_1 \supseteq T_0 \subseteq T_2$, such that T_0 is complete. Theories of practical interest usually do not have common subtheories which are complete. Following Ghilardi [Ghi03], our goal is to weaken this condition so as to allow larger classes of theories to be combined.

Specifically, instead of requiring that the common subtheory T_0 be complete, we require that T_1 and T_2 be T_0-compatible, and that there are models $\mathcal{M}_1 \models T_1$ and $\mathcal{M}_2 \models T_2$ with a common Σ_0-submodel $\mathcal{A}$.

T_0-compatibility of T_i essentially translates into two ideas. First, T_0 has model-completion T_0^*, which implies that T_0^* admits elimination of quantifiers. This in turn implies that T_0^* is submodel-complete, i.e. if $\mathcal{A}$ is a submodel of a model of T_0^* then $T_0^* \cup \Delta(\mathcal{A})$ is a *complete theory.* Second, if $\mathcal{A}$ is a submodel of a model of T_i then it is also a submodel of a model of T_0^*.

Hence, the idea here is that, if T_1 and T_2 are T_0-compatible theories then one can extend them suitably to T_1' and T_2' respectively, such that there is a complete common subtheory $T_0^* \cup \Delta(\mathcal{A}) \supseteq T_0$ (Theorem 15 in this section, illustrated in Figure 2). By Robinson's Consistency Theorem in the previous section, $T_1' \cup T_2'$ is consistent, which trivially implies the consistency of $T_1 \cup T_2$.

6.1 Submodel Completeness and Quantifier Elimination

We begin by defining the notions of quantifier elimination (QE), submodel-completeness of a theory and subsequently show that QE implies the submodel completeness of a theory. We use the notion of QE, in fact the notion of submodel-completeness of a theory, in Theorem 14 and Theorem 15 from Section 6 to establish that the union of two individually T_0-compatible theories is T_0-compatible.

Definition. Quantifier Elimination, Submodel-Completeness.

Quantifier Elimination: A Σ-theory T is said to *admit elimination of quantifiers* whenever for each Σ-formula $\varphi(x_1, \ldots, x_n)$ there is a quantifier free Σ-formula $\psi(x_1, \ldots, x_n)$ such that:

$$T \models \forall x_1 \ldots \forall x_n (\varphi(x_1, \ldots, x_n) \leftrightarrow \psi(x_1, \ldots, x_n)).$$

We assume that Σ has at least one constant symbol or $n > 0$.

Submodel Completeness: A Σ-theory T is said to be *submodel-complete* whenever $T \cup \Delta(\mathcal{A})$ is a complete $\Sigma_{\mathcal{A}}$-theory for every Σ-submodel $\mathcal{A}$ of any Σ-model of T.

Lemma 11. *(Embedding-Submodel Lemma) Assume there is an embedding from a Σ-model $\mathcal{A}$ into a Σ-model $\mathcal{B}$, where $\Sigma = (P, F, C, S)$. Then there exists a Σ-model $\mathcal{D} \cong \mathcal{B}$ such that $\mathcal{A} \subseteq \mathcal{D}$.*

Proof. Let $\mathcal{A}' \subseteq \mathcal{B}$ be the submodel of $\mathcal{B}$ such that $\mathcal{A}' \cong \mathcal{A}$, by the definition of embedding. Let $H = \{h_s : A_s \to A_s' \mid s \in S\}$ be the isomorphism from $\mathcal{A}$ to $\mathcal{A}'$, and $H' = H^{-1}$. The model $\mathcal{D} = \langle D, I_{\mathcal{D}} \rangle$ is constructed by replacing the elements from the sort-domains of $\mathcal{A}'$ in $\mathcal{B}$ with the corresponding elements of $\mathcal{A}$, and constructing the interpretation $I_{\mathcal{D}}$ as follows:

$$
\begin{aligned}
I_{\mathcal{D}}(c_s) &= \begin{cases} H'(c_s^{\mathcal{B}}), & \text{if } c_s^{\mathcal{B}} \in A'_s \\ c_s^{\mathcal{B}} & \text{otherwise.} \end{cases}, \text{ where } c_s \in C \\
I_{\mathcal{D}}(f_{s_1 \dots s_n, s})(a_1, \dots, a_n) &= H'(f^{\mathcal{B}}_{s_1 \dots s_n, s}(M(a_1), \dots, M(a_n))) \\
& \text{ for all } \langle a_1, \dots, a_n \rangle \in A_{s_1} \times \cdots \times A_{s_n}, \text{and } f \in F \\
I_{\mathcal{D}}(p_{s_1 \dots s_n})(a_1, \dots, a_n) &= p^{\mathcal{B}}_{s_1 \dots s_n}(H(a_1), \dots, H(a_n)) \\
& \text{ for all } \langle a_1, \dots, a_n \rangle \in A_{s_1} \times \cdots \times A_{s_n}, \text{and } p \in P
\end{aligned}
$$

Next we show that $\mathcal{B} \cong \mathcal{D}$ and $\mathcal{A} \subseteq \mathcal{D}$. Define a domain mapping $N = \{n_s : B_s \to D_s \mid s \in S\}$ between $\mathcal{B}$ and $\mathcal{D}$ as follows:

$$
n_s(e_s) = \begin{cases} h'_s(e_s) & \text{if } e_s \in A'_s \\ e_s & \text{otherwise.} \end{cases}
$$

It is easy to see that N is indeed an isomorphism between $\mathcal{B}$ and $\mathcal{D}$, and $\mathcal{A} \subseteq \mathcal{D}$ is by construction of N and $\mathcal{D}$. □

Now, we show that if T admits quantifier elimination then it is submodel-complete.

Theorem 12. *(QE-Submodel-Completeness Theorem) Let T be a Σ-theory for a signature $\Sigma = (P, F, C, S)$. Then statement 1 below implies statement 2 which in turn implies statement 3.*

1. T admits elimination of quantifiers.

2. Whenever $\mathcal{A} \subseteq \mathcal{G}$, $\mathcal{A} \subseteq \mathcal{H}$ for Σ-models $\mathcal{A}$, $\mathcal{G}$ and $\mathcal{H}$ of T, there exists a Σ-model $\mathcal{D}$ of T such that both $\mathcal{G}_{\mathcal{A}}$ and $\mathcal{H}_{\mathcal{A}}$ are elementarily embedded in $\mathcal{D}_{\mathcal{A}}$ (here the existence of $\mathcal{D}_{\mathcal{A}}$ implicitly requires that the domain of $\mathcal{D}$ contains all the elements from the domain of $\mathcal{A}$), where $\mathcal{A}_{\mathcal{A}}$, $\mathcal{G}_{\mathcal{A}}$, $\mathcal{H}_{\mathcal{A}}$ and $\mathcal{D}_{\mathcal{A}}$ are the corresponding canonical $\Sigma_{\mathcal{A}}$-expansions.

3. T is submodel-complete.

The proof of this theorem follows the ideas presented in [WD04].

Proof. ($1 \Rightarrow 2$) Without loss of generality we can assume that $\Sigma_{\mathcal{G}} \cap \Sigma_{\mathcal{H}} = \Sigma_{\mathcal{A}}$. Given that T admits elimination of quantifiers, and for Σ-models $\mathcal{A}$, $\mathcal{G}$, $\mathcal{H}$ of T such that $\mathcal{A} \subseteq \mathcal{G}$, $\mathcal{A} \subseteq \mathcal{H}$, we have to show the existence of a Σ-model $\mathcal{D}$ such that both $\mathcal{G}_{\mathcal{A}}$ and $\mathcal{H}_{\mathcal{A}}$ elementarily embed into $\mathcal{D}_{\mathcal{A}}$.

We construct a Σ-model $\mathcal{D}$ by first finding a $\Sigma_{\mathcal{G}\cup\mathcal{H}}$-model $\mathcal{D}'$ of the theory $Th(\mathcal{G}_{\mathcal{G}}) \cup Th(\mathcal{H}_{\mathcal{H}})$ and then building the Σ-reduct $\mathcal{D}$ of $\mathcal{D}'$. We assume that the constant symbols corresponding to the elements of $\mathcal{A}$ are interpreted as themselves in $\mathcal{D}'$ (otherwise we can always find another model $\mathcal{D}''$ isomorphic to $\mathcal{D}'$ that satisfies this condition by Lemma 11). We then construct the $\Sigma_{\mathcal{A}}$-reduct of $\mathcal{D}'$ which, by the above assumption, is exactly the canonical $\Sigma_{\mathcal{A}}$-expansion $\mathcal{D}_{\mathcal{A}}$ of $\mathcal{D}$. Then we show that $\mathcal{G}_{\mathcal{A}}$ and $\mathcal{H}_{\mathcal{A}}$ elementarily embed into $\mathcal{D}_{\mathcal{A}}$, which completes the proof of ($1 \Rightarrow 2$) .

To show that such a $\mathcal{D}'$ exists we simply need to show that $Th(\mathcal{G}_{\mathcal{G}}) \cup Th(\mathcal{H}_{\mathcal{H}})$ is consistent, where we know trivially that $Th(\mathcal{G}_{\mathcal{G}})$ and $Th(\mathcal{H}_{\mathcal{H}})$ are individually consistent. By the Robinson Joint Consistency Theorem (Theorem 10), it suffices to show that there is no $\Sigma_{\mathcal{A}}$-sentence σ such that

$$Th(\mathcal{G}_{\mathcal{G}}) \models \sigma \quad \text{and} \quad Th(\mathcal{H}_{\mathcal{H}}) \models \neg\sigma.$$

Suppose σ is a $\Sigma_{\mathcal{A}}$-sentence such that $Th(\mathcal{G}_{\mathcal{G}}) \models \sigma$ and $Th(\mathcal{H}_{\mathcal{H}}) \models \neg\sigma$. Let $a_1, \ldots, a_n \in C_{\mathcal{A}}$ be the set of $\Sigma_{\mathcal{A}}$-constant symbols appearing in σ, added to Σ from sort-domains of $\mathcal{A}$. Let $\varphi(x_1, \ldots, x_n)$ be obtained from σ by replacing each a_i for a new variable x_i. Since T admits elimination of quantifiers, there exists a quantifier free Σ-formula $\psi(x_1, \ldots, x_n)$ such that:

$$T \models \forall x_1 \ldots \forall x_n (\varphi \leftrightarrow \psi).$$

Let ψ^* be $\psi(a_1, \ldots, a_n)$ (i.e. the result of substituting a_i for each x_i in ψ). Note that ψ^* is also quantifier free.

Since $\mathcal{G}_{\mathcal{G}} \models \sigma$ we have that $\varphi^{\mathcal{G},\alpha} = \mathsf{true}$ for a variable interpretation α over $\mathcal{A}$ such that $\alpha(x_i) = a_i$ for $i \in \{1 \ldots n\}$ (by construction of φ and Substitution Lemma). We are given that $\mathcal{G} \models T$ and, therefore, $\psi^{\mathcal{G},\alpha} = \mathsf{true}$ for the same α. Hence we conclude that $\mathcal{G}_{\mathcal{A}} \models \psi^*$. Since ψ^* is quantifier free and $\mathcal{A}_{\mathcal{A}} \subseteq \mathcal{G}_{\mathcal{A}}$, we have $\mathcal{A}_{\mathcal{A}} \models \psi^*$ by Theorem 8; similarily, since $\mathcal{A}_{\mathcal{A}} \subseteq \mathcal{H}_{\mathcal{A}}$, $\mathcal{A}_{\mathcal{A}} \models \psi^*$ and ψ^* is quantifier-free, we have $\mathcal{H}_{\mathcal{A}} \models \psi^*$. Hence, we conclude that $\psi^{\mathcal{H},\alpha} = \mathsf{true}$, and since $\mathcal{H} \models T$, we have $\varphi^{\mathcal{H},\alpha} = \mathsf{true}$. This implies that $\mathcal{H}_{\mathcal{A}} \models \sigma$, and consequently, $\mathcal{H}_{\mathcal{H}} \models \sigma$. That is, $Th(\mathcal{H}_{\mathcal{H}}) \models \sigma$, which is a contradiction. Therefore $Th(\mathcal{G}_{\mathcal{G}}) \cup Th(\mathcal{H}_{\mathcal{H}})$ is consistent establishing the existence of the model $\mathcal{D}'$.

Next we show that the Σ-reduct of $\mathcal{D}'$ is the requisite model $\mathcal{D}$. First, observe that $\mathcal{D}_{\mathcal{A}}$, the canonical $\Sigma_{\mathcal{A}}$-expansion of $\mathcal{D}$, is exactly the $\Sigma_{\mathcal{A}}$-reduct of $\mathcal{D}'$ (this is shown in the first paragraph of the proof). Another observation is that $E\Delta(\mathcal{G}_{\mathcal{G}}) = Th(\mathcal{G}_{\mathcal{G}})$ and $E\Delta(\mathcal{H}_{\mathcal{H}}) = Th(\mathcal{H}_{\mathcal{H}})$, because $\Sigma_{\mathcal{G}_{\mathcal{G}}} = \Sigma_{\mathcal{G}}$ and $\Sigma_{\mathcal{H}_{\mathcal{H}}} = \Sigma_{\mathcal{H}}$. It follows that $\mathcal{D}'$ is a model of $E\Delta(\mathcal{G}_{\mathcal{G}})$, and similarly, a model of $E\Delta(\mathcal{H}_{\mathcal{H}})$. Since $\mathcal{A} \subseteq \mathcal{G}$, it trivially follows that $\mathcal{D}'$ is a model of $E\Delta(\mathcal{G}_{\mathcal{A}})$. Similarly, $\mathcal{D}'$ is a model of $E\Delta(\mathcal{H}_{\mathcal{A}})$. Hence, the $\Sigma_{\mathcal{A}}$-reduct $\mathcal{D}_{\mathcal{A}}$ of $\mathcal{D}'$ is also a model of both $E\Delta(\mathcal{G}_{\mathcal{A}})$ and $E\Delta(\mathcal{H}_{\mathcal{A}})$. This implies that $\mathcal{G}_{\mathcal{A}}$ and $\mathcal{H}_{\mathcal{A}}$ elementarily embed into $\mathcal{D}_{\mathcal{A}}$ by the Robinson's Diagram Lemma (Theorem 6, the elementary version).

($2 \Rightarrow 3$) Let $\mathcal{G} \models T$ be any model of T and $\mathcal{A} \subseteq \mathcal{G}$ be any submodel of $\mathcal{G}$; we show that the $\Sigma_{\mathcal{A}}$-model $T \cup \Delta(\mathcal{A})$ is complete. First observe that $\mathcal{G}_{\mathcal{A}} \models T \cup \Delta(\mathcal{A})$. By Theorem 7, it suffices to show that $\mathcal{G}_{\mathcal{A}} \equiv \mathcal{E}$ for each $\Sigma_{\mathcal{A}}$-model $\mathcal{E}$ where $\mathcal{E} \models T \cup \Delta(\mathcal{A})$. Note that $\mathcal{G}_{\mathcal{A}}$ and $\mathcal{E}$ may interpret the constant symbols corresponding to the elements of $\mathcal{A}$ differently. In particular, $\mathcal{G}_{\mathcal{A}}$ interprets constant symbols from the sort-domains of $\mathcal{A}$ as themselves, while $\mathcal{E}$ might not.

Let $\mathcal{E}$ be any $\Sigma_{\mathcal{A}}$-model such that $\mathcal{E} \models T \cup \Delta(\mathcal{A})$. Since $\mathcal{E} \models \Delta(\mathcal{A})$ we conclude that $\mathcal{A}$ embeds into $\mathcal{E}$, by Robinson's Diagram Lemma. Without loss of generality, we assume that $\mathcal{A} \subseteq \mathcal{E}|_{\Sigma}$ (if it is not, then we can always find $\mathcal{E}' \cong \mathcal{E}$ such that $\mathcal{A} \subseteq \mathcal{E}'|_{\Sigma}$, by Lemma 11). For each such $\Sigma_{\mathcal{A}}$-model $\mathcal{E}$, define $\mathcal{H} = \mathcal{E}|_{\Sigma}$ (the Σ-reduct of $\mathcal{E}$). Since $\mathcal{H}$ is a Σ-reduct of $\mathcal{E}$, it follows that $\mathcal{H} \models T$ and $\mathcal{A} \subseteq \mathcal{H}$.

It is easy to construct an isomorphism between $\mathcal{H}_{\mathcal{A}}$ (the canonical $\Sigma_{\mathcal{A}}$-expansion of $\mathcal{H}$) and $\mathcal{E}$, hence $\mathcal{H}_{\mathcal{A}} \cong \mathcal{E}$. We started out with the assumption that $\mathcal{G} \models T$ and $\mathcal{A} \subseteq \mathcal{G}$. We have constructed a Σ-model $\mathcal{H}$, different from $\mathcal{G}$, such that $\mathcal{H} \models T$ and $\mathcal{A} \subseteq \mathcal{H}$. We apply (2) to conclude that there is a $\Sigma_{\mathcal{A}}$-model $\mathcal{D}_{\mathcal{A}}$ into which both $\mathcal{G}_{\mathcal{A}}$ and $\mathcal{H}_{\mathcal{A}}$ embed elementarily.

By Theorem 8, $\mathcal{H}_{\mathcal{A}} \cong \mathcal{E}$ implies $\mathcal{H}_{\mathcal{A}} \equiv \mathcal{E}$, and $\mathcal{H}_{\mathcal{A}} \preceq \mathcal{D}_{\mathcal{A}}$ implies $\mathcal{H}_{\mathcal{A}} \equiv \mathcal{D}_{\mathcal{A}}$. Similarily we conclude that $\mathcal{G}_{\mathcal{A}} \equiv \mathcal{D}_{\mathcal{A}}$. By transitivity of $\equiv$, we conclude that $\mathcal{E} \equiv \mathcal{H}_{\mathcal{A}} \equiv \mathcal{D}_{\mathcal{A}} \equiv \mathcal{G}_{\mathcal{A}}$ and hence $\mathcal{E} \equiv \mathcal{G}_{\mathcal{A}}$. $\square$

6.2 Compatibility and Consistency

One of the most important notions that we use in this paper is T_0-compatibility of a theory T_1, where $T_0 \subseteq T_1$. The idea of T_0-compatibility relies on the notion of model completion of the theory T_0. As explained in the begining of Section 6, T_0-compatibility is a sufficient condition on two individually con-

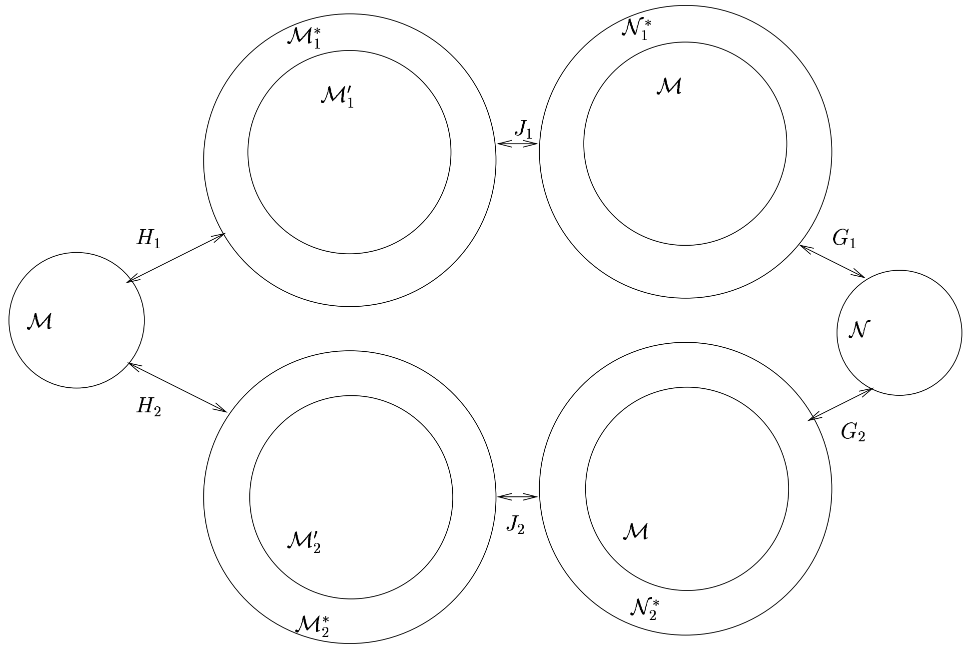

Figure 3: Diagram for the proof of Theorem 14 (T_0-compatibility of $T_1 \cup T_2$)

sistent theories T_1 and T_2 in order for the union $T_1 \cup T_2$ to be T_0-compatible. Here we present a many-sorted version of the definition of model-completion, T_0-compatibility, and prove that the union $T_1 \cup T_2$ is T_0-compatible and consistent, provided the individual theories are.

Definition 13. Model Completion, Compatibility, Stably Infinite theory

Model Completion: Let T be a universal theory (i.e. its axioms are all universal sentences) and $T \subseteq T^*$ for some Σ-theory T^*. We say that T^* is a *model-completion* of T iff

1. Every model of T has an embedding into a model of T^* and
2. T^* admits elimination of quantifiers.

Note. The standard definition of model completion is quite different from the one given here. We refer the reader to the appendix where we show that the standard definition and the defintion presented here are equivalent for MSL.
Compatibility: Let T be a Σ-theory and let T_0 be a universal theory in a subsignature $\Sigma_0 \subseteq \Sigma$. We say that T is T_0-compatible iff

- $T_0 \subseteq T$,
- T_0 has a model completion T_0^* ,
- every model of T embeds into a model of $T \cup T_0^*$.

We now prove two major results of this paper, namely, Union Compatibility Theorem and the Union Consistency Theorem. The proofs are very similar for both theorems.

Theorem 14. *(Union Compatibility Theorem) Let T_1 be a Σ_1-theory and T_2 be a Σ_2-theory; suppose that they are both individually compatible with respect to a universalΣ_0-theory T_0, where $\Sigma_0 = \Sigma_1 \cap \Sigma_2$. Then the $\Sigma_1 \cup \Sigma_2$-theory $T_1 \cup T_2$ is T_0-compatible.*

Proof. We only need to show that every model of $T_1 \cup T_2$ embeds into a model of $T_1 \cup T_2 \cup T_0^*$, since the other conditions for T_0-compatibility are automatically satisfied. If $T_1 \cup T_2$ is inconsistent, then the conclusion follows trivially.

Suppose $T_1 \cup T_2$ is consistent and $\mathcal{M}$ is a $\Sigma_1 \cup \Sigma_2$-model of $T_1 \cup T_2$. We need to show that $\mathcal{M}$ embeds into some $\Sigma_1 \cup \Sigma_2$-model $\mathcal{N}$ of the $\Sigma_1 \cup \Sigma_2$-theory $T_1 \cup T_2 \cup T_0^*$ (The construction is depicted in Figure 3).

First we construct a suitable model for $T_1 \cup T_2 \cup T_0^*$.

Observe that since $\mathcal{M} \models T_1 \cup T_2$ we can conclude $\mathcal{M} \models T_1$ and $\mathcal{M} \models T_2$. By the definition of T_0-compatibility, we have that $\mathcal{M}$ embeds into some Σ_i-model $\mathcal{M}_i^* \models T_i \cup T_0^*$, for $i \in \{1, 2\}$. In other words, $\mathcal{M} \cong \mathcal{M}_i' \subseteq \mathcal{M}_i^*$.

By the Embedding-Submodel Lemma (Lemma 11), the isomorphisms $\mathcal{M} \cong \mathcal{M}_1'$ and $\mathcal{M} \cong \mathcal{M}_2'$ imply that there exist $\mathcal{N}_1^*$ and $\mathcal{N}_2^*$ isomorphic to $\mathcal{M}_1^*$ and $\mathcal{M}_2^*$ respectively, satisfying the following two conditions:

- $\mathcal{M} \subseteq \mathcal{N}_1^*$ and $\mathcal{M} \subseteq \mathcal{N}_2^*$ and
- the sort-domains of $\mathcal{M}_i^*$ and $\mathcal{N}_i^*$ are the same except possibly for elements belonging to the submodels $\mathcal{M}_i'$ and $\mathcal{M}$ respectively.

Since $\mathcal{M} \subseteq \mathcal{N}_i^*$, it follows that $\Delta(\mathcal{M}) \subseteq E\Delta(\mathcal{N}_1^*) \cup E\Delta(\mathcal{N}_2^*)$. It is easy to see that $T_0^* \subseteq E\Delta(\mathcal{N}_1^*) \cup E\Delta(\mathcal{N}_2^*)$ and consequently $T_0^* \cup \Delta(\mathcal{M}) \subseteq E\Delta(\mathcal{N}_1^*) \cup E\Delta(\mathcal{N}_2^*)$ (Note that trivially $T_1 \cup T_2 \cup T_0^* \subseteq E\Delta(\mathcal{N}_1^*) \cup E\Delta(\mathcal{N}_2^*)$).

Observe that $\mathcal{M}$ is trivially a submodel of some model of T_0^* (in fact, $\mathcal{M} \subseteq \mathcal{N}_i^*$ where $\mathcal{N}_i^* \models T_0^*$). Since T_0^* admits elimination of quantifiers, it follows that T_0^* is submodel-complete by QE Submodel-Completeness Theorem (Theorem 12). By the definition of submodel-completeness, we conclude that $T_0^* \cup \Delta(\mathcal{M})$ is a complete theory.

Since $T_0^* \cup \Delta(\mathcal{M})$ is a complete theory and $T_0^* \cup \Delta(\mathcal{M}) \subseteq E\Delta(\mathcal{N}_1^*)$ and $T_0^* \cup \Delta(\mathcal{M}) \subseteq E\Delta(\mathcal{N}_2^*)$, we conclude that $E\Delta(\mathcal{N}_1^*) \cup E\Delta(\mathcal{N}_2^*)$ is consistent, by the many-sorted version of Robinson's Joint Consistency Theorem (Theorem 10). Let $\mathcal{N}'$ be a $(\Sigma_1)_{\mathcal{N}_1^*} \cup (\Sigma_2)_{\mathcal{N}_2^*}$-model of $E\Delta(\mathcal{N}_1^*) \cup E\Delta(\mathcal{N}_2^*)$. We have established that $T_1 \cup T_2 \cup T_0^*$ is consistent.

We now show that $\mathcal{M}$ embeds into $\mathcal{N} = \mathcal{N}' |_{\Sigma_1 \cup \Sigma_2}$.

Let H_i be the embedding from $\mathcal{M}$ into $\mathcal{M}_i^*$, for $i \in \{1, 2\}$. Let J_i be the embedding from M_i^* into $\mathcal{N}_i^*$(in fact J_i are isomorphisms which automatically implies that they are embeddings as well) and K_i be the embedding from $\mathcal{N}_i^*$ into $\mathcal{N}'$ (The existence of K_i follows from the Robinson Diagram Lemma).

Observe that for all elements m_s of any sort-domain M_s of $\mathcal{M}$ we have $K_1(m_s) = K_2(m_s)$. (Recall that $\mathcal{M} = \langle M, I_{\mathcal{M}} \rangle$ where $M = \{M_s \mid s \in S\}$ is an S-indexed family of sort-domains.) The reasoning is as follows: Assume that $K_1(m_s) \neq K_2(m_s)$ and let $K_1(m_s) = a_s$, $K_2(m_s) = b_s$, $a_s \neq b_s$, where a_s, b_s are distinct elements of sort-domain N_s of $\mathcal{N}'$. Note that $\mathcal{N}'$ is a model for $E\Delta(\mathcal{N}_1^*) \cup E\Delta(\mathcal{N}_2^*)$ and that m_s is a constant symbols in the signature $(\Sigma_1)_{\mathcal{N}_1^*} \cup (\Sigma_2)_{\mathcal{N}_2^*}$. By the defintion of embeddings K_1 and K_2, it follows that $m_s^{\mathcal{N}'} = a_s$, $m_s^{\mathcal{N}'} = b_s$ which implies $a_s = b_s$. This is a contradiction and hence $K_1(m_s) = K_2(m_s)$.

Having established that K_i behave the same for elements of $\mathcal{M}$, it is easy to see that the $H_1 \circ J_1 \circ K_1 \cup H_2 \circ J_2 \circ K_2$ is an embedding from $\mathcal{M}$ into $\mathcal{N} = \mathcal{N}' |_{\Sigma_1 \cup \Sigma_2}$. $\square$

Theorem 15. *(Union Consistency Theorem). Let $\mathcal{N}_1$ be a Σ_1-model of T_1 and let $\mathcal{N}_2$ be a Σ_2-model of T_2, where T_1, T_2 are T_0-compatible theories for the Σ_0-theory $T_1 \supseteq T_0 \subseteq T_2$ and $\Sigma_0 = \Sigma_1 \cap \Sigma_2$; suppose also that $\mathcal{N}_1$ and $\mathcal{N}_2$ share a common Σ_0-submodel $\mathcal{A}$. Then there is a $(\Sigma_1 \cup \Sigma_2)_{\mathcal{A}}$-model $\mathcal{M} \models T_1 \cup T_2$ and two $(\Sigma_i)_{\mathcal{A}}$-embeddings $\mathcal{N}_i \to \mathcal{M}$ $(i = 1, 2)$.*

Proof. We are given that T_1 and T_2 are T_0-compatible and consequently we can assume that $\mathcal{N}_1$ is a submodel of a model $\mathcal{M}_1$ of $T_1 \cup T_0^*$ and $\mathcal{N}_2$ is a submodel of a model $\mathcal{M}_2$ of $T_2 \cup T_0^*$(by Lemma 11), where T_0^* is the model-completion of T_0. We can also assume that the sort-domains of $\mathcal{M}_1$ and $\mathcal{M}_2$ corresponding to the non-shared sorts are pairwise disjoint. Also, it is easy to see that $\mathcal{A}$ is a submodel of $\mathcal{M}_1$ and $\mathcal{M}_2$. Our first goal is to show that the elementary diagrams $E\Delta(\mathcal{M}_1)$ and $E\Delta(\mathcal{M}_2)$ of $\mathcal{M}_1$ and $\mathcal{M}_2$ respectively are jointly consistent as a $(\Sigma_1)_{\mathcal{M}_1} \cup (\Sigma_2)_{\mathcal{M}_2}$-theory.

Since $T_0 \subseteq T_1$ and $T_0 \subseteq T_2$ it follows that $\mathcal{A}$ is a Σ_0-submodel of some model of T_0 (in fact, $\mathcal{M}_i \models T_0$ and $\mathcal{A} \subseteq \mathcal{M}_i$ for $i \in \{1, 2\}$). Since every model of T_0 embeds into a model of T_0^*, we have that $\mathcal{A}$ is a Σ_0-submodel of some model of T_0^* (by Lemma 11).

Since T_0^* admits elimination of quantifiers (by definition of model completion), T_0^* is submodel-complete by QE Submodel-Completeness Theorem (Theorem 12). Consequently, we have that $T_0^* \cup \Delta(\mathcal{A})$ is a complete $\Sigma_0^{\mathcal{A}}$-theory. Since $\mathcal{A}$ is a submodel of $\mathcal{M}_1$ and $\mathcal{M}_2$, and both are models of T_0^*, it follows that $E\Delta(\mathcal{M}_1)$ and $E\Delta(\mathcal{M}_2)$ are $(\Sigma_1)_{\mathcal{M}_1}$ and $(\Sigma_2)_{\mathcal{M}_2}$-extensions of the theory $T_0^* \cup \Delta(\mathcal{A})$, respectively. By the Robinson Joint Consistency Theorem (Theorem 10) $E\Delta(\mathcal{M}_1) \cup E\Delta(\mathcal{M}_2)$ is a consistent theory, and any model $\mathcal{M}$ of $E\Delta(\mathcal{M}_1) \cup E\Delta(\mathcal{M}_2)$ is also a model of $T_1 \cup T_2$. The existence of embeddings from $\mathcal{M}_1$ and $\mathcal{M}_2$ into $\mathcal{M}$ follows directly from the Robinson's Diagram Lemma (Lemma 6). □

7 The Combination Method

We now have all the tools necessary to build a complete combination procedure for the $\Sigma_1 \cup \Sigma_2$-theory $T_1 \cup T_2$, i.e. a (semi)-decision procedure to determine whether $T_1 \cup T_2 \models \varphi$ for any universal $\Sigma_1 \cup \Sigma_2$-sentence φ, where T_1 is a Σ_1-theory, T_2 is a Σ_2-theory and Σ_1, Σ_2 may overlap (i.e $\Sigma_0 = \Sigma_1 \cap \Sigma_2$ may have constant, function, predicate and sort symbols common to Σ_1 and Σ_2). We first describe a mathematical method which under certain sufficiency conditions can be turned into a semi-decision procedure. In Subsection 7.2 & Subsection 7.3, we prove the most important theorems of this work, namely the Generated Common Submodel Theorem, Finite Residue Chain Theorem and the Union Completeness Theorem. Although some of the theorems and proofs are similar to the FOL case [Ghi03], there are significant differences. We use these theorems to establish that the combination procedure is complete.

7.1 The Combination Procedure

We present a mathematical construction, which under certain sufficiency conditions can be turned into a (semi)-decision procedure. Let T_1 and T_2 be T_0-compatible where T_0 is a universal Σ_0-theory and $T_1 \supseteq T_0 \subseteq T_2$. The following definitions are needed to describe this construction.

Notation 16. For better readability we shall denote the signature $(\Sigma_i)_{\bar{a}}$ as $\Sigma_i^{\bar{a}}$, the signature obtained by extending Σ_i with constant symbols $\bar{a} = \{a_1, \ldots, a_n\}$.

Definition. Positive residue chain, saturated set of clauses, generated model.

Positive residue chain. A finite list of positive ground Σ_0-clauses

$$C_1, \ldots, C_n,$$

is called a *positive residue chain (or finite residue chain)* if for every $k = 1, \ldots, n$ either

$$T_1 \cup \{C_1, \ldots, C_{k-1}\} \models C_k$$

or

$$T_2 \cup \{C_1, \ldots, C_{k-1}\} \models C_k$$

Saturated set of clauses. A set Γ_0 of positive ground $\Sigma_0^{\bar{a}}$-clauses is *saturated* iff it is closed under the following two rules:

$$T_1 \cup \Gamma_1 \cup \Gamma_0 \models C \quad \Longrightarrow \quad C \in \Gamma_0$$
$$T_2 \cup \Gamma_2 \cup \Gamma_0 \models C \quad \Longrightarrow \quad C \in \Gamma_0$$

for all positive ground $\Sigma_0^{\bar{a}}$-clauses C. Here $\bar{a} = \{a_1, \ldots, a_n\}$ is a finite set of constant symbols not occuring in $\Sigma_1 \cup \Sigma_2$ and decorated by the sorts from Σ_0, the signatures $\Sigma_j^{\bar{a}}$ are the extensions of Σ_j with constant symbols $\bar{a}$ for $j \in \{0, 1, 2\}$, and Γ_i are sets of $\Sigma_i^{\bar{a}}$-sentences, $i \in \{1, 2\}$.

Generated model. Let $\mathcal{A} = \langle A, I_{\mathcal{A}} \rangle$ be a Σ-model, $\Sigma = (P, F, C, S)$, and let $X = \{X_s \mid X_s \subseteq A_s, s \in S\}$ be an S-indexed family of sets, an element-wise subset of A. We say that the Σ-submodel $\mathcal{B} = \langle B, I_{\mathcal{A}} \rangle$ of $\mathcal{A}$ is *generated by* X if for every sort $s \in S$ we have

$$B_s = \{t(x_1, \ldots, x_n)^{\mathcal{A}, \alpha} \mid \alpha(x_{s'}) \in X_{s'} \text{ for all } x_{s'} \in \{x_1, \ldots, x_n\}\},$$

for all Σ-terms $t(x_1, \ldots, x_n)$, where $x_{s'}$ denotes a variable of sort s'.

Note that such a submodel $\mathcal{B}$ does not always exist, since it is possible for a sort-domain B_s to be empty if constructed as above. Therefore, care must be taken to ensure that every sort-domain in $\mathcal{B}$ is non-empty for each sort in Σ. For example, it is sufficient (although too strong in practice) to require that if $X_s = \emptyset$ for some sort $s \in S$, then there exists a constant symbol $c_s \in C$.

We now describe the mathematical construction for converting the problem $T_1 \cup T_2 \models \varphi$ into a problem of checking whether $\mathsf{false} \in \Gamma_0$ for a set of ground formulas Γ_0 (whose construction is described below), where φ is a universal $\Sigma_1 \cup \Sigma_2$-sentence, and theories T_1 and T_2 are T_0-compatible for $T_1 \supseteq T_0 \subseteq T_2$, a universal Σ_0-theory.

Formally, we define the *combination method* D for solving the entailment problem $T_1 \cup T_2 \models \varphi$ as follows:

1. *Purify* the negation of the given $\Sigma_1 \cup \Sigma_2$-sentence $\neg\varphi$ into Γ_1 and Γ_2 (the sets of ground $\Sigma_1^{\bar{a}}$ and $\Sigma_2^{\bar{a}}$-formulas, respectively), which is always possible whenever φ is a universal sentence.

2. Check whether $\mathsf{false} \in \Gamma_0$, where Γ_0 is the set of all positive ground $\Sigma_0^{\bar{a}}$-clauses from all the residue chains of $T_1 \cup \Gamma_1$ and $T_2 \cup \Gamma_2$. If $\mathsf{false} \in \Gamma_0$, then conclude that $T_1 \cup T_2 \models \varphi$. Otherwise conclude that $T_1 \cup T_2 \not\models \varphi$.

The correctness and completeness of this method follows from the Finite Residue Chain Theorem (20 below) and the fact that Γ_0 is saturated (shown below).

Step 2 is the only non-trivial step from the decidability point of view. It is obvious that if Γ_0 is recursively enumerable, then D is a semi-decision procedure; i.e. it will always terminate with the correct answer when $T_1 \cup T_2 \models \varphi$, but might not terminate otherwise. Similarly, D becomes a decision procedure whenever Γ_0 is recursive.

Theorem 17. Γ_0 *is saturated.*

Proof. We prove this by showing that for every positive ground$\Sigma_0^{\bar{a}}$-clause C such that $T_i \cup \Gamma_i \cup \Gamma_0 \models C$, for $i = 1$ or $i = 2$, we have a positive residue chain for C, and hence, by construction of Γ_0, we have that $C \in \Gamma_0$ thus satisfying the conditions required of a saturated set.

Given $T_i \cup \Gamma_i \cup \Gamma_0 \models C$, it follows that there is a finite subset F of Γ_0 such that $T_i \cup \Gamma_i \cup F \models C$, by compactness. By construction of Γ_0, each element of F has a positive residue chain. To get a positive residue chain for C, we simply chain the residue chains of F, and this works because for any two positive residue chains R_1 and R_2 it is easily proved that R_1, R_2 or R_2, R_1 is also a positive residue chain. □

7.1.1 Semi-Decision Procedure for $T_1 \cup T_2$

Under the following additional conditions, the above mathematical construction can be turned into a semi-decision procedure D:

1. There is an algorithm to construct finite sets Γ_1 and Γ_2, such that $T_1 \cup T_2 \cup \{\neg\varphi\} \models \mathsf{false}$ iff $T_1 \cup T_2 \cup \Gamma_1 \cup \Gamma_2 \models \mathsf{false}$, where Γ_i are finite sets of ground $\Sigma_i^{\bar{a}}$-formulas, $\bar{a}$ is a finite set of fresh uninterpreted constant symbols not present in $\Sigma_1 \cup \Sigma_2$ (new Skolem, or *purification* constants), and φ is a universal $\Sigma_1 \cup \Sigma_2$-sentence.

2. There are algorithms D_i, $i = 1, 2$ for recursive enumeration of the sets of positive ground $\Sigma_i^{\bar{a}}$-clauses $\Delta_i(\Gamma) = \{\varphi \mid T_i \cup \Gamma \models \varphi\}$ for a finite set of ground $\Sigma_i^{\bar{a}}$-formulas Γ and a finite set of constant symbols $\bar{a}$.

It is easy to see that the procedure D is correct, that is, whenever D terminates, it reports the correct answer. The completeness of the procedure D (i.e. if $T_1 \cup T_2 \models \phi$, then D is garunteed to terminate and return true) follows from the recursive enumerability of Γ_0 (shown below) and the Finite Residue Chain Theorem.

The proof for the recursive enumerability of Γ_0 is essentially as follows: Each positive ground $\Sigma_0^{\bar{a}}$-clauses $\varphi \in \Gamma_0$ is the last clause of a positive residue chain, by construction of Γ_0. We map each positive residue chain to a tuple of positive numbers. Conversely, each tuple of positive numbers corresponds to only a finite number of positive residue chains. It is known that the set of tuples of positive numbers is recursively enumerable. To recursively enumerate Γ_0, we enumerate tuples of positive numbers, and hence the positive residue chains and consequently the formulas in Γ_0.

Theorem 18. *Under the above conditions, Γ_0 is recursively enumerable.*

Proof. Given the sets of sentences T_1, T_2 and sets of ground formulas Γ_1, Γ_2 satisfying the above conditons, we construct a recursive enumeration of a set of positive ground $\Sigma_0^{\bar{a}}$-clauses, and prove that it is indeed an enumeration of Γ_0.

Let $\Delta_i(\Gamma, k)$ denote the k-th sentence in the recursive enumeration of $\Delta_i(\Gamma)$ (thus, $\Delta_i(\Gamma, k)$ is a recursive function). For a tuple of integers

$$\langle k_1, \ldots, k_n \rangle \in \mathcal{N}^+,$$

define $\Delta(\langle k_1, \ldots, k_n \rangle)$ to be θ_n which is the last element in the sequence $\theta_0 \subseteq \theta_1 \subseteq \cdots \subseteq \theta_n$ such that:

$$\begin{aligned} \theta_0 &= \emptyset \\ \theta_{i+1} &= \theta_i \cup \bigcup_{j=1}^{k_{i+1}} \Delta_{i \bmod 2}(\Gamma_{i \bmod 2} \cup \theta_i, j). \end{aligned}$$

Here Γ_1 and Γ_2 are the given finite sets of ground $\Sigma_0^{\bar{a}}$-formulas obtained from purification of ϕ in step 1. Notice, that θ_n is finite and computable from $\langle k_1, \ldots, k_n \rangle$, hence, $\Delta(\langle k_1, \ldots, k_n \rangle)$ is a recursive function. Since the set of all finite tuples of integers $\mathcal{N}^+$ is recursively enumerable, the set $\Delta = \bigcup_{\bar{k} \in \mathcal{N}^+} \Delta(\bar{k})$ is also recursively enumerable.

Now we claim that $\Delta = \Gamma_0$. The fact that $\Delta \subseteq \Gamma_0$ is obvious by construction of Δ and definition of Γ_0. For the other direction, $\Gamma_0 \subseteq \Delta$, consider an arbitrary clause $\varphi \in \Gamma_0$ and show that $\varphi \in \Delta$. From definition of Γ_0, it follows that there is a finite residue chain $\varphi_1, \ldots, \varphi_n$ such that $\varphi_n = \varphi$. This residue chain can be partitioned into m subsequences $\bar{\varphi}_1 \bar{\varphi}_2 \cdots \bar{\varphi}_m$ such that

$$\begin{aligned} T_1 \cup \Gamma_1 &\models \bar{\varphi}_1 \\ T_2 \cup \Gamma_2 \cup \bar{\varphi}_1 &\models \bar{\varphi}_2 \\ &\vdots \\ T_{m \bmod 2} \cup \Gamma_{m \bmod 2} \cup \bar{\varphi}_1 \cup \cdots \cup \bar{\varphi}_{m-1} &\models \bar{\varphi}_m \end{aligned}$$

Notice, that since $\Delta_1(\Gamma_1)$ is recursively enumerable, there exists $k_1 \geq 1$ such that

$$\bar{\varphi}_1 \subseteq \theta_1 = \bigcup_{j=1}^{k_1} \Delta_1(\Gamma_1, j).$$

Similarly, there exists $k_2 \geq 1$ such that

$$\bar{\varphi}_2 \subseteq \theta_2 = \theta_1 \cup \bigcup_{j=1}^{k_2} \Delta_2(\Gamma_2 \cup \theta_1, j),$$

and so on. Effectively, we have constructed a tuple of natural numbers $\langle k_1, \ldots, k_m \rangle$ and a sequence of sets of clauses $\theta_1, \ldots, \theta_m$ such that $\bar{\varphi}_i \subseteq \theta_i$ and $\theta_i = \Delta(\langle k_1, \ldots, k_i \rangle)$. In particular, $\varphi = \varphi_n \in \theta_m$, and hence, $\varphi \in \Delta$. The proof is complete. □

7.2 Generated Common SubModel Theorem

The Generated Common Submodel Theorem (Theorem 19 below) states that if there is a saturated set of clauses Γ_0 and $\mathsf{false} \notin \Gamma_0$, then there exist models $\mathcal{M}_1 \models T_1 \cup \Gamma_1 \cup \Gamma_0$ and $\mathcal{M}_2 \models T_2 \cup \Gamma_2 \cup \Gamma_0$ which share a common Σ_0-submodel $\mathcal{A}$ generated by $\bar{a}$. The proof of this theorem is essentially the construction of $\Delta(\mathcal{A})$, the diagram of the common Σ_0-submodel $\mathcal{A}$ generated by $\bar{a}$, from the premise that a saturated Γ_0 exists and $\mathsf{false} \notin \Gamma_0$.

Note. The Theorem 19 does not require that T_1 and T_2 be T_0-compatible. This assumption is needed for Finite Residue Chain Theorem (Theorem 20 below), its corollaries, and the Union Consistency and Compatibility Theorems.

The existence of the gernerated common sub-model is one of the crucial assumptions used in establishing the consistency of the union theory $T_1 \cup T_2$ for two T_0-compatible theories $T_1 \supseteq T_0 \subseteq T_2$ in the Union Consistency Theorem (Theorem 15 of Subsection 6.2).

Although the statement of Theorem 19 and its proof are similar to the FOL case [Ghi03], there are some significant differences. In the corresponding theorem for FOL, Ghilardi [Ghi03] allows for the possibility that $\Sigma_0^{\bar{a}}$ may not have any constant symbols, and in particular $\bar{a}$ may be empty. Also, Ghilardi allows for models with empty universes whereas we disallow that.

In the FOL case, if the signature $\Sigma_0^{\bar{a}}$ has no constant symbols then clearly the domain of the common Σ_0-submodel $\mathcal{A}$ generated by $\bar{a}$ is empty ($\mathcal{A}$ with empty universe is also a model according to Ghilardi's definition of a model). Moreover, if $\bar{a}$ is empty it means that no purification of the input $\Sigma_1 \cup \Sigma_2$-sentence φ is necessary. In other words, φ is a Boolean combination of some Σ_1-sentence φ_1 and Σ_2-sentence φ_2, e.g. $\varphi = \varphi_1 \wedge \varphi_2$. In this case, we can determine whether $T_1 \cup T_2 \models \varphi$, by simply determining whether $T_1 \models \varphi_1$and $T_2 \models \varphi_2$. Hence, the combination problem in such cases is trivial, and we don't need to consider the combination procedure (or invoke the associated theorems) to determine if $T_1 \cup T_2 \models \varphi$. The existence of a common submodel generated by $\bar{a}$ with empty universe is irrelevant and the corresponding theorems hold trivially.

Consider a similar scenario for MSL. Suppose that $\Sigma_0^{\bar{a}}$ has no constant symbols of a particular sort $s \in S$, where S is the set of shared sorts, and $\bar{a}$ is otherwise non-empty. This leads us to the possibility that the submodel $\mathcal{A}$ generated by $\bar{a}$ is such that A_s, the sort-domain of sort s in the submodel $\mathcal{A}$, is emtpy. By definition of a model, we do not allow such possibilities. For the sake of argument, assume that models can have empty sort-domains. Since $\bar{a}$ is non-empty we have a non-trivial combination problem. An empty sort domain A_s implies that terms of sort s maybe non-denoting. Taking into account non-denoting terms may require considerable alterations to the existing framework that we have developed, not to mention altering the definition of a model to allow for empty sort-domains. Consequently, we assume that the signature $\Sigma_0^{\bar{a}}$ is such that the common $\Sigma_0^{\bar{a}}$-submodel $\mathcal{A}$ generated by $\bar{a}$ may not have any emtpy sort-domains.

Theorem 19. *(Generated Common Sub-model Theorem): Given Σ_1-theory T_1 and Σ_2-theory T_2, and sets of ground $\Sigma_1^{\bar{a}}$ and $\Sigma_2^{\bar{a}}$-formulas Γ_1 and Γ_2, respectively, for a finite set of fresh constants $\bar{a} = \{a_1, \ldots, a_n\}$. Assume that $\Sigma_0 = \Sigma_1 \cap \Sigma_2$ is such that any Σ_0-model has a Σ_0-submodel generated by $\bar{a}$.*

Suppose that the set of positive $\Sigma_0^{\bar{a}}$-clauses Γ_0 is saturated and does not contain the empty clause (i.e the atom false*). Then there are $\Sigma_i^{\bar{a}}$-models $\mathcal{M}_i$ $(i = 1, 2)$ such that*

$$\begin{aligned} \mathcal{M}_1 &\models T_1 \cup \Gamma_1 \cup \Gamma_0 \\ \mathcal{M}_2 &\models T_2 \cup \Gamma_2 \cup \Gamma_0, \end{aligned}$$

and moreover, the models $\mathcal{M}_1$ and $\mathcal{M}_2$ have a common Σ_0-submodel.

Note that the assumption about the existence of a Σ_0-submodel generated by $\bar{a}$ for any Σ_0-model is necessary for the theorem to hold. The following counterexample satisfies all the conditions in the theorem except this one, violating the theorem.

$$
\begin{aligned}
\Sigma_1 &= \langle \emptyset,\, \{f_{s,s}\},\, \{c^1_s,\, c^2_s\},\, \{s,\, s'\}\rangle \\
\Sigma_2 &= \langle \emptyset,\, \{f_{s,s}\},\, \{d^1_s,\, d^2_s\},\, \{s,\, s'\}\rangle \\
\Sigma_0 &= \langle \emptyset,\, \{f_{s,s}\},\, \emptyset,\, \{s,\, s'\}\rangle \\
\bar{a} &= \langle a_{s'}\rangle \\
Ax(T_i) &= \{\exists v_s \exists w_s.\, v \not\approx w,\ \forall x_s \forall y_s \forall z_s.\, x \approx y \vee x \approx z \vee y \approx z\} \\
\Gamma_1 &= \{c^1 \not\approx c^2,\, f(c^1) \approx c^1,\, f(c^2) \approx c^2\} \\
\Gamma_2 &= \{d^1 \not\approx d^2,\, f(d^1) \approx d^2,\, f(d^2) \approx d^1\}
\end{aligned}
$$

Both theories T_1 and T_2 restrict the sort-domain for s to contain exactly two elements. The clauses in Γ_0 are the ones constructed from the single literal $a \approx a$. In particular, Γ_0 does not contain false. The theories $T_1 \cup \Gamma_1 \cup \Gamma_0$ and $T_2 \cup \Gamma_2 \cup \Gamma_0$ are individually consistent, and therefore, have models $\mathcal{M}_1$ and $\mathcal{M}_2$. However, these models cannot have a common submodel for the following reason. $\mathcal{M}_1$ must interpret the sort s over $M^1_s = \{m_1, m_2\}$ and $f^{\mathcal{M}_1}(m_i) = m_i$, for $i \in \{1, 2\}$. However, in $\mathcal{M}_2$ we have $M^2_s = \{m_1, m_2\}$ and $f^{\mathcal{M}_2}(m_i) \neq m_i$. Therefore, there is no common submodel of $\mathcal{M}_1$ and $\mathcal{M}_2$.

Observe that if $\bar{a}$ contains a constant symbol a_2 of sort s (which guarantees the existence of a Σ_0-submodel generated by $\bar{a}$ for any Σ_0-model), then Γ_0 has both $f(a_2) \approx a_2$ and $f(a_2) \not\approx a_2$ derived from $T_1 \cup \Gamma_1$ and $T_2 \cup \Gamma_2$, respectively, and hence, contains false as well.

Proof. (of Theorem 19) First we prove (by contradiction) that $T_i \cup \Gamma_i \cup \Gamma_0$ is consistent for $i = 1, 2$. Suppose $T_i \cup \Gamma_i \cup \Gamma_0$ is inconsistent. Then it follows that all ground $\Sigma_0^{\overline{a}}$-formulas are entailed by $T_i \cup \Gamma_i \cup \Gamma_0$ including the empty clause false. Since Γ_0 is saturated it follows that false must be in Γ_0, which is a contradiction.

The theorem requires us to show the existence of two $\Sigma_i^{\overline{a}}$-models $(i = 1, 2)$ $\mathcal{M}_1 \models T_1 \cup \Gamma_1 \cup \Gamma_0$ and $\mathcal{M}_2 \models T_2 \cup \Gamma_2 \cup \Gamma_0$ such that they share the same Σ_0-submodel generated by the elements corresponding to $\overline{a}$. To show this we construct an *exhaustive set* of ground $\Sigma_0^{\overline{a}}$-literals Δ, i.e. for every ground $\Sigma_0^{\overline{a}}$-literal φ either it or its negation is in Δ, and Δ is consistent with $T_i \cup \Gamma_i \cup \Gamma_0$ for $i = 1, 2$. Then we show how to construct the required Σ_0-submodel from Δ.

In order to satisfy all clauses in Γ_0, it is sufficient to satisfy at least one literal in each clause of Γ_0. Intuitively, Δ will provide a minimal assignment to the literals in Γ_0 to satisfy all the the clauses.

We choose a strict total terminating order $>$ over the $\Sigma_0^{\overline{a}}$-atoms and extend it to $\Sigma_0^{\overline{a}}$-clauses (treated as sets of atoms) as follows (the extension is also a strict total terminating order). Let S_1 and S_2 be two sets of atoms. We say that $S_1 > S_2$ if for every atom $a_2 \in S_2$ there is some atom $a_1 \in S_1$ such that $a_1 > a_2$.

First, we define a Γ_0-indexed family of sets of atoms Δ_C^+ by transfinite induction. A clause $C \equiv A \vee A_1 \vee \cdots \vee A_n$ from Γ_0 is called *productive* iff $A_1, \ldots, A_n \notin \Delta_{<C}^+$, where A is the largest atom in C, and $\Delta_{<C}^+ = \bigcup_{D<C} \Delta_D^+$. (For convenience, we assume that the leftmost atom in a clause C is always the largest w.r.t. $>$). If C is productive, then we define $\Delta_C^+ = \{A\} \cup \Delta_{<C}^+$, and otherwise $\Delta_C^+ = \Delta_{<C}^+$.

Next, let $\Delta^+ = \bigcup_{C\in\Gamma_0} \Delta_C^+$ and $\Delta = \Delta^+ \cup \{\neg A \mid A \notin \Delta^+\}$, where A is a ground $\Sigma_0^{\overline{a}}$-atom. It is easy to see that $\Delta \models \Gamma_0$ (since $\Delta^+ \models \Gamma_0$), and we simply need to show that $T_i \cup \Gamma_i \cup \Delta$ is consistent for $i = 1, 2$.

Observe that if the clause $C = A \vee A_1 \vee \cdots \vee A_n$ is productive, and A is the maximum atom in C, then $A_1, \ldots, A_n \notin \Delta^+$.

Suppose now that $T_1 \cup \Gamma_1 \cup \Delta$ is not consistent. By compactness, this implies that there is a finite set $\{\neg B_1, \ldots, \neg B_m, A_1, \ldots, A_n\} \subseteq \Delta$ of ground $\Sigma_0^{\overline{a}}$-literals which is inconsistent with $T_1 \cup \Gamma_1$:

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

or, equivalently:

$$T_1 \cup \Gamma_1 \cup \{A_1, \ldots, A_n\} \models B_1 \vee \cdots \vee B_m. \tag{1}$$

By construction of Δ we know that $B_1, \ldots, B_m \notin \Delta^+$, and there are productive clauses in Γ_0:

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

corresponding to the finite set of ground $\Sigma_0^{\overline{a}}$-atoms $\{A_1, \ldots, A_n\}$ in Δ^+. By simple boolean manipulations of (1) we have:

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

Since $C_1, \ldots, C_n$ are clauses in Γ_0, and Γ_0 is saturated, it follows that the clause

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

is also in Γ_0. By the construction of Δ^+ it follows that some of the atoms $\{A_{11}, \ldots, A_{nk_n}, B_1, \ldots, B_m\}$ are in Δ^+. Since $A_1, \ldots, A_n \in \Delta^+$ it follows that $A_{11}, \ldots, A_{nk_n}$ are not in Δ^+, implying that some of $B_1, \ldots, B_m$ are in Δ^+. Contradiction.

The consistency of $T_2 \cup \Gamma_2 \cup \Delta$ is proven similarly. Thus, we have constructed an exhaustive set of atoms Δ consistent with $T_i \cup \Gamma_i$ for $i = 1, 2$. This implies that there exist models $\mathcal{M}_1$ and $\mathcal{M}_2$ such that:

$$\begin{aligned} \mathcal{M}_1 &\models T_1 \cup \Gamma_1 \cup \Gamma_0 \cup \Delta \\ \mathcal{M}_2 &\models T_2 \cup \Gamma_2 \cup \Gamma_0 \cup \Delta. \end{aligned}$$

Next, we show that these have a common Σ_0-submodel generated by $\overline{a}$.

Let $\mathcal{M}_i \mid_{\Sigma_0^{\overline{a}}}$ be the $\Sigma_0^{\overline{a}}$-reduct of $\mathcal{M}_i$ for $i = 1, 2$. For $i = 1, 2$, let $\mathcal{M}_\Delta^i$ denote the $\Sigma_0^{\overline{a}}$-submodel of $\mathcal{M}_i \mid_{\Sigma_0^{\overline{a}}}$ generated by $\overline{a}$ (such submodel exists by the assumption on Σ_0 and $\bar{a}$). We show that $\mathcal{M}_\Delta^1$ is Σ_0-isomorphic to $\mathcal{M}_\Delta^2$.

First observe that $\mathcal{M}^i \mid_{\Sigma_0^{\overline{a}}} \models \Delta$, since Δ is a set of ground $\Sigma_0^{\overline{a}}$-literals. Since $\mathcal{M}_\Delta^i \subseteq \mathcal{M}^i \mid_{\Sigma_0^{\overline{a}}}$, we can apply proposition 3 to conclude that $\mathcal{M}_\Delta^i \models \Delta$.

Let Q be the set of all ground $\Sigma_0^{\overline{a}}$-terms. Consider two arbitrary terms $t_1, t_2 \in Q$. Since $\mathcal{M}_\Delta^1 \models \Delta$ and Δ is exhaustive, if $\mathcal{M}_\Delta^1 \models t_1 = t_2$, then the atom $t_1 = t_2$ must be in Δ, and since $\mathcal{M}_\Delta^2 \models \Delta$, we also have that $\mathcal{M}_\Delta^2 \models t_1 = t_2$. Similarly, $\mathcal{M}_\Delta^2 \models t_1 = t_2$ implies that $\mathcal{M}_\Delta^1 \models t_1 = t_2$. Thus, for any two terms $t_1, t_2 \in Q$ we have that $t_1^{\mathcal{M}_\Delta^1} = t_2^{\mathcal{M}_\Delta^1}$ if and only if $t_1^{\mathcal{M}_\Delta^2} = t_2^{\mathcal{M}_\Delta^2}$. This, in turn, implies that there exists a bijective domain mapping R between $\mathcal{M}_\Delta^1$ and $\mathcal{M}_\Delta^2$ such that $R(t^{\mathcal{M}_\Delta^1}) = t^{\mathcal{M}_\Delta^2}$ for any $t \in Q$, due to the fact that $\mathcal{M}_\Delta^i$ are submodels generated by $\bar{a}$.[1]

We show that for any constant symbol c, function symbol f, predicate symbol p in $\Sigma_0^{\bar{a}}$, and any variable interpretation α over $\mathcal{M}_\Delta^1$ we have:

$$\begin{aligned} R(c^{\mathcal{M}_\Delta^1}) &= c^{\mathcal{M}_\Delta^2} \\ R(f(x_1, \ldots, x_n)^{\mathcal{M}_\Delta^1, \alpha}) &= f(x_1, \ldots, x_n)^{\mathcal{M}_\Delta^2, R(\alpha)} \\ p(x_1, \ldots, x_n)^{\mathcal{M}_\Delta^1, \alpha} &= p(x_1, \ldots, x_n)^{\mathcal{M}_\Delta^2, R(\alpha)}. \end{aligned}$$

[1] A precise way to construct such bijective domain mapping R is the following: for any a_1 from a sort-domain of $\mathcal{M}_\Delta^1$ find a term $t \in Q$ s.t. $t^{\mathcal{M}_\Delta^1} = a_1$, and define $R(a_1) = t^{\mathcal{M}_\Delta^2}$.

The first equality holds by construction of R, since $c \in Q$. Since $\mathcal{M}^1_\Delta$ is generated by $\bar{a}$, there exist terms $t_1, \ldots, t_n \in Q$ such that $\alpha(x_i) = t_i^{\mathcal{M}^1_\Delta}$ for each $i = 1 \ldots n$. Hence, we have:

$$
\begin{array}{rcll}
R(f(x_1, \ldots, x_n)^{\mathcal{M}^1_\Delta, \alpha}) & = & R(f(t_1, \ldots, t_n)^{\mathcal{M}^1_\Delta}) & \text{by subst. lemma} \\
& = & f(t_1, \ldots, t_n)^{\mathcal{M}^2_\Delta} & \text{since } f(t_1, \ldots, t_n) \in Q \\
& = & f(x_1, \ldots, x_n)^{\mathcal{M}^2_\Delta, R(\alpha)} & \text{by subst. lemma}
\end{array}
$$

The property $p(x_1, \ldots, x_n)^{\mathcal{M}^1_\Delta, \alpha} = p(x_1, \ldots, x_n)^{\mathcal{M}^2_\Delta, R(\alpha)}$ is proven similarly. □

7.3 Finite Residue Chain Theorem

The Finite Residue Chain Theorem (Theorem 20 below) states that $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$ is inconsistent iff there exists a positive residue chain which ends in false. A sketch of the proof of the contrapositive of Theorem 20 is as follows. From the assumption that no positive residue chains end in false, a saturated Γ_0 is constructed such that false $\notin \Gamma_0$. Applying the Generated Common Submodel Theorem we conclude the existence of $\mathcal{M}_1 \models T_1 \cup \Gamma_1 \cup \Gamma_0$ and $\mathcal{M}_2 \models T_2 \cup \Gamma_2 \cup \Gamma_0$ which share a common Σ_0-submodel $\mathcal{A}$ generated by $\bar{a}$. Now applying the Union Consistency Theorem we can conclude that $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$ is consistent.

We now state some assumption needed to prove the next few theorems.

1. For a finite set of constant symbols $\overline{a} = \{a_1, \ldots, a_n\}$ not occuring in $\Sigma_1 \cup \Sigma_2$, let Γ_1 and Γ_2 be finite sets of ground formulas over $\Sigma_1^{\bar{a}}$ and $\Sigma_2^{\bar{a}}$, respectively.

2. There is a universal Σ_0-theory T_0 such that both T_1 and T_2 are T_0-compatible where T_1 is a Σ_1-theory and T_2 is a Σ_2-theory and $\Sigma_0 = \Sigma_1 \cap \Sigma_2$.

3. Σ_0 and $\bar{a}$ are such that every Σ_0-model has a Σ_0-submodel generated by $\bar{a}$.

Theorem 20. *(Finite Residue Chain Theorem) In 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* false.

Proof. If there is a positive residue chain ending with false, then it is easy to show that $(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2) \models \mathsf{false}$. Suppose there is no positive residue chain ending up with false. Let Γ_0 be the set of all positive ground $\Sigma_0^{\overline{a}}$-clauses from all finite residue chains. Clearly, Γ_0 is saturated and does not contain false, hence, Theorem 19 applies. This means that there are models $\mathcal{M}_1 \models T_1 \cup \Gamma_1$ and $\mathcal{M}_2 \models T_2 \cup \Gamma_2$ which have a common Σ_0-submodel $\mathcal{A}$ generated by $\overline{a}$. By Theorem 15, there is a model $\mathcal{M} \models T_1 \cup T_2$ which embeds both $\mathcal{M}_1$ and $\mathcal{M}_2$. Since $\mathcal{M}_i \models \Gamma_i$, we also have $\mathcal{M} \models \Gamma_i$, and thus, $\mathcal{M} \models T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$.

Following is an easy corollary. □

Theorem 21. *(Union Completeness Theorem) In the above assumptions,* $(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2)$ *is inconsistent iff there is a quantifier-free ground* $\Sigma_0^{\overline{a}}$*-formula* φ *such that*

$$T_1 \cup \Gamma_1 \models \varphi \quad \textit{and} \quad T_2 \cup \Gamma_2 \models \neg\varphi.$$

Proof. If $(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2)$ is inconsistent, then by Finite Residue Chain Theorem (Theorem 20) there exists a finite positive residue chain $C_1, \ldots, C_n$ such that $C_n = \mathsf{false}$. Let C_k be an *i-residue* $(i = 1, 2)$, that is, $T_i \cup \Gamma_i \cup \{C_1, \ldots, C_{k-1}\} \models C_k$. Let ψ_k be the quantifier-free ground $\Sigma_0^{\overline{a}}$-formula $\neg C_1 \vee \cdots \vee \neg C_{k-1} \vee C_k$ and let φ be the conjunction of all ψ_k such that C_k is a 1-residue. Clearly, $T_1 \cup \Gamma_1 \models \varphi$. Moreover, by induction, it is easy to see that $T_2 \cup \Gamma_2 \cup \{\varphi\} \models C_j$ for all $j = 1, \ldots, n$, and in particular, $T_2 \cup \Gamma_2 \cup \{\varphi\} \models \mathsf{false}$ for $j = n$. Thus, $T_2 \cup \Gamma_2 \models \neg\varphi$.

The other direction is trivial. □

8 The Many-Sorted Nelson-Oppen Method

One of the most interesting application of the results in the previous sections is the many-sorted version of the Nelson-Oppen combination result. As always, we start this section with some useful definitions, followed by theorems needed to establish the MSL Nelson-Oppen method and finally its statement and proof. First, we establish that the theory of infinite sorts T_S is the model completion of the theory of pure equality $T_\approx$. Next, we show that the notion of stably-infiniteness of a theory T follows from the notion of $T_\approx$-compatibility of T, where $T_\approx \subseteq T$ and $T_\approx$ is the theory of pure equality. We then proceed to derive the MSL Nelson-Oppen method from these theorems and the combination results from section 7.

Definition 22. Empty Signature, Empty theory, infinite model, theory of Infinite sorts, T-equivalent formulas, Arrangement.

Empty Signature: By empty signature $\Sigma_{\approx}$ we mean a signature with no constant, function and relation symbols. The set of sort symbols S of $\Sigma_{\approx}$ may not be empty (there could be infinitely many (countably so) sort symbols). As always, the predicate symbol $\approx$ is part of the signature by default.

Empty Theory: Also called *the theory of pure equality* $T_{\approx}$. It is the theory over the empty signature $\Sigma_{\approx}$ with an empty set of axioms. The literals are well-formed equalities and disequalities over sorted variables. Sentences belonging to the empty theory are built out of these literals in the usual way.

Infinite Model: We say that a Σ-model is infinite in the sorts $S = \{s_1, \ldots\}$, where S is a subset of the set of sorts of Σ, if each sort-domain corresponding to the sorts in S is at least countably infinite. Such models are sometimes referred to as S-infinite model.

Theory of the Infinite sorts T_S: The $\Sigma_{\approx}$-theory of the infinite sorts has only those sentences which assert that "there are at least n distinct elements of the sort s_i" for each $n \in \mathbb{N}$ and for each $s_i \in S$, where S is the set of sorts of $\Sigma_{\approx}$. A model for such a theory is an infinite model in the sorts of S (This follows from the well-known fact that for FOL, i.e., any theory that has arbitrarily large finite models has an infinite model. Apply this fact on a per sort basis to derive the above conclusion).

Stably Infinite Theories: Let T be a Σ-theory. Let S denote the set of sorts in $\Sigma_{\approx}$. We say that a theory T *is stably infinite over* S if any quantifier free Σ-formula is satisfiable in some Σ-model of T iff it is satisfiable in a Σ-model of T infinite in the sorts in S.

T-equivalent formulas: Given a Σ-theory T, two Σ-formulas $\varphi(x_1, \ldots, x_n)$ and $\psi(x_1, \ldots, x_n)$ are *T-equivalent* iff $T \models (\forall x_1 \ldots \forall x_n)\varphi \leftrightarrow \psi$.

Partition: A set $P \subseteq 2^V$ is called a *partition* of a set of variables V if P is the set of equivalent classes of some equivalence relation R over V. Notice, that a partition P completely specifies its equivalence relation R.

Arrangement: An arrangement $\mathsf{ar}(V)$ of a finite set of decorated variables $V = \{v_1, \ldots, v_n\}$ given by a partition P is the maximal set of well-formed (w.r.t. sorts) equalities and disequalities consistent with the equivalence relation R corresponding to P:

$$\begin{aligned}\mathsf{ar}(V) \quad = \quad & \{v_i \approx v_j \mid v_i, v_j \in V \text{ and } v_i \, R \, v_j \text{ for } v_i,\, v_j \text{ of sort } s\} \\ & \cup \{v_i \not\approx v_j \mid v_i, v_j \in V \text{ and not } v_i \, R \, v_j \text{ for } v_i,\, v_j \text{ of sort } s\}\end{aligned}$$

for all $i, j \in \{1, \ldots, n\}$. For instance, if $V := \{v_0, v_1, v_2, v_3\}$ and the partition is $P := \{\{v_0, v_1, v_2\}, \{v_3\}\}$ (Note: v_0,v_1,v_2,v_3 i must be of the same sort), then

$$\mathsf{ar}(V) \quad := \quad \{v_0 \approx v_1, v_0 \approx v_2, v_1 \approx v_2, v_0 \not\approx v_3, v_1 \not\approx v_3, v_2 \not\approx v_3\}.$$

The conjunction of formulas in $ar(V)$ is also referred to as an arrangement of V.

Basic Formulas Over $\Sigma_{\approx}$: For basic formulas we take all the well formed equalities and well formed disequalities.

Inconsistent Formula: We say a formula is inconsistent if it is logically equivalent to false.

Assignment: Any function from propositional variables to $\{\mathsf{true}, \mathsf{false}\}$ is called an assignment.

8.1 QE and the Theory of Infinite Sorts

In this subsection we establish that the theory of the infinite sorts admits elimination of quantifiers. This coupled with the theorems in the previous section allow us to conclude that the theory of infinite sorts T_S is submodel-complete and also that T_S is the model completion of the theory of pure equality $T_{\approx}$. In general the method of elimination of quantifiers is as follows: First, depending on the theory T, we pick out an appropriate set of formulas, called *basic formulas*. Then we show that every $\Sigma_{\approx}$-formula is T_S-equivalent to a boolean combination (i.e formulas built out of basic formulas and logical connectives $\neg, \wedge, \vee$) of basic formulas.

Lemma 23. *Every quantifier-free $\Sigma_{\approx}$-formula $\varphi(v_1, \ldots, v_n)$ is either inconsistent (i.e. logically equivalent to* false*) or is equivalent to a disjunction of finitely many arrangements over $V = \{v_1, \ldots, v_n\}$ where V is the set of free variables of $\varphi(v_1, \ldots, v_n)$.*

Proof. Any quantifier-free $\Sigma_{\approx}$-formula $\varphi(x_1, \ldots, x_n)$ is a boolean combination of $\Sigma_{\approx}$-literals ($\Sigma_{\approx}$-atoms and their negations, in particular $\Sigma_{\approx}$-atoms are equalities and false). Consider the set A of all arrangements over V (there are only finitely many). Construct a disjunction D of all those arrangements $a \in A$ which are consistent with φ (i.e. the conjunction of every such arrangement a with φ has an $\Sigma_{\approx}$-model). We now show that this disjunction D of finitely many arrangements over V is equivalent to φ.

Let p_e be a propositional variable corresponding to any $\Sigma_{\approx}$-equality e occuring in φ and $\neg p_d$ be the propositional literal corresponding to any $\Sigma_{\approx}$-disequality $\neg d$, where $\neg d$ occurs in φ, and let Q denote the set of all such propositional literals. Construct φ_{prop} by replacing each equality in φ by the corresponding propositional variable in Q and replacing each disequality by the corresponding literal in Q. For any arrangement a over V construct a_{prop} similarly. Let D be a disjunction of arrangements and D_{prop} denote the disjunction obtained by replacing the arrangements a in D with a_{prop}.

($\varphi_{prop} \Rightarrow D_{prop}$): Suppose φ_{prop} is true under some assignment f. We show that there is an arrangement a such that a_{prop} is true under f. We construct a partition P over the set $\{v_1, \ldots, v_n\}$ as follows. For every equation $e \equiv v_i \approx v_j$ in φ to which f assigns true, add the set $\{v_i, v_j\}$ to P. For every disequation $d \equiv v_i \not\approx v_j$ in φ to which f assigns false, add the set $\{v_i, v_j\}$ to P. For every $e \equiv v_i \approx v_j$ in φ to which f assigns false, add two sets $\{v_i\}, \{v_j\}$ to P. For every disequation $d \equiv v_i \not\approx v_j$ in φ to which f assigns true, add the two sets $\{v_j\}, \{v_j\}$ to P. Merge those sets in P which have variables in common. We have the requisite partition in P to construct a. Construct an arrangement a over the variables in V given by P and clearly a_{prop} is true under f. Add a_{prop} to D_{prop}. It is easy to see that both φ_{prop} and a_{prop} are true under f. Repeat this for every assignment (there are only finitely many) which makes φ_{prop} true and add the resulting a_{prop} as a disjunct to the disjunction D_{prop}. It is easy to check that D_{prop} logically follows from φ_{prop}.

($D_{prop} \Rightarrow \varphi_{prop}$): Let f be any assignment under which D_{prop} is true. This implies that at least one disjunct, say a_{prop}, is true under f. It follows from the construction of D_{prop} that both a_{prop} and φ_{prop} are true under f. This implies that $D_{prop} \Rightarrow \varphi_{prop}$.

It is easy to show that if $\varphi_{prop} \Leftrightarrow D_{prop}$ then $\varphi \Leftrightarrow D$. If D_{prop} is empty then it follows that φ is inconsistent. □

Theorem 24. *Every $\Sigma_{\approx}$-formula φ is T_S-equivalent to a boolean combination ψ of basic formulas. Moreover, if all the free variables of φ are among $V = \{v_1, \ldots, v_n\}$ then ψ may be chosen so that all its free variables are among $v_1, \ldots, v_n$. In particular, if φ is a sentence, then so is ψ.*

Proof. We prove a slightly stronger statement, that every φ is T_S-equivalent to a disjunction of arrangements of V. The proof is by induction over the structure of φ.

The base case is trivial, since an atomic formula is already an arrangement.

For the Boolean connectives, the theorem follows directly from Lemma 23. The only non-trivial case is the existential quantifier; that is, $\varphi \equiv \exists((v_0)_{s0})\, \psi(v_0, \ldots, v_n)$. By inductive hypothesis, $\psi(v_0, \ldots, v_n)$ is equivalent to a disjunction of arrangements over $V' = \{v_0\} \cup V$:

$$\psi(v_0, \ldots, v_n) \quad \Longleftrightarrow \quad \psi_0 \vee \cdots \vee \psi_p.$$

Replacing ψ with this disjunction in φ yields an equivalent formula:

$$\begin{aligned}
\varphi \quad &\equiv \quad \exists((v_0)_{s0})\, \psi(v_0, \ldots, v_n) \\
&\Longleftrightarrow \quad \exists((v_0)_{s0})\, (\psi_0 \vee \cdots \vee \psi_p) \\
&\Longleftrightarrow \quad (\exists((v_0)_{s0})\, \psi_0) \vee \cdots \vee (\exists((v_0)_{s0})\, \psi_p)
\end{aligned}$$

Now we only need to show that a formula of the form $\exists(v_s)\phi(v, V)$ is equivalent to the disjunction of arrangements over V, where ϕ itself is an arrangement over $\{v\} \cup V$.

Procedure: remove all the (dis)equalities from ϕ which have v, and turn that into a disjunction of arrangements by Lemma 23. To prove: this reduction is T_S-equivalent to $\exists(v_s)\, \phi$.

Let $\psi(v_0, \ldots, v_n)$ be an arbitrary boolean combination of basic formulas. First, by inductive hypothesis it follows that ψ is equivalent to a formula of the form (in disjunctive normal form)

$$\psi_0 \vee \cdots \vee \psi_p$$

where each ψ_i is an open formula. Also, it is easy to see that $\exists(v_n : s_n)\, \psi$ is logically equivalent to

$$(\exists((v_n)_{s_n})\, \psi_0) \vee \cdots \vee (\exists((v_n)_{s_n})\, \psi_p)$$

Using Lemma 23, we conclude that each ψ_i, $i \in \{0, \ldots, p\}$, is either equivalent to false or else a disjunction of finitely many arrangements ψ_{ij} over V given by some partition P_{ij} where $j \in \{0, \ldots, k\}$ is an index over the finitely many arrangments which occur in ψ_i. Assume without any loss of generality that all ψ_i, $i \in \{0, \ldots, p\}$, are equivalent to a disjunction of finitely many arrangements over V. For each arrangement ψ_{ij}, construct ψ^*_{ij} by deleting all equations and disequations in which v_n occurs. Then ψ^*_{ij} is an arrangement over the remaining variables $v_1, \ldots, v_{n-1}$ given by the partition P^*_{ij} where P^*_{ij} is obtained by deleting all occurences of v_n from P_{ij}. It is easy to check that each ψ_{ij} asserts the existence of r_{ij} many equivalence classes of over the variables of type s_n in V. Then it is easy to see that $\exists((v_n)_{s_n})\,\psi_{ij}$ is equivalent to $\sigma^{s_n}_{r_{ij}-1} \wedge \psi^*_{ij}$ (The reason is that eliminating the quantifier over v_n leaves one fewer variable of type s_n and thus the number of equivalence classes over variables of type s_n goes down by at most 1. The remaining equivalence classes are captured by the arrangement over the variables $v_1, \ldots, v_{n-1}$ i.e. ψ^*_{ij}). Let ψ^*_i represent the finite disjunction of all the arrangements ψ^*_{ij} where $j \in \{0, \ldots, k\}$. Let $\sigma^{s_n}_{r_i-1}$ represent the sigma formula which asserts the existence of the largest number of elements of type s_n among all the finitely many sentences $\sigma^{s_n}_{r_{ij}-1}$. It follows that $\exists v_n \psi$ is equivalent to $(\sigma^{s_n}_{r_0-1} \wedge \psi^*_0) \vee \cdots \vee (\sigma^{s_n}_{r_p-1} \wedge \psi^*_p)$. Clearly, the formulas $\sigma_{r_0-1}, \ldots, \sigma_{r_p-1}$ are $T_{\mathcal{S}}$-equivalent to true. The resulting formula $\psi^*_0 \vee \cdots \vee \psi^*_p$ is indeed a boolean combination of basic formulas (in fact arrangements over $v_1, \ldots, v_n$) and all its free variables are among $v_1, \ldots, v_n$. We are done. □

From Theorem 24 we can immediately conclude that the theory T_S admits elmination of quantifiers.

8.2 Compatibility and Stably Infinite Theories

In this subsection we show that if a theory T is stably-infinite over the sorts in $\Sigma_{\approx}$ then it is $T_{\approx}$-compatible. The converse is true as well. To show this we first have to show that the theory of infinite sorts T_S is the model completion of the theory of pure equality $T_{\approx}$ using the results in the previous subsection.

Lemma 25. *The theory of infinite sorts T_S is the model completion of the theory of pure equality $T_{\approx}$.*

Proof. Recall that a theory T^* is a model completion of a universal theory T if the following two conditions are satisfied. First, T^* must admit elimination of quantifiers and secondly every model of T must embed into a model of T^*. From Theorem 24 we can conclude that T_S admits elimination of quantifiers thus satisfying the first condition. To see that every model of $T_{\approx}$ embeds in some model of T_S consider this. From Lemma 23 we have that any quantifier free $\Sigma_{\approx}$-formula φ is either false or equivalent to a disjunction of finitely many arrangements and let us say $\psi(x_1, \ldots, x_n)$ denotes this disjunction equivalent to φ. Let $\mathcal{A} \models T_{\approx}$, $\mathcal{B} \models T_S$ and let $\mathcal{A} \models \psi[a_1, \ldots, a_n]$ for some $a_1, \ldots, a_n$ in the appropriate sort-domains of $\mathcal{A}$. By the pigeon-hole principle a finite arrangement (or disjunction thereof) can only assert the existence of only finitely many elements of any sort in the set S. But $\mathcal{B}$ is infinite in all the sorts in S and it follows that $\mathcal{B} \models \psi[a_1, \ldots, a_n]$. Simply apply proposition 3 to conclude that there is an embedding from $\mathcal{A}$ into $\mathcal{B}$. ☐

Lemma 26. *Let S be a subset of the set of sorts in the signature Σ. A Σ-theory T is stably infnite over the sorts in S iff every model of T embeds into an S-infinite model of T.*

Proof. ($\Rightarrow$) Consider the set Q of all sentences which assert that for each natural number n there are at least n elements in each sort-domains of the sorts in S (i.e. asserting that each sort-domain in S has infinite cardinality). Let $\mathcal{M}$ be any model of T. We show that $T \cup Q \cup \Delta(\mathcal{M})$ is consistent. If not, we have $T \cup Q \models \neg\varphi(a_1, \ldots, a_n)$, where $\varphi(a_1, \ldots, a_n)$ is a finite conjunction of formulas from $\Delta(\mathcal{M})$. This means that there is a quantifier-free Σ-formula $\varphi(x_1, \ldots, x_n)$, a finite set of elements $a_1, \ldots, a_n$ from the sort-domains of $\mathcal{M}$, such that $\varphi(a_1, \ldots, a_n)$ is true in $\mathcal{M}$. As the constants $a_1, \ldots, a_n$ do not belong to the signature Σ, we have that $T \cup Q \models \forall x_1 \cdots \forall x_n \neg\varphi(x_1, \ldots, x_n)$. But $\varphi(x_1, \ldots, x_n)$ is a quantifier-free formula satisfiable in a model of T and, by hypothesis, there is a model of $T \cup Q$ (i.e. an S-infinite model of T) in which $\varphi(x_1, \ldots, x_n)$ is satisfiable too, contrary to the fact that $T \cup Q \models \neg\exists x_1 \cdots x_n \varphi(x_1, \ldots, x_n)$.

($\Leftarrow$) Easily follows from definition of stably-infiniteness and embedding. ☐

Lemma 27. *Let T_i be a Σ_i-theory and let T_0 be a $\Sigma_\approx$-theory such that $\Sigma_\approx \subseteq \Sigma_i$ and $T_0 \subseteq T_i$. Then T_i is stably infinite over S if and only if T_i is T_0-compatible.*

Proof. ($\Rightarrow$) Suppose T_i is stably infinite over the sorts in S. From Lemma 26 we have that every model of T_i embeds into a model of T_i which is infinite over S. We are given that $T_0 \subseteq T_i$. Also, from Lemma 25 we know that T_0 has a model completion T_S (recall that sentences in T_S simply assert that there are n elements for each sort in S, for all $n \in \mathbb{N}$. It easily follows that all models of T_S are models infinite over the sorts in S, also referred to as S-infinite models). It is easy to see that $T_i \cup T_S$ is a consistent theory since T_i has models which is infinite over the sorts in S. From the assumption that T_i is stably-infinite in the sorts over S and the fact that $T_i \cup T_S$ is consistent and its models are precisely the models of T_i which are infinite over the sorts in S, it follows that every model of T_i should embed in a model of $T_i \cup T_S$. We have shown that T_i satisfies all three conditions of T_0-compatibility.

($\Leftarrow$) If T_i is T_0-compatible then we have that every model of T_i embeds into a model of $T_i \cup T_S$ where T_S is the model completion of the $\Sigma_\approx$-theory T_0 from Lemma 25 (Note: we can assume that $T_i \cup T_S$ is consistent without loss of generality). Our goal is to show that every model of T_i embeds into an S-infinite model of T_i. Recall that sentences in T_S simply assert that there are n elements for each sort in S, for all $n \in \mathbb{N}$. It easily follows that all models of T_S are models infinite over the sorts in S, also referred to as S-infinite models. This implies that all models of $T_i \cup T_S$ are necessarily infinite over the sorts in S and hence by T_0-compatibility every model of T_i embeds into an S-infinite model of $T_i \cup T_S$ and hence into an S-infinite model of T_i. Therefore T_i is stably-infinite. $\square$

Lemma 28. *Let T be a stably-infinite Σ-theory over the sorts $S' \subseteq S$, where S is the set of sorts in Σ. Let Γ be a set of quantifier-free ground Σ-formulas consistent with T. Then $T \cup \Gamma$ is stably-infinite over S'.*

Proof. By definition of stably-infiniteness, for any model $\mathcal{M} \models T$ there exists an S'-infinite model $\mathcal{M}^\infty \models T$ such that $\mathcal{M}$ embeds in $\mathcal{M}^\infty$ (by Lemma 26). Since Γ is consistent with T, there is a model $\mathcal{A} \models T \cup \Gamma$, and hence, there is an S'-infinite model $\mathcal{A}^\infty \models T$ such that $\mathcal{A}$ embeds into $\mathcal{A}^\infty$. Since embedding preserves the interpretation of quantifier-free formulas, we conclude that $\mathcal{A}^\infty \models \Gamma$. Thus, we have shown that any model of $T \cup \Gamma$ embeds into an S'-infinite model of $T \cup \Gamma$, and by Lemma 26, $T \cup \Gamma$ is also stably-infinite. $\square$

We are now ready to state and prove the MSL Nelson-Oppen combination result.

8.3 Many-Sorted Nelson-Oppen Theorem

Theorem 29. *(Many-Sorted Nelson-Oppen combination). Given two consistent theories T_1 and T_2 over the signatures Σ_1 and Σ_2, such that*

1. *$\Sigma_0 = \Sigma_1 \cap \Sigma_2$ is the empty signature $\Sigma_\approx$(i.e. the signatures are disjoint but they may share sort symbols);*

2. *T_1 and T_2 are stably-infinite theories over the shared sorts $S \in \Sigma_0$;*

3. *The universal fragment of the theories T_1 and T_2 are individually decidable (i.e. $T_i \models \varphi_i$ is decidable for any universal Σ_i-sentence φ_i, where $i \in \{1, 2\}$).*

Then the following hold:

1. *$T_1 \cup T_2$ is consistent;*

2. *$T_1 \cup T_2$ is decidable; that is, $T_1 \cup T_2 \models \varphi$ is decidable for any universal $\Sigma_1 \cup \Sigma_2$-sentence φ.*

Proof. We construct a nondeterministic algorithm for deciding the problem $T_1 \cup T_2 \models \varphi$ and establish that the algorithm is sound, complete and terminating. This algorithm closely follows the semi-decision procedure given in Section 7.

First, notice that the problem $T \models \varphi$ for some $\Sigma_1 \cup \Sigma_2$-theory T is equivalent to the problem of determining $T \cup \{\neg\varphi\} \models \mathsf{false}$.

1. Purify $\neg\varphi$ into Γ_1 and Γ_2 (such that $T_1 \cup T_2 \cup \{\varphi\} \models \mathsf{false}$ iff $T_1 \cup T_2 \cup \Gamma_1 \cup \Gamma_2 \models \mathsf{false}$), where Γ_i are the sets of ground $\Sigma_i^{\bar{c}}$-formulas, and $\bar{c}$ is a finite set of fresh uninterpreted constant symbols (not present in $\Sigma_1 \cup \Sigma_2$); It is easy to show that for quantifier-free $\Sigma_1 \cup \Sigma_2$-formulas efficient purification algorithms exist [Bar03].

2. If there is a $\Sigma_0^{\bar{c}}$-arrangement A such that both $T_1 \cup \Gamma_1 \cup \{A\}$ and $T_2 \cup \Gamma_2 \cup \{A\}$are consistent, then we conclude that $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$ is consistent (or equivalently $T_1 \cup T_2 \cup \{\varphi\}$ is consistent). The reasoning for this step is as follows: First, observe that an arrangement is a saturated set of clauses (in fact a positive residuce chain which does not contain false) and by definition of arrangement it does not contain false. Hence, we can apply the Generated Common Submodel Theorem (Theorem 7.2) to conclude that there exist two models $\mathcal{M}_1 \models T_1 \cup \Gamma_1 \cup \{A\}$ and $\mathcal{M}_2 \models T_2 \cup \Gamma_2 \cup \{A\}$ such that they have a common submodel $\mathcal{A}$ generated by $\bar{c}$. Also, since T_1 and T_2 are stably-infinite, it follows that so are $T_1 \cup \Gamma_1 \cup \{A\}$ and $T_2 \cup \Gamma_2 \cup \{A\}$ by Lemma 28, and hence, they are $T_{\approx}$-compatible as well (Lemma 27). Now we can apply the Union Consistency Theorem (Theorem 15) to conclude that the $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2 \cup A$ is consistent and hence $T_1 \cup T_2 \models \varphi$.

3. Otherwise (if there is no such arrangement A) the sentence φ is inconsistent with $T_1 \cup T_2$. Since there are no arrangements we can only conclude that there exist positive residue chain with false in them. Consequently $(T_1 \cup \Gamma_1) \cup (T_2 \cup \Gamma_2)$ is inconsistent implying that $T_1 \cup T_2 \not\models \varphi$, by the Finite Residue Chain Theorem (Theorem 20) or the Union Completeness Theorem (Theorem 21).

□

The termination of the algorithm follows from the fact that the number of arrangements is finite, and the problem $T_i \cup \Gamma_i \cup \{A\} \models \mathsf{false}$ is decidable by assumption, for $i \in \{1, 2\}$. The soundness follows from the fact that each step is sound (supported by the theorems). The completeness follows from the fact that if $T_1 \cup T_2 \models \varphi$ then algorithm terminates with the conclusion $T_1 \cup T_2 \models \varphi$, by step 2; (and alternatively if $T_1 \cup T_2 \not\models \varphi$ then the algorithm concludes that $T_1 \cup T_2 \not\models \varphi$ in step 3).

9 Decidability Conditions

There are certain sufficient conditions under which the semi-decision procedure presented in Section 7 can be turned into a decision procedure for the universal fragment of the $\Sigma_1 \cup \Sigma_2$-theory $T_1 \cup T_2$, where T_1 and T_2 are universally decidable Σ_1 and Σ_2-theories, respectively, satisfying the conditions of the Finite Residue Chain Theorem (Theorem 20). Assume that φ is a $\Sigma_1 \cup \Sigma_2$-sentence, and Γ_1 and Γ_2 are sets of ground formulas over signatures $\Sigma_1^{\bar{a}}$ and $\Sigma_2^{\bar{a}}$ respectively, obtained by purifying φ, where $\bar{a}$ is a set of fresh Skolem constants.

Let D be a semi-decision procedure for the universal fragment of theory $T_1 \cup T_2$ satisfying the sufficient conditions given in Section 7, and D_i be a decision procedure for the universal fragment of theory T_i, for $i \in \{1, 2\}$.

Local Finiteness.

We adapt the definition of local finiteness from Ghilardi's paper [Ghi03] as follows.

Definition 30. A Σ-theory T is called *locally finite* w.r.t. a finite set of constant symbols $\bar{a}$ (not necessarily in Σ), if Σ is finite, and there exists a finite set of ground $\Sigma^{\bar{a}}$-terms $\mathcal{T} = \{t_1, \ldots, t_n\}$ such that for every ground $\Sigma^{\bar{a}}$-term q we have $T \models t_i = q$ for some $t_i \in \mathcal{T}$. We call the set of terms $\mathcal{T}$ the *canonical terms of T w.r.t. $\bar{a}$.*

Theory T is *effectively locally finite* if the set $\mathcal{T}$ is computable for any $\bar{a}$, and there is an algorithm to compute the canonical term $t_i \in \mathcal{T}$ for any $\Sigma^{\bar{a}}$-term q.

The proof of the fact that effective local finiteness of the common subtheory T_0 implies the universal decidability of the union theory $T_1 \cup T_2$ is a straightforward extension of the corresponding result by Ghilardi [Ghi03].

One-way Communication between Theories

Another scenario under which the problem $T_1 \cup T_2 \models \varphi$ becomes universally decidable, provided the theories T_1 and T_2 are universally decidable, is described below in an intuitive and procedural way.

First, recall the steps of the semi-decision procedure D from Section 7 for the problem $T_1 \cup T_2 \models \varphi$. In step 2, the set Γ_0 is recursively enumerable, and if $\mathsf{false} \in \Gamma_0$, then we have already shown that D will terminate and establish that $T_1 \cup T_2 \models \varphi$. On the other hand, if $T_1 \cup \Gamma_1 \cup T_2 \cup \Gamma_2$ is consistent, then D may not terminate. Intuitively, D enumerates Γ_0 by exchanging positive ground $\Sigma_0^{\bar{a}}$-clauses between D_1 and D_2 (semi-decision procedures for the problem $T_i \models \varphi_i$ where φ_i is a universal Σ_i-sentence, for $i = 1, 2$ respectively), and in general this exchange may go on forever. However, if at any point in this exchange one of the decision procedures, e.g. D_1, starts producing only those clauses which D_2 can deduce from $T_2 \cup \Gamma_2 \cup \Pi_j$, where Π_j is the set of clauses exchanged so far (i.e. D_1 does not produce any new information), then the problem $T_1 \cup T_2 \models \varphi$ is equivalent to checking $T_2 \cup \Gamma_2 \cup \Pi_j \models \mathsf{false}$, which is decidable. We refer to this scenario as *one-way communication* from D_1 to D_2. Formally, this idea is captured by the following definitions and the subsequent theorem.

Definition 31. For a set of ground $\Sigma_i^{\bar{a}}$-formulas Γ, let Δ be the set of all $\Sigma_0^{\bar{a}}$-clauses such that $T_i \cup \Gamma \models \Delta$, for $i \in \{1, 2\}$. Assume that there is an algorithm to construct a finite set of $\Sigma_0^{\bar{a}}$-clauses θ such that $T_i \cup \Gamma \models \theta$ and the set of all $\Sigma_0^{\bar{a}}$-clauses entailed by $T_i \cup \theta$ is exactly Δ. We say that θ *finitely characterizes* Δ. We denote by $\xi_i(\Gamma)$ the recursive functions which construct θ for a given Γ.

We always assume that the finite characterization functions ξ_i are *monotonic*, that is, $\xi_i(\Gamma_1) \subseteq \xi_i(\Gamma_2)$ whenever $\Gamma_1 \subseteq \Gamma_2$.

Definition 32. For every finite set of ground $\Sigma_i^{\bar{a}}$-formulas Λ_i for $i \in \{1, 2\}$, let Π_1, Π_2 and Π_3 be the sets of all $\Sigma_0^{\bar{a}}$-clauses such that

$$T_1 \cup \Lambda_1 \models \Pi_1 \qquad T_2 \cup \Lambda_2 \cup \Pi_1 \models \Pi_2 \qquad T_1 \cup \Lambda_1 \cup \Pi_2 \models \Pi_3.$$

If $\Pi_3 = \Pi_2$, we say that there is *one way communication* from theory T_1 to T_2.

In other words, the saturated set Γ_0 is reached in one step, and is equal to Π_2.

Theorem 33. *Assume there is one way communication from theory T_1 to T_2. Also, for every set of ground $\Sigma_1^{\bar{a}}$-formulas Λ, let there be a finite characterization $\theta = \xi_1(\Lambda)$. Under the above conditions the problem $T_1 \cup T_2 \models \varphi$ is universally decidable for any universal $\Sigma_1 \cup \Sigma_2$-sentence φ.*

Proof. Let Γ_1 and Γ_2 be the sets of ground $\Sigma_1^{\bar{a}}$ and $\Sigma_2^{\bar{a}}$-formulas from the purification of φ, which are already known to be computable for any universal φ. By the one way communication condition, we know that the set Π_2 of all clauses entailed by $T_2 \cup \Gamma_2 \cup \xi_1(\Gamma_1)$ is saturated. Hence, checking whether $\mathsf{false} \in \Gamma_0$ is equivalent to checking whether $T_2 \cup \Gamma_2 \cup \xi_1(\Gamma_1) \models \mathsf{false}$, which is decidable. □

Strictly Decreasing Measure on Clauses

Assume there is a well-founded measure ρ over positive ground $\Sigma_0^{\bar{a}}$-clauses, and T_i,Λ_i, D and the notion of finite characterization are as described above. Consider a sequence of sets of positive ground $\Sigma_0^{\bar{a}}$-clauses $\emptyset = \Pi_0 \subseteq \Pi_1 \subseteq \cdots$ such that $T_i \cup \Gamma_i \cup \Pi_j \models \Pi_{j+1}$ for some $i \in \{1, 2\}$ and every $j \geq 0$. Let Δ_j be the set of newly added $\Sigma_0^{\bar{a}}$-clauses at the j-th step: $\Delta_j = \Pi_i - \Pi_{i-1}$, and let $\rho(\Delta_j)$ be the measure of the maximum clause in Δ_j w.r.t. ρ. If $\rho(\Delta_j)$ is strictly decreasing after a finite number of steps, i.e. $\rho(\Delta_{j+1}) < \rho(\Delta_j)$ for every $j \geq k$ for some natural number $k \geq 1$, then there is $n \geq k$ such that $\Pi_m = \Pi_n$ for every $m \geq n$ (since ρ is well-founded). If every Π_j is recursive, $\Pi_n = \Gamma_0$ is also recursive[2], and hence, the original problem $T_1 \cup T_2 \models \varphi$ is decidable. We formalize this idea in the following definitions and theorem.

1. There is a well-founded complexity measure ρ over positive ground $\Sigma_0^{\bar{a}}$-clauses.

2. For finite sets of ground $\Sigma_i^{\bar{a}}$-formulas Λ and positive ground $\Sigma_i^{\bar{a}}$-clauses θ, let $\gamma = \xi_i(\Lambda)$ and $\delta = \xi_i(\Lambda \cup \theta)$ be the finite characterizations of all

[2] The reasoning for $\Pi_n = \Gamma_0$ is as follows: Π_n is the set of all positive ground $\Sigma_0^{\bar{a}}$-clauses from all finite residue chains of $T_1 \cup \Gamma_1$ and $T_2 \cup \Gamma_2$, by construction. This is precisely Γ_0, the saturated set that we enumerate in our semi-decision procedure D.

clauses entailed by $T_i \cup \Lambda$ and $T_i \cup \Lambda \cup \theta$, respectively, where $i \in \{1, 2\}$. We assume that the functions ξ_i are *monotonic*, that is, $\xi_i(\Lambda_1) \subseteq \xi_i(\Lambda_2)$ whenever $\Lambda_1 \subseteq \Lambda_2$. If $\rho(\delta - (\gamma \cup \theta)) < \rho(\theta)$, then we say that the *finite characterizations entailed by the theory T_i monotonically decrease w.r.t ρ.*

Theorem 34. *Let ρ denote a well-founded measure over positive ground $\Sigma_0^{\bar{a}}$-clauses, and assume that the finite characterizations entailed by theory T_i monotonically decrease w.r.t ρ, for $i \in \{1, 2\}$. Under the above conditions the union theory $T_1 \cup T_2$ is universally decidable.*

Proof. We establish the decidability of the problem $T_1 \cup T_2 \models \varphi$ for a universal sentence φ by constructing an algorithm.

1. Purify φ into sets Γ_1 and Γ_2.

2. Construct a sequence of finite characterizations:

$$\begin{aligned}\Pi_0 &= \emptyset \\ \Pi_{j+1} &= \xi_1(\Gamma_1 \cup \Pi_j) \cup \xi_2(\Gamma_2 \cup \Pi_j)\end{aligned}$$

   for $j \geq 1$. This construction is algorithmically possible due to the existence of finite characterization (Definition 31).

By Condition 2 and the well-foundedness of the measure ρ, there is $n \geq 1$ such that $\Pi_{n+1} = \Pi_n$. This Π_n is a finite characterization of Γ_0. Since each Π_j is finitely characterizable and hence recursive, Γ_0 is also recursive. This converts D into a decision procedure. □

10 Practical Applications

Our work on combination results for many-sorted theories with overlapping signatures has primarily been motivated by the need to develop a theoretical foundation for our validity checking tools, Cooperating Validity Checker (CVC [SBD02]) and its successor CVC Lite [BB04]. These tools have been based on a presumed extension of unsorted combination methods [Bar03] (similar to Nelson-Oppen) to sorted quantifier-free first-order theories. The combination result in Section 7 reduces in a relatively straightforward way to the many-sorted version of the Nelson-Oppen method, providing a formal basis for the correctness of such an extension.

In the case of CVC Lite, there is a demand for combining theories with overlapping signatures. For example, consider a *theory of bit-vectors*: finite strings of bits with concatenation and substring extraction operators. Bit-vectors also represent integers (in binary encoding) which can be added, subtracted, and compared for equality and inequality. Thus, the signature of the bit-vector theory must include arithmetic operators and integer constants, making it overlap with the theory of linear arithmetic already implemented in CVC Lite.

The bit-vector theory and linear arithmetic happen to satisfy the T_0-compatibility condition for their common subtheory T_0, the universal fragment of Presburger arithmetic. More precisely, the signature Σ_0 of T_0 is

$$\Sigma_0 = \langle <_{\mathbb{Z},\mathbb{Z}},\ +_{\mathbb{Z}\times\mathbb{Z}\to\mathbb{Z}},\ (\cong_m)_{\mathbb{Z},\mathbb{Z}},\ C_{\mathbb{Z}},\ \mathbb{Z}\rangle,$$

where $C_{\mathbb{Z}}$ is the set of all integer constants, $\mathbb{Z}$ is the sort of integer numbers, and $\cong_m$ is a family of congruences modulo m for all natural numbers $m \geq 1$.

The theory of linear arithmetic extends Σ_0 with the sort of real numbers $\mathbb{R}$, the set of real constants $C_{\mathbb{R}}$, and the operators $+_{\mathbb{R}\times\mathbb{R}\to\mathbb{R}}$, $<_{\mathbb{R},\mathbb{R}}$, and $\mathsf{int2real}_{\mathbb{Z}\to\mathbb{R}}$ (the conversion of integers to reals). Bit-vector theory adds a sort $\mathbb{B}_n$ of bit-vectors of length n for every integer $n \geq 1$ (so, its signature has infinitely many sorts), a family of concatenation $@_{\mathbb{B}_i\times\mathbb{B}_j\to\mathbb{B}_{i+j}}$ and extraction $\mathsf{sub}^{l,u}_{\mathbb{B}_n\to\mathbb{B}_{u-l+1}}$ operators for every $i, j, n, u, l \geq 1$, where $0 \leq l < u < n$, and conversion operators $\mathsf{bv2int}_{\mathbb{B}_n\to\mathbb{Z}}$ from bit-vectors to integers.

The model completion of T_0 in this case is Presburger arithmetic (since Presburger arithmetic admits elimination of quantifiers). It is easy to see that both bit-vector and linear arithmetic theories are T_0-compatible. We have already established the first two conditions of T_0-compatibility (i.e. T_0 is a subtheory of both, and T_0 has a model completion T_0^*). The third condition also holds simply because Presburger arithmetic is a subtheory of both bit-vectors (T_1) and linear arithmetic (T_2); that is, $T_i \cup T_0^* = T_i$, since $T_0^* \subseteq T_i$.

Finally, every ground bit-vector formula can be equivalently translated into a ground Σ_0-formula, and therefore, the combination of bit-vectors and linear arithmetic satisfies the one-way communication decidability condition stated in Section 9. This makes the combination of the two theories decidable. Similar arguments can be made for a theory of lists with the length operator, which also shares symbols with the theory of linear arithmetic.

11 Conclusions

We have presented a new combination result for many-sorted first-order theories with overlapping signature, a non-trivial extension of Ghilardi's work [Ghi03]. Besides the completeness results, we have also given new practical decidability conditions and illustrated their use by examples of theories relevant to CVC Lite [BB04]. As a bonus, the many-sorted version of the Nelson-Oppen combination directly follows from our results.

The combination conditions imposed on the individual theories (in particular, T_0-compatibility) are still too strong for many practical theories, and also quite involved for most tool developers. There is a lot of work to be done to make these conditions more practical and easier to check. We intend to use these results to combine theories with overlapping signatures in CVC Lite. Extensions to order-sorted logics (many-sorted logics with subsorts) combining the results presented here and in [TZ04] is another interesting direction of research.

Index

12 Appendix A

We establish the equivalence between standard definition of model-completion [CK98] and the one used in this paper for universal theories.

Definition 35. Let T be a Σ-theory and let $T^* \supseteq T$ be a Σ-theory as well; 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 Σ-model $\mathcal{M}$ of T, we have that $T^* \cup \Delta(\mathcal{M})$ is a complete $(\Sigma)_{\mathcal{M}}$-theory.

First, we prove the following lemma.

Lemma 36. *Let T be a Σ-theory and let $T^* \supseteq T$ be a Σ-theory as well; 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^* admits elimination of quantifiers.*

Proof. Suppose T^* satisifies the conditions above. We show that for every Σ-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 $T^* \cup \Delta(\mathcal{M})$, a Σ-formula $\varphi(x_1, \ldots, x_n)$ and a tuple of elements $\langle a_1, \ldots, a_n \rangle$ from the appropriate sort-domains of $\mathcal{M}$.We show that the $(\Sigma)_{\mathcal{M}}$-sentence $\varphi(a_1, \ldots, a_n)$ is true in $\mathcal{N}_1$ iff it is true in $\mathcal{N}_2$. This shows that the theory $T^* \cup \Delta(\mathcal{M})$ is complete since the models of a complete theory are elementarily equivalent.

By the Robinson's Diagram Lemma $\mathcal{M}$ is a common substructure of $\mathcal{N}_1$ and $\mathcal{N}_2$; moreover φ is T^*-equivalent to a quatifier-free formula $\varphi'(x_1, \ldots, x_n)$, hence if $\mathcal{N}_1 \models \varphi(a_1, \ldots, a_n)$ then $\mathcal{N}_1 \models \varphi'(a_1, \ldots, a_n)$. Consequently $\mathcal{M} \models \varphi'(a_1, \ldots, a_n)$ and $\mathcal{N}_2 \models \varphi'(a_1, \ldots, a_n)$ thus establishing that $\varphi(a_1, \ldots, a_n)$ is true in $\mathcal{N}_2$.

If T is universal then the converse is easy to show. □

References

[Bar03] C. Barrett. *Checking Validity of Quantifier-Free Formulas in Combinations of First -Order Theories.* PhD thesis, Stanford University, 2003.

[BB03] Clark Barrett and Sergey Berezin. A proof-producing boolean search engine. In *PDPAR'03 Workshop, Miami, Florida,* July 2003.

[BB04] Clark Barret and Sergey Berezin. CVC Lite: A new implementation of the cooperating validity checker. In Rajeev Alur and Doron A. Peled, editors, *Computer-Aided Verification (CAV'04)*, LNCS. Springer Verlag, July 2004.

[BT97] F. Baader and C. Tinelli. A new approach for combining decision procedures for the word problem, and its connection to the Nelson-Oppen combination method. In W. McCune, editor, *14th International Conference on Computer Aided Deduction*, LNCS, pages 19–33. Springer-Verlag, 1997.

[CK98] C. C. Chang and H. J. Keisler. *Model Theory*. Elsevier Science Ltd, 3rd edition, 1998.

[DNS] D. Detlefs, G. Nelson, and J. Saxe. Simplify theorem prover. `http://research.compaq.com/SRC/esc/Simplify.html`.

[Fef68] S. Feferman. Lectures on proof theory. In M. Löb, editor, *proc. of Summer School in Logic, Leeds 67*, volume 70 of *LNM*, pages 1–107. Springer, 1968.

[Fef74] S. Feferman. Applications of many-sorted interpolation theorems. In L. Henkin at al., editor, *Tarski Symposium*, volume XXV of *AMS Proc. of Symposia in Pure Mathematics*, pages 205–223, 1974.

[FJOS03] C. Flanagan, R. Joshi, X. Ou, and J. Saxe. Theorem proving using lazy proof explication. In Warren Hunt and Fabio Somenzi, editors, *Computer-Aided Verification (CAV'03)*, volume 2725 of *LNCS*, pages 355–367. Springer Verlag, July 2003.

[FORS01] Jean-Christophe Filliâtre, Sam Owre, Harald Rueß, and N. Shankar. ICS: integrated canonizer and solver. In *Computer-Aided Verification (CAV'01)*, volume 2101 of *LNCS*, pages 246–249. Springer Verlag, 2001.

[Ghi03] S. Ghilardi. Reasoners' cooperation and quantifier elimination. Technical report, Dipartimento di Scienze dell'Informazione, Università degli Studi di Milano, 2003.

[Man96] Maria Manzano. *Extensions of First Order Logic*. Cambrige University Press, 1996.

[NO79] G. Nelson and D. Oppen. Simplification by cooperating decision procedures. *ACM Transactions on Programming Languages and Systems*, 1(2):245–57, 1979.

[Rob56] A. Robinson. A result on consistency and its applications to the theory of definition. In *Proc. of the Royal Academy of Sciences, Amsterdam*, A. 59, pages 47–58, 1956.

[RS01] H. Ruess and N. Shankar. Deconstructing Shostak. In *16th IEEE Symposium on Logic in Computer Science*, 2001.

[SBD02] A. Stump, C. Barrett, and D. Dill. CVC: a Cooperating Validity Checker. In *14th International Conference on Computer-Aided Verification*, 2002.

[Sho84] R. Shostak. Deciding combinations of theories. *Journal of the Association for Computing Machinery*, 31(1):1–12, 1984.

[TZ04] Cesare Tinelli and Calogero Zarba. Combining decision procedures for theories in sorted logics. Technical Report 04-01, Department of Computer Science, The University of Iowa, February 2004.

[WD04] William Weiss and Cherie D'Mello. *Fundamentals of Model Theory*. `http://at.yorku.ca/i/a/a/i/10.htm`, 2004.