Combinations of theories for decidable fragments of first-order logic*

Pascal Fontaine

Université de Nancy, Loria
Nancy, France
Pascal.Fontaine@loria.fr

Abstract. The design of decision procedures for first-order theories and their combinations has been a very active research subject for thirty years; it has gained practical importance through the development of SMT (satisfiability modulo theories) solvers. Most results concentrate on combining decision procedures for data structures such as theories for arrays, bitvectors, fragments of arithmetic, and uninterpreted functions. In particular, the well-known Nelson-Oppen scheme for the combination of decision procedures requires the signatures to be disjoint and each theory to be stably infinite; every satisfiable set of literals in a stably infinite theory has an infinite model.
In this paper we consider some of the best-known decidable fragments of first-order logic with equality, including the Löwenheim class (monadic FOL with equality, but without functions), Bernays-Schönfinkel-Ramsey theories (finite sets of formulas of the form $\exists^*\forall^*\varphi$, where φ is a function-free and quantifier-free FOL formula), and the two-variable fragment of FOL. In general, these are not stably infinite, and the Nelson-Oppen scheme cannot be used to integrate them into SMT solvers. Noticing some elementary results about the cardinalities of the models of these theories, we show that they can nevertheless be combined with almost any other decidable theory.

1 Introduction

Among automated deduction techniques for the verification of computer systems, SMT solvers (Satisfiability Modulo Theories) are nowadays attracting a lot of interest. These solvers are built on top of SAT solvers for propositional logic and include decision procedures for different first-order theories, thus providing more expressive input languages. Usually, SMT solvers implement a combination of a fixed number of theories such as linear arithmetic, uninterpreted symbols, list operators, bit vectors, etc., based on the classical Nelson-Oppen framework [16, 21] for combining decidable theories. This framework covers combinations of disjoint theories provided they are stably infinite: if a set of quantifier-free formulas has a model with respect to a theory, it should also have an infinite model. For instance, a combination of decision procedures for integer linear arithmetic

* This work is partly supported by the ANR project DECERT.

and for the empty theory (equality and uninterpreted symbols) can detect the unsatisfiability of the formula

$$x \leq y \wedge y \leq x + f(x) \wedge P(h(x) - h(y)) \wedge \neg P(0) \wedge f(x) = 0.$$

The Bernays-Schönfinkel-Ramsey (BSR) class [4, 17] is certainly the most well-known decidable class of first-order theories. A BSR theory is a finite set (conjunction) of formulas of the form $\exists^*\forall^*\varphi$, where the first-order formula φ is function-free and quantifier-free. Many verification problems generate formulas in this class (see for instance [11]). The CASC competition [20] for first-order theorem provers has a dedicated division (EPR, Effectively Propositional) for this class. BSR theories are in general not stably infinite. As a trivial example, consider the BSR theory $\forall x \forall y \,.\, x = y$ that only accepts models with singleton domains. The Nelson-Oppen framework does not apply to combinations including BSR theories.

A Löwenheim theory with equality is a finite set of closed formulas in a language containing only unary predicates, and no function except constants. This class is also known as first-order relational monadic logic, and it is decidable. The theory $\forall x \forall y \,.\, x = y$ also belongs to the Löwenheim class, and hence the Nelson-Oppen framework does not apply to this class.

The last decidable class we study in this paper is the class of finitely axiomatized first-order theories built in a language with equality, only two variables, and no functions (except constants). Again $\forall x \forall y \,.\, x = y$ belongs to this class, and the Nelson-Oppen framework is not appropriate.

The objective of the present paper is to lay the ground for incorporating theories from these three well-known classes into SMT solvers.

We are not aware of previous combination results about the full Löwenheim class with equality or the full two-variable fragment with equality. However, it has already been observed [23] that, thanks to its finite model property, a BSR theory can be combined with a theory $\mathcal{T}$ provided the following conditions hold:

- if a set of ground literals L is $\mathcal{T}$-satisfiable, then the minimal cardinality of $\mathcal{T}$-models for L can be computed;
- $\mathcal{T}$ only has finite models.

The second requirement is quite strong. In particular, it is not satisfied by combinations including decidable fragments of arithmetic, which admit only infinite models. For example, the combination scheme of [23] cannot be used to decide the satisfiability of the set of literals such as

$$\{a > 0, a < 2, a + b = 2, b > 0, A(f(a)), \neg C(f(b))\}$$

(where $a, b, f(a), f(b)$ are integers and $+$, $<$, $>$, 0, 2 have their usual meaning over integers) with respect to the BSR theory

$$\mathcal{T} = \{\forall x \, [(A(x) \vee B(x)) \equiv (C(x) \vee D(x))]\}.$$

The classical Nelson-Oppen combination scheme and that of [23] introduce rather strong requirements on the theories in the combination, and these requirements ensure that component theories agree on model cardinalities. For

instance, the stably infinite requirement ensures that both theories will agree on the cardinality $\aleph_0$ for their models. But essentially, the combination process is a matter of matching the interpretation of shared symbols (by exchanging disjunction of equalities), and cardinalities of the models of the theories [12, 23, 9].

We observe in this paper that it is possible to compute all the cardinalities of models admitted by a theory in the BSR, Löwenheim, or two-variable classes with equality. The set of cardinalities accepted by such theories even has a very particular structure. In section 3 we characterize this structure, and show that any decidable theory that verifies this property can be combined with a decidable theory $\mathcal{T}$ provided $\mathcal{T}$ fulfils very liberal constraints. These constraints are trivially met in most practical cases.

For convenience, the results in this paper are presented in an *unsorted* framework, although most SMT-solvers work in a many-sorted logic framework (see for instance [8]). Our results could be transferred to a many-sorted framework, at the expense of heavier notations.

The remainder of this paper is structured as follows: Section 2 introduces basic concepts and notations. Section 3 presents the general scheme for combining (not necessarily stably infinite) theories, and introduces the required notions for the new combination results with the considered first-order decidable classes. Sections 4, 5 and 6 respectively present essential cardinality results about the Löwenheim, BSR, and two-variables classes. We do not claim that the results in those three sections are original. Some of them can be found in classical Model Theory books [5–7]. But some of them are less known. This paper thus presents them together, and relates them to the combination scheme. Section 7 presents a simple example, and Section 8 concludes the paper.

2 Notations

A first-order language is a tuple $\mathcal{L} = \langle \mathcal{V}, \mathcal{F}, \mathcal{P} \rangle$ such that $\mathcal{V}$ is an enumerable set of variables, $\mathcal{F}$ and $\mathcal{P}$ are sets of function and predicate symbols. Every function and predicate symbol is assigned an arity. Nullary predicates are propositions, and nullary functions are constants. Terms and formulas over the language $\mathcal{L}$ are defined in the usual way. A ground term is a term without variables. An atomic formula is either $t = t'$ where t and t' are terms, or a predicate symbol applied to the right number of terms. Formulas are built from atomic formulas, Boolean connectives ($\neg$, $\wedge$, $\vee$, $\Rightarrow$, $\equiv$), and quantifiers ($\forall$, $\exists$). A formula with no free variables is closed. A theory is a set of closed formulas. Two theories are disjoint if no predicate symbol in $\mathcal{P}$ or function symbol in $\mathcal{F}$ appears in both theories. A finite theory or a finitely axiomatized theory is a finite set of formulas.

An interpretation $\mathcal{I}$ for a first-order language provides a domain D, a total function $\mathcal{I}[f] : D^r \to D$ of appropriate arity for every function symbol f, a predicate $\mathcal{I}[p] : D^r \to \{\top, \bot\}$ of appropriate arity for every predicate symbol p, and an element $\mathcal{I}[x] \in D$ for every variable x. By extension, an interpretation defines a value in D for every term, and a truth value for every formula. The

notation $\mathcal{I}_{x_1/d_1,\ldots,x_n/d_n}$ stands for the interpretation that agrees with $\mathcal{I}$, except that it associates the elements d_i to the variables x_i.

A model of a formula (or a theory) is an interpretation in which the formula (resp., every formula in the theory) evaluates to true. A formula or theory is satisfiable if it has a model, and it is unsatisfiable otherwise. A formula G is $\mathcal{T}$-satisfiable if it is satisfiable in the theory $\mathcal{T}$, that is, if $\mathcal{T} \cup \{G\}$ is satisfiable. A $\mathcal{T}$-model of G is a model of $\mathcal{T} \cup \{G\}$. A formula G is $\mathcal{T}$-unsatisfiable if it has no $\mathcal{T}$-models.

The cardinality of an interpretation is the cardinality of its domain. The restriction of a predicate p on domain D to domain $D' \subseteq D$ is the predicate p' with domain D' such that p and p' have the same truth value for all arguments in D'.

A formula is universal if it is of the form $\forall x_1 \ldots \forall x_n.\varphi$ where φ is quantifier-free. A Skolem formula is a formula where all universal quantifiers appear with a positive polarity, and all existential quantifiers appear with a negative polarity. It is always possible to transform a given formula into an equisatisfiable Skolem formula, using Skolemization. We refer to [2] for Skolemization.

3 Combination of theories

Assume we want to study the satisfiability of the set of literals

$$L = \{a \leq b, b \leq a + f(a), P(h(a) - h(b)), \neg P(0), f(a) = 0\}$$

in the combination of the integer linear arithmetic theory $\mathcal{T}_1$ and the empty theory (i.e. the theory of uninterpreted symbols) $\mathcal{T}_2$. First, a *separation* is built by introducing fresh uninterpreted constants[1], to produce the equisatisfiable problem

$$L_1 = \big\{a \leq b,\, b \leq a + v_1,\, v_1 = 0,\, v_2 = v_3 - v_4,\, v_5 = 0\big\}$$
$$L_2 = \big\{P(v_2),\, \neg P(v_5),\, v_1 = f(a),\, v_3 = h(a),\, v_4 = h(b)\big\}.$$

The set L_1 only contains arithmetic symbols and uninterpreted constants. The symbols in L_2 are all uninterpreted. The only shared symbols are the uninterpreted constants in the set $S = \{a, b, v_1, v_2, v_3, v_4, v_5\}$. Notice that although L is unsatisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$, L_1 is $\mathcal{T}_1$-satisfiable, and L_2 is $\mathcal{T}_2$-satisfiable; it is not sufficient for the decision procedures for $\mathcal{T}_1$ and $\mathcal{T}_2$ to only examine the satisfiability of their part of the separation. Indeed, the decision procedures also have to "agree on the common part". This can be captured using the notion of arrangement:

Definition 1. *An arrangement $\mathcal{A}$ for a set of constant symbols S is a maximal satisfiable set of equalities and inequalities $a = b$ or $a \neq b$, with $a, b \in S$.*

[1] Traditionally combination schemes use variables for this role. Since variables will be used in quantifiers in the following sections, for consistency and clarity we will rather use uninterpreted constants here.

The following theorem (other formulations can be found in [22, 23, 12]) then states the completeness of the combination of decision procedures:

Theorem 1. *Assume $\mathcal{T}_1$ and $\mathcal{T}_2$ are theories over the disjoint languages $\mathcal{L}_1$ and $\mathcal{L}_2$, and L_i ($i = 1, 2$) is a set of literals in $\mathcal{L}_i$ augmented by a finite set of fresh constant symbols S. Then $L_1 \cup L_2$ is $\mathcal{T}_1 \cup \mathcal{T}_2$-satisfiable if and only if there exist an arrangement $\mathcal{A}$ of S, a cardinality k, and a $\mathcal{T}_i$-model $\mathcal{M}_i$ of $\mathcal{A} \cup L_i$ with cardinality k for $i = 1, 2$.*

Proof. Assume $\mathcal{I}$ is an interpretation on domain D for a language $\mathcal{L}$, and $\mathcal{L}'$ is a sub-language of $\mathcal{L}$, i.e. the set of variable, function, and predicate symbols in $\mathcal{L}'$ are subsets of their counterpart in $\mathcal{L}$. We say that the interpretation $\mathcal{I}'$ on domain D for language $\mathcal{L}'$ is the restriction of $\mathcal{I}$ if $\mathcal{I}'$ and $\mathcal{I}$ give the same interpretation for the symbols in $\mathcal{L}'$.

The condition is necessary. Assume $\mathcal{M}$ is a $\mathcal{T}_1 \cup \mathcal{T}_2$-model for $L_1 \cup L_2$. $\mathcal{M}$ perfectly defines an arrangement $\mathcal{A}$ of S: indeed $a = b \in \mathcal{A}$ with $a, b \in S$ iff $a = b$ is true according to $\mathcal{M}$. The restriction of $\mathcal{M}$ to $\mathcal{L}_i$ augmented with the constant symbols S is a $\mathcal{T}_i$-model for $\mathcal{A} \cup L_i$, $i = 1, 2$.

The condition is sufficient. Assume that $\mathcal{A}$ is an arrangement for S, $\mathcal{M}_1$ on domain D_1 is a $\mathcal{T}_1$-model for $\mathcal{A} \cup L_1$, $\mathcal{M}_2$ on domain D_2 is a $\mathcal{T}_2$-model for $\mathcal{A} \cup L_2$, and $|D_1| = |D_2|$. Since both $\mathcal{M}_1$ and $\mathcal{M}_2$ are models of $\mathcal{A}$, there exist two interpretations $\mathcal{M}'_1$ and $\mathcal{M}'_2$ on the same domain that are respectively isomorphic to $\mathcal{M}_1$ and $\mathcal{M}_2$ and such that $\mathcal{M}'_1[a] = \mathcal{M}'_2[a]$ for every $a \in S$. It is then possible to build an interpretation $\mathcal{M}$ such that its restriction to the language $\mathcal{L}_i$ augmented with S is $\mathcal{M}'_i$, $i = 1, 2$. $\mathcal{M}$ is a $\mathcal{T}_1 \cup \mathcal{T}_2$-model of $L_1 \cup L_2$. □

Checking the existence of a model is the task of the decision procedures for the decidable theories in the combination. The previous theorem however also imposes a restriction on cardinalities: the two decision procedures should exhibit a model with the same cardinality. A theory $\mathcal{T}$ is said to be stably infinite when every $\mathcal{T}$-satisfiable set of literals has a model with cardinality $\aleph_0$. Combining only stably infinite theories is a radical solution to the cardinality requirement in the previous theorem; k can always be $\aleph_0$. Since the empty theory and the theory of integer linear arithmetic are both stably infinite, the set of literals L in our example is $\mathcal{T}_1 \cup \mathcal{T}_2$-satisfiable if and only if there exists an arrangement $\mathcal{A}$ of the seven variables in S such that $\mathcal{A} \cup L_i$ is $\mathcal{T}_i$-satisfiable for $i = 1$ and $i = 2$. No such arrangements exist, and indeed, L is $\mathcal{T}_1 \cup \mathcal{T}_2$-unsatisfiable.

The first-order decidable classes considered in this paper contain theories that are not stably infinite. For instance the formula $\forall x\,(x = a \vee x = b)$ belongs to the BSR, Löwenheim and two variable classes, and it only accepts models with at most two elements. A combination scheme to handle such theories requires to carefully examine cardinalities. The notion of spectrum is helpful for this task:

Definition 2. *The* spectrum *of a theory $\mathcal{T}$ is the set of cardinalities k such that $\mathcal{T}$ is satisfiable in a model of cardinality k.*[2]

[2] The spectrum of a theory is usually defined as the set of the *finite* cardinalities of its models. We here slightly extend the definition for convenience.

Using this definition and Theorem 1, a combination scheme for disjoint theories (not necessarily stably infinite) can thus be easily expressed:

Corollary 1. *Given two theories $\mathcal{T}_1$ and $\mathcal{T}_2$ over the disjoint languages $\mathcal{L}_1$ and $\mathcal{L}_2$, the $\mathcal{T}_1 \cup \mathcal{T}_2$-satisfiability problem for sets of literals (written in the union of the languages $\mathcal{L}_1$ and $\mathcal{L}_2$) is decidable if, for any sets of literals L_1 and L_2 (respectively written in the languages $\mathcal{L}_1$ and $\mathcal{L}_2$ augmented with a finite set of fresh uninterpreted constants) it is possible to compute if the intersection of the spectrums for $\mathcal{T}_1 \cup L_1$ and for $\mathcal{T}_2 \cup L_2$ is non-empty.*

In the case of stably infinite decidable theories, it is guaranteed that, if $\mathcal{T}_1 \cup L_1$ and $\mathcal{T}_2 \cup L_2$ are satisfiable, both spectrums contain cardinality $\aleph_0$, and so their intersection is trivially non-empty.

To characterize the spectrum of the decidable classes considered in this paper, we introduce the following property:

Definition 3. *A theory $\mathcal{T}$ is* gentle *if, for every set L of literals in the language of $\mathcal{T}$ (augmented by a finite number of fresh constants), the spectrum of $\mathcal{T} \cup L$ can be computed and is either*

- *a finite set of finite cardinalities*
- *the union of a finite set of finite cardinalities and all the (finite and infinite) cardinalities greater than a computable finite cardinality; it is thus co-finite.*

A gentle theory is decidable. In the following sections, we show that the BSR theories, the Löwenheim theories, and finite theories with only two variables are gentle. The empty theory, as a special case of a BSR theory, is gentle. Shiny theories in general (see [23]) are gentle. We also have the following result:

Theorem 2. *The union of disjoint gentle theories is a gentle theory.*

Proof. The case for the union of any number of disjoint gentle theories can be proved by induction, and using the case for two gentle theories.

The intersection of two spectrums of gentle theories is also either a finite set of finite cardinalities, or the union of a finite set of finite cardinalities and all the (finite and infinite) cardinalities greater than a (computable) finite cardinality. The case for two gentle theories is thus a direct consequence of Theorem 1. □

We point out that a theory $\mathcal{T}$ taking part in a combination of theories has some interesting property about its spectrum. Since the $\mathcal{T}$-satisfiability problem for sets of literals (written in the language of the theory plus fresh constants) is decidable, it is also possible to assess for any set of literals L if $\mathcal{T} \cup L$ has a model of cardinality greater than a given number k. Indeed it suffices to introduce k new constants $a_1, \ldots, a_k$ and check the satisfiability of $\mathcal{T} \cup L \cup \{a_i \neq a_j \mid i \neq j, \quad i, j = 1, \ldots, k\}$. Also notice that it is always possible to decide if a finite first-order theory admits a model of a given finite cardinality. Indeed there are only a finite number of interpretations for a finite language, and it takes a finite time to check if a given finite interpretation is a model of the finite theory.

Some widely used theories are not gentle, but in practical cases they can be combined with gentle theories:

Theorem 3. *Given a gentle theory $\mathcal{T}$ and another disjoint theory $\mathcal{T}'$, the $\mathcal{T} \cup \mathcal{T}'$-satisfiability problem for sets of literals written in the union of their language is decidable if one of the following cases holds:*

- *$\mathcal{T}'$ is gentle;*
- *$\mathcal{T}'$ is a decidable finitely axiomatized first-order theory;*
- *$\mathcal{T}'$ is a decidable theory that only admits a fixed finite (possibly empty) known set of finite cardinalities for its models, and possibly infinite models.*

Proof. Assume $L \cup L'$ is the separation to check for $\mathcal{T} \cup \mathcal{T}'$-satisfiability. If an arrangement $\mathcal{A}$ is such that $\mathcal{A} \cup L$ is $\mathcal{T}$-satisfiable, and $\mathcal{A} \cup L'$ is $\mathcal{T}'$-satisfiable, then it is possible to compute the spectrum $\mathcal{S}$ of $\mathcal{T} \cup \mathcal{A} \cup L$. Either $\mathcal{S}$ is a finite set of finite cardinalities, or it is a union of a finite set of finite cardinalities and the set of all cardinalities greater than a number k.

If $\mathcal{T}'$ is also gentle, it is possible to compute the spectrum of $\mathcal{T}' \cup \mathcal{A} \cup L'$, and the intersection of the two spectrums can easily be computed.

If $\mathcal{T}'$ is a decidable finite first-order theory, it is possible to check if $\mathcal{T}' \cup \mathcal{A} \cup L'$ admits a cardinality in the finite part of $\mathcal{S}$, and, if $\mathcal{S}$ is infinite, it is possible to check if $\mathcal{T}' \cup \mathcal{A} \cup L'$ admits a cardinality greater than k.

If $\mathcal{T}'$ is a decidable theory that only admits a fixed finite known set of cardinalities for its models, it suffices to check if one of these cardinalities is in the spectrum $\mathcal{S}$. The considered theories are first-order, and the Löwenheim-Skolem theorem states that, if a theory has an infinite model, it has models for every infinite cardinality. Infinite cardinalities can thus be understood as one cardinality. □

For instance, the real or integer linear arithmetic theories (or combinations involving real or integer linear arithmetic) fall into the last case, and the usual theories for arrays fall into the second one.

4 The Löwenheim class with equality

A Löwenheim theory is a finite set of closed formulas in a language containing only unary predicates, and no functions except constants. This class is also known as first-order relational monadic logic. Usually one distinguishes the Löwenheim class with and without equality. The Löwenheim class has the finite model property (and is thus decidable) even with equality. Full monadic logic *without equality*, i.e. the class of finite theories over a language containing symbols (predicates and functions) of arity at most 1, also has the finite model property. Considering monadic logic with equality, the class of finite theories over a language containing only unary predicates and just two unary functions is already undecidable. With only one unary function however the class remains decidable, but does not have the finite model property anymore. Since the spectrum for this last class is significantly more complicated [14] than for the Löwenheim class we will here only concentrate on the Löwenheim class with equality (only classes with equality are relevant in our context). More can be found about monadic first-order logic in [5, 6]. In particular, the following Theorem can be found in [6]:

Theorem 4. *Assume $\mathcal{T}$ is a Löwenheim theory with equality with n distinct unary predicates. Let q be the number of constants plus the maximum number of nested quantifiers in $\mathcal{T}$. If $\mathcal{T}$ has a model of some cardinality $\geq q\,2^n$, then $\mathcal{T}$ has models of every cardinality $\geq q\,2^n$.*

Proof. For simplicity, assume $\mathcal{T}$ is constant-free and is a single formula. Because $\mathcal{T}$ is finite, it is always possible to get back to such a case by taking the conjunction of all formulas in $\mathcal{T}$, and then quantify existentially over all constants in the formula.

Let $p_1, \ldots, p_n$ be the unary predicates used in $\mathcal{T}$. Given an interpretation $\mathcal{I}$ on domain D for $\mathcal{T}$, every element $d \in D$ has a color $c(d) = c_1 \ldots c_n \in \{\top, \bot\}^n$ where $c_i = \mathcal{I}[p_i](d)$. We denote by $D_c \subseteq D$ the set of elements with color c.

Two interpretations $\mathcal{I}$ (on domain D) and $\mathcal{I}'$ (on domain D') for a formula ψ are *similar* if

- either $D_c = D'_c$ or $|D_c \cap D'_c| \geq q$ for every color $c \in \{\top, \bot\}^n$;
- $D_c \cap D'_{c'} = \emptyset$ for any two distinct colors $c, c' \in \{\top, \bot\}^n$;
- $\mathcal{I}[x] = \mathcal{I}'[x]$ for every variable free in ψ.

We first prove that, given a formula ψ, two similar interpretations for ψ give the same truth value to ψ and to every sub-formula of ψ.

This is proved by induction on the structure of the (sub-)formula ψ. It is obvious if ψ is atomic, since similar interpretations assign the same value to variables, and since ψ is variable-free. If ψ is $\neg\varphi_1$, $\varphi_1 \vee \varphi_2$, $\varphi_1 \wedge \varphi_2$ or $\varphi_1 \Rightarrow \varphi_2$, the result holds if it also holds for φ_1 and φ_2.

Assume $\mathcal{I}$ makes true the formula $\psi = \exists x\, \varphi(x)$. Then there exists some $d \in D$ such that $\mathcal{I}_{x/d}$ is a model of $\varphi(x)$. If $d \in D'$, then $\mathcal{I}'_{x/d}$ is similar to $\mathcal{I}_{x/d}$ and, by the induction hypothesis, it is a model of $\varphi(x)$; $\mathcal{I}'$ is thus a model of ψ. If $d \notin D'$, it means that $D_{c(d)} \cap D'_{c(d)} \geq q$. Furthermore, since the whole formula contains at most q nested quantifiers, $\varphi(x)$ contains at most $q-1$ free variables. Let $x_1, \ldots, x_m$ be those variables. There exists some $d' \in D_{c(d)} \cap D'_{c(d)}$ such that $d' \neq \mathcal{I}[x_i]$ for every $i \in \{1, \ldots, m\}$. By structural induction, it is easy to show that $\mathcal{I}_{x/d}$ and $\mathcal{I}_{x/d'}$ give the same truth value to $\varphi(x)$. Furthermore $\mathcal{I}_{x/d'}$ and $\mathcal{I}'_{x/d'}$ are similar. $\mathcal{I}'$ is thus a model of ψ. To summarize, if $\mathcal{I}$ is a model of ψ, $\mathcal{I}'$ is also a model of ψ. By symmetry, if $\mathcal{I}'$ is a model of ψ, $\mathcal{I}$ is also a model of ψ. Thus, if $\psi = \exists x\, \varphi(x)$, the results hold if it also holds for $\varphi(x)$. The proof for formulas of the form $\forall x\, \varphi(x)$ is dual.

If $\mathcal{M}$ on domain D is a model for $\mathcal{T}$ with cardinality $\geq q\,2^n$, then there exists a color c such that $|D_c| \geq q$. For any cardinality $k \geq q\,2^n$ one can build a model $\mathcal{M}'$ of cardinality k for $\mathcal{T}$, similar to $\mathcal{M}$. □

Corollary 2. *The Löwenheim class has the finite model property.*

Proof. Assume $\mathcal{T}$ is a Löwenheim theory, with n distinct unary predicates. Let q be the maximum number of nested quantifiers in $\mathcal{T}$. Let $\mathcal{I}$ be a model of $\mathcal{T}$. According to Theorem 4, if $\mathcal{I}$ has an infinite cardinality ($\geq q\,2^n$), $\mathcal{T}$ also has a finite model (e.g. of cardinality $q\,2^n$). □

Corollary 3. *The satisfiability problem for the Löwenheim class is decidable.*

Proof. It is well-known that any class of finite first-order theories that has the finite model property is also decidable. The decidability of the Löwenheim class can also be easily proved directly. Assume $\mathcal{T}$ is a Löwenheim theory, with n distinct unary predicates. Let q be the maximum number of nested quantifiers in $\mathcal{T}$. There exist only a finite number of interpretations of a finite theory for a given cardinality. It is thus decidable to check if $\mathcal{T}$ has a model of cardinality $q\,2^n$. If such a model exists $\mathcal{T}$ is satisfiable. If no such models exist, Theorem 4 states that $\mathcal{T}$ has no models of cardinality $\geq q\,2^n$. It remains to decide if $\mathcal{T}$ has a model of cardinality $< q\,2^n$, i.e. it remains to examine a finite number of interpretations. □

Corollary 4. *The spectrum of a Löwenheim theory can be computed and expressed either as a finite set of naturals, or as the union of a finite set of naturals with the set of all the (finite or infinite) cardinalities greater than a natural. The Löwenheim theories are gentle.*

5 The Bernays-Schönfinkel-Ramsey class

A Bernays-Schönfinkel-Ramsey theory (BSR) is a finite set of formulas of the form $\exists^*\forall^*\varphi$, where φ is a first-order formula which is function-free (but constants are allowed) and quantifier-free. Bernays and Schönfinkel first proved the decidability of this class without equality; Ramsey later proved that it remains decidable with equality. The results about the spectrum of BSR theories are less known, but were also originally found by Ramsey.

For simplicity, we will assume that existential quantifiers are Skolemized. In the following, a BSR theory is thus a finite closed set of universal function-free first-order formulas.

Theorem 5. *Let $\mathcal{T}$ be a BSR theory, and let k_c be the number of constants in $\mathcal{T}$, or $k_c = 1$ if $\mathcal{T}$ is constant-free. If $\mathcal{T}$ has a model with cardinality $k \geq k_c$, then $\mathcal{T}$ has a model for every cardinality i, with $k \geq i \geq k_c$.*

Proof. Given a model $\mathcal{M}$ for a BSR theory $\mathcal{T}$ with domain D, then any interpretation $\mathcal{M}'$ such that

- the domain of $\mathcal{M}'$ is a non-empty set $D' \subseteq D$ such that $\mathcal{M}'[a] = \mathcal{M}[a] \in D'$ for every constant a in $\mathcal{T}$, and
- for every predicate p, $\mathcal{M}'[p]$ is the restriction of $\mathcal{M}[p]$ to the domain D'

is also a model of $\mathcal{T}$. Intuitively, this states that the elements in the domain that are not assigned to ground terms (i.e. the constants) can be eliminated in a model of a BSR theory. Since $\mathcal{M}$ is a model of $\mathcal{T}$, for each closed formula $\forall x_1 \ldots x_n.\varphi$ in $\mathcal{T}$ (where φ is function-free and quantifier-free), and for all $d_1, \ldots, d_n \in D' \subseteq D$, $\mathcal{M}_{x_1/d_1,\ldots,x_n/d_n}$ is a model of φ. This also means that, for all $d_1, \ldots, d_n \in D'$, $\mathcal{M}'_{x_1/d_1,\ldots,x_n/d_n}$ is a model of φ, and finally that $\mathcal{M}'$ is a model of $\forall x_1 \ldots x_n.\varphi$. □

Theorem 6. *There exists a computable function f such that, for any BSR theory $\mathcal{T}$, if $\mathcal{T}$ has a model of some cardinality $\geq f(\mathcal{T})$, then it has a model for every cardinality $\geq f(\mathcal{T})$.*

Proof. The proof is quite long and requires a non trivial theorem on hypergraph coloring. A partial proof can be found in [6], and a full self-contained proof can be found in the full version of the paper [10]. □

The proofs of the following corollaries are similar to the corresponding proofs for the Löwenheim class.

Corollary 5. *The BSR class has the finite model property.*

Corollary 6. *The satisfiability problem for the BSR class is decidable.*

Corollary 7. *The spectrum of a BSR theory can be computed and expressed either as a finite set of naturals, or as the union of a finite set of naturals with the set of all the (finite or infinite) cardinalities greater than a natural. BSR theories are gentle.*

6 First-order logic with two variables

Following [7], we will denote by FO^2 the class of finite theories built over a language with only two variables, and no functions (except constants). The satisfiability problem for FO^2 is known to be decidable with and without equality (see for instance [5, 7, 13]). Again, we will only concentrate here on the language with equality. This class has the finite model property, and also has very nice properties concerning the cardinalities of its models.

The Scott class is a subset of FO^2: it is the class of finite theories over a language with only two variables, and no functions (except constants) such that every formula in the theory is of the form $\forall x \forall y\, \varphi(x, y)$ or $\forall x \exists y\, \varphi(x, y)$ where $\varphi(x, y)$ is quantifier-free. The satisfiability problem for FO^2 (with equality) is traditionally translated into the satisfiability problem for the Scott class, using the following theorem (see [5, 7] for equivalent theorems):

Theorem 7. *There exists an algorithm that, for each finite theory $\mathcal{T}$ of FO^2, constructs a theory $\mathcal{T}'$ in the Scott class such that $\mathcal{T}$ has a model of a given cardinality if and only if $\mathcal{T}'$ has a model of the same cardinality. The size of $\mathcal{T}'$ is linear with respect to the size of $\mathcal{T}$.*

Proof. First notice that formula $\forall x\, (R(x) \equiv Qy\, \varphi(x, y))$ where Q is either $\exists$ or $\forall$ can be rewritten as a set of formulas in the required form:

- $\forall x\, (R(x) \equiv \forall y\, \varphi(x, y)) \longleftrightarrow \forall x \forall y\, (R(x) \Rightarrow \varphi(x, y)) \wedge \forall x \exists y\, (\varphi(x, y) \Rightarrow R(x))$
- $\forall x\, (R(x) \equiv \exists y\, \varphi(x, y)) \longleftrightarrow \forall x \exists y\, (R(x) \Rightarrow \varphi(x, y)) \wedge \forall x \forall y\, (\varphi(x, y) \Rightarrow R(x))$

The theory $\mathcal{T}$ can thus be rewritten into a suitable theory $\mathcal{T}'$ by iteratively applying the following step until no more formulas of unsuitable form exist in the theory:

– select a formula ψ in the theory that does not have the required form;
– choose a sub-formula of form $Qy\,\varphi(x,y)$ of ψ where Q is $\exists$ or $\forall$ and $\varphi(x,y)$ is quantifier-free;
– take a new unary predicate R not used in the theory;
– define the formula ψ' as ψ where $Qy\,\varphi(x,y)$ has been substituted by $R(x)$;
– remove ψ from the theory, and add ψ', and the formulas in the required form for $\forall x\,(R(x) \equiv Qy\,\varphi(x,y))$.

□

The following theorem is left as an exercice in [7]. For completeness we here give the full proof.

Theorem 8. *There exists a computable function f such that, for any Scott theory $\mathcal{T}$, if $\mathcal{T}$ has a model of some cardinality $\geq f(\mathcal{T})$, then $\mathcal{T}$ has models for every cardinality $\geq f(\mathcal{T})$.*

Proof. We first assume that every formula ψ_i in $\mathcal{T}$ $(i = 1, \ldots, m)$ of the form $\forall x \exists y\, \varphi(x,y)$ is such that every model of $\varphi(x,y)$ is a model of $x \neq y$. This assumption is acceptable if $f(\mathcal{T}) \geq 2$ for all Scott theories $\mathcal{T}$ since for all models with at least two elements $\forall x \exists y\, \varphi(x,y)$ is equivalent to $\forall x \exists y\,.\, x \neq y \wedge (\varphi(x,y) \vee \varphi(x,x))$.

For the rest of the proof, we assume that the Scott theory $\mathcal{T}$ has a model $\mathcal{M}$ on domain D. We define the sets $A = \{\mathcal{M}[a] : a \text{ is a constant in } \mathcal{T}\}$ and $B = D \setminus A$. We establish that if B is larger than a computable cardinality $\geq f(\mathcal{T})$, one can build a model for every cardinality $\geq f(\mathcal{T})$.

Given a first-order language $\mathcal{L}$, a k-table[3] $T[x_1, \ldots, x_k]$ over the variables $x_1, \ldots, x_k$ is a maximal satisfiable set of atomic formulas and negation of atomic formulas using only variables $x_1, \ldots, x_k$. Given an interpretation $\mathcal{I}$ on domain D and k elements $d_1, \ldots, d_k$ of D, the k-table of $d_1, \ldots, d_k$ (denoted $T_{\mathcal{I}}[d_1, \ldots, d_k]$) is the unique k-table $T[x_1, \ldots, x_k]$ such that the interpretation $\mathcal{I}_{x_1/d_1,\ldots,x_k/d_k}$ is a model of $T[x_1, \ldots, x_k]$. Notice that there are only a finite number of k-tables, for a finite language with no functions except constants. In particular if A is the set of constants, a 1-table is determined by at most $b = \sum_p (|A|+1)^{\text{arity}(p)}$ Boolean values, where the sum ranges over all predicates in the language. Indeed, given a predicate p of arity r, there are at most $(|A|+1)^r$ terms that can be built with p and $A \cup \{x\}$. Thus the number of different 1-tables is bounded by $C = 2^b$.

For every formula $\psi_i = \forall x \exists y\, \varphi_i(x,y)$ in $\mathcal{T}$ $(i = 1, \ldots, m)$, there exists a total function g_i on domain D ranging on D such that $\mathcal{M}[\varphi_i](d, g_i(d))$ is true for every $d \in D$. The set K (commonly referred as the set of kings) is defined as the union of A and of the possibly empty set of all elements of $d \in D$ such that the 1-table of d is unique, i.e. $T_{\mathcal{M}}[d'] \neq T_{\mathcal{M}}[d]$ for every $d' \in D$ such that $d' \neq d$. The set C (commonly referred as the court) is the possibly empty set $C = K \cup \{g_i(d) \mid d \in K, i = 1, \ldots, m\}$. The set S is defined as $T_{\mathcal{M}}[D] \setminus T_{\mathcal{M}}[C]$ where $T_{\mathcal{M}}[D]$ is the set of all 1-tables of elements in D (and similarly for $T_{\mathcal{M}}[C]$). We choose a function h on domain S that ranges on D such that $T_{\mathcal{M}}[h(t)] = t$.

[3] We here adopt the notation of [5]. The same notion is also called (atomic) k-type, for instance in [13].

The set D' is defined as $C \cup (S \times \{1, \ldots, m\} \times \{0, 1, 2\})$. A model $\mathcal{M}'$ on D' for $\mathcal{T}$ is defined such that:

- $T_{\mathcal{M}'}[d_1, \ldots, d_k] = T_{\mathcal{M}}[d_1, \ldots, d_k]$ for $d_1, \ldots d_k \in C$, $k \in \mathbb{N}$;
- $T_{\mathcal{M}'}[(t, i, j)]$ is t, for every $(t, i, j) \in D' \setminus C$;
- if $g_i(h(t)) \in K$ then $T_{\mathcal{M}'}[(t, i, j), g_i(h(t))] = T_{\mathcal{M}}[h(t), g_i(h(t))]$;
- if $g_i(h(t)) \notin K$ then $T_{\mathcal{M}'}[(t, i, j), (T_{\mathcal{M}}(g_i(h(t))), i, (j + 1) \mod 3)]$ is equal to $T_{\mathcal{M}}[h(t), g_i(h(t)]$
- if not yet defined $T_{\mathcal{M}'}[d'_1, d'_2]$ is $T_{\mathcal{M}}[d_1, d_2]$, where d_i is chosen such that $T_{\mathcal{M}}[d_i] = T_{\mathcal{M}'}[d_i]$ $(i = 1, 2)$.

The undefined interpretations are not relevant for interpreting the theory and can be arbitrarily defined. The previous assignments are non-conflicting, i.e. 2-tables are never defined twice inconsistently.

Assume $\forall x \forall y\, \varphi(x, y)$ belongs to $\mathcal{T}$. Then $\mathcal{M}'_{x/d'_1, y/d'_2} \varphi(x, y) = \top$ since there exists d_1 and d_2 such that $T_{\mathcal{M}}[d_1, d_2] = T_{\mathcal{M}'}[d'_1, d'_2]$. It remains to prove that $\mathcal{M}'$ is a model of every formula $\forall x \exists y\, \varphi_i(x, y)$ in $\mathcal{T}$, or equivalently, that for every $d \in D'$, $\mathcal{M}'_{x/d}$ is a model of $\exists y\, \varphi_i(x, y)$:

- if $d \in K$, $g_i(d) \in C \subseteq D'$, and $\mathcal{M}'_{x/d, y/g_i(d)}$ is a model of $\varphi_i(x, y)$;
- if $d \in C \setminus K$, if $g_i(d) \in C$ then $\mathcal{M}'_{x/d, y/g_i(d)}$ is a model of $\varphi_i(x, y)$;
- if $d \in C \setminus K$, if $g_i(d) \notin C$ then $T_{\mathcal{M}'}[d, (T_{\mathcal{M}}(g_i(d)), i, 0)] = T_{\mathcal{M}}[d, g_i(d)]$, and thus $\mathcal{M}'_{x/d, y/(T_{\mathcal{M}}(g_i(d)), i, 0)}$ is a model of $\varphi_i(x, y)$;
- if $d = (t, i, j) \in D' \setminus C$, if $g_i(h(t)) \in K$ then $T_{\mathcal{M}'}[(t, i, j), g_i(h(t))] = T_{\mathcal{M}}[h(t), g_i(h(t))]$, and thus $\mathcal{M}'_{x/d, y/g_i(h(t))}$ is a model of $\varphi_i(x, y)$;
- if $d = (t, i, j) \in D' \setminus C$, if $g_i(h(t)) \notin K$ then $T_{\mathcal{M}'}[(t, i, j), (T_{\mathcal{M}}(g_i(h(t))), i, (j + 1) \mod 3)] = T_{\mathcal{M}}[h(t), g_i(h(t)]$, and thus $\mathcal{M}'_{x/d, y/(T_{\mathcal{M}}(g_i(h(t))), i, (j+1) \mod 3)}$ is a model of $\varphi_i(x, y)$.

Finally notice that $D \setminus K$ is necessarily non-empty if $|D| \geq 2^b + 1 + |A|$. In the process of building $\mathcal{M}'$, any element $(t, i, 0)$ may be duplicated, thus creating models of arbitrary size $\geq 3m\, 2^b + (m+1)|A|$ where m is the number of formulas of the form $\forall x \exists y\, \varphi(x, y)$ in $\mathcal{T}$. □

Corollary 8. *There exists a computable function f such that, for any finite theory $\mathcal{T}$ of* FO^2*, if $\mathcal{T}$ has a model of some cardinality $\geq f(\mathcal{T})$, then $\mathcal{T}$ has models for every cardinality $\geq f(\mathcal{T})$.*

Corollary 9. *The class of finite theories of* FO^2 *has the finite model property.*

Corollary 10. *The satisfiability problem for finite theories of* FO^2 *is decidable.*

Corollary 11. *The spectrum of a finite theory of* FO^2 *can be computed and expressed as a finite set of naturals, or as the union of a finite set of naturals with the set of all the (finite or infinite) cardinalities greater than a natural. The finite theories of* FO^2 *are gentle.*

7 An example

Assume that one wants to study the satisfiability of the simple example given in the introduction:

$$\{a > 0, a < 2, a + b = 2, b > 0, A(f(a)), \neg C(f(b))\}$$

in the combination of the theories

$$\mathcal{T}_1 = \forall x \left[(A(x) \vee B(x)) \equiv (C(x) \vee D(x))\right]$$

and $\mathcal{T}_2$, where $\mathcal{T}_2$ is itself the combination of the theory of uninterpreted functions and linear arithmetic over the integers. The theory $\mathcal{T}_2$ is decidable, and a decision procedure can be built using the standard Nelson-Oppen scheme since both components are stably infinite. The domain of the models of $\mathcal{T}_2$ is always the set of integers, thus all models have cardinality $\aleph_0$. The theory $\mathcal{T}_1$ belongs to the BSR, Löwenheim, and two-variables classes and is thus gentle.[4] The third case of Theorem 3 is fulfilled. First, a separation is built, to produce the equisatisfiable problem $L_1 \cup L_2$ with

$$L_1 = \{A(t), \neg C(u)\}$$
$$L_2 = \{a > 0, a < 2, a + b = 2, b > 0, t = f(a), u = f(b)\}.$$

The set $L_1 \cup L_2$ is $\mathcal{T}_1 \cup \mathcal{T}_2$-satisfiable if and only if there exists a $\mathcal{T}_1$-model $\mathcal{M}_1$ for L_1 and a $\mathcal{T}_2$-model $\mathcal{M}_2$ for L_2, such that $\mathcal{M}_1$ and $\mathcal{M}_2$ agree on which shared constant symbols (i.e. t and u) are equal, and agree on cardinalities (Theorem 1). The first requirement is fulfilled by checking every arrangement of the variables (here: $t = u$ or $t \neq u$): $\{t \neq u\} \cup L_2$ is $\mathcal{T}_2$-unsatisfiable, but $\{t = u\} \cup L_1$ and $\{t = u\} \cup L_2$ are both satisfiable in their respective theory. It remains to check if it is possible for both models to agree on cardinalities. The theory of integer linear arithmetic only accepts models of cardinality $\aleph_0$, therefore $L_1 \cup L_2$ is $\mathcal{T}_1 \cup \mathcal{T}_2$-satisfiable if and only if $\mathcal{T}_1 \cup \{t = u\} \cup L_1$ has a model of cardinality $\aleph_0$.

The theory $\mathcal{T}_1 \cup \{t = u\} \cup L_1$ uses only one quantified variable, four predicate symbols (A, B, C, D), and two constants (t, u). Using for instance the fact that this theory is a Löwenheim theory, one can use Theorem 4 to check if it has an infinite model. The theory contains two constants, at most one "nested" quantifier, and four unary predicates. If there is a model with cardinality 3×2^4, then there is an infinite model. It can easily be showed that $\mathcal{T}_1 \cup \{t = u\} \cup L_1$ indeed accepts such a model with cardinality 48. Similar bounds exist for BSR and two-variable theories, but unfortunately they are also large compared to the size of this toy example.

There exists another criteria to check if a BSR theory has an infinite model. Indeed, a BSR theory with n variables has an infinite model if and only if it has a n-repetitive model (see the full version of the paper [10]). Checking if

[4] $\mathcal{T}_1$ is also stably infinite, but we ignore this fact to illustrate the generic approach.

$\mathcal{T}_1 \cup \{t = u\} \cup L_1$ has a 1-repetitive model simply amounts to check if $\mathcal{T}_1 \cup \{t = u\} \cup L_1 \cup \{v \neq t, v \neq u\}$ is satisfiable.

As a final remark, notice that the example used in this section encodes the set of formulas

$$\{a > 0, a < 2, a + b = 2, b > 0, f(a) \in A, f(b) \notin C, A \cup B = C \cup D\}$$

in a language that combines integer linear arithmetic, uninterpreted function symbols, and elementary set-theoretic operations. One motivation for the work reported in this paper is indeed to augment the languages accepted by SMT solvers with certain operators on sets or relations, which can conveniently be represented by BSR theories over their characteristic predicates.

8 Conclusion

In this paper we observed that one can express completely the spectrum, i.e. the set of the cardinalities of the models, for Löwenheim theories, BSR theories, and finite theories in the two-variables fragment. We characterise those theories as *gentle*. Gentle theories can be combined with almost any decidable theory, including gentle theories, integer or real linear arithmetic, finite first-order theories, and some combinations of these.

It remains to develop algorithmic techniques to make this combination work in practice. The results presented here are prohibitively expensive, the finite cardinalities that guarantee the existence of an infinite model grow very rapidly with the size of the theories. In that sense, the combination scheme presented in this paper is really at the frontiers of combining decision procedures. It is certainly not practical to first extract all cardinalities of the gentle theories in the combination, just like, in the Nelson-Oppen combination scheme, it is not practical to check every arrangement one by one. Rather than guessing arrangements, SMT solvers use, among other techniques, equality propagation. Equality propagation can thus be seen as the negotiation of an arrangement. A practical way to agree on cardinality could also rely on negotiation. This negotiation would often be trivial, for instance if one theory puts very strong constraints on cardinalities, or if most theories are on the contrary very permissive. Another approach to handle the same classes of theories can be found in [24]: it consists in reducing each part of the separation to a formula in a common decidable language including Presbruger arithmetics; this approach has the drawback of being much more complex, but as a counterpart the language handled is in some aspects much more expressive.

Usually, SMT solvers implement a combination of a fixed set of theories, which are known *a priori*, and are also known to have the right properties according to cardinalities (typically, being stably infinite). Here, we show that every theory in the major well-known first-order decidable classes can be integrated in a combination. Since it can be shown that assertions over sets or relations and elementary set-theoretic operations like $\cup$, $\cap$, etc. just introduce one more BSR theory in the combination, the problem remains decidable even if this theory is

not fixed a priori. We mainly target formal methods based on set theory such as B [1] and TLA$^+$ [15]. We believe that the results in this paper can help automating the proof of some parts of the verification conditions, which often mix arithmetic symbols, uninterpreted functions, and elementary set theory. Verification conditions generated within those formal methods are usually small and should be most of the time within reach of a decision procedure, even if it is inefficient.

An interesting direction for further research is to investigate how to use the techniques embedded in state-of-the-art first order provers (for instance [3, 18, 19]) to efficiently handle the first-order theories within a combination of decision procedures.

Acknowledgments: We are grateful to Yves Guiraud, Yuri Gurevich, Stephan Merz, Silvio Ranise, Christophe Ringeissen, Michaël Rusinowitch, and Duc-Khanh Tran for helpful remarks and interesting discussions. We also thank the anonymous reviewers for their comments.

References

1. J.-R. Abrial. *The B-Book: Assigning Programs to Meanings*. Cambridge University Press, 1996.
2. M. Baaz, U. Egly, and A. Leitsch. Normal form transformations. In J. A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume I, chapter 5, pages 273–333. Elsevier Science B.V., 2001.
3. P. Baumgartner, A. Fuchs, and C. Tinelli. Implementing the Model Evolution Calculus. In S. Schulz, G. Sutcliffe, and T. Tammet, editors, *Special Issue of the International Journal of Artificial Intelligence Tools (IJAIT)*, volume 15 of *International Journal of Artificial Intelligence Tools*, 2005.
4. P. Bernays and M. Schönfinkel. Zum Entscheidungsproblem der mathematischen Logik. *Math. Annalen*, 99:342–372, 1928.
5. E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer-Verlag, Berlin, 1997.
6. B. Dreben and W. D. Goldfarb. *The Decision Problem: Solvable Classes of Quantificational Formulas*. Addison-Wesley, Reading, Massachusetts, 1979.
7. H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Perspectives in Mathematical Logic. Springer-Verlag, Berlin, 1995.
8. H. B. Enderton. *A Mathematical Introduction to Logic*. Academic Press, Inc., Orlando, Florida, 1972.
9. P. Fontaine. Combinations of theories and the Bernays-Schönfinkel-Ramsey class. In B. Beckert, editor, *4th International Verification Workshop - VERIFY'07, Bremen, 15/07/07-16/07/07*, July 2007.
10. P. Fontaine. Combinations of theories for decidable fragments of first-order logic, 2009. Available at `http://www.loria.fr/~fontaine/Fontaine12b.pdf`.
11. P. Fontaine and E. P. Gribomont. Decidability of invariant validation for parameterized systems. In H. Garavel and J. Hatcliff, editors, *Tools and Algorithms for Construction and Analysis of Systems (TACAS)*, volume 2619 of *Lecture Notes in Computer Science*, pages 97–112. Springer-Verlag, 2003.

12. P. Fontaine and E. P. Gribomont. Combining non-stably infinite, non-first order theories. In W. Ahrendt, P. Baumgartner, H. de Nivelle, S. Ranise, and C. Tinelli, editors, *Selected Papers from the Workshops on Disproving and the Second International Workshop on Pragmatics of Decision Procedures (PDPAR 2004)*, volume 125 of *Electronic Notes in Theoretical Computer Science*, pages 37–51, July 2005.
13. E. Grädel, P. G. Kolaitis, and M. Y. Vardi. On the decision problem for two-variable first-order logic. *The Bulletin of Symbolic Logic*, 3(1):53–69, 1997.
14. Y. Gurevich and S. Shelah. Spectra of monadic second-order formulas with one unary function. In *LICS '03: Proceedings of the 18th Annual IEEE Symposium on Logic in Computer Science*, pages 291–300, Washington, DC, USA, 2003. IEEE Computer Society.
15. L. Lamport. *Specifying Systems*. Addison-Wesley, Boston, Mass., 2002.
16. G. Nelson and D. C. Oppen. Simplifications by cooperating decision procedures. *ACM Transactions on Programming Languages and Systems*, 1(2):245–257, Oct. 1979.
17. F. P. Ramsey. On a Problem of Formal Logic. *Proceedings of the London Mathematical Society*, 30:264–286, 1930.
18. A. Riazanov and A. Voronkov. The design and implementation of Vampire. *AI Communications*, 15(2):91–110, 2002.
19. S. Schulz. System Abstract: E 0.61. In R. Goré, A. Leitsch, and T. Nipkow, editors, *International Joint Conference on Automated Reasoning (IJCAR)*, number 2083 in Lecture Notes in Artificial Intelligence, pages 370–375. Springer, 2001.
20. G. Sutcliffe and C. Suttner. The State of CASC. *AI Communications*, 19(1):35–48, 2006.
21. C. Tinelli and M. T. Harandi. A new correctness proof of the Nelson–Oppen combination procedure. In F. Baader and K. U. Schulz, editors, *Frontiers of Combining Systems (FroCoS)*, Applied Logic, pages 103–120. Kluwer Academic Publishers, Mar. 1996.
22. C. Tinelli and C. Ringeissen. Unions of non-disjoint theories and combinations of satisfiability procedures. *Theoretical Computer Science*, 290(1):291–353, Jan. 2003.
23. C. Tinelli and C. G. Zarba. Combining non-stably infinite theories. *Journal of Automated Reasoning*, 34(3):209–238, 2005.
24. T. Wies, R. Piskac, and V. Kuncak. Combining theories with shared set operations. In *Frontiers of Combining Systems (FroCoS)*, 2009. To appear in this volume.