A Gentle Non-Disjoint Combination of Satisfiability Procedures

Paula Chocron[1,3], Pascal Fontaine[2], and Christophe Ringeissen[3⋆]

[1] Universidad de Buenos Aires, Argentina
[2] INRIA, Université de Lorraine & LORIA, Nancy, France
[3] INRIA & LORIA, Nancy, France

Abstract. A satisfiability problem is often expressed in a combination of theories, and a natural approach consists in solving the problem by combining the satisfiability procedures available for the component theories. This is the purpose of the combination method introduced by Nelson and Oppen. However, in its initial presentation, the Nelson-Oppen combination method requires the theories to be signature-disjoint and stably infinite (to guarantee the existence of an infinite model). The notion of gentle theory has been introduced in the last few years as one solution to go beyond the restriction of stable infiniteness, but in the case of disjoint theories. In this paper, we adapt the notion of gentle theory to the non-disjoint combination of theories sharing only unary predicates (plus constants and the equality). Like in the disjoint case, combining two theories, one of them being gentle, requires some minor assumptions on the other one. We show that major classes of theories, i.e. Löwenheim and Bernays-Schönfinkel-Ramsey, satisfy the appropriate notion of gentleness introduced for this particular non-disjoint combination framework.

1 Introduction

The design of satisfiability procedures has attracted a lot of interest in the last decade due to their ubiquity in SMT (Satisfiability Modulo Theories [4]) solvers and automated reasoners. A satisfiability problem is very often expressed in a combination of theories, and a very natural approach consists in solving the problem by combining the satisfiability procedures available for each of them. This is the purpose of the combination method introduced by Nelson and Oppen [15]. In its initial presentation, the Nelson-Oppen combination method requires the theories in the combination to be (1) signature-disjoint and (2) stably infinite (to guarantee the existence of an infinite model). These are strong limitations, and many recent advances aim to go beyond disjointness and stable infiniteness. Both corresponding research directions should not be opposed. In both cases, the problems are similar, i.e. building a model of $\mathcal{T}_1 \cup \mathcal{T}_2$ from a model of $\mathcal{T}_1$ and

⋆ This work has been partially supported by the project ANR-13-IS02-0001-01 of the Agence Nationale de la Recherche, by the European Union Seventh Framework Programme under grant agreement no. 295261 (MEALS), and by the STIC AmSud MISMT

a model of $\mathcal{T}_2$. This is possible if and only if there exists an isomorphism between the restrictions of the two models to the shared signature [24]. The issue is to define a framework to enforce the existence of this isomorphism. In the particular case of disjoint theories, the isomorphism can be obtained if the domains of the models have the same cardinality, for instance infinite; several classes of *kind* theories (shiny [25], polite [19], gentle [9]) have been introduced to enforce a (same) domain cardinality on both sides of the combination. For extensions of Nelson-Oppen to non-disjoint cases, e.g. in [24,27], cardinality constraints also arise. In this paper, we focus on non-disjoint combinations for which the isomorphism can be simply constructed by satisfying some cardinality constraints. More precisely, we extend the notion of gentle theory to the non-disjoint combination of theories sharing only unary predicates (plus constants and the equality). Some major classes of theories fit in our non-disjoint combination framework.

Contributions. The first contribution is to introduce a class of $\mathcal{P}$*-gentle* theories, to combine theories sharing a finite set of unary predicates symbols $\mathcal{P}$. The notion of $\mathcal{P}$-gentle theory extends the one introduced for the disjoint case [9]. Roughly speaking, a $\mathcal{P}$-gentle theory has nice cardinality properties not only for domains of models but also more locally for all Venn regions of shared unary predicates. We present a combination method for unions of $\mathcal{P}$-gentle theories sharing $\mathcal{P}$. The proposed method can also be used to combine a $\mathcal{P}$-gentle theory with another arbitrary theory for which we assume the decidability of satisfiability problems with cardinality constraints. This is a natural extension of previous works on combining non-stably infinite theories, in the straight line of combination methods à la Nelson-Oppen. Two major classes of theories are $\mathcal{P}$-gentle, namely the Löwenheim and Bernays-Schönfinkel-Ramsey (BSR) classes.

We characterize precisely the cardinality properties satisfied by Löwenheim theories. As a side contribution, bounds on cardinalities given in [8] have been improved, and we prove that our bounds are optimal. Our new result establishes that Löwenheim theories are $\mathcal{P}$-gentle.

We prove that BSR theories are also $\mathcal{P}$-gentle. This result relies on a non-trivial extension of Ramsey's Theorem on hypergraphs. This extension should be considered as another original contribution, since it may be helpful as a general technique to construct a model preserving the regions.

Related Work. Our combination framework is a way to combine theories with sets. The relation between (monadic) logic and sets is as old as logic itself, and this relation is particularly clear for instance considering Aristotle Syllogisms. It is however useful to again study monadic logic, and more particularly the Löwenheim class, and in view of the recent advances in combinations with non-disjoint and non-stably infinite theories.

In [26], the authors focus on the satisfiability problem of unions of theories sharing set operations. The basic idea is to reduce the combination problem into a satisfiability problem in a fragment of arithmetic called $BAPA$ (Boolean Algebra and Presburger Arithmetic). Löwenheim and BSR classes are also considered, but infinite cardinalities were somehow defined out of their reduction scheme,

whilst infinite cardinalities are smoothly taken into account in our combination framework. In [26], BSR was shown to be reducible to Presburger. We here give a detailed proof. We believe such a proof is useful since it is more complicated that it may appear. In particular, our proof is based on an original (up to our knowledge) extension of Ramsey's Theorem to accommodate a domain partitioned into (Venn) regions. Finally, the notion of $\mathcal{P}$-gentleness defined and used here is stronger than semi-linearity of Venn-cardinality, and allows non-disjoint combination with more theories, e.g. the guarded fragment.

In [21,22], a locality property is used to properly instantiate axioms connecting two disjoint theories. Hence, the locality is a way to reduce (via instantiation) a non-disjoint combination problem to a disjoint one. In that context, cardinality constraints occur when considering bridging functions over a data structure with some cardinality constraints on the underlying theory of elements [28,21,23].

In [12], Ghilardi proposed a very general model-theoretic combination framework to obtain a combination method à la Nelson-Oppen when $\mathcal{T}_1$ and $\mathcal{T}_2$ are two *compatible* extensions of the same shared theory (satisfying some properties). This framework relies on an application of the Robinson Joint Consistency Theorem (roughly speaking, the union of theories is consistent if the intersection is complete). Using this framework, several shared fragments of arithmetic have been successfully considered [12,16,17]. Due to its generality, Ghilardi's approach is free of cardinality constraints.

It is also possible to consider a general semi-decision procedure for the unsatisfiability problem modulo $\mathcal{T}_1 \cup \mathcal{T}_2$, e.g. a superposition calculus. With the rewrite-based approach initiated in [3], the problem reduces to proving the termination of this calculus. General criteria have been proposed to get modular termination results for superposition, when $\mathcal{T}_1$ and $\mathcal{T}_2$ are either disjoint [2] or non-disjoint [20]. Notice that the superposition calculus can also be used as a deductive engine to entail some cardinality constraints, as shown in [5].

Structure of the paper. Section 2 introduces some classical notations and definitions. In Section 3, we introduce the notion of $\mathcal{P}$-gentle theory and we present the related combination method for unions of theories sharing a (non-empty finite) set $\mathcal{P}$ of unary predicate symbols. All the theories in the Löwenheim class and in the BSR class are $\mathcal{P}$-gentle, as shown respectively in Section 4 and in Section 5. A simple example is given in Section 6. The conclusion (Section 7) discusses the current limitations of our approach and mentions some possible directions to investigate. Our extension of Ramsey's Theorem can be found in Appendix A.

2 Notation and Basic Definitions

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, while $\mathcal{F}$ and $\mathcal{P}$ are sets of function and predicate symbols. Every function and predicate symbol is assigned an arity. Nullary predicate symbols are called proposition symbols, and nullary function symbols are called constant

symbols. A first-order language is called relational if it only contains function symbols of arity zero. A relational formula is a formula in a relational language. Terms, atomic formulas and first-order formulas over the language $\mathcal{L}$ are defined in the usual way. In particular an atomic formula is either an equality, 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 literal is an atomic formula or the negation of an atomic formula. Free variables are defined in the usual way. A formula with no free variables is closed, and a formula without variables is ground. A universal formula is a closed formula $\forall x_1 \ldots \forall x_n.\varphi$ where φ is quantifier-free. A (finite) theory is a (finite) set of closed formulas. Two theories are disjoint if no predicate symbol in P or function symbol in F appears in both theories, except constants and equality.

An interpretation $\mathcal{I}$ for a first-order language $\mathcal{L}$ provides a non empty domain D, a total function $\mathcal{I}[f] : D^r \to D$ for every function symbol f of arity r, a predicate $\mathcal{I}[p] \subseteq D^r$ for every predicate symbol p of arity r, and an element $\mathcal{I}[x] \in D$ for every variable x. The cardinality of an interpretation is the cardinality of its domain. The notation $\mathcal{I}_{x_1/d_1,\ldots,x_n/d_n}$ for $x_1, \ldots, x_n$ different variables stands for the interpretation that agrees with $\mathcal{I}$, except that it associates $d_i \in D$ to the variable x_i, $1 \leq i \leq n$. By extension, an interpretation defines a value in D for every term, and a truth value for every formula. We may write $\mathcal{I} \models \varphi$ whenever $\mathcal{I}[\varphi] = \top$. Given an interpretation $\mathcal{I}$ on domain D, the *restriction* $\mathcal{I}'$ of $\mathcal{I}$ on $D' \subseteq D$ is the unique interpretation on D' such that $\mathcal{I}$ and $\mathcal{I}'$ interpret predicates, functions and variables the same way on D'. An *extension* $\mathcal{I}'$ of $\mathcal{I}$ is an interpretation on a domain D' including D such that $\mathcal{I}'$ restricted to D is $\mathcal{I}$.

A model of a formula (theory) is an interpretation that evaluates the formula (resp. all formulas in the theory) to true. A formula or theory is satisfiable if it has a model; 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. In our context, a theory $\mathcal{T}$ is *decidable* if the $\mathcal{T}$-satisfiability problem for sets of (ground) literals is decidable in the language of $\mathcal{T}$ (extended with fresh constants).

Consider an interpretation $\mathcal{I}$ on a language with unary predicates $p_1, \ldots, p_n$ and some elements D in the domain of this interpretation. Every element $d \in D$ belongs to a *Venn region* $v(d) = v_1 \ldots v_n \in \{\top, \bot\}^n$ where $v_i = \mathcal{I}[p_i](d)$. We denote by $D_v \subseteq D$ the set of elements of D in the Venn region v. Notice also that, for a language with n unary predicates, there are 2^n Venn regions. Given an interpretation $\mathcal{I}$, D^c denotes the subset of elements in D associated to constants by $\mathcal{I}$. Naturally, D^c_v denotes the set of elements associated to constants that are in the Venn region v.

3 Gentle Theories Sharing Unary Predicates

From now on, we assume that $\mathcal{P}$ is a non-empty finite set of unary predicates. A $\mathcal{P}$-*union* of two theories $\mathcal{T}_1$ and $\mathcal{T}_2$ is a union sharing only $\mathcal{P}$, a set of constants and the equality.

Definition 1. *An arrangement $\mathcal{A}$ for finite sets of constant symbols S and unary predicates $\mathcal{P}$ is a maximal satisfiable set of equalities and inequalities $a = b$ or $a \neq b$ and literals $p(a)$ or $\neg p(a)$, with $a, b \in S$, $p \in \mathcal{P}$.*

There are only a finite number of arrangements for given sets S and $\mathcal{P}$.

Given a theory $\mathcal{T}$ whose signature includes $\mathcal{P}$ and a model $\mathcal{M}$ of $\mathcal{T}$ on domain D, the *$\mathcal{P}$-cardinality* $\boldsymbol{\kappa}$ is the tuple of cardinalities of all Venn regions of $\mathcal{P}$ in $\mathcal{M}$ (κ_v will denote the cardinality of the Venn region v). The following theorem (specialization of general combination lemmas in e.g. [24,25]) states the completeness of the combination procedure for $\mathcal{P}$-unions of theories:

Theorem 1. *Consider a $\mathcal{P}$-union of theories $\mathcal{T}_1$ and $\mathcal{T}_2$ whose respective languages $\mathcal{L}_1$ and $\mathcal{L}_2$ share a finite set S of constants, and let L_1 and L_2 be sets of literals, respectively in $\mathcal{L}_1$ and $\mathcal{L}_2$. Then $L_1 \cup L_2$ is $\mathcal{T}_1 \cup \mathcal{T}_2$-satisfiable if and only if there exist an arrangement $\mathcal{A}$ for S and $\mathcal{P}$, and a $\mathcal{T}_i$-model $\mathcal{M}_i$ of $\mathcal{A} \cup L_i$ with the same $\mathcal{P}$-cardinality for $i = 1, 2$.*

The *spectrum* of a theory $\mathcal{T}$ is the set of $\mathcal{P}$-cardinalities of its models. The above theorem can thus be restated as:

Corollary 1. *The $\mathcal{T}_1 \cup \mathcal{T}_2$-satisfiability problem for sets of literals is decidable if, for any sets of literals $\mathcal{A} \cup L_1$ and $\mathcal{A} \cup L_2$ it is possible to decide if the intersection of the spectrums of $\mathcal{T}_1 \cup \mathcal{A} \cup L_1$ and of $\mathcal{T}_2 \cup \mathcal{A} \cup L_2$ is non-empty.*

To characterize the spectrum of the decidable classes considered in this paper, we introduce the notion of *cardinality constraint*. A *finite* cardinality constraint is simply a $\mathcal{P}$-cardinality with only finite cardinalities. An *infinite* cardinality constraint is given by a $\mathcal{P}$-cardinality $\boldsymbol{\kappa}$ with only finite cardinalities and a non-empty set of Venn regions V, and stands for all the $\mathcal{P}$-cardinalities $\boldsymbol{\kappa}'$ such that $\kappa'_v \geq \kappa_v$ if $v \in V$, and $\kappa'_v = \kappa_v$ otherwise. The *spectrum* of a finite set of cardinality constraints is the union of all $\mathcal{P}$-cardinalities represented by each cardinality constraint. It is now easy to define the class of theories we are interested in:

Definition 2. *A theory $\mathcal{T}$ is* $\mathcal{P}$-gentle *if, for every set L of literals in the language of $\mathcal{T}$, the spectrum of $\mathcal{T} \cup L$ is the spectrum of a computable finite set of cardinality constraints.*

Notice that a $\mathcal{P}$-gentle theory is (by definition) decidable. To relate the above notion with the gentleness in the disjoint case [9], observe that if p is a unary predicate symbol not occurring in the signature of the theory $\mathcal{T}$, then $\mathcal{T} \cup \{\forall x.p(x)\}$ is $\{p\}$-gentle if and only if $\mathcal{T}$ is gentle.

If a theory is $\mathcal{P}$-gentle, then it is $\mathcal{P}'$-gentle for any non-empty subset $\mathcal{P}'$ of $\mathcal{P}$. It is thus interesting to have $\mathcal{P}$-gentleness for the largest possible $\mathcal{P}$. Hence, when $\mathcal{P}$ is not explicitly given for a theory, we assume that $\mathcal{P}$ denotes the set of unary predicates symbols occurring in its signature. In the following sections we show that the Löwenheim theories and the BSR theories are $\mathcal{P}$-gentle.

The union of two $\mathcal{P}$-gentle theories is decidable, as a corollary of the following modularity result:

Theorem 2. *The class of $\mathcal{P}$-gentle theories is closed under $\mathcal{P}$-union.*

Proof. If we consider the $\mathcal{P}$-union of two $\mathcal{P}$-gentle theories with respective spectrums $\mathcal{S}_1$ and $\mathcal{S}_2$, then we can build some finite set of cardinality constraints whose spectrum is $\mathcal{S}_1 \cap \mathcal{S}_2$. □

Some very useful theories are not $\mathcal{P}$-gentle, but in practical cases they can be combined with $\mathcal{P}$-gentle theories. To define more precisely the class of theories $\mathcal{T}'$ that can be combined with a $\mathcal{P}$-gentle one, let us introduce the *$\mathcal{T}'$-satisfiability problem with cardinality constraints*: given a formula and a finite set of cardinality constraints, the problem amounts to check whether the formula is satisfiable in a model of $\mathcal{T}$ whose $\mathcal{P}$-cardinality is in the spectrum of the cardinality constraints. As a direct consequence of Corollary 1:

Theorem 3. *$\mathcal{T} \cup \mathcal{T}'$-satisfiability is decidable if the theory $\mathcal{T}$ is $\mathcal{P}$-gentle and $\mathcal{T}'$-satisfiability with cardinality constraints is decidable.*

Notice that $\mathcal{T}$-satisfiability with cardinality constraints is decidable for most common theories, e.g. the theories handled in SMT solvers. This gives the theoretical ground to add to the SMT solvers any number of $\mathcal{P}$-gentle theories sharing unary predicates.

From the results in the rest of the paper, it will also follow that the non-disjoint union (sharing unary predicates) of BSR and Löwenheim theories with one decidable theory accepting further constraints of the form $\forall x\,.\,((\neg)p_1(x) \wedge \ldots (\neg)p_n(x)) \Rightarrow (x = a_1 \vee \ldots x = a_m)$ is decidable. For instance, the guarded fragment with equality accepts such further constraints and the superposition calculus provides a decision procedure [11]. Thus any theory in the guarded fragment can be combined with Löwenheim and BSR theories sharing unary predicates.

In the disjoint case, any decidable theory expressed as a finite set of first-order axioms can be combined with a gentle theory [9]. Here this is not the case anymore. Indeed, consider the theory $\psi = \varphi \vee \exists x\, p(x)$ where p does not occur in φ; any set of literals is satisfiable in the theory ψ if and only if it is satisfiable in the theory of equality. If the satisfiability problem of literals in the theory φ is undecidable, the $\mathcal{P}$-union of ψ and the Löwenheim theory $\forall x\, \neg p(x)$ will also be undecidable.

4 The Löwenheim Class

We first review some classical results about this class and refer to [6] for more details. A Löwenheim theory is a finite set of closed formulas in a relational 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 [6], but does not have the finite model property anymore. Since the spectrum for this last class is significantly more complicated [13] than for the Löwenheim class we will here only focus on the Löwenheim class with equality (only classes with equality are relevant in our context), that is, without functions. More can be found about monadic first-order logic in [6,8]. In particular, a weaker version of Corollary 2 (given below) can be found in [8].

Previously [9,1], combining theories with non-stably infinite theories took advantage of "pumping" lemmas, allowing — for many decidable fragments — to build models of arbitrary large cardinalities. The following theorem is such a pumping lemma, but it considers the cardinalities of the Venn regions and not only the global cardinality.

Lemma 1. *Assume $\mathcal{T}$ is a Löwenheim theory with equality. Let q be the number of variables in $\mathcal{T}$. If there exists a model $\mathcal{M}$ on domain D with $|D_v \setminus D^c| \geq q$, then, for each cardinality $q' \geq q$, there is a model extension or restriction $\mathcal{M}'$ of $\mathcal{M}$ on domain D' such that $|D'_v \setminus D^c| = q'$ and $D'_{v'} = D_{v'}$ for all $v' \neq v$.*

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

- $|(D_v \cap D'_v) \setminus D^c| \geq q$;
- $D_{v'} = D'_{v'}$ for each Venn region v' distinct from v;
- $\mathcal{I}[a] = \mathcal{I}'[a]$ for each constant in ψ;
- $\mathcal{I}[x] = \mathcal{I}'[x]$ for each variable free in ψ.

Considering $\mathcal{M}$ as above, we can build a model $\mathcal{M}'$ as stated in the theorem, such that $\mathcal{M}$ and $\mathcal{M}'$ are similar. Indeed similarity perfectly defines a model with respect to another, given the cardinalities of the Venn regions.

We now prove that, given a Löwenheim formula ψ (or a set of formulas), two similar interpretations for ψ give the same truth value to ψ and to each sub-formula of ψ.

The proof is 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 constants. 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 $d \notin D'$, then $d \in D_v$ and $|(D_v \cap D'_v) \setminus D^c| \geq q$. Furthermore, since the whole formula contains at most q variables, $\varphi(x)$ contains at most $q-1$ free variables besides x. Let $x_1, \ldots, x_m$ be those variables. There exists some $d' \in (D_v \cap D'_v) \setminus D^c$ such that $d' \neq \mathcal{I}[x_i]$ for all $i \in \{1, \ldots, m\}$. 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. The proof for formulas of the form $\forall x\, \varphi(x)$ is dual. □

Lemma 1 has the following consequence on the acceptable cardinalities for the models of a Löwenheim theory:

Corollary 2. *Assume $\mathcal{T}$ is a Löwenheim theory with equality with n distinct unary predicates. Let r and q be respectively the number of constants and variables in $\mathcal{T}$. If $\mathcal{T}$ has a model of some cardinality κ strictly larger than $r + 2^n \max(0, q-1)$, then $\mathcal{T}$ has models of each cardinality equal or larger than $\min(\kappa, r + q\,2^n)$.*

Proof. If a model with such a cardinality exists, then there are Venn regions v such that $|D_v \setminus D^c| \geq q$. Then the number of elements in these Venn regions can be increased to any arbitrary larger cardinality, thanks to Lemma 1. If $\kappa > r + q\,2^n$, it means some Venn regions v are such that $|D_v \setminus D^c| > q$, and by eliminating elements in such Venn regions (using again Lemma 1), it is possible to obtain a model of cardinality $r + q\,2^n$. □

In [8], the limit is $q\,2^n$, q being the number of constants plus the maximum number of nested quantifiers. Now q is more precisely set to the number of variables, and the constants are counted separately. Moreover, $\max(0, q-1)$ replaces the factor q.

The case where q and r are both 0 corresponds to pure propositional logic (Löwenheim theories without variables and constants), where the size of the domain is not relevant. With $q = 1$ (one variable), there is no way to compare two elements (besides the ones associated to constants) and enforce them to be equal. It is still possible to constrain the domain to be of size at most r, using constraints like $\forall x\,.\,x = c_1 \vee \ldots \vee x = c_r$, but any model with one element not associated to a constant can be extended to a model of arbitrary cardinality (by somehow duplicating any number of time this element). Notice also that it is possible to set a lower bound on the size of the domain that can be $r + 2^n$. Consider for instance a set of sentences of the form $\exists x.(\neg)p_1(x) \vee \ldots (\neg)p_n(x)$; there are 2^n such formulas, each enforcing one Venn region to be non-empty.

Using several variables, a Löwenheim formula can enforce upper bounds larger than r on cardinalities. For $q = 2$, it is indeed easy to build a formula that has only models of cardinality at most $(q-1)\,2^n = 2^n$:

$$\forall x \forall y\,.\,\Big[\bigwedge_{0<i<j\leq n} p_i(x) = p_j(y)\Big] \Rightarrow x = y.$$

With a larger number of variables, the following formula ($q \geq 2$)

$$\forall x_1 \ldots \forall x_q\,.\,\Big[\bigwedge_{\substack{0<i<j\leq n \\ 0<i'<j'\leq q}} p_i(x_{i'}) = p_j(x_{j'})\Big] \Rightarrow \bigvee_{0<i'<j'\leq q} x_{i'} = x_{j'}$$

enforces the cardinality of the domain to be at most $(q-1)\,2^n$. To obtain a formula with constants that accepts only models of cardinality up to $r + 2^n \max(0, q-1)$, it suffices to add as a guard in the above formula the conjunctive sets of atoms expressing that the variables are disjoint from the r constants. So the above condition in Corollary 2 is the strongest one.

Besides the finite model property and the decidability of Löwenheim theories, Corollary 2 also directly entails the $\mathcal{P}$-gentleness:

Theorem 4. *Löwenheim theories on a language with unary predicates in $\mathcal{P}$ are $\mathcal{P}$-gentle.*

5 The Bernays-Schönfinkel-Ramsey Class

A Bernays-Schönfinkel-Ramsey (BSR for short) theory 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. More can be found about BSR theories in [6]. Ramsey also gave some (less known) results about the spectrum of BSR theories [18]. We here give a proof that BSR theories are $\mathcal{P}$-gentle.

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

Lemma 2. *Let $\mathcal{T}$ be a BSR theory, and $\mathcal{M}$ be a model of $\mathcal{T}$ on domain D. Then any restriction $\mathcal{M}'$ of $\mathcal{M}$ on domain D' with $D^c \subseteq D' \subseteq D$ is a model of $\mathcal{T}$.*

Proof. Consider $\mathcal{M}$ and $\mathcal{M}'$ as above. Since $\mathcal{M}$ is a model of $\mathcal{T}$, for each closed formula $\forall x_1 \dots x_n \,.\, \varphi$ in $\mathcal{T}$ (where φ is function-free and quantifier-free), and for all $d_1, \dots, d_n \in D' \subseteq D$, $\mathcal{M}_{x_1/d_1,\dots,x_n/d_n}$ is a model of φ. This also means that, for all $d_1, \dots, d_n \in D'$, $\mathcal{M}'_{x_1/d_1,\dots,x_n/d_n}$ is a model of φ, and finally that $\mathcal{M}'$ is a model of $\forall x_1 \dots x_n \,.\, \varphi$. □

Intuitively, this states that the elements not assigned to ground terms (i.e. the constants) can be eliminated from a model of a BSR theory. It is known [18,9] that for any BSR theory $\mathcal{T}$ there is a computable finite number k such that if $\mathcal{T}$ has a model of cardinality greater or equal to k, then it has a model of any cardinality larger than k. Later in this section, we prove that the same occurs locally for each Venn region.

The notion of n-repetitive models, which we now define, is instrumental for this. Informally, a model is n-repetitive if it is symmetric for those elements of its domain that are not assigned to constants in the theory.

Definition 3. *An interpretation $\mathcal{I}$ on domain D for a BSR theory $\mathcal{T}$ is* n-repetitive for a set V of Venn regions *if, for each $v \in V$, $|D_v \setminus D^c| \geq n$ and there exists a total order $\prec$ on elements in $D_v \setminus D^c$ such that*

- *for every r-ary predicate symbol p in $\mathcal{T}$*
- *for all $d_1, \dots, d_r \in D$, and $d'_1, \dots, d'_r \in D$ with*
  - $|\{d_1, \dots, d_r\} \setminus D^c| \leq n$
  - $d'_i = d_i$ *if* d_i *or* $d'_i \in D^c \cup \bigcup_{v \notin V} D_v$
  - $v(d'_i) = v(d_i)$

- $d'_i \prec d'_j$ *iff* $d_i \prec d_j$*, if for some* $v' \in V$*,* $d_i, d_j \in D_{v'} \setminus D^c$

we have $\mathcal{I}[p](d_1, \ldots, d_r) = \mathcal{I}[p](d'_1, \ldots, d'_r)$.

Notice that a same interpretation can be n-repetitive for several Venn regions at the same time. Also, the above definition allows $D_v \setminus D^c$ to be empty for every $v \notin V$. Previously [9] (without distinguishing regions) we showed that one can decide if a BSR theory $\mathcal{T}$ is n-repetitive by building another BSR theory that is satisfiable if and only if $\mathcal{T}$ is n-repetitive. The same occurs to n-repetitiveness for Venn regions.

Theorem 5. *Consider a BSR theory* $\mathcal{T}$ *with* n *variables and a model* $\mathcal{M}$ *on domain* D*. If* $\mathcal{M}$ *is* n*-repetitive for the Venn regions* V *then, for any (finite or infinite) cardinalities* $\kappa_v \geq |D_v|$ *(*$v \in V$*),* $\mathcal{T}$ *has a model* $\mathcal{M}'$ *extension of* $\mathcal{M}$ *on domain* D' *such that* $|D'_v| = \kappa_v$ *if* $v \in V$ *and* $D'_{v'} = D_{v'}$ *for all* $v' \notin V$*.*

Proof. Assume that $\prec$ are the total orders mentioned in Definition 3. We first build an extension $\mathcal{M}'$ of $\mathcal{M}$ as specified in the theorem, and later prove it is a model of $\mathcal{T}$.

Let E be the set of new elements $E = D' \setminus D$, and fix arbitrary total orders (again denoted by $\prec$) on $D'_v \setminus D^c$ for all $v \in V$ that extend the given orders on $D_v \setminus D^c$. Since $\mathcal{M}'$ is an extension of $\mathcal{M}$, the interpretation of the predicate symbols is already defined when all arguments belong to D. When some arguments belong to E, the truth value of an r-ary predicate p is defined as follows:

- $(d'_1, \ldots, d'_r) \notin \mathcal{M}'[p]$ for $|\{d'_1, \ldots, d'_r\} \setminus D^c| > n$: the interpretation of p over tuples with more than n elements outside D^c is fixed arbitrarily. Indeed, such tuples are irrelevant for the evaluation of the formulas of $\mathcal{T}$: terms occurring as arguments of a predicate are either variables or constants, and no more than n variables occur in any formula of $\mathcal{T}$.
- otherwise, to determine $\mathcal{M}'[p](d'_1, \ldots, d'_r)$, first choose $d_1, \ldots d_r \in D$ such that $d'_1, \ldots, d'_r$ and $d_1, \ldots d_r$ are related to each other just like in Definition 3. This is possible since, for every Venn region v for which the interpretation is repetitive, there are at least n elements in $D_v \setminus D^c$. Then $(d'_1, \ldots, d'_r) \in \mathcal{M}'[p]$ iff $(d_1, \ldots, d_r) \in \mathcal{M}[p]$. Observe that all possible choices of $d_1, \ldots, d_n$ lead to the same definition because $\mathcal{M}$ is n-repetitive.

The construction is such that $\mathcal{M}'$ is also n-repetitive for the same regions. It is also a model of $\mathcal{T}$: all formulas in $\mathcal{T}$ are of the form $\forall x_1 \ldots x_m \,.\, \varphi(x_1, \ldots, x_m)$, with $m \leq n$. For all $d'_1 \ldots, d'_m \in D'$, if $\{d'_1, \ldots, d'_m\} \subseteq D$ then

$$\mathcal{M}'_{x_1/d'_1, \ldots, x_m/d'_m}[\varphi(x_1, \ldots, x_m)] = \mathcal{M}_{x_1/d'_1, \ldots, x_m/d'_m}[\varphi(x_1, \ldots, x_m)]$$

since $\mathcal{M}'$ is an extension of $\mathcal{M}$. Otherwise, let $d_1, \ldots, d_m \in D$ be some elements related to $d'_1, \ldots, d'_m$ like in Definition 3. Since $\mathcal{M}'$ is n-repetitive,

$$\begin{aligned} \mathcal{M}'_{x_1/d'_1, \ldots, x_m/d'_m}[\varphi(x_1, \ldots, x_m)] &= \mathcal{M}'_{x_1/d_1, \ldots, x_m/d_m}[\varphi(x_1, \ldots, x_m)] \\ &= \mathcal{M}_{x_1/d_1, \ldots, x_m/d_m}[\varphi(x_1, \ldots, x_m)]. \end{aligned}$$

In both cases, $\mathcal{M}'_{x_1/d'_1, \ldots, x_m/d'_m}[\varphi(x_1, \ldots, x_m)]$ evaluates to true, and therefore $\mathcal{M}'$ is a model of $\forall x_1 \ldots x_n \,.\, \varphi(x_1, \ldots, x_m)$. □

Now it is possible to state that the full spectrum of a BSR theory only depends on (a finite set of) $\mathcal{P}$-cardinalities $\boldsymbol{\kappa}$ such that, for all Venn region v, $\kappa_v \leq k$ for some finite cardinality k only depending on the theory. The proof requires an extension of Ramsey's Theorem which can be found in the appendix A.

Theorem 6. *Given a BSR theory $\mathcal{T}$ with n variables, there exists a number k computable from the theory, such that, if $\mathcal{T}$ has a model $\mathcal{M}$ on domain D such that $|D_v \setminus D^c| \geq k$ for Venn regions $v \in V$, then it has a model which is n-repetitive for Venn regions V.*

Proof. Using Lemma 2, we can assume that $\mathcal{T}$ has a (sufficiently large) finite model $\mathcal{M}$ on domain D. We can assume without loss of generality that $\mathcal{M}$ is such that, for every predicate p of the language, $(d_1, \dots d_r) \notin \mathcal{M}[p]$ whenever there are more than n elements in $\{d_1, \dots d_r\} \setminus D^c$; indeed, these interpretations play no role in the truth value of a formula with n variables.

Let $\prec$ be an order on $D \setminus D^c$. Given two ordered (with respect to $\prec$) sequences $e_1, \dots, e_n$ and $e'_1, \dots, e'_n$ of elements in $D \setminus D^c$ such that $v(e_i) = v(e'_i)$ $(1 \leq i \leq n)$, we say that the configurations for $e_1, \dots, e_n$ and $e'_1, \dots, e'_n$ agree if for every r-ary predicate p, and for every $d_1, \dots, d_r \in D^c \cup \{e_1, \dots, e_n\}$, $(d_1, \dots, d_r) \in \mathcal{M}[p]$ iff $(d'_1, \dots, d'_r) \in \mathcal{M}[p]$, with $d'_i = e'_j$ if $d_i = e_j$ for some j, and $d'_i = d_i$ otherwise. Notice that there are only a finite number of disagreeing configurations for n elements in $D \setminus D^c$: more precisely a configuration is determined by at most $b = \sum_p (n + |D^c|)^{\text{arity}(p)}$ Boolean values, where the sum ranges over all predicates in the theory. Thus the number of disagreeing configurations is bounded by $C = 2^b$.

Interpreting configurations as colors, one can use the extension of Ramsey's Theorem given in Appendix A: according to Theorem 7, there is a computable function f such that, for any $N \in \mathbb{N}$, if $|D \setminus D^c|_V \geq f(n, N, C)$, then there exists a model on $D' \subseteq D$ with $|D' \setminus D^c|_V \geq N$ for which configurations agree if they have the same number of elements in each Venn region of V. Taking $N = n$, this is actually building a n-repetitive restriction of $\mathcal{M}$. □

The BSR class obviously has the finite model property, and is decidable. Lemma 2 and Theorems 5 and 6 above also prove that BSR theories are (gentle and) $\mathcal{P}$-gentle:

Corollary 3. *BSR theories on a language including unary predicates in $\mathcal{P}$ are $\mathcal{P}$-gentle.*

A simple constructive proof of this corollary would consider the finite number of all $\mathcal{P}$-cardinalities $\boldsymbol{\kappa}$ such that $\boldsymbol{\kappa}_v \leq k$ (where k comes from Theorem 6). All such $\mathcal{P}$-cardinalities can be understood as cardinality constraints, the extendable Venn regions being the ones for which $\boldsymbol{\kappa}_v > k$. Of course this construction is highly impractical, since it uses some kind of Ramsey numbers, known to be extremely large. In practice, we believe there are much better constructions: the important elements of the domain are basically only the ones associated to constants, and theoretical upper bounds are not met in non-artificial cases.

6 Example: Non-Disjoint Combination of Order and Sets

To illustrate the kind of theories that can be handled in our framework, consider a simple yet informative example with a BSR theory defining an ordering $<$ and augmented with clauses connecting the ordering $<$ and the sets p and q (we do not distinguish sets and their related predicates):

$$\mathcal{T}_1 = \begin{cases} \forall x.\ \neg(x < x) \\ \forall x, y, z.\ (x < y \wedge y < z) \Rightarrow x < z \\ \forall x, y.\ (p(x) \wedge \neg p(y)) \Rightarrow x < y \\ \forall x, y.\ (q(x) \wedge \neg q(y)) \Rightarrow x < y \end{cases}$$

and a Löwenheim theory

$$\mathcal{T}_2 = \begin{cases} \exists y \forall x.\ (p(x) \wedge q(x)) \equiv x = y \\ \forall x \exists y.\ p(x) \Rightarrow (x \neq y \wedge q(y)) \end{cases}$$

Notice that $\mathcal{T}_2$ is not a BSR theory due to the $\forall\exists$ quantification of its second axiom, but both theories $\mathcal{T}_1$ and $\mathcal{T}_2$ are actually $\mathcal{P}$-gentle. The theory $\mathcal{T}_1$ imposes either $p \cap \overline{q}$ or $\overline{p} \cap q$ to be empty (we will assume that the domain is non-empty and simplify the cardinality constraints accordingly). The theory $\mathcal{T}_2$ imposes the cardinality of $p \cap q$ to be exactly 1, and the cardinality of $\overline{p} \cap q$ to be at least 1. The following table collects the cardinality constraints:

	$\mathcal{T}_1$		$\mathcal{T}_2$
$\overline{p} \cap \overline{q}$	≥ 0	≥ 0	≥ 0
$\overline{p} \cap q$	0	≥ 0	≥ 1
$p \cap \overline{q}$	≥ 0	0	≥ 0
$p \cap q$	≥ 0	≥ 0	1

The theory $\mathcal{T}_1 \cup \mathcal{T}_2$ imposes $p \cap \overline{q}$ to be empty, in other words $p \subseteq q$. Moreover, the cardinality of $p \cap q$ is 1, and so it implies that the cardinality of p is 1. Hence, the set

$$\mathcal{T}_1 \cup \mathcal{T}_2 \cup \{p(a), p(b), a \neq b\}$$

is unsatisfiable. As a final comment, there could be theories using directly the Venn cardinalities as integer variables. For instance, imagine a constraint stating $|p| > 1$ in a theory including linear arithmetic on integers. This would of course be unsatisfiable with $\mathcal{T}_1 \cup \mathcal{T}_2$.

7 Conclusion

The notion of gentleness was initially presented as a tool to combine non-stably infinite disjoint theories. In this paper, we have introduced a notion of $\mathcal{P}$-gentleness which is well-suited for combining theories sharing (besides constants and the equality) only unary predicates in a set $\mathcal{P}$. The major contributions of this paper are that the Löwenheim theories and BSR theories are $\mathcal{P}$-gentle. A

corollary is that the non-disjoint union (sharing unary predicates) of Löwenheim theories, BSR theories, and decidable theories accepting further constraints of the form $\forall x \,.\, ((\neg)p_1(x) \wedge \ldots (\neg)p_n(x)) \Rightarrow (x = a_1 \vee \ldots x = a_m)$ is decidable.

Our combination method is limited to shared unary predicates. Unfortunately, the theoretical limitations are strong for a framework sharing predicates with larger arities: for instance even the guarded fragment with two variables and transitivity constraints is undecidable [10], although the guarded fragment (or first-order logic with two variables) is decidable, and transitivity constraints can be expressed in BSR. The problem of combining theories with only a shared dense order has however been successfully solved [12,14]. In that specific case, there is again an implicit infiniteness argument that could be possibly expressed as a form of extended gentleness, to reduce the isomorphism construction problem into solving some appropriate extension of cardinality constraints. A clearly challenging problem is to identify an appropriate extended notion of gentleness for some particular binary predicates.

Also in future works, the reduction approach (Löwenheim and BSR theories can be simplified to a subset of Löwenheim) may be useful as a simplification procedure for sets of formulas that can be seen as non-disjoint (sharing unary predicates only) combinations of BSR, Löwenheim theories and an arbitrary first-order theory: this would of course not provide a decision procedure, but refutational completeness can be preserved. More generally we also plan to study how superposition-based satisfiability procedures could benefit from a non-disjoint (sharing unary predicates) combination point of view. In particular, superposition-based satisfiability procedures could be used as deductive engines with the capability to exchange constraints à la Nelson-Oppen.

The results here are certainly too combinatorially expensive to be directly applicable. However, this paper paves the theoretical grounds for mandatory further works that would make such combinations practical. There are important incentives since the BSR and Löwenheim fragments are quite expressive: for instance, it is possible to extend the language of SMT solvers with sets and cardinalities. Many formal methods are based on logic languages with sets. Expressive decision procedures (even if they are not efficient) including e.g. sets and cardinalities will help proving the often small but many verification conditions stemming from these applications.

References

1. Areces, C., Fontaine, P.: Combining theories: The Ackerman and Guarded fragments. In Tinelli, C., Sofronie-Stokkermans, V., eds.: Frontiers of Combining Systems (FroCoS). Volume 6989 of LNCS., Springer (2011) 40–54
2. Armando, A., Bonacina, M.P., Ranise, S., Schulz, S.: New results on rewrite-based satisfiability procedures. ACM Trans. Comput. Log. **10**(1) (2009)
3. Armando, A., Ranise, S., Rusinowitch, M.: A rewriting approach to satisfiability procedures. Inf. Comput. **183**(2) (2003) 140–164
4. Barrett, C., Sebastiani, R., Seshia, S.A., Tinelli, C.: Satisfiability modulo theories. In Biere, A., Heule, M.J.H., van Maaren, H., Walsh, T., eds.: Handbook of Satis-

fiability. Volume 185 of Frontiers in Artificial Intelligence and Applications. IOS Press (February 2009) 825–885
5. Bonacina, M.P., Ghilardi, S., Nicolini, E., Ranise, S., Zucchelli, D.: Decidability and undecidability results for Nelson-Oppen and rewrite-based decision procedures. In Furbach, U., Shankar, N., eds.: International Joint Conference on Automated Reasoning (IJCAR). Volume 4130 of LNCS., Springer (2006) 513–527
6. Börger, E., Grädel, E., Gurevich, Y.: The Classical Decision Problem. Perspectives in Mathematical Logic. Springer-Verlag, Berlin (1997)
7. Chocron, P., Fontaine, P., Ringeissen, C.: A Gentle Non-Disjoint Combination of Satisfiability Procedures (Extended Version). Research Report 8529, Inria (2014) `http://hal.inria.fr/hal-00985135`.
8. Dreben, B., Goldfarb, W.D.: The Decision Problem: Solvable Classes of Quantificational Formulas. Addison-Wesley, Reading, Massachusetts (1979)
9. Fontaine, P.: Combinations of theories for decidable fragments of first-order logic. In Ghilardi, S., Sebastiani, R., eds.: Frontiers of Combining Systems (FroCoS). Volume 5749 of LNCS., Springer (2009) 263–278
10. Ganzinger, H., Meyer, C., Veanes, M.: The two-variable guarded fragment with transitive relations. In: Logic In Computer Science (LICS), IEEE Computer Society (1999) 24–34
11. Ganzinger, H., Nivelle, H.D.: A superposition decision procedure for the guarded fragment with equality. In: Logic In Computer Science (LICS), IEEE Computer Society Press (1999) 295–303
12. Ghilardi, S.: Model-theoretic methods in combined constraint satisfiability. Journal of Automated Reasoning **33**(3-4) (2004) 221–249
13. Gurevich, Y., Shelah, S.: Spectra of monadic second-order formulas with one unary function. In: Logic In Computer Science (LICS), Washington, DC, USA, IEEE Computer Society (2003) 291–300
14. Manna, Z., Zarba, C.G.: Combining decision procedures. In Aichernig, B.K., Maibaum, T.S.E., eds.: Formal Methods at the Crossroads. From Panacea to Foundational Support, 10th Anniversary Colloquium of UNU/IIST, Revised Papers. Volume 2757 of LNCS., Springer (2003) 381–422
15. Nelson, G., Oppen, D.C.: Simplification by cooperating decision procedures. ACM Trans. on Programming Languages and Systems **1**(2) (October 1979) 245–257
16. Nicolini, E., Ringeissen, C., Rusinowitch, M.: Combinable extensions of Abelian groups. In Schmidt, R.A., ed.: Proc. Conference on Automated Deduction (CADE). Volume 5663 of LNCS., Springer (2009) 51–66
17. Nicolini, E., Ringeissen, C., Rusinowitch, M.: Combining satisfiability procedures for unions of theories with a shared counting operator. Fundam. Inform. **105**(1-2) (2010) 163–187
18. Ramsey, F.P.: On a Problem of Formal Logic. Proceedings of the London Mathematical Society **30** (1930) 264–286
19. Ranise, S., Ringeissen, C., Zarba, C.G.: Combining data structures with nonstably infinite theories using many-sorted logic. In Gramlich, B., ed.: Frontiers of Combining Systems (FroCoS). Volume 3717 of LNCS., Springer (2005) 48–64
20. Ringeissen, C., Senni, V.: Modular termination and combinability for superposition modulo counter arithmetic. In Tinelli, C., Sofronie-Stokkermans, V., eds.: Frontiers of Combining Systems (FroCoS). Volume 6989 of LNCS., Springer (2011) 211–226
21. Sofronie-Stokkermans, V.: Locality results for certain extensions of theories with bridging functions. In Schmidt, R.A., ed.: Proc. Conference on Automated Deduction (CADE). Volume 5663 of LNCS., Springer (2009) 67–83

22. Sofronie-Stokkermans, V.: On combinations of local theory extensions. In Voronkov, A., Weidenbach, C., eds.: Programming Logics - Essays in Memory of Harald Ganzinger. Volume 7797 of LNCS., Springer (2013) 392–413
23. Suter, P., Dotta, M., Kuncak, V.: Decision procedures for algebraic data types with abstractions. In Hermenegildo, M.V., Palsberg, J., eds.: Principles of Programming Languages (POPL), ACM (2010) 199–210
24. Tinelli, C., Ringeissen, C.: Unions of non-disjoint theories and combinations of satisfiability procedures. Theoretical Computer Science **290**(1) (2003) 291–353
25. Tinelli, C., Zarba, C.G.: Combining non-stably infinite theories. Journal of Automated Reasoning **34**(3) (April 2005) 209–238
26. Wies, T., Piskac, R., Kuncak, V.: Combining theories with shared set operations. In Ghilardi, S., Sebastiani, R., eds.: Frontiers of Combining Systems (FroCoS). Volume 5749 of LNCS., Springer (2009) 366–382
27. Zarba, C.G.: Combining sets with cardinals. J. Autom. Reasoning **34**(1) (2005) 1–29
28. Zhang, T., Sipma, H.B., Manna, Z.: Decision procedures for term algebras with integer constraints. Inf. Comput. **204**(10) (2006) 1526–1574

A An Extension of Ramsey's Theorem

We define an *n-subset* of S to be a subset of n elements of S. An *n-hypergraph* of S is a set of n-subsets of S. In particular, a 2-hypergraph is an (undirected) graph. The *complete* n-hypergraph of S is the set of all n-subsets of S, and its *size* is the cardinality of S. An n-hypergraph G is *colored* with c colors if there is a coloring function that assigns one color to every n-subset in G. In particular, a colored 2-hypergraph (that is, a colored graph), is a graph where all edges are assigned a color. Consider a set S of elements partitioned into disjoint regions $R = \{R_1, \dots R_m\}$. We say that a set $S' \subseteq S$ has *region size* larger than x and note $|S'|_R \geq x$ if $|S' \cap R_i| \geq x$ for all $i \in \{1, \dots, m\}$. We also say that an n-hypergraph is *region-monochromatic* if the color of each hyperedge only depends on the number of elements belonging to each region. Two hyperedges are said of the *same kind* if they have the same number of elements in each region; all hyperedges of the same kind of a region-monochromatic hypergraph thus have the same color. The following extension[4] of Ramsey's Theorem holds:

Theorem 7. *There exists a computable function f such that,*

- *for every number of colors c*
- *for every $n, N \in \mathbb{N}$*
- *for every complete n-hypergraph G on S colored with c colors*

if $|S|_R \geq f(n, N, c)$, then there exists a complete region-monochromatic n-sub-hypergraph of G on some $S' \subseteq S$ with $|S'|_R \geq N$.

Proof. The full proof can be found in [7]. □

[4] The classical Ramsey's Theorem is the case with only one region.