Deciding First-Order Satisfiability when Universal and Existential Variables are Separated

Thomas Sturm

CNRS,* Inria,† and University of Lorraine,‡ Nancy, France

Marco Voigt

Max Planck Institute for Informatics, Saarbrücken, Germany, and
Saarbrücken Graduate School of Computer Science, Saarland University

Christoph Weidenbach

Max Planck Institute for Informatics, Saarbrücken, Germany

Abstract

We introduce a new decidable fragment of first-order logic with equality, which strictly generalizes two already well-known ones—the Bernays–Schönfinkel–Ramsey (BSR) Fragment and the Monadic Fragment. The defining principle is the syntactic separation of universally quantified variables from existentially quantified ones at the level of atoms. Thus, our classification neither rests on restrictions on quantifier prefixes (as in the BSR case) nor on restrictions on the arity of predicate symbols (as in the monadic case). We demonstrate that the new fragment exhibits the finite model property and derive a non-elementary upper bound on the computing time required for deciding satisfiability in the new fragment. For the subfragment of prenex sentences with the quantifier prefix $\exists^*\forall^*\exists^*$ the satisfiability problem is shown to be complete for NEXPTIME. Finally, we discuss how automated reasoning procedures can take advantage of our results.

1 Introduction

The question of whether satisfiability of first-order sentences of a certain syntactic form is decidable or not has a long tradition in computational logic. Over the decades different dimensions have been introduced along which decidable first-order fragments can be separated from undecidable ones. Löwenheim's pioneering work [26] shows that confinement to unary predicate symbols (i.e. to the *Relational Monadic Fragment*) leads to decidability, while the set of sentences in which binary predicate symbols may be used without further restriction yields a *reduction class* for first-order logic. Another dimension is the confinement to certain quantifier prefixes for formulas in *prenex normal form*, e.g. to $\exists^*\forall^*$—the *Bernays–Schönfinkel (BS) Fragment* [5]—consisting of decidable sentences. The inclusion or exclusion of the distinguished equality predicate yields yet another possibility. Löwenheim had already considered it for the Relational Monadic Fragment, and Ramsey [30] extended the BS Fragment in this way, leading to the *Bernays–Schönfinkel–Ramsey (BSR) Fragment*, which also remains decidable. Already in the nineties, the classification along those dimensions was considered to be mainly completed, cf. [12, 7].

In the present paper we introduce another dimension of classification: syntactic separation of existentially quantified variables from universally quantified ones at the level of atoms. We disallow atoms $P(\ldots, x, \ldots, y, \ldots)$ that contain both a universally quantified variable x and an

---

*CNRS, LORIA, UMR 7503, 54506 Vandoeuvre-lès-Nancy, France

†Inria, 54600 Villers-lès-Nancy, France

‡University of Lorraine, LORIA, UMR 7503, 54506 Vandoeuvre-lès-Nancy, France

existentially quantified variable y at the same time. We call the class of first-order sentences built from atoms with thus separated variables the *Separated Fragment (SF)*. The Separated Fragment is decidable. It properly includes both the BSR Fragment and the Relational Monadic Fragment (monadic formulas without non-constant function symbols) without equality and is to the best of our knowledge the first known decidable fragment enjoying this property. Consequently, separation of differently quantified variables constitutes a unifying principle which underlies both the BSR Fragment and the Relational Monadic Fragment. More concretely, our contributions are:

(i) We introduce the new decidable first-order fragment SF (Sections 3, 4). Its sentences can be transformed into equivalent BSR sentences and thus enjoy the finite model property. Our current upper bound on the size of small models is non-elementary in the length of the formula (Theorem 17).

(ii) For SF formulas in which blocks of existentially quantified variables are pairwise disjoint the size of small models is at most double exponential (Theorem 15).

(iii) For SF sentences of the form $\exists^*\forall^*\exists^*\varphi$, where φ is quantifier-free, the size of small models is single exponential in the length of the formula. Therefore, satisfiability is NEXPTIME-complete in this setting (Theorem 14).

(iv) Sentences from the Monadic Fragment with unary function symbols or from the Relational Monadic Fragment with equality can be translated into SF sentences. The translation preserves satisfiability and increases the length of the formulas only polynomially (Section 4.2).

(v) We provide a methodology for the translation of SF sentences of the form $\exists^*\forall^*\exists^*\varphi$ into the BS(R) fragment, which facilitates automated reasoning (Section 5).

The paper concludes with a discussion of related and future work in Section 6.

The present paper is an extended version of [32].

2 Preliminaries

We consider first-order logic formulas. The underlying signature shall not be mentioned explicitly, but will become clear from the current context. For the distinguished *equality* predicate (whose semantics is fixed to be the identity relation) we use $\approx$. If not explicitly excluded, we allow equality in our investigations. As usual, we interpret a formula φ with respect to given structures. A *structure* $\mathcal{A}$ consists of a nonempty *universe* $\mathfrak{U}_\mathcal{A}$ (also: *domain*) and interpretations $f^\mathcal{A}$ and $P^\mathcal{A}$ of all considered function and predicate symbols, respectively, in the usual way. Given a formula φ, a structure $\mathcal{A}$, and a variable assignment β, we write $\mathcal{A}, \beta \models \varphi$ if φ evaluates to *true* under $\mathcal{A}$ and β. We write $\mathcal{A} \models \varphi$ if $\mathcal{A}, \beta \models \varphi$ holds for every β. The symbol $\models\!\mid$ denotes *semantic equivalence* of formulas, i.e. $\varphi \models\!\mid \psi$ holds whenever for every structure $\mathcal{A}$ and every variable assignment β we have $\mathcal{A}, \beta \models \varphi$ if and only if $\mathcal{A}, \beta \models \psi$. We call two sentences φ and ψ *equisatisfiable* if φ has a model if and only if ψ has one.

We use $\varphi(x_1, \ldots, x_m)$ to denote a formula φ whose free variables form a subset of $\{x_1, \ldots, x_m\}$. In all formulas we tacitly assume that no variable occurs freely and bound at the same time and that all quantifiers bind distinct variables. For convenience, we sometimes identify tuples $\vec{\mathbf{x}}$ of variables with the set containing all the variables that occur in $\vec{\mathbf{x}}$.

A structure $\mathcal{A}$ is a *substructure* of a structure $\mathcal{B}$ (over the same signature) if (1) $\mathfrak{U}_\mathcal{A} \subseteq \mathfrak{U}_\mathcal{B}$, (2) $c^\mathcal{A} = c^\mathcal{B}$ for every constant symbol c, (3) $P^\mathcal{A} = P^\mathcal{B} \cap \mathfrak{U}_\mathcal{A}^m$ for every m-ary predicate symbol P, and (4) $f^\mathcal{A}(\mathfrak{a}_1, \ldots, \mathfrak{a}_m) = f^\mathcal{B}(\mathfrak{a}_1, \ldots, \mathfrak{a}_m)$ for every m-ary function symbol f and every m-tuple $\langle \mathfrak{a}_1, \ldots, \mathfrak{a}_m \rangle \in \mathfrak{U}_\mathcal{A}^m$. The following is a standard lemma, see, e.g., [13] for a proof.

Lemma 1 (Substructure Lemma)**.** Let φ be a first-order sentence in prenex normal form without existential quantifiers and let $\mathcal{A}$ be a substructure of $\mathcal{B}$. $\mathcal{B} \models \varphi$ entails $\mathcal{A} \models \varphi$.

We denote *substitution* by $\varphi[x/t]$, where every free occurrence of x in φ is to be substituted by the term t. For *simultaneous substitution* of pairwise distinct variables $x_1, \ldots, x_n$ with $t_1, \ldots, t_n$, respectively, we use the notation $\varphi[x_1/t_1, \ldots, x_n/t_n]$. For example, $P(x, y)[x/f(y), y/g(x)]$ results in $P(f(y), g(x))$. We also write $[\vec{x}/\vec{t}]$ to abbreviate $[x_1/t_1, \ldots, x_n/t_n]$.

Lemma 2 (Miniscoping)**.** Let φ, ψ, χ be first-order formulas, and assume that x does not occur freely in χ.

(i) $\exists x.(\varphi \vee \psi) \ \models\mid\ (\exists x_1.\varphi) \vee (\exists x_2.\psi)$

(ii) $\exists x.(\varphi \circ \chi) \ \models\mid\ (\exists x.\varphi) \circ \chi$ $\qquad$ with $\circ \in \{\wedge, \vee\}$

(iii) $\forall x.(\varphi \wedge \psi) \ \models\mid\ (\forall x_1.\varphi) \wedge (\forall x_2.\psi)$

(iv) $\forall x.(\varphi \circ \chi) \ \models\mid\ (\forall x.\varphi) \circ \chi$ $\qquad$ with $\circ \in \{\wedge, \vee\}$

Consequently, if $x_1 \notin \mathrm{vars}(\chi)$ and $x_2 \notin \mathrm{vars}(\varphi)$ holds for two first-order formulas φ and χ, we get $(\exists x_1.\varphi) \wedge (\exists x_2.\chi) \models\mid \exists x_1 x_2.(\varphi \wedge \chi)$ and dually $(\forall x_1.\varphi) \vee (\forall x_2.\chi) \models\mid \forall x_1 x_2.(\varphi \vee \chi)$.

3 Separated Variables and Transposition of Quantifier Blocks

Let φ be a first-order formula. We call two disjoint sets of variables $\vec{x}$ and $\vec{y}$ *separated in* φ if and only if for every atom A occurring in φ we have $\mathrm{vars}(A) \cap \vec{x} = \emptyset$ or $\mathrm{vars}(A) \cap \vec{y} = \emptyset$.

We first show how we can transpose quantifier blocks if the variables they bind are separated. Throughout this section we admit equality or other predicates with fixed semantics in the formulas we consider. Moreover, function symbols may occur, even in an arbitrarily nested fashion.

Proposition 3. Let $\varphi(\vec{x}, \vec{y}, \vec{z})$ be a quantifier-free first-order formula in which $\vec{x}$ and $\vec{y}$ are separated. There exists some $m \geq 1$ and a quantifier-free first-order formula $\varphi'(\vec{x}, \vec{y}_1, \ldots, \vec{y}_m, \vec{z})$ such that $\forall \vec{x} \exists \vec{y}.\varphi(\vec{x}, \vec{y}, \vec{z})$ and $\exists \vec{y}_1 \ldots \exists \vec{y}_m \forall \vec{x}.\varphi'(\vec{x}, \vec{y}_1, \ldots, \vec{y}_m, \vec{z})$ are semantically equivalent, and the length of each of the tuples $\vec{y}_k$ is identical to the length of $\vec{y}$. Moreover, all atoms in φ' are variants of atoms in φ, i.e. for every atom A' in φ' there is an atom A in φ that is identical to A' up to renaming of variables.

Proof. We first transform the matrix φ into a disjunction of conjunctions of literals. Since the sets $\vec{x}$ and $\vec{y}$ are separated in φ, the literals in every conjunction can be grouped into three parts:

(1) $\psi_i(\vec{x}, \vec{z})$, containing none of the variables from $\vec{y}$,

(2) $\chi_i(\vec{y}, \vec{z})$, containing none of the variables from $\vec{x}$, and

(3) $\eta_i(\vec{z})$, containing neither variables from $\vec{x}$ nor from $\vec{y}$.

Employing Lemma 2, we move the existential quantifier block $\exists \vec{y}$ inwards such that it only binds the (sub-)conjunctions $\chi_i(\vec{y}, \vec{z})$. The emerging subformulas $(\exists \vec{y}.\chi_i(\vec{y}, \vec{z}))$ shall be treated as indivisible units in the further process.

After moving the $\exists \vec{y}$ block inwards, the resulting formula has the form $\forall \vec{x}.\varphi''$. Next, we transform φ'' into a conjunction of disjunctions, and move the universal quantifier block $\forall \vec{x}$ inwards in a way analogous to the procedure described for the $\exists \vec{y}$ block. The shape of the result then allows to move the existential quantifiers outwards first (renaming variables where necessary) and the universal ones afterwards so that we again obtain a prenex formula, this time with an $\exists^* \forall^*$ prefix.

$$\forall \vec{x} \, \exists \vec{y}.\varphi(\vec{x}, \vec{y}, \vec{z})$$

$$\models\mid\ \forall \vec{x}.\exists \vec{y}. \bigvee_i \psi_i(\vec{x}, \vec{z}) \wedge \chi_i(\vec{y}, \vec{z}) \wedge \eta_i(\vec{z})$$

$$\models\mid\ \forall \vec{x}. \bigvee_i \psi_i(\vec{x}, \vec{z}) \wedge \left(\exists \vec{y}.\chi_i(\vec{y}, \vec{z}) \right) \wedge \eta_i(\vec{z})$$

$$\models\mid\ \forall \vec{x}. \bigwedge_{k=1}^{m} \psi'_k(\vec{x}, \vec{z}) \vee \bigvee_\ell \left(\exists \vec{y}.\chi'_{k,\ell}(\vec{y}, \vec{z}) \right) \vee \eta'_k(\vec{z})$$

$$\models \bigwedge_{k=1}^{m} \left(\forall \vec{\mathbf{x}}.\psi_k'(\vec{\mathbf{x}}, \vec{\mathbf{z}}) \right) \vee \left(\exists \vec{\mathbf{y}}. \bigvee_{\ell} \chi_{k,\ell}'(\vec{\mathbf{y}}, \vec{\mathbf{z}}) \right) \vee \eta_k'(\vec{\mathbf{z}})$$

$$\models \exists \vec{\mathbf{y}}_1 \ldots \exists \vec{\mathbf{y}}_m \forall \vec{\mathbf{x}}. \bigwedge_{k=1}^{m} \psi_k'(\vec{\mathbf{x}}, \vec{\mathbf{z}}) \vee \bigvee_{\ell} \chi_{k,\ell}'(\vec{\mathbf{y}}_k, \vec{\mathbf{z}}) \vee \eta_k'(\vec{\mathbf{z}})$$

The subformulas ψ_i, χ_i, η_i, and $\chi_{k,\ell}'$ are conjunctions of literals, and the ψ_k' and η_k' are disjunctions of literals. $\qquad\square$

The following example illustrates how quantifiers can be transposed in accordance with Proposition 3 and what the limits are.

Example 4. Consider the sentence $\forall x \exists y. P(x) \leftrightarrow Q(y)$. On the one hand, it is easy to see that direct transposition of $\forall x$ and $\exists y$ does not preserve semantics: whilst the structure $\mathcal{A}$ with $\mathfrak{U}_{\mathcal{A}} := \{\mathfrak{a}, \mathfrak{b}\}$ and $P^{\mathcal{A}} := Q^{\mathcal{A}} := \{\mathfrak{a}\}$ is a model of $\forall x \exists y. P(x) \leftrightarrow Q(y)$, it does not satisfy the version with transposed quantifiers $\exists y \forall x. P(x) \leftrightarrow Q(y)$.

On the other hand, we can show equivalence of $\forall x \exists y. P(x) \leftrightarrow Q(y)$ and $\exists y_1 y_2 \forall x. \big(P(x) \rightarrow Q(y_1) \big) \wedge \big(Q(y_2) \rightarrow P(x) \big)$:

$$\forall x \exists y. P(x) \leftrightarrow Q(y)$$
$$\models \forall x \exists y. \big(\neg P(x) \vee Q(y) \big) \wedge \big(P(x) \vee \neg Q(y) \big)$$
$$\models \forall x. \big(\neg P(x) \wedge (\exists y_2. \neg Q(y_2)) \big) \vee \big((\exists y_1. Q(y_1)) \wedge P(x) \big)$$
$$\models \big((\forall x. \neg P(x)) \vee (\exists y_1. Q(y_1)) \big)$$
$$\qquad\qquad \wedge \big((\exists y_2. \neg Q(y_2)) \vee (\forall x. P(x)) \big)$$
$$\models \exists y_1 y_2 \forall x. \big(\neg P(x) \vee Q(y_1) \big) \wedge \big(\neg Q(y_2) \vee P(x) \big)$$
$$\models \exists y_1 y_2 \forall x. \big(P(x) \rightarrow Q(y_1) \big) \wedge \big(Q(y_2) \rightarrow P(x) \big) \ .$$

Next, we demonstrate how the transposition of quantifiers affects the length of formulas.

Proposition 5. In the worst case transposing quantifier blocks in accordance with Proposition 3 leads to a blow-up in the number of existentially quantified variables that is exponential in the length of the original formula.

Proof. Consider the following first-order sentence and how we can transpose the quantifier blocks therein. $\varphi := \forall x \exists y. (P_1(x) \leftrightarrow Q_1(y)) \wedge \ldots \wedge (P_n(x) \leftrightarrow Q_n(y))$ can be transformed into the equivalent $\varphi' :=$

$$\exists \underbrace{y_{\langle 0,\ldots,0 \rangle} \cdots y_{\langle 1,\ldots,1 \rangle}}_{2^n \text{ variables}} \forall x.$$

$$\bigwedge_{\bar{b} \in \{0,1\}^n} \left(\left(\bigwedge_{\substack{1 \le i \le n, \\ b_i = 1}} P_i(x) \quad \wedge \bigwedge_{\substack{1 \le j \le n, \\ b_j = 0}} \neg P_j(x) \right) \right.$$

$$\left. \longrightarrow \left(\bigwedge_{\substack{1 \le i \le n, \\ b_i = 1}} Q_i(y_{\bar{b}}) \quad \wedge \bigwedge_{\substack{1 \le j \le n, \\ b_j = 0}} \neg Q_j(y_{\bar{b}}) \right) \right)$$

Consider the following model $\mathcal{A}$ of φ and φ':

$$\mathfrak{U}_{\mathcal{A}} := \big\{ \mathfrak{a}_{\bar{b}} \mid \bar{b} \in \{0,1\}^n \big\} \cup \big\{ \mathfrak{a}_{\bar{b}}' \mid \bar{b} \in \{0,1\}^n \big\} \ ,$$
$$P_i^{\mathcal{A}} := \big\{ \mathfrak{a}_{\langle b_1,\ldots,b_n \rangle} \mid b_i = 1 \big\} \qquad \text{for } i = 1, \ldots, n \ ,$$
$$Q_i^{\mathcal{A}} := \big\{ \mathfrak{a}_{\langle b_1,\ldots,b_n \rangle}' \mid b_i = 1 \big\} \qquad \text{for } i = 1, \ldots, n \ .$$

We make the following observation and shall keep it in mind: $(*)$ removing any of the $\mathfrak{a}'_{\bar{b}}$ with $\bar{b} \neq \langle 0, \ldots, 0 \rangle$ from $\mathcal{A}$ does not lead to a model of φ. We will argue that any sentence φ_* (in prenex normal form), which is semantically equivalent to φ and starts with a quantifier prefix of the form $\exists^* \forall^*$, contains at least $2^n - 1$ existential quantifiers.

Let $\varphi_* := \exists y_1 \ldots y_k \forall x_1 \ldots x_\ell . \chi_*$ (with χ_* being quantifier-free) be a sentence with minimal k that is semantically equivalent to φ. Since $\mathcal{A}$ is also a model of φ_*, we know that there is a sequence of elements $\mathfrak{c}_1, \ldots, \mathfrak{c}_k$ taken from the domain $\mathfrak{U}_\mathcal{A}$ such that $\mathcal{A}, [y_1 \mapsto \mathfrak{c}_1, \ldots, y_k \mapsto \mathfrak{c}_k] \models \forall x_1 \ldots x_\ell . \chi_*$. Consequently, we can extend $\mathcal{A}$ to a model $\mathcal{A}_*$ (over the same domain) of the Skolemized formula $\varphi_{\mathrm{Sk}} := \forall x_1 \ldots x_\ell . \chi_* [y_1/c_1, \ldots, y_k/c_k]$ by adding $c_j^{\mathcal{A}_*} := \mathfrak{c}_j$ for $j = 1, \ldots, k$. On the other hand, every model of the Skolemized formula φ_{Sk} immediately yields a model of φ_*.

The signature underlying φ_{Sk} comprises exactly the constant symbols $c_1, \ldots, c_k$ and does not contain any other function symbols. Suppose $k < 2^n - 1$. Hence, there is some $\mathfrak{a}'_{\bar{b}}$ with $\bar{b} \neq \langle 0, \ldots, 0 \rangle$ such that for every j it holds $c_j^{\mathcal{A}_*} \neq \mathfrak{a}'_{\bar{b}}$. By the Substructure Lemma, the following substructure $\mathcal{B}$ of $\mathcal{A}_*$ constitutes a model of φ_{Sk}: $\mathfrak{U}_\mathcal{B} := \mathfrak{U}_{\mathcal{A}_*} \setminus \{\mathfrak{a}'_{\bar{b}}\}$, $P_i^\mathcal{B} := P_i^{\mathcal{A}_*} \cap \mathfrak{U}_\mathcal{B}$ and $Q_i^\mathcal{B} := Q_i^{\mathcal{A}_*} \cap \mathfrak{U}_\mathcal{B}$ for every i, and $c_j^\mathcal{B} := c_j^{\mathcal{A}_*}$ for every j.

However, then $\mathcal{B}$ must also be a model of both φ_* and φ, since every model of φ_{Sk} is a model of φ_*, and because we assumed φ_* and φ to be semantically equivalent. This contradicts observation $(*)$, and thus we must have $k \geq 2^n - 1$.

This evidently shows that φ' is (almost) optimal regarding the number of existentially quantified variables it contains. $\qquad\square$

It is possible to generalize Proposition 3 to the case of several quantifier alternations as long as all universally quantified variables are separated from all existentially quantified ones.

Lemma 6. Let $\varphi(\vec{\mathbf{x}}_1, \ldots, \vec{\mathbf{x}}_n, \vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_n, \vec{\mathbf{z}})$ be a quantifier-free first-order formula in which the sets $\vec{\mathbf{x}}_1 \cup \ldots \cup \vec{\mathbf{x}}_n$ and $\vec{\mathbf{y}}_1 \cup \ldots \cup \vec{\mathbf{y}}_n$ are separated. There exists a quantifier-free first-order formula $\varphi'(\vec{\mathbf{u}}, \vec{\mathbf{v}}, \vec{\mathbf{z}})$ such that $\forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n . \varphi(\vec{\mathbf{x}}_1, \ldots, \vec{\mathbf{x}}_n, \vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_n, \vec{\mathbf{z}})$ and $\exists \vec{\mathbf{u}} \forall \vec{\mathbf{v}} . \varphi'(\vec{\mathbf{u}}, \vec{\mathbf{v}}, \vec{\mathbf{z}})$ are semantically equivalent and all atoms in φ' are variants of atoms in φ. Notice that $\vec{\mathbf{x}}_1$ and $\vec{\mathbf{y}}_n$ may be empty.

Proof. We apply a syntactic transformation following the strategy of the proof of Proposition 3, but in an iterated fashion:

$$\forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n . \varphi(\vec{\mathbf{x}}_1, \ldots, \vec{\mathbf{x}}_n, \vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_n, \vec{\mathbf{z}})$$

$$\models \forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n . \bigvee_i \psi_i(\vec{\mathbf{x}}_1, \ldots, \vec{\mathbf{x}}_n, \vec{\mathbf{z}})$$
$$\wedge \chi_i(\vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_n, \vec{\mathbf{z}}) \wedge \eta_i(\vec{\mathbf{z}})$$

$$\models \forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \forall \vec{\mathbf{x}}_n . \bigvee_i \psi_i(\vec{\mathbf{x}}_1, \ldots, \vec{\mathbf{x}}_n, \vec{\mathbf{z}})$$
$$\wedge \underbrace{\left(\exists \vec{\mathbf{y}}_n . \chi_i(\vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_n, \vec{\mathbf{z}}) \right)}_{=: \chi_i^{(1)}(\vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_{n-1}, \vec{\mathbf{z}})} \wedge \eta_i(\vec{\mathbf{z}})$$

$$\models \forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \forall \vec{\mathbf{x}}_n . \bigwedge_k \psi'_k(\vec{\mathbf{x}}_1, \ldots, \vec{\mathbf{x}}_n, \vec{\mathbf{z}})$$
$$\vee \chi_k'^{(1)}(\vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_{n-1}, \vec{\mathbf{z}}) \vee \eta'_k(\vec{\mathbf{z}})$$

$$\models \forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \exists \vec{\mathbf{y}}_{n-1} . \bigwedge_k \underbrace{\left(\forall \vec{\mathbf{x}}_n . \psi'_k(\vec{\mathbf{x}}_1, \ldots, \vec{\mathbf{x}}_n, \vec{\mathbf{z}}) \right)}_{=: \psi_k^{(1)}(\vec{\mathbf{x}}_1, \ldots, \vec{\mathbf{x}}_{n-1}, \vec{\mathbf{z}})}$$
$$\vee \chi_k'^{(1)}(\vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_{n-1}, \vec{\mathbf{z}}) \vee \eta'_k(\vec{\mathbf{z}})$$

$$\models \forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \exists \vec{\mathbf{y}}_{n-1} . \bigvee_i \psi_i'^{(1)}(\vec{\mathbf{x}}_1, \ldots, \vec{\mathbf{x}}_{n-1}, \vec{\mathbf{z}})$$
$$\wedge \chi_i''^{(1)}(\vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_{n-1}, \vec{\mathbf{z}}) \wedge \eta_i''(\vec{\mathbf{z}})$$

5

$$\models \forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \forall \vec{\mathbf{x}}_{n-1} . \bigvee_i \psi_i'^{(1)}(\vec{\mathbf{x}}_1, \ldots, \vec{\mathbf{x}}_{n-1}, \vec{\mathbf{z}})$$
$$\wedge \underbrace{\left(\exists \vec{\mathbf{y}}_{n-1} . \chi_i''^{(1)}(\vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_{n-1}, \vec{\mathbf{z}}) \right)}_{=: \, \chi_i^{(2)}(\vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_{n-2}, \vec{\mathbf{z}})} \wedge \eta_i''(\vec{\mathbf{z}})$$
$$\vdots$$
$$\models \forall \vec{\mathbf{x}}_1 . \bigwedge_k \psi_k''^{(n-1)}(\vec{\mathbf{x}}_1, \vec{\mathbf{z}}) \vee \chi_k'^{(n)}(\vec{\mathbf{z}}) \vee \eta_k^*(\vec{\mathbf{z}})$$
$$\models \bigwedge_k \underbrace{\left(\forall \vec{\mathbf{x}}_1 . \psi_k''^{(n-1)}(\vec{\mathbf{x}}_1, \vec{\mathbf{z}}) \right)}_{=: \, \psi_k^{(n)}(\vec{\mathbf{z}})} \vee \chi_k'^{(n)}(\vec{\mathbf{z}}) \vee \eta_k^*(\vec{\mathbf{z}})$$
$$\models \exists \vec{\mathbf{u}} \forall \vec{\mathbf{v}} . \bigwedge_k \widehat{\psi}_k^{(n)}(\vec{\mathbf{u}}, \vec{\mathbf{z}}) \vee \widehat{\chi}_k^{(n)}(\vec{\mathbf{v}}, \vec{\mathbf{z}}) \vee \eta_k^*(\vec{\mathbf{z}}) \ .$$

The $\chi_k'^{(\ell)}$ are disjunctions of subformulas $\chi_j^{(\ell)}$, whereas the $\chi_k''^{(\ell)}$ are conjunctions of such formulas. Dually, the $\psi_k'^{(\ell)}$ are conjunctions of subformulas $\psi_j^{(\ell)}$ and the $\psi_k''^{(\ell)}$ are disjunctions of such formulas. The $\widehat{\psi}_k^{(n)}(\vec{\mathbf{u}}, \vec{\mathbf{z}})$ and the $\widehat{\chi}_k^{(n)}(\vec{\mathbf{v}}, \vec{\mathbf{z}})$ are quantifier-free variants of $\psi_k^{(n)}(\vec{\mathbf{z}})$ and $\chi_k'^{(n)}(\vec{\mathbf{z}})$, respectively, i.e. quantifiers have been moved outwards and variables have been renamed appropriately. □

Notice that every transformation into a disjunction of conjunctions or a conjunction of disjunctions in the above proof causes at most an exponential increase in the length of the formula. Moving quantifier blocks inwards causes at most a factor of two per moved block. Roughly speaking, the length of φ' is thus at most $2n$-fold exponential in the double length of φ. This constitutes a very crude upper bound, and it does not take into account that redundant subformulas may be removed. The complexity results in Section 4.1 will be based on semantic arguments, and thus do not rely on this rough estimate.

4 The Separated Fragment of First-Order Logic

We next explore a striking consequence of Lemma 6. It is well-known that the Bernays–Schönfinkel–Ramsey (BSR) Fragment is decidable. This fragment comprises all first-order sentences in prenex normal form with quantifier prefixes of the form $\exists^* \forall^*$. Moreover, equality may occur in BSR sentences but non-constant function symbols may not. By virtue of Lemma 6, we can now extend this decidability result to the following class of first-order sentences.

Definition 7 (Separated Fragment (SF)). The *Separated Fragment (SF)* of first-order logic shall consist of all existential closures of prenex formulas in which existentially quantified variables are separated from universally quantified ones.

More precisely, it consists of all first-order sentences with equality but without non-constant function symbols that are of the form $\exists \vec{\mathbf{z}} \forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n . \varphi(\vec{\mathbf{x}}_1, \ldots, \vec{\mathbf{x}}_n, \vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_n, \vec{\mathbf{z}})$ in which φ is quantifier-free and the sets $\vec{\mathbf{x}}_1 \cup \ldots \cup \vec{\mathbf{x}}_n$ and $\vec{\mathbf{y}}_1 \cup \ldots \cup \vec{\mathbf{y}}_n$ are separated.

As already stated, Lemma 6 shows that every sentence in SF can be transformed into an equivalent one which belongs to the BSR Fragment: just replace the subformula $\forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n . \varphi$ with an equivalent one of the form $\exists \vec{\mathbf{u}} \forall \vec{\mathbf{v}} . \varphi'$. This proves our main result:

Theorem 8. Satisfiability of sentences in SF is decidable.

It is worth noticing that SF is not only a proper superset of the BSR Fragment but also of one more well-known decidable first-order fragment, namely the Relational Monadic Fragment *without equality*, i.e. the set of first-order sentences without non-constant function symbols in which only predicate symbols of arity one are allowed.

Theorem 9. SF properly contains the BSR fragment and the Relational Monadic Fragment without equality.

Proof. Let $\varphi := \exists \vec{z} \forall \vec{x}.\psi$ be a BSR sentence, i.e. it may contain equality but non-constant function symbols may not appear. Moreover, assume that ψ is quantifier free. The sentence φ may be considered as the existential closure of the formula $\forall \vec{x}.\psi$. Since $\forall \vec{x}.\psi$ does not contain any existentially quantified variables, the separation criterion in Definition 7 is trivially fulfilled. Hence, φ lies in SF.

Let φ' be a relational monadic sentence without equality and without non-constant function symbols. Since all predicate symbols in φ' have an arity of at most one, any two disjoint sets of variables are trivially separated in φ'. Therefore, φ' belongs to SF. □

In Section 4.2 we show how SF can be extended so that the Relational Monadic Fragment *with equality* and also the Full Monadic Fragment *without equality* (i.e. Relational Monadic plus unary function symbols) become proper subsets of the extension.

4.1 Range-Restricted Skolemization

In this section we shall demonstrate another approach to showing decidability of SF. This approach emphasizes the *small model property* of SF, i.e. we can give a computable function h_{SF} which takes any SF sentence φ as input and yields a positive integer $h_{\mathrm{SF}}(\varphi)$ such that whenever φ has a model then it also has a model based on a universe with at most $h_{\mathrm{SF}}(\varphi)$ elements.

It is well-known that BSR sentences exhibit the small model property, where the size of small models is linear in the number of occurring constant symbols plus the number of occurrences of existential quantifiers in the sentence at hand. Hence, the number of constant symbols and existentially quantified variables in any BSR sentence φ' that is semantically equivalent to an SF sentence φ would yield an upper bound on the size of small models of φ. However, the transformations carried out in the proof of Lemma 6 do not immediately admit a reasonable estimate on the number of variables in the resulting BSR sentence. This is why we tackle the problem in a different way. As it turns out, this alternative approach does not only facilitate the derivation of tighter upper bounds on the size of small models in subfragments of SF. In addition, automated reasoning procedures may benefit from the developed methods and results. We shall assess this potential in Section 5.

On an abstract level, our semantic approach is akin to proofs of the small model property of relational monadic sentences. Usually, the central argument goes as follows: any sentence φ without equality and without non-constant function symbols which contains exactly k predicate symbols $P_1, \ldots, P_k$—all of them unary—cannot distinguish more than 2^k domain elements. To formalize this, we associate with every domain element $\mathfrak{a} \in \mathfrak{U}_{\mathcal{A}}$ of a given structure $\mathcal{A}$ a *fingerprint* with respect to the predicates $P_1^{\mathcal{A}}, \ldots, P_k^{\mathcal{A}}$, namely the set $\lambda(\mathfrak{a}) := \{i \mid \mathfrak{a} \in P_i^{\mathcal{A}}\}$. Hence, $\lambda(\mathfrak{a})$ contains exactly the indices i, $1 \leq i \leq k$, for which $\mathfrak{a}$ belongs to $\mathcal{A}$'s interpretation of P_i. In a certain sense, this fingerprint of a domain element is all that matters for a relational monadic sentence φ under $\mathcal{A}$, i.e., if two elements $\mathfrak{a}, \mathfrak{b}$ have the same fingerprint $\lambda(\mathfrak{a}) = \lambda(\mathfrak{b})$, then φ cannot distinguish the two. As a consequence, given a model of φ, domain elements with identical fingerprints can be merged, and the resulting structure is still a model of φ. Since there are at most 2^k distinct fingerprints, this entails the small model property for relational monadic sentences.

In order to treat SF sentences, we need to modify the just described idea of fingerprints in several ways:

(a) Since the definition of SF does not pose any restriction on the arity of predicate symbols, we have to generalize the idea of a fingerprint from single elements to tuples of elements.

(b) We shall concentrate on the parts of sentences which contain existentially quantified variables—by definition of SF these can be isolated from the ones containing universally quantified variables. The rationale behind this approach is rooted in the idea underlying the Substructure Lemma, namely that only the part of the domain is of interest, which is generated by the interpretations of function symbols. And since in the SF setting non-constant function symbols are only introduced by Skolemization, we focus on existentially quantified variables.

7

(c) Instead of regarding the membership of a tuple $\vec{\mathfrak{a}}$ in predicates $P^{\mathcal{A}} \subseteq \mathfrak{U}_{\mathcal{A}}^m$ as the characteristic feature to define fingerprints, we consider whether $\vec{\mathfrak{a}}$ satisfies certain subformulas of normal forms of φ. More precisely, we refer to the *disjunctive normal form (DNF)* $\bigvee_i \psi_i$ or the *conjunctive normal form (CNF)* $\bigwedge_j \psi_j'$ and ask the question whether the part $\eta_i(\vec{y})$ of a conjunction (or disjunction) $\psi_i = \chi_i(\vec{x}) \wedge \eta_i(\vec{y})$ is true under $\mathcal{A}, [\vec{y} \mapsto \vec{\mathfrak{a}}]$ or not. If (and only if) it is, then the index i belongs to the fingerprint of $\vec{\mathfrak{a}}$.

(d) In general settings with several quantifier alternations we cannot only rely on a single fingerprint function, but rather construct one for every existential quantifier block occurring in the quantifier prefix. The fingerprints associated with the earlier quantifier blocks will be nested in the sense that they comprise all the potential fingerprints that may be produced starting from the current point. For example, consider the sentence $\forall x_1 \exists y_1 \forall x_2 \exists y_2. \big(Q_1(x_1, x_2) \wedge R_1(y_1, y_2)\big) \vee \big(Q_2(x_1, x_2) \wedge R_2(y_1, y_2)\big)$ and the structure $\mathcal{A}$ with $\mathfrak{U}_{\mathcal{A}} := \{\mathfrak{a}, \mathfrak{a}', \mathfrak{b}\}$ and $R_1^{\mathcal{A}} := \{\langle \mathfrak{a}, \mathfrak{b} \rangle, \langle \mathfrak{a}', \mathfrak{a} \rangle, \langle \mathfrak{a}', \mathfrak{a}' \rangle, \langle \mathfrak{b}, \mathfrak{b} \rangle\}$, $R_2^{\mathcal{A}} := \{\langle \mathfrak{a}, \mathfrak{a} \rangle, \langle \mathfrak{a}, \mathfrak{a}' \rangle, \langle \mathfrak{a}, \mathfrak{b} \rangle, \langle \mathfrak{a}', \mathfrak{a} \rangle, \langle \mathfrak{a}', \mathfrak{a}' \rangle, \langle \mathfrak{a}', \mathfrak{b} \rangle, \langle \mathfrak{b}, \mathfrak{a} \rangle\}$. ($Q_1^{\mathcal{A}}$ and $Q_2^{\mathcal{A}}$ may be defined arbitrarily.) First, we define a fingerprint function λ_2 for pairs of elements from $\mathfrak{U}_{\mathcal{A}}$ such that for every pair $\langle \mathfrak{c}, \mathfrak{d} \rangle \in \mathfrak{U}_{\mathcal{A}}^2$ and every index $i = 1, 2$ we have $i \in \lambda_2(\mathfrak{c}, \mathfrak{d})$ if and only if $\mathcal{A}, [y_1 \mapsto \mathfrak{c}, y_2 \mapsto \mathfrak{d}] \models R_i(y_1, y_2)$. Concretely, λ_2 assigns fingerprints as follows:

λ_2	$\mathfrak{a}$	$\mathfrak{a}'$	$\mathfrak{b}$
$\mathfrak{a}$	$\{2\}$	$\{2\}$	$\{1, 2\}$
$\mathfrak{a}'$	$\{1, 2\}$	$\{1, 2\}$	$\{2\}$
$\mathfrak{b}$	$\{2\}$	$\emptyset$	$\{1\}$

Based on λ_2 we next define the fingerprint function λ_1 for single domain elements such that for every element $\mathfrak{c} \in \mathfrak{U}_{\mathcal{A}}$ and every λ_2-fingerprint $S \subseteq \{1, 2\}$ it holds $S \in \lambda_1(\mathfrak{c})$ if and only if there is another element $\mathfrak{d} \in \mathfrak{U}_{\mathcal{A}}$ such that $S = \lambda_2(\mathfrak{c}, \mathfrak{d})$. In the above example, this means $\lambda_1(\mathfrak{a}) = \lambda_1(\mathfrak{a}') = \big\{\{2\}, \{1, 2\}\big\}$, and $\lambda_1(\mathfrak{b}) = \big\{\emptyset, \{1\}, \{2\}\big\}$. We will see later that the fact that λ_2 assigns the same fingerprint to $\mathfrak{a}$ and $\mathfrak{a}'$ entails that the quantifier $\exists y_1$ does not have to take both $\mathfrak{a}$ and $\mathfrak{a}'$ into account. It suffices to consider only one of them. Hence, after Skolemizing $\exists y_1$, we may restrict the range of the Skolem term $f_{y_1}(x_1)$ under $\mathcal{A}$ to $\{\mathfrak{a}, \mathfrak{b}\}$ or $\{\mathfrak{a}', \mathfrak{b}\}$, but we do not have to consider the full range $\{\mathfrak{a}, \mathfrak{a}', \mathfrak{b}\}$. This is what we call *range-restricted Skolemization* (cf. Lemma 12 and similar results in Section 4.1.2).

In what follows we shall use the notation $[k]$ to abbreviate the set $\{1, \ldots, k\}$ for any positive integer k. Moreover, $\mathfrak{P}$ shall be used as the power set operator, i.e. $\mathfrak{P}S$ denotes the set of all subsets of a given set S.

4.1.1 Sentences with prefix $\exists^* \forall^* \exists^*$

We develop the following for sentences with the $\forall^* \exists^*$ quantifier prefix, because leading existential quantifiers may be replaced by constant symbols (under preservation of satisfiability). Let $\psi(\vec{x}, \vec{y})$ be a quantifier-free first-order formula without non-constant function symbols in which $\vec{x}$ and $\vec{y}$ are separated. We can transform $\varphi := \forall \vec{x} \exists \vec{y}.\psi(\vec{x}, \vec{y})$ into an equivalent sentence $\varphi_{\mathrm{DNF}} := \forall \vec{x} \exists \vec{y}. \bigvee_{k=1}^{m_{\mathrm{DNF}}} \chi_k(\vec{x}) \wedge \eta_k(\vec{y})$, such that χ_k and η_k are conjunctions of literals and m_{DNF} is some non-negative integer.

Intuitively speaking, we combine two complementary fingerprint functions in this setting, where the second one will only appear implicitly. Given a structure $\mathcal{A}$, for any tuple $\vec{\mathfrak{b}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{x}|}$ of domain elements assigned to $\vec{x}$ it is only of importance which subformulas $\chi_k(\vec{x})$ it fulfills. Hence, the fingerprint function $\lambda(\vec{\mathfrak{b}}) := \big\{k \in [m_{\mathrm{DNF}}] \mid \mathcal{A}, [\vec{x} \mapsto \vec{\mathfrak{b}}] \models \chi_k(\vec{x})\big\}$ is a reasonable choice. Complementary to every $\vec{\mathfrak{b}}$ a tuple $\vec{\mathfrak{a}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{y}|}$ is necessary, for which $\mathcal{A}, [\vec{y} \mapsto \vec{\mathfrak{a}}] \models \eta_k(\vec{y})$ holds for at least one $k \in \lambda(\vec{\mathfrak{b}})$. Note that empty η_k are treated as the Boolean constant *true*. Together, two such matching tuples satisfy the formula in the structure $\mathcal{A}$. In order to formally classify tuples $\vec{\mathfrak{a}}$ in accordance with the fingerprint feature just described, we define the sets $\widehat{\mathfrak{U}}_1, \ldots, \widehat{\mathfrak{U}}_{m_{\mathrm{DNF}}} \subseteq \mathfrak{U}_{\mathcal{A}}^{|\vec{y}|}$ by $\widehat{\mathfrak{U}}_k := \big\{\vec{\mathfrak{a}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{y}|} \mid \mathcal{A}, [\vec{y} \mapsto \vec{\mathfrak{a}}] \models \eta_k\big\}$ for $k = 1, \ldots, m_{\mathrm{DNF}}$. Invoking the axiom of choice, we pick one representative $\alpha_k \in \widehat{\mathfrak{U}}_k$ for every nonempty $\widehat{\mathfrak{U}}_k$ and fix it. In addition, we define a mapping

$\mu : \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{x}}|} \to [m_{\mathrm{DNF}}]$ such that $\mu(\vec{\mathfrak{b}})$ yields some fixed index $k \in \lambda(\vec{\mathfrak{b}})$, for which $\widehat{\mathfrak{U}}_k$ is nonempty. If no such k exists for $\vec{\mathfrak{b}}$, $\mu(\vec{\mathfrak{b}})$ shall be undefined. If $\mathcal{A} \models \forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})$, then for every $\vec{\mathfrak{b}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{x}}|}$ the value $\mu(\vec{\mathfrak{b}})$ must be defined, and we have $\mathcal{A}, [\vec{\mathbf{x}} \mapsto \vec{\mathfrak{b}}, \vec{\mathbf{y}} \mapsto \alpha_{\mu(\vec{\mathfrak{b}})}] \models \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})$, by definition of the representative $\alpha_{\mu(\vec{\mathfrak{b}})}$. This means, no matter which values are assigned to the variables in $\vec{\mathbf{x}}$, for the valuation of the variables in $\vec{\mathbf{y}}$ it is sufficient to consider exclusively values from $\{\alpha_k \mid k \in [m_{\mathrm{DNF}}]\}$. Hence, we can restrict the choice for the $\vec{\mathbf{y}}$ to the representatives $\alpha_1, \ldots, \alpha_{m_{\mathrm{DNF}}}$. Therefore, the sentences $\forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})$ and $\forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}) \wedge \bigvee_{\ell=1}^{m_{\mathrm{DNF}}} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} y_i \approx c_{\ell,i}$ are equisatisfiable, where the $c_{\ell,i}$ are fresh constant symbols. If there is a model at all, then there is one, under which the tuples of constant symbols $\vec{c}_1, \ldots, \vec{c}_{m_{\mathrm{DNF}}}$ are interpreted by the representatives $\alpha_1, \ldots, \alpha_{m_{\mathrm{DNF}}}$, respectively.

Remark 10. For this simple setting there is also a shorter argument leading to an even stronger result. We have used the more complicated one, however, since it is easily applicable in the more general case in Section 4.1.2, where the simple argument does not work in a straight-forward fashion.

Due to Lemma 2, $\varphi_{\mathrm{DNF}} = \forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}. \bigvee_{k=1}^{m_{\mathrm{DNF}}} \chi_k(\vec{\mathbf{x}}) \wedge \eta_k(\vec{\mathbf{y}})$ is equivalent to $\forall \vec{\mathbf{x}}. \bigvee_{k=1}^{m_{\mathrm{DNF}}} \chi_k(\vec{\mathbf{x}}) \wedge \exists \vec{\mathbf{y}}_k.\eta_k(\vec{\mathbf{y}}_k)$. Now *inner Skolemization* (cf. [27]) leads to $\forall \vec{\mathbf{x}}. \bigvee_{k=1}^{m_{\mathrm{DNF}}} \chi_k(\vec{\mathbf{x}}) \wedge \eta_k(\vec{\mathbf{y}}_k)[\vec{\mathbf{y}}_k/\vec{\mathfrak{c}}_k]$. This immediately entails equisatisfiability of the latter sentence and φ_{DNF}.

Dually, we can transform the sentence $\varphi = \forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})$ into an equivalent one $\varphi_{\mathrm{CNF}} := \forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}. \bigwedge_{j=1}^{m_{\mathrm{CNF}}} \chi_j'(\vec{\mathbf{x}}) \vee \eta_j'(\vec{\mathbf{y}})$ in which the χ_j' and η_j' denote disjunctions of literals and m_{CNF} is some non-negative integer. Given a structure $\mathcal{A}$, we define the fingerprint function $\lambda : \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}|} \to \mathfrak{P}[m_{\mathrm{CNF}}]$ by setting $\lambda(\vec{\mathfrak{a}}) := \{j \mid \mathcal{A}, [\vec{\mathbf{y}} \mapsto \vec{\mathfrak{a}}] \models \eta_j'(\vec{\mathbf{y}})\}$. The fingerprints assigned by λ induce a partition of $\mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}|}$ into at most $2^{m_{\mathrm{CNF}}}$ equivalence classes: two tuples $\vec{\mathfrak{a}}, \vec{\mathfrak{a}}'$ are equivalent if and only if they have the same fingerprint $\lambda(\vec{\mathfrak{a}}) = \lambda(\vec{\mathfrak{a}}')$. For every fingerprint $S \subseteq [m_{\mathrm{CNF}}]$, for which some tuple $\vec{\mathfrak{a}}$ with $\lambda(\vec{\mathfrak{a}}) = S$ exists, we fix a representative $\alpha_S \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}|}$ of $\vec{\mathfrak{a}}$'s equivalence class, i.e. $\lambda(\alpha_{\lambda(\vec{\mathfrak{a}})}) = \lambda(\vec{\mathfrak{a}})$.

Clearly, two tuples $\vec{\mathfrak{a}}, \vec{\mathfrak{a}}'$ are indistinguishable by the formulas $\eta_j'(\vec{\mathbf{y}})$, if they are associated with the same fingerprint. Analogous to our previous result, we could immediately derive equisatisfiability of φ and the formula $\forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}) \wedge \bigvee_{j=1}^{2^{m_{\mathrm{CNF}}}} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} y_i \approx c_{j,i}$, where the $c_{j,i}$ are fresh.

However, it turns out that being indistinguishable is more than we need. In fact, even in the worst case we do not need to consider $2^{m_{\mathrm{CNF}}}$ different representatives. A tuple $\vec{\mathfrak{a}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}|}$ *covers* a tuple $\vec{\mathfrak{a}}' \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}|}$, denoted $\vec{\mathfrak{a}} \sqsupseteq \vec{\mathfrak{a}}'$, if and only if $\lambda(\vec{\mathfrak{a}}') \subseteq \lambda(\vec{\mathfrak{a}})$. Of two representatives α, α' with $\alpha \sqsupseteq \alpha'$ we do actually only need one, namely α. More formally speaking, it is sufficient to partition $\mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}|}$ into parts $\widetilde{\mathfrak{U}}_1, \ldots, \widetilde{\mathfrak{U}}_{\kappa_{\mathrm{CNF}}}$ such that (i) in every part $\widetilde{\mathfrak{U}}_\ell$ for all distinct $\vec{\mathfrak{a}}, \vec{\mathfrak{a}}' \in \widetilde{\mathfrak{U}}_\ell$ either $\vec{\mathfrak{a}} \sqsupseteq \vec{\mathfrak{a}}'$ or $\vec{\mathfrak{a}}' \sqsupseteq \vec{\mathfrak{a}}$, and (ii) for all distinct parts $\widetilde{\mathfrak{U}}_\ell, \widetilde{\mathfrak{U}}_{\ell'}$ all tuples $\vec{\mathfrak{a}} \in \widetilde{\mathfrak{U}}_\ell$ and $\vec{\mathfrak{a}}' \in \widetilde{\mathfrak{U}}_{\ell'}$ are pairwise non-covering, i.e. we neither have $\vec{\mathfrak{a}} \sqsupseteq \vec{\mathfrak{a}}'$ nor $\vec{\mathfrak{a}}' \sqsupseteq \vec{\mathfrak{a}}$. But now, we have to choose the representatives more carefully: a representative $\widetilde{\alpha}_S \in \widetilde{\mathfrak{U}}_\ell$ has to cover all tuples $\vec{\mathfrak{a}} \in \widetilde{\mathfrak{U}}_\ell$. Formulated differently, we have $\widetilde{\alpha}_{\lambda(\vec{\mathfrak{a}})} \sqsupseteq \vec{\mathfrak{a}}$ for every $\vec{\mathfrak{a}}$. Putting things together, we observe for all tuples $\vec{\mathfrak{b}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{x}}|}$ and $\vec{\mathfrak{a}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}|}$ that $\mathcal{A}, [\vec{\mathbf{x}} \mapsto \vec{\mathfrak{b}}, \vec{\mathbf{y}} \mapsto \vec{\mathfrak{a}}] \models \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})$ entails $\mathcal{A}, [\vec{\mathbf{x}} \mapsto \vec{\mathfrak{b}}, \vec{\mathbf{y}} \mapsto \widetilde{\alpha}_{\lambda(\vec{\mathfrak{a}})}] \models \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})$.

It remains to pinpoint the value of κ_{CNF}. Put differently, we need to determine the maximal number of pairwise non-inclusive fingerprints. The answer is provided by Sperner's Theorem.

Theorem 11 (Sperner's Theorem [31])**.** Let m be a positive integer. Consider the lattice formed by the partially ordered set $\langle \mathfrak{P}[m], \subseteq \rangle$. There are $\kappa = \binom{m}{\lfloor m/2 \rfloor}$ sets $M_1, \ldots, M_\kappa \subseteq [m]$ such that for all distinct i, j we have $M_i \not\subseteq M_j$. There are no more than κ such sets.

As a consequence, $\forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})$ is equisatisfiable to $\forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}. \ \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}) \wedge \bigvee_{j=1}^{\kappa_{\mathrm{CNF}}} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} y_i \approx c_{j,i}$, where $\kappa_{\mathrm{CNF}} := \binom{m_{\mathrm{CNF}}}{\lfloor m_{\mathrm{CNF}}/2 \rfloor}$ and the $c_{j,i}$ are fresh.

Together with the dual case, we observe that it is possible to restrict the range of the quantifier block $\exists \vec{\mathbf{y}}$ even further to $m_* := \min(\kappa_{\mathrm{CNF}}, m_{\mathrm{DNF}})$ different tuples of domain elements (preserving satisfiability). After Skolemization these restrictions apply to the freshly introduced Skolem functions and affect their range. These results are summarized in the following lemma.

9

Lemma 12 (Range-restricted Skolemization). The following sentences are pairwise equisatisfiable:

(1) $\forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})$,

(2) $\forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}) \wedge \bigvee_{\ell=1}^{m_*} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} y_i \approx c_{\ell,i}$,

(3) $\forall \vec{\mathbf{x}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})\big[y_1/f_1(\vec{\mathbf{x}}), \ldots, y_{|\vec{\mathbf{y}}|}/f_{|\vec{\mathbf{y}}|}(\vec{\mathbf{x}})\big]$
$\quad \wedge \bigvee_{\ell=1}^{m_*} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} f_i(\vec{\mathbf{x}}) \approx c_{\ell,i}$,

where $m_* = \min\left(m_{\mathrm{DNF}}, \binom{m_{\mathrm{CNF}}}{\lfloor m_{\mathrm{CNF}}/2\rfloor}\right)$ and the $c_{\ell,i}$ are fresh constant symbols and the f_i are Skolem functions of appropriate arity.

Please note that the just stated result nicely fits together with Proposition 3, which states the equivalence of $\forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})$ and some sentence $\exists \vec{\mathbf{y}}_1 \ldots \vec{\mathbf{y}}_m \forall \vec{\mathbf{x}}.\psi'(\vec{\mathbf{x}}, \vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_m)$, which results in $\forall \vec{\mathbf{x}}.\psi'(\vec{\mathbf{x}}, \vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_m)\big[\vec{\mathbf{y}}_1/\vec{c}_1, \ldots, \vec{\mathbf{y}}_m/\vec{c}_m\big]$, when existentially quantified variables are Skolemized.

Small model property

Let us call the sentence (1) in Lemma 12 φ and number (3) φ'. Employing the Substructure Lemma, we can derive the small model property of SF sentences with the $\exists^*\forall^*\exists^*$ prefix from the equisatisfiability of φ and φ'. Suppose φ is a satisfiable sentence belonging to SF. In particular, this means φ does not contain any non-constant function symbols. Let $\mathcal{B}$ be a model of φ'. We define the structure $\mathcal{A}$ by $\mathfrak{U}_{\mathcal{A}} := \big\{\mathfrak{a} \in \mathfrak{U}_{\mathcal{B}} \mid$ there is some constant c in φ' such that $\mathfrak{a} = c^{\mathcal{B}}\big\}$; $c^{\mathcal{A}} := c^{\mathcal{B}}$ for every constant c in φ'; $f^{\mathcal{A}}(\vec{\mathfrak{a}}) := f^{\mathcal{B}}(\vec{\mathfrak{a}})$ for every m-ary function symbol f in φ' and every tuple $\vec{\mathfrak{a}} \in \mathfrak{U}_{\mathcal{A}}^m$; $P^{\mathcal{A}} := P^{\mathcal{B}} \cap \mathfrak{U}_{\mathcal{A}}^m$ for every m-ary predicate symbol P occurring in φ'. Since $\mathcal{A}$ is a substructure of $\mathcal{B}$, it is a model of φ'. It is easy to see that $\mathcal{A}$ is also a model of φ. Moreover, we have defined $\mathcal{A}$'s universe so that it contains at most $|\mathrm{consts}(\varphi')| = |\mathrm{consts}(\varphi)| + m_* \cdot |\vec{\mathbf{y}}|$ elements.

We next bound m_* from above in terms of the length of φ. Again, consider $\varphi_{\mathrm{DNF}} = \forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.$ $\bigvee_{k=1}^{m_{\mathrm{DNF}}} \chi_k(\vec{\mathbf{x}}) \wedge \eta_k(\vec{\mathbf{y}})$, which is equivalent to φ. Without loss of generality, we may assume the following: (i) the conjunctions $\chi_k \wedge \eta_k$ contain only literals which appear in ψ after transformation into negation normal form, and (ii) there are no distinct indices k_1, k_2 such that the set of literals occurring in $\chi_{k_1} \wedge \eta_{k_1}$ is a subset of the set of literals occurring in $\chi_{k_2} \wedge \eta_{k_2}$ (otherwise, $\chi_{k_2} \wedge \eta_{k_2}$ would be redundant and could be removed from the disjunction). Consequently, m_{DNF} is bounded from above by the maximal number of pairwise non-inclusive subsets of the set of all literals occurring in φ. Hence, by virtue of Sperner's Theorem, an upper bound for m_{DNF} is $\binom{\mathrm{len}(\varphi)}{\lfloor \mathrm{len}(\varphi)/2\rfloor} \leq 2^{\mathrm{len}(\varphi)}$, where $\mathrm{len}(\cdot)$ shall denote a reasonable measure of length of formulas (taking into account occurrences of quantifiers, Boolean connectives, variables, and occurrences of predicate, function, and constant symbols; we assume $\mathrm{len}(\varphi \rightarrow \psi) = \mathrm{len}(\neg\varphi \vee \psi)$ and $\mathrm{len}(\varphi \leftrightarrow \psi) = \mathrm{len}((\neg\varphi \vee \psi) \wedge (\varphi \vee \neg\psi))$). All in all, we may conclude $|\mathfrak{U}_{\mathcal{A}}| \leq \mathrm{len}(\varphi) + 2^{\mathrm{len}(\varphi)} \cdot \mathrm{len}(\varphi) \leq 2^{3\cdot\mathrm{len}(\varphi)}$. This means that if φ is satisfiable, then it has a model of size at most $2^{3\cdot\mathrm{len}(\varphi)}$.

Worst-case time complexity

Lewis employed the following lemma in [24] to find upper bounds on the required time for Bernays–Schönfinkel sentences and relational monadic sentences without equality.

Lemma 13 ([24], Proposition 3.2). Let φ be a first-order sentence in prenex normal form containing n universally quantified variables. The question whether φ has a model of cardinality m can be decided nondeterministically in time $p\big(m^n \cdot \mathrm{len}(\varphi)\big)$ for some polynomial p.

Together with our previous results and known lower bounds on time complexity (cf. [24]) this yields the following theorem.

Theorem 14. Satisfiability of sentences in SF with the quantifier prefix $\exists^*\forall^*\exists^*$ is NEXPTIME-complete.

Proof. Let $\varphi := \exists \vec{z} \forall \vec{x} \exists \vec{y}.\psi$ be an SF sentence, in which ψ is quantifier-free. Due to previous observations we know that the sentence $\varphi'' := \forall \vec{x} \exists \vec{y}.\psi[\vec{z}/\vec{d}\,] \wedge \bigvee_{\ell=1}^{m_*} \bigwedge_{i=1}^{|\vec{y}|} y_i \approx c_{\ell,i}$ (for Skolem constants $d_1, \ldots, d_{|\vec{z}|}$) has a model (based on a universe with at most $2^{3 \cdot \mathrm{len}(\varphi)}$ elements) if and only if φ has one. Clearly, every model of φ'' is also a model of φ. By Lemma 13, we can nondeterministically check whether φ has a model of size $2^{3 \cdot \mathrm{len}(\varphi)}$ in at most $p\big(2^{(3 \cdot \mathrm{len}(\varphi)^2)} \cdot \mathrm{len}(\varphi)\big)$ computational steps for some polynomial p.

Lewis [24] has shown NEXPTIME-hardness of satisfiability of BS sentences, i.e. of $\exists^* \forall^*$ sentences. By Theorem 9, these are included in the $\exists^* \forall^* \exists^*$ subfragment of SF. $\square$

It is worth noting that satisfiability of $\exists^* \forall^3 \exists^*$ sentences, in which variables are not separated, is known to be undecidable for several subcases, see [7] for references.

4.1.2 Sentences with several blocks of quantifiers

As in the previous section we replace leading existential quantifiers with constant symbols in this section.

The special case of strong separation

Consider a sentence $\varphi := \forall \vec{x}_1 \exists \vec{y}_1 \ldots \forall \vec{x}_n \exists \vec{y}_n.\psi$ for some quantifier-free first-order formula ψ without non-constant function symbols. We assume that the sets $\vec{y}_1, \ldots, \vec{y}_n$ and $\vec{x}_1 \cup \ldots \cup \vec{x}_n$ are all pairwise separated in ψ. Hence, we can transform ψ into a disjunction of conjunctions of the form $\psi_{\mathrm{DNF}} := \bigvee_{j=1}^{m_{\mathrm{DNF}}} \chi_j(\vec{x}_1, \ldots, \vec{x}_n) \wedge \bigwedge_{k=1}^{n} \eta_{j,k}(\vec{y}_k)$, in which the χ_j and the $\eta_{j,k}$ are conjunctions of literals. The additional requirement of $\vec{y}_1, \ldots, \vec{y}_n$ being pairwise separated relieves us from the need to use nested fingerprints, since at the level of atoms every tuple $\vec{y}_k$ occurs isolated from the others, i.e. the values assigned to one tuple $\vec{y}_{k_1}$ do not influence the truth values of the subformulas $\eta_{j,k_2}(\vec{y}_{k_2})$ under $\mathcal{A}$ for $k_2 \neq k_1$. Nested fingerprints will be necessary in the general case later on.

Let $\mathcal{A}$ be an arbitrary structure over the signature of φ. For every index $k \leq n$ we define a fingerprint function $\lambda_k : \mathfrak{U}_{\mathcal{A}}^{|\vec{y}_k|} \to \mathfrak{P}[m_{\mathrm{DNF}}]$ such that for every tuple $\vec{a}_k \in \mathfrak{U}_{\mathcal{A}}^{|\vec{y}_k|}$ it holds $\lambda_k(\vec{a}_k) := \big\{j \mid \mathcal{A}, [\vec{y}_k \mapsto \vec{a}_k] \models \eta_{j,k}\big\}$. If two tuples $\vec{a}_k, \vec{a}_k'$ are assigned the same fingerprint by λ_k, then they result in the same truth value for $\eta_{j,k}(\vec{y}_k)$ under $\mathcal{A}$, i.e. for every j we have $\mathcal{A}, [\vec{y}_k \mapsto \vec{a}_k] \models \eta_{j,k}$ if and only if $\mathcal{A}, [\vec{y}_k \mapsto \vec{a}_k'] \models \eta_{j,k}$. Since the variables in $\vec{y}_k$ do not occur in other subformulas than the $\eta_{j,k}$, we conclude the following for every k and an arbitrary variable assignment β: $\mathcal{A}, \beta[\vec{y}_k \mapsto \vec{a}_k] \models \forall \vec{x}_{k+1} \exists \vec{y}_{k+1} \ldots \forall \vec{x}_n \exists \vec{y}_n.\psi_{\mathrm{DNF}}$ holds if and only if $\mathcal{A}, \beta[\vec{y}_k \mapsto \vec{a}_k'] \models \forall \vec{x}_{k+1} \exists \vec{y}_{k+1} \ldots \forall \vec{x}_n \exists \vec{y}_n.\psi_{\mathrm{DNF}}$. In other words, $\vec{a}_k$ and $\vec{a}_k'$ are interchangeable as values for $\vec{y}_k$, whenever they are assigned the same fingerprint by λ_k.

Every λ_k induces a partition of $\mathfrak{U}_{\mathcal{A}}^{|\vec{y}_k|}$ into at most $2^{m_{\mathrm{DNF}}}$ equivalence classes of tuples with identical fingerprints with respect to λ_k. As we have done before, we can define representatives $\alpha_{k,S}$ for every fingerprint $S \subseteq [m_{\mathrm{DNF}}]$, for which there is some tuple $\vec{a}$ with $\lambda_k(\vec{a}) = S$. In the end, in analogy to the simpler case, we may derive equisatisfiability of φ and $\forall \vec{x}_1 \exists \vec{y}_1 \ldots \forall \vec{x}_n \exists \vec{y}_n.\psi \wedge \bigwedge_{k=1}^{n} \bigvee_{j=1}^{2^{m_{\mathrm{DNF}}}} \bigwedge_{i=1}^{|\vec{y}_k|} y_{k,i} \approx c_{k,j,i}$, where all the $c_{k,j,i}$ are fresh constant symbols.

Using the techniques we have seen earlier, we can improve this result in two ways. For one thing, the last quantifier block $\exists \vec{y}_n$ does not need to range over $2^{m_{\mathrm{DNF}}}$ tuples of constants, but m_{DNF} are sufficient, as we have already seen in Section 4.1.1. Secondly, the $\vec{y}_1, \ldots, \vec{y}_{n-1}$ do not need to range over $2^{m_{\mathrm{DNF}}}$ tuples either, since we can stick to *covering* representatives instead of representatives with *exactly the same fingerprint* and then apply Sperner's Theorem again. Thus, we need to consider at most $\kappa_{\mathrm{DNF}} := \binom{m_{\mathrm{DNF}}}{\lfloor m_{\mathrm{DNF}}/2 \rfloor} \leq 2^{m_{\mathrm{DNF}}}$ representatives for every $k < n$.

Consequently, we may conclude that the original φ is equisatisfiable to $\forall \vec{x}_1 \exists \vec{y}_1 \ldots \forall \vec{x}_n \exists \vec{y}_n.\psi \wedge \big(\bigwedge_{k=1}^{n-1} \bigvee_{j=1}^{\kappa_{\mathrm{DNF}}} \bigwedge_{i=1}^{|\vec{y}_k|} y_{k,i} \approx c_{k,j,i}\big) \wedge \big(\bigvee_{j=1}^{m_{\mathrm{DNF}}} \bigwedge_{i=1}^{|\vec{y}_n|} y_{n,i} \approx c_{n,j,i}\big)$ where all the $c_{k,j,i}$ are fresh.

Applying a similar analysis as in Section 4.1.1 leads to the observation that if φ has a model, then it has one with at most $|\mathrm{consts}(\varphi)| + \kappa_{\mathrm{DNF}} \cdot \sum_{k=1}^{n-1} |\vec{y}_k| + m_{\mathrm{DNF}} \cdot |\vec{y}_n| \leq 2^{3 \cdot 2^{\mathrm{len}(\varphi)}}$ domain elements.

Theorem 15. Let $\varphi := \exists \vec{\mathbf{z}} \forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n.\psi$ be a sentence in SF for some quantifier-free ψ without non-constant function symbols. If the sets $\vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_n$ are pairwise separated in ψ, then satisfiability of φ can be decided nondeterministically in time that is at most double exponential in $\mathrm{len}(\varphi)$.

The general case

Consider a sentence $\varphi := \forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n.\psi$ for some quantifier-free ψ without function symbols in which the sets $\vec{\mathbf{x}}_1 \cup \ldots \cup \vec{\mathbf{x}}_n$ and $\vec{\mathbf{y}}_1 \cup \ldots \cup \vec{\mathbf{y}}_n$ are separated. We can transform ψ into a disjunction of conjunctions $\psi_{\mathrm{DNF}} := \bigvee_{j=1}^{m_{\mathrm{DNF}}} \chi_j(\vec{\mathbf{x}}_1, \ldots, \vec{\mathbf{x}}_n) \wedge \eta_j(\vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_n)$ in which the χ_j and the η_j are conjunctions of literals. As we have already announced in the beginning of Section 4.1, we need to deal with a nested form of fingerprints in the general case of several quantifier alternations. The reason is that the truth values of the η_j depend on the values assigned to variables across multiple existential quantifier blocks.

For the following definitions we enhance the power set operator $\mathfrak{P}$ by allowing for iteration: $\mathfrak{P}^k S$ shall denote the k-fold application of $\mathfrak{P}$ to S. Let $\mathcal{A}$ be an arbitrary structure over the signature of φ. We inductively define fingerprint functions $\lambda_{n-1}, \ldots, \lambda_1$ with signatures $\lambda_k :$ $\mathfrak{U}_{\mathcal{A}}^{\sum_{j=1}^{k} |\vec{\mathbf{y}}_j|} \to \mathfrak{P}^{n-k}[m_{\mathrm{DNF}}]$ as follows.

- $\lambda_{n-1} : \mathfrak{U}_{\mathcal{A}}^{\sum_{j=1}^{n-1} |\vec{\mathbf{y}}_j|} \to \mathfrak{P}[m_{\mathrm{DNF}}]$ such that for all tuples $\vec{\mathfrak{a}}_1 \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}_1|}, \ldots, \vec{\mathfrak{a}}_{n-1} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}_{n-1}|}$ and every $j \in [m_{\mathrm{DNF}}]$ it holds $j \in \lambda_{n-1}(\vec{\mathfrak{a}}_1, \ldots, \vec{\mathfrak{a}}_{n-1})$ if and only if there exists a tuple $\vec{\mathfrak{a}}_n$ such that $\mathcal{A}, [\vec{\mathbf{y}}_1 \mapsto \vec{\mathfrak{a}}_1, \ldots, \vec{\mathbf{y}}_{n-1} \mapsto \vec{\mathfrak{a}}_{n-1}, \vec{\mathbf{y}}_n \mapsto \vec{\mathfrak{a}}_n] \models \eta_j$;

- $\lambda_{n-2} : \mathfrak{U}_{\mathcal{A}}^{\sum_{j=1}^{n-2} |\vec{\mathbf{y}}_j|} \to \mathfrak{P}^2[m_{\mathrm{DNF}}]$ such that for all tuples $\vec{\mathfrak{a}}_1 \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}_1|}, \ldots, \vec{\mathfrak{a}}_{n-2} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}_{n-2}|}$ and every $S \in \mathfrak{P}[m_{\mathrm{DNF}}]$ it holds $S \in \lambda_{n-2}(\vec{\mathfrak{a}}_1, \ldots, \vec{\mathfrak{a}}_{n-2})$ if and only if there exists a tuple $\vec{\mathfrak{a}}_{n-1}$ such that $\lambda_{n-1}(\vec{\mathfrak{a}}_1, \ldots, \vec{\mathfrak{a}}_{n-2}, \vec{\mathfrak{a}}_{n-1}) = S$;

$\vdots$

- $\lambda_1 : \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}_1|} \to \mathfrak{P}^{n-1}[m_{\mathrm{DNF}}]$ such that for every tuple $\vec{\mathfrak{a}}_1 \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}_1|}$ and every $S \in \mathfrak{P}^{n-2}[m_{\mathrm{DNF}}]$ it holds $S \in \lambda_1(\vec{\mathfrak{a}}_1)$ if and only if there exists a tuple $\vec{\mathfrak{a}}_2$ such that $\lambda_2(\vec{\mathfrak{a}}_1, \vec{\mathfrak{a}}_2) = S$.

The following lemma expresses that sequences of tuples with identical fingerprints may be interchanged without affecting semantics.

Lemma 16. For every $k < n$ and all tuples $\vec{\mathfrak{a}}_1, \vec{\mathfrak{a}}_1' \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}_1|}, \ldots, \vec{\mathfrak{a}}_k, \vec{\mathfrak{a}}_k' \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}_k|}, \vec{\mathfrak{b}}_1 \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{x}}_1|}, \ldots, \vec{\mathfrak{b}}_k \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{x}}_k|}$ if $\lambda_k(\vec{\mathfrak{a}}_1, \ldots, \vec{\mathfrak{a}}_k) = \lambda_k(\vec{\mathfrak{a}}_1', \ldots, \vec{\mathfrak{a}}_k')$ then it holds
(i) $\mathcal{A}, [\vec{\mathbf{x}}_1 \mapsto \vec{\mathfrak{b}}_1, \ldots, \vec{\mathbf{x}}_k \mapsto \vec{\mathfrak{b}}_k, \vec{\mathbf{y}}_1 \mapsto \vec{\mathfrak{a}}_1, \ldots, \vec{\mathbf{y}}_k \mapsto \vec{\mathfrak{a}}_k] \models \forall \vec{\mathbf{x}}_{k+1} \exists \vec{\mathbf{y}}_{k+1} \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n.\psi$
if and only if
(ii) $\mathcal{A}, [\vec{\mathbf{x}}_1 \mapsto \vec{\mathfrak{b}}_1, \ldots, \vec{\mathbf{x}}_k \mapsto \vec{\mathfrak{b}}_k, \vec{\mathbf{y}}_1 \mapsto \vec{\mathfrak{a}}_1', \ldots, \vec{\mathbf{y}}_k \mapsto \vec{\mathfrak{a}}_k'] \models \forall \vec{\mathbf{x}}_{k+1} \exists \vec{\mathbf{y}}_{k+1} \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n.\psi$.

Proof. We proceed by induction from $k = n-1$ to $k = 1$. For the sake of readability, we abbreviate sequences such as $\vec{\mathbf{x}}_1 \mapsto \vec{\mathfrak{b}}_1, \ldots, \vec{\mathbf{x}}_k \mapsto \vec{\mathfrak{b}}_k$ by $\vec{\mathbf{x}}_{\leq k} \mapsto \vec{\mathfrak{b}}_{\leq k}$.

Suppose $k = n-1$. Let S be the value of $\lambda_{n-1}(\vec{\mathfrak{a}}_1, \ldots, \vec{\mathfrak{a}}_{n-1})$ (which we assume to be identical to $\lambda_{n-1}(\vec{\mathfrak{a}}_1', \ldots, \vec{\mathfrak{a}}_{n-1}')$). By definition of λ_{n-1}, S is a subset of $[m_{\mathrm{DNF}}]$. Assume (i) holds and let $\vec{\mathfrak{c}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{x}}_n|}$ be arbitrary. There must be some index $j \in [m_{\mathrm{DNF}}]$ and some $\vec{\mathfrak{d}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}_n|}$ such that $\mathcal{A}, [\vec{\mathbf{x}}_{\leq n-1} \mapsto \vec{\mathfrak{b}}_{\leq n-1}, \vec{\mathbf{x}}_n \mapsto \vec{\mathfrak{c}}] \models \chi_j$ and $\mathcal{A}, [\vec{\mathbf{y}}_{\leq n-1} \mapsto \vec{\mathfrak{a}}_{\leq n-1}, \vec{\mathbf{y}}_n \mapsto \vec{\mathfrak{d}}] \models \eta_j$. Hence, $j \in S$ and thus there is some $\vec{\mathfrak{d}}' \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}_n|}$ such that $\mathcal{A}, [\vec{\mathbf{y}}_{\leq n-1} \mapsto \vec{\mathfrak{a}}_{\leq n-1}', \vec{\mathbf{y}}_n \mapsto \vec{\mathfrak{d}}'] \models \eta_j$. Therefore, (ii) must hold too. Since this argument is completely symmetric, (i) holds if and only if (ii) does.

Suppose that $k < n-1$ and let $S := \lambda_k(\vec{\mathfrak{a}}_1, \ldots, \vec{\mathfrak{a}}_k) = \lambda_k(\vec{\mathfrak{a}}_1', \ldots, \vec{\mathfrak{a}}_k')$. Assume (i) and let $\vec{\mathfrak{c}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{x}}_{k+1}|}$ be arbitrary. Then there must be some tuple $\vec{\mathfrak{d}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}_{k+1}|}$ such that $\mathcal{A}, [\vec{\mathbf{x}}_{\leq k} \mapsto \vec{\mathfrak{b}}_{\leq k}, \vec{\mathbf{x}}_{k+1} \mapsto \vec{\mathfrak{c}}, \vec{\mathbf{y}}_{\leq k} \mapsto \vec{\mathfrak{a}}_{\leq k}, \vec{\mathbf{y}}_{k+1} \mapsto \vec{\mathfrak{d}}] \models \forall \vec{\mathbf{x}}_{k+2} \exists \vec{\mathbf{y}}_{k+2} \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n.\psi$. Because of $\lambda_{k+1}(\vec{\mathfrak{a}}_1, \ldots, \vec{\mathfrak{a}}_k, \vec{\mathfrak{d}}) \in S$, there must be some tuple $\vec{\mathfrak{d}}' \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}_{k+1}|}$ such that $\lambda_{k+1}(\vec{\mathfrak{a}}_1', \ldots, \vec{\mathfrak{a}}_k', \vec{\mathfrak{d}}') = \lambda_{k+1}(\vec{\mathfrak{a}}_1, \ldots, \vec{\mathfrak{a}}_k, \vec{\mathfrak{d}})$. Hence,

by induction, we get $\mathcal{A}, [\vec{\mathbf{x}}_{\leq k} \mapsto \vec{\mathfrak{b}}_{\leq k}, \vec{\mathbf{x}}_{k+1} \mapsto \vec{\mathfrak{c}}, \vec{\mathbf{y}}_{\leq k} \mapsto \vec{\mathfrak{a}}'_{\leq k}, \vec{\mathbf{y}}_{k+1} \mapsto \vec{\mathfrak{d}}'] \models \forall \vec{\mathbf{x}}_{k+2} \exists \vec{\mathbf{y}}_{k+2} \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n.\psi$. Consequently, (ii) must hold, too. Again, by symmetry, (ii) does also entail (i). $\qquad\square$

Having defined the functions $\lambda_1, \ldots, \lambda_{n-1}$ and having shown that tuples of elements, which are associated with the same fingerprint, are interchangeable, we can now employ the same methods as before: (i) partition the sets $\mathfrak{U}_{\mathcal{A}}^{\sum_{j=1}^{k} |\vec{\mathbf{y}}_k|}$ in accordance with the fingerprints assigned by the λ_k, (ii) fix representatives $\alpha_{k,S}^{\ell}$ for every k and every occurring fingerprint S (the ℓ accounts for multiple representatives with fingerprint S at level k), (iii) restrict the range of quantifier blocks $\exists \vec{\mathbf{y}}_k$ to (parts of) the representatives $\alpha_{k,S}^{\ell}$. Step (ii) is slightly more complicated in this setting, because the truth value of the subformulas $\eta_j(\vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_n)$ under $\mathcal{A}$ depends on all the values assigned to $\vec{\mathbf{y}}_1, \ldots, \vec{\mathbf{y}}_n$. The consequence is not only a nesting of fingerprints but also a nesting of representatives: the $\alpha_{k,S}^{\ell}$ at level k are extensions of some $\alpha_{k-1,S'}^{\ell'}$ at level $k-1$ with $S \in S'$, respectively. More precisely, starting from some representative $\alpha_{k-1,S'}^{\ell'} = \langle \vec{\mathfrak{a}}_1, \ldots, \vec{\mathfrak{a}}_{k-1} \rangle$ with the fingerprint $\lambda_{k-1}(\vec{\mathfrak{a}}_1, \ldots, \vec{\mathfrak{a}}_{k-1}) = S'$, we pick one representative $\alpha_{k,S}^{\ell} = \langle \vec{\mathfrak{a}}_1, \ldots, \vec{\mathfrak{a}}_{k-1}, \vec{\mathfrak{a}}_k \rangle$ with $\lambda_k(\vec{\mathfrak{a}}_1, \ldots, \vec{\mathfrak{a}}_{k-1}, \vec{\mathfrak{a}}_k) = S$ for every $S \in S'$. Obviously, this approach might produce more than one representative with the fingerprint S at level k. In order to account for such multiplicities, we formally annotate the $\alpha_{k,S}$ with indices ℓ.

In the end, we can derive equisatisfiability of φ and $\varphi' :=$

$$\forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n.\psi$$

$$\wedge \bigvee_{j_1=1}^{|\mathfrak{P}^{n-1}[m_{\mathrm{DNF}}]|} \left(\vec{\mathbf{y}}_1 \approx \vec{c}_{\langle j_1 \rangle} \wedge \bigvee_{j_2=1}^{|\mathfrak{P}^{n-2}[m_{\mathrm{DNF}}]|} \left(\vec{\mathbf{y}}_2 \approx \vec{c}_{\langle j_1, j_2 \rangle} \right. \right.$$

$$\wedge \quad \ldots \quad \left(\ldots \bigvee_{j_n=1}^{|[m_{\mathrm{DNF}}]|} \vec{\mathbf{y}}_n \approx \vec{c}_{\langle j_1, \ldots, j_n \rangle} \right) \ldots \left. \left. \right) \right),$$

where $\vec{\mathbf{y}}_k \approx \vec{c}_{\langle j_1, \ldots, j_k \rangle}$ stands for $\bigwedge_{i=1}^{|\vec{\mathbf{y}}_k|} y_{k,i} \approx c_{\langle j_1, \ldots, j_k \rangle; i}$ and all the $c_{\langle j_1, \ldots, j_k \rangle; i}$ are fresh constant symbols. In accordance with the approach described above, we introduce a nested form of range-restricting constraint this time.

In order to compute the number of constant symbols in φ', we first define the notation $2^{\uparrow k}(m)$ inductively: $2^{\uparrow 0}(m) := m$ and $2^{\uparrow k+1}(m) := 2^{(2^{\uparrow k}(m))}$. The number of constants that occur in φ' is $|\mathrm{consts}(\varphi)| + \sum_{k=0}^{n-1} \left(\prod_{\ell=k}^{n-1} 2^{\uparrow \ell}(m_{\mathrm{DNF}}) \right) \cdot |\vec{\mathbf{y}}_{n-k}| \leq \mathrm{len}(\varphi) + n \cdot \mathrm{len}(\varphi) \cdot \left(2^{\uparrow n}(\mathrm{len}(\varphi)) \right)^n$.

Theorem 17. Let $\varphi := \exists \vec{\mathbf{z}} \forall \vec{\mathbf{x}}_1 \exists \vec{\mathbf{y}}_1 \ldots \forall \vec{\mathbf{x}}_n \exists \vec{\mathbf{y}}_n.\psi$ be a sentence in SF for some quantifier-free ψ without non-constant function symbols. Satisfiability of φ can be decided nondeterministically in time that is at most n-fold exponential in $\mathrm{len}(\varphi)$.

4.1.3 Open formulas and dependencies

Let $\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}})$ be a quantifier-free first-order formula in which $\vec{\mathbf{x}}$ and $\vec{\mathbf{y}}$ are separated. We proceed analogously to the case of closed formulas treated in Section 4.1.1. The only difference is that the employed set constructions and the underlying fingerprint function depend on the valuation of the parameter tuple $\vec{\mathbf{z}}$.

We transform $\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}})$ into an equivalent formula $\bigvee_{k=1}^{m_{\mathrm{DNF}}} \chi_k(\vec{\mathbf{x}}, \vec{\mathbf{z}}) \wedge \eta_k(\vec{\mathbf{y}}, \vec{\mathbf{z}})$ such that χ_k and η_k denote conjunctions of literals and m_{DNF} is some non-negative integer. Given a structure $\mathcal{A}$ and a tuple $\vec{\mathfrak{c}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{z}}|}$, we define the fingerprint function $\lambda_{\vec{\mathfrak{c}}}$ by setting $\lambda_{\vec{\mathfrak{c}}}(\vec{\mathfrak{b}}) := \{ k \in [m_{\mathrm{DNF}}] \mid \mathcal{A}, [\vec{\mathbf{x}} \mapsto \vec{\mathfrak{b}}, \vec{\mathbf{z}} \mapsto \vec{\mathfrak{c}}] \models \chi_k(\vec{\mathbf{x}}, \vec{\mathbf{z}}) \}$ for every $\vec{\mathfrak{b}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{x}}|}$. In order to complement $\lambda_{\vec{\mathfrak{c}}}$, we construct the sets $\widehat{\mathfrak{U}}_{\vec{\mathfrak{c}},1}, \ldots, \widehat{\mathfrak{U}}_{\vec{\mathfrak{c}},m_{\mathrm{DNF}}} \subseteq \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}|}$ such that $\widehat{\mathfrak{U}}_{\vec{\mathfrak{c}},k} := \{ \vec{\mathfrak{a}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}|} \mid \mathcal{A}, [\vec{\mathbf{y}} \mapsto \vec{\mathfrak{a}}, \vec{\mathbf{z}} \mapsto \vec{\mathfrak{c}}] \models \eta_k \}$. For every nonempty $\widehat{\mathfrak{U}}_{\vec{\mathfrak{c}},k}$ we pick one representative $\alpha_{\vec{\mathfrak{c}},k} \in \widehat{\mathfrak{U}}_{\vec{\mathfrak{c}},k}$. Based on these sets, we define the mapping $\mu_{\vec{\mathfrak{c}}} : \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{x}}|} \to [m_{\mathrm{DNF}}]$ such that $\mu_{\vec{\mathfrak{c}}}(\vec{\mathfrak{b}})$ yields some fixed index $k \in \lambda_{\vec{\mathfrak{c}}}(\vec{\mathfrak{b}})$ for which $\mathfrak{U}_{\vec{\mathfrak{c}},k}$ is nonempty. If no such k exists, $\mu_{\vec{\mathfrak{c}}}(\vec{\mathfrak{b}})$ remains undefined.

If $\mathcal{A}, [\vec{\mathbf{z}} \mapsto \vec{\mathfrak{c}}] \models \forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}. \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}})$, then for every $\vec{\mathfrak{b}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{x}}|}$ we have $\mathcal{A}, [\vec{\mathbf{x}} \mapsto \vec{\mathfrak{b}}, \vec{\mathbf{y}} \mapsto \alpha_{\vec{\mathfrak{c}}, \mu_{\vec{\mathfrak{c}}}(\vec{\mathfrak{b}})}, \vec{\mathbf{z}} \mapsto \vec{\mathfrak{c}}] \models$
$\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}})$. Hence, for every quantifier prefix $\mathcal{Q}\vec{\mathbf{z}}$—some of the variables in $\vec{\mathbf{z}}$ may be existentially quantified and some universally—the following sentences are mutually equisatisfiable:

(1) $\mathcal{Q}\vec{\mathbf{z}} \forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}. \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}})$,

(2) $\mathcal{Q}\vec{\mathbf{z}} \forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}. \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}}) \wedge \bigvee_{\ell=1}^{m_{\mathrm{DNF}}} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} y_i \approx g_{\ell,i}(\vec{\mathbf{z}})$,

(3) $\mathcal{Q}\vec{\mathbf{z}} \forall \vec{\mathbf{x}}. \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}}) \big[y_1, \ldots, y_{|\vec{\mathbf{y}}|} \big] f_1(\vec{\mathbf{x}}, \vec{\mathbf{z}}), \ldots, f_{|\vec{\mathbf{y}}|}(\vec{\mathbf{x}}, \vec{\mathbf{z}})$
   $\wedge \bigvee_{\ell=1}^{m_{\mathrm{DNF}}} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} f_i(\vec{\mathbf{x}}, \vec{\mathbf{z}}) \approx g_{\ell,i}(\vec{\mathbf{z}})$,

where the $g_{\ell,i}$ are fresh function symbols of appropriate arity and the f_i are Skolem functions of appropriate arity.

Dually, we can transform $\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}})$ into an equivalent formula $\bigwedge_{i=1}^{m_{\mathrm{CNF}}} \chi_i'(\vec{\mathbf{x}}, \vec{\mathbf{z}}) \vee \eta_i'(\vec{\mathbf{y}}, \vec{\mathbf{z}})$ in which the χ_i' and η_i' are disjunctions of literals. Given a structure $\mathcal{A}$ and a fixed tuple $\vec{\mathfrak{c}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{z}}|}$ of domain elements, we define the fingerprint function $\lambda_{\vec{\mathfrak{c}}} : \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}|} \to \mathfrak{P}[m_{\mathrm{CNF}}]$ such that for every $\vec{\mathfrak{a}}$ it holds $\lambda_{\vec{\mathfrak{c}}}(\vec{\mathfrak{a}}) := \big\{ i \in [m_{\mathrm{CNF}}] \mid \mathcal{A}, [\vec{\mathbf{y}} \mapsto \vec{\mathfrak{a}}, \vec{\mathbf{z}} \mapsto \vec{\mathfrak{c}}] \models \eta_i'(\vec{\mathbf{y}}, \vec{\mathbf{z}}) \big\}$. The function $\lambda_{\vec{\mathfrak{c}}}$ induces a partition of $\mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}|}$ into $2^{m_{\mathrm{CNF}}}$ equivalence classes. For every fingerprint $S \subseteq [m_{\mathrm{CNF}}]$ such that there is some tuple $\vec{\mathfrak{a}}$ with $\lambda_{\vec{\mathfrak{c}}}(\vec{\mathfrak{a}}) = S$ we fix a representative $\alpha_{\vec{\mathfrak{c}}, S} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}|}$ of $\vec{\mathfrak{a}}$'s equivalence class, i.e. $\lambda_{\vec{\mathfrak{c}}}(\alpha_{\vec{\mathfrak{c}}, \lambda_{\vec{\mathfrak{c}}}(\vec{\mathfrak{a}})}) = \lambda_{\vec{\mathfrak{c}}}(\vec{\mathfrak{a}})$.

Then for all tuples $\vec{\mathfrak{b}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{x}}|}$ and $\vec{\mathfrak{a}} \in \mathfrak{U}_{\mathcal{A}}^{|\vec{\mathbf{y}}|}$, we have $\mathcal{A}, [\vec{\mathbf{x}} \mapsto \vec{\mathfrak{b}}, \vec{\mathbf{y}} \mapsto \vec{\mathfrak{a}}, \vec{\mathbf{z}} \mapsto \vec{\mathfrak{c}}] \models \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}})$ if and only if $\mathcal{A}, [\vec{\mathbf{x}} \mapsto \vec{\mathfrak{b}}, \vec{\mathbf{y}} \mapsto \alpha_{\vec{\mathfrak{c}}, \lambda_{\vec{\mathfrak{c}}}(\vec{\mathfrak{a}})}, \vec{\mathbf{z}} \mapsto \vec{\mathfrak{c}}] \models \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}})$. Consequently, for every quantifier prefix $\mathcal{Q}\vec{\mathbf{z}}$ the following sentences are mutually equisatisfiable:

(1) $\mathcal{Q}\vec{\mathbf{z}} \forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}. \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}})$,

(2) $\mathcal{Q}\vec{\mathbf{z}} \forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}. \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}}) \wedge \bigvee_{j=1}^{2^{m_{\mathrm{CNF}}}} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} y_i \approx g_{j,i}(\vec{\mathbf{z}})$,

(3) $\mathcal{Q}\vec{\mathbf{z}} \forall \vec{\mathbf{x}}. \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}}) \big[y_1, \ldots, y_{|\vec{\mathbf{y}}|} \big] f_1(\vec{\mathbf{x}}, \vec{\mathbf{z}}), \ldots, f_{|\vec{\mathbf{y}}|}(\vec{\mathbf{x}}, \vec{\mathbf{z}})$
   $\wedge \bigvee_{j=1}^{2^{m_{\mathrm{CNF}}}} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} f_i(\vec{\mathbf{x}}, \vec{\mathbf{z}}) \approx g_{j,i}(\vec{\mathbf{z}})$,

where the $g_{j,i}$ are fresh function symbols of appropriate arity and the f_i are Skolem functions of appropriate arity.

Together with the dual case, we have inferred equisatisfiability of the sentences $\mathcal{Q}\vec{\mathbf{z}} \forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.$ $\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}})$, and $\mathcal{Q}\vec{\mathbf{z}} \forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}. \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}, \vec{\mathbf{z}}) \wedge \bigvee_{j=1}^{m_*} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} y_i \approx g_{j,i}(\vec{\mathbf{z}})$, where $m_* := \min(2^{m_{\mathrm{CNF}}}, m_{\mathrm{DNF}})$ and $\mathcal{Q}\vec{\mathbf{z}}$ can be any quantifier prefix closing the formula—some of the variables in $\vec{\mathbf{z}}$ may be existentially quantified and some universally.

Relation to Henkin quantifiers

In [19] Henkin introduced a generalized notion of quantifiers, sometimes called *finite partially ordered quantifiers* or *branching quantifiers* or *nonlinear quantifiers*. We have seen that separation of variables leads to a weaker dependency of existentially quantified variables on universally quantified ones (cf. Proposition 3 and Lemma 12). The arity of Skolem functions may even be decreased at the price of a possibly exponential increase in the length of the formula (cf. Proposition 5). However, we have also seen that this does not lead to complete independence, as it would in the case of quantification in Henkin's style.

Example 18. Consider the following equivalent sentences

- $\forall z \forall x \exists y. \, Q(z, y) \leftrightarrow P(x)$,

- $\forall z \forall x \exists y. \big(\neg Q(z, y) \wedge \neg P(x) \big) \vee \big(Q(z, y) \wedge P(x) \big)$,

- $\forall z \exists y_1 y_2 \forall x. \big(Q(z, y_1) \to P(x) \big) \wedge \big(P(x) \to Q(z, y_2) \big)$,

which we shall address by φ, φ_{DNF} and φ', respectively.

Standard Skolemization of φ introduces a binary Skolem function f_y and replaces the single occurrence of y with the term $f_y(z,x)$, whose value fully depends on the value assigned to x (distinct values for x may cause distinct values for $f_y(z,x)$). In contrast to this, Skolemization of φ' replaces the two variables y_1, y_2 with the terms $f_{y_1}(z)$ and $f_{y_2}(z)$, respectively, for two unary Skolem functions f_{y_1}, f_{y_2}. Interestingly, neither of the new terms depends on x.

On the other hand, range-restricted Skolemization reconciles the two previous approaches by introducing three terms $f_y(z,x)$ and $g_1(z)$, $g_2(z)$—one depending on x and two which are independent of x—and by adding the restriction $\forall z \forall x. f_y(z,x) \approx g_1(z) \vee f_y(z,x) \approx g_2(z)$. This limits the dependence of the value of $f_y(z,x)$ on the value assigned to x to a finite degree (over infinite domains, distinct values for x cannot always result in distinct values for $f_y(z,x)$). However, it still does not lead to complete independence from x.

Henkin quantifiers can explicitly express dependencies of existentially quantified variables on universally quantified ones. For instance, we could write $\psi := \forall z \forall x \exists_z y. Q(z,y) \leftrightarrow P(x)$ to express that the value of y depends on the value of z but not on x's value. Then ψ is equisatisfiable to $\psi_{\mathrm{Sk}} := \forall z \forall x. Q(z, f_y'(z)) \leftrightarrow P(x)$ for some Skolem function f_y'. Due to the enforced independence of y from x in ψ, it is easy to construct a model $\mathcal{A}$ of φ which cannot be extended to a model $\mathcal{B}$ of ψ_{Sk} (e.g., set $\mathfrak{U}_{\mathcal{A}} := \{0,1\}$, $P^{\mathcal{A}} := \{0\}$ and $Q^{\mathcal{A}} := \{\langle 0,0 \rangle, \langle 1,1 \rangle\}$). Finding a satisfying extension $\mathcal{B}$ of $\mathcal{A}$ is not a problem in any other case of Skolemization that we have described above (for example, set $\mathfrak{U}_{\mathcal{B}} := \mathfrak{U}_{\mathcal{A}}$, $P^{\mathcal{B}} := P^{\mathcal{A}}$, $Q^{\mathcal{B}} := Q^{\mathcal{A}}$, $f_{y_1}^{\mathcal{B}}(0) := g_1^{\mathcal{B}}(0) := 1$, $f_{y_1}^{\mathcal{B}}(1) := g_1^{\mathcal{B}}(1) := 0$, $f_{y_2}^{\mathcal{B}}(0) := g_2^{\mathcal{B}}(0) := 0$, $f_{y_2}^{\mathcal{B}}(1) := g_2^{\mathcal{B}}(1) := 1$, and $f_y^{\mathcal{B}}(0,\mathfrak{a}) := g_2^{\mathcal{B}}(\mathfrak{a})$, $f_y^{\mathcal{B}}(1,\mathfrak{a}) := g_1^{\mathcal{B}}(\mathfrak{a})$ for every $\mathfrak{a} \in \{0,1\}$).

Altogether, the example illustrates that separation of existentially quantified variables from universally quantified ones does lead to a certain degree of independence, but it does not reach the level of independence Henkin quantifiers can guarantee. This is not at all surprising, because Henkin quantifiers increase the expressiveness of first-order logic.

4.2 Extensions of the Separated Fragment

In this section we describe methods extending SF into a proper superset of the Full Monadic Fragment (with unary function symbols but without equality)—shown to be decidable independently by Löb [25] and Gurevich [18]—and the Relational Monadic Fragment with equality (but without non-constant function symbols). We will show how sentences from both fragments can be transformed into ones pursuant to Definition 7 under an at most quadratic increase in the length of the formulas.

Adopting a method already used by Löb in [25] and also by Grädel (cf. proof of Proposition 6.2.7 in [7]), we can handle unary function symbols under certain restrictions.

Proposition 19. Let φ be a first-order sentence without non-unary function symbols (constants are allowed). If the unary function symbols exclusively occur in atoms starting with a unary predicate symbol, then we can find an equisatisfiable sentence φ' without non-constant function symbols such that any model $\mathcal{B}$ of φ' can be transformed into a model $\mathcal{A}$ of φ over the same domain. The length of φ' lies in $\mathcal{O}(\mathrm{len}(\varphi))$. Moreover, if φ belongs to SF, then φ' is also an SF sentence.

Proof. Let $f_1, \ldots, f_k$ be the unary function symbols occurring in φ. We apply the following transformation iteratively. Assume φ contains the atom $P(f_i(t))$ for some term t. We may transform φ into $\varphi\bigl[P(f_i(t)) \,/\, R(t)\bigr] \wedge \forall x. P(f_i(x)) \leftrightarrow R(x)$, where the R is a fresh unary predicate symbol and $\varphi\bigl[P(f_i(t)) \,/\, R(t)\bigr]$ is the formula we obtain from φ by replacing every occurrence of $P(f_i(t))$ by $R(t)$. Exhaustive application of this transformation to φ yields a sentence φ'' of the form $\psi \wedge \bigwedge_{i=1}^{k} \bigwedge_j \forall x. P_j(f_i(x)) \leftrightarrow R_{i,j}(x)$, where ψ does not contain any of the f_i anymore. If we conceive the f_i in φ' as Skolem functions and revert the Skolemization, the f_i vanish completely and we end up with the equisatisfiable sentence $\varphi' := \psi \wedge \forall x \exists y_1 \ldots y_k. \bigwedge_{i=1}^{k} \bigwedge_j P_j(y_i) \leftrightarrow R_{i,j}(x)$.

Because of $\mathrm{len}(\psi) \leq \mathrm{len}(\varphi)$ and since for any occurrence of an f_i in φ at most one new conjunct of a fixed length is introduced, it holds $\mathrm{len}(\varphi') \in \mathcal{O}(\mathrm{len}(\varphi))$. □

If we now allow for unary function symbols occurring in monadic atoms in SF, as Proposition 19 suggests, this extended fragment becomes a proper superset of the Full Monadic fragment without equality.

Equality in SF, as defined in Definition 7, is subject to the separation condition. However, there is no such restriction in monadic formulas with equality. For instance, while the sentence $\forall x \exists y.x \approx y$ is admitted for the Relational Monadic Fragment with equality, the sets $\{x\}$ and $\{y\}$ are not separated in this sentence. Next, we show why the separation restriction may be discarded for equations, if the sentence at hand exhibits the small model property. We start by treating the case of monadic sentences with equality but without non-constant function symbols.

Proposition 20. For every sentence φ in the Relational Monadic Fragment with equality we can construct an equisatisfiable relational monadic sentence φ' without equality. Moreover, the length of φ' is of order $\mathcal{O}(\mathrm{len}(\varphi)^2)$.

Proof. Let φ be a first-order sentence without non-constant function symbols and containing at most the unary predicate symbols $P_1, \ldots, P_n$ (besides the equality predicate). Let k be the number of occurrences of quantifiers in φ plus the number of constant symbols in φ, and set $\kappa := \lceil \log_2 k \rceil$. We extend the underlying signature with fresh unary predicate symbols $Q_1, \ldots, Q_\kappa$ and define the formula

$$\psi_\approx(x, y) := \bigwedge_{i=1}^{n} \big(P_i(x) \leftrightarrow P_i(y)\big) \wedge \bigwedge_{i=1}^{\kappa} \big(Q_i(x) \leftrightarrow Q_i(y)\big) .$$

The length of $\psi_\approx(x, y)$ lies in $\mathcal{O}(\mathrm{len}(\varphi))$. Let φ' be constructed from φ by replacing every occurrence of an equation $s \approx t$ with $\psi_\approx(x, y)\big[x/s, \ y/t\big]$. Since s, t cannot contain non-constant function symbols, the length of $\psi_\approx(x, y)\big[x/s, y/t\big]$ is the same as that of $\psi_\approx(x, y)$. We show that φ is satisfiable if and only if φ' is so.

Let $\mathcal{A}$ be a model of φ over the domain $\mathfrak{U}_\mathcal{A}$. We construct a fingerprint function $\lambda : \mathfrak{U}_\mathcal{A} \to \mathfrak{P}[n]$ for which $\lambda(\mathfrak{a}) := \big\{i \in [n] \mid \mathfrak{a} \in P_i^\mathcal{A}\big\}$ for every $\mathfrak{a} \in \mathfrak{U}_\mathcal{A}$. We now partition $\mathfrak{U}_\mathcal{A}$ into parts $\mathfrak{U}_S$ such that every $\mathfrak{U}_S$ contains exactly the elements $\mathfrak{a}$ with the fingerprint $\lambda(\mathfrak{a}) = S$.

Starting from $\mathcal{A}$ we construct an interpretation $\mathcal{A}'$ by arbitrarily choosing subsets $\mathfrak{U}'_S \subseteq \mathfrak{U}_S$ for every $S \subseteq [n]$ such that $|\mathfrak{U}'_S| = \min(|\mathfrak{U}_S|, k)$, and by defining $\mathcal{A}'$'s domain to be $\mathfrak{U}_{\mathcal{A}'} := \bigcup_{S \subseteq [n]} \mathfrak{U}'_S$. Moreover, we set $P_i^{\mathcal{A}'} := P_i^A \cap \mathfrak{U}_{\mathcal{A}'}$ for all P_i and define the predicates $Q_i^{\mathcal{A}'}$ so that for every $S \in [n]$ and for all elements $\mathfrak{a}, \mathfrak{a}' \in \mathfrak{U}'_S$ we can find an index j such that $\mathfrak{a} \in Q_j^{\mathcal{A}'}$ and $\mathfrak{a}' \notin Q_j^{\mathcal{A}'}$ or vice versa. Clearly, for every variable assignment $\beta : X \to \mathfrak{U}_{\mathcal{A}'}$ it then holds $\mathcal{A}', \beta \models x \approx y$ if and only if $\mathcal{A}', \beta \models \psi_\approx(x, y)$

Since a formula with k different variables and constant symbols cannot distinguish more than k elements of the same color, $\mathcal{A}'$ must also be a model of φ. This together with the equivalence of $x \approx y$ and $\psi_\approx(x, y)$ under $\mathcal{A}'$ entails that $\mathcal{A}'$ is a model of φ'. This proves the "only if"-direction.

Now let $\mathcal{B}'$ be an arbitrary model of φ'. We now construct a fingerprint function $\lambda' : \mathfrak{U}_{\mathcal{B}'} \to \mathfrak{P}[n{+}\kappa]$ analogously to λ, but now also taking the predicates $Q_i^{\mathcal{B}'}$ into account. Again, we partition $\mathfrak{U}_{\mathcal{B}'}$ into parts $\mathfrak{U}_S$ so that all elements in a set $\mathfrak{U}_S$ are assigned the same color by λ'. We now define the new universe $\mathfrak{U}_\mathcal{B}$ by arbitrarily picking exactly one element from every part $\mathfrak{U}_S$, $S \subseteq [n + \kappa]$. In addition, we set $P_i^\mathcal{B} := P_i^{\mathcal{B}'} \cap \mathfrak{U}_\mathcal{B}$ and $Q_i^\mathcal{B} := Q_i^{\mathcal{B}'} \cap \mathfrak{U}_\mathcal{B}$.

Since φ' can only distinguish elements of $\mathfrak{U}_\mathcal{B}$ by means of their belonging to predicates $P_i^\mathcal{B}$ and $Q_i^\mathcal{B}$, it is clear that $\mathcal{B}$ is a model of φ'. Moreover, for any variable assignment $\beta : X \to \mathfrak{U}_\mathcal{B}$ we get $\mathcal{B}, \beta \models \psi_\approx(x, y)$ if and only if $\mathcal{B}, \beta \models x \approx y$. Put together, these two facts entail $\mathcal{B} \models \varphi$. $\qquad\square$

More generally, we can formulate the above result for all sentences which exhibit the small model property.

Proposition 21. Let φ be a first-order sentence with equality. Suppose we can compute a positive integer k such that if φ is satisfiable, then there is a model $\mathcal{A} \models \varphi$ over a universe of cardinality at most k. Let $\kappa := \lceil \log_2 k \rceil$. We can transform φ into an equisatisfiable sentence φ' without equality using only the vocabulary of φ plus κ fresh unary predicate symbols $Q_1, \ldots, Q_\kappa$. The length of φ' lies in $\mathcal{O}(\kappa \cdot \mathrm{len}(\varphi)^3)$. Moreover, if existentially quantified variables are separated from universally quantified ones in all atoms in φ except for equations, then they are separated in *all* atoms in φ'.

Proof. We generalize the argument from the proof of Proposition 20. Since φ may also contain non-unary predicate symbols and non-constant function symbols, the encoding of equality requires additional congruence axioms. First, $\psi_{\approx}(x, y)$ becomes slightly simpler:

$$\psi_{\approx}(x, y) := \bigwedge_{i=1}^{\kappa} \big(Q_i(x) \leftrightarrow Q_i(y) \big) \ .$$

In addition, we need the following congruence axioms:

$$\psi_{\text{pred}} := \forall xy. \psi_{\approx}(x, y) \ \longrightarrow \ \bigwedge_{P} \bigwedge_{i=0}^{\text{arity}(P)-1} \big(\forall \vec{\mathbf{u}}_i \vec{\mathbf{u}}_i'. \ P(\vec{\mathbf{u}}_i, x, \vec{\mathbf{u}}_i') \to P(\vec{\mathbf{u}}_i, y, \vec{\mathbf{u}}_i') \big) \ ,$$

where P ranges over all predicate symbols in φ, except for the equality predicate, and $\vec{\mathbf{u}}_i, \vec{\mathbf{u}}_i'$ are disjoint sets of variables of cardinality $|\vec{\mathbf{u}}_i| = i$ and $|\vec{\mathbf{u}}_i'| = \text{arity}(P) - i - 1$ for the respective P, and

$$\psi_{\text{func}} := \forall xy. \psi_{\approx}(x, y) \ \longrightarrow \ \bigwedge_{f} \bigwedge_{i=0}^{\text{arity}(f)-1} \big(\forall \vec{\mathbf{v}}_i \vec{\mathbf{v}}_i'. \ \psi_{\approx}(x, y) \big[x/f(\vec{\mathbf{v}}_i, x, \vec{\mathbf{v}}_i'), y/f(\vec{\mathbf{v}}_i, y, \vec{\mathbf{v}}_i') \big] \big) \ ,$$

where f ranges over all non-constant function symbols in φ, and $\vec{\mathbf{v}}_i, \vec{\mathbf{v}}_i'$ are disjoint sets of variables of cardinality $|\vec{\mathbf{v}}_i| = i$ and $|\vec{\mathbf{v}}_i'| = \text{arity}(f) - i - 1$ for the respective f.

We construct φ' from φ by replacing every occurrence of an equation $s \approx t$ with the formula $\psi_{\approx}(x, y)[x/s, y/t]$, and by conjunctively connecting ψ_{pred} and ψ_{func} to it.

Let $\mathcal{A}$ be a model of φ over the finite domain $\mathfrak{U}_{\mathcal{A}} := \{\mathfrak{a}_1, \ldots, \mathfrak{a}_k\}$. Starting from $\mathcal{A}$, we construct a structure $\mathcal{B}$ such that $\mathfrak{U}_{\mathcal{B}} := \mathfrak{U}_{\mathcal{A}}$, $P^{\mathcal{B}} := P^{\mathcal{A}}$ for every predicate symbol P occurring in φ, $c^{\mathcal{B}} := c^{\mathcal{A}}$ for every constant symbol c occurring in φ, $f^{\mathcal{B}} := f^{\mathcal{A}}$ for every function symbol f occurring in φ. Moreover, we define $Q_1^{\mathcal{B}}, \ldots, Q_{\kappa}^{\mathcal{B}}$ such that for all distinct $\mathfrak{a}_{i_1}, \mathfrak{a}_{i_2} \in \mathfrak{U}_{\mathcal{B}}$ we find some j, $1 \le j \le \kappa$, for which $\mathfrak{a}_{i_1} \in Q_j^{\mathcal{B}}$ if and only if $\mathfrak{a}_{i_2} \notin Q_j^{\mathcal{B}}$. This is possible, because, having $\kappa = \lceil \log_2 k \rceil$ unary predicates, one can distinguish $2^{\kappa} \ge k$ domain elements. Hence, for any variable assignment β it holds $\mathcal{B}, \beta \models \psi_{\approx}(x, y)$ if and only if $\beta(x) = \beta(y)$ if and only if $\mathcal{B}, \beta \models x \approx y$. Consequently, $\mathcal{B}$ is a model of φ as well as of φ'. The finishes the "only if"-part of the proof.

Now let $\mathcal{B}$ be a model of φ'. Hence, $\mathcal{B}$ is also a model of ψ_{pred} and ψ_{func}. We define a fingerprint function $\lambda : \mathfrak{U}_{\mathcal{B}} \to \mathfrak{P}[\kappa]$ where we set $\lambda(\mathfrak{a}) := \big\{ i \in [\kappa] \ \big| \ \mathfrak{a} \in Q_i^{\mathcal{B}} \big\}$ for every $\mathfrak{a} \in \mathfrak{U}_{\mathcal{B}}$. Using λ, we decompose $\mathfrak{U}_{\mathcal{B}}$ into parts $\mathfrak{U}_S := \{\mathfrak{a} \in \mathfrak{U}_{\mathcal{B}} \mid \lambda(\mathfrak{a}) = S \subseteq [\kappa]\}$ and fix for every such set one representative $\alpha_S \in \mathfrak{U}_S$.

We now define a new structure $\mathcal{A}$ based on the representatives α_S. To this end, we set $\mathfrak{U}_{\mathcal{A}} := \{\alpha_{\lambda(\mathfrak{a})} \mid \mathfrak{a} \in \mathfrak{U}_{\mathcal{B}}\}$, $P^{\mathcal{A}} := P^{\mathcal{B}} \cap \mathfrak{U}_{\mathcal{A}}^m$ for every m-ary predicate symbol P occurring in φ, $c^{\mathcal{A}} := \alpha_{\lambda(c^{\mathcal{B}})}$ for every constant symbol c occurring in φ, $f^{\mathcal{A}}(\vec{\mathfrak{a}}) := \alpha_{\lambda(f^{\mathcal{B}}(\vec{\mathfrak{a}}))}$ for every m-ary function symbol f occurring in φ and every tuple $\vec{\mathfrak{a}} \in \mathfrak{U}_{\mathcal{A}}^m$.

We make the following observations:

(1) $\lambda(\alpha_{\lambda(\mathfrak{a})}) = \lambda(\mathfrak{a})$ for every $\mathfrak{a} \in \mathfrak{U}_{\mathcal{B}}$.

(2) $\alpha_{\lambda_B(\mathfrak{a})} = \mathfrak{a}$ for every $\mathfrak{a} \in \mathfrak{U}_{\mathcal{A}}$.

(3) Given $\mathfrak{a}, \mathfrak{b} \in \mathfrak{U}_{\mathcal{B}}$, $\lambda(\mathfrak{a}) = \lambda(\mathfrak{b})$ implies $\mathcal{B}, \beta \models \psi_{\approx}(x, y)$ for every variable assignment β that fulfills $\beta(x) = \mathfrak{a}$ and $\beta(y) = \mathfrak{b}$.

(4) Given $\mathfrak{a}_1, \ldots, \mathfrak{a}_m, \mathfrak{b}_1, \ldots, \mathfrak{b}_m \in \mathfrak{U}_{\mathcal{B}}$ such that $\lambda_B(\mathfrak{a}_i) = \lambda(\mathfrak{b}_i)$ for $i = 1, \ldots, m$, it holds

    (a) $\lambda(f^{\mathcal{B}}(\mathfrak{a}_1, \ldots, \mathfrak{a}_m)) = \lambda(f^{\mathcal{B}}(\mathfrak{b}_1, \ldots, \mathfrak{b}_m))$ and

    (b) $\langle \mathfrak{a}_1, \ldots, \mathfrak{a}_m \rangle \in P^{\mathcal{B}}$ if and only if $\langle \mathfrak{b}_1, \ldots, \mathfrak{b}_m \rangle \in P^{\mathcal{B}}$

for every m-ary function symbol f and every m-ary predicate symbol P occurring in φ', respectively.

All of these observations are consequences of the definition of the coloring λ and our assumption $\mathcal{B} \models \psi_{\text{pred}} \wedge \psi_{\text{func}}$.

<u>Claim I:</u> Let t be some term occurring in φ', and let β, β' be two variable assignments such that $\beta'(x) = \alpha_{\lambda(\beta(x))}$ holds for every variable x occurring in t. We observe $\lambda\big(\mathcal{A}(\beta')(t)\big) = \lambda\big(\mathcal{B}(\beta)(t)\big)$.

<u>Proof:</u> The proof proceeds by induction on the structure of t.

   Base cases: Suppose $t = c$ for some constant symbol c. By definition of $\mathcal{A}$, it holds $c^{\mathcal{A}} = \alpha_{\lambda(c^{\mathcal{B}})}$. Observation (1) thus entails $\lambda\big(\mathcal{A}(\beta')(c)\big) = \lambda\big(c^{\mathcal{A}}\big) = \lambda\big(\alpha_{\lambda(c^{\mathcal{B}})}\big) = \lambda\big(c^{\mathcal{B}}\big) = \lambda\big(\mathcal{B}(\beta)(c)\big)$.

   Suppose $t = x$ for some variable x. By assumption, we have $\beta'(x) = \alpha_{\lambda(\beta(x))}$. Observation (1) thus entails $\lambda\big(\mathcal{A}(\beta')(x)\big) = \lambda\big(\beta'(x)\big) = \lambda\big(\alpha_{\lambda(\beta(x))}\big) = \lambda\big(\beta(x)\big) = \lambda\big(\mathcal{B}(\beta)(x)\big)$.

   Inductive case: Let $t = f(s_1, \ldots, s_m)$ for some m-ary function symbol f and arbitrary terms $s_1, \ldots, s_m$. By definition of $f^{\mathcal{A}}$ and Observation (1), we get

   $$\lambda\big(\mathcal{A}(\beta')(f(s_1, \ldots, s_m))\big)$$
   $$= \lambda\big(f^{\mathcal{A}}(\mathcal{A}(\beta')(s_1), \ldots, \mathcal{A}(\beta')(s_m))\big)$$
   $$= \lambda\big(\alpha_{\lambda(f^{\mathcal{B}}(\mathcal{A}(\beta')(s_1), \ldots, \mathcal{A}(\beta')(s_m)))}\big)$$
   $$= \lambda\big(f^{\mathcal{B}}(\mathcal{A}(\beta')(s_1), \ldots, \mathcal{A}(\beta')(s_m))\big) \ .$$

   By induction, we conclude $\lambda\big(\mathcal{A}(\beta')(s_i)\big) = \lambda\big(\mathcal{B}(\beta)(s_i)\big)$ for $i = 1, \ldots, m$. Hence, Observation (4a) leads to

   $$\lambda\big(f^{\mathcal{B}}(\mathcal{A}(\beta')(s_1), \ldots, \mathcal{A}(\beta')(s_m))\big)$$
   $$= \lambda\big(f^{\mathcal{B}}(\mathcal{B}(\beta)(s_1), \ldots, \mathcal{B}(\beta)(s_m))\big)$$
   $$= \lambda\big(\mathcal{B}(\beta)(f(s_1, \ldots, s_m))\big) \ .$$

   Consequently, $\lambda\big(\mathcal{A}(\beta')(f(s_1, \ldots, s_m))\big)$ equals $\lambda\big(\mathcal{B}(\beta)(f(s_1, \ldots, s_m))\big)$. $\diamond$

<u>Claim II:</u> Let A be some atom occurring in φ', and let β, β' be two variable assignments such that $\beta'(x) = \alpha_{\lambda(\beta(x))}$ holds for every variable x occurring in A. $\mathcal{B}, \beta \models A$ holds if and only if $\mathcal{A}, \beta' \models A$.

<u>Proof:</u> Since φ' does not contain equality, A must be of the shape $P(t_1, \ldots, t_m)$ for some m-ary predicate symbol P and terms $t_1, \ldots, t_m$. $\mathcal{B}, \beta \models P(t_1, \ldots, t_m)$ holds if and only if $\big\langle \mathcal{B}(\beta)(t_1), \ldots, \mathcal{B}(\beta)(t_m) \big\rangle \in P^{\mathcal{B}}$. By Claim I, we know $\lambda\big(\mathcal{B}(\beta)(t_i)\big) = \lambda\big(\mathcal{A}(\beta')(t_i)\big)$ for $i = 1, \ldots, m$, and thus Observation (4b) entails that $\big\langle \mathcal{B}(\beta)(t_1), \ldots, \mathcal{B}(\beta)(t_m) \big\rangle \in P^{\mathcal{B}}$ holds true if and only if $\big\langle \mathcal{A}(\beta')(t_1), \ldots, \mathcal{A}(\beta')(t_m) \big\rangle \in P^{\mathcal{B}}$ does. And since all the domain elements $\mathcal{A}(\beta')(t_i)$ belong to $\mathfrak{U}_{\mathcal{A}}$, it follows that $\big\langle \mathcal{A}(\beta')(t_1), \ldots, \mathcal{A}(\beta')(t_m) \big\rangle \in P^{\mathcal{B}}$ is equivalent to $\big\langle \mathcal{A}(\beta')(t_1), \ldots, \mathcal{A}(\beta')(t_m) \big\rangle \in P^{\mathcal{B}} \cap \mathfrak{U}_{\mathcal{A}}^m = P^{\mathcal{A}}$. Consequently, we have $\mathcal{A}, \beta' \models P(t_1, \ldots, t_m)$ if and only if $\mathcal{B}, \beta \models P(t_1, \ldots, t_m)$ holds. $\diamond$

<u>Claim III:</u> $\mathcal{B} \models \varphi'$ implies $\mathcal{A} \models \varphi'$.

<u>Proof:</u> The proof proceeds by induction on the structure of φ'. However, we need to slightly generalize the claim first: given any two variable assignments β, β' for which $\beta'(x) = \alpha_{\lambda(\beta(x))}$ for every variable x occurring freely in φ', we observe that $\mathcal{B}, \beta \models \varphi'$ entails $\mathcal{A}, \beta' \models \varphi'$.

The base case is already handled by Claim II. Of the inductive cases only the ones with quantifiers are of interest, the others are trivial.

Suppose φ' is of the form $\forall x.\psi$. In order to show $\mathcal{A}, \beta' \models \forall x.\psi$, we must show that $\mathcal{A}, \beta'[x \mapsto \mathfrak{a}] \models \psi$ holds for every $\mathfrak{a} \in \mathfrak{U}_{\mathcal{A}}$. By Observation 2 and our assumption about

18

β and β', we have $\big(\beta'[x\mapsto\mathfrak{a}]\big)(v) = \alpha_{\lambda((\beta[x\mapsto\mathfrak{a}])(v))}$ for every variable v occurring freely in ψ, and thus the updated variable assignments $\beta[x\mapsto\mathfrak{a}]$ and $\beta'[x\mapsto\mathfrak{a}]$ fulfill the requirements of the inductive hypothesis. By induction, we derive $\mathcal{A},\beta'[x\mapsto\mathfrak{a}] \models \psi$ for any $\mathfrak{a} \in \mathfrak{U}_{\mathcal{A}}$ from $\mathcal{B},\beta[x\mapsto\mathfrak{a}] \models \psi$.

Suppose φ' is of the form $\exists y.\psi$. We need to find some $\mathfrak{a} \in \mathfrak{U}_{\mathcal{A}}$ for which $\mathcal{A},\beta'[y\mapsto\mathfrak{a}] \models \psi$ holds true. Since we assume $\mathcal{B},\beta \models \exists y.\psi$, there is some $\mathfrak{b} \in \mathfrak{U}_{\mathcal{B}}$ such that $\mathcal{B},\beta[y\mapsto\mathfrak{b}] \models \psi$. We set $\mathfrak{a} := \alpha_{\lambda(\mathfrak{b})}$. Hence, we have $\big(\beta'[y\mapsto\mathfrak{a}]\big)(v) = \alpha_{\lambda((\beta[y\mapsto\mathfrak{b}])(v))}$ for every variable v occurring freely in ψ. By induction, we obtain $\mathcal{A},\beta'[y\mapsto\mathfrak{a}] \models \psi$ from $\mathcal{B},\beta[y\mapsto\mathfrak{b}] \models \psi$.

This finishes the proof of $\mathcal{B},\beta \models \varphi'$ implying $\mathcal{A},\beta' \models \varphi'$. Since φ' is assumed to be a sentence and thus does not contain free occurrences of any variable, it immediately follows that $\mathcal{B} \models \varphi'$ entails $\mathcal{A} \models \varphi'$ $\diamond$

Due to the fact that all domain elements in $\mathfrak{U}_{\mathcal{A}}$ are assigned a different fingerprint by λ, we may conclude for every β that $\mathcal{A},\beta \models \psi_{\approx}(x,y)$ if and only if $\mathcal{A},\beta \models x \approx y$. Consequently, $\mathcal{A} \models \varphi'$ entails $\mathcal{A} \models \varphi$.

An inspection of the described construction leads to the following upper bounds on the length of the resulting (sub)formulas: $\text{len}\big(\psi_{\approx}(x,y)\big) \in \mathcal{O}(\kappa)$, $\text{len}\big(\psi_{\approx}(x,y)\big[x/s,y/t\big]\big) \in \mathcal{O}(\kappa \cdot \text{len}(\varphi))$, $\text{len}\big(\psi_{\text{pred}}\big) \in \mathcal{O}(\kappa + \text{len}(\varphi)^3)$, and $\text{len}\big(\psi_{\text{func}}\big) \in \mathcal{O}(\kappa \cdot \text{len}(\varphi)^3)$. All in all, this yields $\text{len}(\varphi') \in \mathcal{O}(\kappa \cdot \text{len}(\varphi)^3)$. □

The just proven proposition means that some sentences, which are *almost* SF sentences, can be transformed into equisatisfiable SF sentences, if they exhibit the small model property. More precisely, being *almost* SF in this context means they do not contain non-constant function symbols and fulfill the separation condition only for non-equational atoms.

In fact, we can even allow certain occurrences of non-constant function symbols. For instance, the setting of unary function symbols described in Proposition 19 is possible, i.e. equations of the form $f(g(f(x))) \approx g(h(y))$ with x being universally quantified and y existentially, may be allowed, as long as the sentence at hand has a finite model, for which we can computer an upper bound on its size. In this case, first applying the transformation given in the proof of Proposition 20 and subsequently the construction from the proof of Proposition 19 finally leads to an SF sentence. Combinations with other translation methods are also conceivable.

Remark 22. Other methods for the elimination of the distinguished equality predicate have been proposed. For instance, Dreben and Goldfarb [12] (Chapter 8, Sections 1 and 2) describe *negative identity reduction*, which introduces a binary predicate I to capture equality. However, this approach can obviously not resolve non-separated equations into separated atoms, and therefore does not exactly meet our particular needs. On the other hand, this method is not restricted to sentences possessing the small model property.

For an automated-reasoning perspective on the elimination of equality, consult [4, 2, 8, 23], for instance.

5 Separation and Automated Reasoning

The size of the search space of a first-order automated reasoning procedure is related to the size of the relevant subset of the Herbrand base that is actually explored by the procedure. Automated reasoning on first-order formulas with an explicit or implicit finite Herbrand base has attracted a lot of attention in recent years [16, 3, 28, 11, 20, 21, 6, 1]. In particular, all these procedures are decision procedures for the BS(R) Fragment. But also a clause set with an explicit finite domain axiom or a clause set enjoying an explicit or implicit acyclic atom structure generates only a finite relevant subset of the overall Herbrand base. If the relevant subset of the Herbrand base is infinite, implying the presence of non-constant function symbols, then the search space of an automated reasoning procedure becomes infinite and it does not terminate anymore, in general. Even for a

finite relevant subset generated over a finite number of constants the actual size of the set has an important influence on the explored search space.

In this section we apply our results to the benefit of the above mentioned automated reasoning procedures. We show the techniques for one quantifier block alternation, but they can be generalized to several blocks analogously to the results in Section 4. Note that given a set of sentences where not all sentences are separated, our results are still applicable to the separated sentences. Consider the separated sentence $\forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})$. This formula is satisfiable if and only if it is satisfiable over a domain of $m_* |\vec{\mathbf{y}}| + |\mathrm{consts}(\psi)|$ (see Lemma 12) elements

$$\forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}) \wedge \bigvee_{\ell=1}^{m_*} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} y_i \approx c_{\ell,i}.$$

Skolemizing the $\vec{\mathbf{y}}$ and explicitly representing the substitution in ψ by equations yields

$$\forall \vec{\mathbf{x}}\vec{\mathbf{y}}.\left(\left(\bigwedge_{i=1}^{|\vec{\mathbf{y}}|} f_i(\vec{\mathbf{x}}) \approx y_i \right) \to \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}) \right) \wedge \bigvee_{\ell=1}^{m_*} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} f_i(\vec{\mathbf{x}}) \approx c_{\ell,i}$$

that is a first-order sentence with an explicit finite domain axiom, so the relevant Herbrand base is finite. Still, reasoning with equality on the f_i is not needed here and we can further simplify the formula to the equisatisfiable sentence

$$\forall \vec{\mathbf{x}}\vec{\mathbf{y}}.\left(\left(\bigwedge_{i=1}^{|\vec{\mathbf{y}}|} R_i(\vec{\mathbf{x}}, y_i) \right) \to \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}) \right) \wedge \bigvee_{\ell=1}^{m_*} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} R_i(\vec{\mathbf{x}}, c_{\ell,i})$$

where we replaced the f_i by relations R_i without the need to add further axioms. The transformation preservers satisfiability. Totality of the f_i after translation to R_i is guaranteed by the finite domain axiom $\bigvee_{\ell=1}^{m_*} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} f_i(\vec{\mathbf{x}}) \approx c_{\ell,i}$. If some R_i interpretation contains more than one value for an $\vec{\mathbf{x}}$ assignment, all values except one can simply be dropped, preserving satisfiability. This is a consequence of the fact that all positive f_i equational occurrences (R_i occurrences) are exactly in the finite domain axiom. Eventually, if $\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})$ does not contain any non-constant function symbols, we moved a separated sentence that would have resulted after CNF generation in an infinite Herbrand base to a clause set of the BS(R) Fragment.

The size of the finite domain axiom is worst-case exponential in the number m_{CNF} of clauses generated out of $\forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})$ by a CNF procedure without renaming [27]. Thus, if the number of clauses can be reduced, it improves the size of the finite domain axiom. Redundant clauses do not contribute to m_{CNF}. For example, if two clauses subsume each other, only one needs to be considered for m_{CNF}. Note that m_{CNF} without considering redundancy can be computed without actually generating the CNF [27]. Renaming, i.e. the replacement of subformulas via fresh predicates, cannot be applied for determining m_{CNF}, because it may generate atoms in which variables $\vec{\mathbf{x}}$ and $\vec{\mathbf{y}}$ are no longer separated. An analogous argument holds for the transformation into DNF.

The only formula whose length may grow exponentially in the length of $\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})$ in the above transformation is the finite domain axiom. Instead of adding this axiom to the formula it could as well be built into a decision procedure for the BS(R) Fragment. Although this is subject to future work, one line of extension for resolution refutation building procedures, for example, is to start with a finite domain axiom for a small number of constants. In case of a refutation, extend the refutation derivation and the involved finite domain axiom by literals with further constants, until the overall limit $m_* |\vec{\mathbf{y}}| + |\mathrm{consts}(\psi)|$ is reached or a model is found. In case the finite domain axiom is a priori small, an explicit instantiation of $\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})$ can support the efficiency of an automated reasoning procedure. Explicit instantiation is the typical method used by SMT solvers when confronted with quantification [17]. In general, an exponential number of domain elements can be necessary for finding a model of a separated sentence. So the exponential growth cannot be escaped.

Lemma 23. The following sentences are equisatisfiable. (We assume that the f_i are fresh Skolem functions of appropriate arity, the R_i are fresh predicate symbols, and the constant symbols $c_{\ell,i}$ do not occur in ψ.)

(1) $\forall \vec{\mathbf{x}} \exists \vec{\mathbf{y}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}) \wedge \bigvee_{\ell=1}^{m_*} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} y_i \approx c_{\ell,i}$,

(2) $\forall \vec{\mathbf{x}}.\psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}) \big[y_1/f_1(\vec{\mathbf{x}}), \ldots, y_{|\vec{\mathbf{y}}|}/f_{|\vec{\mathbf{y}}|}(\vec{\mathbf{x}}) \big] \wedge \bigvee_{\ell=1}^{m_*} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} f_i(\vec{\mathbf{x}}) \approx c_{\ell,i}$,

(3) $\forall \vec{\mathbf{x}}\vec{\mathbf{y}}.\left(\left(\bigwedge_{i=1}^{|\vec{\mathbf{y}}|} f_i(\vec{\mathbf{x}}) \approx y_i \right) \to \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}}) \right) \wedge \bigvee_{\ell=1}^{m_*} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} f_i(\vec{\mathbf{x}}) \approx c_{\ell,i}$,

(4) $\forall \vec{\mathbf{x}}\vec{\mathbf{y}}.\left(\left(\bigwedge_{i=1}^{|\vec{\mathbf{y}}|} R_i(\vec{\mathbf{x}}, y_i)\right) \to \psi(\vec{\mathbf{x}}, \vec{\mathbf{y}})\right) \wedge \bigvee_{\ell=1}^{m_*} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} R_i(\vec{\mathbf{x}}, c_{\ell,i})$.

Proof. (2) is the Skolemization of (1). Equisatisfiability of the sentences in (2) and (3) is obvious—they are even equivalent. It remains to prove equisatisfiability of the sentences in (3) and (4). Let us denote the former by φ_f and the latter by φ_R.

It is easy to convert any model $\mathcal{A}$ of φ_f into a model $\mathcal{A}'$ of φ_R: simply take the function graphs of the $f_i^{\mathcal{A}}$ as the interpretation of the R_i under $\mathcal{A}'$.

Conversely, let $\mathcal{B}'$ be a model of φ_R. We define a structure $\mathcal{B}$ exactly like $\mathcal{B}'$, except for the interpretations of the f_i and R_i. For every $\vec{\mathfrak{a}} \in \mathfrak{U}_{\mathcal{B}}^{|\vec{\mathbf{x}}|}$ we pick one index $\ell \in [m_*]$ such that for every $i = 1, \ldots, |\vec{\mathbf{y}}|$ it holds $\langle \vec{\mathfrak{a}}, c_{\ell,i}^{\mathcal{B}'} \rangle \in R_i^{\mathcal{B}'}$, and set $f_i^{\mathcal{B}}(\vec{\mathfrak{a}}) := c_{\ell,i}^{\mathcal{B}'}$ (which equals $c_{\ell,i}^{\mathcal{B}}$). Such an index ℓ must exists because of $\mathcal{B}' \models \forall \vec{\mathbf{x}}. \bigvee_{\ell=1}^{m_*} \bigwedge_{i=1}^{|\vec{\mathbf{y}}|} R_i(\vec{\mathbf{x}}, c_{\ell,i})$. To finish the construction of $\mathcal{B}$, we define $R_i^{\mathcal{B}}$ to be the function graph of the newly defined function $f_i^{\mathcal{B}}$ for every index i. It is easy to check that $\mathcal{B}$ is a model of both φ_R and φ_f. $\square$

6 Related and Future Work

Dreben and Goldfarb [12] (page 65) extend the Relational Monadic Fragment to a certain extent in the direction of the BS Fragment and call the result *Initially-extended Essentially Monadic Class*. In essence, they allow constants and discard the restriction to unary predicate symbols. However, they require that every atom contains at most one variable (possibly with multiple occurrences). Consequently, their fragment does not fully include BS.

A broad overview over decidable standard fragments of first-order logic is given in [7]. More recent decidability results are often formulated as syntactic restrictions on clause sets [33, 14, 22] that are incomparable to or subsumed by SF.

There is recent work [15, 34, 10, 29, 9] which considers the BS(R) Fragment in settings beyond pure first-order logic. It might be of interest to investigate all those combinations and extensions based on SF instead of the BS(R) Fragment. Furthermore, SF might also be meaningful to first-order theories over fixed structures such as arithmetic.

Our results on the complexity of satisfiability of SF sentences left some gaps open, which remain to be closed in the future. Moreover, the syntactic restrictions on SF sentences may be weakened and still lead to a decidable fragment.

References

[1] G. Alagi and C. Weidenbach. NRCL – A Model Building Approach to the Bernays–Schönfinkel Fragment. In *Frontiers of Combining Systems (FroCoS'15)*, LNCS 9322, pages 69–84. Springer, 2015.

[2] L. Bachmair, H. Ganzinger, and A. Voronkov. Elimination of Equality via Transformation with Ordering Constraints. In *Automated Deduction (CADE-15)*, LNCS 1421, pages 175–190. Springer, 1998.

[3] P. Baumgartner, A. Fuchs, and C. Tinelli. Lemma Learning in the Model Evolution Calculus. In *Logic for Programming, Artificial Intelligence, and Reasoning (LPAR-13)*, LNCS 4246, pages 572–586. Springer, 2006.

[4] P. Baumgartner, A. Fuchs, H. de Nivelle, and C. Tinelli. Computing finite models by reduction to function-free clause logic. *J. Applied Logic*, 7(1):58–74, 2009.

[5] P. Bernays and M. Schönfinkel. Zum Entscheidungsproblem der mathematischen Logik. *Mathematische Annalen*, 99(1):342–372, 1928.

[6] M. P. Bonacina and D. A. Plaisted. SGGS Theorem Proving: an Exposition. In *Workshop on Practical Aspects in Automated Reasoning (PAAR'14)*, EPiC 31, pages 25–38. EasyChair, 2014.

[7] E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Springer, 1997.

[8] D. Brand. Proving Theorems with the Modification Method. *SIAM J. Comput.*, 4(4):412–430, 1975.

[9] D. Bresolin, D. Della Monica, A. Montanari, and G. Sciavicco. The light side of interval temporal logic: the Bernays–Schönfinkel fragment of CDT. *Annals of Mathematics and Artificial Intelligence*, 71(1-3):11–39, 2014.

[10] W. Charatonik and P. Witkowski. On the Complexity of the Bernays–Schönfinkel Class with Datalog. In *Logic for Programming, Artificial Intelligence, and Reasoning (LPAR-17)*, LNCS 6397, pages 187–201. Springer, 2010.

[11] K. Claessen and A. Lillieström. Automated Inference of Finite Unsatisfiability. *J. Autom. Reasoning*, 47(2):111–132, 2011.

[12] B. Dreben and W. D. Goldfarb. *The Decision Problem: Solvable Classes of Quantificational Formulas*. Addison-Wesley, 1979.

[13] H. Ebbinghaus, J. Flum, and W. Thomas. *Mathematical Logic*. Springer, 1994.

[14] C. G. Fermüller, A. Leitsch, U. Hustadt, and T. Tammet. Resolution Decision Procedures. In *Handbook of Automated Reasoning*, pages 1791–1849. Elsevier and MIT Press, 2001.

[15] P. Fontaine. Combinations of theories for decidable fragments of first-order logic. In *Frontiers of Combining Systems (FroCoS'09)*, LNCS 5749, pages 263–278. Springer, 2009.

[16] H. Ganzinger and K. Korovin. New Directions in Instantiation-Based Theorem Proving. In *Logic in Computer Science (LICS'03)*, pages 55–64. IEEE, 2003.

[17] Y. Ge and L. M. de Moura. Complete Instantiation for Quantified Formulas in Satisfiabiliby Modulo Theories. In *Computer Aided Verification (CAV'09)*, LNCS 5643, pages 306–320. Springer, 2009.

[18] Y. Gurevich. The Decision Problem for the Logic of Predicates and Operations. *Algebra i Logika*, 8:284–308, 1969.

[19] L. Henkin. Some remarks on infinitely long formulas. In *Infinistic Methods*, pages 167–183. Pergamon Press, 1961.

[20] T. Hillenbrand and C. Weidenbach. Superposition for bounded domains. In *Automated Reasoning and Mathematics – Essays in Memory of William W. McCune*, LNCS 7788, pages 68–100. Springer, 2013.

[21] K. Korovin. From Resolution and DPLL to Solving Arithmetic Constraints. In *Frontiers of Combining Systems (FroCoS'13)*, LNCS 8152, pages 261–262. Springer, 2013.

[22] K. Korovin. Non-cyclic Sorts for First-Order Satisfiability. In *Frontiers of Combining Systems (FroCoS'13)*, LNCS 8152, pages 214–228. Springer, 2013.

[23] R. Letz and G. Stenz. Integration of Equality Reasoning into the Disconnection Calculus. In *Automated Reasoning with Analytic Tableaux and Related Methods (TABLEAUX'02)*, LNAI 2381, pages 176–190. Springer, 2002.

[24] H. R. Lewis. Complexity Results for Classes of Quantificational Formulas. *J. Comput. Syst. Sci.*, 21(3):317–353, 1980.

[25] M. Löb. Decidability of the monadic predicate calculus with unary function symbols. *J. Symb. Logic*, 32:563, 1967.

[26] L. Löwenheim. Über Möglichkeiten im Relativkalkül. *Mathematische Annalen*, 76:447–470, 1915.

[27] A. Nonnengart and C. Weidenbach. Computing Small Clause Normal Forms. In *Handbook of Automated Reasoning*, pages 335–367. Elsevier and MIT Press, 2001.

[28] R. Piskac, L. M. de Moura, and N. Bjørner. Deciding Effectively Propositional Logic Using DPLL and Substitution Sets. *J. Autom. Reasoning*, 44(4):401–424, 2010.

[29] A. Policriti and E. Omodeo. The Bernays–Schönfinkel–Ramsey class for set theory: decidability. *J. Symb. Logic*, 77:896–918, 2012.

[30] F. P. Ramsey. On a Problem of Formal Logic. *Proceedings of The London Mathematical Society*, s2-30:264–286, 1930.

[31] E. Sperner. Ein Satz über Untermengen einer endlichen Menge. *Mathematische Zeitschrift*, 27(1):544–548, 1928.

[32] T. Sturm, M. Voigt, and C. Weidenbach. Deciding First-Order Satisfiability when Universal and Existential Variables are Separated. In *Logics in Computer Science (LICS'16), to appear*. IEEE/ACM, 2016.

[33] C. Weidenbach. Towards an Automatic Analysis of Security Protocols in First-Order Logic. In *Automated Deduction (CADE-16)*, LNCS 1632, pages 314–328. Springer, 1999.

[34] T. Wies, R. Piskac, and V. Kuncak. Combining Theories with Shared Set Operations. In *Frontiers of Combining Systems (FroCoS'09)*, LNCS 5749, pages 366–382. Springer, 2009.