

Separateness of Variables

A Novel Perspective on Decidable First-Order Fragments

Marco Voigt

*Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany,
Saarbrücken Graduate School of Computer Science*

Abstract

The classical decision problem, as it is understood today, is the quest for a delineation between the decidable and the undecidable parts of first-order logic based on elegant syntactic criteria. In this paper, we treat the concept of separateness of variables and explore its applicability to the classical decision problem. Two disjoint sets of first-order variables are separated in a given formula if variables from the two sets never co-occur in any atom of that formula. This simple notion facilitates extending many well-known decidable first-order fragments significantly and in a way that preserves decidability. We will demonstrate that for several prefix fragments, several guarded fragments, the two-variable fragment, and for the fluted fragment. Altogether, we will investigate nine such extensions more closely. Interestingly, each of them contains the relational monadic first-order fragment without equality. Although the extensions exhibit the same expressive power as the respective originals, certain logical properties can be expressed much more succinctly. In three cases the succinctness gap cannot be bounded using any elementary function.

1 Introduction

In the early twentieth century David Hilbert initiated his famous program striving for a formalization of the foundations of mathematics. At its core lay the *classical decision problem* of first-order logic: Find an algorithm that determines the validity of any given first-order sentence. Nowadays, the classical decision problem is understood as the problem of classifying first-order logic into fragments with a decidable or undecidable satisfiability problem.¹ This quest has produced a wealth of positive and also negative results, see [18, 11, 21] for references. We will give definitions of a number of well-known decidable fragments in later sections. More recently identified decidable fragments that are neither covered in the present paper nor in the mentioned texts are treated in [62, 41, 49, 61, 12]. A much more detailed but still brief overview of past and recent developments in the area is given in the beginning of Chapter 3 of [69].


In the present paper we explore the concept of *separateness of variables* in the context of the classical decision problem. Two sets X, Y of variables are *separated in a formula*, if there are no co-occurrences of any $x \in X$ and $y \in Y$ in any atom. Put more formally, we get the following definition.

Definition 1.1 (Separateness of variables). Let φ be any first-order formula and let X, Y be two disjoint sets of first-order variables. We say that X and Y are *separated in φ* if for every atom A occurring in φ we have $\text{vars}(A) \cap X = \emptyset$ or $\text{vars}(A) \cap Y = \emptyset$ or both.

This simple concept enables us to elegantly define nontrivial extensions of well-known decidable fragments of first-order logic. In essence, the definitions of the new fragments are careful

¹For convenience, we will be less precise every now and then and speak of *(un)decidable fragments*.

combinations of the concepts used in the original definitions with the concept of separateness of quantified variables. All new fragments still have a decidable satisfiability problem. Figure 1 provides a schematic overview of the most important fragments we will introduce. All of them



MFO – monadic first-order fragment
 BSR – Bernays–Schönfinkel–Ramsey fragment
 FO² – two-variable fragment
 AF – Ackermann fragment
 GKS – Gödel–Kalmár–Schütte fragment
 FL – fluted fragment
 GF – guarded fragment
 LGF – loosely guarded fragment
 GNFO – guarded negation fragment

SF – separated fragment
 SBSR – separated BSR
 SFO² – separated FO²
 SAF – separated AF
 SGKS – separated GKS
 SFL – separated FL
 SGF – separated GF
 SLGF – separated LGF
 SGNFO – separated GNFO

Figure 1: Left-hand side: Schematic overview of well-known decidable fragments of first-order logic. Only the partial overlaps between MFO and the other fragments are depicted. We neglect any other partial overlaps. Moreover, the containment of AF in GKS and of GF in LGF is shown. Right-hand side: Schematic overview of the extended fragments (in green) that are treated in the present paper. Notice that MFO is properly contained in all extended fragments. The focus is again on the overlaps with MFO and on the proper containment relations. Other depicted overlaps might be unsubstantiated.

enjoy the *finite model property*, i.e. every satisfiable sentence in such a fragment has a finite model. This is a sufficient condition for decidability of the associated satisfiability problem. If we can derive a computable upper bound regarding the size of smallest models, we speak of a *small model property*.

From a qualitative point of view, the extended fragments do not come with an increased expressiveness compared to the original fragments. This is witnessed by the existence of equivalence-preserving translations from each extended fragment into the respective original. However, we will show for several cases that there may be significant gaps regarding the length of shortest formulas that express one and the same property. For example, the succinctness gap between the well-known *Bernays–Schönfinkel–Ramsey fragment* (the $\exists^*\forall^*$ prefix class) and an extension of it called the *separated fragment* [65] (all relational sentences $\exists z\forall\bar{x}_1\exists\bar{y}_1\dots\forall\bar{x}_n\exists\bar{y}_n.\varphi$ in which $\bar{x}_1\cup\dots\cup\bar{x}_n$ and $\bar{y}_1\cup\dots\cup\bar{y}_n$ are separated) is as follows [68]. For every positive natural number n we can find some property that can be expressed in the separated fragment with a formula of

length $k \cdot n^2$ for some fixed natural number k , whereas expressing the very same property in the Bernays–Schönfinkel–Ramsey fragment requires a formula whose length is at least

$$\left. \begin{array}{c} \vdots^{2^2} \\ 2^2 \end{array} \right\} \text{height } n .$$

Hence, some of the extended fragments enable us to describe certain logical properties much more succinctly and elegantly. Table 1 summarizes the succinctness gaps that we will derive.

Table 1: Summary of the succinctness gaps that are explored in the present paper. The abbreviations for fragments are spelled out in Figure 1. Notice that $\text{SF} \subseteq \text{SBSR}$. The gap between SAF and AF is conditional on $\text{EXPTIME} \neq \text{NEXPTIME}$. All other gaps are unconditional.

More succinct fragments	Less succinct fragments	Succinctness gaps	Reference
SF / SBSR	BSR	non-elementary	Theorem 3.8
SF / SBSR	Gaifman-local first-order fragment	non-elementary	Prop. 3.12
SAF	AF	super-polynomial	Prop. 5.12
SGKS	GKS	exponential	Theorem 6.8
SGF	LGF	non-elementary	Theorem 7.5
SGNFO	GNFO	non-elementary	Theorem 8.6
SFO ²	FO ²	exponential	Theorem 9.5

Non-trivial cases of separateness appear, for instance, in formulas where universal and existential quantifiers are nested and the variables they bind are separated. Consider the sentence $\varphi := \forall x \exists y. P(x) \leftrightarrow Q(y)$ in which the singleton sets $\{x\}$ and $\{y\}$ are obviously separated. It expresses a certain symmetry in structures \mathcal{A} . For every domain element \mathbf{a} there is some element \mathbf{b} such that \mathbf{a} belongs to $P^{\mathcal{A}}$ if and only if \mathbf{b} belongs to $Q^{\mathcal{A}}$. It turns out that the same property can be expressed without any nesting of alternating quantifiers. Indeed, the distributivity laws of Boolean algebra and quantifier shifting (cf. Proposition 2.1) facilitate a transformation of φ into the equivalent sentence $\psi := ((\exists x. P(x)) \rightarrow (\exists y_1. Q(y_1))) \wedge ((\exists x. \neg P(x)) \rightarrow (\exists y_2. \neg Q(y_2)))$ — the symbol \models denotes semantic equivalence:

$$\begin{aligned} \varphi &= \forall x \exists y. P(x) \leftrightarrow Q(y) \\ &\models \forall x \exists y. (\neg P(x) \vee Q(y)) \wedge (P(x) \vee \neg Q(y)) \\ &\models \forall x \exists y. (\neg P(x) \wedge \neg Q(y)) \vee (Q(y) \wedge P(x)) \\ &\models \forall x. (\neg P(x) \wedge (\exists y_2. \neg Q(y_2))) \vee ((\exists y_1. Q(y_1)) \wedge P(x)) \\ &\models ((\forall x. \neg P(x)) \vee (\exists y_1. Q(y_1))) \wedge ((\exists y_2. \neg Q(y_2)) \vee (\forall x. P(x))) \\ &\models ((\exists x. P(x)) \rightarrow (\exists y_1. Q(y_1))) \wedge ((\exists x. \neg P(x)) \rightarrow (\exists y_2. \neg Q(y_2))) = \psi \end{aligned}$$

We could even shift quantifiers outwards again and finally obtain an equivalent sentence with a $\exists \forall$ quantifier prefix: $\psi' := \exists y_1 y_2 \forall x. (P(x) \rightarrow Q(y_1)) \wedge (\neg P(x) \rightarrow \neg Q(y_2))$. In this example we can not only transform nested quantification of separated variables into quantification that is not nested. In addition, we can replace $\forall \exists$ alternations in exchange for $\exists \forall$ alternations, or vice versa. In general, succinct representations of certain logical properties can be unfolded into more verbose ones that require a lower *quantifier rank* (i.e. a lower nesting depth of quantifiers) or even use fewer quantifier alternations. The sentence ψ can, using a $\forall \exists$ alternation, represent the symmetry property of structures more succinctly than the sentence ψ' can with an $\exists \forall$ quantifier alternation. This is a key feature that we will observe for several of the extended fragments. We will also see that such succinctness gaps can become k -fold exponential, if we start from k nested $\forall \exists$ alternations.

The following list summarizes the main contributions of the present paper:

- (1) Eight novel decidable fragments of first-order logic are introduced that extend well-known decidable first-order fragments: the *separated Bernays–Schönfinkel–Ramsey fragment (SBSR)*, the *separated Ackermann fragment (SAF)*, the *separated Gödel–Kalmár–Schütte fragment (SGKS)*, the *separated guarded fragment (SGF)*, the *separated loosely guarded fragment (SLGF)*, the *separated guarded-negation fragment (SGNFO)*, the *separated two-variable fragment (SFO²)*, and the *separated fluted fragment (SFL)*. Moreover, it is proved that the qualitative expressiveness of each extended fragment coincides with the respective original.
- (2) Significant gaps regarding succinctness are derived for several of the extended fragments: SBSR, SGKS, SGF, SLGF, SGNFO, SFO², cf. Table 1. This evidently shows that several of the extended fragments constitute a substantial quantitative improvement regarding expressiveness compared to the original fragments.

Note to editors and reviewers: The present paper is (except for Theorem 8.6) a condensation of material taken from the author’s PhD thesis [69], mainly Chapter 3. Moreover, results from previous conference publications are recapitulated [65, 68] (with citations).

2 Preliminaries

We mainly rely on the standard notions from first-order logic. Nevertheless, we need to agree on some notation.

Syntax

A *vocabulary* Σ comprises a finite set of predicate symbols and function symbols, each equipped with its *arity*. A vocabulary Σ is *relational* if it exclusively contains predicate symbols. We define Σ -*terms* and Σ -*formulas* as usual, allowing the logical connectives \neg, \wedge, \vee and the first-order quantifiers \forall, \exists . The two additional connectives $\rightarrow, \leftrightarrow$ are used as shortcuts: $\varphi \rightarrow \psi$ abbreviates $(\neg\varphi) \vee \psi$ and $\varphi \leftrightarrow \psi$ abbreviates $(\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$. If not explicitly excluded, we allow equality. A Σ -*sentence* is a *closed* Σ -*formula*, i.e. one without free variables. A Σ -*formula* / Σ -*sentence* is *relational*, if Σ is relational. If there is no danger of confusion, we drop the explicit reference to Σ and just speak of *terms*, *formulas*, and *sentences*. In order to save parentheses, we follow the convention that negation binds strongest, that conjunction binds stronger than disjunction, and that all of the aforementioned bind stronger than implication and equivalence. Equivalence, in turn, binds weaker than implication. The scope of quantifiers will stretch as far to the right as admitted by parentheses. In all formulas we tacitly assume, if not explicitly stated otherwise, that no variable occurs freely and bound at the same time and that all distinct occurrences of quantifiers bind distinct variables. We use $\varphi(v_1, \dots, v_m)$ to denote a formula φ whose free first-order variables form a subset of $\{v_1, \dots, v_m\}$. The variables v_1, \dots, v_m are assumed to be pairwise distinct. For any formula φ we denote by $\text{vars}(\varphi)$ the set of all variables occurring freely or bound in φ .

Given a set Φ of Σ -formulas, we call a Σ -formula φ a *Boolean combination of formulas from* Φ , if φ consist of formulas from Φ , possibly connected via the Boolean connectives $\neg, \wedge, \vee, \rightarrow, \leftrightarrow$. If we restrict the set of Boolean connectives even further to \wedge, \vee , we speak of a *\wedge - \vee -combination of formulas from* Φ .

A *quantifier block* is a maximal sequence $Q\bar{v} = Qv_1.Qv_2.\dots.Qv_n$ of quantifiers of the same kind occurring in a given formula. For convenience, we often identify tuples \bar{v} of variables with the set containing all the variables that occur in \bar{v} . We occasionally also use regular expressions to describe sequences of quantifiers. For example, for any positive integer k the expression $\exists^*\forall^k\exists\exists$ stands for the set of all prefixes of the form $\exists y_1 \dots y_m \forall x_1 \dots x_k \exists z_1 z_2$, where $m = 0$ is allowed.

A formula is in *prenex normal form* if all quantifiers are lined up in front of the formula, i.e. it has the shape $Q_1v_1 \dots Q_nv_n.\psi$ with quantifier-free ψ and $Q_i \in \{\forall, \exists\}$. A formula is in *negation normal form* if it exclusively contains the connectives \wedge, \vee, \neg and every negation sign occurs immediately in front of an atom; quantifiers are of course admitted.

We use a standard notion of term and formula length, denoted $\text{len}(\varphi)$. Notice that we have $\text{len}(\varphi \rightarrow \psi) = \text{len}(\neg\varphi \vee \psi)$ and $\text{len}(\varphi \leftrightarrow \psi) = \text{len}((\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi))$.

Semantics

As usual, we interpret (Σ) -formulas with respect to (Σ) -structures \mathcal{A} , consisting of a nonempty domain A and interpretations $f^{\mathcal{A}}$ and $P^{\mathcal{A}}$ of all function and predicate symbols in the underlying vocabulary. Given a term s , a structure \mathcal{A} , and a variable assignment β over \mathcal{A} 's domain, we denote the *evaluation of s under \mathcal{A} and β* by $\mathcal{A}(\beta)(s)$. It is defined as usual. We simply write $\mathcal{A}(s)$ if s is variable free. We write $\mathcal{A}, \beta \models \varphi$ if φ is *satisfied under \mathcal{A} and β* in the usual sense. When there is no danger of confusion, we conveniently abbreviate expressions of the form $\mathcal{A}, [v_1 \mapsto \mathbf{a}_1, \dots, v_m \mapsto \mathbf{a}_m] \models \varphi(v_1, \dots, v_m)$ by $\mathcal{A} \models \varphi(\mathbf{a}_1, \dots, \mathbf{a}_m)$. We write $\mathcal{A} \models \varphi$ if $\mathcal{A}, \beta \models \varphi$ holds for every variable assignment β over \mathcal{A} 's domain. In such cases, we say that \mathcal{A} is a *model* of φ . For sentences φ we say that \mathcal{A} *satisfies* φ if $\mathcal{A}, \beta \models \varphi$ for any β . Two sentences φ and ψ are considered *equisatisfiable* if φ has a model if and only if ψ has one.

We also use the symbol \models to denote *semantic entailment* of two Σ -formulas, i.e. we have $\varphi \models \psi$ whenever for every Σ -structure \mathcal{A} and every variable assignment β , $\mathcal{A}, \beta \models \varphi$ implies $\mathcal{A}, \beta \models \psi$. The symbol \equiv denotes *semantic equivalence* of formulas, i.e. $\varphi \equiv \psi$ holds whenever $\varphi \models \psi$ and $\psi \models \varphi$.

We will frequently use the fact that quantifiers can be *shifted* in certain ways within formulas under preservation of the formula's semantics.

Proposition 2.1 (Quantifier shifting). Let φ, ψ, χ be formulas, and assume that x and y do not occur freely in χ . We have the following equivalences, where $\circ \in \{\wedge, \vee\}$:

$$\begin{array}{ll} \text{(i)} & \exists y. (\varphi \vee \psi) \equiv (\exists y. \varphi) \vee (\exists y. \psi) & \text{(ii)} & \forall x. (\varphi \wedge \psi) \equiv (\forall x. \varphi) \wedge (\forall x. \psi) \\ \text{(iii)} & \exists y. (\varphi \circ \chi) \equiv (\exists y. \varphi) \circ \chi & \text{(iv)} & \forall x. (\varphi \circ \chi) \equiv (\forall x. \varphi) \circ \chi \\ \text{(v)} & \exists y_1 \exists y_2. \varphi \equiv \exists y_2 \exists y_1. \varphi & \text{(vi)} & \forall x_1 \forall x_2. \varphi \equiv \forall x_2 \forall x_1. \varphi \end{array}$$

A structure \mathcal{A} is a *substructure* of a structure \mathcal{B} (over the same vocabulary) if (1) $A \subseteq B$, (2) $c^{\mathcal{A}} = c^{\mathcal{B}}$ for every constant symbol c , (3) $P^{\mathcal{A}} = P^{\mathcal{B}} \cap A^m$ for every m -ary predicate symbol P , and (4) $f^{\mathcal{A}}(\bar{\mathbf{a}}) = f^{\mathcal{B}}(\bar{\mathbf{a}})$ for every m -ary function symbol f and every m -tuple $\bar{\mathbf{a}} \in A^m$. Given a structure \mathcal{A} and some subset S of \mathcal{A} 's domain, the *substructure of \mathcal{A} induced by S* is the unique substructure \mathcal{B} of \mathcal{A} with the domain $B := S$. The following is a standard lemma (see, e.g., [19], Lemma 5.7 in Chapter III).

Lemma 2.2 (Substructure Lemma). Let φ be a first-order sentence without existential quantifiers and in which no universal quantifier lies within the scope of any negation sign — we treat any subformula $\varphi_1 \rightarrow \varphi_2$ as abbreviation for $\neg\varphi_1 \vee \varphi_2$ and any subformula $\varphi_1 \leftrightarrow \varphi_2$ as abbreviation for $(\neg\varphi_1 \vee \varphi_2) \wedge (\varphi_1 \vee \neg\varphi_2)$ to account for implicit negation signs as well. Moreover, let \mathcal{A} be a substructure of \mathcal{B} . If $\mathcal{B} \models \varphi$, then $\mathcal{A} \models \varphi$.

Additional Notation

We use the notation $[k]$ to abbreviate the set $\{1, \dots, k\}$ for any positive integer k . The *power set* of a set S , i.e. the set of all subsets of S , is denoted by $\mathcal{P}(S)$. The iterated application of \mathcal{P} is given by $\mathcal{P}^0(S) := S$ and $\mathcal{P}^{k+1}(S) := \mathcal{P}^k(\mathcal{P}(S))$ for $k \geq 0$. Furthermore, we also define the *tetration operation* inductively by $2^{\uparrow 0}(m) := m$ and $2^{\uparrow k+1}(m) := 2^{(2^{\uparrow k}(m))}$.

3 The Separated Fragment (SF)

Before we define the separated fragment, we first recall some details about the two well-known decidable fragments it extends. The *monadic first-order fragment (MFO)* comprises all relational first-order sentences without equality that contain only unary predicate symbols. When we refer to the *monadic first-order fragment with equality*, we use the abbreviation MFO_{\approx} . The decidability of

the respective satisfiability problems MFO-Sat and MFO_{\approx} -Sat was proved in [45, 64, 8]. Several decades later, Löb [44] and Gurevich [34] extended the positive result to monadic first-order sentences with unary function symbols but without equality, the *Löb–Gurevich fragment*. MFO and MFO_{\approx} possess the exponential model property: every satisfiable MFO_{\approx} sentence φ has a model whose domain size is at most exponential in φ 's length. Moreover, satisfiability for MFO and MFO_{\approx} is NEXPTIME -complete, cf. [11].

The *Bernays–Schönfinkel–Ramsey fragment (BSR)* comprises all relational first-order sentences in prenex normal form with an $\exists^*\forall^*$ quantifier prefix and with equality. Bernays and Schönfinkel [9] showed that satisfiability for the relational $\exists^*\forall^*$ prefix class without equality is decidable. Following up, Ramsey [57] obtained a positive decidability result in the presence of equality. This extended class is known to possess a linear model property (cf. Proposition 3.5), and BSR-Sat is known to be complete for NEXPTIME .

Now we are ready for defining the *separated fragment, SF* for short. Technically, it is defined as a class of prenex sentences, but this is not an essential property. The defining principle of SF sentences is simply that co-occurrences of universally and existentially quantified variables in atoms are forbidden. Existential variables quantified by *leading* existential quantifiers are exempt from this rule. We consider an existential quantifier *leading* if it does not lie within the scope of any universal quantifier.

Definition 3.1 (Separated fragment (SF), [65]). The *separated fragment (SF)* consists of all relational first-order sentences φ with equality of the form $\exists \bar{z} \forall \bar{x}_1 \exists \bar{y}_1 \dots \forall \bar{x}_n \exists \bar{y}_n. \psi$ with quantifier-free ψ in which the sets $\bar{x} := \bar{x}_1 \cup \dots \cup \bar{x}_n$ and $\bar{y} := \bar{y}_1 \cup \dots \cup \bar{y}_n$ are separated. The tuples \bar{z} and \bar{y}_n may be empty, i.e. the quantifier prefix does not have to start with an existential quantifier and it does not have to end with an existential quantifier either.

Recall that \bar{x} and \bar{y} are separated in φ if and only if for every atom A occurring in φ we either have $\text{vars}(A) \cap \bar{x} = \emptyset$ or $\text{vars}(A) \cap \bar{y} = \emptyset$. Moreover, notice that the variables in \bar{z} are not subject to any restriction regarding their occurrences.

It is not hard to see that SF is a proper syntactic extension of BSR. Clearly, the quantified variables in every BSR sentence $\varphi := \exists \bar{z} \forall \bar{x}. \psi$ with quantifier-free ψ trivially satisfy the separateness conditions imposed by Definition 3.1, as \bar{y} is empty. Similarly, every monadic sentence without equality in prenex normal form trivially satisfies the conditions of Definition 3.1, because any monadic atom contains at most one first-order variable. Since any MFO sentence can easily be transformed into an equivalent sentence in prenex normal form, it is fair to say that SF also contains MFO. On the other hand, it is an easy task to find sentences that belong to SF but neither to BSR nor to MFO, e.g. $\forall x_1 x_2 \exists y_1 y_2. P(x_1, x_2) \leftrightarrow Q(y_1, y_2)$.

Proposition 3.2. SF properly contains BSR and MFO.

Another interesting question is whether MFO_{\approx} is subsumed by SF. For instance, the sentence $\forall x \exists y. x \approx y$ is in MFO_{\approx} but violates the separateness conditions of SF. Therefore, from the syntactic point of view, there are MFO_{\approx} sentences whose variables are not sufficiently separated for SF. However, the sentence $\forall x \exists y. x \approx y$ is equivalent to $\forall x. x \approx x$, which even belongs to BSR. Similarly, we have the MFO_{\approx} sentence $\forall x \exists y. x \not\approx y$, which is not in SF but equivalent to the BSR sentence $\exists y_1 y_2. y_1 \not\approx y_2$. The following proposition witnesses that this is by no means a coincidence. As one consequence, speaking in terms of expressiveness, MFO_{\approx} is subsumed by BSR and, hence, also by SF.

Proposition 3.3. For every MFO_{\approx} sentence there is an equivalent BSR sentence.

The proof of this result can be found in [69] (Theorem 3.1.5). It is based on techniques described by Behmann [8] in the context of second-order quantifier elimination for the monadic second-order fragment — see [70], Section 13.2, for a modern account.

3.1 Translation of SF into BSR: Upper and Lower Bounds

It was first proved in [65] that every SF sentence can be transformed into an equivalent BSR sentence. We will present such a translation for the more general case of SBSR sentences in Section 4.1 (Lemma 4.4 and Theorem 4.5). For SF we simply state the result here.

Theorem 3.4 ([68], Lemma 12). Every SF sentence φ with k $\forall\exists$ alternations is equivalent to some BSR sentence whose length is at most k -fold exponential in φ 's length.

It is worth noting that [68] uses a much more fine-grained measure than the number of $\forall\exists$ alternations. Even more details can be found in [69], Section 3.2.

Translations like the one from SF sentences to equivalent BSR sentences lead to large blowups in the worst case. In general, the length of formulas has a significant effect on the size of smallest models. For BSR this relation is linear.

Proposition 3.5 (cf. Proposition 6.2.17 in [11]). Let $\varphi := \exists\bar{z}\forall\bar{x}.\psi$ be a satisfiable BSR sentence with quantifier-free ψ , containing (at most) k constant symbols. There is a model $\mathcal{A} \models \varphi$ such that $|\mathbf{A}| \leq \max(|\bar{z}| + k, 1)$.

It was proved in [68] (Lemma 24) that for satisfiable SF sentences φ the size of smallest models cannot be bounded by any tower of exponents $2^{\cdot^{\cdot^{\cdot}}}$ of a fixed height. In other words, the asymptotic growth of the size of smallest models is non-elementary in the length of the regarded SF sentence.

The analysis conducted in [68] yields matching upper and lower bounds that are not formulated in terms of the number of quantifier alternations. Rather, the upper bounds are based on a measure of how much quantified variables interact in atoms. The motivation for this fine-grained analysis is that the upper bound is intended to also give a tight estimate for MFO sentences. It is well known that the size of small models for satisfiable MFO sentences is exponential in their length no matter how many quantifier alternations are present (cf. Proposition 6.2.1 in [11]). Due to space limitations, we will not elaborate any further on this topic. The interested reader will find more details in [68] and in [69], Sections 3.2 and 3.5.

Put together, Theorem 3.4 and Proposition 3.5 immediately entail the following small model property for SF.

Proposition 3.6 (Small model property for SF). Every satisfiable SF sentence φ with $k \geq 1$ $\forall\exists$ alternations has a model with at most $\text{len}(\varphi) + k \cdot (\text{len}(\varphi))^2 \cdot (2^{\uparrow k}(\text{len}(\varphi)))^k$ domain elements.

It is well known that a small model property leads to decidability of the associated satisfiability problem (see, e.g. Proposition 6.0.4 in [11]). Since BSR enjoys a linear model property, even if constant symbols are allowed in the syntax, the separated fragment immediately inherits this property. Hence, the satisfiability problem for SF (*SF-Sat*) is decidable.

Proposition 3.7. *SF-Sat* is decidable, even if we allow constant symbols in SF sentences.

The unrestricted presence of function symbols in SF would lead to an undecidable satisfiability problem. Nevertheless, SF could easily be extended so far that it also subsumes the Löb-Gurevich fragment (see the beginning of Section 3 for the definition) while retaining decidability (see [69], Section 3.14.1).

Next, we complement the upper bound from Theorem 3.4 with a corresponding non-elementary lower bound.

Theorem 3.8 ([68], Theorem 16). There is a class of satisfiable SF sentences such that for every positive integer n the class contains a sentence φ with n $\forall\exists$ quantifier alternations and with a length polynomial in n for which any equivalent BSR sentence contains at least $\sum_{k=1}^n 2^{\uparrow k}(n)$ leading existential quantifiers.

Proof. Let $n \geq 1$ be some positive integer. Consider the following first-order sentence in which the sets $\{x_1, \dots, x_n\}$ and $\{y_1, \dots, y_n\}$ are separated:

$$\varphi := \forall x_n \exists y_n \dots \forall x_1 \exists y_1 \cdot \bigwedge_{i=1}^{4n} (P_i(x_1, \dots, x_n) \leftrightarrow Q_i(y_1, \dots, y_n)) .$$

Notice the orientation of the indices in the quantifier prefix! Moreover, recall that $[m]$ with $m \geq 1$ denotes the set $\{1, \dots, m\}$. In order to construct a particular model of φ , we inductively define the following sets: $\mathcal{S}_1 := \{S \in \mathcal{P}([4n]) \mid |S| = 2n\}$, $\mathcal{S}_{k+1} := \{S \in \mathcal{P}(\mathcal{S}_k) \mid |S| = \frac{1}{2} \cdot |\mathcal{S}_k|\}$ for every $k > 1$. Then, we observe that

$$|\mathcal{S}_1| = \binom{4n}{2n} \geq \left(\frac{4n}{2n}\right)^{2n} = 2^{2n},$$

$$|\mathcal{S}_2| = \binom{|\mathcal{S}_1|}{|\mathcal{S}_1|/2} \geq \left(\frac{|\mathcal{S}_1|}{|\mathcal{S}_1|/2}\right)^{|\mathcal{S}_1|/2} = 2^{|\mathcal{S}_1|/2} \geq 2^{2^{2n-1}},$$

⋮

$$|\mathcal{S}_n| = \binom{|\mathcal{S}_{n-1}|}{|\mathcal{S}_{n-1}|/2} \geq 2^{|\mathcal{S}_{n-1}|/2} \geq 2^{2^{2^{2^{n-1}-1}}} \geq 2^{\uparrow n}(2n - (n-1)) = 2^{\uparrow n}(n+1),$$

where the inequality $\binom{n}{k} \geq (n/k)^k$ can be found in [14] (page 1097), for example.

Having the sets \mathcal{S}_k , we now define the structure \mathcal{A} as follows:

$$\mathbf{A} := \{\mathbf{a}_S^{(k)}, \mathbf{b}_S^{(k)} \mid 1 \leq k \leq n \text{ and } S \in \mathcal{S}_k\},$$

$$P_i^{\mathcal{A}} := \{\langle \mathbf{a}_{S_1}^{(1)}, \dots, \mathbf{a}_{S_n}^{(n)} \rangle \in \mathbf{A}^n \mid i \in S_1 \in S_2 \in \dots \in S_n\} \text{ for } i = 1, \dots, 4n, \text{ and}$$

$$Q_i^{\mathcal{A}} := \{\langle \mathbf{b}_{S_1}^{(1)}, \dots, \mathbf{b}_{S_n}^{(n)} \rangle \in \mathbf{A}^n \mid i \in S_1 \in S_2 \in \dots \in S_n\} \text{ for } i = 1, \dots, 4n.$$

For any choice of S_1, \dots, S_n and every i , $1 \leq i \leq 4n$, we then have

$$\mathcal{A} \models P_i(\mathbf{a}_{S_1}^{(1)}, \dots, \mathbf{a}_{S_n}^{(n)}) \leftrightarrow Q_i(\mathbf{b}_{S_1}^{(1)}, \dots, \mathbf{b}_{S_n}^{(n)}) .$$

For any other choice of tuples $\langle \mathbf{c}_1, \dots, \mathbf{c}_n \rangle$, i.e. if there do not exist sets $S_1 \in \mathcal{S}_1, \dots, S_n \in \mathcal{S}_n$ such that $\langle \mathbf{c}_1, \dots, \mathbf{c}_n \rangle$ equals $\langle \mathbf{a}_{S_1}^{(1)}, \dots, \mathbf{a}_{S_n}^{(n)} \rangle$ or $\langle \mathbf{b}_{S_1}^{(1)}, \dots, \mathbf{b}_{S_n}^{(n)} \rangle$, we observe $\mathcal{A} \not\models P_i(\mathbf{c}_1, \dots, \mathbf{c}_n)$ and $\mathcal{A} \not\models Q_i(\mathbf{c}_1, \dots, \mathbf{c}_n)$ for every i . Hence,

$$\mathcal{A} \models \bigwedge_{i=1}^{4n} P_i(\mathbf{c}_1, \dots, \mathbf{c}_n) \leftrightarrow Q_i(\mathbf{c}_1, \dots, \mathbf{c}_n) .$$

Consequently, \mathcal{A} is a model of φ .

Consider the following simple two-player game with Players \mathfrak{A} and \mathfrak{B} . In the first round \mathfrak{A} moves first by picking some domain element $\mathbf{a}_{S_{\mathfrak{A},n}}^{(n)}$ for some set $S_{\mathfrak{A},n} \in \mathcal{S}_n$. Player \mathfrak{B} answers by picking a domain element $\mathbf{b}_{S_{\mathfrak{B},n}}^{(n)}$ for some set $S_{\mathfrak{B},n} \in \mathcal{S}_n$. The game continues for $n-1$ more rounds, where in every round Player \mathfrak{A} picks a domain element $\mathbf{a}_{S_{\mathfrak{A},j}}^{(j)}$ with $S_{\mathfrak{A},j} \in \mathcal{S}_{\mathfrak{A},j+1}$ and \mathfrak{B} answers by picking some $\mathbf{b}_{S_{\mathfrak{B},j}}^{(j)} \in \mathcal{S}_{\mathfrak{B},j+1}$. Hence, in the last round the chosen domain elements $\mathbf{a}_{S_{\mathfrak{A},1}}^{(1)}$ and $\mathbf{b}_{S_{\mathfrak{B},1}}^{(1)}$ are such that $S_{\mathfrak{A},1}$ and $S_{\mathfrak{B},1}$ are both nonempty subsets of $[4n]$. Player \mathfrak{A} wins if and only if $\mathcal{A} \not\models P_i(\mathbf{a}_{S_1}^{(1)}, \dots, \mathbf{a}_{S_n}^{(n)}) \leftrightarrow Q_i(\mathbf{b}_{S_1}^{(1)}, \dots, \mathbf{b}_{S_n}^{(n)})$ for some $i \in [4n]$, and Player \mathfrak{B} wins if and only if $\mathcal{A} \models P_i(\mathbf{a}_{S_1}^{(1)}, \dots, \mathbf{a}_{S_n}^{(n)}) \leftrightarrow Q_i(\mathbf{b}_{S_1}^{(1)}, \dots, \mathbf{b}_{S_n}^{(n)})$ for every $i \in [4n]$. Since \mathcal{A} is a model of φ , there must exist a winning strategy for \mathfrak{B} .

Claim I: There is exactly one winning strategy for \mathfrak{B} , namely, for every $j = n, \dots, 1$ Player \mathfrak{B} picks the element $\mathbf{b}_{S_{\mathfrak{B},j}}^{(j)}$ in round $n-j+1$, i.e. for every j we have $S_{\mathfrak{B},j} = S_{\mathfrak{A},j}$.

Proof: It is easy to see that the described strategy is a winning strategy for \mathfrak{B} .

Assume \mathfrak{B} deviates from this strategy. More precisely, suppose there exists some j_* , $1 \leq j_* \leq n$, such that \mathfrak{B} did not adhere to the described strategy in the $(n - j_* + 1)$ st round, i.e. $S_{\mathfrak{B},j_*} \neq S_{\mathfrak{A},j_*}$. It can be shown by induction on j_* that \mathfrak{A} has a winning strategy from this deviation point on. ◇

Claim I would still hold true if we allowed \mathfrak{B} to freely pick any element of the domain A at every round. The reason is that for any choice of elements $a_{S_{\mathfrak{A},n}}^{(n)}, \dots, a_{S_{\mathfrak{A},1}}^{(1)}$ made by \mathfrak{A} with $S_{\mathfrak{A},1} \in \dots \in S_{\mathfrak{A},n} \in \mathcal{S}_n$ we know that $S_{\mathfrak{A},1}$ is nonempty. Hence, we can always find some $i_* \in S_{\mathfrak{A},1}$ such that $\mathcal{A} \models P_{i_*}(a_{S_{\mathfrak{A},1}}^{(1)}, \dots, a_{S_{\mathfrak{A},n}}^{(n)})$. On the other hand, for any sequence c_n, \dots, c_1 picked by \mathfrak{B} that does not comply with the rules of the described game, we have $\mathcal{A} \not\models Q_{i_*}(c_1, \dots, c_n)$. This result leads to the following observation.

Claim II: For any of the $\mathfrak{b}_S^{(k)}$ the substructure of \mathcal{A} induced by the domain $A \setminus \{\mathfrak{b}_S^{(k)}\}$ does not satisfy φ .

Proof: The reason is simply that in this case player \mathfrak{A} can always prevent \mathfrak{B} from reaching a state of the game where \mathfrak{B} can apply the described winning strategy. ◇

We have already analyzed the size of the sets \mathcal{S}_k . Due to the observed lower bounds, we know that A contains at least $\sum_{k=1}^n 2^{\uparrow k}(n)$ elements of the form $\mathfrak{b}_S^{(k)}$.

Next, we argue that any $\exists^* \forall^*$ -sentence φ_* that is semantically equivalent to φ must contain at least $\sum_{k=1}^n 2^{\uparrow k}(n)$ leading existential quantifiers. Let

$$\varphi_* := \exists y_1 \dots y_m \forall x_1 \dots x_\ell. \chi_*(y_1, \dots, y_m, x_1, \dots, x_n)$$

with quantifier-free χ_* be a sentence with minimal m that is semantically equivalent to φ . Since \mathcal{A} is also a model of φ_* , we know that there is a sequence of elements c_1, \dots, c_m taken from the domain A such that $\mathcal{A} \models \forall x_1 \dots x_\ell. \chi_*(c_1, \dots, c_m, x_1, \dots, x_n)$. Consequently, we can extend \mathcal{A} to a model \mathcal{A}_* (over the same domain) of the Skolemized sentence $\varphi_{\text{Sk}} := \forall x_1 \dots x_\ell. \chi_*[y_1/d_1, \dots, y_m/d_m]$ by adding $d_j^{A_*} := c_j$ for $j = 1, \dots, m$. On the other hand, every model of φ_{Sk} is also a model of φ_* . The vocabulary underlying φ_{Sk} comprises exactly the constant symbols d_1, \dots, d_m and does not contain any other function symbols. Suppose $m < \sum_{k=1}^n 2^{\uparrow k}(n)$. Hence, there is some $\mathfrak{b}_S^{(k)}$ with $S \in \mathcal{S}_k$ such that for none of the d_j we have $d_j^{A_*} = \mathfrak{b}_S^{(k)}$. Then, by the Substructure Lemma, the following substructure \mathcal{B} of \mathcal{A}_* constitutes a model of φ_{Sk} : $B := A_* \setminus \{\mathfrak{b}_S^{(k)}\}$, $P_i^{\mathcal{B}} := P_i^{A_*} \cap B^n = P_i^{A_*}$ and $Q_i^{\mathcal{B}} := Q_i^{A_*} \cap B^n$ for every i , and $d_j^{\mathcal{B}} := d_j^{A_*}$ for every j . But then \mathcal{B} must also be a model of both φ_* and φ , since every model of φ_{Sk} is a model of φ_* which we, in turn, assume to be equivalent to φ . This contradicts Claim II, and thus we must have $m \geq \sum_{k=1}^n 2^{\uparrow k}(n)$. □

Theorem 3.8 entails that there is no elementary upper bound on the length of the BSR sentences that result from any equivalence-preserving transformation of SF sentences into BSR. On the other hand, there is an elementary upper bound, if we only consider SF sentences with a bounded number of quantifier alternations (cf. Theorem 4.5). A special case of Theorem 3.8 highlights the difference in succinctness between BSR and MFO. The following proposition states that, in the worst case, there is an unavoidable exponential gap between the two fragments.

Proposition 3.9. There is a class of MFO sentences such that for every positive integer n the class contains a sentence φ of a length polynomial in n for which any equivalent BSR sentence contains at least 2^n leading existential quantifiers.

One witnessing class of sentences comprises $\forall x \exists y. \bigwedge_{i=1}^{2^n} (P_i(x) \leftrightarrow Q_i(y))$ for $n \geq 1$.

3.2 Expressiveness of SF

We have already seen that SF is at least as expressive as BSR, MFO, and MFO_{\approx} (cf. Propositions 3.2 and 3.3). Moreover, every \wedge - \vee -combination of sentences from BSR and/or MFO_{\approx} is equivalent to some SF sentence. On the other hand, Theorem 3.8 shows that SF sentences can be considerably more succinct than their BSR equivalents.

Whenever it is possible to restrict our attention to bounded-size models, then SF is as expressive as full (relational) first-order logic. This alone is not a very interesting result, as already the *existential fragment* of relational first-order logic, i.e. the class of relational \exists^* prefix sentences, possesses this property (universal quantification can be replaced by finite conjunctions). What makes the case of SF special is that the incurred blowup in formula length is not linear in the bound but may be significantly lower.

Due to space limitations, we can only state the result here. The interested reader will find the full details in [69], Section 3.3.3. Abstractly speaking, when restricted to models of the size $2^{\uparrow n}(m)$, any first-order sentence can be translated into an equisatisfiable SF sentence whose length is polynomial in n , m , and the length of the original sentence.

Proposition 3.10. Let $m \geq 1$ and $n \geq 2$ be two integers. There exists an efficiently computable SF sentence $\chi_{m,n}$ and an effective translation $T_{m,n}$ mapping any relational first-order sentence φ to some SF sentence $T_{m,n}(\varphi)$ that satisfies the following properties.

- (a) $\mathcal{A} \models \chi_{m,n}$ entails $|\mathcal{A}| \leq 2^{\uparrow n}(m)$.
- (b) $\chi_{m,n} \wedge \varphi$ is equivalent to $\chi_{m,n} \wedge T_{m,n}(\varphi)$.
- (c) The formula length of $T_{m,n}(\varphi)$ is at most $p(m, n) \cdot \text{len}(\varphi)$ for some polynomial $p(m, n)$.
- (d) $T_{m,n}(\varphi)$ is computable in time $q(m, n, \text{len}(\varphi))$ for some polynomial $q(m, n, k)$.

Proposition 3.10 entails that SF-Sat is computationally at least as hard as the satisfiability problem for any first-order fragment that enjoys a small model property with an elementary upper bound on the size of small models. For instance, the *Ackermann fragment* (cf. Section 5), the *Gödel–Kalmár–Schütte fragment* (cf. Section 6), the *guarded fragment* (cf. Section 7), the *guarded-negation fragment* (cf. Section 8), and the *two-variable fragment* (cf. Section 9) fall into this category. Even the satisfiability problem for first-order fragments enjoying a small model property with bounds $2^{\uparrow \lceil c \cdot \text{len}(\varphi) \rceil}(\lceil d \cdot \text{len}(\varphi) \rceil)$ for constants c, d , such as the *fluted fragment* (cf. Section 10), can be polynomially reduced to SF-Sat. Although this latter observation already yields a non-elementary lower bound regarding the computational complexity of SF-Sat, a more accurate lower bound was presented in [68] by encoding *bounded domino problems* (see also [69], Section 5.3).

Proposition 3.11. Let \mathcal{C} be any class of relational first-order sentences for which we know constants $c, d \geq 1$ such that every satisfiable φ in \mathcal{C} has a model whose domain size is bounded by $2^{\uparrow \lceil c \cdot \text{len}(\varphi) \rceil}(\lceil d \cdot \text{len}(\varphi) \rceil)$. Then, \mathcal{C} -Sat is polynomial-time reducible to SF-Sat.

The proof of this result employs the translation $T_{m,n}$ from Proposition 3.10 for $m := \lceil d \cdot \text{len}(\varphi) \rceil$ and $n := \lceil c \cdot \text{len}(\varphi) \rceil$ for any given \mathcal{C} -sentence φ .

Employing the ideas underlying Proposition 3.10, one can also derive other lower bounds, similar to Theorem 3.8, regarding the length of sentences that are equivalent to SF sentences but adhere to certain syntactic restrictions. A classical result by Gaifman [22] states that every first-order formula is equivalent to some formula that is *local* in a certain sense (see, e.g. [43]). It has been shown [16] that there is a non-elementary gap between the length of first-order sentences and their shortest equivalents in *Gaifman normal form*. Using a proof approach inspired by the one in [16], one can show that this gap also applies to SF.

Proposition 3.12. There is some vocabulary Σ and some polynomial $p(h)$ such that for every $h \geq 0$ there is an SF Σ -sentence $\varphi_{\text{SF},h}$ of length $p(h)$ satisfying the following property. Every first-order Σ -sentence ψ in Gaifman normal form that is equivalent to $\varphi_{\text{SF},h}$ has length at least $2^{\uparrow h}(1)$.

Again, the interested reader will find the full details in [69], Section 3.3.3.

4 The Separated Bernays–Schönfinkel–Ramsey Fragment

In this section we introduce the separated extension of the Bernays–Schönfinkel–Ramsey fragment (BSR). The separated fragment can be conceived as an intermediate step, as it lies properly between BSR and the *separated Bernays–Schönfinkel–Ramsey fragment* (SBSR).

Recall that SF contains relational sentences $\exists z \forall \bar{x}_1 \exists \bar{y}_1 \dots \forall \bar{x}_n \exists \bar{y}_n. \psi$ in which the sets $\bar{x}_1 \cup \dots \cup \bar{x}_n$ and $\bar{y}_1 \cup \dots \cup \bar{y}_n$ are separated. The exemption of the leading existential quantifier block from the separateness conditions may lead to *benign* co-occurrences of existentially and universally quantified variables in atoms, which do not pose an obstacle to deciding the satisfiability problem. Extrapolating this emerging pattern of benign co-occurrences leads to the definition of the *separated Bernays–Schönfinkel–Ramsey fragment*. Intuitively speaking, an SBSR sentence φ has the form $\forall \bar{x}_1 \exists \bar{y}_1 \dots \forall \bar{x}_n \exists \bar{y}_n. \psi$ with quantifier-free ψ that may contain equality and possesses the following properties. Each atom in φ only contains variables from an $\exists^* \forall^*$ subsequence of φ 's quantifier prefix. If two atoms share a universally quantified variable, the same quantifier subsequence is used for both atoms.

Definition 4.1 (Separated Bernays–Schönfinkel–Ramsey fragment (SBSR)). Let Y, X_1, X_2, X_3, \dots be pairwise-disjoint, countable sets of first-order variables. The *separated Bernays–Schönfinkel–Ramsey fragment* (SBSR) comprises all relational first-order sentences with equality having the shape $\varphi := \forall \bar{x}_1 \exists \bar{y}_1 \dots \forall \bar{x}_n \exists \bar{y}_n. \psi$ with quantifier-free ψ where (a) $\bar{x}_i \subseteq X_1 \cup \dots \cup X_i$ and $\bar{y}_i \subseteq Y$ for every i , and (b) for every atom A in φ there is some i , $0 \leq i \leq n$, such that $\text{vars}(A) \subseteq \bar{y}_1 \cup \dots \cup \bar{y}_i \cup X_{i+1}$. (Notice that this entails separateness of X_1, \dots, X_n , and for every A there is some i with $\text{vars}(A) \subseteq \bar{y}_1 \cup \dots \cup \bar{y}_i \cup X_{i+1} \cup \dots \cup X_n$.)

Suppose we are given an arbitrary first-order sentence for which we intend to check membership in SBSR without knowing a-priori how the set of occurring variables is to be partitioned. Checking for the existence of such a suitable partition can be done deterministically in polynomial time. A corresponding procedure can be based on a graph algorithm (see [69], Theorem 3.4.3).

The main difference between SF and SBSR lies in the concession policy regarding benign co-occurrences of existential and universal variables. The following example gives a first impression of SBSR sentences and how they can be translated into BSR.

Example 4.2. Consider the first-order sentence $\varphi := \exists u \forall x \exists y \forall z. (P(u, z) \wedge Q(u, x)) \vee (P(y, z) \wedge Q(u, y))$. It belongs to SBSR, as witnessed by the following partition of its variables: $Y = \{u, y\}$, $X_1 = \{x\}$, $X_2 = \{z\}$, $X_3 = \emptyset$. Obviously, φ neither belongs to BSR nor to SF. As universal quantification does not distribute over disjunction, the quantifier $\forall z$ cannot be shifted inwards with the standard quantifier shifting rules from Lemma 2.1 alone. However, it turns out that the transformation methods that we have first met in Section 1 also facilitate translations of SBSR sentences into BSR sentences. We will elaborate on this in Section 4.1. For φ we get the equivalent BSR sentence

$$\begin{aligned} \varphi' := \exists u y \forall x z v. & ((P(u, x) \vee P(y, x)) \wedge P(u, x) \wedge Q(u, x)) \\ & \vee ((P(u, z) \vee P(y, z)) \wedge Q(u, y) \wedge Q(u, z)) \\ & \vee ((P(u, v) \vee P(y, v)) \wedge Q(u, y) \wedge P(y, v)) . \end{aligned}$$

We have advertised SBSR as an extension of SF and, hence, also of BSR and MFO. Indeed, given an SF sentence $\chi := \exists \bar{v}_1 \forall \bar{u}_2 \exists \bar{v}_2 \dots \forall \bar{u}_n \exists \bar{v}_n. \chi'$, we can partition the set of χ 's variables into $Y := \bar{v}_1 \cup \dots \cup \bar{v}_n$ and $X_1 := \emptyset$, $X_2 := \bar{u}_2 \cup \dots \cup \bar{u}_n$, and $X_j := \emptyset$ for $j \geq 3$. We then observe for every atom A in χ that either $\text{vars}(A) \subseteq Y$ or $\text{vars}(A) \subseteq \bar{v}_1 \cup X_2$. This partition complies with Definition 4.1. On the other hand, the sentence φ from Example 4.2 belongs to SBSR but not to SF. Hence, SBSR is a proper extension of SF.

Proposition 4.3. SBSR properly contains SF and, hence, BSR and MFO.

By Proposition 3.3, SBSR in addition semantically subsumes MFO_{\approx} , like SF does.

4.1 Translation of SBSR into BSR

We have already advertised several times that there is an effective equivalence-preserving translation from SBSR into BSR. It is mainly based on the standard axioms of Boolean algebra and quantifier shifting. Roughly speaking, we iteratively (re-)transform a given SBSR sentence into particular syntactic shapes and apply quantifier shifting so that we eventually obtain a formula in which no existential quantifier occurs within the scope of any universal quantifier. We then shift all quantifiers outwards again — existential quantifiers first —, renaming bound variables as necessary. The final result is a BSR sentence. Since SBSR contains SF, Theorem 3.8 entails that there is no elementary upper bound on the blowup that we incur in any equivalence-preserving translation from SBSR into BSR. On the other hand, the blowup for SBSR-BSR translations is not significantly worse than in the case of SF-BSR translations. It seems that in this sense SBSR does not offer much more succinctness when describing logical properties than SF does.

The following lemma formulates the technical invariants for transforming SBSR sentences into BSR sentences via shifting inwards quantifier blocks one after the other.

Lemma 4.4. Let Y, X_1, X_2, \dots be pairwise-disjoint, countable sets of first-order variables. Fix a quantifier sequence $\forall \bar{x}_1 \exists \bar{y}_1 \dots \forall \bar{x}_n \exists \bar{y}_n$ (without double occurrences of variables) such that $\bar{x}_i \subseteq X_1 \cup \dots \cup X_i$ and $\bar{y}_i \subseteq Y$ for every i . Moreover, fix a set of atoms At over the variables in $\bar{x}_1, \bar{y}_1, \dots, \bar{x}_n, \bar{y}_n$ such that for every atom $A \in \text{At}$ there is some i , $0 \leq i \leq n$, with $\text{vars}(A) \subseteq \bar{y}_1 \cup \dots \cup \bar{y}_i \cup X_{i+1}$.

Let $\varphi := \mathcal{Q}v.\psi$ be any first-order formula over the atoms in At such that

- (i) φ is in negation normal form, no variable is bound by two distinct quantifier occurrences, no variable occurs freely and bound,
- (ii) any sequence of nested quantifiers $\mathcal{Q}_1 u_1 \dots \mathcal{Q}_k u_k$ occurring in φ (from left to right) is a subsequence of $\forall \bar{x}_1 \exists \bar{y}_1 \dots \forall \bar{x}_n \exists \bar{y}_n$,
- (iii) for every subformula $\forall x.\chi$ of ψ we have $\text{vars}(\chi) \subseteq \bar{y}_1 \cup \dots \cup \bar{y}_{j_*-1} \cup X_{j_*}$ for some j_* ,
- (iv) for every subformula $\exists y.\chi$ in ψ with $y \in \bar{y}_i$ we have that χ does not contain any occurrences of variables from $\bar{x}_1 \cup \dots \cup \bar{x}_i$, and
- (v) any sequence of nested quantifiers $\mathcal{Q}_1 u_1 \dots \mathcal{Q}_k u_k$ occurring in ψ (from left to right) is a subsequence of $\exists \bar{y}_1 \dots \exists \bar{y}_i \forall \bar{x}_{i+1} \dots \forall \bar{x}_n$ for some i .

Then, φ can be transformed into an equivalent formula φ' such that Conditions (i) and (ii) still apply to φ' and Conditions (iii) to (v) apply to φ' instead of ψ only.

Proof. First of all, we remove from φ (and any later formulas) any quantifiers that bind variables which do not occur in any atom. A *basic formula* is any atom and any subformula $(\mathcal{Q}u \dots)$ in ψ that does not lie within the scope of any quantifier in ψ .

Suppose $\mathcal{Q}v$ is a universal quantifier. Then, there are indices i_*, j_* such that $v \in \bar{x}_{i_*} \cap X_{j_*}$ and $j_* \leq i_*$. We transform ψ into an equivalent conjunction of disjunctions of negated or non-negated basic formulas. This is always possible. Due to our assumptions, the constituents of the k -th disjunction can be grouped into i_*+1 parts: $\chi_{k,1} \vee \dots \vee \chi_{k,i_*} \vee \eta_k$ with $\text{vars}(\chi_{k,\ell}) \subseteq \bar{y}_1 \cup \dots \cup \bar{y}_\ell \cup X_{\ell+1}$ for $0 \leq \ell \leq i_*-1$ and $\text{vars}(\eta_k) \subseteq \bar{y}_1 \cup \dots \cup \bar{y}_n \cup X_{i_*+1} \cup \dots \cup X_n$. Due to Conditions (ii) and (iv), η_k can be defined so that it contains all basic formulas from ψ that have the form $\exists y.\eta'$ for any $y \subseteq \bar{y}_{i_*} \cup \dots \cup \bar{y}_n$. Hence, φ is equivalent to some formula of the form

$$\forall v. \bigwedge_k \chi_{k,1}(X_1) \vee \chi_{k,2}(\bar{y}_1, X_2) \vee \dots \vee \chi_{k,j_*}(\bar{y}_1, \dots, \bar{y}_{j_*-1}, X_{j_*}) \vee \dots \\ \vee \chi_{k,i_*}(\bar{y}_1, \dots, \bar{y}_{i_*-1}, X_{i_*}) \vee \eta_k(\bar{y}_1, \dots, \bar{y}_{i_*-1}).$$

We shift the universal quantifier $\forall v$ inwards so that it only binds the (sub-)conjunctions $\chi_{k,j_*}(\bar{y}_1, \dots, \bar{y}_{j_*-1}, \bar{x}_{j_*})$. The resulting formula

$$\bigwedge_k \chi_{k,1}(X_1) \vee \chi_{k,2}(\bar{y}_1, X_2) \vee \dots \vee (\forall v. \chi_{k,j_*}(\bar{y}_1, \dots, \bar{y}_{j_*-1}, X_{j_*})) \vee \dots \vee \chi_{k,i_*}(\bar{y}_1, \dots, \bar{y}_{i_*-1}, X_{i_*}) \vee \eta_k(\bar{y}_1, \dots, \bar{y}_{i_*-1}).$$

is the sought φ' . It is easy to see that φ' (after renaming bound variables) still satisfies Conditions (i) and (ii). Moreover, we now have established Conditions (iii), (iv), and (v) for the whole formula φ' instead of only the subformula ψ .

Now suppose $\mathcal{Q}v$ is an existential quantifier. Then, there is an index i_* such that $v \in \bar{y}_{i_*}$. We transform ψ into an equivalent disjunction of conjunctions of negated or non-negated basic formulas. Due to our assumptions, the constituents of the k -th conjunction can be grouped into two parts χ_k and η_k such that the variables occurring freely in η_k are a subset of $\bar{y}_1 \cup \dots \cup \bar{y}_{i_*}$ and χ_k may contain free occurrences of variables from $\bar{x}_1, \dots, \bar{x}_{i_*-1}$ but none from \bar{y}_{i_*} . After shifting the quantifier $\exists v$ inwards, we obtain the following formula φ' that is equivalent to the original φ :

$$\bigvee_k \chi_k(\bar{x}_1, \dots, \bar{x}_{i_*-1}, \bar{y}_1, \dots, \bar{y}_{i_*-1}) \wedge (\exists v. \eta_k(\bar{y}_1, \dots, \bar{y}_{i_*-1}, \bar{y}_{i_*})).$$

It is easy to see that φ' (after renaming bound variables) still satisfies Conditions (i) and (ii). Moreover, we now have established the Conditions (iii), (iv), and (v) for the whole formula φ' instead of only the subformula ψ . \square

Lemma 4.4 provides a tool to shift single quantifiers in SBSR sentences inwards in a fashion that, after applying the lemma iteratively, yields a sentence in which no existential quantifier lies inside the scope of any universal quantifier. Afterwards quantifiers may be shifted outward again so that we in the end obtain a BSR sentence.

Theorem 4.5. Every SBSR sentence $\varphi := \forall \bar{x}_1 \exists \bar{y}_1 \dots \forall \bar{x}_n \exists \bar{y}_n. \psi$ is equivalent to some BSR sentence whose length is at most n -fold exponential in the length of φ .

Proof Sketch. Let $\mathcal{Q}_1 v_1 \mathcal{Q}_2 v_2$ be the two rightmost quantifiers in φ 's quantifier prefix. It is easy to check that $\mathcal{Q}_2 v_2. \psi$ satisfies the prerequisites of Lemma 4.4 and, hence, can be transformed into an equivalent formula ψ' in accordance with that same lemma. Then, it is again easy to verify that $\mathcal{Q}_1 v_1. \psi'$ satisfies the lemma again. Proceeding further, quantifier by quantifier, we in the end obtain a formula in which all sequences of nested quantifiers have the form $\exists^* \forall^*$. Notice that the nesting depth of distinct quantifier blocks in the final result is at most n . Since the set of occurring atoms does not change (modulo variable renaming), the final formula cannot contain more than n -fold exponentially many distinct, non-redundant subformulas $\mathcal{Q}\bar{v}. \chi$ that do not occur within the scope of another quantifier. Consequently, shifting outwards all quantifier blocks in an existential-quantifiers-first fashion yields a BSR sentence equivalent to φ whose length is at most n -fold exponential in the length of φ . \square

On page 7 we mentioned that an analysis of the complexity of the SF-BSR translation process can be based on a measure that is more fine-grained than the number of occurring quantifier alternations. There is a similar measure for the setting of SBSR, which yields an improved upper bound. Due to space limitations, we will not go any further into the details at this point. The interested reader will find more material in [69], Sections 3.2 and 3.5.

Theorem 4.5 also holds in the presence of constant symbols: every SBSR sentence φ with constant symbols is equivalent to some BSR sentence φ' with the same constant symbols. Due to Proposition 3.5, Theorem 4.5 entails the following small model property for SBSR.

Corollary 4.6. Every satisfiable SBSR sentence φ with $n \forall \exists$ quantifier alternations and with or without constant symbols has a model whose size is at most n -fold exponential in the length of φ . Hence, *SBSR-Sat* is decidable nondeterministically in n -fold exponential time, even if we allow constant symbols to occur.

Regarding lower bounds, the result formulated in Theorem 3.8 immediately entails that there are SBSR sentences that inevitably lead to a non-elementary blowup when translating them into equivalent BSR sentences. Moreover, Proposition 3.12 is also relevant for SBSR. It means that for every natural number k there are SBSR sentences whose shortest equivalent in Gaifman normal form is k -fold exponentially longer than the original.

4.2 Taking Boolean Structure into Account

It is possible to liberalize the definition of SBSR while retaining decidability of the satisfiability problem, if one takes Boolean structure into account. For instance, since every \wedge - \vee -combination of BSR sentences is equivalent to some BSR sentence, we also observe that every \wedge - \vee -combination of SBSR sentences is equivalent to some SBSR sentence and, ultimately, also to some BSR sentence. Beyond such trivial observations, one may use approximations of conjunctive and disjunctive normal forms in order to predict when quantifier shifting combined with CNF/DNF-like normal-form transformations of the respective quantifier scopes would result in \wedge - \vee -combinations of SBSR sentences.

Example 4.7. Consider the sentence

$$\varphi := \forall x_1 x_2 \exists y \forall z_1 z_2. ((P(x_1, z_1) \vee P(z_2, x_2)) \wedge P(y, z_1)) \vee (P(x_2, z_2) \wedge P(x_1, x_2)),$$

which does not satisfy the conditions of SBSR, as x_1, z_1 and z_1, y co-occur in $P(x_1, z_1)$, $P(y, z_1)$, respectively. However, the sentence φ can be transformed into the equivalent

$$\left(\forall x_1 x_2 z_1 z_2. (P(x_1, z_1) \vee P(z_2, x_2) \vee P(x_2, z_2)) \wedge (P(x_1, z_1) \vee P(z_2, x_2) \vee P(x_1, x_2)) \right) \\ \wedge \left(\forall x_1 x_2 \exists y \forall z_1 z_2. (P(y, z_1) \vee P(x_2, z_2)) \wedge (P(y, z_1) \vee P(x_1, x_2)) \right).$$

Evidently, each of the two constituents of the topmost conjunction is an SBSR sentence.

More details regarding the liberalization of SBSR based on a suitable analysis of Boolean structure can be found in [69], Section 3.6.

5 The Separated Ackermann Fragment (SAF)

The *Ackermann fragment (AF)* comprises all relational first-order sentences in prenex normal form with an $\exists^* \forall \exists^*$ quantifier prefix and without equality. Ackermann derived the finite model property for AF in [2]. Ackermann's decidability proof in [3] proceeds via a reduction to MFO-Sat without any reference to AF's finite model property. The satisfiability problem for AF is EXPTIME-complete (see [11], Section 6.3). In [18] the finite model property of AF with equality is derived.

Gurevich [33] and Maslov and Orevkov [48] studied AF sentences with arbitrary function symbols, which yields the *Gurevich–Maslov–Orevkov fragment*. While Gurevich proved the finite model property for this fragment, Orevkov and Maslov took a proof-theoretic route based on Maslov's *inverse method*. Another extension of AF is the *Shelah fragment*: $\exists^* \forall \exists^*$ -sentences with equality and a single unary function symbol [63]. This class contains *infinity axioms* and, hence, does not possess the finite model property. A more detailed version of Shelah's proof can be found in Section 7.3 in [11].

In the present section, we generalize AF to the *separated Ackermann fragment (SAF)*. Moreover, we devise an effective translation procedure from SAF sentences into equivalent AF sentences. It will turn out that this procedure is compatible with function symbols and equality. That is, our results will show that SAF with equality is equivalent to AF with equality, SAF with arbitrary function symbols is equivalent to the Gurevich–Maslov–Orevkov fragment, and SAF with equality and a single unary function symbol is equivalent to the Shelah fragment. Hence, all these extensions of SAF are decidable.

Intuitively speaking, an SAF sentence is of the form $\varphi := \forall \bar{x}_1 \exists \bar{y}_1 \bar{u}_1 \dots \forall \bar{x}_n \exists \bar{y}_n \bar{u}_n. \psi$ with quantifier-free ψ and it satisfies the following properties. Each atom in φ contains only variables

from some subsequence of φ 's quantifier prefix of the form $\exists^*\forall\exists^*$. If two atoms share a universally quantified variable or some variable from the trailing \exists^* -block of their respective quantifier subsequence, then they have the same $\exists^*\forall\exists^*$ -subsequence as source of all their variables. For the formal definition, we define the *index of a variable* $v \in \text{vars}(\varphi)$ to be $\text{idx}_\varphi(v) := k$ if and only if $v \in \bar{x}_k \cup \bar{y}_k \cup \bar{u}_k$. We set $\text{idx}_\varphi(v) := \infty$ for all variables not occurring in φ . For convenience, we drop the reference to φ , if φ is clear from the context.

Definition 5.1 (Separated Ackermann fragment (SAF)). Let $Y, X, U_1, U_2, U_3, \dots$ be pairwise-disjoint, countable sets of first-order variables. The *separated Ackermann fragment (SAF)* comprises all relational first-order sentences without equality having the shape $\varphi := \forall \bar{x}_1 \exists \bar{y}_1 \bar{u}_1 \dots \forall \bar{x}_n \exists \bar{y}_n \bar{u}_n. \psi$ with quantifier-free ψ where (a) $\bar{x}_i \subseteq X$, $\bar{y}_i \subseteq Y$, and $\bar{u}_i \subseteq U_1 \cup U_2 \cup U_3 \cup \dots$ for every i , (b) $X = \{x_1, x_2, \dots\}$, (c) for every j and every $u \in U_j$ we have $\text{idx}(u) \geq \text{idx}(x_j)$, and (d) for every atom A in φ either (d.1) $\text{vars}(A) \subseteq Y$ or (d.2) there exists some j such that $\text{vars}(A) \subseteq \bar{y}_1 \cup \dots \cup \bar{y}_{\text{idx}(x_j)-1} \cup \{x_j\} \cup U_j$.

(Notice that this entails separateness of the sets $\{x_1\} \cup U_1, \{x_2\} \cup U_2, \{x_3\} \cup U_3, \dots$ in φ and every atom A in φ either contains exclusively variables from Y or there is some j such that $\text{vars}(A) \subseteq \bar{y}_1 \cup \dots \cup \bar{y}_{\text{idx}(x_j)-1} \cup \{x_j\} \cup \bar{u}_{\text{idx}(x_j)} \cup \dots \cup \bar{u}_n$.)

The tuples \bar{x}_i and \bar{y}_i, \bar{u}_i in φ may be empty, i.e. φ 's quantifier prefix does not have to start with a universal quantifier and it does not have to end with an existential quantifier. Moreover, variable $u \in \bar{u}$ that occurs in φ is associated with exactly one *reference variable* $x \in \bar{x}$, determined by the set U_i to which u belongs. Intuitively speaking, using suitable equivalence-preserving transformations, any quantifier $\exists u$ with $u \in \bar{u}$ can be shifted out of the scope of any universal quantifier but the one binding u 's reference variable. This is the essence of the first step of the effective translation procedure from SAF into AF.

The following example gives a first impression of SAF sentences and how they can be translated into AF.

Example 5.2. Consider the first-order sentence

$$\begin{aligned} \varphi := & \exists y \forall x_1 \exists u_1 \forall x_2 \exists u_2 u_3. (\neg P(y, x_1) \vee (Q(x_1, u_1) \wedge R(y, x_2, u_2))) \\ & \wedge (P(y, x_1) \vee (\neg Q(x_1, u_1) \wedge \neg R(y, x_2, u_3))) . \end{aligned}$$

The partition of the variables in φ into $Y := \{y\}$, $X := \{x_1, x_2\}$, and $U_1 := \{u_1\}$, $U_2 := \{u_2, u_3\}$ is a witness for the belonging of φ to SAF. Due to the Boolean structure of φ , the quantifiers $\exists u_3$, $\exists u_2$, and $\forall x_2$ can be shifted inwards immediately but $\exists u_1$ cannot. This yields the equivalent sentence

$$\begin{aligned} \exists y \forall x_1 \exists u_1. (\neg P(y, x_1) \vee (Q(x_1, u_1) \wedge \forall x_2 \exists u_2. R(y, x_2, u_2))) \\ \wedge (P(y, x_1) \vee (\neg Q(x_1, u_1) \wedge \forall x_2 \exists u_3. \neg R(y, x_2, u_3))) . \end{aligned}$$

Because of the two universal quantifiers $\forall x_1$ and $\forall x_2$, which are even interspersed with an existential one, φ does not belong to AF. However, there exists an equivalent sentence φ' in which every atom lies within the scope of at most one universal quantifier:

$$\begin{aligned} \exists y. \left(\forall x_1. (\neg P(y, x_1) \vee \exists u_1. Q(x_1, u_1)) \right) \wedge \left((\forall x_1. \neg P(y, x_1)) \vee \forall x_2 \exists u_2. R(y, x_2, u_2) \right) \\ \wedge \left(\forall x_1. (\exists u_1. \neg Q(x_1, u_1)) \vee P(y, x_1) \right) \wedge \left((\forall x_1 \exists u_1. \neg Q(x_1, u_1)) \vee \forall x_2 \exists u_2. R(y, x_2, u_2) \right) \\ \wedge \left((\forall x_2 \exists u_3. \neg R(y, x_2, u_3)) \vee \forall x_1. P(y, x_1) \right) \wedge \left((\forall x_2 \exists u_3. \neg R(y, x_2, u_3)) \vee \forall x_1 \exists u_1. Q(x_1, u_1) \right) \\ \wedge \left((\forall x_2 \exists u_3. \neg R(y, x_2, u_3)) \vee \forall x_2 \exists u_2. R(y, x_2, u_2) \right) \end{aligned}$$

Transforming φ into φ' requires only basic logical laws and is very similar to approaches we have seen before: first, we shift the quantifiers $\exists u_3, \exists u_2, \forall x_2$ inwards as far as possible. Then, we construct a disjunction of conjunctions of certain subformulas using distributivity. This allows us

to shift the quantifier $\exists u_1$ inwards. Afterwards, we apply the laws of distributivity again to obtain a conjunction of disjunctions of certain subformulas. This step enables us to shift the universal quantifier $\forall x_1$ inwards. Although the resulting sentence is not in AF, it is reasonably close to AF. The only difference is that we get more than only one universally quantified variable in the sentence as a whole, but at most one per atom. We will show later (Lemma 5.6) that each such sentence is indeed equivalent to some AF sentence.

Another example of a simple SAF sentence is the sentence

$$\psi := \exists u \forall x \exists y \forall z. (P(u, z) \wedge Q(u, x)) \vee (P(y, z) \wedge Q(u, y)).$$

It belongs to SBSR and SAF at the same time, while it does not belong to AF, SF, BSR, or MFO. Hence, even the intersection of SBSR and SAF contains sentences which do not fall into the syntactic categories offered by the standard fragments.

Suppose we intend to check membership in SAF given an arbitrary first-order sentence without knowing a-priori how the set of occurring variables is to be partitioned. Like in the SBSR setting, there is a deterministic polynomial-time procedure for this task based on a graph algorithm (see [69], Theorem 3.7.3).

The next proposition confirms that SAF extends AF and MFO. Since the sentence φ from Example 5.2 belongs to SAF but lies in neither of the other two fragments, it is immediately clear that SAF constitutes a proper extension of both.

Proposition 5.3. SAF properly contains AF and MFO.

Proof. Let $\varphi := \exists \bar{z} \forall x \exists \bar{v}. \psi$ be an AF sentence with quantifier-free ψ . We simply set $Y := \bar{z}$, $X := \{x\}$, $U_1 = \emptyset$, $U_2 := \bar{v}$, and $U_i := \emptyset$ for every $i \geq 3$. Then, φ trivially satisfies the conditions for SAF.

Let $\varphi' := \forall \bar{x}_1 \exists \bar{y}_1 \dots \forall \bar{x}_n \exists \bar{y}_n. \psi'$ be an MFO sentence. We set $Y := \bar{y}_1 \cup \dots \cup \bar{y}_n$, $X := \bar{x}_1 \cup \dots \cup \bar{x}_n$, and $U_i := \emptyset$ for every $i \geq 1$. Obviously, this partition of φ' 's atoms meets all the conditions for SAF sentences. \square

5.1 Translation of SAF into the Ackermann Fragment

In this section we devise an effective equivalence-preserving translation from SAF into AF that proceeds in two stages. The first stage unfolds nestings of quantifiers that bind separated sets of variables. This results in a sentence in which every subformula lies within the scope of at most one universal quantifier. Such sentences can easily be converted into a special syntactic form, called *SAF special form*. Then, in the second stage of the translation process, a sentence in SAF special form is transformed into an equivalent AF sentence.

The next lemma provides the central tool (including technical invariants) for the first stage of the translation process.

Lemma 5.4. Let Y, X, U_1, U_2, \dots be pairwise-disjoint, countable sets of first-order variables. Fix a quantifier sequence $\sigma := \forall \bar{x}_1 \exists \bar{y}_1 \bar{u}_1 \dots \forall \bar{x}_n \exists \bar{y}_n \bar{u}_n$ (without double occurrences of variables) such that (a) $\bar{x}_i \subseteq X = \{x_1, x_2, \dots\}$, $\bar{y}_i \subseteq Y$, and $\bar{u}_i \subseteq U_1 \cup U_2 \cup \dots$ for every i and (b) every $u \in U_j$ that occurs in σ occurs somewhere right of x_j . Moreover, fix a set of atoms At over the variables in $\bar{x}_1, \bar{y}_1, \bar{u}_1, \dots, \bar{x}_n, \bar{y}_n, \bar{u}_n$ such that for every atom $A \in \text{At}$ either (a) $\text{vars}(A) \subseteq Y$ or (b) there exists some j such that $\text{vars}(A) \subseteq \bar{y}_1 \cup \dots \cup \bar{y}_{\text{id}_X(x_j)-1} \cup \{x_j\} \cup U_j$.

Let $\varphi := \mathcal{Q}v. \psi$ be any first-order formula over the atoms from At such that

- (i) φ is in negation normal form, no variable is bound by two distinct quantifier occurrences, no variable occurs freely and bound,
- (ii) any sequence of nested quantifiers $\mathcal{Q}_1 v_1 \dots \mathcal{Q}_k v_k$ occurring in φ (from left to right) is a subsequence of $\forall \bar{x}_1 \exists \bar{y}_1 \bar{u}_1 \dots \forall \bar{x}_n \exists \bar{y}_n \bar{u}_n$,
- (iii) for every subformula $\forall x_j. \chi$ of ψ we have $\text{vars}(\chi) \subseteq \bar{y}_1 \cup \dots \cup \bar{y}_{\text{id}_X(x_j)-1} \cup \{x_j\} \cup U_j$,
- (iv) for every subformula $\exists u. \chi$ of ψ with $u \in U_j$ we have $\text{vars}(\chi) \subseteq \bar{y}_1 \cup \dots \cup \bar{y}_{\text{id}_X(x_j)-1} \cup \{x_j\} \cup U_j$,

- (v) for every subformula $\exists y. \chi$ of ψ with $y \in \bar{y}_i$ we have that all variables occurring freely in χ stem from Y , and
- (vi) any sequence of nested quantifiers $\mathcal{Q}_1 v_1 \dots \mathcal{Q}_k v_k$ occurring in ψ (from left to right) is either a subsequence of $\exists \bar{y}_1 \dots \exists \bar{y}_n$ or a subsequence of $\exists \bar{y}_1 \dots \exists \bar{y}_{\text{idx}(x_j)-1} \forall x_j \exists (\bar{u}_{\text{idx}(x_j)} \cap U_j) \dots \exists (\bar{u}_n \cap U_j)$ for some j .

Then, φ can be transformed into an equivalent formula φ' such that Conditions (i) and (ii) still apply to φ' and Conditions (iii) to (vi) apply to φ' instead of ψ only.

Proof Sketch. The proof works in analogy to the proof of Lemma 4.4. At any step, we remove any quantifiers that bind variables which do not occur in their scope. A *basic formula* is, again, any atom and any subformula $(\mathcal{Q}v \dots)$ in ψ that does not lie within the scope of any quantifier in ψ .

Suppose $\mathcal{Q}v$ is a universal quantifier and, hence, $v = x_{j_*} \in X$ for some j_* . We transform ψ into an equivalent conjunction of disjunctions of negated or non-negated basic formulas. Due to Conditions (iii) and (v), the constituents of the k -th disjunction can be grouped into two parts: $\chi_k \vee \eta_k$ with $\text{vars}(\chi_k) \subseteq \bar{y}_1 \cup \dots \cup \bar{y}_{\text{idx}(x_{j_*})-1} \cup \{x_{j_*}\} \cup U_{j_*}$ and $\text{vars}(\eta_k) \cap (\{x_{j_*}\} \cup U_{j_*}) = \emptyset$. Hence, we may shift the universal quantifier $\forall v$ inwards so that it only binds the (sub-)conjunctions χ_k . The resulting formula

$$\bigwedge_k (\forall x_{j_*} \cdot \chi_k(\bar{y}_1, \dots, \bar{y}_{\text{idx}(x_{j_*})-1}, x_{j_*})) \\ \vee \eta_k(\bar{x}_1, \bar{y}_1, \bar{u}_1, \dots, \bar{x}_{\text{idx}(x_{j_*})-1}, \bar{y}_{\text{idx}(x_{j_*})-1}, \bar{u}_{\text{idx}(x_{j_*})-1}, (\bar{x}_{\text{idx}(x_{j_*})} \setminus \{x_{j_*}\}))$$

is the sought φ' . It is easy to see that φ' still satisfies Conditions (i) and (ii). Moreover, we now have established the Conditions (iii), (iv), and (v) for the whole formula φ' instead of only the subformula ψ .

The cases where $\mathcal{Q}v$ is an existential quantifier and where either $v \in Y$ or $v \in U_j$ for some j can be handled in a similar fashion. \square

Lemma 4.4 provides a tool to shift single quantifiers in SAF sentences inwards in a fashion that, when applying the lemma iteratively, yields a sentence in which no universal quantifier lies inside the scope of another universal quantifier. Any first-order sentence in this form can be further transformed into a particular shape to which we refer as *SAF special form*:

$$\exists \bar{z}. \bigwedge_i \left(\bigvee_j \forall x_{i,j} \exists \bar{y}_{i,j} \cdot \chi_{i,j}(\bar{z}, x_{i,j}, \bar{y}_{i,j}) \right) \vee \eta_i(\bar{z})$$

where the $\chi_{i,j}$ and the η_i are quantifier free.

Lemma 5.5 (SAF special form). If $\varphi := \forall \bar{x}_1 \exists \bar{y}_1 \bar{u}_1 \dots \forall \bar{x}_n \exists \bar{y}_n \bar{u}_n \cdot \psi$ belongs to SAF, then we can effectively construct an equivalent sentence of the form

$$\exists \bar{z}. \bigwedge_i \left(\bigvee_j \forall x_{i,j} \exists \bar{y}_{i,j} \cdot \chi_{i,j}(\bar{z}, x_{i,j}, \bar{y}_{i,j}) \right) \vee \eta_i(\bar{z}),$$

where the $\chi_{i,j}$ and the η_i are quantifier free.

Proof. Let $\mathcal{Q}_1 v_1 \mathcal{Q}_2 v_2$ be the two rightmost quantifiers in φ 's quantifier prefix. It is easy to check that $\mathcal{Q}_2 v_2 \cdot \psi$ satisfies the prerequisites of Lemma 5.4 and, hence, can be transformed into an equivalent formula ψ' in accordance with that same lemma. Then, it is again easy to verify that $\mathcal{Q}_1 v_1 \cdot \psi'$ satisfies the lemma again. Proceeding further, quantifier by quantifier, we in the end obtain a formula in which no subformula lies within the scope of two distinct universal quantifiers. We may restore the property that no two quantifiers in φ' bind the same variables by appropriately renaming bound variables in φ' . Hence, we can effectively construct a sentence φ' that is equivalent to φ and that does not contain any nested occurrences of universal quantifiers.

We now transform φ' into a sentence φ'' in SAF special form. First, we shift all existential quantifiers in φ' to the front that do not lie within the scope of any universal quantifier. In the resulting sentence $\exists \bar{z} \cdot \psi'$ every existential quantifier in ψ' lies within the scope of exactly one

universal quantifier. We treat every subformula of the form $\forall x. \chi$ in ψ' as indivisible unit while transforming ψ' into an equivalent conjunction of disjunctions of literals and such indivisible units. The resulting formula can be brought into the desired shape by shifting existential quantifiers that lie in the scope of a universal quantifier outwards until they form an existential quantifier block directly right of the corresponding universal quantifier. \square

It is interesting to note that a sentence in SAF special form is not merely a Boolean combination of AF sentences. The difference is that distinct subformulas $\forall x \exists \bar{y}. \chi$ and $\forall x' \exists \bar{y}'. \chi'$ may share existentially quantified variables. However, one can show that every such sentence is indeed equivalent to some AF sentence. Therefore, every SAF sentence is equivalent to an AF sentence. Before we make this claim precise (cf. Lemma 5.7), we develop an auxiliary result that we will reuse later.

Lemma 5.6. Let $\psi(\bar{z})$ be a first-order formula of the form $\psi(\bar{z}) := \bigvee_j \forall \bar{x} \exists \bar{y}. \chi_j(\bar{z}, \bar{x}, \bar{y})$ with quantifier-free $\chi_j(\bar{z}, \bar{x}, \bar{y})$. Then, $\psi(\bar{z})$ is equivalent to the formula $\psi'(\bar{z})$ of the form

$$\exists \bar{v}_1 \dots \bar{v}_q \exists \bar{y}_1 \dots \bar{y}_q. \left(\bigvee_j \bigwedge_{k=1}^q \chi_j(\bar{z}, \bar{v}_k, \bar{y}_k) \right) \wedge \forall \bar{x} \exists \bar{y}. \bigvee_{k=1}^q \bigwedge_{A \in \text{At}} (A(\bar{z}, \bar{x}, \bar{y}) \leftrightarrow A(\bar{z}, \bar{v}_k, \bar{y}_k)) ,$$

where At denotes the set of all atoms occurring in $\psi(\bar{z})$ and $q := 2^{|\text{At}|}$. In addition, we have $|\bar{v}_k| = |\bar{x}|$ for every k and $|\bar{y}_\ell| = |\bar{y}|$ for every ℓ .

Proof. We first prove $\psi(\bar{z}) \models \psi'(\bar{z})$. Let \mathcal{A} be any structure, \bar{e} be any tuple of elements from \mathbf{A} , and j be any index such that $\mathcal{A} \models \forall \bar{x} \exists \bar{y}. \chi_j(\bar{e}, \bar{x}, \bar{y})$. For every set $S \subseteq \text{At}$ we define

$$D_S := \{ \langle \bar{a}, \bar{c} \rangle \mid \text{for every } A(\bar{z}, \bar{x}, \bar{y}) \in \text{At} \text{ we have } \mathcal{A} \models A(\bar{e}, \bar{a}, \bar{c}) \text{ if and only if } A \in S \} .$$

We write $S \models \chi_j(\bar{e}, \bar{x}, \bar{y})$ if D_S is nonempty and if we have $\mathcal{A} \models \chi_j(\bar{e}, \bar{a}, \bar{c})$ for every tuple $\langle \bar{a}, \bar{c} \rangle$ in D_S . Let S_1, \dots, S_r be an enumeration of all the sets S_k with $S_k \models \chi_j(\bar{e}, \bar{x}, \bar{y})$. Notice that $1 \leq r \leq q$. Let $\langle \bar{b}_1, \bar{c}_1 \rangle, \dots, \langle \bar{b}_r, \bar{c}_r \rangle$ be some sequence with $\langle \bar{b}_k, \bar{c}_k \rangle \in D_{S_k}$ for every k . Then, for every k the assumption $S_k \models \chi_j(\bar{e}, \bar{x}, \bar{y})$ entails $\mathcal{A} \models \chi_j(\bar{e}, \bar{b}_k, \bar{c}_k)$. Hence,

$$\mathcal{A} \models \exists \bar{y}_1 \dots \exists \bar{y}_q. \bigwedge_{k=1}^r \chi_j(\bar{e}, \bar{b}_k, \bar{y}_k) \wedge \bigwedge_{k=r+1}^q \chi_j(\bar{e}, \bar{b}_1, \bar{y}_k) . \quad (1)$$

Let $\bar{a} \in \mathbf{A}^{|\bar{x}|}$ be any tuple of length $|\bar{x}|$. Because of $\mathcal{A} \models \forall \bar{x} \exists \bar{y}. \chi_j(\bar{e}, \bar{x}, \bar{y})$, there is some S_k , $1 \leq k \leq r$, and some \bar{c} such that $\langle \bar{a}, \bar{c} \rangle \in D_{S_k}$ and $S_k \models \chi_j(\bar{e}, \bar{x}, \bar{y})$. Therefore, we get the following for $\langle \bar{b}_k, \bar{c}_k \rangle$:

$$\mathcal{A} \models \exists \bar{y}. \bigwedge_{A \in \text{At}} (A(\bar{e}, \bar{a}, \bar{y}) \leftrightarrow A(\bar{e}, \bar{b}_k, \bar{c}_k)) . \quad (2)$$

Put together, (1) and (2) entail

$$\begin{aligned} \mathcal{A} \models \exists \bar{y}_1 \dots \exists \bar{y}_q. & \left(\bigvee_j \bigwedge_{k=1}^r \chi_j(\bar{e}, \bar{b}_k, \bar{y}_k) \wedge \bigwedge_{k=r+1}^q \chi_j(\bar{e}, \bar{b}_1, \bar{y}_k) \right) \\ & \wedge \forall \bar{x} \exists \bar{y}. \bigvee_{k=1}^r \bigwedge_{A \in \text{At}} (A(\bar{e}, \bar{x}, \bar{y}) \leftrightarrow A(\bar{e}, \bar{b}_k, \bar{y}_k)) . \end{aligned}$$

This proves $\mathcal{A} \models \psi'(\bar{e})$. Hence, we have shown that $\mathcal{A} \models \psi(\bar{e})$ implies $\mathcal{A} \models \psi'(\bar{e})$.

Next, we show $\psi(\bar{z})' \models \psi(\bar{z})$. Let \mathcal{A} be any structure and let $\bar{e}, \bar{b}_1, \dots, \bar{b}_q, \bar{c}_1, \dots, \bar{c}_q$ be tuples for which

$$\mathcal{A} \models \left(\bigvee_j \bigwedge_{k=1}^q \chi_j(\bar{e}, \bar{b}_k, \bar{c}_k) \right) \wedge \forall \bar{x} \exists \bar{y}. \bigvee_{k=1}^q \bigwedge_{A \in \text{At}} (A(\bar{e}, \bar{x}, \bar{y}) \leftrightarrow A(\bar{e}, \bar{b}_k, \bar{c}_k)) . \quad (3)$$

Then, there is some j such that $\mathcal{A} \models \bigwedge_{k=1}^q \chi_j(\bar{e}, \bar{b}_k, \bar{c}_k)$. Let D_1, \dots, D_q be sets defined by

$$D_k := \left\{ \bar{a} \mid \text{there is some } \bar{c} \text{ such that for every } A(\bar{z}, \bar{x}, \bar{y}) \in \text{At} \right. \\ \left. \text{we have } \mathcal{A} \models A(\bar{e}, \bar{a}, \bar{c}) \text{ if and only if } \mathcal{A} \models A(\bar{e}, \bar{b}_k, \bar{c}_k) \right\}.$$

Note that the sets D_k are all nonempty but not necessarily pairwise disjoint. Then, because of Assumption (3), for every $\bar{a} \in A^{|\bar{x}|}$ there is some k , $1 \leq k \leq q$, such that $\bar{a} \in D_k$. Because of $\mathcal{A} \models \chi_j(\bar{e}, \bar{b}_k, \bar{c}_k)$, we therefore have $\mathcal{A} \models \chi_j(\bar{e}, \bar{a}, \bar{c})$ for some \bar{c} . In other words, we have $\mathcal{A} \models \forall \bar{x} \exists \bar{y}. \chi_j(\bar{e}, \bar{x}, \bar{y})$ which entails $\mathcal{A} \models \psi(\bar{e})$. Hence, $\mathcal{A} \models \psi'(\bar{e})$ implies $\mathcal{A} \models \psi(\bar{e})$. \square

Lemma 5.6 is essential for the second stage in the transformation process between SAF and AF. With this tool at hand, the following lemma is easy to prove.

Lemma 5.7. For every SAF sentence φ we can effectively construct an equivalent sentence φ' over the same vocabulary that has the shape $\exists \bar{v} \forall x \exists \bar{w}. \psi$ with quantifier-free ψ .

Proof sketch. By virtue of Lemma 5.5, we can transform φ into an equivalent sentence φ'' in SAF special form, i.e. $\varphi'' = \exists \bar{z}. \bigwedge_i \left(\bigvee_j \forall x_{i,j} \exists \bar{y}_{i,j}. \chi_{i,j}(\bar{z}, x_{i,j}, \bar{y}_{i,j}) \right) \vee \eta_i(\bar{z})$, where the $\chi_{i,j}$ and the η_i are quantifier free. Consider any subformula of the form $\psi'(\bar{z}) := \bigvee_j \forall x \exists \bar{y}. \chi_j(\bar{z}, x, \bar{y})$. By Lemma 5.6, $\psi'(\bar{z})$ is equivalent to some formula $\exists \bar{v}' \bar{y}'. \chi'(\bar{z}, \bar{v}', \bar{y}') \wedge \forall x \exists \bar{y}. \chi''(\bar{z}, x, \bar{y}, \bar{v}', \bar{y}')$ with quantifier-free χ', χ'' . Hence, φ'' is equivalent to some sentence that, after shifting quantifiers outwards, is of the form $\exists \bar{z}. \bigwedge_i \left(\exists \bar{u}_i \forall x_i \exists \bar{w}_i. \psi''_i(\bar{z}, \bar{u}_i, x_i, \bar{w}_i) \right) \vee \eta_i(\bar{z})$, where the ψ''_i and the η_i are quantifier free. A prenex version of this sentence yields the sought φ' , since the universal quantifiers distribute over the topmost conjunction. \square

Notice that the proofs of Lemmas 5.4 to 5.7 still work in the presence of the equality predicate or function symbols. Therefore, we obtain the following result.

Theorem 5.8. Every SAF sentence φ is equivalent to some AF sentence ψ . Moreover, we get the following for relaxed restrictions on the syntax.

- (a) Every SAF sentence with equality is equivalent to some AF sentence with equality.
- (b) Every SAF sentence with arbitrary function symbols and without equality is equivalent to some Gurevich–Maslov–Orevkov sentence (cf. page 14).
- (c) Every SAF sentence with equality and with a single unary function symbol is equivalent to some Shelah sentence (cf. page 14).

In addition, constant symbols are admissible in all of the above cases.

Since AF possesses the finite model property, so does SAF, even in the first two syntactically extended cases mentioned in Theorem 5.8. On the other hand, it is known that the Shelah fragment contains infinity axioms. One example is the sentence $\forall x \exists y. f(f(y)) \approx f(x) \wedge f(y) \not\approx x$ ([11], proof of Proposition 6.5.5). Still, the satisfiability problem for the Shelah fragment is known to be decidable (cf. [11], Section 7.3). Therefore, we get the following positive results regarding the decidability of SAF-Sat.

Corollary 5.9. SAF-Sat is decidable, even in the syntactically more liberal cases given in Theorem 5.8. The syntactic extensions of SAF described in items (a) and (b) of Theorem 5.8 enjoy the finite model property.

Remark 1. The Löb–Gurevich fragment (cf. page 5) is subsumed by SAF when we in addition allow unary function symbols. By Lemma 5.7, every such sentence is equivalent to some $\exists^* \forall \exists^*$ -sentence over the same vocabulary. The latter kind of sentences constitutes a subclass of the Gurevich–Maslov–Orevkov fragment. Hence, Lemmas 5.4 to 5.7 also establish a translation from the Löb–Gurevich fragment to the Gurevich–Maslov–Orevkov fragment.

At this point we have settled the question concerning decidability of SAF-Sat, also under certain syntactic extensions. In fact, decidability of SAF-Sat without any syntactic extensions is already a corollary of the decidability of the satisfiability problem for Maslov’s fragment K [47, 39]. The reason is that the latter syntactically subsumes SAF.

Definition 5.10 (Maslov’s fragment K, adapted from [39]). Let φ be any relational sentence in negation normal form and let $\psi(u_1, \dots, u_m)$ be any subformula of φ in which u_1, \dots, u_m are exactly the freely occurring variables. The φ -*prefix* of ψ is the sequence $\mathcal{Q}_1 v_1 \dots \mathcal{Q}_m v_m$ of quantifiers in φ (read from left to right) that bind the free variables of ψ , in particular, we have $\{v_1, \dots, v_m\} = \{u_1, \dots, u_m\}$. The *terminal φ -prefix* of ψ is the longest contiguous suffix of $\mathcal{Q}_1 v_1 \dots \mathcal{Q}_m v_m$ starting with a universal quantifier. Put differently, if $\mathcal{Q}_1 v_1 \dots \mathcal{Q}_m v_m$ is of the form $\exists v_1 \dots v_k \forall v_{k+1} \mathcal{Q}_{k+2} v_{k+2} \dots \mathcal{Q}_m v_m$, then the terminal φ -prefix of ψ is $\forall v_{k+1} \mathcal{Q}_{k+2} v_{k+2} \dots \mathcal{Q}_m v_m$. Notice that the terminal prefix may be empty. The sentence φ belongs to *Maslov’s fragment K* if there are $k \geq 0$ universal quantifiers $\forall x_1, \dots, \forall x_k$ in φ that are not interspersed with existential quantifiers such that for every atom A in φ the terminal φ -prefix of A either (a) is at most of length one, or (b) ends with an existential quantifier, or (c) is of the form $\forall x_1 \dots \forall x_k$.

Proposition 5.11. SAF is contained in Maslov’s fragment K.

Proof. Let $\varphi := \forall \bar{x}_1 \exists \bar{y}_1 \bar{u}_1 \dots \forall \bar{x}_n \exists \bar{y}_n \bar{u}_n. \psi$ be any SAF sentence with quantifier-free ψ . Recall that φ is relational and does not contain equality. In the terminology of the definition of Maslov’s fragment K, the terminal φ -prefix of any atom A in φ is either empty or it is a subsequence of $\forall x_j \exists (\bar{u}_{\text{id}_x(x)} \cap U_j) \dots (\bar{u}_n \cap U_j)$ for some j . Therefore, the terminal φ -prefix of A either is either empty or it ends with an existential quantifier or it is of length one. \square

Of course, Proposition 5.11 fails for any extensions of SAF with either equality or non-constant function symbols. We will see in the next section, how SAF can be extended in such a way that we obtain a generalization of the Gödel–Kalmár–Schütte fragment (GKS). Although GKS is syntactically contained in Maslov’s fragment K as well, its extension SGKS is not contained (cf. Proposition 6.3).

We have not yet given any lower bounds on the blowup that we incur when translating SAF sentences into AF. However, the computational complexity of AF-Sat (in EXPTIME, cf. [11], Theorems 6.3.26 and 6.3.1) and MFO-Sat (NEXPTIME-complete, cf. [11], Theorem 6.2.13) provide some evidence that this blowup is at least exponential. Since MFO is a subfragment of SAF, this entails the following conditional lower bound.

Proposition 5.12. In the worst case, there is at least a super-polynomial blowup in formula length when translating SAF sentences into equivalent AF sentences in a uniform algorithmic way, unless EXPTIME = NEXPTIME

6 The Separated Gödel–Kalmár–Schütte Fragment (SGKS)

The *Gödel–Kalmár–Schütte fragment (GKS)* comprises all relational first-order sentences in prenex normal form with an $\exists^* \forall \exists^*$ quantifier prefix and without equality. Gödel [23, 24], Kalmár [40], and Schütte [58, 59] independently showed that the satisfiability problem for GKS is decidable. Gödel and Kalmár established the finite model property. A probabilistic proof was later given by Gurevich and Shelah [35], see also Section 6.2.3 in [11]. Although Gödel claimed that his proof methods could also be applied for GKS sentences with equality, Goldfarb refuted this claim [25]. However, decidable subclasses are known, e.g. the syntactic subfragments described in [26] and in [71], Section 12. A decidable subclass described in semantic terms is mentioned in Section 6.2.3 in [11] — AF with equality falls into this category, for instance. Satisfiability for GKS is NEXPTIME-complete ([11], Section 6.2).

It is only a tiny step from the Ackermann fragment to the Gödel–Kalmár–Schütte fragment: simply allow two consecutive universal quantifiers in the quantifier prefix instead of only one. If one views the definition of SAF from the right angle, it is a similarly small step to go from SAF to a

separated generalization of GKS, which we will call the *separated Gödel–Kalmár–Schütte fragment* (SGKS). Intuitively speaking, an SGKS sentence is of the form $\varphi := \forall \bar{x}_1 \exists \bar{y}_1 \bar{u}_1 \dots \forall \bar{x}_n \exists \bar{y}_n \bar{u}_n. \psi$ with quantifier-free ψ and satisfies the following properties. Each atom in φ contains only variables from some subsequence of φ 's quantifier prefix of the form $\exists^* \forall \forall \exists^*$. We allow only fixed pairs of universally quantified variables to co-occur in atoms. Any two atoms that are associated with the same pair have the same $\exists^* \forall \forall \exists^*$ -subsequence as source of all their variables. The same applies to any two atoms that share some variable from the trailing \exists^* -block of their respective quantifier subsequence. We use the same notion of *index of a variable* like in Definition 5.1 (cf. page 15).

Definition 6.1 (Separated Gödel–Kalmár–Schütte fragment (SGKS)). Let $Y, X, U_1, U_2, U_3, \dots$ be pairwise-disjoint, countable sets of first-order variables. The *separated Gödel–Kalmár–Schütte fragment* (SGKS) comprises all relational first-order sentences without equality having the shape $\varphi := \forall \bar{x}_1 \exists \bar{y}_1 \bar{u}_1 \dots \forall \bar{x}_n \exists \bar{y}_n \bar{u}_n. \psi$ with quantifier-free ψ where the following conditions are satisfied. (a) $\bar{x}_i \subseteq X$, $\bar{y}_i \subseteq Y$, and $\bar{u}_i \subseteq U_1 \cup U_2 \cup U_3 \cup \dots$ for every i , (b) $\{x_1, x'_1\}, \{x_2, x'_2\}, \dots$ is a partition of X into sets of one or two variables each ($x_j = x'_j$ is allowed) such that for every j we have $\text{idx}(x_j) \leq \text{idx}(x'_j)$, (c) for every j and every $u \in U_j$ we have $\text{idx}(u) \geq \text{idx}(x'_j)$, and (d) for every atom A in φ either (d.1) $\text{vars}(A) \subseteq Y$ or (d.2) there exists some j such that $\text{vars}(A) \subseteq \bar{y}_1 \cup \dots \cup \bar{y}_{\text{idx}(x_j)-1} \cup \{x_j, x'_j\} \cup U_j$.

(Notice that this entails separateness of the sets $\{x_1, x'_1\} \cup U_1, \{x_2, x'_2\} \cup U_2, \{x_3, x'_3\} \cup U_3, \dots$ in φ and every atom A in φ either contains exclusively variables from Y or there is some j such that $\text{vars}(A) \subseteq \bar{y}_1 \cup \dots \cup \bar{y}_{\text{idx}(x_j)-1} \cup \{x_j, x'_j\} \cup \bar{u}_{\text{idx}(x'_j)} \cup \dots \cup \bar{u}_n$.)

Like for the other fragments the quantifier prefix of SGKS sentences does not have to start with a universal quantifier and it does not have to end with an existential quantifier either. Another analogy to SAF sentences is that every variable $u \in \bar{u}_1 \cup \dots \cup \bar{u}_n$ that occurs in φ is associated with a set $\{x_i, x'_i\} \subseteq X$ containing at least one and at most two *reference variables* determined by the set U_i in which u occurs. Intuitively speaking, like in the case of SAF, any quantifier $\exists u$ with $u \in \bar{u}_1 \cup \dots \cup \bar{u}_n$ can be shifted out of the scope of any universal quantifier that does not bind one of u 's reference variables.

Like for SAF and SBSR, deciding membership in SGKS for a given first-order sentence can be done deterministically in polynomial time (see [69], Theorem 3.9.2 for details).

SGKS obviously contains sentences that SAF does not, e.g. $\forall x_1 x_2. P(x_1, x_2)$. It is also easy to see that SGKS is an extension of SAF: if we restrict the sets $\{x_i, x'_i\}$ to singleton sets, then we essentially obtain SAF. Hence, also AF and MFO are subsets of SGKS. Finally, consider any GKS sentence $\exists \bar{y} \forall x x' \exists \bar{u}. \psi$ with quantifier-free ψ . We define $Y := \bar{y}$, $X := \{x, x'\}$, $U_1 = \emptyset$, $U_2 := \bar{u}$, and $U_i := \emptyset$ for every $i \geq 3$. Then, the sets Y, X, U_1, U_2, \dots satisfy the conditions of Definition 6.1 and thus witness that φ belongs to SGKS.

Proposition 6.2. SGKS properly contains GKS, SAF, AF, and MFO.

In the previous section we have seen that SAF is contained in Maslov's fragment K. We will see now that we have left the realm of the latter class when going from SAF to SGKS.

Proposition 6.3. SGKS and Maslov's fragment K are syntactically incomparable.

Proof. The following sentence witnesses that SGKS is not contained in Maslov's fragment K: $\forall x_1 x'_1 x_2 x'_2. P(x_1, x'_1) \vee Q(x_2, x'_2)$ belongs to SGKS but not to Maslov's K. On the other hand, it is easy to find sentences that belong to K but not to SGKS, e.g. $\forall x_1 x_2 x_3 \exists y. P(x_1, x_2, y) \wedge Q(x_1, x_3, y) \wedge R(x_1, x_2, x_3)$. \square

Next, we sketch an equivalence-preserving translation from SGKS into GKS, similar to the one in Section 5.1. Again, we proceed in two stages, first transforming a given SGKS sentence into *SGKS special form* and, afterwards, into an equivalent GKS sentence.

Lemma 6.4. If φ belongs to SGKS, we can effectively construct an equivalent sentence φ' in which every subformula lies within the scope of at most two universal quantifiers, and the scope of every universal quantifier contains at most one more universal quantifier. Moreover, all literals in φ' occur in φ (modulo variable renaming).

The transformation mentioned in the lemma is based on a straightforward adaptation of Lemma 5.4 (the singly occurring variables x_j have to be replaced with $\{x_j, x'_j\}$ and the rest needs to be adapted accordingly). The sentence φ' can easily be further transformed into a particular shape to which we will refer as *SGKS special form*:

$$\exists \bar{z}. \bigwedge_i \left(\bigvee_j (\forall x_{i,j} x'_{i,j} \exists \bar{y}_{i,j} \cdot \chi_{i,j}(\bar{z}, x_{i,j}, x'_{i,j}, \bar{y}_{i,j})) \vee \eta_i(\bar{z}) \right)$$

where the $\chi_{i,j}$ and the η_i are quantifier free.

Lemma 6.5. Every SGKS sentence φ in SGKS special form can be effectively transformed into an equivalent sentence φ' that has the shape $\exists \bar{z} \forall x x' \exists \bar{y} \cdot \psi$ with quantifier-free ψ .

Proof. Since φ is in SGKS special form, it has the shape

$$\varphi'' := \exists \bar{z}. \bigwedge_i \left(\bigvee_j (\forall x_{i,j} x'_{i,j} \exists \bar{y}_{i,j} \cdot \chi_{i,j}(\bar{z}, x_{i,j}, x'_{i,j}, \bar{y}_{i,j})) \vee \eta_i(\bar{z}) \right),$$

where the $\chi_{i,j}$ and the η_i are quantifier free. Consider any subformula of the form $\psi' := \bigvee_j \forall x x' \exists \bar{y} \cdot \chi_j(\bar{z}, x, x', \bar{y})$, possibly containing free variables from \bar{z} . By virtue of Lemma 5.6, ψ' is equivalent to some formula of the form $\exists \bar{v}' \bar{y}' \cdot \chi'(\bar{z}, \bar{v}', \bar{y}') \wedge \forall x x' \exists \bar{y} \cdot \chi''(\bar{z}, x, x', \bar{y}, \bar{v}', \bar{y}')$ with quantifier-free χ', χ'' . Hence, φ'' is equivalent to some sentence that, after shifting some quantifiers outwards, is of the form $\exists \bar{z}. \bigwedge_i (\exists \bar{u}_i \forall x_i x'_i \exists \bar{w}_i \cdot \psi''_i(\bar{z}, \bar{u}_i, x_i, x'_i, \bar{w}_i)) \vee \eta_i(\bar{z})$, where the ψ''_i and the η_i are quantifier free. A prenex version yields the sought φ' . \square

Theorem 6.6. Every SGKS sentence is equivalent to some GKS sentence.

Since we know that GKS enjoys the finite model property, this result immediately entails decidability of the satisfiability problem for SGKS (*SGKS-Sat*).

Corollary 6.7. The satisfiability problem for SGKS is decidable, and SGKS enjoys the finite model property.

We finish the present section by proving that SGKS sentences can be substantially more succinct than equivalent GKS sentences. The following theorem formulates a lower bound regarding the incurred blowup that comes along with any equivalence-preserving translation from SGKS to GKS.

Theorem 6.8. There is a class of SGKS sentences and some positive integer n_0 such that for every integer $n \geq n_0$ the class contains a sentence φ with a length linear in n for which any equivalent GKS sentence has a length that is at least exponential in n .

Proof sketch. Let $n \geq 1$ be some positive integer. Consider the following first-order sentence in which the sets $\{x_1, x_2\}$ and $\{y_1, y_2\}$ are separated:

$$\varphi := \forall x_2 \exists y_2 \forall x_1 \exists y_1 \cdot \bigwedge_{i=1}^{8n} (P_i(x_1, x_2) \leftrightarrow Q_i(y_1, y_2)) .$$

In analogy to the proof of Theorem 3.8, we construct the following model \mathcal{A} using the sets $\mathcal{S}_1 := \{S \subseteq [8n] \mid |S| = 2n\}$ and $\mathcal{S}_2 := \{S \subseteq \mathcal{S}_1 \mid |S| = \frac{1}{2}|\mathcal{S}_1|\}$. We observe that

$$|\mathcal{S}_1| = \binom{8n}{2n} \geq \left(\frac{8n}{2n}\right)^{2n} = 2^{4n} \quad \text{and} \quad |\mathcal{S}_2| = \binom{|\mathcal{S}_1|}{|\mathcal{S}_1|/2} \geq \left(\frac{|\mathcal{S}_1|}{|\mathcal{S}_1|/2}\right)^{|\mathcal{S}_1|/2} \geq 2^{2^{4n-1}} .$$

Having the sets $\mathcal{S}_1, \mathcal{S}_2$, we now define the structure \mathcal{A} as follows:

$$\mathbf{A} := \{\mathbf{a}_S^{(1)}, \mathbf{b}_S^{(1)} \mid S \in \mathcal{S}_1\} \cup \{\mathbf{a}_S^{(2)}, \mathbf{b}_S^{(2)} \mid S \in \mathcal{S}_2\},$$

$$P_i^{\mathcal{A}} := \{\langle \mathbf{a}_{S_1}^{(1)}, \mathbf{a}_{S_2}^{(2)} \rangle \in \mathbf{A}^2 \mid i \in S_1 \in S_2\} \text{ for } i = 1, \dots, 8n, \text{ and}$$

$$Q_i^{\mathcal{A}} := \{\langle \mathbf{b}_{S_1}^{(1)}, \mathbf{b}_{S_2}^{(2)} \rangle \in \mathbf{A}^2 \mid i \in S_1 \in S_2\} \text{ for } i = 1, \dots, 8n.$$

Like in the proof of Theorem 3.8, it is easy to show that \mathcal{A} is a model of φ .

For every $S \in \mathcal{S}_1 \cup \mathcal{S}_2$ we define the structure \mathcal{A}_{-S} as the substructure of \mathcal{A} induced by the domain $\mathbf{A}_{-S} := \mathbf{A} \setminus \{\mathbf{b}_S^{(k)}\}$, where $k = 1$ if $S \in \mathcal{S}_1$ and $k = 2$ if $S \in \mathcal{S}_2$. Like in the proof of Theorem 3.8 (Claim II), we can prove the following claim.

Claim I: For every $S \in \mathcal{S}_1 \cup \mathcal{S}_2$ the substructure \mathcal{A}_{-S} of \mathcal{A} does not satisfy φ . \diamond

Let $\varphi_* := \exists \bar{z} \forall x_1 x_2 \exists \bar{y}. \psi_*$ with quantifier-free ψ_* be a shortest GKS sentence equivalent to φ . Suppose that the length of φ_* is less than 2^n . Let $\psi' := \bigvee_{i \in I} \chi_i(\bar{z}, x_1, x_2, \bar{y})$ be a shortest disjunction of conjunctions χ_i of literals that is equivalent to ψ_* . We observe that the index set I contains fewer than 2^{2^n} indices and that every conjunction χ_i contains fewer than 2^n literals, for otherwise we could find a shorter formula with the desired properties.

Let $\bar{\mathbf{d}}$ be some tuple for which we have

$$\mathcal{A} \models \forall x_1 x_2 \exists \bar{y}. \bigvee_{i \in I} \chi_i(\bar{\mathbf{d}}, x_1, x_2, \bar{y}).$$

Let $\mathbf{D} := \{\mathbf{b}_S^{(2)} \mid S \in \mathcal{S}_2 \text{ and } \mathbf{b}_S^{(2)} \notin \bar{\mathbf{d}}\}$. Because of $|\bar{\mathbf{d}}| \leq |\bar{z}| \leq \text{len}(\varphi_*) \leq 2^n$ and $|\mathcal{S}_2| \geq 2^{2^{4n-1}} \geq 2^{2^{3n}}$, we have $|\mathbf{D}| \geq 2^{2^{2n}}$ for sufficiently large n . By Claim I, we observe

$$\mathcal{A}_{-S} \not\models \forall x_1 x_2 \exists \bar{y}. \bigvee_{i \in I} \chi_i(\bar{\mathbf{d}}, x_1, x_2, \bar{y})$$

for every S with $\mathbf{b}_S^{(2)} \in \mathbf{D}$. Hence, for every $\mathbf{b}_S^{(2)} \in \mathbf{D}$ there is some pair $\mathbf{c}_1, \mathbf{c}_2 \in \mathbf{A} \setminus \{\mathbf{b}_S^{(2)}\}$, some tuple $\bar{\mathbf{b}}$ containing $\mathbf{b}_S^{(2)}$ and some index $i_S \in I$ such that $\mathcal{A} \models \chi_{i_S}(\bar{\mathbf{d}}, \mathbf{c}_1, \mathbf{c}_2, \bar{\mathbf{b}})$ and $\mathcal{A}_{-S} \not\models \exists \bar{y}. \chi_{i_S}(\bar{\mathbf{d}}, \mathbf{c}_1, \mathbf{c}_2, \bar{y})$. Because of $|I| < 2^{2^n}$ and $|\mathbf{D}| \geq 2^{2^{2n}}$, there must be some index i_* that appears in the role of i_S for at least

$$\frac{|\mathbf{D}|}{|I|} \geq \frac{2^{2^{2n}}}{2^{2^n}} = 2^{2^{2n} - 2^n} \geq 2^{(2^n)^2 / 2^n} = 2^{2^n}$$

distinct $\mathbf{b}_S^{(2)} \in \mathbf{D}$, if n is sufficiently large. Let $\mathbf{D}_* \subseteq \mathbf{D}$ be the set that comprises exactly those elements. In other words, we have $|\mathbf{D}_*| \geq 2^{2^n}$ and for every $\mathbf{b}_S^{(2)} \in \mathbf{D}_*$ there is some pair $\mathbf{c}_1, \mathbf{c}_2$ and some tuple $\bar{\mathbf{b}}$ containing $\mathbf{b}_S^{(2)}$ such that

$$\mathcal{A} \models \chi_{i_*}(\bar{\mathbf{d}}, \mathbf{c}_1, \mathbf{c}_2, \bar{\mathbf{b}}) \quad \text{and} \quad \mathcal{A}_{-S} \not\models \exists \bar{y}. \chi_{i_*}(\bar{\mathbf{d}}, \mathbf{c}_1, \mathbf{c}_2, \bar{y}). \quad (4)$$

Consider some $\mathbf{b}_S^{(2)} \in \mathbf{D}_*$ with $S \in \mathcal{S}_2$ and fix it. The only atoms in χ_{i_*} that could possibly contribute to the effect described in (4) for $\mathbf{b}_S^{(2)}$ have the form $Q_j(z, y')$, $Q_j(y, y')$, $Q_j(x_1, y')$, or $Q_j(x_2, y')$ for $z \in \bar{z}$, $y, y' \in \bar{y}$ and, moreover, the variables z, y, x_1, x_2 need to be assigned values $\mathbf{b}_T^{(1)}$ with $T \in \mathcal{S}_1$. Let \mathcal{S}'_1 be the set collecting all the T from \mathcal{S}_1 that are assigned to such variables occurring in atoms of the mentioned kind. As χ_{i_*} contains at most 2^n such variables, $|\mathcal{S}'_1| \leq 2^n$. Recall that S contains $\frac{1}{2}|\mathcal{S}_1| \geq 2^{4n-1}$ sets of indices. By construction of \mathcal{S}_2 , there must be some $S' \in \mathcal{S}_2$ such that for every $T \in \mathcal{S}'_1$ we have $T \in S'$ if and only if $T \in S$. Let $\bar{\mathbf{b}}'$ be the tuple that results from $\bar{\mathbf{b}}$ by replacing every occurrence of $\mathbf{b}_S^{(2)}$ in the tuple by $\mathbf{b}_{S'}^{(2)}$. Then, we get $\mathcal{A}_{-S} \models \chi_{i_*}(\bar{\mathbf{d}}, \mathbf{c}_1, \mathbf{c}_2, \bar{\mathbf{b}}')$, which contradicts (4). Consequently, the length of the sentence φ_* cannot be less than 2^n . \square

7 Separateness and Guarded Quantification

The *guarded fragment (GF)* comprises all relational first-order sentences with equality that satisfy the following properties. An *atomic guard* $\gamma(\bar{u}, \bar{v})$ is an atom A such that all $u \in \bar{u} \cup \bar{v}$ occur in A . We define the *guarded fragment (GF)* inductively: (i) every relational atom is a GF formula

(equality is allowed); (ii) every Boolean combination of GF formulas is a GF formula; (iii) for all tuples \bar{u}, \bar{v} , every atomic guard $\gamma(\bar{u}, \bar{v})$, and every GF formula $\psi(\bar{u}, \bar{v})$ the following formulas belong to GF: $\forall \bar{u}. (\gamma(\bar{u}, \bar{v}) \rightarrow \psi(\bar{u}, \bar{v}))$, abbreviated by $(\forall \bar{u}. \gamma(\bar{u}, \bar{v}))\psi(\bar{u}, \bar{v})$, and $\exists \bar{u}. (\gamma(\bar{u}, \bar{v}) \wedge \psi(\bar{u}, \bar{v}))$, abbreviated by $(\exists \bar{u}. \gamma(\bar{u}, \bar{v}))\psi(\bar{u}, \bar{v})$. Notice that we assume in any GF formula $(\mathcal{Q}\bar{u}. \gamma(\bar{u}, \bar{v}))\psi(\bar{u}, \bar{v})$ that all variables that occur freely in ψ also occur in γ .

The guarded fragment was introduced by Andr eka, N emeti, and van Benthem [4] as one characterization of the fragment of first-order logic in which propositional modal logic can be embedded via the so-called *standard translation* (cf. Section 2.4 in [10]). Van Benthem [67] also proposed a more liberal form of guards, *loose guards*. A *loose guard* $\gamma(\bar{u}, \bar{v})$ is a nonempty conjunction of atoms $\gamma(\bar{u}, \bar{v}) := A_1(\bar{u}, \bar{v}) \wedge \dots \wedge A_k(\bar{u}, \bar{v})$ such that all u, v with $u \in \bar{u}$ and $v \in \bar{u} \cup \bar{v}$ co-occur in at least one A_j . The *loosely guarded fragment (LGF)* is then defined by liberalizing (iii) such that loose guards are used instead of atomic guards. In particular, we assume in any LGF formula $(\mathcal{Q}\bar{u}. \gamma(\bar{u}, \bar{v}))\psi(\bar{u}, \bar{v})$ that (a) all variables that occur freely in ψ also occur in γ and (b) every variable that is bound by $\mathcal{Q}\bar{u}$ co-occurs with every free variable from ψ in some atom in γ . We will occasionally use sloppy language and speak of LGF *formulas* when we mean loosely guarded formulas that are not necessarily closed. Formally, LGF exclusively contains sentences. The same applies to GF *formulas*.

Gr adel [28] derived the tree-like model property for GF and LGF and the finite model property for GF. Moreover, the computational complexity of the associated satisfiability problems is pinpointed in the same article: both are complete for deterministic doubly exponential time. More variants of guards and guarded quantification have been proposed, which lead to the definition of the *clique-guarded fragment* [27] and the *packed guarded fragment* [46], for instance. Hodkinson [38] showed that also the loosely guarded fragment, the clique-guarded fragment, and the packed guarded fragment enjoy the finite model property.

At first glance it seems that guarded quantification and separateness of quantified variables are two opposite properties. In particular, any guard $\gamma(\bar{u}, \bar{v})$ in a formula $(\forall \bar{u}. \gamma(\bar{u}, \bar{v}))\varphi(\bar{u}, \bar{v})$ has to ensure that every $u \in \bar{u}$ co-occurs with each $v \in \bar{v}$ in at least one atom in γ . This destroys any separateness of variables from \bar{u} and \bar{v} which might be separated in φ . However, it turns out that guardedness and separateness can indeed be combined in a beneficial way.

Definition 7.1 (Separated loosely guarded fragment (SLGF)). Two tuples \bar{x}, \bar{y} are *guard-separated* in a formula ψ with guarded quantification if \bar{x} and \bar{y} are separated in ψ and, in addition, for every guard γ in ψ either $\text{vars}(\gamma) \cap \bar{x} = \emptyset$ or $\text{vars}(\gamma) \cap \bar{y} = \emptyset$.

We define the set of *SLGF formulas* inductively as follows.

- (i) Every relational atom is an SLGF formula, equality is admitted.
- (ii) Every Boolean combination of SLGF formulas is an SLGF formula.
- (iii) Let $\bar{u}, \bar{v}, \bar{z}$ be tuples of variables and let $\gamma(\bar{u}, \bar{v})$ be any loose guard. The following are SLGF formulas: $\forall \bar{u}. (\gamma(\bar{u}, \bar{v}) \rightarrow \psi(\bar{u}, \bar{v}, \bar{z}))$ and $\exists \bar{u}. (\gamma(\bar{u}, \bar{v}) \wedge \psi(\bar{u}, \bar{v}, \bar{z}))$, where the sets \bar{u} and \bar{z} are guard-separated in ψ .

The *separated loosely guarded fragment (SLGF)* is the class of all SLGF sentences. When we restrict guards to be atomic, we obtain the *separated guarded fragment (SGF)*.

Remark 2. Notice that every formula of the form $\forall \bar{u}\bar{x}. \gamma(\bar{u}, \bar{v}) \wedge \delta(\bar{x}, \bar{y}) \rightarrow \psi(\bar{u}, \bar{v}, \bar{x}, \bar{y}, \bar{z})$ where \bar{u}, \bar{x} , and \bar{z} are pairwise distinct and guard-separated in ψ is equivalent to the SLGF formula $(\forall \bar{u}. \gamma(\bar{u}, \bar{v}))((\forall \bar{x}. \delta(\bar{x}, \bar{y}))\psi(\bar{u}, \bar{v}, \bar{x}, \bar{y}, \bar{z}))$. A dual observation holds true for existential quantification. This means that, under certain restrictions, we may mix variables that are subject to distinct guards in a single quantifier block. One could incorporate this idea into the definition of SLGF and, hence, obtain a syntactically slightly extended version. However, for the sake of simplicity, we adhere to the simpler definition given above.

It is easy to see that SLGF is indeed a proper syntactic extension of LGF, and that the same applies to SGF and GF. A simple sentence witnessing the strictness of these containment relations

is the sentence $(\forall x. x \approx x)(\exists y. y \approx y)(P(y) \leftrightarrow \neg P(x))$. It belongs neither to GF nor to LGF, but to both SGF and SLGF. Moreover, the sentence is a witness of the following observation: Every MFO sentence can be easily turned into an equivalent SGF sentence with a length linear in the original. We just need to add trivial equations $v \approx v$ as guards to subformulas $\mathcal{Q}v. \chi$. The result of this transformation lies in the intersection of SGF and MFO_{\approx} .

Proposition 7.2. SGF properly contains GF and SLGF properly contains SGF, LGF, and GF. Moreover, every MFO sentence can be turned into an equivalent SGF sentence with a length linear in the original.

For MFO_{\approx} sentences the matter seems to be more complicated. The sentence $\forall xy. x \approx y$, for instance, is not an SLGF sentence and cannot be directly transformed into an equivalent SLGF sentence in the described way.

Like for all the other novel first-order fragments we have defined, there exists an effective translation procedure that transforms SLGF sentences into equivalent LGF sentences.

Lemma 7.3. Every SLGF formula is equivalent to some LGF formula.

Proof. We prove an auxiliary result from which the lemma follows: Consider any SLGF formula $\varphi := (\mathcal{Q}\bar{u}. \gamma(\bar{u}, \bar{v}))\psi(\bar{u}, \bar{v}, \bar{z})$ where ψ is an LGF formula, $\bar{u}, \bar{v}, \bar{z}$ are pairwise disjoint, and φ 's free variables are exactly the ones in \bar{v}, \bar{z} . Then, φ is equivalent to some LGF formula $\varphi'(\bar{v}, \bar{z})$. Moreover, any two guard-separated sets of variables in φ are also guard-separated in φ' .

Suppose \mathcal{Q} is a universal quantifier (the case for existential quantification is dual). Recall that, by definition of SLGF, the tuples \bar{u} and \bar{z} need to be guard-separated in ψ . Since ψ is an LGF formula and since we assume that no variable occurs freely and bound in φ at the same time, we know that in every subformula $\chi := (\mathcal{Q}\bar{x}. \delta(\bar{x}, \bar{y}))\eta(\bar{x}, \bar{y})$ of ψ we either have $\text{vars}(\chi) \cap \bar{u} = \emptyset$ or $\text{vars}(\chi) \cap \bar{z} = \emptyset$ (or both). Moreover, since φ is an SLGF formula, we have $\text{vars}(A) \cap \bar{u} = \emptyset$ or $\text{vars}(A) \cap \bar{z} = \emptyset$ for every atom A in ψ . Hence, φ is equivalent to some formula of the form $\varphi'' := \forall \bar{u}. \gamma(\bar{u}, \bar{v}) \rightarrow \bigwedge_i (\chi_i(\bar{u}, \bar{v}) \vee \eta_i(\bar{v}, \bar{z}))$, where the χ_i and η_i are disjunctions of literals or LGF formulas of the form $(\mathcal{Q}\bar{x}. \delta(\bar{x}, \bar{y}))\eta(\bar{x}, \bar{y})$. Applying distributivity and shifting the quantifier $\forall \bar{u}$ in φ'' , it is easy to show equivalence to $\varphi' := \bigwedge_i ((\forall \bar{u}. \gamma(\bar{u}, \bar{v}) \rightarrow \chi_i(\bar{u}, \bar{v})) \vee \eta_i(\bar{v}, \bar{z}))$. This is the sought LGF formula. \square

Notice that the proof works irrespectively of the structure of guards. Hence, we also observe that every SGF formula is equivalent to some GF formula.

In connection with the fact that GF and LGF possess the finite model property [28, 38], the obvious consequence of Lemma 7.3 is that the satisfiability problems associated with SGF and SLGF (*SGF-Sat* and *SLGF-Sat*) are decidable.

Theorem 7.4. Both SGF and SLGF possess the finite model property. Moreover, SGF-Sat and SLGF-Sat are decidable.

We conclude this section with an investigation of the succinctness gap between SLGF and LGF. The following theorem entails that there is no elementary upper bound on the length of the LGF sentences that result from any equivalence-preserving transformation of SLGF sentences (and even SGF sentences) into LGF.

Theorem 7.5. There is a class of SGF sentences such that for every integer $n \geq 3$ the class contains a sentence φ with n $\forall\exists$ alternations and with a length polynomial in n for which any equivalent LGF sentence has at least $(n - 1)$ -fold exponential length in n .

Proof sketch. Let $n \geq 3$. Consider the following SGF sentence in which the sets $\{x_1, \dots, x_n\}$ and $\{y_1, \dots, y_n\}$ are separated:

$$\varphi := (\forall x_n. R_n(x_n))(\exists y_n. T_n(y_n)) \dots$$

$$(\forall x_1. R_1(x_1, \dots, x_n))(\exists y_1. T_1(y_1, \dots, y_n)) \cdot \bigwedge_{i=1}^{4n} (P_i(x_1, \dots, x_n) \leftrightarrow Q_i(y_1, \dots, y_n)).$$

In order to construct a particular model of φ , we reuse the sets $\mathcal{S}_1, \dots, \mathcal{S}_n$ from the proof of Theorem 3.8 and define the structure \mathcal{A} as follows:

$$\begin{aligned} \mathbf{A} &:= \{ \langle \mathbf{a}_S^{(k)}, \mathbf{b}_S^{(k)} \rangle \mid 1 \leq k \leq n \text{ and } S \in \mathcal{S}_k \}, \\ P_i^{\mathcal{A}} &:= \{ \langle \mathbf{a}_{S_1}^{(1)}, \dots, \mathbf{a}_{S_n}^{(n)} \rangle \in \mathbf{A}^n \mid i \in S_1 \in S_2 \in \dots \in S_n \} \text{ for } i = 1, \dots, 4n, \\ Q_i^{\mathcal{A}} &:= \{ \langle \mathbf{b}_{S_1}^{(1)}, \dots, \mathbf{b}_{S_n}^{(n)} \rangle \in \mathbf{A}^n \mid i \in S_1 \in S_2 \in \dots \in S_n \} \text{ for } i = 1, \dots, 4n, \\ R_j^{\mathcal{A}} &:= \{ \langle \mathbf{a}_{S_j}^{(j)}, \dots, \mathbf{a}_{S_n}^{(n)} \rangle \in \mathbf{A}^n \mid S_j \in \dots \in S_n \} \text{ for } j = 1, \dots, n, \text{ and} \\ T_j^{\mathcal{A}} &:= \{ \langle \mathbf{b}_{S_j}^{(j)}, \dots, \mathbf{b}_{S_n}^{(n)} \rangle \in \mathbf{A}^n \mid S_j \in \dots \in S_n \} \text{ for } j = 1, \dots, n. \end{aligned}$$

Once again, it is easy to show that \mathcal{A} is a model of φ .

In analogy to the proof of Theorem 3.8, one may use a game-based argument to show the following claim. For every $S \in \mathcal{S}_k$, $1 \leq k \leq n$, we define the structure \mathcal{A}_{-S} as the substructure of \mathcal{A} induced by the domain $\mathbf{A}_{-S} := \mathbf{A} \setminus \{ \mathbf{b}_S^{(k)} \}$.

Claim I: For every $S \in \mathcal{S}_k$, $1 \leq k \leq n$, the substructure \mathcal{A}_{-S} of \mathcal{A} does not satisfy φ . \diamond

A detailed proof can be found in [69], Section 3.10 (Theorem 3.10.8).

We have already analyzed the size of the sets \mathcal{S}_k in the proof of Theorem 3.8. Due to the observed lower bounds, we know that \mathbf{A} contains at least $\sum_{k=1}^n 2^{\uparrow k}(n)$ elements $\mathbf{b}_S^{(k)}$.

Let φ_{LGF} be a shortest LGF sentence that is semantically equivalent to φ . Next, we argue that $\text{len}(\varphi_{\text{LGF}})$ is at least $(n-1)$ -fold exponential in n . We start by introducing some additional notation. We divide the domain \mathbf{A} into two disjoint parts $\mathbf{A}_a := \{ \mathbf{a}_S^{(k)} \mid 1 \leq k \leq n \text{ and } S \in \mathcal{S}_k \}$ and $\mathbf{A}_b := \{ \mathbf{b}_S^{(k)} \mid 1 \leq k \leq n \text{ and } S \in \mathcal{S}_k \}$. Moreover, we subdivide \mathbf{A}_b into parts $\mathbf{A}_{b,k} := \{ \mathbf{b}_S^{(k)} \mid S \in \mathcal{S}_k \}$ with $1 \leq k \leq n$. We define the following vocabularies

$$\begin{aligned} \Sigma &:= \{ P_i, Q_i \mid 1 \leq i \leq 4n \} \cup \{ R_j, T_j \mid 1 \leq j \leq n \}, \\ \Sigma_{PR} &:= \{ P_i \mid 1 \leq i \leq 4n \} \cup \{ R_j \mid 1 \leq j \leq n \}, \text{ and} \\ \Sigma_{QT} &:= \{ Q_i \mid 1 \leq i \leq 4n \} \cup \{ T_j \mid 1 \leq j \leq n \}. \end{aligned}$$

Moreover, let Σ'_{PR} and Σ'_{QT} be disjoint extensions of the vocabularies Σ_{PR} and Σ_{QT} , respectively, each extended by a countably infinite number of nullary predicate symbols.

An atom is called *linear* if every variable in it occurs at most once. Any occurrence of a variable v in a non-equational Σ -atom A is called a *column- k -occurrence*, if v is the $(n-k+1)$ -st argument from the right in A . For example, if we fix n to be 6, then v has a column-5-occurrence in each of the atoms $Q_i(x_1, x_2, x_3, x_4, v, x_6)$, $T_3(x_3, x_4, v, x_6)$, $T_5(v, x_6)$, but v has no column-5-occurrence in the atoms $T_6(x_6)$ or $Q_i(v, v, v, v, x_5, v)$.

Starting from φ_{LGF} we construct the sentence ψ_{LGF} with the following properties.

- (a) The sentence ψ_{LGF} is a Boolean combination of loosely-guarded Σ'_{PR} -sentences and loosely-guarded Σ'_{QT} -sentences. Moreover, ψ_{LGF} is in negation normal form and none of the constituent sentences of ψ_{LGF} properly contains a non-atomic subsentence.
- (b) The vocabulary of ψ_{LGF} is $\Sigma'_{PR} \cup \Sigma'_{QT}$, i.e. Σ plus fresh nullary predicate symbols.
- (c) The structure \mathcal{A} can be uniquely expanded to some model \mathcal{B} of ψ_{LGF} over the same domain and conserving the interpretations of all predicate symbols in Σ ; for every \mathcal{B}_{-S} — which is defined to be the substructure of \mathcal{B} induced by the domain $\mathbf{B} \setminus \{ \mathbf{b}_S^{(k)} \}$ — we have $\mathcal{B}_{-S} \not\models \psi_{\text{LGF}}$.
- (d) $\text{len}(\psi_{\text{LGF}}) \in \mathcal{O}(\text{len}(\varphi_{\text{LGF}}))$.
- (e) Equations in Σ'_{QT} -subsentences of ψ_{LGF} occur only in guards, and these consist of a single trivial equation, say $v = v$, and no other atoms, and belong to top-most quantifiers in the subsentence.

- (f) Every non-equational Σ_{QT} -atom is linear and for every variable v occurring in it there is some k such that all occurrences of v in non-equational Σ_{QT} -atoms are column- k -occurrences.
- (g) For all distinct variables v, v' that occur freely in a Σ'_{QT} -subformula χ and have column- k -occurrences and column- k' -occurrences, respectively, we know that $k \neq k'$.

Notice that every Σ'_{QT} -subsentence that is part of ψ_{LGF} is actually a variable-renamed version of a loosely guarded FO^n sentence (see Section 9 for a definition of n -variable logics).

The details of the construction of ψ_{LGF} can be found in [69], Theorem 3.10.8. It roughly amounts to the following steps. Certain subformulas are simply replaced with Boolean constants **true** or **false**, other subformulas χ are replaced with fresh nullary predicate symbols, say M , for which definitions of the form $M \leftrightarrow \chi$ are conjoined. The final result has a length that is linear in the length of the original.

Now, suppose that ψ_{LGF} has fewer than $2^{\uparrow n-1}(n)$ subformulas. We observed earlier that $\mathcal{B} \models \psi_{\text{LGF}}$ and $\mathcal{B}_{-S} \not\models \psi_{\text{LGF}}$ for every $S \in \mathcal{S}_n$. Hence, for every $S \in \mathcal{S}_n$ there is some Σ'_{QT} -subformula ψ_S in ψ_{LGF} of the form $(\exists \bar{y}. \gamma_S(\bar{y}, \bar{z})) \chi_S(\bar{y}, \bar{z})$ and some variable assignment β_S such that the following properties hold. We have $\beta_S(y_*) = \mathbf{b}_S^{(n)}$ for exactly one $y_* \in \bar{y}$ and for every $v \in \bar{y} \cup \bar{z}$ different from y_* we have $\beta_S(v) \in \mathbf{A}_b \setminus \mathbf{A}_{b,n}$. Moreover, we have

- (*) $\mathcal{B}, \beta_S \models \gamma_S(\bar{y}, \bar{z}) \wedge \chi_S(\bar{y}, \bar{z})$ and $\mathcal{B}, \beta' \not\models \gamma_S(\bar{y}, \bar{z}) \wedge \chi_S(\bar{y}, \bar{z})$ for every β' that differs from β_S only in the value assigned to y_* .

The tuple $\beta_S(\bar{z})$ represents a sequence \bar{c}_S of domain elements from \mathbf{A}_b that can be completed to a chain $\mathbf{b}_{T_1}^{(1)}, \dots, \mathbf{b}_{T_{n-1}}^{(n-1)}, \mathbf{b}_S^{(n)}$ with $T_1 \in \dots \in T_{n-1} \in S$.

Fix any $S_* \in \mathcal{S}_n$ and consider the formula $\psi_{S_*}(\bar{z})$. There is a nonempty set \widehat{S}_* such that $\psi_{S_*}(\bar{z})$ coincides with every $\psi_S(\bar{z})$ with $S \in \widehat{S}_*$. For any distinct $S, S' \in \widehat{S}_*$ the sequences $\bar{c}_S := \beta_S(\bar{z})$ and $\bar{c}_{S'} := \beta_{S'}(\bar{z})$ must differ, for otherwise (*) would be violated. As there are at most $\prod_{k=1}^{n-1} 2^{\uparrow k}(n)$ distinct sequences \bar{c}_S , \widehat{S}_* can contain at most $\prod_{k=1}^{n-1} 2^{\uparrow k}(n) < (2^{\uparrow n-1}(n))^n$ sets. Recall that there are fewer than $2^{\uparrow n-1}(n)$ subformulas in ψ_{LGF} . We have just inferred that each of these can only serve as ψ_S for at most $(2^{\uparrow n-1}(n))^n$ sets $S \in \mathcal{S}_n$. Hence, only

$$(2^{\uparrow n-1}(n))^n \cdot 2^{\uparrow n-1}(n) = 2^{(n+1) \cdot 2^{\uparrow n-2}(n)} < 2^{2^{\uparrow n-1}(n)} = 2^{\uparrow n}(n)$$

sets S have a corresponding subformula ψ_S . But then $|\mathcal{S}_n| \geq 2^{\uparrow n}(n+1)$ implies that there are $S \in \mathcal{S}_n$ such that $\mathcal{B}_{-S} \models \psi_{\text{LGF}}$, which contradicts our assumptions. Consequently, ψ_{LGF} must have more than $2^{\uparrow n-1}(n)$ subformulas. \square

8 Separateness and Guarded Negation

Bárány, ten Cate, and Segoufin [6, 7] have discovered that guards can be shifted from quantification to negation, see also [61]. This leads to the *guarded-negation first-order fragment (GNFO)*. GNFO comprises all relational first-order formula with equality over the Boolean connectives \neg, \wedge, \vee and existential quantification. Universal quantification has to be simulated using negation, based on the equivalence $\forall x. \psi \models \neg \exists x. \neg \psi$. Every occurrence of the negation sign is accompanied by a guard, i.e. negation may only occur in the form $\gamma(\bar{u}, \bar{v}) \wedge \neg \varphi(\bar{v})$, where γ is an atomic guard and φ is a GNFO formula. In terms of expressiveness, GNFO strictly subsumes GF [7] (Proposition 2.2 and Example 2.3). Moreover, in the same article it is shown that GNFO enjoys the tree-like model property and the finite model property. The associated satisfiability problem is complete for deterministic doubly exponential time. Clique-guarded variants of GNFO have also been studied [7].

Similar to guarded quantification, guarded negation can be made compatible with separateness of variables in a way that allows us to syntactically extend GNFO while retaining its expressive power and the decidability of the associated satisfiability problem (*GNFO-Sat*).

Definition 8.1 (Separated guarded-negation fragment (SGNFO)). Given any sequence $\bar{u}_1, \dots, \bar{u}_n, \bar{v}$ of pairwise-disjoint tuples of first-order variables, a *separated negation guard* $\gamma(\bar{u}_1, \dots, \bar{u}_n, \bar{v})$ is a conjunction of n atoms $A_1(\bar{u}_1, \bar{v}) \wedge \dots \wedge A_n(\bar{u}_n, \bar{v})$ (possibly equations) such that for every i , $1 \leq i \leq n$, all variables from \bar{u}_i occur at least once in $A_i(\bar{u}_i, \bar{v})$.

We define the set of *SGNFO formulas* inductively:

- (i) every relational atom is an SGNFO formula, equality is admitted;
- (ii) every \wedge - \vee -combination of SGNFO formulas is an SGNFO formula;
- (iii) for every tuple \bar{y} and every SGNFO formula $\psi(\bar{y})$ the formula $\exists \bar{y}. \psi(\bar{y})$ is an SGNFO formula;
- (iv) for every separated negation guard $\gamma(\bar{u}_1, \dots, \bar{u}_n, \bar{v})$, and every SGNFO formula $\psi(\bar{u}_1, \dots, \bar{u}_n)$ the formula $\gamma(\bar{u}_1, \dots, \bar{u}_n, \bar{v}) \wedge \neg \psi(\bar{u}_1, \dots, \bar{u}_n)$ is an SGNFO formula if the following conditions are met. Let Z be the set of variables that are quantified in $\psi(\bar{u}_1, \dots, \bar{u}_n)$. We require that Z can be divided into pairwise disjoint, possibly empty subsets Z_1, \dots, Z_n such that the sets $Z_1 \cup \bar{u}_1, \dots, Z_n \cup \bar{u}_n$ are all pairwise separated in $\psi(\bar{u}_1, \dots, \bar{u}_n)$.

The *separated guarded-negation fragment (SGNFO)* is the set of all SGNFO sentences.

It is obvious that GNFO is contained in SGNFO and that there are SGNFO sentences that do not belong to GNFO. Moreover, every MFO sentence φ can be easily turned into an equivalent SGNFO sentence with a length linear in the original. We first transform φ into negation normal form and add trivial equations $v \approx v$ as guards to negated atomic subformulas $\neg P(v)$. The result lies in the intersection of SGNFO and MFO_{\approx} . For MFO_{\approx} sentences the matter seems to be more complicated. The sentence $\exists xy. x \not\approx y$, for instance, is not an SGNFO sentence and does not seem to have an SGNFO equivalent.

Proposition 8.2. SGNFO properly contains GNFO. Moreover, every MFO sentence φ can be turned into an equivalent SGNFO sentence of length $\mathcal{O}(\text{len}(\varphi))$.

After we have seen the results obtained for the other novel first-order fragments, it should not come as a surprise that there is an effective translation from SGNFO to GNFO.

Definition 8.3 (Strict separateness). Let φ be any first-order formula and let X, Y be two disjoint sets of first-order variables. We say that X, Y are *strictly separated in φ* if X and Y are separated in φ and, in addition, for every subformula $\chi := (Qv. \dots)$ of φ we either have $\text{vars}(\chi) \cap X = \emptyset$ or $\text{vars}(\chi) \cap Y = \emptyset$.

Lemma 8.4. Every SGNFO formula is equivalent to some GNFO formula.

Proof. We infer two auxiliary results from which the lemma follows:

Claim I: Consider any SGNFO formula $\varphi(\bar{u}_1, \dots, \bar{u}_n, \bar{v}) := \gamma(\bar{u}_1, \dots, \bar{u}_n, \bar{v}) \wedge \neg \psi(\bar{u}_1, \dots, \bar{u}_n)$ where $\psi(\bar{u}_1, \dots, \bar{u}_n)$ is any GNFO formula, and the $\bar{u}_1, \dots, \bar{u}_n$ are pairwise strictly separated in $\psi(\bar{u}_1, \dots, \bar{u}_n)$. Then, $\varphi(\bar{u}_1, \dots, \bar{u}_n, \bar{v})$ is equivalent to some GNFO formula $\varphi'(\bar{u}_1, \dots, \bar{u}_n, \bar{v})$ in which $\bar{u}_1, \dots, \bar{u}_n$ are pairwise strictly separated.

Proof: Let *basic formulas* in ψ be subformulas that do not lie in the scope of any quantifier or negation sign in ψ and that are either guarded negation formulas $\delta(\bar{x}_1, \dots, \bar{x}_k, \bar{y}) \wedge \neg \chi(\bar{x}_1, \dots, \bar{x}_k)$, quantified formulas $\exists \bar{y}. \chi(\bar{y}, \bar{x})$, or atoms. Transform ψ into a conjunction $\psi' := \bigwedge_{i \in I} \eta_i(\bar{u}_1, \dots, \bar{u}_n)$ of disjunctions η_i of basic formulas. Since we assumed $\psi(\bar{u}_1, \dots, \bar{u}_n)$ to be a GNFO formula in which the $\bar{u}_1, \dots, \bar{u}_n$ are pairwise strictly separated, we conclude that every basic formula $\chi(\bar{x})$ in ψ satisfies $\bar{x} \cap \bar{u}_\ell \neq \emptyset$ for at most one ℓ , $1 \leq \ell \leq n$. Hence, the disjuncts η_i in ψ' can be regrouped such that ψ' has the form

$$\bigwedge_{i \in I} \eta_{i,1}(\bar{u}_1) \vee \dots \vee \eta_{i,n}(\bar{u}_n) .$$

Therefore, $\gamma(\bar{u}_1, \dots, \bar{u}_n, \bar{v}) \wedge \neg\psi(\bar{u}_1, \dots, \bar{u}_n)$ is equivalent to the following sentence, where $A_1(\bar{u}_1, \bar{v}), \dots, A_n(\bar{u}_n, \bar{v})$ is the list of atoms that $\gamma(\bar{u}_1, \dots, \bar{u}_n, \bar{v})$ comprises:

$$\begin{aligned} & A_1(\bar{u}_1, \bar{v}) \wedge \dots \wedge A_n(\bar{u}_n, \bar{v}) \wedge \neg\psi'(\bar{u}_1, \dots, \bar{u}_n) \\ & \models A_1(\bar{u}_1, \bar{v}) \wedge \dots \wedge A_n(\bar{u}_n, \bar{v}) \wedge \neg \bigwedge_{i \in I} \eta_{i,1}(\bar{u}_1) \vee \dots \vee \eta_{i,n}(\bar{u}_n) \\ & \models A_1(\bar{u}_1, \bar{v}) \wedge \dots \wedge A_n(\bar{u}_n, \bar{v}) \wedge \bigvee_{i \in I} \neg\eta_{i,1}(\bar{u}_1) \wedge \dots \wedge \neg\eta_{i,n}(\bar{u}_n) \\ & \models \bigvee_{i \in I} (A_1(\bar{u}_1, \bar{v}) \wedge \neg\eta_{i,1}(\bar{u}_1)) \wedge \dots \wedge (A_n(\bar{u}_n, \bar{v}) \wedge \neg\eta_{i,n}(\bar{u}_n)) \end{aligned}$$

This is the sought GNFO formula. \diamond

Claim II: Consider any SGNFO formula $\varphi(\bar{x}, \bar{v}) := \exists \bar{y}. \psi(\bar{y}, \bar{x}, \bar{v})$ where $\psi(\bar{y}, \bar{x}, \bar{v})$ is any GNFO formula in which the sets $\bar{y} \cup \bar{x}$ and \bar{v} are strictly separated. Then, $\varphi(\bar{x}, \bar{v})$ is equivalent to some GNFO formula $\varphi(\bar{x}, \bar{v})$ in which $\bar{y} \cup \bar{x}$ and \bar{v} are strictly separated.

Proof: Let *basic formulas* in ψ be defined like in the proof of Claim I. Transform ψ into a disjunction $\psi' := \bigvee_{i \in I} \eta_i(\bar{y}, \bar{x}, \bar{v})$ of conjunctions $\eta_i(\bar{y}, \bar{x}, \bar{v})$ of basic formulas. Since we assumed ψ to be a GNFO formula in which the sets $\bar{y} \cup \bar{x}$ and \bar{v} are strictly separated, every basic formula $\chi(\bar{u})$ in ψ satisfies $\bar{u} \cap (\bar{y} \cup \bar{x}) = \emptyset$ or $\bar{u} \cap \bar{v} = \emptyset$. Hence, the conjuncts η_i in ψ' can be regrouped such that ψ' has the form

$$\bigvee_{i \in I} \eta_{i,1}(\bar{y}, \bar{x}) \wedge \eta_{i,2}(\bar{v}).$$

Therefore, $\exists \bar{y}. \psi(\bar{y}, \bar{x}, \bar{v})$ is equivalent to the sentence

$$\exists \bar{y}. \bigvee_{i \in I} \eta_{i,1}(\bar{y}, \bar{x}) \vee \eta_{i,2}(\bar{v}) \quad \models \quad \bigvee_{i \in I} (\exists \bar{y}. \eta_{i,1}(\bar{y}, \bar{x})) \wedge \eta_{i,2}(\bar{v}).$$

This is the sought GNFO formula. \diamond

Now consider any SGNFO formula φ that is not a GNFO formula. Let $\chi(\bar{u}_1, \dots, \bar{u}_n, \bar{v}) := \gamma(\bar{u}_1, \dots, \bar{u}_n, \bar{v}) \wedge \neg\chi'(\bar{u}_1, \dots, \bar{u}_n)$ be a smallest subformula of φ that violates the conditions of guarded negation in GNFO. Hence, the set Z of variables quantified in $\chi'(\bar{u}_1, \dots, \bar{u}_n)$ can be divided into pairwise disjoint sets Z_1, \dots, Z_n such that $Z_1 \cup \bar{u}_1, \dots, Z_n \cup \bar{u}_n$ are pairwise separated in χ' . Further suppose that in χ' the sets $Z_1 \cup \bar{u}_1, \dots, Z_n \cup \bar{u}_n$ are not strictly separated. Let $\eta(\bar{x}) := \exists \bar{y}. \eta'(\bar{y}, \bar{x})$ be a smallest subformula of χ' that violates this strict-separateness condition. Then, we can subdivide \bar{y} into pairwise disjoint parts $\bar{y}_1, \dots, \bar{y}_n$ such that $\bar{y}_i \subseteq Z_i$ for every i . Moreover, we can subdivide \bar{x} into pairwise disjoint parts $\bar{x}_1, \dots, \bar{x}_n$ such that $\bar{x}_i \subseteq Z_i \cup \bar{u}_i$ for every i . Then, $\eta(\bar{x})$ can be rewritten into $\exists \bar{y}_1 \exists \bar{y}_2 \dots \exists \bar{y}_n. \eta'(\bar{y}_1, \bar{x}_1, \dots, \bar{y}_n, \bar{x}_n)$. Since we assume $\eta(\bar{x})$ to be minimal, the sets $\bar{y}_1 \cup \bar{x}_1, \dots, \bar{y}_n \cup \bar{x}_n$ are pairwise strictly separated in $\eta'(\bar{y}_1, \bar{x}_1, \dots, \bar{y}_n, \bar{x}_n)$. By Claim II, $\eta(\bar{x})$ is equivalent to some $\eta''(\bar{x})$ in which the $Z_1 \cup \bar{u}_1, \dots, Z_n \cup \bar{u}_n$ are pairwise strictly separated. Therefore, the formula $\chi'(\bar{u}_1, \dots, \bar{u}_n)$ can be transformed into an equivalent formula $\chi''(\bar{u}_1, \dots, \bar{u}_n)$ in which $Z_1 \cup \bar{u}_1, \dots, Z_n \cup \bar{u}_n$ are pairwise strictly separated. By Claim I, $\gamma(\bar{u}_1, \dots, \bar{u}_n, \bar{v}) \wedge \neg\chi''(\bar{u}_1, \dots, \bar{u}_n)$ can be transformed into an equivalent formula that belongs to GNFO and in which the sets $\bar{u}_1, \dots, \bar{u}_n$ are pairwise strictly separated.

By iterative and exhaustive application of the outlined transformation, we can derive a GNFO formula that is equivalent to the SGNFO formula φ . \square

Since GNFO is known to possess the finite model property [7], Lemma 8.4 entails the same for SGNFO. Of course, this also means that *SGNFO-Sat* is decidable.

Theorem 8.5. SGNFO possesses the finite model property and, hence, SGNFO-Sat is decidable.

We conclude this section with an investigation of the succinctness gap between SGNFO and GNFO. The following theorem entails that there is no elementary upper bound on the length of the GNFO sentences that result from any equivalence-preserving transformation of SGNFO sentences into GNFO.

Theorem 8.6. There is a class of SGNFO sentences such that for every integer $n \geq 3$ the class contains a sentence φ with a length polynomial in n for which any equivalent GNFO sentence has at least $(n - 1)$ -fold exponential length in n .

Proof sketch. Let $n \geq 3$. The following SGNFO sentence is equivalent to the sentence φ given in the proof of Theorem 7.5. In the sentence the sets $\{x_1, \dots, x_n\}$ and $\{y_1, \dots, y_n\}$ and $\{z\}$ are separated:

$$\begin{aligned} & \exists z. z = z \\ & \wedge \neg \exists x_n. R_n(x_n) \wedge \neg \exists y_n. R_n(x_n) \wedge T_n(y_n) \\ & \wedge \neg \exists x_{n-1}. R_{n-1}(x_{n-1}, x_n) \wedge T_n(y_n) \wedge \neg \exists y_{n-1}. R_{n-1}(x_{n-1}, x_n) \wedge T_{n-1}(y_{n-1}, y_n) \\ & \dots \\ & \wedge \neg \exists x_1. R_1(x_1, \dots, x_n) \wedge T_2(y_2, \dots, y_n) \wedge \neg \exists y_1. R_1(x_1, \dots, x_n) \wedge T_1(y_1, \dots, y_n) \\ & \wedge \bigwedge_{i=1}^{4n} (R_i(x_1, \dots, x_n) \wedge T_i(x_1, \dots, x_n) \wedge P_i(x_1, \dots, x_n) \wedge Q_i(y_1, \dots, y_n)) \\ & \vee (R_1(x_1, \dots, x_n) \wedge T_1(x_1, \dots, x_n) \wedge \neg (P_i(x_1, \dots, x_n) \vee Q_i(y_1, \dots, y_n))). \end{aligned}$$

The subformulas $z = z$, $R_n(x_n)$, and the more complex $R_i(x_i, \dots, x_n) \wedge T_i(y_i, \dots, y_n)$ and $R_i(x_i, \dots, x_n) \wedge T_{i+1}(y_{i+1}, \dots, y_n)$ serve as separated negation guards for negated subformulas. We need to introduce a bit of redundancy in order to meet the syntactic requirements of SGNFO. Nevertheless, it is easy to see that the above sentence is equivalent to the sentence φ used in the proof of Theorem 7.5. The rest of the proof works in analogy to that proof. Details can be found in Appendix A. \square

9 Separateness and Finite-Variable First-Order Logic

Classes of first-order formulas over a fixed finite set of variables have been studied extensively (see, e.g., [52, 32, 15, 31, 42] and also [43], Section 11, and [29], Sections 1.1.3, 2.7, and 2.8). If the number of variables is restricted to two, we even obtain a decidable fragment. The *two-variable fragment* (FO^2) comprises all relational first-order sentences with equality that are build up using at most two variables. It is important to understand that this restriction allows reusing variable names in quantifiers — be they nested or not. Therefore, in the formulas in the present section we explicitly allow variables to occur free and bound in a formula, and to reappear in distinct occurrences of quantifiers in the same formula. For example, the sentence $\forall x \exists y. (E(x, y) \wedge \exists x. (E(y, x) \wedge \exists y. E(x, y)))$ belongs to FO^2 .

Scott gave a reduction of FO^2 Sat to GKS-Sat [60]. This reduction works only for sentences without equality. In 1975 Moritmer [50] proved that FO^2 with equality possesses the finite model property. The computational complexity of the satisfiability problem for FO^2 has been determined by Grädel, Kolaitis, and Vardi [30]: it is NEXPTIME-complete. A recent survey of FO^2 and various extensions is [42].

We will see in the present section that also in the context of finite-variable logics separateness can give us more syntactic freedom and the ability to express certain properties in a substantially more succinct way.

Definition 9.1 (Separated finite-variable formulas). For every positive integer k we define FO^k to be the set of all relational first-order formulas with equality in which all variables are taken from a finite sequence x_1, \dots, x_k of pairwise distinct first-order variables.

For every $k \geq 1$ we define the class SFO^k of relational first-order formulas as follows. Let V_1, V_2, V_3, \dots be a sequence of pairwise disjoint sets V_i of first-order variables, each containing exactly k pairwise distinct variables. For every $m \geq 1$ we define the set $\text{SFO}^{k,m}$ to be the set of all relational first-order formulas φ with equality in which all variables are taken from $V_1 \cup \dots \cup V_m$ and in which all sets V_1, \dots, V_m are pairwise separated. The class SFO^k is the union $\bigcup_{m \geq 1} \text{SFO}^{k,m}$.

It is easy to see that FO^k is a special case of SFO^k . Moreover, MFO is a proper subset of SFO^k for $k = 1$. In contrast, for every positive integer k the MFO $_{\approx}$ sentence $\forall x_1 \dots x_k \exists y. \bigwedge_k y \not\approx x_k$ does not belong to SFO^k .

Proposition 9.2. For every positive k , SFO^k contains FO^k and MFO.

In the following lemma we establish the equivalence between SFO^k and FO^k for every positive k by devising an equivalence-preserving translation between the two fragments.

Lemma 9.3. Every SFO^k sentence is equivalent to some FO^k sentence.

Proof. Let m be any positive integer and consider any sentence φ from $\text{SFO}^{k,m}$. Then, $\text{vars}(\varphi) \subseteq V_1 \cup \dots \cup V_m$ and all V_1, \dots, V_m are pairwise separated in φ . Without loss of generality, we assume that φ is in negation normal form.

We prove an auxiliary result from which the lemma follows.

Claim I: Consider any subformula $\psi = \mathcal{Q}\bar{v}. \chi$ of φ with $\bar{v} \subseteq V_i$ for some i . If the sets V_1, \dots, V_m are pairwise strictly separated in χ (cf. Definition 8.3), then we can construct a formula ψ' that is equivalent to ψ and in which all sets V_1, \dots, V_m are pairwise strictly separated.

Proof: A *basic formula* is any atom and any subformula $(\mathcal{Q}'v \dots)$ in χ that does not lie within the scope of any quantifier in χ . Suppose \mathcal{Q} is an existential quantifier. (The case of $\mathcal{Q} = \forall$ can be treated in an analogous way.)

Let \bar{z} be the tuple collecting all variables that occur freely in ψ . We first transform χ into an equivalent disjunction of conjunctions of negated or non-negated basic formulas. This is always possible. Since the sets V_1, \dots, V_m are pairwise strictly separated in χ , the constituents of the j -th conjunction can be grouped into m parts: $\eta_{j,1}(V_1 \cap (\bar{v} \cup \bar{z})), \dots, \eta_{j,m}(V_m \cap (\bar{v} \cup \bar{z}))$ with $\text{vars}(\eta_{j,\ell}) \subseteq V_\ell$. This is possible because of our assumption that the sets V_1, \dots, V_m are all pairwise strictly separated in χ . Hence, since $\bar{v} \subseteq V_i$, ψ is equivalent to a formula of the form

$$\exists \bar{v}. \bigvee_j \eta_{j,1}(V_1 \cap \bar{z}) \wedge \dots \wedge \eta_{j,i}(V_i \cap (\bar{v} \cup \bar{z})) \wedge \dots \wedge \eta_{j,m}(V_m \cap \bar{z}).$$

We shift the existential quantifier block $\exists \bar{v}$ inwards so that it only binds the (sub-)conjunctions $\eta_{j,i}(V_i \cap (\bar{v} \cup \bar{z}))$. The resulting formula

$$\bigvee_j \eta_{j,1}(V_1 \cap \bar{z}) \wedge \dots \wedge \left(\exists \bar{v}. \eta_{j,i}(V_i \cap (\bar{v} \cup \bar{z})) \right) \wedge \dots \wedge \eta_{j,m}(V_m \cap \bar{z}).$$

is the sought ψ' in which the sets V_1, \dots, V_m are all pairwise strictly separated. \diamond

The sets V_1, \dots, V_m are pairwise strictly separated in any quantifier-free subformula of φ . Hence, applying Claim I iteratively, we can transform φ into an equivalent sentence φ' in which the sets V_1, \dots, V_m are pairwise strictly separated. Since φ' is a sentence, the strict separateness condition leads to the observation that for every subformula $\mathcal{Q}\bar{v}. \chi$ in φ' there is some j such that $\text{vars}(\mathcal{Q}\bar{v}. \chi) \subseteq V_j$. As each of the V_i contains exactly k variables, we can rename the bound variables in φ' such that φ' is an \wedge - \vee -combination of FO^k sentences. Since FO^k is closed under Boolean combinations, φ' is an FO^k sentence. \square

Since FO^2 possess the finite model property [50, 30], Lemma 9.3 entails the same for the class of SFO^2 sentences and, hence, SFO^2 -*Sat* is decidable.

Theorem 9.4. The class of SFO^2 sentences possesses the finite model property and, hence, the satisfiability problem for SFO^2 sentences is decidable.

Having established the equivalence between FO^k and SFO^k regarding expressiveness, it remains to investigate the succinctness gap between the two fragments. We will do this for the special case of SFO^2 sentences compared to the class of FO^2 sentences.

Theorem 9.5. There is a class of SFO^2 sentences and some positive integer n_0 such that for every integer $n \geq n_0$ the class contains a sentence φ with a length linear in n for which any equivalent FO^2 sentence has a length that is at least exponential in n .

Proof. Let $n \geq 1$ be some positive integer. Consider the following first-order sentence in which the sets $\{x_1, x_2\}$ and $\{y_1, y_2\}$ are separated:

$$\varphi := \forall x_2 \exists y_2 \forall x_1 \exists y_1. \bigwedge_{i=1}^{2(n+1)} (P_i(x_1, x_2) \leftrightarrow Q_i(y_1, y_2)).$$

In analogy to the proof of Theorem 3.8, we construct a model \mathcal{A} using the sets $\mathcal{S}_1 := \{S \subseteq [2(n+1)] \mid |S| = n+1\}$ and $\mathcal{S}_2 := \{S \subseteq \mathcal{S}_1 \mid |S| = \frac{1}{2}|\mathcal{S}_1|\}$. We observe

$$|\mathcal{S}_1| = \binom{2(n+1)}{n+1} \geq \left(\frac{2(n+1)}{n+1}\right)^{n+1} = 2^{n+1} \text{ and } |\mathcal{S}_2| = \binom{|\mathcal{S}_1|}{|\mathcal{S}_1|/2} \geq \left(\frac{|\mathcal{S}_1|}{|\mathcal{S}_1|/2}\right)^{|\mathcal{S}_1|/2} \geq 2^{2^n}.$$

Claim I: Let $\widehat{\mathcal{S}}$ be any subset of \mathcal{S}_2 such that for every $S \in \widehat{\mathcal{S}}$ there is some $T \in S \subseteq \mathcal{S}_1$ which does not belong to any $S' \in \widehat{\mathcal{S}} \setminus \{S\}$. Then, $\widehat{\mathcal{S}}$ contains at most $|\mathcal{S}_1| \leq 2^{2(n+1)}$ sets as elements.

Proof: Obvious. ◇

Let \mathcal{A} be the structure with

$$\mathbf{A} := \{\langle \mathbf{a}_S^{(1)}, \mathbf{b}_S^{(1)} \mid S \in \mathcal{S}_1 \rangle \cup \{\langle \mathbf{a}_S^{(2)}, \mathbf{b}_S^{(2)} \mid S \in \mathcal{S}_2 \rangle\},$$

$$P_i^{\mathcal{A}} := \{\langle \mathbf{a}_{S_1}^{(1)}, \mathbf{a}_{S_2}^{(2)} \rangle \in \mathbf{A}^2 \mid i \in S_1 \in S_2\} \text{ for } i = 1, \dots, 2(n+1), \text{ and}$$

$$Q_i^{\mathcal{A}} := \{\langle \mathbf{b}_{S_1}^{(1)}, \mathbf{b}_{S_2}^{(2)} \rangle \in \mathbf{A}^2 \mid i \in S_1 \in S_2\} \text{ for } i = 1, \dots, 2(n+1).$$

Like in the proof of Theorem 3.8, it is easy to show that \mathcal{A} is a model of φ , and that the following claim holds.

Claim II: For every $S \in \mathcal{S}_2$ the substructure \mathcal{A}_{-S} of \mathcal{A} induced by $\mathbf{A}_{-S} := \mathbf{A} \setminus \{\mathbf{b}_S^{(2)}\}$ does not satisfy φ . ◇

Let φ_{FO^2} be a shortest FO^2 sentence that is semantically equivalent to φ . Next, we argue that $\text{len}(\varphi_{\text{FO}^2})$ is at least exponential in n . In [60] a normal form for FO^2 sentences was introduced, which is referred to as *Scott normal form* in the literature. For instance, Lemma 8.1.2 in [11] states that there is some relational FO^2 sentence ψ_{FO^2} that has the following properties:

- (a) ψ_{FO^2} is of the form $(\forall uv. \chi(u, v)) \wedge \bigwedge_{i=1}^m \forall x \exists y. \eta_i(x, y)$ with quantifier-free χ and η_i ,
- (b) the vocabulary underlying ψ_{FO^2} is that of φ_{FO^2} extended with fresh unary predicate symbols R_1, \dots, R_κ with $\kappa \in \mathcal{O}(\text{len}(\varphi_{\text{FO}^2}))$,
- (c) $\psi_{\text{FO}^2} \models \varphi_{\text{FO}^2}$,
- (d) every model of φ_{FO^2} can be uniquely expanded to a model of ψ_{FO^2} over the same domain and conserving the interpretations of all predicate symbols in φ_{FO^2} , and
- (e) $\text{len}(\psi_{\text{FO}^2}) \in \mathcal{O}(\text{len}(\varphi_{\text{FO}^2}))$.

Let \mathcal{B} be the unique expansion of \mathcal{A} for which $\mathcal{B} \models \psi_{\text{FO}^2}$ and $\mathbf{B} = \mathbf{A}$. Claim II can be extended to \mathcal{B} , because of $\psi_{\text{FO}^2} \models \varphi_{\text{FO}^2}$. The set $\{\mathbf{b}_S^{(2)} \mid S \in \mathcal{S}_2\}$ can be partitioned into at most 2^κ parts, each containing elements that are indistinguishable by their belonging to the sets $R_k^\mathcal{B}$. Let $\widehat{\mathbf{D}}$ be the largest of these parts and let $\widehat{\mathcal{S}} := \{S \mid \mathbf{b}_S^{(2)} \in \widehat{\mathbf{D}}\}$. Hence, for all $\mathbf{b}, \mathbf{b}' \in \widehat{\mathbf{D}}$ and every k with $1 \leq k \leq \kappa$ we have $\mathbf{b} \in R_k^\mathcal{B}$ if and only if $\mathbf{b}' \in R_k^\mathcal{B}$.

Claim III: Let n be sufficiently large. If κ is polynomial in n , then there is some $S_* \in \widehat{\mathcal{S}}$ such that $\mathbf{b}_{S_*}^{(2)} \in \widehat{\mathbf{D}}$ and for every $T \in S_*$ there is some $S' \in \widehat{\mathcal{S}} \setminus \{S_*\}$ that also contains T and we have $\mathbf{b}_{S'}^{(2)} \in \widehat{\mathbf{D}}$.

Proof: $\widehat{\mathbf{D}}$ contains at least $2^{2^n}/2^\kappa = 2^{2^n - \kappa}$ domain elements. Hence, $|\widehat{\mathcal{S}}| \geq 2^{2^n - \kappa}$. Moreover, we observe $2^{2^n - \kappa} > 2^{2(n+1)}$ for sufficiently large n , if κ is polynomial in n . By Claim I, there is some $S_* \in \widehat{\mathcal{S}}$ such that for every $T \in S_*$ there is some $S' \in \widehat{\mathcal{S}} \setminus \{S_*\}$ with $T \in S'$. Claim III follows by definition of $\widehat{\mathbf{D}}$ and $\widehat{\mathcal{S}}$. \diamond

We fix some $S_* \in \widehat{\mathcal{S}}$ as specified in Claim III. Let \mathcal{B}_{-S_*} be the substructure of \mathcal{B} induced by the domain $\mathbf{B}_* := \mathbf{B} \setminus \{\mathbf{b}_{S_*}^{(2)}\}$. By Claim II (extended to \mathcal{B}), there is some maximal nonempty set $J \subseteq [m]$ such that for every $j \in J$ we have $\mathcal{B}_{-S_*} \not\models \forall x \exists y. \eta_j(x, y)$. Consequently, for every $j \in J$ there is some domain element $\mathbf{c} \in \mathbf{B}_*$ such that $\mathcal{B} \models \eta_j(\mathbf{c}, \mathbf{b}_{S_*}^{(2)})$ and $\mathcal{B} \not\models \eta_j(\mathbf{c}, \mathbf{d})$ for every $\mathbf{d} \in \mathbf{B}_*$. Regarding the domain element \mathbf{c} , we distinguish two cases.

Consider any $j \in J$ and any $\mathbf{c} \in \{\mathbf{b}_S^{(2)} \mid S \in \mathcal{S}_2, S \neq S_*\} \cup \{\mathbf{a}_S^{(1)} \mid S \in \mathcal{S}_1\} \cup \{\mathbf{a}_S^{(2)} \mid S \in \mathcal{S}_2\}$ for which we have $\mathcal{B} \not\models \eta_j(\mathbf{c}, \mathbf{d})$ for every $\mathbf{d} \in \mathbf{B}_*$. Let S' be some set from $\widehat{\mathcal{S}}$ that is different from S_* and for which $\mathbf{c} \neq \mathbf{b}_{S'}^{(2)}$. Notice that η_j is quantifier free and, hence, exclusively contains atoms over two variables x, y . Moreover, for every binary atom $A(x, y)$ of the form $P_i(x, y)$, $P_i(y, x)$, $Q_i(x, y)$, or $Q_i(y, x)$ we have $\mathcal{B} \not\models A(\mathbf{c}, \mathbf{d})$ for every $\mathbf{d} \in \{\mathbf{b}_S^{(2)} \mid S \in \mathcal{S}_2\}$, including $\mathbf{b}_{S_*}^{(2)}$ and $\mathbf{b}_{S'}^{(2)}$. Since all other non-equational atoms occurring in η_j are monadic and because of $\mathbf{b}_{S_*}^{(2)}, \mathbf{b}_{S'}^{(2)} \in \widehat{\mathbf{D}}$, we conclude the following. For every non-equational atom A occurring in η_j we have $\mathcal{B} \models A(\mathbf{c}, \mathbf{b}_{S_*}^{(2)})$ if and only if $\mathcal{B} \models A(\mathbf{c}, \mathbf{b}_{S'}^{(2)})$ (more precisely: $\mathcal{B} \not\models A(\mathbf{c}, \mathbf{b}_{S_*}^{(2)})$ and $\mathcal{B} \not\models A(\mathbf{c}, \mathbf{b}_{S'}^{(2)})$). Consider any equation $x \approx y$. Because of $\mathbf{c} \in \mathbf{B}_{-S_*}$, we have $\mathcal{B} \not\models \mathbf{c} \approx \mathbf{b}_{S_*}^{(2)}$. On the other hand, we also have $\mathcal{B}_{-S_*} \not\models \mathbf{c} \approx \mathbf{b}_{S'}^{(2)}$. But then, we all in all get $\mathcal{B}_{-S_*} \models \eta_j(\mathbf{c}, \mathbf{b}_{S'}^{(2)})$, which entails $\mathcal{B}_{-S_*} \models \exists y. \eta_j(\mathbf{c}, y)$. This leads to a contradiction and, hence, there cannot be any pair j, \mathbf{c} as described.

Consider any $j \in J$ and any $\mathbf{c} = \mathbf{b}_T^{(1)}$ with $T \in \mathcal{S}_1$ for which $\mathcal{B} \not\models \eta_j(\mathbf{c}, \mathbf{d})$ for every $\mathbf{d} \in \mathbf{B}_*$. Suppose $T \notin S_*$. Then, for every binary atom $A(x, y)$ of the form $P_i(x, y)$, $P_i(y, x)$, $Q_i(x, y)$, $Q_i(y, x)$, or $x \approx y$ we have $\mathcal{B} \not\models A(\mathbf{c}, \mathbf{b}_{S_*}^{(2)})$ but also $\mathcal{B} \not\models A(\mathbf{c}, \mathbf{b}_S^{(2)})$ for every $S \in \mathcal{S}_1 \setminus \{T\}$. Like in the above case we conclude $\mathcal{B}_{-S_*} \models \exists y. \eta_j(\mathbf{c}, y)$, yielding a contradiction. Suppose $T \in S_*$. By Claim III, there is some $S' \in \widehat{\mathcal{S}} \setminus \{S_*\}$ such that $T \in S'$ and $\mathbf{b}_{S'}^{(2)} \in \widehat{\mathbf{D}} \setminus \{\mathbf{b}_{S_*}^{(2)}\} \subseteq \mathbf{B}_*$. Then, we have

$$\mathcal{B} \models Q_i(\mathbf{b}_T^{(1)}, \mathbf{b}_{S_*}^{(2)}) \text{ if and only if } i \in T$$

and

$$\mathcal{B} \models Q_i(\mathbf{b}_T^{(1)}, \mathbf{b}_{S'}^{(2)}) \text{ if and only if } i \in T.$$

For every other binary atom $A(x, y)$ of the form $Q_i(y, x)$, $P_i(x, y)$, $P_i(y, x)$, or $x \approx y$ we have $\mathcal{B} \not\models A(\mathbf{b}_T^{(1)}, \mathbf{b}_{S_*}^{(2)})$ and $\mathcal{B} \not\models A(\mathbf{b}_T^{(1)}, \mathbf{b}_{S'}^{(2)})$. For every monadic atom $A(x, y)$ occurring in η_j we have $\mathcal{B} \models A(\mathbf{b}_T^{(1)}, \mathbf{b}_{S_*}^{(2)})$ if and only if $\mathcal{B} \models A(\mathbf{b}_T^{(1)}, \mathbf{b}_{S'}^{(2)})$. All in all, this leads to

$$\mathcal{B} \models \eta_j(\mathbf{b}_T^{(1)}, \mathbf{b}_{S_*}^{(2)}) \text{ if and only if } \mathcal{B} \models \eta_j(\mathbf{b}_T^{(1)}, \mathbf{b}_{S'}^{(2)}).$$

Therefore, we get $\mathcal{B}_{-S_*} \models \exists y. \eta_j(\mathbf{b}_T^{(1)}, y)$, which constitutes a contradiction.

This means, the number κ of unary predicate symbols occurring in ψ_{FO^2} cannot be polynomial in n , for otherwise we get $\mathcal{B}_{-S_*} \models \psi_{\text{FO}^2}$ and $\mathcal{A}_{-S_*} \models \varphi_{\text{FO}^2}$. Since $\kappa \in \mathcal{O}(\text{len}(\varphi_{\text{FO}^2}))$, it follows that $\text{len}(\varphi_{\text{FO}^2})$ cannot be polynomial in n but must be at least exponential, in order to satisfy $2^{2^n} \leq 2^{2(n+1)+\kappa}$ for growing n . \square

10 Separateness and Fluted Formulas

The *fluted fragment* (FL) comprises all relational first-order sentences without equality that satisfy the following properties. Let x_1, x_2, x_3, \dots be a fixed ordered sequence of pairwise distinct variables. For every nonnegative integer k we define the set $FL^{(k)}$ inductively as follows. Any atom $P(x_\ell, \dots, x_k)$ with $\ell \geq 1$ belongs to $FL^{(k)}$ — notice that x_ℓ, \dots, x_k is asserted to be a gap-free subsequence of x_1, x_2, x_3, \dots . The set $FL^{(k)}$ is closed under Boolean combinations, i.e. if φ and ψ belong to $FL^{(k)}$, then so do $\neg\varphi$, $\varphi \wedge \psi$, $\varphi \vee \psi$, $\varphi \rightarrow \psi$, $\varphi \leftrightarrow \psi$. Given any $FL^{(k+1)}$ formula $\varphi(x_1, \dots, x_{k+1})$, then $\forall x_{k+1}.\varphi$ and $\exists x_{k+1}.\varphi$ belong to $FL^{(k)}$. The *fluted fragment* (FL) is the set $FL^{(0)}$, which contains exclusively sentences. Notice that every sentence φ from $FL^{(k)}$ can be turned into an equivalent $FL^{(0)}$ sentence $\forall x_1 \dots x_k.\varphi$ by adding vacuous quantifiers $\forall x_1 \dots x_k$.

The fluted fragment was introduced by Quine in two steps [55, 56]. In an attempt to extrapolate an extension of MFO, Quine [55] considered so-called *homogeneous k -adic* sentences, i.e. FL sentences in which all predicate symbols have arity k . Decidability of *FL-Sat* was shown via an extension of Herbrand's decidability proof for MFO [36]. Later on, namely at the very end of [56], Quine claimed that the same proof would work for full FL. This claim turned out to be wrong [51]. Only recently, a proof including a correct complexity analysis was published [53, 54], showing that FL-Sat is non-elementary.

Herzig [37] considered a class of relational first-order sentences that is very similar to the fluted fragment. *Herzig's ordered fragment* consists of all relational first-order sentences without equality in which every atom $P(v_1, \dots, v_m)$ satisfies the following property. For every i the (unique) quantifier Qv_i binding v_i lies within the scope of any quantifier $Q'u$ if and only if $Q'u$ binds one of the v_j with $j < i$, i.e. $u \in \{v_1, \dots, v_{i-1}\}$. Notice that the definition implies that the v_1, \dots, v_m are pairwise distinct. While atoms in fluted formulas $\varphi \in FL^{(k)}$ need to contain a contiguous suffix of the variable sequence x_1, \dots, x_k , any atom A in Herzig's ordered formulas must contain a contiguous prefix of the variables bound by the quantifier sequence governing A .

As the two fragments seem to be so similar, one could ask whether they are equivalent in expressiveness. Indeed, using the concept of separateness of variables, we can reconcile the two fragments while, at the same time, extending both of them to a common superclass, called the *separated fluted fragment* (SFL).

In the first-order formulas in this section we allow bound variables to reappear in distinct occurrences of quantifiers in the same formula. Before we formulate the definition of SFL, we adapt the following notation from the definition of Maslov's fragment K (Definition 5.10). Let $\psi(u_1, \dots, u_m)$ be any subformula of a first-order sentence φ . We assume that u_1, \dots, u_m are exactly the variables occurring freely in ψ and that they are pairwise distinct. The φ -*prefix* of ψ is the sequence $Q_1v_1 \dots Q_mv_m$ of quantifiers in φ (read from left to right) that bind the free variables of ψ , in particular, we have $\{v_1, \dots, v_m\} = \{u_1, \dots, u_m\}$.

Definition 10.1 (Separated fluted fragment (SFL)). Let $\mathcal{V}_1, \mathcal{V}_2, \mathcal{V}_3, \dots$ be disjoint ordered sequences of pairwise distinct variables $\mathcal{V}_i = x_1^i, x_2^i, x_3^i, \dots$. In what follows, we occasionally treat the sequences \mathcal{V}_i as sets.

The *separated fluted fragment* (SFL) comprises all relational first-order sentences φ without equality in which every atom A satisfies the following properties.

- (a) A is of the form $P(x_\ell^i, \dots, x_k^i)$ for some predicate symbol P and certain integers i, k, ℓ with $i \geq 1$, $k \geq 0$, and $1 \leq \ell \leq k$.
- (b) The φ -prefix of A is of the form $Q_\ell x_\ell^i, \dots, Q_k x_k^i$ with $Q_j \in \{\exists, \forall\}$.

Although separateness is not explicitly mentioned in the definition of SFL, it implicitly plays an important role. For every atom A in any SFL sentence φ , we find one sequence \mathcal{V}_i from which all variables in A stem, i.e. $\text{vars}(A) \subseteq \mathcal{V}_i$. Since the $\mathcal{V}_1, \mathcal{V}_2, \mathcal{V}_3, \dots$ are pairwise disjoint, they are, hence, also pairwise separated in φ .

Example 10.2. A fluted sentence:

$$\forall x_1 \exists x_2. (\forall x_3. P(x_1, x_2, x_3)) \wedge (\exists x_3 \forall x_4. Q(x_2, x_3, x_4)) .$$

An SFL sentence that is not fluted:

$$\forall x_1 \exists x_2. (\forall x_2 \forall x_3. P(x_1, x_2, x_3)) \wedge (\exists x_3 \forall x_4. Q(x_2, x_3, x_4)) .$$

It is not hard to see that every FL sentence also belongs to SFL. The simple monadic sentence $\forall x_1^1 \exists x_1^2. P(x_1^1) \leftrightarrow Q(x_1^2)$ is neither fluted nor does it belong to Herzig's fluted fragment. However, it belongs to SFL. Indeed, every MFO sentence can be turned into an SFL sentence by renaming bound variables. Consider any MFO sentence φ and suppose that all quantifiers in φ bind distinct variables. Let u_1, \dots, u_k be an enumeration of all the first-order variables occurring in φ . Let φ' be the sentence that results from φ by renaming every u_i into x_1^i . This sentence φ' clearly belongs to SFL.

Finally, consider any sentence ψ that belongs to Herzig's ordered fragment. Let $P(u_1, \dots, u_m)$ and $Q(v_1, \dots, v_{m'})$ be two atoms in ψ . Let j, j' be any two indices with $1 \leq j \leq m$ and $1 \leq j' \leq m'$ such that u_j and $v_{j'}$ are bound by the same quantifier $Q_j u_j = Q'_{j'} v_{j'}$ in ψ . By definition of Herzig's ordered fragment, these quantifiers $Q_j u_j$ and $Q'_{j'} v_{j'}$ are exactly in the scopes of $Q_1 u_1, \dots, Q_{j-1} u_{j-1}$ and $Q'_1 v_1, \dots, Q'_{j'-1} v_{j'-1}$, respectively, and no other quantifier scopes. As the quantifiers $Q_j u_j$ and $Q'_{j'} v_{j'}$ coincide, the sets $\{u_1, \dots, u_j\}$ and $\{v_1, \dots, v_{j'}\}$ must be equal. Applying this argument iteratively, we infer $j = j'$ and that the sequences u_1, \dots, u_j and $v_1, \dots, v_{j'}$ coincide. Suppose $j_* \geq 1$ is the maximal index such that u_{j_*} and v_{j_*} are bound by the same quantifier. For any indices $\ell, \ell' > j_*$ we have that neither of the quantifiers $Q u_\ell$ and $Q' v_{\ell'}$ binding the variables u_ℓ and $v_{\ell'}$, respectively, lies in the scope of the other. For otherwise, assume that $Q u_\ell$ were in the scope of $Q' v_{\ell'}$. Hence, $\ell' < \ell$ and there is some $u_{\ell''}$ with $\ell'' < \ell$ such that $u_{\ell''}$ is also bound by the quantifier $Q' v_{\ell'}$. By the above argument, we have that $\ell' = \ell''$ and that the sequences $u_1, \dots, u_{\ell'}$ and $v_1, \dots, v_{\ell'}$ must coincide. But since j_* is maximal and $j_* < \ell'$, we get a contradiction. Consequently, we can rename the bound variables in ψ in such a way that every atom A has the form $P(x_1^1, \dots, x_k^1)$ for some k and the ψ -prefix of A is of the form $Q_1 x_1^1, \dots, Q_k x_k^1$.

Proposition 10.3. SFL properly contains (modulo renaming of bound variables) FL, MFO, and Herzig's ordered fragment.

The following lemma stipulates that every SFL sentence has an equivalent in FL. As usual, this result is established via an effective equivalence-preserving translation from SFL to FL.

Lemma 10.4. Every SFL sentence is equivalent to some FL sentence.

Proof. For every nonnegative integer k and every positive integer i we define the set $\text{FL}^{(k)}(\mathcal{V}_i)$ inductively as follows. Any atom $P(x_\ell^i, \dots, x_k^i)$ with $1 \leq \ell \leq k$ belongs to $\text{FL}^{(k)}(\mathcal{V}_i)$. The set $\text{FL}^{(k)}(\mathcal{V}_i)$ is closed under Boolean combinations. Given any $\text{FL}^{(k+1)}(\mathcal{V}_i)$ formula $\varphi(x_\ell^i, \dots, x_{k+1}^i)$, then $\forall x_{k+1}. \varphi$ and $\exists x_{k+1}. \varphi$ belong to $\text{FL}^{(k)}(\mathcal{V}_i)$.

Consider any SFL sentence φ and let m be the smallest integer such that $\text{vars}(\varphi) \subseteq \mathcal{V}_1 \cup \dots \cup \mathcal{V}_m$. Then, all $\mathcal{V}_1, \dots, \mathcal{V}_m$ are pairwise separated in φ . Without loss of generality, we assume that φ is in negation normal form.

Claim I: Consider any subformula $\psi = Q x_k^i. \chi$ of φ with $Q \in \{\forall, \exists\}$ that satisfies the following properties:

- (a) χ is a Boolean combination of formulas from $\bigcup_{k', i'} \text{FL}^{(k')}(\mathcal{V}_{i'})$ — which we will call *basic formulas* in what follows;
- (b) each of these basic formulas that contains x_k^i is an $\text{FL}^{(k)}(\mathcal{V}_i)$ formula;
- (c) every subformula of χ that is of the form $Q'' x_k^i. \chi''$ is an $\text{FL}^{(k-1)}(\mathcal{V}_i)$ formula.

Then, we can construct a formula ψ' such that

- (1) ψ' is equivalent to ψ ,
- (2) ψ' is a Boolean combination of formulas from $\bigcup_{k', i'} \text{FL}^{(k')}(\mathcal{V}_{i'})$, and
- (3) every subformula $Q x_k^i. \chi'$ occurring in ψ' belongs to $\text{FL}^{(k-1)}(\mathcal{V}_i)$.

Proof: We treat the case where \mathcal{Q} is an existential quantifier; the case of $\mathcal{Q} = \forall$ can be treated dually.

First, we transform χ into an equivalent disjunction of conjunctions of basic formulas that is of the form

$$\bigvee_j \eta_{j,i,k}(x_1^i, \dots, x_k^i) \wedge \bigwedge_{i' \neq i} \bigwedge_{k'} \eta'_{j,i',k'}(x_1^{i'}, \dots, x_{k'}^{i'}),$$

where we group the basic formulas in accordance with their belonging to the sets $\text{FL}^{(k')}(\mathcal{V}_{i'})$. More precisely, the conjunctions $\eta_{j,i,k}$ contain exactly those basic formulas from the j -th disjunct in which the variable bound by $\mathcal{Q}x_k^i$ occurs freely. Moreover, any basic formula from $\text{FL}^{(k')}(\mathcal{V}_{i'})$ that occurs in the j -th disjunct and does not contain x_k^i as free variable is a conjunct of $\eta'_{j,i',k'}$. By assumption, each $\eta_{j,i,k}$ belongs to $\text{FL}^{(k)}(\mathcal{V}_i)$, as we assumed that every basic formula in which x_k^i occurs is an $\text{FL}^{(k)}(\mathcal{V}_i)$ formula.

Hence, ψ is equivalent to a formula of the form

$$\exists x_k^i. \bigvee_j \eta_{j,i,k}(x_1^i, \dots, x_k^i) \wedge \bigwedge_{i' \neq i} \bigwedge_{k'} \eta_{j,i',k'}(x_1^{i'}, \dots, x_{k'}^{i'}).$$

We shift the existential quantifier $\exists x_k^i$ inwards so that it only binds the (sub-)conjunctions $\eta_{j,i,k}$. The emerging subformula $\exists x_k^i. \eta_{j,i,k}$ belongs to $\text{FL}^{(k-1)}(\mathcal{V}_i)$. The result

$$\bigvee_j \left(\exists x_k^i. \eta_{j,i,k}(x_1^i, \dots, x_k^i) \right) \wedge \bigwedge_{i' \neq i} \bigwedge_{k'} \eta_{j,i',k'}(x_1^{i'}, \dots, x_{k'}^{i'})$$

is the sought ψ' that is a Boolean combination of formulas from $\bigcup_{k',i'} \text{FL}^{(k')}(\mathcal{V}_{i'})$. \diamond

By Definition 10.1, every atom in φ is an $\text{FL}^{(k')}(\mathcal{V}_{i'})$ formula for certain k', i' . Hence, every subformula $\mathcal{Q}x_k^i. \chi$ of φ with quantifier-free χ satisfies the conditions of Claim I. Consider any subformula $\psi := \mathcal{Q}x_k^i. \chi$ of φ such that χ is a Boolean combination of atoms and of formulas $\psi' := \mathcal{Q}'x_{k'}^{i'}. \chi'$ that satisfy the preconditions of Claim I. By Claim I, we can transform all these ψ' into equivalent formulas ψ'' in such a way that ψ , after all these transformations, satisfies the preconditions of Claim I. Due to this observation, we can iteratively apply Claim I to transform the sentence φ into an equivalent sentence φ' that is a Boolean combination of sentences from $\bigcup_{k',i'} \text{FL}^{(k')}(\mathcal{V}_{i'})$. Since every sentence $\chi \in \text{FL}^{(k')}(\mathcal{V}_{i'})$ is equivalent to the sentence $\forall x_1^{i'} \dots x_{k'}^{i'}. \chi$, we can transform φ' into an equivalent sentence φ'' that is a Boolean combination of sentences from $\bigcup_{i'} \text{FL}^{(0)}(\mathcal{V}_{i'})$. In φ'' the sets $\mathcal{V}_1, \dots, \mathcal{V}_m$ are pairwise strictly separated. Hence, we can rename bound variables in φ'' in such a way that the result φ''' is a Boolean combination of sentences from $\text{FL}^{(0)}(\mathcal{V}_1)$. This sentence φ''' belongs to the fluted fragment. \square

Since FL enjoys the finite model property [53], Lemma 10.4 implies that the same holds true for SFL. Hence, *SFL-Sat* is decidable.

Theorem 10.5. SFL possesses the finite model property and, hence, SFL-Sat is decidable.

11 Conclusion

In the present paper we have treated separateness of first-order variables in the context of the classical decision problem. Separateness turned out to be an enabler for the definition of significant syntactic extensions of nine of the best-known decidable first-order fragments — see Figure 1 in Section 1. Hence, separateness opens a new perspective on the landscape that research activity around the classical decision problem has revealed over the course of about a hundred years. It seems likely that separateness facilitates more extensions of decidable first-order fragments. For instance, Maslov's fragment K may be an interesting candidate for being extended, as may be the more recent unary-negation fragment [62] and the uniform one-dimensional fragment [41].

Interestingly, each and every of the novel fragments discussed in the present paper subsumes MFO. The inclusion of MFO yields a litmus test for the generality of definitions based on separateness: if MFO is not covered, then the definition is not yet liberal enough. Another peculiarity is that each of our extended fragments exhibits the same expressiveness as the underlying original fragment does. This is witnessed by the equivalence-preserving translation procedures that we have devised for each extended fragment, say F , into the respective original fragment, say G . From this perspective, the syntax of G could be conceived as a normal form of F -sentences: there is a procedure bringing any F -sentence into G -normal form, so to speak. Furthermore, we have seen that this translation in some cases inevitably leads to a super-polynomial or even non-elementary blowup of the formula length in the worst case — take a look at Table 1 again for a summary. This shows that separateness provides the ability to express certain logical properties in significantly more succinct ways.

In the present paper the main method for proving decidability of the newly introduced fragments is based on syntactic transformations. In [69], Chapter 4, this syntactic point of view is complemented with a semantic perspective, based on an investigation of dependences between existentially and universally quantified variables in sentences. In general, nested first-order quantification introduces such dependences. For example, consider the first-order sentence $\varphi := \forall x \exists y. P(x) \leftrightarrow Q(y)$. Skolemization of the quantifier $\exists y$ removes the quantifier and replaces every occurrence of y with the *Skolem term* $f(x)$. The result is the equivalent second-order sentence $\varphi_{\text{Sk}} := \exists f. \forall x. P(x) \leftrightarrow Q(f(x))$, where the dependence of the argument of Q on x is explicit. We distinguish two kinds of dependences that occur between existentially quantified variables and universally quantified variables. Roughly speaking, such a dependence in a sentence φ is *strong*, if there are models (with an infinite domain) where the range of the Skolem function f has to be infinite. On the other hand, we speak of *weak* dependence of existentially quantified y on universally quantified x in the following case. If the values of all other variables are fixed and only x does not have a fixed value, then the range of y (and the Skolem function f) can always, i.e. in all models, be restricted to a finite set of values. The following definition makes this intuition more precise.

Definition 11.1 (Weak dependence). Consider any satisfiable relational first-order sentence ψ that contains some subformula $\chi := \exists y. \chi'(\bar{u}, \bar{v}, \bar{x}, y)$ such that the variables from \bar{u} and \bar{x} are universally quantified in ψ , and the variables from \bar{v} are existentially quantified in ψ . Let ψ_{Sk} be the result of replacing every occurrence of y in ψ with the Skolem term $f(\bar{u}, \bar{x})$ for some fresh Skolem function f . Then, y *depends weakly* on the variables in \bar{x} , if every model $\mathcal{A} \models \psi_{\text{Sk}}$ can be turned into a model $\mathcal{B} \models \psi_{\text{Sk}}$ by replacing $f^{\mathcal{A}}$ with some mapping $f^{\mathcal{B}}$ satisfying the following property. There exists a *finite* family of mappings $(g_i : A^{|\bar{u}|} \rightarrow A)_{i \in I}$ and some mapping $h : A^{|\bar{x}|} \rightarrow I$ such that $f^{\mathcal{B}}(\bar{a}, \bar{b}) := g_{h(\bar{b})}(\bar{a})$ for all $\bar{a} \in A^{|\bar{u}|}$ and $\bar{b} \in A^{|\bar{x}|}$.

Regarding weak dependences, BSR, SF, and SBSR are special fragments, since all dependences in such sentences are weak. In addition, and less obvious, every sentence in which all dependences are weak is equivalent to some BSR sentence [69] (Theorem 4.2.1 and Remark 4.2.12). Hence, it is fair to say that the property of containing exclusively weak dependences between existential and universal variables is a semantic characterization of BSR (and also of its syntactic extensions SF and SBSR).

Integrating the analysis of weak dependences into Skolemization strategies could make them sensitive to weak dependences. This might offer interesting and valuable directions for research in the fields of automated reasoning and proof complexity, where improvements to Skolemization strategies can have tremendous effects on the length of proofs [5, 20].

Possible connections of weakness of dependences in the sense of Definition 11.1 to the field of dependence logic (broadly construed) [66, 1] remain to be investigated.

Separateness may turn out to be even more versatile in future investigations. Some hints are given in [69] (e.g. in Chapter 7). We will conclude with a sketch of one more idea. We have emphasized time and again that, compared to BSR sentences, SF sentences can express certain logical properties much more succinctly. This holds true in particular for properties that exhibit

a high degree of *structural regularity*. An example for such a property is the one described by the family of SF sentences $(\varphi_n)_{n \geq 1}$ with

$$\varphi_n := \forall x_n \exists y_n \dots \forall x_1 \exists y_1. \bigwedge_{i=1}^n (P_i(x_1, \dots, x_n) \leftrightarrow Q_i(y_1, \dots, y_n)).$$

We have already encountered variants of this class of sentences in all the proofs of succinctness gaps (Theorems 3.8, 6.8, 7.5, 8.6, 9.5). Although the domain size of the following family of models $(\mathcal{A}_n)_{n \geq 1}$ with $\mathcal{A}_n \models \varphi_n$ grows massively with increasing n , its interpretation of the predicate symbols P_i and Q_i is given by a rather simple pattern and, hence, each \mathcal{A}_n is very regular in an intuitive sense — the latter is witnessed by the shortness of the following definition of \mathcal{A}_n :

$$\mathcal{A}_n := \bigcup_{k=1}^n \{ \mathbf{a}_S^{(k)}, \mathbf{b}_S^{(k)} \mid S \in \mathcal{P}^k([n]) \}, \text{ where } \mathcal{P}^k \text{ is the } k\text{-fold power set operator,}$$

$$P_i^{\mathcal{A}_n} := \{ \langle \mathbf{a}_{S_1}^{(1)}, \dots, \mathbf{a}_{S_n}^{(n)} \rangle \in \mathcal{A}^n \mid i \in S_1 \in S_2 \in \dots \in S_n \} \text{ for } i = 1, \dots, n, \text{ and}$$

$$Q_i^{\mathcal{A}_n} := \{ \langle \mathbf{b}_{S_1}^{(1)}, \dots, \mathbf{b}_{S_n}^{(n)} \rangle \in \mathcal{A}^n \mid i \in S_1 \in S_2 \in \dots \in S_n \} \text{ for } i = 1, \dots, n.$$

Any of the structures \mathcal{A}_n neatly captures the essence of the logical property described by φ_n , as every domain element $\mathbf{a}_S^{(k)}$ has a corresponding twin element $\mathbf{b}_S^{(k)}$ that mirrors in the predicates $Q_i^{\mathcal{A}_n}$ exactly the role that $\mathbf{a}_S^{(k)}$ plays in the predicates $P_i^{\mathcal{A}_n}$.

More generally, consider any logical property π_n that is parameterized by some positive integer n and that can be expressed by a (uniform) family of BSR sentences. Let $f(n)$ be the function representing the length of a shortest BSR sentence ψ that describes π_n . Let $g(n)$ be the function that denotes the length of a shortest SF sentence describing π_n . By Theorem 3.8, we know that there are properties π_n such that $g(n)$ can be bounded from above by some polynomial but we cannot find any integer k such that $f(n)$ is bounded from above by any k -fold exponential function. In such a case we would say that π_n is structurally fairly regular, as we can describe it with an SF sentence of polynomial length. Now imagine a property π'_n accompanied with corresponding functions $f'(n)$ and $g'(n)$ for which we have $g'(n) \in \Omega(f'(n))$, i.e. the length of shortest SF sentences describing π'_n is asymptotically of the same order as the length of shortest BSR sentences describing π'_n . On an intuitive level, this means that the relaxed syntactic conditions of SF do not provide a significant edge over BSR when π'_n is to be described. It seems that π'_n requires a more sophisticated and lengthy description than, for instance, π_n does, or, viewed from the opposite angle, π'_n exhibits a lower degree of structural regularity than π_n . A possible measure for this lack of regularity might be provided by the gap between $f'(n)$ and $g'(n)$: the smaller the gap, the higher the *structural irregularity* of π'_n .

Instead of the comparison SF versus BSR, one could also use the comparison between SF sentences and equivalent Gaifman-local sentences. Of course, the above said is also relevant to other fragments and not exclusively to SF — for instance, to SLGF versus LGF, to SFO² versus FO², or to the full class of relational first-order sentences versus relational Gaifman-local sentences.

The general idea of measuring structural regularity by means of the asymptotic length of shortest logical descriptions appears to have some similarity to concepts investigated in the field of algorithmic information theory and Kolmogorov complexity (see [17, 13] for introductory material). Potential connections and interrelations remain to be studied.

References

- [1] Samson Abramsky, Juha Kontinen, Jouko Väänänen, and Heribert Vollmer, editors. *Dependence Logic, Theory and Applications*. Springer, 2016.
- [2] Wilhelm Ackermann. Über die Erfüllbarkeit gewisser Zählausdrücke. *Mathematische Annalen*, 100:638–649, 1928.
- [3] Wilhelm Ackermann. *Solvable Cases of the Decision Problem*. North-Holland, 1954.
- [4] Hajnal Andréka, István Németi, and Johan van Benthem. Modal Languages and Bounded Fragments of Predicate Logic. *Journal of Philosophical Logic*, 27(3):217–274, 1998.
- [5] Matthias Baaz and Alexander Leitsch. On Skolemization and Proof Complexity. *Fundamenta Informaticae*, 20(4):353–379, 1994.

- [6] Vince Bárány, Balder ten Cate, and Luc Segoufin. Guarded negation. In *Automata, Languages and Programming (ICALP'11), Part II*, pages 356–367, 2011.
- [7] Vince Bárány, Balder ten Cate, and Luc Segoufin. Guarded Negation. *Journal of the ACM*, 62(3):22, 2015.
- [8] Heinrich Behmann. Beiträge zur Algebra der Logik, insbesondere zum Entscheidungsproblem. *Mathematische Annalen*, 86(3–4):163–229, 1922.
- [9] Paul Bernays and Moses Schönfinkel. Zum Entscheidungsproblem der mathematischen Logik. *Mathematische Annalen*, 99(1):342–372, 1928.
- [10] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2002.
- [11] Egon Börger, Erich Grädel, and Yuri Gurevich. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer, 1997.
- [12] Simone Bova and Fabio Mogavero. Herbrand property, finite quasi-Herbrand models, and a Chandra–Merlin theorem for quantified conjunctive queries. In *Logic in Computer Science (LICS'17)*, pages 1–12, 2017.
- [13] Cristian Calude. *Information and Randomness : An Algorithmic Perspective*. Springer, second, revised and extended edition, 2002.
- [14] Thomas H. Cormen, Clifford Stein, Ronald L. Rivest, and Charles E. Leiserson. *Introduction to Algorithms*. McGraw-Hill Higher Education, second edition, 2001.
- [15] Anuj Dawar. Finite models and finitely many variables. In D. Niwinski and R. Maron, editors, *Logic, Algebra and Computer Science*, volume 46 of *Banach Center Publications*, pages 93–117. Polish Academy of Sciences, 1999.
- [16] Anuj Dawar, Martin Grohe, Stephan Kreutzer, and Nicole Schweikardt. Model Theory Makes Formulas Large. In *Automata, Languages and Programming (ICALP'07)*, pages 913–924, 2007.
- [17] Rodney G. Downey and Denis R. Hirschfeldt. *Algorithmic Randomness and Complexity*. Springer, 2010.
- [18] Burton Dreben and Warren D. Goldfarb. *The Decision Problem: Solvable Classes of Quantificational Formulas*. Addison-Wesley, 1979.
- [19] Heinz-Dieter Ebbinghaus, Jörg Flum, and Wolfgang Thomas. *Mathematical Logic*. Springer, second edition, 1994.
- [20] Uwe Egly. On the Value of Antiprenexing. In *Logic Programming and Automated Reasoning (LPAR'94)*, pages 69–83, 1994.
- [21] Christian G. Fermüller, Alexander Leitsch, Ullrich Hustadt, and Tanel Tammet. Resolution Decision Procedures. In Alan Robinson and Andrei Voronkov, editors, *Handbook of Automated Reasoning*, volume II, pages 1791–1849. Elsevier and MIT Press, 2001.
- [22] Haim Gaifman. On Local and Non-Local Properties. In J. Stern, editor, *Proceedings of the Herbrand Symposium, Logic Colloquium '81*, pages 105–135. North-Holland, 1982.
- [23] Kurt Gödel. Ein Spezialfall des Entscheidungsproblems der theoretischen Logik. *Ergebnisse eines mathematischen Kolloquiums*, (2):27–28, 1932. English translation in Solomon Feferman, John W. Dawson, Jr., Stephen C. Kleene, Gregory H. Moore, Robert M. Solovay, and Jean van Heijenoort, *Kurt Gödel: Collected Works: Volume I: Publications 1929-1936*, Oxford University Press, 1986, pp. 230–233.

- [24] Kurt Gödel. Zum Entscheidungsproblem des logischen Funktionenkalküls. *Monatshefte für Mathematik und Physik*, 40:433–443, 1933. Reprinted in Solomon Feferman, John W. Dawson, Jr., Stephen C. Kleene, Gregory H. Moore, Robert M. Solovay, and Jean van Heijenoort, *Kurt Gödel: Collected Works: Volume I: Publications 1929-1936*, Oxford University Press, 1986, pp. 306–326.
- [25] Warren D. Goldfarb. The Unsolvability of the Gödel Class with Identity. *Journal of Symbolic Logic*, 49(4):1237–1252, 1984.
- [26] Warren D. Goldfarb, Yuri Gurevich, and Saharon Shelah. A Decidable Subclass of the Minimal Gödel Class with Identity. *Journal of Symbolic Logic*, 49(4):1253–1261, 1984.
- [27] Erich Grädel. Invited Talk: Decision procedures for guarded logics. In *Automated Deduction (CADE-16)*, pages 31–51, 1999.
- [28] Erich Grädel. On the Restraining Power of Guards. *Journal of Symbolic Logic*, 64:1719–1742, 12 1999.
- [29] Erich Grädel, Phokion G. Kolaitis, Leonid Libkin, Maarten Marx, Joel Spencer, Moshe Y. Vardi, Yde Venema, and Scott Weinstein. *Finite Model Theory and Its Applications*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2007.
- [30] Erich Grädel, Phokion G. Kolaitis, and Moshe Y. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic*, 3(1):53–69, 1997.
- [31] Erich Grädel and Martin Otto. On Logics with Two Variables. *Theoretical Computer Science*, 224(1-2):73–113, 1999.
- [32] Martin Grohe. Finite variable logics in descriptive complexity theory. *Bulletin of Symbolic Logic*, 4(4):345–398, 1998.
- [33] Y. Gurevich. Formulas with one \forall . In *Selected Questions in Algebra and Logic; in memory of A. Malcev*, pages 97–110. Nauka, Moscow, 1973. In Russian. A German translation is available at TIB Universität Hannover, Germany.
- [34] Yuri Gurevich. The Decision Problem for the Logic of Predicates and Operations. *Algebra i Logika*, 8:284–308, 1969.
- [35] Yuri Gurevich and Saharon Shelah. Random Models and the Gödel Case of the Decision Problem. *Journal of Symbolic Logic*, 48(4):1120–1124, 1983.
- [36] Jacques Herbrand. *Recherches sur la théorie de la démonstration*. PhD thesis, L’Université de Paris, 1930. English translation in Jacques Herbrand, *Logical Writings* (edited by Warren D. Goldfarb), D. Reidel Publishing Company, 1971.
- [37] Andreas Herzig. A new decidable fragment of first order logic. In *Abstracts of the Third Logical Biennial, Summer School & Conference in Honour of S. C. Kleene, Varna, Bulgaria*, 1990.
- [38] Ian M. Hodkinson. Loosely Guarded Fragment of First-Order Logic has the Finite Model Property. *Studia Logica*, 70(2):205–240, 2002.
- [39] Ullrich Hustadt and Renate A. Schmidt. Maslov’s Class K Revisited. In *Automated Deduction (CADE-16)*, pages 172–186, 1999.
- [40] László Kalmár. Über die Erfüllbarkeit derjenigen Zählausdrücke, welche in der Normalform zwei benachbarte Allzeichen enthalten. *Mathematische Annalen*, 108:466–484, 1933.

- [41] Emanuel Kieronski and Antti Kuusisto. Complexity and Expressivity of Uniform One-Dimensional Fragment with Equality. In *Mathematical Foundations of Computer Science (MFCS'14)*, pages 365–376, 2014.
- [42] Emanuel Kieroński, Ian Pratt-Hartmann, and Lidia Tendera. Two-variable logics with counting and semantic constraints. *SIGLOG News*, 5(3):22–43, 2018.
- [43] Leonid Libkin. *Elements of Finite Model Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2004.
- [44] Martin H. Löb. Decidability of the monadic predicate calculus with unary function symbols. *Journal of Symbolic Logic*, 32:563, 1967.
- [45] Leopold Löwenheim. Über Möglichkeiten im Relativkalkül. *Mathematische Annalen*, 76:447–470, 1915. English translation in Jean van Heijenoort. *From Frege to Gödel – A Source Book in Mathematical Logic, 1879–1931*, Harvard University Press, 2002, pp. 228–251.
- [46] Maarten Marx. Tolerance Logic. *Journal of Logic, Language and Information*, 10(3):353–374, 2001.
- [47] Sergei Yu. Maslov. The Inverse Method for Establishing Deducibility for Logical Calculi (in Russian). *Trudy Matem. Inst. AN SSSR*, 98:26–87, 1968. English translation in *Proceedings of the Steklov Institute of Mathematics* 98:25–95, 1968. American Mathematical Society, 1971.
- [48] Sergei Yu. Maslov and V. P. Orevkov. Decidable classes reducing to a one-quantifier class. *Trudy Matem. Inst. AN SSSR*, 121, 1972. In Russian. English translation in *Proceedings of the Steklov Institute of Mathematics* 121:61–72, 1972. American Mathematical Society, 1974.
- [49] Fabio Mogavero and Giuseppe Perelli. Binding Forms in First-Order Logic. In *Computer Science Logic (CSL'15)*, pages 648–665, 2015.
- [50] Michael Mortimer. On Languages with Two Variables. *Mathematical Logic Quarterly*, 21(1):135–140, 1975.
- [51] Aris Noah. Predicate-functors and the limits of decidability in logic. *Notre Dame Journal of Formal Logic*, 21(4):701–707, 1980.
- [52] Martin Otto. *Bounded Variable Logics and Counting: A Study in Finite Models*, volume 9 of *Lecture Notes in Logic*. Springer, 1997. Reprinted by Cambridge University Press in 2017.
- [53] Ian Pratt-Hartmann, Wiesław Szwaśt, and Lidia Tendera. Quine’s Fluted Fragment is Non-Elementary. In *Computer Science Logic (CSL'16)*, pages 39:1–39:21, 2016.
- [54] Ian Pratt-Hartmann, Wiesław Szwaśt, and Lidia Tendera. Quine’s Fluted Fragment Revisited. 2019. To appear in the *Journal of Symbolic Logic*.
- [55] Willard Van Orman Quine. On the Limits of Decision. In *14th International Congress of Philosophy*, volume III, pages 57–62, 1969. An extended version appeared in *W.V. Quine Theories and Things*, Harvard University Press, 1981.
- [56] Willard Van Orman Quine. The Variable. In Willard Van Orman Quine, editor, *The Ways of Paradox and other essays. Revised and enlarged edition*. Harvard University Press, 1976.
- [57] Frank Plumpton Ramsey. On a Problem of Formal Logic. *Proceedings of The London Mathematical Society*, s2-30:264–286, 1930.
- [58] Kurt Schütte. Über die Erfüllbarkeit einer Klasse von logischen Formeln. *Mathematische Annalen*, 110(2):161–194, 1934.

- [59] Kurt Schütte. Untersuchungen zum Entscheidungsproblem der mathematischen Logik. *Mathematische Annalen*, 109(4):572–603, 1934.
- [60] Dana Scott. A decision method for validity of sentences in two variables. *Journal of Symbolic Logic*, 27:477, 1962.
- [61] Luc Segoufin. A survey on guarded negation. *SIGLOG News*, 4(3):12–26, 2017.
- [62] Luc Segoufin and Balder ten Cate. Unary negation. *Logical Methods in Computer Science*, 9(3), 2013.
- [63] Saharon Shelah. Decidability of a portion of the predicate calculus. *Israel Journal on Mathematics*, 28:32–44, 1977.
- [64] Thoralf Skolem. Untersuchungen über die Axiome des Klassenkalküls und über Produktions- und Summationsprobleme welche gewisse Klassen von Aussagen betreffen. *Videnskapsselskapets Skrifter I. Mat.-Nat. Klasse (3)*, 1919.
- [65] Thomas Sturm, Marco Voigt, and Christoph Weidenbach. Deciding First-Order Satisfiability when Universal and Existential Variables are Separated. In *Logic in Computer Science (LICS'16)*, pages 86–95. IEEE/ACM, 2016. An extended version is available at the arXiv preprint server ([arXiv.org](https://arxiv.org)) under the signature arXiv:1511.08999 [cs.LO].
- [66] Jouko A. Väänänen. *Dependence Logic – A New Approach to Independence Friendly Logic*, volume 70 of *London Mathematical Society student texts*. Cambridge University Press, 2007.
- [67] Johan van Benthem. Dynamic bits and pieces. Technical Report LP-97-01, Institute for Logic, Language and Computation (ILLC), University of Amsterdam, 1997.
- [68] Marco Voigt. A Fine-Grained Hierarchy of Hard Problems in the Separated Fragment. In *Logic in Computer Science (LICS'17)*, pages 1–12. IEEE/ACM, 2017. An extended version is available at the arXiv preprint server ([arXiv.org](https://arxiv.org)) under the signature arXiv:1704.02145 [cs.LO].
- [69] Marco Voigt. *Decidable Fragments of First-Order Logic and of First-Order Linear Arithmetic with Uninterpreted Predicates*. PhD thesis, Department of Computer Science, Saarland University, 2019. Available under <https://dx.doi.org/10.22028/D291-28428>.
- [70] Christoph Wernhard. Heinrich Behmann’s Contributions to Second-Order Quantifier Elimination from the View of Computational Logic. Technical report, TU Dresden, 2015.
- [71] Martin Wirsing. *Das Entscheidungsproblem der Prädikatenlogik 1. Stufe mit Identität und Funktionszeichen in Herbrandformeln*. PhD thesis, Ludwig-Maximilians-Universität München, 1976.

A Proof of Theorem 8.6

Proof sketch. Let $n \geq 3$. The following SGNFO sentence is equivalent to the sentence φ given in the proof of Theorem 7.5. In the sentence the sets $\{x_1, \dots, x_n\}$ and $\{y_1, \dots, y_n\}$ and $\{z\}$ are

separated:

$$\begin{aligned}
& \exists z. z = z \\
& \wedge \neg \exists x_n. R_n(x_n) \wedge \neg \exists y_n. R_n(x_n) \wedge T_n(y_n) \\
& \wedge \neg \exists x_{n-1}. R_{n-1}(x_{n-1}, x_n) \wedge T_n(y_n) \wedge \neg \exists y_{n-1}. R_{n-1}(x_{n-1}, x_n) \wedge T_{n-1}(y_{n-1}, y_n) \\
& \dots \\
& \wedge \neg \exists x_1. R_1(x_1, \dots, x_n) \wedge T_2(y_2, \dots, y_n) \wedge \neg \exists y_1. R_1(x_1, \dots, x_n) \wedge T_1(y_1, \dots, y_n) \\
& \wedge \bigwedge_{i=1}^{4n} (R_1(x_1, \dots, x_n) \wedge T_1(x_1, \dots, x_n) \wedge P_i(x_1, \dots, x_n) \wedge Q_i(y_1, \dots, y_n)) \\
& \vee (R_1(x_1, \dots, x_n) \wedge T_1(x_1, \dots, x_n) \wedge \neg (P_i(x_1, \dots, x_n) \vee Q_i(y_1, \dots, y_n))).
\end{aligned}$$

The subformulas $z = z$, $R_n(x_n)$, and the more complex $R_i(x_i, \dots, x_n) \wedge T_i(y_i, \dots, y_n)$ and $R_i(x_i, \dots, x_n) \wedge T_{i+1}(y_{i+1}, \dots, y_n)$ serve as separated negation guards for negated subformulas. We need to introduce a bit of redundancy in order to meet the syntactic requirements of SGNFO. Nevertheless, it is easy to see that the above sentence is equivalent to the sentence φ used in the proof of Theorem 7.5. We will refer to it as φ as well.

The rest of the proof works in analogy to the proof of Theorem 7.5. We therefore take over definitions and notation from that proof, in particular the sets \mathcal{S}_k , the model \mathcal{A} with the domain elements $\mathbf{a}_S^{(j)}$, $\mathbf{b}_S^{(j)}$ and subdomains $\mathbf{A}_a, \mathbf{A}_b, \mathbf{A}_{b,k}$, the notation \mathcal{A}_{-S} , the vocabularies $\Sigma, \Sigma_{PR}, \Sigma_{QT}$ and their extensions $\Sigma'_{PR}, \Sigma'_{QT}$ with nullary predicate symbols, and the notion of column- k -occurrence.

Let φ_{GNFO} be a shortest GNFO sentence that is semantically equivalent to φ . In analogy to the proof of Theorem 7.5, our goal is to show that $\text{len}(\varphi_{\text{GNFO}})$ is at least $(n-1)$ -fold exponential in n . Indeed, we can transform φ_{GNFO} into an equivalent GNFO sentence ψ_{GNFO} with the following properties.

- (a) The sentence ψ_{GNFO} is a Boolean combination of GNFO Σ'_{PR} -sentences and GNFO Σ'_{QT} -sentences. Moreover, none of the constituent sentences of ψ_{GNFO} properly contains a non-atomic subsentence.
- (b) The vocabulary of ψ_{GNFO} is $\Sigma'_{PR} \cup \Sigma'_{QT}$, i.e. Σ plus fresh nullary predicate symbols.
- (c) The structure \mathcal{A} can be uniquely expanded to some model \mathcal{B} of ψ_{GNFO} over the same domain and conserving the interpretations of all predicate symbols in Σ ; for every \mathcal{B}_{-S} — the substructure of \mathcal{B} induced by the domain $\mathbf{B} \setminus \{\mathbf{b}_S^{(k)}\}$ — we have $\mathcal{B}_{-S} \not\models \psi_{\text{GNFO}}$.
- (d) $\text{len}(\psi_{\text{GNFO}}) \in \mathcal{O}(\text{len}(\varphi_{\text{GNFO}}))$.
- (e) Equations in Σ'_{QT} -subsentsences of ψ_{GNFO} occur only in guards, and these consist of a single trivial equation, say $v = v$, and no other atoms.
- (f) Every non-equational Σ_{QT} -atom is linear and for every variable v occurring in it there is some k such that all occurrences of v in non-equational Σ_{QT} -atoms are column- k -occurrences.
- (g) For all distinct variables v, v' that occur freely in a Σ'_{QT} -subformula χ and have column- k -occurrences and column- k' -occurrences, respectively, we know that $k \neq k'$.

The transformation of φ_{GNFO} into ψ_{GNFO} proceeds in analogy to the corresponding transformation in the proof of Theorem 7.5, except that the transformation into negation normal form is postponed. When we next transform ψ_{GNFO} into negation normal form, we obtain the equivalent sentence ψ , which does not necessarily belong GNFO anymore, but still shares the properties that are essential for the rest of the proof.

Now, suppose that ψ has fewer than $2^{\uparrow^{n-1}}(n)$ subformulas. By (c), we have $\mathcal{B} \models \psi$ and $\mathcal{B}_{-S} \not\models \psi$ for every $S \in \mathcal{S}_n$. Hence, for every $S \in \mathcal{S}_n$ there is some Σ'_{QT} -subformula ψ_S in ψ of the form $\exists \bar{y}. \chi_S(\bar{y}, \bar{z})$ and some variable assignment β_S such that the following properties hold.

We have $\beta_S(y_*) = \mathbf{b}_S^{(n)}$ for exactly one $y_* \in \bar{y}$ and for every $v \in \bar{y} \cup \bar{z}$ different from y_* we have $\beta_S(v) \in \mathbf{A}_b \setminus \mathbf{A}_{b,n}$. Moreover, we have

(*) $\mathcal{B}, \beta_S \models \chi_S(\bar{y}, \bar{z})$ and $\mathcal{B}, \beta' \not\models \chi_S(\bar{y}, \bar{z})$ for every β' that differs from β_S only in the value assigned to y_* .

The tuple $\beta_S(\bar{z})$ represents a sequence \bar{c}_S of domain elements from \mathbf{A}_b that can be completed to a chain $\mathbf{b}_{T_1}^{(1)}, \dots, \mathbf{b}_{T_{n-1}}^{(n-1)}, \mathbf{b}_S^{(n)}$ with $T_1 \in \dots \in T_{n-1} \in S$.

Fix any $S_* \in \mathcal{S}_n$ and consider the formula $\psi_{S_*}(\bar{z})$. There is a nonempty set \widehat{S}_* such that $\psi_{S_*}(\bar{z})$ coincides with every $\psi_S(\bar{z})$ with $S \in \widehat{S}_*$. For any distinct $S, S' \in \widehat{S}_*$ the sequences $\bar{c}_S := \beta_S(\bar{z})$ and $\bar{c}_{S'} := \beta_{S'}(\bar{z})$ must differ, for otherwise (*) would be violated. As there are at most $\prod_{k=1}^{n-1} 2^{\uparrow k}(n)$ distinct sequences \bar{c}_S , \widehat{S}_* can contain at most $\prod_{k=1}^{n-1} 2^{\uparrow k}(n) < (2^{\uparrow n-1}(n))^n$ sets. Recall that there are fewer than $2^{\uparrow n-1}(n)$ subformulas in ψ . We have just inferred that each of these can only serve as ψ_S for at most $(2^{\uparrow n-1}(n))^n$ sets $S \in \mathcal{S}_n$. Hence, only

$$(2^{\uparrow n-1}(n))^n \cdot 2^{\uparrow n-1}(n) = 2^{(n+1) \cdot 2^{\uparrow n-2}(n)} < 2^{2^{\uparrow n-1}(n)} = 2^{\uparrow n}(n)$$

sets S have a corresponding subformula ψ_S . But $|\mathcal{S}_n| \geq 2^{\uparrow n}(n+1)$ implies that there are $S \in \mathcal{S}_n$ such that $\mathcal{B}_{-S} \models \psi_{\text{LGF}}$, which contradicts our assumptions. Consequently, there must be more than $2^{\uparrow n-1}(n)$ subformulas in ψ and in ψ_{GNFO} and, by (d), also in φ_{GNFO} . \square