



A fine-grained hierarchy of hard problems in the separated fragment

Marco Voigt

► To cite this version:

Marco Voigt. A fine-grained hierarchy of hard problems in the separated fragment. LICS 2017 - 32nd Annual ACM/IEEE Symposium on Logic in Computer Science, Jun 2017, Reykjavik, Iceland. pp.1 - 12, 10.1109/LICS.2017.8005094 . hal-01592172

HAL Id: hal-01592172

<https://inria.hal.science/hal-01592172>

Submitted on 27 Sep 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Fine-Grained Hierarchy of Hard Problems in the Separated Fragment

Marco Voigt

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

Abstract—Recently, the separated fragment (SF) has been introduced and proved to be decidable. Its defining principle is that universally and existentially quantified variables may not occur together in atoms. The known upper bound on the time required to decide SF’s satisfiability problem is formulated in terms of quantifier alternations: Given an SF sentence $\exists z \forall x_1 \exists y_1 \dots \forall x_n \exists y_n. \psi$ in which ψ is quantifier free, satisfiability can be decided in non-deterministic n -fold exponential time. In the present paper, we conduct a more fine-grained analysis of the complexity of SF-satisfiability. We derive an upper and a lower bound in terms of the *degree ∂ of interaction of existential variables* (short: *degree*)—a novel measure of how many separate existential quantifier blocks in a sentence are connected via joint occurrences of variables in atoms. Our main result is the k -NEXPTIME-completeness of the satisfiability problem for the set $SF_{\partial \leq k}$ of all SF sentences that have degree k or smaller. Consequently, we show that SF-satisfiability is non-elementary in general, since SF is defined without restrictions on the degree. Beyond trivial lower bounds, nothing has been known about the hardness of SF-satisfiability so far.

1. Introduction

In [17] the separated fragment (SF) of first-order logic with equality is introduced. Its defining principle is that universally and existentially quantified variables may not occur together in atoms. (Topmost existential quantifier blocks are exempt from this rule.) SF properly generalizes both the Bernays–Schönfinkel–Ramsey (BSR) fragment ($\exists^* \forall^*$ -sentences with equality) and the relational monadic fragment without equality (MFO). Still, the satisfiability problem for SF is decidable.

In computational logic formulas are often classified based on the shape of quantifier prefixes. There is a wealth of results that separate decidable first-order formulas from undecidable ones in this

fashion, see [3] for references. The definition of the BSR fragment is only one example. In the context of computational complexity, hierarchies are defined, such as the polynomial hierarchy, where the hardness of problems is assumed to grow with the number of quantifier alternations that are allowed to occur.

Although the definition of SF breaks with the paradigm of restricting quantifier prefixes, the known upper bound on the complexity of SF-satisfiability is again based on quantifier prefixes: Deciding whether an SF sentence $\varphi := \exists z \forall x_1 \exists y_1 \dots \forall x_n \exists y_n. \psi$ with quantifier-free ψ is satisfiable requires a nondeterministic computing time that is at most n -fold exponential in the length of φ (cf. Theorem 17 in [17]). On the one hand, we complement this result with a corresponding lower bound in the present paper. That is, we show that SF-satisfiability is indeed non-elementary. On the other hand, we derive a refined upper bound that is based on the *degree ∂ of interaction of existential variables*. An overview of the resulting hierarchy of complete problems is depicted in Figure 1. Intuitively, φ exhibits a degree $\partial_\varphi = k$, if variables from k distinct existential quantifier blocks interact. We say that two variables x, y interact, if they occur together in at least one atom or if there is a third variable z that interacts with both x and y (i.e. the property is transitive). For instance, in the SF sentence $\forall x_1 \exists y_1 \forall x_2 \exists y_2 \forall x_3 \exists y_3 v_3. (P(x_1, x_2, x_3) \wedge \neg Q(y_1, y_3)) \vee P(y_2, v_2, v_3) \vee \neg Q(y_3, v_1)$ the sets $\{y_1, y_3, v_1\}$ and $\{y_2, v_2, v_3\}$ form the maximal sets of interacting existential variables. Since each of these sets contains variables from at most two distinct quantifier blocks, the formula exhibits a degree $\partial = 2$.

In Section 4.1, and in particular in Theorem 13, we observe that the satisfiability problem for $SF_{\partial \leq k}$ —the set of all SF sentences φ with $\partial_\varphi \leq k$ —lies in k -NEXPTIME. It is worth mentioning that this result adequately accounts for the known complexity of MFO-satisfiability. For

every MFO sentence φ we trivially have $\partial_\varphi = 1$, since all occurring predicate symbols have an arity of at most one. Theorem 13 entails that MFO-satisfiability is in NEXPTIME, which is well known. Still, this bound is not reproducible with the analysis of the complexity of SF-satisfiability conducted in [17].

Apparently, non-elementary satisfiability problems are not very widespread among the decidable fragments of classical first-order logic known today. We show in Section 5 that SF falls into this category. To the present author's knowledge, the only known companion in this respect is the fluted fragment (FL). Indeed, Pratt-Hartmann, Szust, and Tendera show in [14] that satisfiability of fluted sentences with at most $2k$ variables is k -NEXPTIME-hard. Moreover, they argue that satisfiability of fluted sentences with at most k variables lies in k -NEXPTIME. Although a significant gap between these lower and upper bounds remains to be closed, the fluted fragment seems to comprise a similar hierarchy of hard problems as SF does.

Another characteristic of SF is that it enjoys a *small model property*. More precisely, given an SF sentence φ , one can compute a positive integer n that depends on the degree ∂_φ and the length of φ such that, if there is a model of φ at all, then there also is a model whose domain contains at most n elements. Many first-order fragments are known to enjoy a small model property. The BSR fragment and MFO are among the classical ones (see [3] for references). More recently defined fragments include the two-variable fragment (FO₂)

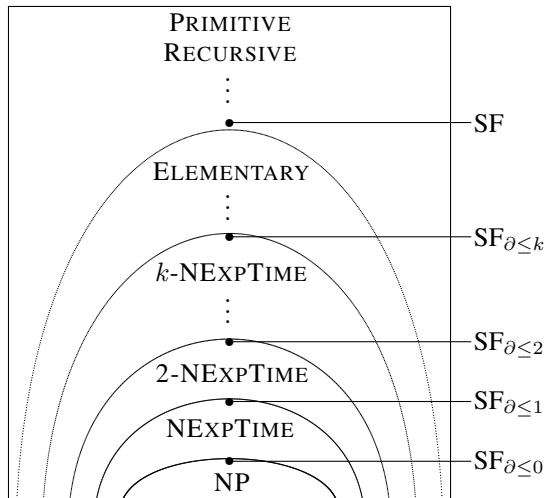


Figure 1. The subfragments of SF scale over the major nondeterministic complexity classes in ELEMENTARY, while SF itself goes even beyond.

[12], [8], the fluted fragment (FL) [15], [16], [14], the guarded fragment (GF) [1], [7], the guarded negation fragment (GNF) [2], and the uniform one-dimensional fragment (UF₁) [11]. While GNF and UF₁ are incomparable, GNF extends GF, and UF₁ can be considered as a generalization of FO₂. Guarded fragments and the two-variable fragment have received quite some attention due to the fact that modal logics have natural translations into them. As a continuation of that theme, we shall see in Section 3.1 how classes of sentences enjoying a small model property can be effectively translated into subclasses of SF. During the translation process the length of formulas increases by a factor that is logarithmic in the size of small models of the original. One benefit of translating non-SF sentences into SF sentences is that in SF one can natively express concepts such as transitivity and basic counting quantifiers (Proposition 5). This is not always possible in other fragments enjoying a small model property. For example, transitivity cannot be expressed in FO₂, GF, and FL.

Summing up, the main contributions are:

- (i) Based on the novel concept of the degree of interaction of existential variables, we substantially refine the existing analysis of the time required to decide SF-satisfiability. More concretely, we show that a satisfiable SF sentence φ with $\partial_\varphi = k$ has a model whose domain is of a size that is at most k -fold exponential in the length of φ (Section 4, Theorem 13). With this refined approach we can close the complexity gap for the class of *strongly separated* sentences (Corollary 15) that was left open in [17]. Moreover, the complexity of MFO can be explained in the refined framework.
- (ii) We complement the complexity analysis with corresponding lower bounds in two respects. We first derive a lower bound on the length of shortest BSR sentences that are equivalent to a given SF sentence (Section 4.2, Theorem 16). In Section 5, we prove k -NEXPTIME-hardness of satisfiability for the class of SF sentences φ with $\partial_\varphi = k$ (Theorem 21). Since SF is in general defined without restrictions on the degree ∂_φ , our result implies that SF-satisfiability is non-elementary.
- (iii) We devise a simple translation from classes of first-order sentences that enjoy a small model property into SF (Proposition 4). Moreover, we argue that SF can express basic counting quantifiers (Proposition 5).

Due to space limitations, we mostly resort to sketches of proofs. The interested reader is referred to the extended version of the present paper [18].

2. Preliminaries

We mainly reuse the basic notions from [17]. We repeat the definition of necessary concepts and notation for the sake of completeness.

We consider first-order logic formulas. The underlying signature shall not be mentioned explicitly, but will become clear from the current context. For the distinguished *equality* predicate we use \approx . We follow the convention that negation binds strongest, that conjunction binds stronger than disjunction, and that all of the aforementioned bind stronger than implication. The scope of quantifiers shall stretch as far to the right as possible. By $\text{len}(\cdot)$ we denote a reasonable measure of the length of formulas satisfying $\text{len}(\varphi \rightarrow \psi) = \text{len}(\neg\varphi \vee \psi)$ and $\text{len}(\varphi \leftrightarrow \psi) = \text{len}((\neg\varphi \vee \psi) \wedge (\varphi \vee \neg\psi))$.

We write $\varphi(x_1, \dots, x_m)$ to denote a formula φ whose free variables form a subset of $\{x_1, \dots, x_m\}$. In all formulas we tacitly assume that no variable occurs freely and bound at the same time and that no variable is bound by two different occurrences of quantifiers, unless explicitly stated otherwise. For convenience, we sometimes identify tuples \vec{x} of variables with the set containing all the variables that occur in \vec{x} . We write $\text{vars}(\varphi)$ to address the set of all variable symbols that occur in φ . Similarly, $\text{consts}(\varphi)$ denotes the set of all constant symbols in φ . We denote *substitution* by $\varphi[x/t]$ if every free occurrence of x in φ is to be substituted with the term t .

A *literal* is an atom or a negated atom, and a *clause* is a disjunction of literals. We say that a formula is in *conjunctive normal form* (CNF), if it is a conjunction of clauses, possibly preceded by a quantifier prefix. A formula in CNF is *Horn* if every clause contains at most one non-negated literal. It is *Krom* if every clause contains at most two literals at all.

A sentence $\varphi := \forall \vec{x}_1 \exists \vec{y}_1 \dots \forall \vec{x}_n \exists \vec{y}_n. \psi$ is in *standard form*, if it is in *negation normal form* (i.e. every negation symbol occurs directly before an atom) and ψ is quantifier free, contains exclusively the Boolean connectives \wedge, \vee, \neg , and does not contain non-constant function symbols. The tuples \vec{x}_1 and \vec{y}_n may be empty, i.e. the quantifier prefix does not have to start with a universal quantifier, and it does not have to end with an existential quantifier. Moreover, we require that every variable occurring in the quantifier prefix does also occur in ψ .

As usual, we interpret a formula φ with respect to given structures. A *structure* \mathcal{A} consists of a

nonempty *universe* $\mathcal{U}_{\mathcal{A}}$ (also: *domain*) and interpretations $f^{\mathcal{A}}$ and $P^{\mathcal{A}}$ of all considered function and predicate symbols, respectively, in the usual way. Given a formula φ , a structure \mathcal{A} , and a variable assignment β , we write $\mathcal{A}, \beta \models \varphi$ if φ evaluates to *true* under \mathcal{A} and β . We write $\mathcal{A} \models \varphi$ if $\mathcal{A}, \beta \models \varphi$ holds for every β . The symbol \models denotes (*semantic*) *equivalence* of formulas, i.e. $\varphi \models \psi$ holds whenever for every structure \mathcal{A} and every variable assignment β we have $\mathcal{A}, \beta \models \varphi$ if and only if $\mathcal{A}, \beta \models \psi$. We call two sentences φ and ψ *equisatisfiable* if φ has a model if and only if ψ has one.

A structure \mathcal{A} is a *substructure* of a structure \mathcal{B} (over the same signature) if (1) $\mathcal{U}_{\mathcal{A}} \subseteq \mathcal{U}_{\mathcal{B}}$, (2) $c^{\mathcal{A}} = c^{\mathcal{B}}$ for every constant symbol c , (3) $P^{\mathcal{A}} = P^{\mathcal{B}} \cap \mathcal{U}_{\mathcal{A}}^m$ for every m -ary predicate symbol P , and (4) $f^{\mathcal{A}}(\mathbf{a}_1, \dots, \mathbf{a}_m) = f^{\mathcal{B}}(\mathbf{a}_1, \dots, \mathbf{a}_m)$ for every m -ary function symbol f and every m -tuple $\langle \mathbf{a}_1, \dots, \mathbf{a}_m \rangle \in \mathcal{U}_{\mathcal{A}}^m$. The following is a standard lemma, see, e.g., [5] for a proof.

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

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

$$\begin{aligned} \exists x.(\varphi \vee \psi) &\models (\exists x_1.\varphi) \vee (\exists x_2.\psi), \\ \exists x.(\varphi \circ \chi) &\models (\exists x.\varphi) \circ \chi, \\ \forall x.(\varphi \wedge \psi) &\models (\forall x_1.\varphi) \wedge (\forall x_2.\psi), \\ \forall x.(\varphi \circ \chi) &\models (\forall x.\varphi) \circ \chi, \end{aligned}$$

where $\circ \in \{\wedge, \vee\}$.

We use the notation $[k]$ to abbreviate the set $\{1, \dots, k\}$ for any positive integer k . Moreover, \mathfrak{P} shall be used as the power set operator, i.e. $\mathfrak{P}S$ denotes the set of all subsets of a given set S . Finally, we need some notation for the *tetration operation*. We define $2^{\uparrow k}(m)$ inductively: $2^{\uparrow 0}(m) := m$ and $2^{\uparrow k+1}(m) := 2^{(2^{\uparrow k}(m))}$.

3. The separated fragment

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

Definition 3 (Separated fragment (SF), [17]). The *separated fragment* (SF) of first-order logic consists of all existential closures of prenex formulas without non-constant function symbols in which existentially quantified variables are separated from

universally quantified ones. More precisely, SF consists of all first-order sentences with equality but without non-constant function symbols that are of the form $\exists \vec{z} \forall \vec{x}_1 \exists \vec{y}_1 \dots \forall \vec{x}_n \exists \vec{y}_n. \psi$, in which ψ is quantifier-free, and in which the sets $\vec{x}_1 \cup \dots \cup \vec{x}_n$ and $\vec{y}_1 \cup \dots \cup \vec{y}_n$ are separated.

The tuples \vec{z} and \vec{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.

Notice that the variables in \vec{z} are not subject to any restriction concerning their occurrences.

In [17] the authors show that the satisfiability problem for SF sentences (*SF-satisfiability*) is decidable. Before we start investigating the complexity issues related to SF-satisfiability, we elaborate on the expressiveness of SF.

3.1. Expressiveness

Every SF sentence is equivalent to a BSR sentence ([17], Lemma 6). We shall outline in Section 4 how to analyze the blow-up that we have to incur during this translation process and how it depends on the degree of interaction of existential variables. Since the BSR fragment enjoys a small model property (cf. Proposition 6), SF inherits the small model property from BSR. However, regarding the size of minimal models of satisfiable formulas, SF sentences are much more compact. While satisfiable BSR sentences have models whose domain is linear in the length of the formula, satisfiable SF sentences can enforce domains of a size that cannot be bounded from above by a finite tower of exponentials. We provide first evidence for this fact in Theorem 16, where we give a non-elementary lower bound on the length of equivalent BSR sentences. This lower bound even applies to the SF-Horn-Krom subfragment of SF. Moreover, we exploit the capability of SF sentences φ to enforce models of ∂_φ -fold exponential size in the proof of the k -NEXPTIME-hardness of SF-satisfiability (for every $k \geq 1$).

Apart from compactness of representation, and from the perspective of satisfiability, all first-order fragments that enjoy small model properties share a common ground of expressiveness. Neglecting efficiency, every sentence φ from such a fragment can be effectively translated into a (finite) propositional formula ϕ in such a way that from a satisfying variable assignment for ϕ one can straightforwardly reconstruct a (Herbrand) model of φ . The reason is simply that universal quantification can then be understood as finite conjunction (over a finite domain)

and existential quantification can be conceived as finite disjunction.

The following proposition illustrates why SF is to some extent prototypical for first-order fragments that enjoy a small model property.

Proposition 4. Consider any nonempty class \mathcal{C} of first-order formulas without non-constant function symbols for which we know a computable mapping $f : \mathcal{C} \rightarrow \mathbb{N}$ such that every satisfiable φ in \mathcal{C} has a model of size at most $f(\varphi)$. Then there exists an effective translation T from \mathcal{C} into SF such that for every $\varphi \in \mathcal{C}$, (a) every model of $T(\varphi)$ is also a model of φ , (b) every model of φ whose size is at most $f(\varphi)$ can be extended to a model of $T(\varphi)$ over the same domain, and (c) the length of $T(\varphi)$ lies in $\mathcal{O}(\text{len}(\varphi) \cdot \log f(\varphi) \cdot \log \log f(\varphi))$.

Proof. We outline the translation T for some given input sentence φ , which we assume to be in negation normal form (without loss of generality). Let $m := \lceil \log_2 f(\varphi) \rceil$ and let Q_1, \dots, Q_m be unary predicate symbols that do not occur in φ . For all terms s, t we define $s \approx t$ as abbreviation of $\bigwedge_{i=1}^m Q_i(s) \leftrightarrow Q_i(t)$. In order to restrict the domain to 2^m elements, we conjoin the formula $\chi_{\text{fin}} := \forall xy. x \approx y \rightarrow x \approx y$. Since in any structure \mathcal{A} there are at most 2^m domain elements that can be distinguished by their membership in the sets $Q_1^{\mathcal{A}}, \dots, Q_m^{\mathcal{A}}$, it is clear that $\mathcal{A} \models \chi_{\text{fin}}$ entails $|\mathcal{U}_{\mathcal{A}}| \leq 2^m$.

(*) Let \mathcal{A} be any structure, let β be any variable assignment over \mathcal{A} 's domain, and let s, t be two terms. If $\mathcal{A} \models \chi_{\text{fin}}$ holds, then we get $\mathcal{A}, \beta \models s \approx t$ if and only if $\mathcal{A}, \beta \models s \approx t$.

This means, if we restrict our attention to domains with at most 2^m domain elements, we can now use a separated form of equality.

(**) Let ψ be any first-order formula and let v be some variable that does not occur in ψ . Then ψ is equivalent to $\forall v. u \approx v \rightarrow \psi[u/v]$.

We can transform φ into an equivalent sentence φ' by consecutively replacing each subformula of the form $\exists y. \psi$ in φ with $\exists y \forall v. y \approx v \rightarrow \psi[y/v]$, where we assume v to be fresh (one fresh variable for every replaced subformula). Consequently, every atom in φ' that is not an equation contains exclusively universally quantified variables. Moreover, (**) implies that φ and φ' are equivalent.

Let φ'' be the result of replacing all equations $y \approx v, v \approx y$ in φ' in which y is existentially quantified and v universally quantified with the formula $y \approx v$. We then set $\varphi_{\text{SF}} := \chi_{\text{fin}} \wedge \varphi''$. By (*), any model of φ_{SF} is also a model of φ .

Conversely, any model \mathcal{A} of φ that has at most 2^m domain elements can be converted into a model \mathcal{B} of φ_{SF} by defining the relations $Q_1^{\mathcal{B}}, \dots, Q_m^{\mathcal{B}}$ in an appropriate way. \square

The φ_{SF} in the above proof belongs to a subfragment of SF that we call *strongly separated* (cf. Definition 14) and whose satisfiability problem is complete for NEXPTIME (cf. Corollary 15).

Unfortunately, the translation methodology of Proposition 4 does not help in the quest for new decidable first-order fragments. The reason is simply that we already need arguments leading to a small model property before we can start the translation process, as we need information about the size of the models that have to be considered. Nevertheless, such translations can be useful in view of the expressiveness of SF that other first-order fragments, such as FO₂, the fluted fragment, and GF, lack. For instance, SF sentences can naturally express the axioms of equivalence, most prominently, transitivity. Hence, fundamental and interesting properties of predicates that have to be assumed at the meta-level when dealing with less expressive logics can be formalized directly in SF. Moreover, basic counting quantifiers can be defined natively in SF and do not have to be introduced via special operators. More precisely, given any formula $\exists^{\geq n} y. \varphi$ with positive n and without non-constant function symbols, its standard translation $\exists y_1 \dots y_n. \bigwedge_{i=1}^n \varphi[y/y_i] \wedge \bigwedge_{i < j} y_i \not\approx y_j$ is not in conflict with the separateness conditions of SF's definition, if the set $\{y_1, \dots, y_n\}$ is separated in φ from the set of universally quantified variables.

Proposition 5. Counting quantifiers $\exists^{\geq n}$ with positive integer n are expressible in SF.

3.2. Basic complexity considerations

We first recall the well-known small model properties of SF's subfragments BSR and MFO (see [3] for references).

Proposition 6. Let $\varphi := \exists \vec{z} \forall \vec{x}. \psi$ be a satisfiable BSR sentence, i.e. ψ is quantifier free and does not contain non-constant function symbols. There is a model $\mathcal{A} \models \varphi$ such that $|\mathcal{U}_{\mathcal{A}}| \leq \max(|\vec{z}| + |\text{consts}(\varphi)|, 1)$.

We make use of this property when we derive an upper bound on the size of small models for satisfiable SF sentences, as our approach will be based on an effective translation of SF sentences into equivalent BSR sentences.

Proposition 7. Let $\varphi := \exists \vec{z} \forall \vec{x}_1 \exists \vec{y}_1 \dots \forall \vec{x}_n \exists \vec{y}_n. \psi$ be a satisfiable monadic sentence without equality and without non-constant function symbols, i.e. all predicate symbols in φ are of arity 1. Moreover, assume that φ contains k distinct predicate symbols. There is a model $\mathcal{A} \models \varphi$ such that $|\mathcal{U}_{\mathcal{A}}| \leq 2^k$.

Notice that the shape of the quantifier prefix does not contribute to the upper bound.

The following lemma links bounds on the size of models with the computing time that is required to decide satisfiability.

Lemma 8 (cf. [3], Proposition 6.0.4). Let φ be a first-order sentence in prenex normal form containing n universally quantified variables. The question whether φ has a model of cardinality m can be decided nondeterministically in time $p(m^n \cdot \text{len}(\varphi))$ for some polynomial p .

With this lemma at hand, it is enough to prove a small model property for a given class of first-order sentences, in order to bound the worst-case time complexity of the corresponding satisfiability problem from above. This is exactly what the authors of [17] have done for SF.

Proposition 9 ([17], Theorem 17). Let $\varphi := \exists \vec{z} \forall \vec{x}_1 \exists \vec{y}_1 \dots \forall \vec{x}_n \exists \vec{y}_n. \psi$ be an SF sentence for some quantifier-free ψ . There is some equivalent BSR sentence $\exists \vec{u} \forall \vec{v}. \psi'$ in which the number of occurring constant symbols plus the number of existential quantifiers is at most $\text{len}(\varphi) + n \cdot \text{len}(\varphi) \cdot (2^{\uparrow n(\text{len}(\varphi))})^n$. As a result, satisfiability of φ can be decided nondeterministically in time that is at most n -fold exponential in $\text{len}(\varphi)$.

Clearly, applying this result to an MFO sentence substantially overshoots the actual worst-case time requirements. To stress it again, the notion of the degree of interaction is a remedy to this sort of inaccuracies, as we shall see in Section 4.

A special case that is worth considering, before we investigate the complexity of full SF, is the class of SF sentences that do not contain universal quantifiers. This species of formulas coincides with the purely existential fragment of first-order logic without non-constant function symbols, and it is a close relative of propositional logic. Recall that SAT is NP-complete [4], Horn-SAT is P-complete [10], [13], and 2SAT is NL-complete [9].

Proposition 10.

- (i) Satisfiability for the class of SF sentences without universal quantifiers is NP-complete.

- (ii) Satisfiability for the class of SF-Horn sentences without universal quantifiers is P-complete.
- (iii) Satisfiability for the class of SF-Krom sentences without universal quantifiers and without equality is NL-complete.

4. Translation of SF sentences into BSR sentences

In this section, we analyze the transformation process from SF into the BSR fragment from the perspective of the *degree of interaction of existential variables*. Our aim is to derive upper and lower bounds on the length of the resulting BSR-formulas. Roughly speaking, in the first phase of the translation process all quantifiers are moved inwards as far as possible (cf. the proof of Lemma 12). In order to do so, we first transform the sentence in question into a formula in CNF. After that, we employ the well-known rules of miniscoping (cf. Lemma 2), supplemented by the rule formulated in the following lemma.

Lemma 11. Let I and K_i , $i \in I$, be sets that are finite, nonempty, and pairwise disjoint. The elements of these sets serve as indices. Let

$$\varphi := \exists \vec{y}. \bigwedge_{i \in I} \left(\chi_i(\vec{z}) \vee \bigvee_{k \in K_i} \eta_k(\vec{y}, \vec{z}) \right)$$

be some first-order formula where the χ_i and the η_k denote arbitrary subformulas that we treat as indivisible units in what follows. We say that $f : I \rightarrow \bigcup_{i \in I} K_i$ is a *selection function* if for every $i \in I$ we have $f(i) \in K_i$. We denote the set of all selection functions of this form by \mathcal{F} .

Then φ is equivalent to $\varphi' :=$

$$\bigwedge_{\substack{S \subseteq I \\ S \neq \emptyset}} \left(\bigvee_{i \in S} \chi_i(\vec{z}) \right) \vee \bigvee_{f \in \mathcal{F}} \left(\exists \vec{y}. \bigwedge_{i \in S} \eta_{f(i)}(\vec{y}, \vec{z}) \right).$$

Proof sketch. The proof of this lemma follows a conceptually simple strategy. Using the distributivity of \wedge over \vee , we first transform φ into a disjunction of conjunctions of the indivisible units $\chi_i(\vec{z})$ and $\eta_k(\vec{y}, \vec{z})$. Then, exploiting the equivalences in Lemma 2, we push the existential quantifier block $\exists \vec{y}$ inwards such that it only binds conjunctions of units $\eta_k(\vec{y}, \vec{z})$. This is possible, because none of the variables in \vec{y} occurs in any of the $\chi_i(\vec{z})$. From this point on, we treat the newly emerged subformulas $\exists \vec{y}. \bigwedge_{k'} \eta_{k'}(\vec{y}, \vec{z})$ as if they were indivisible. We then transform the formula back into a conjunction of disjunctions of

indivisible units, this time using the distributivity of \vee over \wedge . It then remains to show that the result of this transformation exhibits a highly redundant structure and is actually equivalent to φ' . \square

4.1. Degree of interaction of existential variables and the size of small models

Consider the formula $\varphi := \exists \vec{z} \forall \vec{x}_1 \exists \vec{y}_1 \dots \forall \vec{x}_n \exists \vec{y}_n. \psi$ in standard form in which ψ is quantifier free and in which the sets $\vec{x} := \vec{x}_1 \cup \dots \cup \vec{x}_n$ and $\vec{y} := \vec{y}_1 \cup \dots \cup \vec{y}_n$ are separated. In addition, we assume that \vec{x}_1 and \vec{y}_1 are nonempty. The tuple \vec{z} , on the other hand, may be empty.

For any $j \in [n]$ and any variable $y \in \vec{y}_j$ we say that y is a *level- j variable*, denoted $\text{lvl}(y) = j$. For any nonempty set $Y \subseteq \vec{y}$ of existentially quantified variables and any positive integer k we say that Y has *degree k in φ* , denoted $\partial_{Y, \varphi} = k$, if k is the maximal number of distinct variables $y_1, \dots, y_k \in Y$ that belong to different levels in φ , i.e. $\text{lvl}(y_1) < \dots < \text{lvl}(y_k)$. We say that φ 's *degree of interaction of existential variables* (short: *degree*) is k , denoted $\partial_\varphi = k$, if k is the smallest positive integer such that we can partition \vec{y} into $m > 0$ parts Y_1, \dots, Y_m that are all pairwise separated in φ and for which $k = \max\{k_j \mid \partial_{Y_j, \varphi} = k_j, 1 \leq j \leq m\}$. Sentences $\varphi := \exists \vec{z} \forall \vec{x}. \psi$ in standard form with quantifier-free ψ are said to have *degree zero*, i.e. $\partial_\varphi = 0$, if \vec{x} is empty and we define $\partial_\varphi = 1$ if \vec{x} is nonempty.

Lemma 12. Let $\varphi := \exists \vec{z} \forall \vec{x}_1 \exists \vec{y}_1 \dots \forall \vec{x}_n \exists \vec{y}_n. \psi$ be an SF sentence of positive degree ∂_φ in standard form. Let $\mathcal{L}_\varphi(\vec{y})$ denote the set of all literals in φ that contain at least one variable $y \in \vec{y} := \vec{y}_1 \cup \dots \cup \vec{y}_n$. There exists a sentence $\varphi_{\text{BSR}} = \exists \vec{z} \exists \vec{u} \forall \vec{v}. \psi_{\text{BSR}}$ in standard form with quantifier-free ψ_{BSR} that is equivalent to φ and contains at most $|\vec{z}| + |\vec{y}| \cdot \partial_\varphi \cdot (2^{\uparrow \partial_\varphi}(|\mathcal{L}_\varphi(\vec{y})|))^{|\partial_\varphi|}$ leading existential quantifiers.

Proof sketch. Let $\vec{x} := \vec{x}_1 \cup \dots \cup \vec{x}_n$. We transform φ into CNF and then apply Lemma 11 and the rules of miniscoping given in Lemma 2 to push all quantifier blocks inwards. Since the sets \vec{x} and \vec{y} are separated in φ , these operations can be performed in such a way that in the resulting formula φ' no universal quantifier lies within the scope of any existential quantifier (other than the ones in $\exists \vec{z}$) and vice versa. After removing redundant parts from φ' , the depth of nestings of existential quantifier blocks (interspersed with conjunctive connectives in φ' 's

syntax tree) can be upper bounded by ∂_φ . As a consequence, φ' contains at most $2^{\uparrow \partial_\varphi}(|\mathcal{L}_\varphi(\vec{y})|)$ distinct subformulas that are of the form $\exists y.\psi'$ and do not lie within the scope of any quantifier. After further transformations, we obtain a formula $\varphi'' := \bigvee_k (\chi_k(\vec{x}) \wedge \bigwedge_{r_k} \eta_{r_k}(\vec{y}))$ where the r_k range over at most $2^{\uparrow \partial_\varphi}(|\mathcal{L}_\varphi(\vec{y})|)$ indices. Moreover, every constituent $\bigwedge_{r_k} \eta_{r_k}$ in φ'' contains at most $|\vec{y}| \cdot \sum_{k'=1}^{\partial_\varphi} \prod_{d=k'}^{\partial_\varphi} 2^{\uparrow d}(|\mathcal{L}_\varphi(\vec{y})|)$ occurrences of existential quantifiers. Since these existential quantifiers distribute over the topmost disjunction when we move them outwards to the front of the sentence φ'' , and since the universal quantifiers in the χ_k may also be moved back outwards, one can show that φ is equivalent to some BSR sentence with at most $|\vec{y}| \cdot \partial_\varphi \cdot (2^{\uparrow \partial_\varphi}(|\mathcal{L}_\varphi(\vec{y})|))^{\partial_\varphi}$ leading existential quantifiers. \square

Proposition 6 now entails that any satisfiable SF-sentence φ has a model of size at most

$$\text{len}(\varphi) + \text{len}(\varphi) \cdot \partial_\varphi \cdot (2^{\uparrow \partial_\varphi}(\text{len}(\varphi)))^{\partial_\varphi}. \quad (1)$$

Theorem 13. Let k be any positive integer. The satisfiability problem for the class of SF sentences φ in standard form with degree $\partial_\varphi \leq k$ can be decided in nondeterministic k -fold exponential time.

Together with Proposition 10(i), this establishes the upper bounds depicted in Figure 1.

In cases where $\partial_\varphi = 1$, Expression (1) simplifies to $\text{len}(\varphi) + \text{len}(\varphi) \cdot 2^{\text{len}(\varphi)}$. The syntactic class of sentences satisfying this property is called *strongly separated* in [17].

Definition 14 ([17]). Let $\varphi := \forall \vec{x}_1 \exists \vec{y}_1 \dots \forall \vec{x}_n \exists \vec{y}_n. \psi$ be an SF sentence and assume that ψ is quantifier free. We say that φ belongs to the *strongly separated fragment (SSF)* if and only if the sets $\vec{x} := \vec{x}_1 \cup \dots \cup \vec{x}_n$ and $\vec{y}_1, \dots, \vec{y}_n$ are all pairwise separated in φ .

Since MFO and BSR sentences fall into this syntactic category, and since their decision problem is known to be NEXPTIME-hard, we obtain the following corollary.

Corollary 15. The satisfiability problem for SSF is NEXPTIME-complete.

Notice that the presented method can explain the asymptotic complexity of MFO-satisfiability and yields a reasonable upper bound on the size of small models of satisfiable MFO sentences. This works in spite of the fact that monadic sentences may contain arbitrarily nested alternating

quantifiers. This is a considerable improvement compared to the methods used in [17].

Let φ be any SF sentence with the maximally possible degree $\partial_\varphi = n$, where n is the number of occurring $\forall\exists$ -alternations. Then the upper bound shown in Expression (1) regarding the number of elements in small models fits the corresponding result entailed by Proposition 9. As one consequence, Theorem 13 in the present paper subsumes Theorem 17 in [17]. Moreover, Corollary 15 improves the double exponential upper bound on SSF-satisfiability given in Theorem 15 in [17]. Finally, it is worth noticing that all SF sentences with the quantifier prefix $\exists^* \forall^* \exists^* \forall^*$ belong to the strongly separated fragment. Hence, Corollary 15 subsumes Theorem 14 in [17]. The latter stipulates NEXPTIME-completeness of SF sentences with quantifier prefix $\exists^* \forall^* \exists^*$. Clearly, the refined analysis based on the degree of interaction of existential variables, rather than the number of quantifier alternations, yields significantly tighter results in many cases.

4.2. Lower bounds on the length of equivalent BSR formulas

Before we derive lower bounds on the time that is required to decide SF-satisfiability in the worst case, we establish lower bounds on the length of the results of the translation from SF into the BSR fragment.

Theorem 16. There is a class of SF sentences that are Horn and Krom such that for every positive integer n the class contains a sentence φ of degree $\partial_\varphi = n$ and with a length linear in n for which any equivalent BSR sentence contains at least $\sum_{k=1}^n 2^{\uparrow k}(n)$ leading existential quantifiers.

Proof sketch. Recall that $[n]$ abbreviates the set $\{1, \dots, n\}$ and that $\mathfrak{P}S$ denotes the power set of a given set S . 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: $\varphi := \forall x_n \exists y_n \dots \forall x_1 \exists y_1. \bigwedge_{i=1}^{4n} (P_i(x_1, \dots, x_n) \leftrightarrow Q_i(y_1, \dots, y_n))$. Notice that we change the orientation of the indices in the quantifier prefix in this proof.

In order to construct a particular model of φ , we inductively define the following sets: $\mathcal{S}_1 := \{S \subseteq [4n] \mid |S| = 2n\}$, $\mathcal{S}_{k+1} := \{S \in \mathfrak{P}\mathcal{S}_k \mid |S| = \frac{1}{2} \cdot |\mathcal{S}_k|\}$ for every k , $1 < k \leq n$. Hence, we observe $|\mathcal{S}_1| = \binom{4n}{2n} \geq \left(\frac{4n}{2n}\right)^{2n} = 2^{2n}$, $|\mathcal{S}_2| = \binom{|\mathcal{S}_1|}{|\mathcal{S}_1|/2} \geq \left(\frac{|\mathcal{S}_1|}{|\mathcal{S}_1|/2}\right)^{|\mathcal{S}_1|/2} \geq 2^{2^{2n-1}}$,

$$|S_n| = \binom{|S_{n-1}|}{\lfloor |S_{n-1}|/2 \rfloor} \geq 2^{2^{\lfloor |S_{n-1}|/2 \rfloor - 1}} \geq 2^{\uparrow n(n+1)}.$$

We now define the structure \mathcal{A} as follows:

- $\mathcal{U}_{\mathcal{A}} := \bigcup_{k=1}^n \{a_S^{(k)}, b_S^{(k)} \mid S \in S_k\},$
- $P_i^{\mathcal{A}} := \{ \langle a_{S_1}^{(1)}, \dots, a_{S_n}^{(n)} \rangle \in \mathcal{U}_{\mathcal{A}}^n \mid i \in S_1 \in S_2 \in \dots \in S_n \}$ for $i = 1, \dots, 4n$, and
- $Q_i^{\mathcal{A}} := \{ \langle b_{S_1}^{(1)}, \dots, b_{S_n}^{(n)} \rangle \in \mathcal{U}_{\mathcal{A}}^n \mid i \in S_1 \in S_2 \in \dots \in S_n \}$ for $i = 1, \dots, 4n$.

One can easily show that \mathcal{A} is a model of φ . Moreover, employing a game-theoretic argument, one can show the following property:

- (*) the substructure induced by \mathcal{A} 's domain after removing at least one of the $b_S^{(k)}$ does not satisfy φ .

We know that $\mathcal{U}_{\mathcal{A}}$ contains at least $\sum_{k=1}^n 2^{\uparrow k(n)}$ elements of the form $b_S^{(k)}$.

Using (*) and the substructure lemma, one can argue that any BSR sentence φ_* that is semantically equivalent to φ must contain at least $\sum_{k=1}^n 2^{\uparrow k(n)}$ leading existential quantifiers.

The key idea is that φ_* , which is satisfied by \mathcal{A} , must contain one existential quantifier for each and every $b_S^{(k)}$. Otherwise, there would be one $b_S^{(k)}$, call it b_* , such that we could remove b_* from \mathcal{A} 's domain and any tuple $\langle \dots, b_*, \dots \rangle$ from the sets $Q_i^{\mathcal{A}}$, and the resulting structure would then still be a model of φ_* . But this would contradict (*).

Since every atom $Q_i(y_1, \dots, y_n)$ contains n variables from existential quantifier blocks that are separated by universal ones, the degree ∂_{φ} of φ is n . Moreover, φ can easily be transformed into a CNF that is Horn and Krom at the same time.

Hence, the theorem holds. \square

Theorem 16 entails that there is no elementary upper bound on the length of the BSR sentences that result from an equivalence-preserving transformation of SF sentences into BSR. On the other hand, by Lemma 12, there is an elementary upper bound, if we only consider SF sentences up to a certain degree.

5. Lower bounds on the computational complexity of SF-satisfiability

In this section we establish lower bounds on the worst-case time complexity of SF-satisfiability. Our arguments will be based on a particular form of bounded domino (or tiling) problems developed by Grädel (see [6] and [3], Section 6.1.1). By \mathbb{Z}_t

we denote the set of integers $\{0, \dots, t-1\}$ for any positive $t \geq 1$.

Definition 17 ([3], Definition 6.1.1). A *domino system* $\mathfrak{D} := \langle \mathcal{D}, \mathcal{H}, \mathcal{V} \rangle$ is a triple where \mathcal{D} is a finite set of tiles and $\mathcal{H}, \mathcal{V} \subseteq \mathcal{D} \times \mathcal{D}$ are binary relations determining the allowed horizontal and vertical neighbors of tiles, respectively. Consider the torus $\mathbb{Z}_t^2 := \mathbb{Z}_t \times \mathbb{Z}_t$ and let $\overline{D} := D_0 \dots D_{n-1}$ be a word over \mathcal{D} of length $n \leq t$. The letters of \overline{D} represent tiles. We say that \mathfrak{D} *tiles the torus* \mathbb{Z}_t^2 *with initial condition* \overline{D} if and only if there exists a mapping $\tau : \mathbb{Z}_t^2 \rightarrow \mathcal{D}$ such that for every $\langle x, y \rangle \in \mathbb{Z}_t^2$ the following conditions hold, where “+1” denotes increment modulo t .

- (a) If $\tau(x, y) = D$ and $\tau(x+1, y) = D'$, then $\langle D, D' \rangle \in \mathcal{H}$.
- (b) If $\tau(x, y) = D$ and $\tau(x, y+1) = D'$, then $\langle D, D' \rangle \in \mathcal{V}$.
- (c) $\tau(i, 0) = D_i$ for $i = 0, \dots, n-1$.

Definition 18 ([3], Definition 6.1.5). Let $T : \mathbb{N} \rightarrow \mathbb{N}$ be a function and let $\mathfrak{D} := \langle \mathcal{D}, \mathcal{H}, \mathcal{V} \rangle$ be a domino system. The problem $\text{DOMINO}(\mathfrak{D}, T(n))$ is the set of those words \overline{D} over the alphabet \mathcal{D} for which \mathfrak{D} tiles $\mathbb{Z}_{T(n)}^2$ with initial condition \overline{D} .

Domino problems provide a convenient way of deriving lower bounds via reductions. Suppose we are given some well-behaved time bound $T(n)$ that grows sufficiently fast. Further assume there is a reasonable translation from $\text{DOMINO}(\mathfrak{D}, T(n))$ into some problem \mathcal{L} where the length of the results is upper bounded by a function $g(n)$. It follows that the time required to solve the hardest instances of \mathcal{L} lies in $\Omega(T(h(n)))$, where $h(n)$ may be conceived as an inverse of $g(n)$ from an asymptotic point of view. The next proposition formalizes this observation.

Proposition 19 ([3], Theorem 6.1.8). Let $T : \mathbb{N} \rightarrow \mathbb{N}$ be a time-constructible function with $T(c'n)^2 \in o(T(n))$ for some constant $c' > 0$ and let \mathcal{L} be a problem such that for every domino system \mathfrak{D} we have $\text{DOMINO}(\mathfrak{D}, T(n)) \leq_{g(n)} \mathcal{L}$, i.e. $\text{DOMINO}(\mathfrak{D}, T(n))$ is *polynomially reducible* to \mathcal{L} via length order $g(n)$ (cf. Definition 6.1.7 in [3]). Moreover, let $h : \mathbb{N} \rightarrow \mathbb{N}$ be a function such that $h(d \cdot g(n)) \in \mathcal{O}(n)$ for any positive d . There exists a positive constant $c > 0$ such that $\mathcal{L} \notin \text{NTIME}(T(c \cdot h(n)))$.

Subsections 5.1 and 5.2 are devoted to the purpose of outlining the following reductions.

Lemma 20.

- (i) Fix some positive integer $k > 0$ and let \mathfrak{D} be an arbitrary domino system. Let $\text{Sat}(\text{SF}_{\partial \leq k})$ be the set containing all satisfiable SF sentences whose degree ∂ is at most k . We have $\text{DOMINO}(\mathfrak{D}, 2^{\uparrow k}(n)) \leq_{n \cdot \log n} \text{Sat}(\text{SF}_{\partial \leq k})$.
- (ii) Fix some positive integer $m > 1$ and let \mathfrak{D} be an arbitrary domino system. Let $\text{Sat}(\text{SF})$ be the set containing all satisfiable SF sentences. We have $\text{DOMINO}(\mathfrak{D}, 2^{\uparrow n}(m)) \leq_{n^2 \cdot \log n} \text{Sat}(\text{SF})$.

Having these reduction results at hand, Proposition 19 implies the sought lower bounds on SF-satisfiability for classes of sentences with bounded degree and the class of unbounded SF sentences.

Theorem 21. There are positive constants $c, d > 0$ for which $\text{Sat}(\text{SF}_{\partial \leq k}) \notin \text{NTIME}(2^{\uparrow k}(cn/\log n))$ and $\text{Sat}(\text{SF}) \notin \text{NTIME}(2^{\uparrow d \cdot \sqrt{n/\log n}}(2))$.

These lower bounds also hold if we do not allow equality in SF, see Section 5.3. The remainder of Section 5 is devoted to the formalization of sufficiently large tori in SF and to the translation from a given domino system $\mathfrak{D} = \langle \mathcal{D}, \mathcal{H}, \mathcal{V} \rangle$ (for nonempty $\mathcal{D}, \mathcal{H}, \mathcal{V}$) plus an initial condition \bar{D} into an SF sentence φ such that φ is satisfiable if and only if $\bar{D} \in \text{DOMINO}(\mathfrak{D}, T_i(|\bar{D}|))$ with $T_1(n) = 2^{\uparrow \kappa}(n)$ for any given $\kappa > 0$ and $T_2(n) = 2^{\uparrow n}(\mu)$ for any given $\mu > 1$.

5.1. Enforcing a large domain

The following description gives a somewhat simplified view. Technical details will follow.

A crucial part in the reduction is that a grid of size $t \times t$ has to be encoded, where t defines the required computing time and we assume $t := 2^{\uparrow \kappa}(\mu)$ for positive integers κ and $\mu > 1$ that we consider as parameters of the construction.

Every point p on the grid is represented by a pair $p = \langle x, y \rangle$, where each of the coordinates x and y is represented by a bit string of length $\log(2^{\uparrow \kappa}(\mu)) = 2^{\uparrow \kappa-1}(\mu)$. Given a bit string \bar{b} , we represent the j -th bit b_j by the truth value of the atom $J(\underline{\kappa}, \bar{b}, j)$, where $\underline{\kappa}$ is the constant used to address the topmost level of a hierarchy of $\kappa + 1$ sets of indices. The crux of our approach is that we have to enforce the existence of sufficiently many indices j , namely $2^{\uparrow \kappa-1}(\mu)$ many, to address the single bits of \bar{b} . Again, we address each of these indices as a bit string, this time of length $2^{\uparrow \kappa-2}(\mu)$.


Thus, we proceed in an inductive fashion, building up a hierarchy of indices with $\kappa + 1$

levels. The lowest level, level zero, is inhabited by μ indices, which we represent as constants with distinct values. For every $\ell \geq 1$ any index j on the ℓ -th level is represented by a bit string consisting of $2^{\uparrow \ell-1}(\mu)$ bits, i.e. the ℓ -th level of the index hierarchy contains $2^{\uparrow \ell}(\mu)$ indices. The i -th bit of an ℓ -th-level index j corresponds to the truth value of the atom $J(\ell, j, i)$.

Example 22. Assume $\mu = 2$ and $\kappa = 3$.

index level	set of indices	number of indices
0	$\{c_1, c_2\}$	2
1	$\{00, 01, 10, 11\}$	4
2	$\{0000, 0001, \dots, 1111\}$	16
3	$\{0, 1\}^{16}$	65536

On every index level, the bits of one index are indexed by the indices from the previous level. We illustrate this for the word 1010 on all levels from 2 down to 0. The bits of 1010 on level two are indexed by bit strings from level one, each of them having a length of two. The bits of the indices of level one are themselves indexed by objects of level zero which are some values c_1, c_2 assigned to the constants c_1, c_2 . To improve readability, we connect the bits of words by dashes.



For technical reasons the number of indices per level grows slightly slower than described above (cf. Lemma 24). The described index hierarchies can be encoded by SF formulas with the quantifier prefix $\exists^*(\forall \exists)^\kappa$ that have a length that is polynomial in κ and μ . We use the following constant and predicate symbols with the indicated meaning:

$\underline{0}, \underline{1}, \dots, \underline{\kappa}$	constants denoting the levels from 0 to κ ,
c_1, \dots, c_μ	denote the indices at level 0,
d_1, \dots, d_κ	d_ℓ is the min. index at level ℓ ,
e_1, \dots, e_κ	e_ℓ is the max. index at level ℓ ,
$L(\ell, j)$	index j belongs to level ℓ ,
$\text{MinIdx}(\ell, j)$	j is a min. index at level ℓ ,
$\text{MaxIdx}(\ell, j)$	j is a max. index at level ℓ ,
$J(\ell, j, i, b)$	the i -th bit of the index j at level ℓ is b ,
$J^*(\ell, j, i, b)$	$b = 1$ indicates that all the bits of the index j that are less significant than j 's i -th bit are 1,
$\text{Succ}(\ell, j, j')$	j' is the successor index of j at level ℓ .

On every level we establish an ordering over the

indices of that level. We use the usual ordering on natural numbers encoded in binary. Moreover, we formalize the usual successor relation on these numbers by the predicate Succ.

One difficulty that we encounter is that we cannot assert the existence of successors simply by adding $\forall j \exists j'. \text{Succ}(\underline{\ell}, j, j')$, as j and j' would not be separated. Therefore, we fall back on a trick: we start from the equivalent formula $\forall j \exists \tilde{j} j'. j \approx \tilde{j} \wedge \text{Succ}(\underline{\ell}, \tilde{j}, j')$, and replace the atom $j \approx \tilde{j}$ by a subformula $\text{eq}_{j, \tilde{j}}^\ell$ in which j and \tilde{j} are separated and which expresses a certain similarity between j and \tilde{j} instead of identity. However, it turns out that we can specify the hierarchy of indices in a sufficiently strong way such that the similarity expressed by $\text{eq}_{j, \tilde{j}}^\ell$ actually coincides with identity.

Next, we formalize the described index hierarchies in $\text{SF}_{\partial \leq \kappa}$. Every formula is accompanied by a brief description of its purpose. We shall try to use as few non-Horn sentences as possible.

$$\psi_1 := \bigwedge_{\ell=0}^{\kappa} \bigwedge_{\substack{\ell'=0 \\ \ell' \neq \ell}}^{\kappa} \forall j. L(\underline{\ell}, j) \rightarrow \neg L(\underline{\ell}', j)$$

Every index belongs to at most one level.

$$\psi_2 := \bigwedge_{\ell=0}^{\kappa} (\forall j. \text{MinIdx}(\underline{\ell}, j) \rightarrow L(\underline{\ell}, j)) \wedge (\forall j j'. \text{MinIdx}(\underline{\ell}, j) \rightarrow \neg \text{Succ}(\underline{\ell}, j', j))$$

A min. index of level ℓ belongs to level ℓ .
A min. index does not have a predecessor.

$$\psi_3 := \bigwedge_{\ell=0}^{\kappa} \text{MinIdx}(\underline{\ell}, d_\ell) \wedge (\forall j. \text{MinIdx}(\underline{\ell}, j) \rightarrow j \approx d_\ell)$$

There is a unique min. index on every level.

$$\psi_4 := \bigwedge_{\ell=0}^{\kappa} (\forall j. \text{MaxIdx}(\underline{\ell}, j) \rightarrow L(\underline{\ell}, j)) \wedge (\forall j j'. \text{MaxIdx}(\underline{\ell}, j) \rightarrow \neg \text{Succ}(\underline{\ell}, j, j'))$$

A max. index of level ℓ belongs to level ℓ .
A max. index does not have a successor.

$$\psi_5 := \bigwedge_{\ell=0}^{\kappa} \text{MaxIdx}(\underline{\ell}, e_\ell) \wedge (\forall j. \text{MaxIdx}(\underline{\ell}, j) \rightarrow j \approx e_\ell)$$

There is a unique max. index on every level.

$$\psi_6 := \bigwedge_{\ell=0}^{\kappa} \forall j j'. \text{Succ}(\underline{\ell}, j, j') \rightarrow L(\underline{\ell}, j) \wedge L(\underline{\ell}, j')$$

If j' is the successor of j at level ℓ , then both j and j' belong to level ℓ .

$$\psi_7 := \bigwedge_{\ell=0}^{\kappa} \forall j j' j''. \neg \text{Succ}(\underline{\ell}, j, j) \wedge (\text{Succ}(\underline{\ell}, j, j') \wedge \text{Succ}(\underline{\ell}, j, j'') \rightarrow j' \approx j'') \wedge (\text{Succ}(\underline{\ell}, j', j) \wedge \text{Succ}(\underline{\ell}, j'', j) \rightarrow j' \approx j'')$$

The successor relation is irreflexive.

Every index j has at most one successor and at most one predecessor.

$$\psi_8 := \text{MinIdx}(\underline{0}, c_1) \wedge \text{MaxIdx}(\underline{0}, c_\mu) \wedge \bigwedge_{i=1}^{\mu-1} \text{Succ}(\underline{0}, c_i, c_{i+1})$$

At level zero we have the sequence c_1, \dots, c_μ of successors, where c_1 is min. and c_μ max.

$$\psi_9 := \bigwedge_{\ell=1}^{\kappa} \forall j j' i. \text{Succ}(\underline{\ell}, j, j') \wedge L(\underline{\ell-1}, i) \rightarrow \left((J^*(\underline{\ell}, j, i, 1) \wedge J(\underline{\ell}, j, i, 1) \rightarrow J(\underline{\ell}, j', i, 0)) \wedge (J^*(\underline{\ell}, j, i, 1) \wedge J(\underline{\ell}, j, i, 0) \rightarrow J(\underline{\ell}, j', i, 1)) \wedge (J^*(\underline{\ell}, j, i, 0) \wedge J(\underline{\ell}, j, i, 1) \rightarrow J(\underline{\ell}, j', i, 1)) \wedge (J^*(\underline{\ell}, j, i, 0) \wedge J(\underline{\ell}, j, i, 0) \rightarrow J(\underline{\ell}, j', i, 0)) \right)$$

Define what it means to be a successor at level ℓ , $\ell > 0$, in terms of the binary increment operation modulo $2^{\uparrow \ell}(\mu)$.

$$\psi_{10} := \bigwedge_{\ell=1}^{\kappa} \forall j i. \text{MinIdx}(\underline{\ell}, j) \wedge L(\underline{\ell-1}, i) \rightarrow J(\underline{\ell}, j, i, 0)$$

All bits of a minimal index j are 0.

$$\psi_{11} := \bigwedge_{\ell=1}^{\kappa} \forall j i. \text{MaxIdx}(\underline{\ell}, j) \wedge \text{MaxIdx}(\underline{\ell-1}, i) \rightarrow J(\underline{\ell}, j, i, 1)$$

Define what it means to be max. (part 1):
the most significant bit is 1.

$$\psi_{12} := \bigwedge_{\ell=1}^{\kappa} \forall j i. L(\underline{\ell}, j) \wedge \text{MaxIdx}(\underline{\ell-1}, i) \wedge J(\underline{\ell}, j, i, 1) \rightarrow \text{MaxIdx}(\underline{\ell}, j)$$

Define what it means to be max. (part 2):
any index with 1 as its most significant bit is max.

$$\psi_{13} := \bigwedge_{\ell=1}^{\kappa} \forall j i. L(\underline{\ell}, j) \wedge L(\underline{\ell-1}, i) \rightarrow (J(\underline{\ell}, j, i, 0) \rightarrow \neg J(\underline{\ell}, j, i, 1)) \wedge (J^*(\underline{\ell}, j, i, 0) \rightarrow \neg J^*(\underline{\ell}, j, i, 1))$$

No bit of an index is 0 and 1 at the same time.
An analogous requirement is stipulated for J^* .

$$\psi_{14} := \bigwedge_{\ell=1}^{\kappa} \forall j i. L(\underline{\ell}, j) \wedge \text{MinIdx}(\underline{\ell-1}, i) \rightarrow J^*(\underline{\ell}, j, i, 1)$$

$J^*(\underline{\ell}, j, d_{\ell-1}, 1)$ holds for every index j .

$$\psi_{15} := \bigwedge_{\ell=1}^{\kappa} \forall j i i'. L(\underline{\ell}, j) \wedge \text{Succ}(\underline{\ell-1}, i, i') \rightarrow (J^*(\underline{\ell}, j, i', 1) \leftrightarrow (J^*(\underline{\ell}, j, i, 1) \wedge J(\underline{\ell}, j, i, 1))) \wedge (J(\underline{\ell}, j, i, 0) \rightarrow J^*(\underline{\ell}, j, i', 0)) \wedge (J^*(\underline{\ell}, j, i, 0) \rightarrow J^*(\underline{\ell}, j, i', 0))$$

Define the semantics of J^* as indicating that all bits strictly less significant than the i -th bit are 1.

$$\begin{aligned} \text{eq}_{j,\tilde{j}}^1 &:= L(\underline{1}, j) \wedge L(\underline{1}, \tilde{j}) \\ &\quad \wedge \bigwedge_{i=1}^{\mu} (J(\underline{1}, j, c_i, 0) \leftrightarrow J(\underline{1}, \tilde{j}, c_i, 0)) \\ &\quad \wedge (J(\underline{1}, j, c_i, 1) \leftrightarrow J(\underline{1}, \tilde{j}, c_i, 1)) \end{aligned}$$

Base case of equality of indices.

$$\begin{aligned} \text{eq}_{j,\tilde{j}}^\ell &:= L(\underline{\ell}, j) \wedge L(\underline{\ell}, \tilde{j}) \wedge \forall i. L(\underline{\ell-1}, i) \\ &\quad \rightarrow \exists \tilde{i}. L(\underline{\ell-1}, \tilde{i}) \wedge \text{eq}_{i,\tilde{i}}^{\ell-1} \\ &\quad \wedge (J(\underline{\ell}, j, i, 0) \leftrightarrow J(\underline{\ell}, \tilde{j}, \tilde{i}, 0)) \\ &\quad \wedge (J(\underline{\ell}, j, i, 1) \leftrightarrow J(\underline{\ell}, \tilde{j}, \tilde{i}, 1)) \end{aligned}$$

Inductive case of equality of indices for $\ell > 1$.

$$\begin{aligned} \psi_{16} &:= \\ &\quad \bigwedge_{\ell=1}^{\kappa} \forall j i. L(\underline{\ell}, j) \wedge \text{MaxIdx}(\underline{\ell-1}, i) \wedge J(\underline{\ell}, j, i, 0) \\ &\quad \rightarrow \exists \tilde{j} \tilde{j}'. \text{eq}_{j,\tilde{j}}^\ell \wedge \text{Succ}(\underline{\ell}, \tilde{j}, \tilde{j}') \end{aligned}$$

For every index at level ℓ that is not maximal, i.e. whose most significant bit is 0, there exists a successor index.

Until now, we have only introduced sentences that can easily be transformed into SF sentences in Horn form in which existential variables are separated from universal ones, as all quantifiers occur with positive polarity and as consequents of implications are (conjunctions of) literals. Regarding the length of the sentences, we observe $\text{len}(\psi_1 \wedge \dots \wedge \psi_{16}) \in \mathcal{O}(\kappa^2 \log \kappa + \kappa \mu (\log \kappa + \log \mu))$.

The following three sentences do not produce Horn formulas when transformed into CNF. They serve the purpose of removing spurious elements from the model. In particular, χ_3 is essential to enforce large models for $\kappa \geq 2$.

$$\chi_1 := \forall j. L(\underline{0}, j) \rightarrow \bigvee_{i=1}^{\mu} j \approx c_i$$

On level 0 there are no indices but c_1, \dots, c_μ .

$$\begin{aligned} \chi_2 &:= \bigwedge_{\ell=1}^{\kappa} \forall j i. L(\underline{\ell}, j) \wedge L(\underline{\ell-1}, i) \\ &\quad \rightarrow J(\underline{\ell}, j, i, 0) \vee J(\underline{\ell}, j, i, 1) \end{aligned}$$

We stipulate totality for the predicate J .

$$\begin{aligned} \chi_3 &:= \bigwedge_{\ell=1}^{\kappa} \forall j j'. L(\underline{\ell}, j) \wedge L(\underline{\ell}, j') \\ &\quad \rightarrow \exists \tilde{j} \tilde{j}'. \text{eq}_{j,\tilde{j}}^\ell \wedge \text{eq}_{j',\tilde{j}'}^\ell \\ &\quad \wedge \left(\left(\forall \tilde{i}. L(\underline{\ell-1}, \tilde{i}) \right. \right. \\ &\quad \left. \left. \rightarrow (J(\underline{\ell}, \tilde{j}, \tilde{i}, 0) \leftrightarrow J(\underline{\ell}, \tilde{j}', \tilde{i}, 0)) \right) \rightarrow j \approx j' \right) \end{aligned}$$

Two indices at the same level that agree on all of their bits are required to be identical.

Notice that χ_3 is (almost) an SF sentence, since the $\forall \tilde{i}$ turns into a $\exists \tilde{i}$ as soon as we bring the sentence into prenex normal form. We observe $\text{len}(\chi_1 \wedge \chi_2 \wedge \chi_3) \in \mathcal{O}(\kappa^2 \log \kappa + \kappa \mu (\log \kappa + \log \mu))$.

Consider any model \mathcal{A} of $\psi_1 \wedge \dots \wedge \psi_{16} \wedge \chi_1 \wedge \chi_2 \wedge \chi_3$.

Definition 23. We define the following sets and relations: $\mathcal{I}_\ell := \{\mathbf{a} \in \mathcal{U}_{\mathcal{A}} \mid \mathcal{A}, [j \mapsto \mathbf{a}] \models L(\underline{\ell}, j)\}$ for every $\ell = 0, \dots, \kappa$; $\prec_\ell \subseteq \mathcal{I}_\ell \times \mathcal{I}_\ell$ for every $\ell = 0, \dots, \kappa$ such that $\mathbf{a} \prec_\ell \mathbf{a}'$ holds if and only if $\mathcal{A}, [j \mapsto \mathbf{a}, j' \mapsto \mathbf{a}'] \models \text{Succ}(\underline{\ell}, j, j')$.

Lemma 24. For every $\ell = 1, \dots, \kappa$ we have $|\mathcal{I}_\ell| = p$ where $p := 2^{|\mathcal{I}_{\ell-1}|-1} + 1 = 2^{\uparrow \ell}(\mu - 1) + 1$. Moreover, there is a unique chain $\mathbf{a}_1 \prec_\ell \dots \prec_\ell \mathbf{a}_p$ comprising all elements in \mathcal{I}_ℓ .

Leaving out the non-Horn parts χ_1, χ_2, χ_3 renders the lemma invalid for $\ell > 1$. On the other hand, for $\kappa = 1$ the sentence $\psi_1 \wedge \dots \wedge \psi_{16}$ —which can be transformed into an equivalent Horn sentence—has only models \mathcal{A} for which \mathcal{I}_1 contains at least $2^{\mu-1} + 1$ elements. Notice that this could be used to derive EXPTIME-hardness of satisfiability for the class of Horn SF sentences of degree 1. But such lower bounds are already known for the Horn subfragments of MFO and of the BSR fragment, which are proper subsets of SF's Horn subfragment.

5.2. Formalizing a tiling of the torus

In order to formalize a given domino problem $\mathfrak{D} = \langle \mathcal{D}, \mathcal{H}, \mathcal{V} \rangle$ and an initial condition \overline{D} , we introduce the following constant and predicate symbols:

$H(x, y, x', y')$	$\langle x', y' \rangle$ is the horiz. neighbor of $\langle x, y \rangle$, i.e. $x' = x + 1 \pmod{2^{\uparrow \kappa}(\mu - 1) + 1}$ and $y' = y$,
$V(x, y, x', y')$	$\langle x', y' \rangle$ is the vert. neighbor of $\langle x, y \rangle$,
$\underline{D}(x, y)$	$\langle x, y \rangle$ is tiled with $D \in \mathcal{D}$,
$f_1, \dots, f_{ \overline{D} }$	constants addressing points $\langle 0, 0 \rangle, \dots, \langle \overline{D} - 1, 0 \rangle$.

With the ideas we have seen when formalizing the index hierarchy, it is now fairly simple to formalize the torus. For instance, the following sentence makes sure that every point that is not on the “edge” of the torus has a horizontal neighbor.

$$\begin{aligned} \eta_3 &:= \forall x y i. L(\underline{\kappa}, x) \wedge L(\underline{\kappa}, y) \\ &\quad \wedge \text{MaxIdx}(\underline{\kappa-1}, i) \wedge J(\underline{\kappa}, x, i, 0) \\ &\quad \rightarrow \exists \tilde{x} \tilde{y} \tilde{x}'. \text{eq}_{x,\tilde{x}}^\kappa \wedge \text{eq}_{y,\tilde{y}}^\kappa \\ &\quad \wedge \left(\bigwedge_{D \in \mathcal{D}} \underline{D}(x, y) \leftrightarrow \underline{D}(\tilde{x}, \tilde{y}) \right) \wedge H(\tilde{x}, \tilde{y}, \tilde{x}', \tilde{y}) \end{aligned}$$

The next sentence, on the other hand, makes sure that the rules of the domino system \mathfrak{D} are obeyed.

$$\eta_{15} := \forall x x' y. H(x, y, x', y) \rightarrow \bigvee_{\langle D, D' \rangle \in \mathcal{H}} \underline{D}(x, y) \wedge \underline{D}'(x', y)$$

Proceeding this way, the formalization η of a domino system in SF requires a length in $\mathcal{O}(\hat{n} \log \hat{n})$, where $\hat{n} := \max\{\kappa, \mu, |\overline{D}|, |\mathcal{D}|^2\}$.

Lemma 25. Assume that \mathcal{D} , \mathcal{H} , and \mathcal{V} are nonempty and let \mathcal{A} be a model of the sentence $\psi_1 \wedge \dots \wedge \psi_{16} \wedge \chi_1 \wedge \chi_2 \wedge \chi_3 \wedge \eta$. \mathcal{A} induces a tiling τ of \mathbb{Z}_t^2 with initial condition $\overline{D} := D_1, \dots, D_n$, where $t := 2^{\uparrow \kappa}(\mu - 1) + 1$.

5.3. Replacing the equality predicate

Since SF can express reflexivity, symmetry, transitivity, and congruence properties, it is easy to formulate an SF sentence without equality that is equisatisfiable to $\psi_1 \wedge \dots \wedge \psi_{16} \wedge \chi_1 \wedge \chi_2 \wedge \chi_3 \wedge \eta$ and uses atoms $E(s, t)$ instead of $s \approx t$. In addition to replacing equational atoms as indicated, we add the usual axioms concerning the fresh predicate symbol E . Overall, the additional formulas have a length that lies in $\mathcal{O}(\kappa \log \kappa + |\mathcal{D}| \log |\mathcal{D}|)$.

Consequently, the hardness result that we have obtained for SF with equality can be directly transferred to SF without equality.

6. Conclusion

We stress in this paper that an analysis of the computational complexity of satisfiability problems can greatly benefit from an analysis of how variables occur together in atoms instead of exclusively considering the number of occurring quantifier alternations. What we have not yet taken into account is the Boolean structure of sentences. This may widen the scope of our methods considerably and may moreover help understand where the hardness of satisfiability problems stems from.

Consider a quantified Boolean formula $\varphi := \forall \bar{x}_1 \exists \bar{y}_1 \dots \forall \bar{x}_n \exists \bar{y}_n. \psi$ with quantifier-free ψ . All satisfiable formulas of this shape together form a hard problem residing on the n -th level of the polynomial hierarchy. But what if, for instance, ψ has the form $(\bigwedge_i K_i) \wedge (\bigvee_j L_j)$, where the K_i and the L_j are literals and none of the existential variables in $\bigwedge_i K_i$ occurs in $\bigvee_j L_j$? Since Boolean variables cannot jointly occur in atoms, φ can be transformed into the equivalent formula $\exists \bar{y}_1 \dots \exists \bar{y}_n \forall \bar{x}_1 \dots \forall \bar{x}_n. \psi$ by application of the rules of miniscoping (cf. Lemma 2). Apparently, φ belongs to a class of QBF sentences that resides on the first level of the polynomial hierarchy rather than on the n -th.

Perhaps it is time to reconsider some of the definitions that are based on the shape of quantifier prefixes alone.

References

- [1] H. Andréka, I. Németi, and J. van Benthem. Modal Languages and Bounded Fragments of Predicate Logic. *J. Philosophical Logic*, 27(3):217–274, 1998.
- [2] V. Bárány, B. ten Cate, and L. Segoufin. Guarded Negation. *J. ACM*, 62(3):22, 2015.
- [3] E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer, 1997.
- [4] S. A. Cook. The Complexity of Theorem-Proving Procedures. In *Theory of Computing (STOC'71)*, pages 151–158, 1971.
- [5] H. Ebbinghaus, J. Flum, and W. Thomas. *Mathematical Logic*. Springer, 1994.
- [6] E. Grädel. On solvable cases of Hilbert’s ‘Entscheidungsproblem’. Habilitationsschrift, Universität Basel, 1990.
- [7] E. Grädel. On the Restraining Power of Guards. *J. Symb. Log.*, 64(4):1719–1742, 1999.
- [8] E. Grädel, P. G. Kolaitis, and M. Y. Vardi. On the Decision Problem for Two-Variable First-Order Logic. *Bulletin of Symbolic Logic*, 3(1):53–69, 1997.
- [9] N. D. Jones, Y. E. Lien, and W. T. Laaser. New Problems Complete for Nondeterministic Log Space. *Mathematical systems theory*, 10(1):1–17, 1976.
- [10] S. Kasif. On the Parallel Complexity of Some Constraint Satisfaction Problems. In *Artificial Intelligence (AAAI'86)*, volume 1, pages 349–353, 1986.
- [11] E. Kieronski and A. Kuusisto. Complexity and Expressivity of Uniform One-Dimensional Fragment with Equality. In *Mathematical Foundations of Computer Science (MFCS'14)*, pages 365–376, 2014.
- [12] M. Mortimer. On Languages with Two Variables. *Mathematical Logic Quarterly*, 21(1):135–140, 1975.
- [13] D. A. Plaisted. Complete Problems in the First-Order Predicate Calculus. *J. Comput. Syst. Sci.*, 29(1):8–35, 1984.
- [14] I. Pratt-Hartmann, W. Szwaast, and L. Tendera. Quine’s Fluted Fragment is Non-Elementary. In *Computer Science Logic (CSL'16)*, pages 39:1–39:21, 2016.
- [15] W. V. O. 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.
- [16] W. V. O. Quine. The Variable. In W. V. O. Quine, editor, *The Ways of Paradox and other essays. Revised and enlarged edition*. Harvard University Press, 1976.
- [17] T. Sturm, M. Voigt, and C. Weidenbach. Deciding First-Order Satisfiability when Universal and Existential Variables are Separated. In *Logic in Computer Science (LICS'16)*, pages 86–95, 2016.
- [18] M. Voigt. A Fine-Grained Hierarchy of Hard Problems in the Separated Fragment. *ArXiv*, (arXiv:1704.02145 [cs.LO]), 2017.