Gaifman Normal Forms for Counting Extensions of First-Order Logic

Dietrich Kuske
Technische Universität Ilmenau, Germany
dietrich.kuske@tu-ilmenau.de

Nicole Schweikardt[1]
Humboldt-Universität zu Berlin, Germany
schweikn@informatik.hu-berlin.de

Abstract

We consider the extension of first-order logic FO by unary counting quantifiers and generalise the notion of Gaifman normal form from FO to this setting. For formulas that use only ultimately periodic counting quantifiers, we provide an algorithm that computes equivalent formulas in Gaifman normal form. We also show that this is not possible for formulas using at least one quantifier that is not ultimately periodic.

Now let d be a degree bound. We show that for any formula φ with arbitrary counting quantifiers, there is a formula γ in Gaifman normal form that is equivalent to φ on all finite structures of degree $\leqslant d$. If the quantifiers of φ are decidable (decidable in elementary time, ultimately periodic), γ can be constructed effectively (in elementary time, in worst-case optimal 3-fold exponential time).

For the setting with unrestricted degree we show that by using our Gaifman normal form for formulas with only ultimately periodic counting quantifiers, a known fixed-parameter tractability result for FO on classes of structures of bounded local tree-width can be lifted to the extension of FO with ultimately periodic counting quantifiers (a logic equally expressive as FO+MOD, i.e., first-oder logic with modulo-counting quantifiers).

2012 ACM Subject Classification Theory of computation → Logic

Keywords and phrases Finite model theory, Gaifman locality, modulo-counting quantifiers, fixed parameter tractable model-checking

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.133

1 Introduction

As database specialists know very well, when evaluating a query (i.e., a formula) in a database (i.e., a relational structure), it is often advantageous to first transform the formula into an equivalent one and then evaluate this new formula in the given structure. Using this approach, one also gets algorithmic meta-theorems stating that the evaluation of formulas from a certain logic in structures from a certain class is fixed-parameter tractable. For example, this is known for formulas from monadic second-order logic MSO and its extension CMSO with modulo-counting predicates and the class of labeled trees [5, 27] or classes of bounded tree-width [3, 1]. For first-order logic FO, it is known for classes of structures of bounded degree [25], for the class of planar graphs and, more generally, for classes of bounded local tree-width [9], for classes of locally bounded expansion [6], and for classes that are effectively nowhere dense [12].

[1] Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SCHW 837/5-1.

Gaifman's normal form theorem [10] provides an approach to this formula transformation, and it has been applied for obtaining several of the results on FO mentioned above. It is the aim of this paper to demonstrate that this approach does not only work for first-order logic, but also for its extension FO+MOD by modulo-counting quantifiers (these quantifiers allow to make statements of the form "the number of witnesses x for a formula φ is congruent r modulo m", for fixed integers r and m). The logic FO+MOD has been well-studied, see e.g., [26, 18, 24, 23, 14, 17]. Its expressivity lies strictly between that of FO and CMSO. It is known that all FO+MOD-queries are Gaifman-local [18]. But an extension of Gaifman's normal form theorem [10] from FO to FO+MOD has not been achieved in the literature.

Recall that a first-order formula is in Gaifman normal form if it is a Boolean combination of (1) first-order properties of the neighbourhood of the free variables and (2) statements that express the existence of mutually far-apart elements whose neighbourhoods share a first-order property. Gaifman's normal form theorem states that every first-order formula is effectively equivalent to such a formula in Gaifman normal form.

We propose a notion of Gaifman normal form for FO+MOD and, more generally, for the extension of first-order logic by unary counting quantifiers FO($\mathbf{Q}$): it is a Boolean combination of (1) FO($\mathbf{Q}$)-properties of the neighbourhood of the free variables, (2) statements that express the existence of mutually far-apart elements whose neighbourhoods share an FO($\mathbf{Q}$)-property, and (3) statements that depend on the total number of elements whose neighbourhoods share an FO($\mathbf{Q}$)-property. We show that if a formula uses only ultimately periodic counting quantifiers (and therefore is equivalent to a formula of FO+MOD), then Gaifman's theorem holds *mutatis mutandis*: any such formula can be transformed effectively into an equivalent formula in Gaifman normal form using the same counting quantifiers. The proof of this result extends the original proof for first-order logic from [10]; a crucial ingredient is an effective Feferman-Vaught decomposition [8] for FO($\mathbf{Q}$) that we prove first. Adapting [4], we show that the size of the resulting formula cannot be bounded by an elementary function. Furthermore, we prove that formulas with non-ultimately periodic counting quantifiers (e.g., the set of primes) do not have equivalent formulas in Gaifman normal form.

The situation changes when we restrict attention to classes of finite structures of bounded degree. Call two formulas "finitely d-equivalent" if they are equivalent on all finite structures of degree $\leqslant d$. We show that (1) for a formula with ultimately periodic counting quantifiers, one can compute in (worst-case optimal) 3-fold exponential time a finitely d-equivalent formula in Gaifman normal form; (2) from a formula with computable counting quantifiers, we can effectively compute a finitely d-equivalent formula in Gaifman normal form; and (3) if we allow arbitary counting quantifiers, then we get at least the existence of finitely d-equivalent Gaifman normal forms. In other words, by restricting the class of structures, the complexity drops from non-elementary to 3-fold exponential (for ultimately periodic counting quantifiers), from non-existent to computable (for computable quantifiers), and from non-existent to existent (for arbitrary quantifiers). The proofs of these results do not follow Gaifman's original proof, but generalise a proof for first-order logic from [16] that, in turn, builds on [2]. In the present setting of FO($\mathbf{Q}$), we first transform the original FO($\mathbf{Q}$)-formula in elementary time into a formula in "(weak) Hanf normal form" [17, 19], and afterwards we transform this formula into Gaifman normal form by a construction similar to the one in [16].

We also provide an algorithmic application that demonstrates the usefulness of our normal form: By applying our Gaifman normal form algorithm, we lift the result of [9] from FO to FO+MOD, showing that the model-checking problem for FO+MOD-sentences on classes of finite relational structures of bounded local tree-width is fixed-parameter tractable.

The rest of the paper is structured as follows. Section 2 provides the basic definitions and introduces our notion of Gaifman normal form. Section 3 provides effective Feferman-Vaught decompositions for the extension of first-order logic by ultimately periodic counting quantifiers. The sections 4 and 5 present our results in the setting without and with a degree bound, respectively. Section 6 provides an algorithmic application.

2 Preliminaries

We write $\mathcal{P}(S)$ to denote the power set of a set S. For an n-tuple $\overline{x} = (x_1, \ldots, x_n)$ we write $|\overline{x}|$ to denote the tuple's length n. We write $\mathbb{N}$ for the set of non-negative integers, and we let $\mathbb{N}_{\geqslant 1} = \mathbb{N} \setminus \{0\}$. For $m, n \in \mathbb{N}$ with $m \leqslant n$, we write $[m, n]$ for the set $\{i \in \mathbb{N} : m \leqslant i \leqslant n\}$. For a real number $r > 0$, we write $\log(r)$ to denote the logarithm of r with respect to base 2. We use the standard $\mathcal{O}$-notation, and by $\mathrm{poly}(n)$ we mean $n^{\mathcal{O}(1)}$. We say that a function f from $\mathbb{N}$ to the set $\mathbb{R}_{\geqslant 0}$ of non-negative reals is *at most k-fold exponential*, for some $k \in \mathbb{N}$, if there exists a number $c > 0$ such that for all sufficiently large $n \in \mathbb{N}$ we have $f(n) \leqslant \exp_k(n^c)$, where $\exp_k(m)$ is a tower of 2*s* of height k with an m on top, i.e., $\exp_0(m) = m$ and $\exp_{k+1}(m) = 2^{\exp_k(m)}$ for all $k, m \geqslant 0$. A function f is *elementary* if it is at most k-fold exponential for some $k \geqslant 0$. The function tower: $\mathbb{N} \to \mathbb{N}$, defined via $\mathrm{tower}(h) := \exp_h(1)$ for all $h \in \mathbb{N}$, is not elementary.

Structures and formulas. A *signature* σ is a *finite* set of relation symbols and constant symbols. Associated with every relation symbol R is a positive integer $\mathrm{ar}(R)$ called the *arity* of R. We call a signature *relational* if it only contains relation symbols. A *σ-structure* $\mathcal{A}$ consists of a non-empty set A called the *universe* of $\mathcal{A}$, a relation $R^{\mathcal{A}} \subseteq A^{\mathrm{ar}(R)}$ for each relation symbol $R \in \sigma$, and an element $c^{\mathcal{A}} \in A$ for each constant symbol $c \in \sigma$. Note that according to this definition, all signatures considered in this paper are finite while structures can be infinite. We write $\mathcal{A} \cong \mathcal{B}$ to indicate that two σ-structures $\mathcal{A}$ and $\mathcal{B}$ are isomorphic.

We use the standard notation concerning first-order logic and extensions thereof, cf. [7, 20]. By $\mathrm{FO}[\sigma]$ we denote the class of all first-order formulas of signature σ, and by FO we denote the union of all $\mathrm{FO}[\sigma]$ for arbitrary signatures σ. By $\mathrm{free}(\varphi)$ we denote the set of all *free variables* of the formula φ. A *sentence* is a formula φ with $\mathrm{free}(\varphi) = \emptyset$. We write $\varphi(\overline{x})$, for $\overline{x} = (x_1, \ldots, x_n)$ with $n \geqslant 0$, to indicate that $\mathrm{free}(\varphi) \subseteq \{x_1, \ldots, x_n\}$. If $\mathcal{A}$ is a σ-structure and $\overline{a} = (a_1, \ldots, a_n) \in A^n$, by $\mathcal{A} \models \varphi(\overline{a})$ or $(\mathcal{A}, \overline{a}) \models \varphi$ we indicate that the formula $\varphi(\overline{x})$ is satisfied in $\mathcal{A}$ when interpreting the free occurrences of the variables $x_1, \ldots, x_n$ with the elements $a_1, \ldots, a_n$.

Unary counting quantifiers. In addition to the existential quantifier $\exists$ we consider *unary counting quantifiers* (for short: counting quantifiers), which are defined as subsets of $\mathbb{N}$. We will use the terms "set (of natural numbers)" and "counting quantifier" interchangeably. For a set $\mathbf{Q} \subseteq \mathcal{P}(\mathbb{N})$ of counting quantifiers we write $\mathrm{FO}(\mathbf{Q})[\sigma]$ to denote the *extension of* $\mathrm{FO}[\sigma]$ *with the quantifiers from* $\mathbf{Q}$. Precisely, we add the following formation rule for formulas:

If $\varphi(\overline{x}, y) \in \mathrm{FO}(\mathbf{Q})[\sigma]$, $\mathsf{Q} \in \mathbf{Q}$, and $k \in \mathbb{N}$, then also $(\mathsf{Q}{+}k)y\,\varphi$ belongs to $\mathrm{FO}(\mathbf{Q})[\sigma]$.

For $(\mathsf{Q}{+}0)y\,\varphi$ we write the more succinct $\mathsf{Q}y\,\varphi$. The formula $(\mathsf{Q}{+}k)y\,\varphi(\overline{x}, y)$ expresses that the number of witnesses y for $\varphi(\overline{x}, y)$ belongs to the set $(\mathsf{Q}{+}k) := \{q + k : q \in \mathsf{Q}\}$. Equivalently, this means that the formula $(\mathsf{Q}{+}k)y\,\varphi(\overline{x}, y)$ is satisfied by a σ-structure $\mathcal{A}$ and an interpretation $\overline{a}$ of the variables $\overline{x}$ iff $|\{b \in A : \mathcal{A} \models \varphi(\overline{a}, b)\}| - k \in \mathsf{Q}$. Here, for an *infinite* set B we use the convention that $|B| = \infty \notin \mathbb{N}$, where ∞ is larger than any integer, and $\infty - k = \infty$ for all integers k. Every formula can be transformed into an equivalent exponentially larger *displacement-free* formula, meaning that counting quantifiers appear only in the form $(\mathsf{Q}+0)$.

▶ **Example 2.1.** For $m \geqslant 2$, the quantifier $\mathsf{D}_m = m \cdot \mathbb{N}$ contains the multiples of m. Let $\mathbf{D} = \{\mathsf{D}_m : m \in \mathbb{N}, m \geqslant 2\}$ denote the collection of all these *divisibility quantifiers*. Then the logic $\mathrm{FO}(\mathbf{D})$ is equally expressive as the logic FO+MOD (cf. [26, 24]).

The quantifier rank $\mathrm{qr}(\varphi)$ of an $\mathrm{FO}(\mathbf{Q})$-formula φ is defined as the maximal nesting depth of *all* quantifiers. For a number $\ell \in \mathbb{N}_{\geqslant 1}$ and a formula $\varphi(\overline{x}, y)$, we write $\exists^{\geqslant \ell} y\, \varphi$ to denote a formula expressing that there are at least ℓ witnesses y which satisfy φ.

Gaifman graph. The *Gaifman graph* $G_{\mathcal{A}}$ of a σ-structure $\mathcal{A}$ is the undirected, loop-free graph with vertex set A and an edge between two distinct vertices $a, b \in A$ iff there exist a relation symbol $R \in \sigma$ and a tuple $(a_1, \ldots, a_{\mathrm{ar}(R)}) \in R^{\mathcal{A}}$ such that $a, b \in \{a_1, \ldots, a_{\mathrm{ar}(R)}\}$. The *degree* of a σ-structure $\mathcal{A}$ is the degree of its Gaifman graph $G_{\mathcal{A}}$. If this degree is at most d, then we call $\mathcal{A}$ *d-bounded*. Two formulas $\varphi(\overline{x})$ and $\psi(\overline{x})$ of signature σ are called *finitely d-equivalent* if $\mathcal{A} \models \forall \overline{x}\, (\varphi \leftrightarrow \psi)$ holds for every finite d-bounded σ-structure $\mathcal{A}$.

The *distance* $\mathrm{dist}^{\mathcal{A}}(a, b)$ between two elements $a, b \in A$ is the minimal length (i.e., the number of edges) of a path from a to b in $G_{\mathcal{A}}$, and if no such path exists, we set $\mathrm{dist}^{\mathcal{A}}(a, b) = \infty$. For a tuple $\overline{a} \in A^m$ and an element $b \in A$, we let $\mathrm{dist}^{\mathcal{A}}(\overline{a}, b) = \min\{\mathrm{dist}^{\mathcal{A}}(a_i, b) : 1 \leqslant i \leqslant m\}$.

For any signature σ and any $k, r \in \mathbb{N}$, there exists a formula $\mathrm{dist}_{<r}(\overline{x}, y) \in \mathrm{FO}[\sigma]$ with $\overline{x} = (x_1, \ldots, x_k)$ such that for any σ-structure $\mathcal{A}$, any $\overline{a} \in A^k$, and any $b \in A$ we have $(\mathcal{A}, \overline{a}, b) \models \mathrm{dist}_{<r}(\overline{x}, y)$ iff $\mathrm{dist}^{\mathcal{A}}(\overline{a}, b) < r$. We write $\mathrm{dist}(\overline{x}, y) < r$ for the formula $\mathrm{dist}_{<r}(\overline{x}, y)$, and $\mathrm{dist}(\overline{x}, y) > r$ for the formula $\neg\, \mathrm{dist}_{<r+1}(\overline{x}, y)$.

Gaifman normal forms. Let $\mathbf{Q}$ be a set of counting quantifiers and σ a signature. A formula $\lambda(\overline{x}) \in \mathrm{FO}(\mathbf{Q})[\sigma]$ is *local* if all quantifications in λ are restricted to the neighborhood of the variables $\overline{x}$. The precise inductive definition proceeds as follows:

- Atomic formulas are *r-local* for any $r \in \mathbb{N}$.
- If $\lambda(\overline{x}, y)$ is r-local, $\overline{z}$ is a non-empty sub-tuple of $\overline{x}$, and $r' \in \mathbb{N}$, then the formula $\exists y\, \big(\mathrm{dist}(\overline{z}, y) \leqslant r' \wedge \lambda(\overline{x}, y)\big)$ is *$(r'+r)$-local*, and for all $\mathsf{Q} \in \mathbf{Q}$ and $k \in \mathbb{N}$, the formula $(\mathsf{Q}{+}k)y\, \big(\mathrm{dist}(\overline{z}, y) \leqslant r' \wedge \lambda(\overline{x}, y)\big)$ is *$(r'+r)$-local* as well.
- If $\lambda_1(\overline{x})$ and $\lambda_2(\overline{x})$ are r-local, then $(\lambda_1 \wedge \lambda_2)$ and $\neg\lambda_1$ are *r-local* as well.

An r-local formula is also $(r+1)$-local. A formula is *local* if it is r-local for some $r \in \mathbb{N}$.

For a σ-structure $\mathcal{A}$, a tuple $\overline{a} = (a_1, \ldots, a_n) \in A^n$ and a number $r \in \mathbb{N}$, the *r-ball* $N_r^{\mathcal{A}}(\overline{a})$ of $\overline{a}$ is the set of all $b \in A$ with $\mathrm{dist}^{\mathcal{A}}(\overline{a}, b) \leqslant r$. If σ is a *relational* signature, then the *r-neighbourhood* $\mathcal{N}_r^{\mathcal{A}}(\overline{a})$ of $\overline{a}$ is the substructure of $\mathcal{A}$ induced on the set $N_r^{\mathcal{A}}(\overline{a})$.

Let $\lambda(\overline{x})$ be an r-local formula and let $\overline{a} \in A^{|\overline{x}|}$. When determining whether $\mathcal{A} \models \lambda(\overline{a})$, quantification is restricted to elements of distance at most r from $\overline{a}$, i.e., to elements in $N_r^{\mathcal{A}}(\overline{a})$. Hence we get $\mathcal{A} \models \lambda(\overline{a}) \iff \mathcal{N}_r^{\mathcal{A}}(\overline{a}) \models \lambda(\overline{a})$.

Let $\mathrm{L} \subseteq \mathrm{FO}(\mathbf{Q})$. A *counting sentence* over L is a sentence of the form $(\mathsf{Q}{+}k)x\, \lambda(x)$, where $\lambda \in \mathrm{L}$ is local, $\mathsf{Q} \in \mathbf{Q}$, and $k \in \mathbb{N}$. A *basic local sentence* over L is a sentence of the form $\exists x_1 \cdots \exists x_m \Big(\bigwedge_{1 \leqslant i < j \leqslant m} \mathrm{dist}(x_i, x_j) > 2r \;\wedge \bigwedge_{1 \leqslant i \leqslant m} \lambda(x_i) \Big)$, where $m \in \mathbb{N}_{\geqslant 1}$, $r \in \mathbb{N}$, and $\lambda \in \mathrm{L}$ is r-local. Such a basic local sentence expresses that there are m witnesses for λ of mutual distance at least $2r+1$.

▶ **Definition 2.2.** A formula $\varphi(\overline{x}) \in \mathrm{FO}(\mathbf{Q})[\sigma]$ is *in Gaifman normal form* if it is a Boolean combination of local formulas $\lambda(\overline{x})$, of counting sentences $(\mathsf{Q}{+}k)y\, \lambda(y)$, and of basic local sentences over $\mathrm{FO}(\mathbf{Q})[\sigma]$.

If $\mathbf{Q} = \emptyset$, then a formula in Gaifman normal form consists of local formulas and basic local sentences over FO. In other words, our definition of Gaifman normal form for formulas from $\mathrm{FO}(\emptyset)$ coincides with the traditional one for formulas from FO [10, 11]. The following is known about the existence and computability of formulas in Gaifman normal form:

▶ **Theorem 2.3.** *Let $\mathbf{Q} = \emptyset$, i.e., $\mathrm{FO}(\mathbf{Q}) = \mathrm{FO}$.*

(1) *From a formula $\varphi(\overline{x}) \in \mathrm{FO}$, one can compute an equivalent formula $\gamma(\overline{x}) \in \mathrm{FO}$ in Gaifman normal form [10].*

(2) *The size of the equivalent formula in Gaifman normal form cannot be bounded by an elementary function in the size of the formula $\varphi(\overline{x})$ [4].*

(3) *From a formula $\varphi(\overline{x}) \in \mathrm{FO}$ and a degree bound d, one can compute in 3-fold exponential time a finitely d-equivalent formula $\gamma(\overline{x})$ in Gaifman normal form [16].*

It is the aim of this paper to study to what extent this theorem holds for the extension of first-order logic with unary counting quantifiers. A special role is played by ultimately periodic quantifiers that we introduce now.

Ultimately periodic sets. Let $p \in \mathbb{N}_{\geqslant 1}$ and $n_0 \in \mathbb{N}$. A set $\mathsf{Q} \subseteq \mathbb{N}$ is *ultimately p-periodic with offset n_0* if $n \in \mathsf{Q} \iff n{+}p \in \mathsf{Q}$ holds for all $n \geqslant n_0$. A set Q is *ultimately p-periodic* if there exists an n_0 such that Q is ultimately p-periodic with offset n_0, and Q is *ultimately periodic* (cf. e.g. [22]) if it is ultimately p-periodic for some $p \in \mathbb{N}_{\geqslant 1}$. The *period* of Q is the *minimal* p for which Q is ultimately p-periodic.

We write $\mathbf{U}$ to denote the set of all ultimately periodic sets $\mathsf{Q} \subseteq \mathbb{N}$.

The *characteristic sequence* χ_{Q} of a set $\mathsf{Q} \subseteq \mathbb{N}$ is the ω-word $w = w_0 w_1 w_2 \cdots \in \{0,1\}^\omega$ with $\mathsf{Q} = \{n \in \mathbb{N} : w_n = 1\}$. Note that Q is ultimately periodic iff there are finite words $\alpha, \pi \in \{0,1\}^*$ with $\chi_{\mathsf{Q}} = \alpha\pi^\omega$. We represent an ultimately periodic set Q by the shortest word $\mathrm{rep}(\mathsf{Q}) := \alpha\#\pi$ satisfying $\chi_{\mathsf{Q}} = \alpha\,\pi^\omega$. The *size* $\|\mathsf{Q}\|$ of Q is defined as the length of the word $\mathrm{rep}(\mathsf{Q})$. The size of an $\mathrm{FO}(\mathbf{U})$-formula φ of signature σ is its length when viewed as a word over the alphabet $\sigma \cup \mathrm{Var} \cup \{,\} \cup \{=, \exists, \neg, \vee, (,), 0, 1, +, \#\}$, where Var is a countable set of variables, each quantifier $\mathsf{Q} \in \mathbf{U}$ is represented by the word $\mathrm{rep}(\mathsf{Q})$, and each number k is given in binary (in subformulas of the form $(\mathsf{Q}{+}k)y\,\varphi$). The set $\mathbf{D}$ of all divisibility quantifiers (see Example 2.1) is a subset of $\mathbf{U}$. With every set of ultimately periodic quantifiers $\mathbf{Q} \subseteq \mathbf{U}$, we associate the set $\mathbf{D}_{\mathbf{Q}} \subseteq \mathbf{D}$ of divisibility quantifiers which consists of precisely those $\mathsf{D}_p = p \cdot \mathbb{N}$ for which $p \geqslant 2$ is the period of some $\mathsf{Q} \in \mathbf{Q}$.

▶ **Lemma 2.4** ([16, 17])**.** *Let $\mathbf{Q} \subseteq \mathcal{P}(\mathbb{N})$ be a set of counting quantifiers.*

(a) *Let $\mathsf{Q} \in \mathbf{Q}$ be ultimately p-periodic with offset n_0, and let $k \in \mathbb{N}$. Every formula from $\mathrm{FO}(\mathbf{Q} \cup \mathbf{D}_{\mathbf{Q}})$ of the shape $(\mathsf{Q}{+}k)y\,\varphi$ is equivalent to a Boolean combination of formulas of the form $(\mathsf{D}_p{+}\ell)y\,\varphi$ and $\exists^{\geqslant m}y\,\varphi$ with $\ell < p$ and $m < n_0{+}k{+}p$. This Boolean combination can be computed from $\mathrm{rep}(\mathsf{Q}), k, \varphi$ in polynomial time.*

(b) *Let $\mathsf{Q} \in \mathbf{Q}$ be ultimately periodic with period $p \geqslant 2$.*
Every formula from $\mathrm{FO}(\mathbf{Q} \cup \mathbf{D}_{\mathbf{Q}})$ of the shape $(\mathsf{D}_p{+}k)y\,\varphi$, for a $k \geqslant 0$, is equivalent to a Boolean combination of formulas of the shape $(\mathsf{Q}{+}\ell)y\,\varphi$ and $\exists^{\geqslant \ell}y\,\varphi$ with $\ell < \|\mathsf{Q}\| + p + k$. This Boolean combination can be computed from $\mathrm{rep}(\mathsf{Q}), k, \varphi$ in polynomial time.

(c) *Let $\mathrm{L} \subseteq \mathrm{FO}(\mathbf{Q})$ be a set of formulas that contains all atomic formulas and is closed under Boolean combinations and existential quantification.*
Every sentence $\exists^{\geqslant \ell}y\,\lambda(y)$, where $\lambda \in \mathrm{L}$ is a local formula, is equivalent to a Boolean combination of basic local sentences over L.
This Boolean combination can be computed from λ and ℓ in time $\mathcal{O}(\|\lambda\|) \cdot 2^{\mathcal{O}(\ell \log \ell)}$.

As an immediate consequence, we obtain:

▶ **Corollary 2.5.** *Let $\mathbf{Q} \subseteq \mathbf{U}$ and let $\mathbf{D_Q}$ be the associated set of divisibility quantifiers.*

(1) *For every $\mathrm{FO}(\mathbf{Q})$-formula φ, we can compute an equivalent $\mathrm{FO}(\mathbf{D_Q})$-formula and vice versa.*

(2) *For every $\mathrm{FO}(\mathbf{D_Q})$-formula ψ in Gaifman normal form, we can compute an equivalent $\mathrm{FO}(\mathbf{Q})$-formula in Gaifman normal form.*

Feferman-Vaught decompositions. A crucial tool in the construction of Gaifman normal forms for first-order logic (i.e., in the proof of Theorem 2.3(1)) is a result by Feferman and Vaught [8]. In its simplest form (which is all that is needed in this context) it expresses that the first-order theory of the disjoint sum of two structures is determined by the first-order theories of the two structures. Our proof of the generalisation of Theorem 2.3(1) to logics of the form $\mathrm{FO}(\mathbf{Q})$ will proceed similarly to Gaifman's proof for FO, and this requires us to first provide a generalisation of the result by Feferman and Vaught.

Let σ be a relational signature and let $\mathcal{A}$ and $\mathcal{B}$ be disjoint σ-structures (i.e., their universes A and B are disjoint). The *disjoint sum* $\mathcal{A} \oplus \mathcal{B}$ of $\mathcal{A}$ and $\mathcal{B}$ is the structure $\big(A \cup B,\ A,\ B,\ (R^{\mathcal{A}} \cup R^{\mathcal{B}})_{R\in\sigma}\big)$ over the signature σ_2 with two additional unary relation symbols (that we denote A and B). Since this is only defined for disjoint structures, the relations A and B form a partition of the universe of the σ_2-structure $\mathcal{A} \oplus \mathcal{B}$. Furthermore, no edge of the Gaifman graph of $\mathcal{A} \oplus \mathcal{B}$ connects nodes from A with nodes from B.

▶ **Definition 2.6.** Let $\mathbf{Q}$ be a set of counting quantifiers and $\varphi(\overline{x}, \overline{y}) \in \mathrm{FO}(\mathbf{Q})[\sigma_2]$ a formula with $\overline{x} = (x_1, \ldots, x_k)$ and $\overline{y} = (y_1, \ldots, y_\ell)$. Furthermore, let Δ be a finite set of pairs of formulas $(\alpha(\overline{x}), \beta(\overline{y}))$ from $\mathrm{FO}(\mathbf{Q})[\sigma]$. The set Δ is a *decomposition of φ w.r.t. $(\overline{x}; \overline{y})$* if

$$\mathcal{A} \oplus \mathcal{B} \models \varphi(\overline{a}, \overline{b}) \quad \iff \quad \text{there exists } (\alpha, \beta) \in \Delta \text{ with } \mathcal{A} \models \alpha(\overline{a}) \text{ and } \mathcal{B} \models \beta(\overline{b})$$

holds for all disjoint σ-structures $\mathcal{A}$ and $\mathcal{B}$ and all tuples $\overline{a} \in A^k$ and $\overline{b} \in B^\ell$.

The following is known about the existence and computability of decompositions:

▶ **Theorem 2.7.** *Let $\mathbf{Q} = \emptyset$, i.e., $\mathrm{FO}(\mathbf{Q}) = \mathrm{FO}$.*

(1) *From a formula $\varphi(\overline{x}, \overline{y}) \in \mathrm{FO}[\sigma_2]$, one can compute a decomposition w.r.t. $(\overline{x}; \overline{y})$ [8].*

(2) *The size of the decomposition cannot be bounded by an elementary function in the size of the formula $\varphi(\overline{x}, \overline{y})$ [4].*

A more general definition of the notion "decomposition" replaces the condition "there exists $(\alpha, \beta) \in \Delta$ with ..." by a Boolean combination of statements of the form "$\mathcal{A} \models \alpha(\overline{a})$" and "$\mathcal{B} \models \beta(\overline{b})$"; and with this definition, a "decomposition" can be computed in 3-fold exponential time from any $\varphi \in \mathrm{FO}(\mathbf{D})[\sigma_2]$ and for any fixed $d \geqslant 0$; but this decomposition is only equivalent to φ provided $\mathcal{A}$ and $\mathcal{B}$ are finite and of degree at most d [15, Theorem 5.2.1] (see [13] for the first-order case). Another result in this direction is due to Courcelle who considers the extension CMSO of monadic second-order logic by predicates expressing the size of a set modulo some fixed number. In this context, Courcelle also proves a result analogous to Theorem 2.7(1) [3, Lemma 4.5]. More results in this vein can be found in [21].

3 Feferman-Vaught decompositions for FO(Q)

If $\mathbf{Q}$ is a set of counting quantifiers and $\mathsf{S} \in \mathbf{Q}$ is not ultimately periodic, then there is no decomposition for the sentence $\mathsf{S}x\, x{=}x$ [15, Theorem 8.5.2]. Here, we prove that if Q contains only ultimately periodic counting quantifiers, decompositions exist and can be computed:

▶ **Theorem 3.1.** *Let $\mathbf{Q} \subseteq \mathbf{U}$ and let σ be a relational signature. From a formula $\varphi(\overline{x}, \overline{y}) \in \mathrm{FO}(\mathbf{Q})[\sigma_2]$, one can construct a decomposition for φ w.r.t. $(\overline{x}; \overline{y})$.*

Proof. By Corollary 2.5(1), the logics $\mathrm{FO}(\mathbf{Q})$ and $\mathrm{FO}(\mathbf{D_Q})$ for $\mathbf{Q} \subseteq \mathbf{U}$ are effectively equally expressive. Therefore, it suffices to prove the theorem for sets of divisibility quantifiers, i.e., for the case where $\mathbf{Q} \subseteq \mathbf{D}$. The proof proceeds by induction on the construction of the formula $\varphi(\overline{x}, \overline{y})$. The cases of atomic formulas, Boolean combinations, and existential quantification are as in the first-order case, see e.g. [11, Lemma 2.3]. Here, we sketch the remaining case where $\varphi(\overline{x}, \overline{y})$ is of the form $\mathsf{D}_m z\, \psi(\overline{x}, \overline{y}, z)$ for some $m \geqslant 2$. Let $\overline{x} = (x_1, \ldots, x_k)$ and $\overline{y} = (y_1, \ldots, y_\ell)$. For $n \in \{0, 1, \ldots, m{-}1\}$, consider the formulas

$$\chi_n(\overline{x}, \overline{y}) := (\mathsf{D}_m{+}n)z\, \big(\, A(z) \wedge \psi(\overline{x}, \overline{y}, z)\, \big) \quad \text{and} \quad \xi_n(\overline{x}, \overline{y}) := (\mathsf{D}_m{+}n)z\, \big(\, B(z) \wedge \psi(\overline{x}, \overline{y}, z)\, \big).$$

Let $\varphi'(\overline{x}, \overline{y})$ be the disjunction of all formulas $\chi_{n_1}(\overline{x}, \overline{y}) \wedge \xi_{n_2}(\overline{x}, \overline{y})$ where $n_1, n_2 \in \{0, \ldots, m{-}1\}$ and $n_1 + n_2 \equiv 0 \pmod{m}$. Clearly, $\mathcal{A} \oplus \mathcal{B} \models (\varphi \leftrightarrow \varphi')(\overline{a}, \overline{b})$ holds for all disjoint structures $\mathcal{A}$ and $\mathcal{B}$ and all $\overline{a} \in A^k$ and $\overline{b} \in B^\ell$. Therefore, every decomposition of φ' is also a decomposition of φ. Furthermore, note that a decomposition for φ' can be computed from decompositions for the formulas $\chi_n(\overline{x}, \overline{y})$ and $\xi_n(\overline{x}, \overline{y})$ for $n \in \{0, \ldots, m{-}1\}$. All that remains to be done is to construct decompositions w.r.t. $(\overline{x}; \overline{y})$ for each of the formulas $\chi_n(\overline{x}, \overline{y})$ and $\xi_n(\overline{x}, \overline{y})$. By symmetry, we only consider the formula ξ_n.

By the induction hypothesis, there is a decomposition $\{(\alpha_i(\overline{x}), \beta_i(\overline{y}, z)) : i \in I\}$ of $\psi(\overline{x}, \overline{y}, z)$ w.r.t. $(\overline{x}; \overline{y}z)$. We can, w.l.o.g., assume that the formulas $\alpha_i(\overline{x})$ are mutually exclusive, i.e., $\alpha_i(\overline{x}) \wedge \alpha_j(\overline{x})$ is unsatisfiable for $i \neq j$. Then, the set $\big\{\, \big(\, \alpha_i(\overline{x})\, ,\, (\mathsf{D}_m{+}n)z\, \beta_i(\overline{y}, z)\, \big) \;:\; i \in I\, \big\}$ is a decomposition of $\xi_n(\overline{x}, \overline{y})$ w.r.t. $(\overline{x}; \overline{y})$. ◀

In the inductive proof of Theorem 3.1, the size of the decomposition (i.e., the number of pairs) increases exponentially with every negation and every quantification. It follows that the size of the formulas in the resulting decomposition can be bounded by a tower of 2s whose height is proportional to the size of the formula. We can adopt and simplify the proof of [4, Theorem 3] to also obtain a non-elementary lower bound:

▶ **Proposition 3.2.** *Let $\sigma = \{E\}$ with $\mathrm{ar}(E) = 2$. There is a sequence $(\varphi_h)_{h \geqslant 0}$ of $\mathrm{FO}[\sigma]$-sentences of size $\mathcal{O}(h)$ such that for every elementary function $f\colon \mathbb{N} \to \mathbb{N}$, there is $h \in \mathbb{N}$ such that every decomposition Δ_h in $\mathrm{FO}(\mathbf{U})$ of φ_h contains some sentence of length $> f(h)$.*

We finish this section with a corollary to Theorem 3.1 that will be used in the construction of Gaifman normal forms in the next section.

▶ **Corollary 3.3.** *Let $\mathbf{Q} \subseteq \mathbf{U}$ be a set of ultimately periodic quantifiers and let $r \in \mathbb{N}$. From an r-local formula $\lambda(\overline{x}, y) \in \mathrm{FO}(\mathbf{Q})$, one can compute a finite set Δ' of pairs of r-local $\mathrm{FO}(\mathbf{Q})$-formulas $(\alpha'(\overline{x}), \beta'(y))$ such that the following two formulas are equivalent:*

$$\mathrm{dist}(\overline{x}, y) > 2r{+}1 \;\wedge\; \lambda(\overline{x}, y) \qquad \textit{and} \qquad \mathrm{dist}(\overline{x}, y) > 2r{+}1 \;\wedge \bigvee_{(\alpha', \beta') \in \Delta'} \big(\alpha'(\overline{x}) \wedge \beta'(y)\big)\,.$$

4 Equivalent Gaifman normal forms

The main result of this section is:

▶ **Theorem 4.1.** *Let $\mathbf{Q}$ be a set of unary counting quantifiers.*

(1) *If $\mathbf{Q} \subseteq \mathbf{U}$ contains only ultimately periodic quantifiers then, from a formula $\varphi \in \mathrm{FO}(\mathbf{Q})$, one can compute an equivalent formula $\gamma \in \mathrm{FO}(\mathbf{Q})$ in Gaifman normal form.*

(2) *If* $\mathbf{Q}$ *contains a quantifier* S *that is not ultimately periodic, then there is a sentence* $\eta \in \mathrm{FO}(\mathbf{Q})$ *such that no* $\mathrm{FO}(\mathbf{Q})$*-sentence in Gaifman normal form is equivalent to* η.

The two statements of the theorem are proved in the next two subsections.

4.1 Ultimately periodic quantifiers

Proof of Theorem 4.1(1). In the light of Corollary 2.5, it suffices to consider sets $\mathbf{Q} \subseteq \mathbf{D}$ of divisibility quantifiers.

The construction of γ proceeds by structural induction following the construction of φ. The cases of atomic formulas φ as well as of Boolean combinations are trivial. If φ is of the form $\exists y\, \psi$, we can argue as in the first-order case ([10], see also [11, Sect. 4.1]), but since ψ is from $\mathrm{FO}(\mathbf{Q})$, we use Corollary 3.3 instead of Feferman-Vaught decompositions for FO (cf. [11, Lemma 2.3]). So it remains to consider the case where $\varphi(\overline{x})$ is of the form $\mathsf{D}_m y\, \psi(\overline{x}, y)$ for some $\mathsf{D_m} \in \mathbf{Q}$. By the induction hypothesis we can assume that ψ is in Gaifman normal form. Hence there are a finite set I, sentences χ_i in Gaifman normal form, and local formals $\lambda_i(\overline{x}, y)$ for all $i \in I$, such that φ is equivalent to the formula $\varphi' := \mathsf{D}_m\, y \Big(\bigvee_{i\in I} \big(\chi_i \wedge \lambda_i(\overline{x}, y) \big) \Big)$.

W.l.o.g. we can assume that the sentences χ_i are mutually exclusive, i.e., $\chi_i \wedge \chi_j$ is unsatisfiable for $i \neq j$.

Set $r' := 2r + 1$. Then, $\varphi(\overline{x})$ is equivalent to a Boolean combination of the formulas

$$\begin{aligned} \gamma_n(\overline{x}) \quad &:= \quad (\mathsf{D}_m{+}n)\, y\ \big(\ \mathrm{dist}(\overline{x}, y) \leqslant r' \ \wedge\ \bigvee_{i\in I} (\ \chi_i \ \wedge\ \lambda_i(\overline{x}, y)\ )\ \big) \qquad \text{and} \\ \delta_n(\overline{x}) \quad &:= \quad (\mathsf{D}_m{+}n)\, y\ \big(\ \mathrm{dist}(\overline{x}, y) > r' \ \wedge\ \bigvee_{i\in I} (\ \chi_i \ \wedge\ \lambda_i(\overline{x}, y)\ )\ \big) \end{aligned}$$

with $n \in \{0, \ldots, m{-}1\}$, and it suffices to transform each of these formulas into Gaifman normal form. Since the sentences χ_i are mutually exclusive, $\gamma_n(\overline{x})$ is equivalent to

$$\bigvee_{i\in I} \big(\ \chi_i \ \wedge\ (\mathsf{D}_m{+}n)\, y\ (\ \mathrm{dist}(\overline{x}, y) \leqslant r' \ \wedge\ \lambda_i(\overline{x}, y)\ )\ \big)\,,$$

which is in Gaifman normal form. Similarly, $\delta_n(\overline{x})$ is equivalent to

$$\bigvee_{i\in I} \big(\ \chi_i \ \wedge\ (\mathsf{D}_m{+}n)\, y\ (\ \mathrm{dist}(\overline{x}, y) > r' \ \wedge\ \lambda_i(\overline{x}, y)\ )\ \big)\,.$$

Let $i \in I$. By Corollary 3.3, we can construct a finite set J and r-local formulas $\alpha_j(\overline{x})$ and $\beta_j(y)$ for $j \in J$, such that the formulas

$$\mathrm{dist}(\overline{x}, y) > r' \ \wedge\ \lambda_i(\overline{x}, y) \qquad \text{and} \qquad \mathrm{dist}(\overline{x}, y) > r' \ \wedge\ \bigvee_{j\in J} \big(\ \alpha_j(\overline{x}) \ \wedge\ \beta_j(y)\ \big)$$

are equivalent. Again, we can assume that the formulas $\alpha_j(\overline{x})$ are mutually exclusive. Then, $\delta_n(\overline{x})$ is equivalent to the formula

$$\bigvee_{(i,j)\in I\times J} \big(\ \chi_i \ \wedge\ \alpha_j(\overline{x}) \ \wedge\ (\mathsf{D}_m{+}n)\, y\ (\ \mathrm{dist}(\overline{x}, y) > r' \ \wedge\ \beta_j(y)\ )\ \big)\,.$$

Finally, the formula $(\mathsf{D}_m{+}n)\, y \big(\mathrm{dist}(\overline{x}, y) > r' \wedge \beta_j(y) \big)$ is equivalent to a Boolean combination of formulas of the form $(\mathsf{D}_m{+}n_1)\, y \big(\mathrm{dist}(\overline{x}, y) \leqslant r' \wedge \beta_j(y) \big)$ and $(\mathsf{D}_m{+}n_2)\, y\, \beta_j(y)$ with $n_1, n_2 \in \{0, \ldots, m{-}1\}$. Note that the first formula is $(r' + r)$-local, and the second formula is a counting sentence, i.e., both these formulas are in Gaifman normal form. ◀

In the inductive proof of Theorem 4.1(1), the size of the Gaifman normal form increases at least exponentially with every (existential or counting) quantifier, since we build the disjunctive normal form to get the formula φ' defined at the beginning of the proof. Consequently, we only obtain a non-elementary upper bound on the size of the formula in Gaifman normal form. We can complement this with a non-elementary lower bound (the sequence of formulas is the same as in Proposition 3.2):

▶ **Proposition 4.2.** *Let $\sigma = \{E\}$ with $\mathrm{ar}(E) = 2$. There is a sequence of $\mathrm{FO}[\sigma]$-sentences $(\varphi_h)_{h\geqslant 0}$ of size $\mathcal{O}(h)$ such that for every elementary function $f\colon \mathbb{N} \to \mathbb{N}$, there is $h \in \mathbb{N}$ such that no $\mathrm{FO}(\mathbf{U})$-sentence in Gaifman normal form of length $< f(h)$ is equivalent to φ_h.*

4.2 Non-ultimately periodic quantifiers

Theorem 4.1(2) is an immediate consequence of the following slightly stronger result.

▶ **Proposition 4.3.** *Let $\mathbf{Q}$ be a set of counting quantifiers, let $\mathsf{S} \in \mathbf{Q}$, let $\sigma_\sim = \{\sim\}$ be the signature with $\mathrm{ar}(\sim) = 2$, and let η be the $\mathrm{FO}(\mathbf{Q})$-sentence*

$$\exists x_1\, \exists x_2 \left(\begin{array}{ll} & \neg\,(x_1 \sim x_2) \;\wedge\; \mathsf{S}z\,(x_1 \sim z \,\vee\, x_2 \sim z) \\ \wedge & \forall y\,\big(\,\mathsf{S}z\,(\,x_1 \sim z \,\vee\, y \sim z\,) \,\vee\, \mathsf{S}z\,(\,x_2 \sim z \,\vee\, y \sim z\,)\,\big) \end{array} \right).$$

If S is not ultimately periodic, then no sentence from $\mathrm{FO}(\mathbf{Q})[\sigma_\sim]$ in Gaifman normal form is equivalent to η (not even on the class of finite equivalence structures).

A *finite equivalence structure* is a $\sigma_\sim$-structure $\mathcal{A} = (A, \sim^{\mathcal{A}})$ where A is finite and $\sim^{\mathcal{A}}$ is an equivalence relation on A. In such structures, the formula η expresses that there are two distinct equivalence classes $[a_1]$ and $[a_2]$ with $|[a_1] \cup [a_2]| \in \mathsf{S}$ such that $|[a_1] \cup [b]| \in \mathsf{S}$ or $|[a_2] \cup [b]| \in \mathsf{S}$ for any equivalence class $[b]$. We will prove that this property cannot be expressed by any sentence in Gaifman normal form.

Let $\lambda(x)$ be a local $\mathrm{FO}(\mathbf{Q})[\sigma_\sim]$-formula, let $\mathcal{A}$ be a finite equivalence structure, and let $a \in A$. Since λ is local, the question whether or not $\mathcal{A} \models \lambda(a)$ is determined solely by the size of the equivalence class $[a]$. In case that $\lambda(x)$ is a first-order formula, it is equivalent to a Boolean combination of statements of the form $|[a]| \geqslant k$ for $k \in \{0, \ldots, r\}$, where r is the formula's quantifier rank. If the formula $\lambda(x)$ also uses counting quantifiers from $\mathbf{Q}$, then we can also express that $|[a]| - \ell \in \mathsf{Q}$, for $\mathsf{Q} \in \mathbf{Q}$; but this is only possible for $\ell \in \{0, \ldots, r'\}$, where r' is a number that only depends on the formula λ, but not on the considered equivalence structure $\mathcal{A}$. These observations lead to the following limitation of the expressiveness of local formulas:

▶ **Lemma 4.4.** *For every local $\mathrm{FO}(\mathbf{Q})[\sigma_\sim]$-formula $\lambda(x)$ there exists an $r \in \mathbb{N}$ such that the following is true for all finite equivalence structures $\mathcal{A}$ and $\mathcal{B}$, all $a \in A$, and all $b \in B$.*
If the equivalences

$$|[a]| \geqslant \ell \iff |[b]| \geqslant \ell \qquad \text{and} \qquad |[a]| - \ell \in \mathsf{Q} \iff |[b]| - \ell \in \mathsf{Q}$$

are satisfied for all $\ell \in \{0, \ldots, r\}$ and all counting quantifiers $\mathsf{Q} \in \mathbf{Q}$ that appear in λ, then $\mathcal{A} \models \lambda(a) \iff \mathcal{B} \models \lambda(b)$.

Proof of Proposition 4.3. Suppose S is not ultimately periodic and suppose, for contradiction, that γ is an $\mathrm{FO}(\mathbf{Q})[\sigma_\sim]$-sentence in Gaifman normal form that is equivalent to η. Let $\mathbf{Q}_\gamma$ consist of all counting quantifiers that appear in γ. We apply Lemma 4.4 to all local formulas $\lambda(x)$ that occur in basic-local sentences or in counting sentences which are subformulas of γ, and we let $\hat{r}$ be the maximum of all the numbers r provided by Lemma 4.4 for each such $\lambda(x)$.

Since $\mathsf{Q}_\gamma \cup \{\mathsf{S}\}$ is finite and S is not ultimately periodic, one can prove the existence of natural numbers $m, n \geqslant \hat{r}$ and $k \geqslant 1$ such that $m+k \notin \mathsf{S}$, $n+k \in \mathsf{S}$, and for all $\mathsf{Q} \in \mathbf{Q}_\gamma \cup \{\mathsf{S}\}$ and all $\ell \in \{0, \ldots, \hat{r}\}$ we have $m - \ell \in \mathsf{Q} \iff n - \ell \in \mathsf{Q}$.

We use the numbers $m, n, k, \hat{r}$ to define two finite equivalence structures $\mathcal{B}$ and $\mathcal{C}$:

- $\mathcal{B}$ has $\hat{r}$ many equivalence classes of size k and m many equivalence classes of size n.
- $\mathcal{C}$ has $\hat{r}$ many equivalence classes of size k and n many equivalence classes of size m.

By Lemma 4.4, no local formula appearing in γ can distinguish equivalence classes of size m from equivalence classes of size n. It follows that $\mathcal{B} \models \gamma \iff \mathcal{C} \models \gamma$.

To prove that $\mathcal{B} \models \eta$, one chooses for x_1 and x_2 elements from equivalence classes of size k and n, respectively. When trying to satisfy η in $\mathcal{C}$, one has to choose for x_1 and x_2 elements from equivalence classes of the same size since $m+k \notin \mathsf{S}$. But there is some y whose equivalence class has different size, and hence $|[x_i]| + |[y]|$ cannot belong to S. Thus, η distinguishes $\mathcal{B}$ from $\mathcal{C}$, but γ does not. ◀

5 Finitely d-equivalent Gaifman normal forms

5.1 Ultimately periodic quantifiers

In Section 4.1 we obtained an algorithm that transforms a given FO($\mathbf{U}$)-formula over a relational signature into an equivalent FO($\mathbf{U}$)-formula in Gaifman normal form. Just as in Gaifman's original locality theorem, the algorithm's runtime is non-elementary in the size of the input formula; and from Proposition 4.2 we know that a non-elementary blow-up in formula size (and hence also runtime) cannot be avoided.

In [16] it was shown that for plain first-order logic FO, the non-elementary blow-up can be improved into a (worst-case optimal) 3-fold exponential running time if we drop the requirement that the Gaifman normal form formula has to be equivalent to the original formula on *all* structures and are content with a *finitely d-equivalent* formula in Gaifman normal form. We can generalise this result to FO($\mathbf{U}$) as follows.

▶ **Theorem 5.1.** *Upon input of a number $d \in \mathbb{N}$ and an* FO($\mathbf{U}$)*-formula φ over some relational signature σ, a finitely d-equivalent formula ψ in Gaifman normal form can be computed in time $2^{d^{2^{O(\|\varphi\|)}}}$ for $d \geqslant 3$, and in time $2^{2^{\mathrm{poly}(\|\varphi\|)}}$ for $d < 3$.*
Furthermore, ψ uses at most the quantifiers from φ and the quantifier $\exists$.

We proceed in the same way as the proof of [16], but instead of building upon the Hanf normal form algorithm for FO of [2] we build upon the Hanf normal form algorithm for FO($\mathbf{U}$) of [17]. For the precise statement of the result of [17], we need the following notation.

Let σ be a relational signature. For every $r \in \mathbb{N}$ and $n \in \mathbb{N}_{\geqslant 1}$, a *type with n centres and radius at most r* is structure of the form $\tau = (\mathcal{N}_r^{\mathcal{A}}(\overline{a}), \overline{a})$ where $\mathcal{A}$ is a σ-structure and $\overline{a} \in A^n$. Such a type is called d-bounded if the structure $\mathcal{N}_r^{\mathcal{A}}(\overline{a})$ is d-bounded.

The following is straightforward. The universe of a d-bounded type τ with n centres and radius $\leqslant r$ has size at most $n \cdot d^{r+1}$ (provided that $d \geqslant 2$). Given τ and r, one can construct an FO[σ]-formula[2] $\mathrm{sph}_\tau(\overline{x})$ with n free variables $\overline{x} = (x_1, \ldots, x_n)$ such that for every σ-structure $\mathcal{A}$ and every tuple $\overline{a} \in A^n$ we have $\mathcal{A} \models \mathrm{sph}_\tau(\overline{a}) \iff (\mathcal{N}_r^{\mathcal{A}}(\overline{a}), \overline{a}) \cong \tau$. We

[2] The formula $\mathrm{sph}_\tau(\overline{x})$ also depends on r, although this is not reflected by the notation here.

can assume w.l.o.g. that the formula $\mathrm{sph}_\tau(\overline{x})$ is r-local and has size at most $(n{\cdot}d^{r+1})^{\mathcal{O}(\|\sigma\|)}$, where $\|\sigma\|$ is defined as the sum of the arities of the relation symbols in σ.

Formulas of the form $\mathrm{sph}_\tau(\overline{x})$ are called (d-bounded) *sphere-formulas* of signature σ. Let $\mathbf{Q}$ be a set of counting quantifiers. An $\mathrm{FO}(\mathbf{Q})$-*Hanf-sentence* of signature σ is a sentence of the form $(\mathsf{Q}{+}k)y\,\mathrm{sph}_\rho(y)$ or of the form $\exists^{\geqslant k}y\,\mathrm{sph}_\rho(y)$, where $k \in \mathbb{N}$, $\mathsf{Q} \in \mathbf{Q}$, and ρ is a type of signature σ and with a single centre. An $\mathrm{FO}(\mathbf{Q})$-formula in *Hanf normal form* and of signature σ is a Boolean combination of sphere-formulas and $\mathrm{FO}(\mathbf{Q})$-Hanf-sentences of signature σ. The proof of Theorem 5.1 follows by combining Lemma 2.4(c) and:

▶ **Theorem 5.2** ([17]). *Upon input of a number $d \in \mathbb{N}$ and an $\mathrm{FO}(\mathbf{U})$-formula φ over some relational signature σ, a finitely d-equivalent $\mathrm{FO}(\mathbf{U})$-formula ψ in Hanf normal form and of signature σ can be computed in time $2^{d^{2^{O(\|\varphi\|)}}}$ for $d \geqslant 3$, and in time $2^{2^{\mathrm{poly}(\|\varphi\|)}}$ for $d < 3$.*

5.2 General quantifiers

In Section 4.2 we showed that if a set $\mathbf{Q}$ contains a quantifier that is *not* ultimately periodic, then there is an $\mathrm{FO}(\mathbf{Q})$-sentence that is not equivalent to any $\mathrm{FO}(\mathbf{Q})$-sentence in Gaifman normal form (not even on the class of finite structures). Somewhat surprisingly, it turns out that if we drop the requirement that the Gaifman normal form formula has to be equivalent to the original formula on *all* structures and are content with a *finitely d-equivalent* formula, Gaifman normal forms do exist for *arbitrary* sets $\mathbf{Q}$ of counting quantifiers. Precisely, we obtain the following result, in which the size $\|\varphi\|$ of an $\mathrm{FO}(\mathbf{Q})$-formula φ of signature σ is defined analogously as the size of $\mathrm{FO}(\mathbf{U})$-formulas, but now each quantifier $\mathsf{Q} \in \mathbf{Q}$ is viewed as an abstract symbol of length 1.

▶ **Theorem 5.3.** *Let $\mathbf{Q}$ be an arbitrary set of counting quantifiers and let $d \in \mathbb{N}$. For every $\mathrm{FO}(\mathbf{Q})$-formula φ over some relational signature σ, there exists a finitely d-equivalent $\mathrm{FO}(\mathbf{Q})$-formula ψ in Gaifman normal form. Moreover, if the sets $\mathsf{Q} \in \mathbf{Q}$ are uniformly decidable (in elementary time), then ψ can be computed from φ and d (in elementary time).*

The proof proceeds in a similar way as the proof of Theorem 5.1, but instead of building upon Theorem 5.2, it uses a result of [19] that can be viewed as a generalisation of Theorem 5.2 to $\mathrm{FO}(\mathbf{Q})$ for arbitrary sets $\mathbf{Q}$ of unary counting quantifiers. For the precise statement of this result, we need the following notation. An $\mathrm{FO}(\mathbf{Q})$-*weak-Hanf-sentence* of signature σ is a sentence of the form $(\mathsf{Q}{+}k)y \bigvee_{\theta\in T}\theta(y)$ or of the form $\exists^{\geqslant k}y \bigvee_{\theta\in T}\theta(y)$, where T is a finite set of sphere-formulas of signature σ, each of them with a single centre and all of the same radius r. An $\mathrm{FO}(\mathbf{Q})$-formula in *weak Hanf normal form* and of signature σ is a Boolean combination of sphere-formulas and $\mathrm{FO}(\mathbf{Q})$-weak-Hanf-sentences of signature σ. The proof of Theorem 5.3 follows by combining Lemma 2.4(c) and:

▶ **Theorem 5.4** ([19]). *Let $\mathbf{Q}$ be an arbitrary set of counting quantifiers and let $d \in \mathbb{N}$. For every $\mathrm{FO}(\mathbf{Q})$-formula φ over some relational signature σ, there exists a finitely d-equivalent $\mathrm{FO}(\mathbf{Q})$-formula ψ in weak Hanf normal form and of signature σ. Moreover, if the sets $\mathsf{Q} \in \mathbf{Q}$ are uniformly decidable (in elementary time), then ψ can be computed from φ and d (in elementary time).*

One may wonder if, analogously to the statement of Theorem 5.1, the last statement of Theorem 5.3 can be improved to a 3-fold exponential running time. To refute this, one observes that from a formula in Gaifman normal form, one can construct an equivalent formula in weak Hanf normal form with the same number of counting sentences. Then a lower bound result of [19] for weak Hanf normal forms implies the following:

▶ **Proposition 5.5.** *There exists a* $\mathsf{Q} \subseteq \mathbb{N}$ *such that for* $\mathbf{Q} := \{\mathsf{Q}\}$*, there is a sequence* $(\varphi_n)_{n\geqslant 1}$ *of* $\mathrm{FO}(\mathbf{Q})$*-sentences of the same relational signature and of size* $\mathcal{O}(n)$ *such that, for all* $n \geqslant 1$*, every* $\mathrm{FO}(\mathbf{Q})$*-sentence in Gaifman normal form that is finitely 3-equivalent to* φ_n *contains at least* $\exp_4(n)$ *distinct subformulas of the form* $(\mathsf{Q}{+}k)y\,\lambda(y)$*.*

From our proof it follows that $\{n^n : n \in \mathbb{N}\}$, $\{n! : n \in \mathbb{N}\}$, and $\{\lfloor 2^{n^c} \rfloor : n \in \mathbb{N}\}$ for all reals $c > 1$, are examples of sets Q for which the statement of Proposition 5.5 holds.

6 An algorithmic meta-theorem for $\mathrm{FO}(\mathbf{U})$

The model-checking problem for a logic L and a class $\mathcal{C}$ of finite relational structures receives as input a sentence $\varphi \in \mathrm{L}$ and a structure $\mathcal{A} \in \mathcal{C}$, and the task is to decide if $\mathcal{A} \models \varphi$. This problem is said to be *fixed-parameter tractable* if it can be solved in time $f(\|\varphi\|)\cdot\mathrm{poly}(\|\mathcal{A}\|)$ where f is a computable function, $\|\varphi\|$ is the size of the formula, and $\|\mathcal{A}\|$ is the size of the structure (defined as $\|\mathcal{A}\| := |A| + \sum_{R\in\sigma} \mathrm{ar}(R)\cdot|R^{\mathcal{A}}|$). Recall from Section 1 the list of examples of logics L and classes $\mathcal{C}$ for which the model-checking problem is known to be fixed-parameter tractable. The aim of this section is to demonstrate that by using our Gaifman normal form result for $\mathrm{FO}(\mathbf{U})$ (Theorem 4.1(1)), the model-checking algorithm for classes of bounded local tree-width of [9] can be generalised from FO to $\mathrm{FO}(\mathbf{U})$.

To provide a precise formulation of the result, we need some further notation. We assume that the reader is familiar with the basic concept of a tree-decomposition and the tree-width $\mathrm{tw}(\mathcal{A})$ of a structure $\mathcal{A}$ (precise definitions can be found in [9] and will not be necessary for understanding the remainder of this section). The *local tree-width* of $\mathcal{A}$ is the function $\mathrm{ltw}^{\mathcal{A}} : \mathbb{N} \to \mathbb{N}$ defined by $\mathrm{ltw}^{\mathcal{A}}(r) := \max\{\mathrm{tw}(\mathcal{N}_r^{\mathcal{A}}(a)) : a \in A\}$ for all $r \in \mathbb{N}$. A class $\mathcal{C}$ of structures has *(effectively) bounded local tree-width* if there is a (computable) function $g : \mathbb{N} \to \mathbb{N}$ such that $\mathrm{ltw}^{\mathcal{A}}(r) \leqslant g(r)$ for all $\mathcal{A} \in \mathcal{C}$ and all $r \in \mathbb{N}$. As shown in [9], examples for classes of bounded local tree-width are classes of trees, classes of structures of tree-width at most w (for each fixed $w \in \mathbb{N}$), classes of degree at most d (for each fixed $d \in \mathbb{N}$), the class of planar graphs, and classes of graphs of genus at most g (for each fixed $g \in \mathbb{N}$).

The overall approach of [9] has been described in [12] as follows: *"Using Gaifman's theorem, the problem to decide whether a general first-order formula* φ *is true in a graph can be reduced to testing whether a formula is true in* r*-neighbourhoods in the graph, where the radius* r *only depends on* φ*, and solving a variant of the (distance d) independent set problem. Hence, if* $\mathcal{C}$ *is a class of graphs where* r*-neighbourhoods have a simple structure, such as the class of planar graphs or classes of bounded local tree-width, this method gives an easy way for deciding properties definable in first-order logic."*

Here, the "(distance d) independent set problem" corresponds to the essence of evaluating a basic local sentence. Our Gaifman normal form for $\mathrm{FO}(\mathbf{U})$-sentences consists of basic local sentences (which can be evaluated in the same way as described in [9]) and counting sentences of the form $(\mathsf{Q}{+}k)x\,\lambda(x)$, and evaluating these boils down to (1) computing the set of all nodes x whose r-neighbourhood satisfies $\lambda(x)$ and (2) checking if the size of this set belongs to $(\mathsf{Q}{+}k)$. The task (1) has been solved in [9] for r-local FO-formulas and can easily be gerneralised to r-local $\mathrm{FO}(\mathbf{U})$-formulas, and the task (2) is straightforward. In summary, by combining the approach of [9] with our Theorem 4.1(1) we obtain:

▶ **Corollary 6.1.** *Let* $\mathcal{C}$ *be a class of finite relational structures of bounded local tree-width and let* φ *be an* $\mathrm{FO}(\mathbf{U})$*-sentence. Then, for every* $k \geqslant 1$*, there is an algorithm deciding in time* $\mathcal{O}(\|A\|^{1+(1/k)})$ *whether a given structure* $\mathcal{A} \in \mathcal{C}$ *satisfies* φ*.*

To keep the runtime analysis of Corollary 6.1 simple, we formulated the corollary in a way which uses the $\mathcal{O}$-notation to hide factors that depend on the sentence φ or the number k. A closer inspection of the proof shows that, for any class of effectively bounded local tree-width, the algorithm's runtime can be bounded by $f(\|\varphi\|, k) \cdot \|\mathcal{A}\|^{1+(1/k)}$, for some computable function f. Thus, in particular, we obtain that for every class $\mathcal{C}$ of effectively bounded local tree-width, the model-checking problem for FO($\mathbf{U}$)-sentences on $\mathcal{C}$ is fixed-parameter tractable. To close this paper, let us mention that we believe that by a similar, but substantially more involved construction also the result of [12] for model-checking on nowhere dense classes can be lifted from FO to FO($\mathbf{U}$) – we plan to do this as future work.

References

1 S. Arnborg, J. Lagergren, and D. Seese. Easy problems for tree-decomposable graphs. *J. Algorithms*, 12(2):308–340, 1991. `doi:10.1016/0196-6774(91)90006-K`.

2 B. Bollig and D. Kuske. An optimal construction of Hanf sentences. *J. Applied Logic*, 10(2):179–186, 2012. `doi:10.1016/j.jal.2012.01.002`.

3 B. Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990. `doi:10.1016/0890-5401(90)90043-H`.

4 A. Dawar, M. Grohe, S. Kreutzer, and N. Schweikardt. Model theory makes formulas large. In *Proceedings of the 34th International Colloquium on Automata, Languages and Programming, ICALP 2007, Wroclaw, Poland, July 9-13, 2007*, pages 913–924, 2007. `doi:10.1007/978-3-540-73420-8_78`.

5 J. Doner. Tree acceptors and some of their applications. *J. Comput. Syst. Sci.*, 4(5):406–451, 1970. `doi:10.1016/S0022-0000(70)80041-1`.

6 Z. Dvorak, D. Král, and R. Thomas. Deciding first-order properties for sparse graphs. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 133–142, 2010. `doi:10.1109/FOCS.2010.20`.

7 H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Perspectives in Mathematical Logic. Springer, 1995.

8 S. Feferman and R. L. Vaught. The first-order properties of products of algebraic systems. *Fundamenta Mathematicae*, 47:57–103, 1959.

9 M. Frick and M. Grohe. Deciding first-order properties of locally tree-decomposable structures. *J. ACM*, 48(6):1184–1206, 2001. `doi:10.1145/504794.504798`.

10 H. Gaifman. On local and nonlocal properties. In J. Stern, editor, *Logic Colloquium '81*, pages 105–135. North-Holland, 1982.

11 M. Grohe. Logic, graphs, and algorithms. In *Logic and Automata: History and Perspectives [in Honor of Wolfgang Thomas]*, volume 2 of *Texts in Logic and Games*, pages 357–422. Amsterdam University Press, 2008.

12 M. Grohe, S. Kreutzer, and S. Siebertz. Deciding first-order properties of nowhere dense graphs. *J. ACM*, 64(3):17:1–17:32, 2017. `doi:10.1145/3051095`.

13 F. Harwath, L. Heimberg, and N. Schweikardt. Preservation and decomposition theorems for bounded degree structures. *Logical Methods in Computer Science*, 11(4), 2015. `doi:10.2168/LMCS-11(4:17)2015`.

14 F. Harwath and N. Schweikardt. On the locality of arb-invariant first-order formulas with modulo counting quantifiers. *Logical Methods in Computer Science*, 12(4), 2016. `doi:10.2168/LMCS-12(4:8)2016`.

15 L. Heimberg. *Complexity of normal forms on structures of bounded degree*. PhD thesis, Humboldt-Universität zu Berlin, 2017.

16 L. Heimberg, D. Kuske, and N. Schweikardt. An optimal Gaifman normal form construction for structures of bounded degree. In *Proceedings of the 28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2013, New Orleans, LA, USA, June 25-28, 2013*, pages 63–72, 2013. doi:10.1109/LICS.2013.11.

17 L. Heimberg, D. Kuske, and N. Schweikardt. Hanf normal form for first-order logic with unary counting quantifiers. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2016, New York, NY, USA, July 5-8, 2016*, pages 277–286, 2016. doi:10.1145/2933575.2934571.

18 L. Hella, L. Libkin, and J. Nurmonen. Notions of locality and their logical characterizations over finite models. *J. Symb. Log.*, 64(4):1751–1773, 1999. doi:10.2307/2586810.

19 D. Kuske and N. Schweikardt. First-order logic with counting. In *Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017*, pages 1–12, 2017. (Full version available at https://arxiv.org/abs/1703.01122). doi:10.1109/LICS.2017.8005133.

20 L. Libkin. *Elements of Finite Model Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2004. doi:10.1007/978-3-662-07003-1.

21 J. A. Makowsky. Algorithmic uses of the Feferman-Vaught Theorem. *Ann. Pure Appl. Logic*, 126(1-3):159–213, 2004. doi:10.1016/j.apal.2003.11.002.

22 A. B. Matos. Periodic sets of integers. *Theor. Comput. Sci.*, 127(2):287–312, 1994. doi:10.1016/0304-3975(94)90044-2.

23 H. Niemistö. *Locality and Order-Invariant Logics*. PhD thesis, Department of Mathematics and Statistics, University of Helsinki, 2007.

24 J. Nurmonen. Counting modulo quantifiers on finite structures. *Inf. Comput.*, 160(1-2):62–87, 2000. doi:10.1006/inco.1999.2842.

25 D. Seese. Linear time computable problems and first-order descriptions. *Mathematical Structures in Computer Science*, 6(6):505–526, 1996.

26 H. Straubing. *Finite automata, formal logic, and circuit complexity*. Progress in Theoretical Computer Science. Birkhäuser Boston Inc., Boston, MA, 1994.

27 J. W. Thatcher and J. B. Wright. Generalized finite automata theory with an application to a decision problem of second-order logic. *Mathematical Systems Theory*, 2(1):57–81, 1968. doi:10.1007/BF01691346.