Hanf normal form for first-order logic with unary counting quantifiers☆

Lucas Heimberg[a], Dietrich Kuske[b], Nicole Schweikardt[a]

[a] *Humboldt-Universität zu Berlin, Institut für Informatik, Unter den Linden 6, D-10099 Berlin*
[b] *Technische Universität Ilmenau, Fachgebiet Automaten und Logik, Postfach 100565, 98684 Ilmenau*

Abstract

We study the existence and construction of Hanf normal forms for extensions $\mathrm{FO}(\mathbf{Q})$ of first-order logic by some set $\mathbf{Q}$ of unary counting quantifiers $\mathsf{Q} \subseteq \mathbb{N}$. A formula is in Hanf normal form if it is a Boolean combination of (i) formulas describing the isomorphism type of a local neighbourhood around its free variables and (ii) statements of the form "the number of witnesses y of $\psi(y)$ belongs to $(\mathsf{Q}{+}k)$" where $\mathsf{Q} \in \mathbf{Q}$, $k \in \mathbb{N}$, and ψ describes the isomorphism type of a local neighbourhood around its unique free variable y.

We show that all formulas from $\mathrm{FO}(\mathbf{Q})$ permit a formula in Hanf normal form that is equivalent on all structures of degree $\leqslant d$ if, and only if, all counting quantifiers from $\mathbf{Q}$ are ultimately periodic. Furthermore, for such sets $\mathbf{Q}$ of ultimately periodic counting quantifiers, we present a worst-case optimal, 3-fold exponential time algorithm which upon input of a degree bound d and an $\mathrm{FO}(\mathbf{Q})$-formula φ produces an according formula in Hanf normal form.

In particular, this yields an algorithmic version of Nurmonen's extension of Hanf's theorem for first-order logic with modulo-counting quantifiers. As an immediate consequence, we obtain that on finite structures of degree $\leqslant d$, model checking of first-order logic with modulo-counting quantifiers and, in general, first-order logic with ultimately periodic quantifiers, is fixed-parameter tractable.

Keywords: Hanf locality, normal forms, modulo-counting quantifiers, ultimately periodic sets, structures of bounded degree, model checking
2010 MSC: 03B10, 03C07, 03C13, 03C80, 03C98

☆This is the full version of the conference contribution [12].

Email addresses: `heimberl@informatik.hu-berlin.de` (Lucas Heimberg), `dietrich.kuske@tu-ilmenau.de` (Dietrich Kuske), `schweikn@informatik.hu-berlin.de` (Nicole Schweikardt)

1. Introduction

Two elements of a given graph are indistinguishable by first-order formulas whenever some automorphism maps the first to the second. More generally, if the two elements have isomorphic neighbourhoods of radius 4^q, then they cannot be distinguished by first-order formulas of quantifier depth q. This and similar phenomena are summarised under the slogan "first-order logic can only express local properties" and formalised by the theorems by Hanf, by Gaifman, and by Schwentick and Barthelmann [11, 6, 10, 23]. All these results give rise to normal forms for first-order formulas. Hanf's and Gaifman's theorem have found various applications in algorithms and complexity (cf., e.g., [15, 17]). In particular, there are very general algorithmic meta-theorems stating that first-order model checking is fixed-parameter tractable for various classes of structures [24, 8], and that the results of first-order queries against various classes of databases can be enumerated with constant delay after a linear-time preprocessing phase [4, 14, 25]. In the context of such algorithms, questions about the efficiency of the normal forms have recently attracted interest (cf. e.g., [3, 18, 2]).

Notions of locality have also been developed for extensions of first-order logic, and they have found application in proving inexpressibility results for these logics (cf., e.g., [13, 22, 17]). When restricting attention to classes of finite structures of bounded degree, these locality notions also give rise to normal forms for the respective logics. Let us focus on the particular case of Hanf-locality:

Hanf's locality theorem for first-order logic implies that for every first-order sentence φ over a relational signature σ, and for every degree bound $d \in \mathbb{N}$, there exists a first-order sentence ψ that is equivalent to φ on all σ-structures of degree $\leqslant d$, such that ψ is a Boolean combination of statements of the form "there are $\geqslant k$ elements whose r-neighbourhood has isomorphism type τ". Such a sentence ψ is said to be in *Hanf normal form*. A worst-case optimal algorithm for constructing ψ when given φ and d has been developed in [2].

In [22], Nurmonen extended Hanf's locality theorem to the extension of first-order logic by modulo-counting quantifiers D_p (for positive integers p), where a formula of the form $\mathsf{D}_p\, y\, \psi(\overline{x}, y)$ states that the number of witnesses y for $\psi(\overline{x}, y)$ is divisible by p. As an easy consequence of Nurmonen's theorem, one obtains that for every sentence φ of first-order logic with modulo-counting quantifiers, and for every degree bound $d \in \mathbb{N}$ there exists a first-order sentence with modulo-counting quantifiers ψ that is equivalent to φ on all finite structures of degree $\leqslant d$, such that ψ is a Boolean combination of statements of the form "the number of elements whose r-neighbourhood has isomorphism type τ is congruent k modulo p" and statements of the form "there are $\geqslant k$ elements whose r-neighbourhood has isomorphism type τ". Again, we say that ψ is in *Hanf normal form*.

For algorithmic applications, an effective procedure for computing ψ when given φ and d would be highly desirable (cf., e.g., the use of Nurmonen's theorem in the full version of [21]). The proof of [22], however, does *not* lead to such an effective procedure. The following two questions started the research whose results are presented in this paper.

(1) Is there an algorithmic version of Nurmonen's result?

(2) For which classes of unary counting quantifiers does an analogue of Nurmonen's result hold?

Answering question (2), our first main result provides a precise characterisation: A class $\mathbf{Q}$ of unary counting quantifiers permits "*Hanf normal forms*" (analogous to the ones obtained from Nurmonen's result) if, and only if, all counting quantifiers in $\mathbf{Q}$ are ultimately periodic.

Answering question (1), our second main result provides an algorithm which, when given a degree bound d and a formula φ of the extension of first-order logic with ultimately periodic unary counting quantifiers, transforms φ into a corresponding "Hanf normal form" which is equivalent to φ on all structures of degree $\leqslant d$. This algorithm uses 3-fold exponential time for $d \geqslant 3$ and 2-fold exponential time for $d = 2$, and is worst-case optimal in both cases. As an easy application of our algorithm, we obtain that Seese's [24] fixed-parameter tractability result for the data complexity of first-order model checking (where the degree serves as parameter) can be generalised to first-order logic with ultimately periodic unary counting quantifiers.

The rest of the paper is structured as follows. Section 2 fixes basic notations used throughout the paper. Section 3 gives precise statements of our two main results and their consequence regarding fixed-parameter tractability of model checking. Section 4 and Section 5 are devoted to the proof of the "only if"-direction and the "if"-direction, respectively, of our characterisation of the sets of unary counting quantifiers that permit Hanf normal forms. In Section 5, we also provide a special case of our construction for plain first-order logic. I.e., we adapt the algorithm from [2] to our setting. Section 6 contains the runtime analysis of our algorithm for transforming a given formula into Hanf normal form. Section 7 shows how to use our algorithm to achieve fixed-parameter tractability of the model checking problem for first-order logic with ultimately periodic unary counting quantifiers. Section 8 provides matching lower bounds for the construction of Hanf normal forms. Section 9 concludes the paper.

2. Preliminaries

2.1. Basic notations

We write $\mathcal{P}(S)$ to denote the power set of a set S. By $\mathbb{N}$ we denote the set of non-negative integers, and by $\mathbb{N}_{\geqslant 1} := \mathbb{N} \setminus \{0\}$ we denote the set of all positive integers. For all $m, n \in \mathbb{N}$ with $m \leqslant n$, we write $[m, n]$ for the set $\{i \in \mathbb{N} : m \leqslant i \leqslant n\}$, and we let $[m, n) := [m, n] \setminus \{n\}$. By $\mathbb{R}_{\geqslant 0}$ we denote the set of non-negative reals. For a real number $r > 0$, we write $\log(r)$ to denote the logarithm of r with respect to base 2.

For every function $f \colon \mathbb{N} \to \mathbb{R}_{\geqslant 0}$, we write $\text{poly}(f(n))$ for the class of all functions $g \colon \mathbb{N} \to \mathbb{R}_{\geqslant 0}$ for which there exists a number $c > 0$ such that $g(n) \leqslant (f(n))^c$ for all sufficiently large $n \in \mathbb{N}$.

A function $f\colon \mathbb{N} \to \mathbb{R}_{\geqslant 0}$ is *at most k-fold exponential*, for some $k \in \mathbb{N}$, if we have $f(n) \in T(k, \mathrm{poly}(n))$, where $T(0, m) := m$ and $T(k{+}1, m) := 2^{T(k,m)}$ for all $k, m \geqslant 0$. I.e.,

$$f(n) \quad \in \quad 2^{2^{\cdot^{\cdot^{\cdot^{2^{\mathrm{poly}(n)}}}}}} \quad \Big\} \text{ a tower of 2s of height } k \text{ with } \mathrm{poly}(n) \text{ on top.}$$

For a finite word $w \in \{0,1\}^*$, we write $|w|$ to denote the length of w. For an ω-word $w = w_0 w_1 w_2 \cdots \in \{0,1\}^\omega$ and a number $n \in \mathbb{N}$, we write $w[n]$ to denote the letter w_n in w at position n. For numbers $i, j \in \mathbb{N}$ with $i \leqslant j$, we write $w[i,j]$ for the (finite) word $w_i w_{i+1} \cdots w_j$. Similarly, we write $w(i,j]$ for the (finite) word $w_{i+1} \cdots w_j$. In particular, $w(i,i]$ is the empty word ϵ, and $w(j{-}1, j] = w[j]$.

A finite word v is *primitive* if $v \in u^*$ for a word u implies that $v = u$. It is well-known that every finite non-empty word w can be written as v^n for some primitive word v and some $n \geqslant 1$, and this primitive word v is uniquely determined by w and is called the *primitive root* of w [19].

2.2. 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. The *size* $\|\sigma\|$ of a signature σ is the number of its constant symbols plus the sum of the arities of its relation symbols. We call a signature *relational* if it only contains relation symbols. A *σ-structure* $\mathcal{A}$ consists of a *finite* 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 these definitions, all signatures and all structures considered in this paper are *finite*. To indicate that two σ-structures $\mathcal{A}$ and $\mathcal{B}$ are isomorphic, we write $\mathcal{A} \cong \mathcal{B}$.

We use the standard notation concerning first-order logic and extensions thereof, cf. [5, 17]. By $\mathrm{FO}[\sigma]$ we denote the class of all first-order formulas of signature σ. That is, $\mathrm{FO}[\sigma]$ is built from atomic formulas of the form $x_1{=}x_2$ and $R(x_1, \ldots, x_{\mathrm{ar}(R)})$, for $R \in \sigma$ and variables or constant symbols $x_1, x_2, \ldots, x_{\mathrm{ar}(R)}$, and closed under Boolean connectives $\neg, \vee$ and existential first-order quantifiers $\exists x$ for any variable x.[1] By FO we denote the union of all $\mathrm{FO}[\sigma]$ for arbitrary signatures σ.

The size $\|\varphi\|$ of an $\mathrm{FO}[\sigma]$-formula is its length when viewed as a word over the alphabet $\sigma \cup \mathrm{Var} \cup \{,\} \cup \{=, \exists, \neg, \vee, (,)\}$, where Var is a countable set of variable symbols.

The *quantifier rank* $\mathrm{qr}(\varphi)$ of an FO-formula φ is defined as the maximal nesting depth of its quantifiers. By $\mathrm{free}(\varphi)$ we denote the set of all *free variables* of 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 $\mathrm{free}(\varphi)$ is a subset of $\{x_1, \ldots, x_n\}$.

[1] As usual, $\forall x$, $\wedge$, $\to$, $\leftrightarrow$ will be used as abbreviations when constructing formulas.

If $\mathcal{A}$ is a σ-structure and $\overline{a} = (a_1, \ldots, a_n) \in A^n$, we write $(\mathcal{A}, \overline{a}) \models \varphi(\overline{x})$ or $\mathcal{A} \models \varphi[\overline{a}]$ to 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 $a_1, \ldots, a_n$.

Two formulas $\varphi(\overline{x})$ and $\psi(\overline{x})$ over a signature σ are *equivalent on* $\mathfrak{C}$ (or, $\mathfrak{C}$-*equivalent*, for short: $\varphi \equiv_{\mathfrak{C}} \psi$) for a class $\mathfrak{C}$ of σ-structures, if for every $\mathcal{A} \in \mathfrak{C}$ and $\overline{a} \in A^n$, we have $\mathcal{A} \models \varphi[\overline{a}] \iff \mathcal{A} \models \psi[\overline{a}]$. We call φ and ψ *equivalent* (for short, $\varphi \equiv \psi$), if they are $\mathfrak{C}^\sigma$-equivalent for the class $\mathfrak{C}^\sigma$ of *all* σ-structures.

2.3. Unary counting quantifiers

A *unary counting quantifier* (for short: *quantifier*) Q is a set of natural numbers, i.e., $\mathsf{Q} \subseteq \mathbb{N}$. The *characteristic sequence* χ_{Q} of Q is the ω-word $w = w_0 w_1 w_2 \cdots \in \{0,1\}^\omega$ where

$$\text{for every } i \in \mathbb{N}, \qquad w_i = 1 \quad \iff \quad i \in \mathsf{Q}\,.$$

Example 2.1. Since we only consider *finite* structures, the classical quantifier $\exists$ and the unary counting quantifier $\mathbb{N}_{\geqslant 1}$ are equivalent. From now on, whenever convenient, we will therefore identify the existential quantifier $\exists$ with the unary counting quantifier $\mathbb{N}_{\geqslant 1}$.

2.3.1. Extensions of first-order logic by unary counting quantifiers

For a set $\mathbf{Q} \subseteq \mathcal{P}(\mathbb{N})$ of unary counting quantifiers, we write $\mathrm{FO}(\mathbf{Q})[\sigma]$ to denote the *extension of* $\mathrm{FO}[\sigma]$ *with the quantifiers from* $\mathbf{Q}$. That is, $\mathrm{FO}(\mathbf{Q})[\sigma]$ is built according to the same rules as $\mathrm{FO}[\sigma]$, but is also closed under unary counting quantifiers $\mathsf{Q}\,x$ for any variable x. By $\mathrm{FO}(\mathbf{Q})$ we denote the union of all $\mathrm{FO}(\mathbf{Q})[\sigma]$ for arbitrary signatures σ.

If Q is a quantifier from $\mathbf{Q}$ and $\varphi(\overline{x}, y)$ is a formula from $\mathrm{FO}(\mathbf{Q})[\sigma]$, then

$$(\mathcal{A}, \overline{a}) \;\models\; \mathsf{Q}\,y\,\varphi(\overline{x}, y) \quad \iff \quad |\{\, b \in A \,:\, \mathcal{A} \models \varphi[\overline{a}, b] \,\}| \;\in\; \mathsf{Q}\,.$$

for every σ-structure $\mathcal{A}$ and each interpretation $\overline{a}$ of the free variables $\overline{x}$.

The *quantifier rank* $\mathrm{qr}(\varphi)$ of an $\mathrm{FO}(\mathbf{Q})$-formula φ is defined as the maximal nesting depth of *all* quantifiers.

Example 2.2. The formula $\exists x\, \mathsf{Q}y\, E(x,y)$ (where $\mathsf{Q} \subseteq \mathbb{N}$ is some quantifier) has quantifier rank 2 and expresses that there exists a node whose out-degree belongs to Q.

2.3.2. Displaced unary counting quantifiers

For a number $k \geqslant 0$, a quantifier $\mathsf{Q} \subseteq \mathbb{N}$ and a formula $\varphi(\overline{x}, y)$, let $(\mathsf{Q}{+}k)\, y\, \varphi(\overline{x}, y)$ denote the formula

$$\begin{aligned} \exists y_1 \cdots \exists y_k \Bigg(& \bigwedge_{1 \leqslant i < j \leqslant k} \neg\, y_i{=}y_j \\ & \wedge \quad \forall y \Big(\bigvee_{1 \leqslant i \leqslant k} y{=}y_i \;\to\; \varphi(\overline{x}, y) \Big) \\ & \wedge \quad \mathsf{Q}y \Big(\varphi(\overline{x}, y) \wedge \bigwedge_{1 \leqslant i \leqslant k} \neg\, y{=}y_i \Big) \Bigg). \end{aligned}$$

For $k = 0$, this boils down to $\mathsf{Q}\, y\, \varphi(\overline{x}, y)$. The formula $(\mathsf{Q}{+}k)\, y\, \varphi(\overline{x}, y)$ expresses, in a σ-structure $\mathcal{A}$ and for an interpretation $\overline{a} \in A^{|\overline{x}|}$ of the variables $\overline{x}$, that the number of elements $b \in A$ with $\mathcal{A} \models \varphi[\overline{a}, b]$ belongs to the set

$$(\mathsf{Q}{+}k) \ :=\ \ \{\, n{+}k \ :\ n \in \mathsf{Q}\,\}\,.$$

For every $k \geqslant 1$, we abbreviate the formula $(\exists{+}(k{-}1))\, y\, \varphi$ by the expression $\exists^{\geqslant k} y\, \varphi$. Furthermore, we will write $\exists^{=k} y\, \varphi$ for the formula $\exists^{\geqslant k} y\, \varphi \wedge \neg \exists^{\geqslant k+1} y\, \varphi$. Note that the formulas $\exists^{\geqslant k} y\, \varphi$ and $\exists^{=k} y\, \varphi$ both have size in $\mathcal{O}(k^2{+}\|\varphi\|)$ and that, for an FO-formula φ, both expressions also denote FO-formulas.

The *displacement* of a formula ψ is the smallest number $K \geqslant 0$ such that for every subformula of ψ of shape $(\mathsf{Q}{+}k)\, y\, \varphi$ with $\mathsf{Q} \subseteq \mathbb{N}$ we have $k \leqslant K$.

The *generalised quantifier rank* $\mathrm{gqr}(\varphi)$ of a formula ψ discounts the quantifiers introduced by subformulas of shape $(\mathsf{Q}{+}k)\, y\, \varphi$. I.e., it is defined in the same way as the quantifier rank as the maximal nesting depth of all quantifiers, with the only exception that a formula ψ of shape $(\mathsf{Q}{+}k)\, y\, \varphi$ has *generalised* quantifier rank $\mathrm{gqr}(\psi) = 1 + \mathrm{gqr}(\varphi)$, whereas its quantifier rank is $\mathrm{qr}(\psi) = (k{+}1) + \mathrm{qr}(\varphi)$.

2.4. Ultimately periodic quantifiers

A set $\mathsf{Q} \subseteq \mathbb{N}$ is *ultimately periodic* (cf., e.g., [20]) if there exist numbers $p, n_0 \in \mathbb{N}$ with $p \geqslant 1$, such that

$$\text{for all } n \geqslant n_0 \text{ we have} \quad n \ \in\ \mathsf{Q} \quad \Longleftrightarrow \quad n{+}p \ \in\ \mathsf{Q}\,. \tag{1}$$

The *period* of Q is the minimal $p \geqslant 1$ for which statement (1) is true for some n_0; the number n_0 is called an *offset* of Q.

We write $\mathbf{U}$ to denote the set of all ultimately periodic quantifiers.

Example 2.3. The existential quantifier $\exists = \mathbb{N}_{\geqslant 1}$ and the modulo-counting quantifier $\mathsf{D}_p := \{\, p \cdot m \ :\ m \in \mathbb{N}\,\}$, for an arbitrary integer $p \geqslant 1$, are ultimately periodic sets (with period 1 and offset 1 and with period p and offset 0, respectively). It will be convenient to also allow the modulo-counting quantifier $\mathsf{D}_1 = \mathbb{N}$ since then, $\exists = (\mathsf{D}_1{+}1)$ can be understood as a displaced modulo-counting quantifier.

The following straightforward fact describes an ultimately periodic set in terms of an ultimately periodic characteristic sequence.

Fact 2.4. For every $\mathsf{Q} \subseteq \mathbb{N}$, the following holds:

- If Q is ultimately periodic with period p and offset n_0, then there are words $\alpha \in \{0,1\}^*$ of length n_0 and $\pi \in \{0,1\}^+$ of length p such that $\chi_{\mathsf{Q}} = \alpha \cdot \pi^\omega$. Furthermore, π is primitive.
- If $\chi_{\mathsf{Q}} = \alpha \cdot \pi^\omega$ for finite words $\alpha \in \{0,1\}^*$ and $\pi \in \{0,1\}^+$, then Q is ultimately periodic, its period is the length of the primitive root of π, and $|\alpha|$ is an offset.

By Fact 2.4, we can represent an ultimately periodic set Q by the finite word $rep(\mathsf{Q}) := \alpha\#\pi$, where $\chi_{\mathsf{Q}} = \alpha \cdot \pi^{\omega}$. To make this definition unambiguous, we demand that $p := |\pi|$ is the period of Q, and $n_0 := |\alpha|$ is the *smallest* offset of Q for p. The *size* $\|\mathsf{Q}\|$ of Q is defined as the length of $rep(\mathsf{Q})$.

The size $\|\varphi\|$ of an $\mathrm{FO}(\mathbf{U})[\sigma]$-formula φ 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 variable symbols, and where each quantifier $\mathsf{Q} \in \mathbf{U}$ is represented by the word $rep(\mathsf{Q})$.

It is the aim of this paper to study the locality of extensions $\mathrm{FO}(\mathbf{Q})$ of FO by sets $\mathbf{Q}$ of unary counting quantifiers $\mathsf{Q} \subseteq \mathbb{N}$ in the sense of Hanf's theorem [11, 6, 5]. To define the according locality notion for these logics, we need a few more notations and concepts; these are introduced in the remainder of this section.

2.5. Gaifman graph and bounded structures

Let $\mathcal{A}$ be a σ-structure. Its *Gaifman graph* $G_{\mathcal{A}}$ is the undirected graph with vertex set A and an edge between two distinct vertices $a, b \in A$ iff there exists $R \in \sigma$ and a tuple $(a_1, \dots, a_{\mathrm{ar}(R)}) \in R^{\mathcal{A}}$ such that $a, b \in \{a_1, \dots, a_{\mathrm{ar}(R)}\}$.

The *degree* of a σ-structure $\mathcal{A}$ is the degree of its Gaifman graph $G_{\mathcal{A}}$. A σ-structure $\mathcal{A}$ is *d-bounded*, for a *degree bound* $d \geqslant 0$, if no node in $G_{\mathcal{A}}$ has more than d neighbours.

By $\mathfrak{C}_d^{\sigma}$ we denote the *class of all d-bounded σ-structures*. Two formulas φ and ψ over the signature σ are *d-equivalent* (for short, $\varphi \equiv_d \psi$) if they are $\mathfrak{C}_d^{\sigma}$-equivalent.

The *distance* $dist^{\mathcal{A}}(a, b)$ 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 $G_{\mathcal{A}}$ (if no such path exists, we set $dist^{\mathcal{A}}(a, b) = \infty$).

For every $r \geqslant 0$ and $a \in A$, the *r-neighbourhood of a in $\mathcal{A}$* is the set

$$N_r^{\mathcal{A}}(a) \ := \ \ \{ b \in A : dist^{\mathcal{A}}(a, b) \leqslant r \}.$$

For a tuple $\overline{a} = (a_1, \dots, a_n) \in A^n$ of length $n \geqslant 1$, we write $N_r^{\mathcal{A}}(\overline{a})$ for the union of the sets $N_r^{\mathcal{A}}(a_i)$ for all $i \in [1, n]$.

2.6. Types, spheres, and sphere-formulas

Let σ be a relational signature and let $c_1, c_2, \dots$ be a sequence of pairwise distinct constant symbols. For every $r \geqslant 0$ and $n \geqslant 1$, a *type with n centres and radius at most r* (for short: *r-type with n centres*) is a structure $\tau = (\mathcal{A}, a_1, \dots, a_n)$ over the signature $\sigma \cup \{c_1, \dots, c_n\}$, where $\mathcal{A}$ is a σ-structure and $(a_1, \dots, a_n) \in A^n$ with $A = N_r^{\mathcal{A}}(a_1, \dots, a_n)$. The elements $a_1, \dots, a_n$ are called the *centres* of τ.

For $d, r \geqslant 0$ and $n \geqslant 1$, we denote by $T_r^{d,\sigma}(n)$ a *set of representatives of the isomorphism classes of all d-bounded types with n centres, radius at most r, and signature* $\sigma \cup \{c_1, \dots, c_n\}$. I.e., for every d-bounded type τ with n centres and radius at most r over the signature $\sigma \cup \{c_1, \dots, c_n\}$, there is precisely one $\tau' \in T_r^{d,\sigma}(n)$ with $\tau \cong \tau'$.

Let $m \geqslant 0$ and let $\mathcal{A}$ be a structure of signature σ or of signature $\sigma \cup \{c_1, \ldots, c_m\}$. For every non-empty set $B \subseteq A$, we write $\mathcal{A}[B]$ to denote the σ-reduct of the restriction of the structure $\mathcal{A}$ to the universe $B \subseteq A$. I.e., $\mathcal{A}[B]$ is the σ-structure with universe B, where for each relation symbol $R \in \sigma$, $R^{\mathcal{A}[B]} := R^{\mathcal{A}} \cap B^{\operatorname{ar}(R)}$. This structure $\mathcal{A}[B]$ is called the *induced substructure* of $\mathcal{A}$ on B.

For each tuple $\overline{a} = (a_1, \ldots, a_n) \in A^n$, the *$r$-sphere of $\overline{a}$ in $\mathcal{A}$* is defined as the r-type with n centres

$$\mathcal{N}_r^{\mathcal{A}}(\overline{a}) \;:=\; (\mathcal{A}[N_r^{\mathcal{A}}(\overline{a})], \overline{a})$$

over the signature $\sigma \cup \{c_1, \ldots, c_n\}$. We say that $\overline{a}$ *realises the type τ in $\mathcal{A}$* if $\mathcal{N}_r^{\mathcal{A}}(\overline{a}) \cong \tau$.

It is straightforward to verify that for any d-bounded structure $\mathcal{A}$, any node $a \in A$, and any $r \in \mathbb{N}$, we have

$$|N_r^{\mathcal{A}}(a)| \;\leqslant\; \nu_d(r)\,,$$

for

$$\nu_d(r) \;:=\; 1 \;+\; d \cdot \sum_{0 \leqslant i < r} (d-1)^i\,.$$

Observe that for all $r \geqslant 0$

$$\begin{aligned} \nu_0(r) &\;=\; 1\,, \\ \nu_1(r) &\;\leqslant\; 2\,, \\ \nu_2(r) &\;=\; 2r+1\,, \\ \text{and}\quad (d-1)^r \;\leqslant\; \nu_d(r) &\;\leqslant\; d^{r+1} \quad \text{for } d \geqslant 3\,. \end{aligned}$$

In other words, ν_d is growing linearly for $d \leqslant 2$ and exponentially for $d > 2$.

Note that for all $d, r \geqslant 0$ and $n \geqslant 1$, the universe of every type $\tau \in T_r^{d,\sigma}(n)$ contains at most $n \cdot \nu_d(r)$ elements. Thus, given τ and r, one can construct a *sphere-formula* $\operatorname{sph}_\tau(\overline{x})$, i.e., an $\mathrm{FO}[\sigma]$-formula such that for every σ-structure $\mathcal{A}$ and every tuple $\overline{a} \in A^n$ we have

$$\mathcal{A} \models \operatorname{sph}_\tau[\overline{a}] \quad\iff\quad \mathcal{N}_r^{\mathcal{A}}(\overline{a}) \;\cong\; \tau\,.$$

The formula $\operatorname{sph}_\tau(\overline{x})$ can be constructed in time $\mathcal{O}(\|\sigma\|)$ if $n \cdot \nu_d(r) = 1$, and otherwise in time $(n \cdot \nu_d(r))^{\mathcal{O}(\|\sigma\|)}$.

3. Main results

In the following, we fix a relational signature σ and a set $\mathbf{Q}$ of unary counting quantifiers. We generalise the classical notion of Hanf normal form (see, e.g., [2]) to the extension $\mathrm{FO}(\mathbf{Q})$ of first-order logic by unary counting quantifiers from $\mathbf{Q}$.

A *counting-formula* from $\mathrm{FO}(\mathbf{Q})[\sigma]$ is a formula of the form

$$(\mathsf{Q}{+}k)\, y\, \operatorname{sph}_\tau(\overline{x}, y)\,,$$

where $\mathsf{Q} \in \mathbf{Q} \cup \{\exists\}$, $k \in \mathbb{N}$, and τ is an r-type with $|\overline{x}|+1$ centres. A *counting-sentence* is a counting-formula without free variables, i.e., with $\overline{x}$ the empty tuple. We call r the *locality radius* of the counting-formula. The formula expresses that the number of interpretations for y such that the r-sphere around $\overline{x}, y$ is isomorphic to τ, belongs to the set $(\mathsf{Q}{+}k)$. I.e., for every σ-structure $\mathcal{A}$ and every tuple $\overline{a} \in A^{|\overline{x}|}$, we have

$$(\mathcal{A}, \overline{a}) \models (\mathsf{Q}{+}k)\, y\, \mathrm{sph}_\tau(\overline{x}, y) \iff |\{\, b \in A : \mathcal{N}_r^{\mathcal{A}}(\overline{a}, b) \cong \tau \,\}| \in (\mathsf{Q}{+}k)\,.$$

An $\mathrm{FO}(\mathbf{Q})[\sigma]$-formula $\varphi(\overline{x})$ is said to be in *Hanf normal form* if it is a Boolean combination[2] of counting-*sentences* from $\mathrm{FO}(\mathbf{Q})[\sigma]$ and sphere-formulas from $\mathrm{FO}[\sigma]$. Accordingly, a *sentence* in Hanf normal form is a Boolean combination of counting-sentences. We will speak of hnf-formulas and hnf-sentences when we mean "formula in Hanf normal form" and "sentence in Hanf normal form", respectively. The *locality radius* of an hnf-formula is the maximum of the locality radii of its counting-sentences and its sphere-formulas.

Remark 3.1. For sentences of first-order logic FO, the definition of Hanf normal form in this paper coincides with the one from [2]. On the other hand, the two notions of Hanf normal form handle free variables differently: While in [2] hnf-formulas are allowed to contain counting-*formulas with free variables*, the present paper's definition handles free variables by sphere-formulas. This will turn out as an advantage for model-checking algorithms (see Theorem 3.5 and Section 7).

Definition 3.2. A set $\mathbf{Q} \subseteq \mathcal{P}(\mathbb{N})$ of unary counting quantifiers *permits Hanf normal forms* if for every relational signature σ and every degree bound $d \geqslant 0$, every $\mathrm{FO}(\mathbf{Q})[\sigma]$-formula is d-equivalent to an hnf-formula from $\mathrm{FO}(\mathbf{Q})[\sigma]$.

From Hanf's [11, 6, 5] and Nurmonen's theorems [22] it follows, that the empty set, i.e., plain first-order logic, and every set $\{\mathsf{D}_p\}$ with $p \geqslant 1$, i.e., first-order logic extended by a modulo-counting quantifier, permit Hanf normal forms.

Our first main result gives a precise characterisation of all the sets $\mathbf{Q} \subseteq \mathcal{P}(\mathbb{N})$ that permit Hanf normal forms.

Theorem 3.3. *A set $\mathbf{Q} \subseteq \mathcal{P}(\mathbb{N})$ of unary counting quantifiers permits Hanf normal forms if, and only if, every quantifier $\mathsf{Q} \in \mathbf{Q}$ is ultimately periodic.*

The proof of the "only if"-direction of Theorem 3.3 can be found in Section 4. There, we show that it already holds for a signature with a single unary predicate.

The proof of the "if"-direction of Theorem 3.3 can be found in Section 5; the overall structure of the proof is as follows. Consider a formula $\varphi \in \mathrm{FO}(\mathbf{U})$. In a first step, we transform φ into an equivalent formula from $\varphi_1 \in \mathrm{FO}(\mathbf{D})$ where $\mathbf{D}$ is a suitable set of modulo-counting quantifiers D_p. In a second step,

[2] Throughout this paper, whenever we speak of *Boolean combinations*, we mean *finite* Boolean combinations.

this formula φ_1 is transformed into a d-equivalent formula $\varphi_2 \in \mathrm{FO}(\mathbf{D})$ in Hanf normal form. This step can be viewed as a constructive and algorithmic version of Nurmonen's result [22]; the actual algorithm is an adaptation and non-trivial extension of the algorithm from [2]. In a third (and final) step, the formula $\varphi_2 \in \mathrm{FO}(\mathbf{D})$ is translated into an equivalent one using at most the quantifiers from φ and the quantifier $\exists$. It turns out that, since φ_2 is in Hanf normal form, also the final formula $\varphi_3 \in \mathrm{FO}(\mathbf{U})$ is in Hanf normal form.

While the "if"-direction of Theorem 3.3 only asserts the *existence* of d-equivalent hnf-formulas, its proof demonstrates their computability. Our second main result concerns the runtime of this algorithm.

Theorem 3.4. *There is an algorithm which receives as input a degree bound $d \geqslant 2$, a relational signature σ, and a formula $\varphi \in \mathrm{FO}(\mathbf{U})[\sigma]$ and constructs a d-equivalent hnf-formula $\psi \in \mathrm{FO}(\mathbf{U})[\sigma]$. This hnf-formula ψ uses at most the quantifiers from φ and the quantifier $\exists$, and it has locality radius $\leqslant 4^{\mathrm{qr}(\varphi)}$ and displacement in*

$$2^{poly(\|\varphi\|)} \quad \text{for } d = 2, \qquad \text{and} \qquad d^{2^{\mathcal{O}(\|\varphi\|)}} \quad \text{for } d \geqslant 3.$$

If $\|\sigma\| \leqslant \|\varphi\|$, the algorithm runs in time

$$2^{2^{poly(\|\varphi\|)}} \quad \text{for } d = 2, \qquad \text{and} \qquad 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}} \quad \text{for } d \geqslant 3.$$

The proof of Theorem 3.4 can be found in Section 6. A 3-fold exponential lower bound from [2] shows that for $d = 3$ our algorithm is worst-case optimal, even for the special case of plain first-order logic. In Section 8, we generalise this lower bound proof to arbitrary $d \geqslant 3$ and adapt it to also obtain a matching 2-fold exponential lower bound for $d = 2$.

As an easy application of Theorem 3.4, we obtain that Seese's [24] fixed-parameter tractability result for first-order model checking can be generalised to first-order logic with ultimately periodic unary counting quantifiers. Precisely, we obtain the following, where $\|\mathcal{A}\|$ denotes the size of a reasonable encoding of a σ-structure $\mathcal{A}$ (as defined, e.g., in [7]).

Theorem 3.5. *There is an algorithm which receives as input*

- *a formula $\varphi(\overline{x}) \in \mathrm{FO}(\mathbf{U})$,*
- *a finite σ-structure $\mathcal{A}$ (where σ consists of precisely the relation symbols that occur in φ), and a tuple $\overline{a} \in A^{|\overline{x}|}$,*

and decides whether $\mathcal{A} \models \varphi[\overline{a}]$.

If $d \geqslant 2$ is a bound on the degree of $\mathcal{A}$, then the algorithm runs in time

$$2^{2^{poly(\|\varphi\|)}} \cdot \|\mathcal{A}\| \quad \text{for } d = 2, \qquad \text{and} \qquad 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}} \cdot \|\mathcal{A}\| \quad \text{for } d \geqslant 3.$$

A proof of Theorem 3.5 can be found in Section 7. Note that the lower bounds for plain first-order logic provided in [9] imply (under the complexity theoretic assumption $\mathrm{FPT} \neq \mathrm{AW}[*]$) that this model checking algorithm is worst-case optimal.

4. Only ultimately periodic quantifiers permit Hanf normal forms

In this section, we will prove the "only if"-direction of Theorem 3.3. More precisely, we will provide, for every non-ultimately periodic quantifier S a formula that is not d-equivalent to any hnf-formula from $\mathrm{FO}(\mathcal{P}(\mathbb{N}))$. For this, it suffices to consider a signature consisting of a single unary relation symbol. Technically, we show the following lemma.

Lemma 4.1. *Let $\sigma_P := \{P\}$ be the signature where P is a unary relation symbol. Let $\mathsf{S} \subseteq \mathbb{N}$ be a unary counting quantifier that is not ultimately periodic.*

Then, there is no hnf-sentence $\delta \in \mathrm{FO}(\mathcal{P}(\mathbb{N}))[\sigma_P]$, such that

$$\mathcal{A} \models \delta \quad \iff \quad |A| \in \mathsf{S}$$

holds for all σ_P-structures $\mathcal{A}$.

Proof. Let δ be an hnf-sentence from $\mathrm{FO}(\mathcal{P}(\mathbb{N}))[\sigma_P]$. We will show that δ does not expresses "$|A| \in \mathsf{S}$".

Since P is unary, the universe of any r-type with one centre is a singleton. Consequently, $P(y)$ and $\neg P(y)$ are the only formulas $\mathrm{sph}_\tau(y)$ of signature σ_P, where τ is a type with one centre. Hence, there is a finite set $\mathbf{Q}' \subseteq \mathcal{P}(\mathbb{N})$ and a natural number $k \geqslant 1$ such that δ is a Boolean combination of counting-sentences of the form

$$(\mathsf{Q}{+}\ell)y\ P(y) \qquad \text{or} \qquad (\mathsf{Q}{+}\ell)y\ \neg P(y)\ ,$$

where $\mathsf{Q} \in \mathbf{Q}'$ and $\ell \in [0,k)$.

Let $\mathsf{Q}_1, \mathsf{Q}_2, \dots, \mathsf{Q}_n$ with $n \geqslant 1$ be a list of all $\mathsf{Q} \in \mathbf{Q}'$. For each $a \in \mathbb{N}$ with $a \geqslant k$ consider the word w_a of length $k \cdot n$ defined as the concatenation of the bitstrings $\chi_{\mathsf{Q}_i}(a{-}k, a]$ for $i \in [1,n]$, i.e.,

$$w_a \ := \quad \chi_{\mathsf{Q}_1}(a{-}k,a]\ \chi_{\mathsf{Q}_2}(a{-}k,a]\ \cdots\ \chi_{\mathsf{Q}_n}(a{-}k,a]\,.$$

Since there is only a finite number of bitstrings of length $k \cdot n$, there are natural numbers $b > a > k$ such that $w_a = w_b$, i.e.,

$$\chi_{\mathsf{Q}_i}(a{-}k,a] \quad = \quad \chi_{\mathsf{Q}_i}(b{-}k,b] \tag{2}$$

for all $i \in [1,n]$.

If, for all $c \geqslant 0$, we have $a{+}c \in \mathsf{S} \iff b{+}c \in \mathsf{S}$, then S is ultimately periodic (with period dividing $b{-}a$ and offset a). Since this is not the case, there is a natural number $c \geqslant 0$ such that

$$a + c \ \in\ \mathsf{S} \quad \iff \quad b + c\ \notin\ \mathsf{S}\,. \tag{3}$$

Now consider σ_P-structures $\mathcal{A}$ and $\mathcal{B}$ with $|A| = a{+}c$, $|B| = b{+}c$, and $|P^{\mathcal{A}}| = |P^{\mathcal{B}}| = c$. By (3), we have

$$|A|\ \in\ \mathsf{S} \quad \iff \quad |B|\ \notin\ \mathsf{S}\,. \tag{4}$$

Nevertheless, $\mathcal{A}$ and $\mathcal{B}$ cannot be distinguished by any of the counting-sentences that occur in δ: to verify this, let $\mathsf{Q} \in \mathbf{Q'}$, let $\ell \in [0,k)$, and consider the counting-sentences $(\mathsf{Q}{+}\ell)\, y\, P(y)$ and $(\mathsf{Q}{+}\ell)\, y\, \neg P(y)$.

For the counting-sentence $(\mathsf{Q}{+}\ell)\, y\, P(y)$, we have

$$\begin{array}{rrcll} & \mathcal{A} & \models & (\mathsf{Q}{+}\ell)\, y\, P(y) & \\ \iff & |P^{\mathcal{A}}| & \in & (\mathsf{Q}{+}\ell) & \\ \iff & |P^{\mathcal{B}}| & \in & (\mathsf{Q}{+}\ell) & (\text{since } |P^{\mathcal{A}}| = |P^{\mathcal{B}}|) \\ \iff & \mathcal{B} & \models & (\mathsf{Q}{+}\ell)\, y\, P(y)\,. & \end{array}$$

For the counting-sentence $(\mathsf{Q}{+}\ell)\, y\, \neg P(y)$, we have

$$\begin{array}{rrcll} & \mathcal{A} & \models & (\mathsf{Q}{+}\ell)\, y\, \neg P(y) & \\ \iff & |A \setminus P^{\mathcal{A}}| & \in & (\mathsf{Q}{+}\ell) & \\ \iff & a-\ell & \in & \mathsf{Q} & (\text{since } |A \setminus P^{\mathcal{A}}| = a) \\ \iff & b-\ell & \in & \mathsf{Q} & (\text{by } (2)) \\ \iff & |B \setminus P^{\mathcal{B}}| & \in & (\mathsf{Q}{+}\ell) & (\text{since } |B \setminus P^{\mathcal{B}}| = b) \\ \iff & \mathcal{B} & \models & (\mathsf{Q}{+}\ell)\, y\, \neg P(y)\,. & \end{array}$$

In summary, the structures $\mathcal{A}$ and $\mathcal{B}$ satisfy the same counting-sentences that occur in δ. As δ is a Boolean combination of these counting-sentences, we obtain

$$\mathcal{A} \models \delta \quad \iff \quad \mathcal{B} \models \delta\,.$$

Thus, it follows from (4) that δ does not expresses "$|A| \in \mathsf{S}$". This completes the proof of Lemma 4.1. □

The "only if"-direction of Theorem 3.3 is an immediate consequence:

Proof of Theorem 3.3, "only if"-direction.
Let σ_P be the signature from Lemma 4.1. Since σ_P only contains a unary relation symbol, every σ_P-structure has degree 0.

Let $\mathbf{Q} \subseteq \mathcal{P}(\mathbb{N})$ be a set of unary counting quantifiers that contains a quantifier S which is not ultimately periodic. The $\mathrm{FO}(\mathbf{Q})[\sigma_P]$-sentence $\mathsf{S}\, y\ y{=}y$ expresses that "$|A| \in \mathsf{S}$". By Lemma 4.1 there is no hnf-sentence in $\mathrm{FO}(\mathbf{Q})[\sigma_P]$ that is d-equivalent to this sentence (for any degree bound $d \geqslant 0$). □

5. Constructing Hanf normal forms

The aim of this section is to prove the "if"-direction of Theorem 3.3. In other words, we will show that every set $\mathbf{Q} \subseteq \mathcal{P}(\mathbb{N})$ of ultimately periodic quantifiers permits Hanf normal forms.

Throughout this section, we will use the following notation. We fix a set $\mathbf{Q} \subseteq \mathcal{P}(\mathbb{N})$ of *ultimately periodic* quantifiers, and we let $\mathbf{D}$ be the set

$$\mathbf{D} \ :=\ \ \{\, \mathsf{D}_p \,:\, p \geqslant 1 \text{ and there is a } \mathsf{Q} \in \mathbf{Q} \text{ with period } p \,\}\,.$$

Let $\top$ be a fixed tautological hnf-sentence from $\mathrm{FO}[\sigma]$; e.g., we can choose $\top := \exists y\,\mathrm{sph}_\tau(y) \vee \neg\exists y\,\mathrm{sph}_\tau(y)$ where τ is an arbitrary, fixed type of radius 0 with one centre. Note that the locality radius and the displacement of $\top$ are both 0. Since the formula $\top$ depends on the signature σ, its size is in $\mathcal{O}(\|\sigma\|)$ (and we should write $\top_\sigma$, but we avoid this complication of notation). We let $\bot := \neg\top$ be the corresponding *unsatisfiable* sentence in Hanf normal form.

To prove the "if"-direction of Theorem 3.3, we will construct a d-equivalent hnf-formula from an arbitrary formula $\varphi \in \mathrm{FO}(\mathbf{Q})[\sigma]$. This construction will proceed in three steps:

Step (1): We first transform φ into an equivalent formula from $\mathrm{FO}(\mathbf{D})[\sigma]$, i.e., into a formula that uses modulo-counting quantifiers from $\mathbf{D}$ instead of the ultimately periodic quantifiers from $\mathbf{Q}$ (cf. Proposition 5.1).

Step (2): Generalising the proof from [2], we then construct a d-equivalent hnf-formula from $\mathrm{FO}(\mathbf{D})[\sigma]$. This step can be understood as an algorithmic version of Nurmonen's proof (cf. Proposition 5.7).

Step (3): Finally, the hnf-formula from $\mathrm{FO}(\mathbf{D})[\sigma]$ is translated into a d-equivalent hnf-formula from $\mathrm{FO}(\mathbf{Q})[\sigma]$ (cf. Proposition 5.9).

5.1. *Carrying out Step (1): From* FO(**Q**) *to* FO(**D**)

For a formula $\varphi \in \mathrm{FO}(\mathbf{Q})$, let $w(\varphi) := 1$ if φ is quantifier-free; and otherwise let $w(\varphi) := \max\{\|\mathsf{Q}\| + k : \text{the displaced quantifier } (\mathsf{Q}{+}k) \text{ appears in } \varphi\}$.

Proposition 5.1. *For all relational signatures σ and formulas $\varphi \in \mathrm{FO}(\mathbf{Q})[\sigma]$, there exists an equivalent formula $\psi \in \mathrm{FO}(\mathbf{D})[\sigma]$.*

Let q be the generalised quantifier rank of φ. Then ψ has displacement $\leqslant w(\varphi)$ and generalised quantifier rank q.

Furthermore, ψ can be computed from σ and φ in time

$$\|\varphi\| \cdot w(\varphi)^{\mathcal{O}(q)}\,.$$

Proof of the first two claims. [3]

This proof proceeds by a straightforward induction on the shape of the input formula. The only interesting case is a subformula of the form $(\mathsf{Q}{+}k)\,y\,\varphi$ with $\mathsf{Q} \in \mathbf{Q}$. This case is handled as follows:

Let n_0 and p be the smallest offset and the period of Q, respectively, and let $\mathsf{Q}' := (\mathsf{Q}{+}k)$. By definition, Q' is ultimately periodic with period p and offset $n_0' := n_0{+}k$. Let $n_1 \in \mathbb{N}$ be the (unique) number in $[n_0', n_0'{+}p)$ that is divisible by p. Clearly, Q' is also ultimately periodic with period p and offset n_1.

Let $\mathsf{Q}_1 := \mathsf{Q}' \cap [0, n_1)$ and let $R := \{r \in [0,p) : n_1{+}r \in \mathsf{Q}'\}$. It is straightforward to verify that $(\mathsf{Q}{+}k)\,y\,\varphi$ is equivalent to the formula

$$\bigvee_{\ell \in \mathsf{Q}_1} \exists^{=\ell} y\,\varphi \quad \vee \quad \left(\exists^{\geqslant n_1} y\,\varphi \;\;\wedge\;\; \bigvee_{r \in R} (\mathsf{D}_p{+}r)\,y\,\varphi \right). \tag{5}$$

[3] See Section 6 for the proof of the third claim.

Clearly, this formula has the same generalised quantifier rank as $(\mathsf{Q}{+}k)\, y\, \varphi$ and displacement $\leqslant w(\varphi)$. $\square$

5.2. Carrying out Step (2): Hanf normal form for $\mathrm{FO}(\mathbf{D})$

We proceed by induction on the construction of the formula $\varphi \in \mathrm{FO}(\mathbf{D})$ and thereby follow the same overall approach as that for first-order logic in [2]. Compared to [2], there are three additional difficulties: first and foremost, we have to handle modulo-counting quantifiers. Secondly, we handle the quantifier $\exists^{\geqslant k}$ in a single step in order to obtain better complexity bounds in terms of the *generalised* quantifier rank instead of the plain quantifier rank. Finally, as explained in Remark 3.1, the Hanf normal forms in this paper differ from the ones in [2] in the handling of free variables.

5.2.1. The base case

The first lemma handles the base case, i.e., quantifier-free formulas.

Lemma 5.2. *For all degree bounds $d \geqslant 2$, relational signatures σ, and quantifier-free formulas $\varphi(\overline{x})$ from $\mathrm{FO}[\sigma]$, there exists a d-equivalent hnf-formula $\psi(\overline{x})$ from $\mathrm{FO}[\sigma]$.*

The locality radius and the displacement of ψ are 0.

Furthermore, ψ can be computed from d, σ, and φ in time

$$2^{\|\varphi\|^{\mathtt{a}}}$$

for a suitable number $\mathtt{a} > 0$ of size $\mathcal{O}(\|\sigma\|)$.

Proof of the first two claims. [4]

Let $\overline{x} = (x_1, \ldots, x_n)$ with $1 \leqslant n < \|\varphi\|$ be the tuple of free variables of $\varphi(\overline{x})$. Since φ is quantifier-free, the validity of $\varphi[\overline{a}]$ in a σ-structure $\mathcal{A}$ is determined by the 0-sphere around $\overline{a} \in A^n$. More precisely, $\mathcal{A} \models \varphi[\overline{a}]$ if, and only if, $\mathcal{N}_0^{\mathcal{A}}(\overline{a}) \models \varphi[\overline{a}]$. Therefore, the algorithm proceeds as follows.

Step (i): We let $T \subseteq T_0^{d,\sigma}(n)$ be the set of all types $\tau = (\mathcal{T}, \overline{c})$ from $T_0^{d,\sigma}(n)$ with $\mathcal{T} \models \varphi[\overline{c}]$.

Step (ii): If T is the empty set, then there is no d-bounded σ-structure $\mathcal{A}$ with a tuple $\overline{a} \in A^n$ such that $\mathcal{A} \models \varphi[\overline{a}]$. Hence, $\varphi(\overline{x})$ is d-equivalent to the unsatisfiable hnf-sentence $\bot$. Otherwise, we let

$$\psi(\overline{x}) \ := \ \bigvee_{\tau \in T} \mathrm{sph}_\tau(\overline{x}).$$

In both cases, the hnf-formula $\psi(\overline{x})$ has locality radius and displacement 0. $\square$

[4] See Section 6 for the proof of the third claim.

5.2.2. The inductive step

Recall that we want to construct inductively, from a formula in $\mathrm{FO}(\mathbf{D})[\sigma]$, a d-equivalent hnf-formula. In this inductive procedure, we have to handle formulas of the form $\psi' \vee \psi''$, $\neg\psi'$, $\exists^{\geqslant k} y\, \psi'(\overline{x}, y)$, and $(\mathsf{D}_p{+}k)\, y\, \psi'(\overline{x}, y)$ where ψ' and ψ'' are hnf-formulas. The first two cases do not pose any problems since the set of hnf-formulas is closed under Boolean combinations. Note that $\exists^{\geqslant k} y\, \psi'(\overline{x}, y)$ is equivalent to $(\mathsf{D}_1{+}k)\, y\, \psi'(\overline{x}, y)$. Therefore, all we have to do is transform a formula of the form $(\mathsf{D}_p{+}k)\, y\, \psi'(\overline{x}, y)$ (with $p \geqslant 1$ and $k \geqslant 0$) into a d-equivalent hnf-formula. This is done in two main steps: First, Lemma 5.3 allows to build a d-equivalent Boolean combination of counting-formulas. Then, by Lemma 5.4, we transform every counting-formula into an hnf-sentence.

Lemma 5.3. *Let $d \geqslant 2$ be a degree bound, σ a relational signature, and $\varphi(\overline{x}) = (\mathsf{Q}{+}k)\, y\, \psi'(\overline{x}, y)$ a formula from $\mathrm{FO}(\mathbf{D})[\sigma]$ with ψ' an hnf-formula. There exists a d-equivalent Boolean combination $\psi(\overline{x}) \in \mathrm{FO}(\mathbf{D})[\sigma]$ of counting-formulas.*

Let the locality radius of ψ' be at most r, for $r \geqslant 1$, and let $K \in \mathbb{N}$ be the displacement of ψ'. Let $p \geqslant 1$ be the period of the quantifier $\mathsf{Q} \in \mathbf{D} \cup \{\exists\}$ and let $n = |\overline{x}|$ be the number of free variables of $\varphi(\overline{x})$. The locality radius of $\psi(\overline{x})$ is $\leqslant r$ and its displacement is $\leqslant \max\{K, k, p\}$. In particular, if $n \geqslant 1$, then $\psi(\overline{x})$ is a Boolean combination of hnf-sentences with displacement $\leqslant K$ and of counting-formulas with free variables and with displacement $\leqslant \max\{k, p\}$.

Furthermore, from d, σ, and $\varphi(\overline{x})$, one can compute $\psi(\overline{x})$ in time

$$\|\psi'\| \cdot \max\{2k, 2p\}^{((n+1)\cdot\nu_d(r))^{\mathcal{O}(\|\sigma\|)}} .$$

Proof of the first two claims. [5]

We describe an algorithm that computes ψ from d, σ, and $\varphi(\overline{x})$. Let $n := |\overline{x}|$ be the number of free variables of φ.

Step (i): Compute the set $T_r^{d,\sigma}(n{+}1)$ of all d-bounded r-spheres with $n{+}1$ centres.

Observe that for every d-bounded σ-structure $\mathcal{A}$ and for every tuple $(\overline{a}, a_{n+1}) \in A^{n+1}$, there is a unique $\tau \in T_r^{d,\sigma}(n{+}1)$ such that $\mathcal{N}_r^{\mathcal{A}}(\overline{a}, a_{n+1}) \cong \tau$.

Step (ii): Recall that, being an hnf-formula, $\psi'(\overline{x}, y)$ is a Boolean combination of sphere-formulas with free variables among $\overline{x}, y$ and of counting-sentences. In particular, every sphere-formula in $\psi'(\overline{x}, y)$ has locality radius $\leqslant r$.

For each $\tau = (\mathcal{T}, c_1, \ldots, c_n, c_{n+1}) \in T_r^{d,\sigma}(n{+}1)$, we construct an hnf-*sentence* ψ'_τ such that the following is true for every tuple $\overline{a}, a_{n+1}$ that realises the type τ in $\mathcal{A}$ (i.e., $\mathcal{N}_r^{\mathcal{A}}(\overline{a}, a_{n+1}) \cong \tau$):

$$\mathcal{A} \models \psi'[\overline{a}, a_{n+1}] \quad\Longleftrightarrow\quad \mathcal{A} \models \psi'_\tau .$$

[5] See Section 6 for the proof of the third claim.

To achieve this, we apply the following procedure to each sphere-formula $\mathrm{sph}_\rho(\overline{z})$ in $\psi'(\overline{x}, y)$: Recall that $\mathrm{sph}_\rho(\overline{z})$ has $m := |\overline{z}| \leqslant n{+}1$ free variables which all belong to the tuple $\overline{x}, y$ and let $1 \leqslant i_1 < \cdots < i_m \leqslant n{+}1$ be the indices of these free variables variables in the tuple $\overline{x}, x_{n+1}$ with $x_{n+1} := y$. We check whether $\mathcal{T} \models \mathrm{sph}_\rho[c_{i_1}, \ldots, c_{i_m}]$, i.e., whether $\mathcal{N}^{\mathcal{T}}_{r'}(c_{i_1}, \ldots, c_{i_m}) \cong \rho$, where $r' \leqslant r$ is the radius of the type ρ. If so, we replace every occurrence of the formula $\mathrm{sph}_\rho(\overline{z})$ in $\psi'(\overline{x}, y)$ with the tautological hnf-sentence $\top$; otherwise we replace it with the unsatisfiable hnf-sentence $\bot$.

Note that the resulting formula ψ'_τ is an hnf-*sentence* of locality radius $\leqslant r$ and with displacement $\leqslant K$. Furthermore,

$$\psi'(\overline{x}, y) \quad \equiv_d \quad \bigvee_{\tau \in T^{d,\sigma}_r(n+1)} \underbrace{\Big(\mathrm{sph}_\tau(\overline{x}, y) \ \wedge \ \psi'_\tau\Big)}_{=:\ \alpha_\tau(\overline{x}, y)} . \tag{6}$$

Step (iii): Recall that $\varphi(\overline{x})$ is of the form $(\mathsf{Q}{+}k)\, y\ \psi'(\overline{x}, y)$. Thus, from (6) we obtain that $\varphi(\overline{x})$ is d-equivalent to

$$(\mathsf{Q}{+}k)\ y \bigvee_{\tau \in T^{d,\sigma}_r(n+1)} \alpha_\tau(\overline{x}, y)\,. \tag{7}$$

Furthermore, recall that for every d-bounded σ-structure $\mathcal{A}$, every $\overline{a} \in A^n$ and every $b \in A$ there is exactly one $\tau \in T^{d,\sigma}_r(n{+}1)$ such that $\mathcal{A} \models \mathrm{sph}_\tau[\overline{a}, b]$, and hence there is at most one $\tau \in T^{d,\sigma}_r(n{+}1)$ such that $\mathcal{A} \models \alpha_\tau[\overline{a}, b]$.

Next, we transform the formula from (7) into an equivalent Boolean combination $\xi(\overline{x})$ of formulas of the form

$$(\mathsf{Q}{+}\ell)\, y\, \alpha_\tau(\overline{x}, y) \text{ with } \tau \in T^{d,\sigma}_r(n{+}1) \text{ and } 0 \leqslant \ell \leqslant \max\{k, p\}. \tag{8}$$

First, suppose $\mathsf{Q} = \mathsf{D}_p \in \mathbf{D}$. Then, the construction uses a divide-and-conqueror approach where we recursively subdivide the set $T^{d,\sigma}_r(n{+}1)$. Let $M \subseteq T^{d,\sigma}_r(n{+}1)$ with $|M| \geqslant 2$. Then divide M into two sets of almost equal size, i.e., choose $M_1 \subseteq M$ with $|M_1| = \left\lfloor \frac{|M|}{2} \right\rfloor$ and set $M_2 := M \setminus M_1$. Then, the formula

$$(\mathsf{D}_p{+}\ell)\, y \bigvee_{\tau \in M} \alpha_\tau(\overline{x}, y) \qquad \text{(with } 0 \leqslant \ell \leqslant \max\{k, p\})$$

is equivalent to the disjunction of all formulas

$$(\mathsf{D}_p{+}\ell_1)\, y \bigvee_{\tau \in M_1} \alpha_\tau(\overline{x}, y) \quad \wedge \quad (\mathsf{D}_p{+}\ell_2)\, y \bigvee_{\tau \in M_2} \alpha_\tau(\overline{x}, y)$$

with $0 \leqslant \ell_1, \ell_2 \leqslant \max\{\ell, p\} \leqslant \max\{k, p\}$, $\ell_1 + \ell_2 \equiv \ell \mod p$ and $\ell_1 + \ell_2 \geqslant \ell$. Proceeding recursively, we obtain that the formula

$(\mathsf{D}_p{+}k)\, y \bigvee_\tau \alpha_\tau(\overline{x}, y)$ is equivalent to a Boolean combination $\xi(\overline{x})$ of formulas of the form (8).

It remains to consider the case $\mathsf{Q} = \exists$. Then the formula from (7) is equivalent to

$$(\mathsf{D}_1{+}(k{+}1))\; y \bigvee_{\tau \in T_r^{d,\sigma}(n+1)} \alpha_\tau(\overline{x}, y)\,.$$

Hence the above construction results in a Boolean combination of formulas of the form

$$(\mathsf{D}_1{+}\ell)\; y\; \alpha_\tau(\overline{x}, y) \quad \equiv \quad \begin{cases} \top & \text{if } \ell = 0 \\ (\exists{+}(\ell{-}1))\, y\, \alpha_\tau(\overline{x}, y) & \text{otherwise} \end{cases}$$

with $0 \leqslant \ell \leqslant \max\{k{+}1, 1\}$. In other words, we obtain a Boolean combination $\xi(\overline{x})$ of formulas of the form (8).

Step (iv): Since ψ'_τ is a sentence, the formula

$$(\mathsf{Q}{+}\ell)\; y\; \alpha_\tau(\overline{x}, y) \quad = \quad (\mathsf{Q}{+}\ell)\; y\; \big(\, \mathrm{sph}_\tau(\overline{x}, y) \wedge \psi'_\tau \,\big)$$

is equivalent to the formula

- $\psi'_\tau \;\wedge\; (\mathsf{Q}{+}\ell)\, y\, \mathrm{sph}_\tau(\overline{x}, y)$ if $0 \notin (\mathsf{Q}{+}\ell)$,
- $\neg\psi'_\tau \;\vee\; (\mathsf{Q}{+}\ell)\, y\, \mathrm{sph}_\tau(\overline{x}, y)$ if $0 \in (\mathsf{Q}{+}\ell)$.

By replacing the formulas $(\mathsf{Q}{+}\ell)\, y\, \alpha_\tau(\overline{x}, y)$ in $\xi(\overline{x})$ accordingly, we obtain a Boolean combination $\psi(\overline{x})$ of hnf-sentences ψ'_τ and counting-formulas $(\mathsf{Q}{+}\ell)\, y\, \mathrm{sph}_\tau(\overline{x}, y)$, such that $\psi(\overline{x})$ is d-equivalent to $\varphi(\overline{x})$. The locality radius of $\psi(\overline{x})$ is $\leqslant r$. If $n = 0$ (i.e., $\overline{x}$ is the empty tuple), then $\psi(\overline{x})$ is an hnf-sentence with displacement $\leqslant \max\{k, p, K\}$. Otherwise, if $n \geqslant 1$, then $\psi(\overline{x})$ is a Boolean combination of hnf-sentences with displacement $\leqslant K$ and counting-formulas with free variables and with displacement $\leqslant \{k, p\}$. □

The above lemma allows to construct, from a formula $(\mathsf{D}_p{+}k)\, y\, \psi'$ with ψ' an hnf-formula, a d-equivalent Boolean combination ψ of counting-formulas. To obtain an hnf-formula, i.e., a Boolean combination of sphere formulas and counting-*sentences*, it remains to transform the counting-formulas in ψ into hnf-formulas. This is achieved by the following lemma. Note that it proves the *existence* of a d-equivalent hnf-formula for all quantifiers, but provides an *algorithm* only for ultimately periodic quantifiers (and therefore in particular for the modulo-counting quantifiers that we are really interested in).

Lemma 5.4. *Let $d \geqslant 2$ be a degree bound, σ a relational signature, R a unary counting quantifier, and $\alpha(\overline{x}) \in \mathrm{FO}(\{\mathsf{R}\})[\sigma]$ a counting-formula. There exists an hnf-formula $\beta(\overline{x}) \in \mathrm{FO}(\{\mathsf{R}\})[\sigma]$ that is d-equivalent to $\alpha(\overline{x})$.*

Let the locality radius of α be at most r, for $r \geqslant 1$, and let ℓ be the displacement and $n = |\overline{x}|$ the number of free variables of α. Then, $\beta(\overline{x})$ has locality radius $\leqslant 4r$ and displacement $\leqslant \ell + n \cdot \nu_d(2r{+}1)$.

Furthermore, from d, σ, and $\alpha(\overline{x})$ with R ultimately periodic, one can compute $\beta(\overline{x})$ in time

$$(\ell^2{+}\|\mathsf{R}\|) \cdot 2^{(n\cdot\nu_d(4r))^{\mathcal{O}(\|\sigma\|)}} .$$

Proof of the first two claims. [6]

If $n = 0$ (i.e., $\overline{x}$ is the empty tuple), then $\beta := \alpha$ is a counting-sentence, hence an hnf-sentence, and we are done.

If $n \geqslant 1$, then $\alpha(\overline{x})$ is of the form

$$(\mathsf{R}{+}\ell)\, y \ \mathrm{sph}_\tau(\overline{x}, y)$$

for some type $\tau \in T^{d,\sigma}_r(n{+}1)$. Let $\overline{x} = (x_1, \ldots, x_n)$ be the free variables of $\alpha(\overline{x})$, and let $\overline{c}, c_{n+1} = (c_1, \ldots, c_n, c_{n+1})$ be the $n{+}1$ centres of τ.

For each $\rho \in T^{d,\sigma}_{4r}(n)$, we will provide an hnf-sentence $\beta_\rho \in \mathrm{FO}(\{\mathsf{R}\})[\sigma]$ such that

$$\begin{array}{c}\text{for all } \sigma\text{-structures } \mathcal{A} \text{ and all } \overline{a} \in A^n \text{ with } \mathcal{N}^{\mathcal{A}}_{4r}(\overline{a}) \cong \rho \text{ we have} \\ \mathcal{A} \models \alpha[\overline{a}] \quad \iff \quad \mathcal{A} \models \beta_\rho\, . \end{array} \tag{9}$$

Then, obviously,

$$\beta(\overline{x}) \ := \bigvee_{\rho \in T^{d,\sigma}_{4r}(n)} \Big(\mathrm{sph}_\rho(\overline{x}) \ \wedge \ \beta_\rho \Big)$$

is in Hanf normal form and is d-equivalent to $\alpha(\overline{x})$.

So consider an arbitrary $\rho \in T^{d,\sigma}_{4r}(n)$ and let $\overline{e} = (e_1, \ldots, e_n)$ be the n centres of ρ. Recall that $\alpha(\overline{x})$ is of the form $(\mathsf{R}{+}\ell)\, y \ \mathrm{sph}_\tau(\overline{x}, y)$, $\tau \in T^{d,\sigma}_r(n{+}1)$, and $\overline{c}, c_{n+1} = (c_1, \ldots, c_n, c_{n+1})$ are the $n{+}1$ centres of τ. The construction of β_ρ proceeds by the following case distinction (see Figure 1 for an illustration).

Case (1): If $c_{n+1} \in N^\tau_{2r+1}(\overline{c})$, then $N^\tau_r(\overline{c})$ intersects $N^\tau_r(c_{n+1})$ or some edge of the Gaifman graph of τ connects some element of the former set to some element of the latter set. Hence, for every σ-structure $\mathcal{A}$ and all $\overline{a} \in A^n$ we have

$$\begin{aligned} & |\{\, b \in A \,:\, \mathcal{N}^{\mathcal{A}}_r(\overline{a}, b) \cong \tau \,\}| \\ = \ & |\{\, b \in N^{\mathcal{A}}_{2r+1}(\overline{a}) \,:\, \mathcal{N}^{\mathcal{A}}_r(\overline{a}, b) \cong \tau \,\}|\, . \end{aligned} \tag{10}$$

Since $(2r{+}1){+}r = 3r{+}1$, we have $N^\tau_r(c_{n+1}) \subseteq N^\tau_{3r+1}(\overline{c}) \subseteq N^\tau_{4r}(\overline{c})$. Thus, for every σ-structure $\mathcal{A}$ and every tuple $\overline{a} \in A^n$ of type ρ in $\mathcal{A}$ (i.e., $\mathcal{N}^{\mathcal{A}}_{4r}(\overline{a}) \cong \rho$), we get

$$\begin{aligned} & |\{\, b \in N^{\mathcal{A}}_{2r+1}(\overline{a}) \,:\, \mathcal{N}^{\mathcal{A}}_r(\overline{a}, b) \cong \tau \,\}| \\ = \ & |\{\, f \in N^{\rho}_{2r+1}(\overline{e}) \,:\, \mathcal{N}^{\rho}_r(\overline{e}, f) \cong \tau \,\}| \quad =: \ k_\rho\, . \end{aligned} \tag{11}$$

[6] See Section 6 for the proof of the third claim.

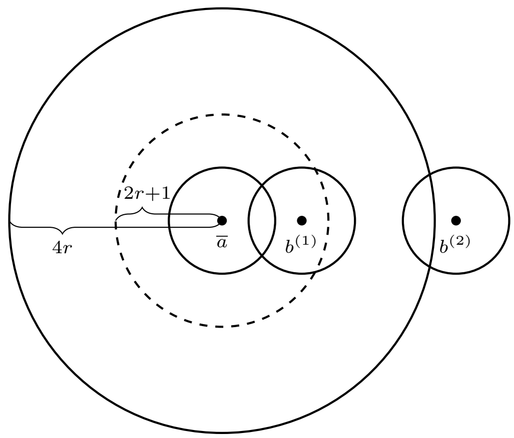

Figure 1: Illustration of the case distinction in the proof of Lemma 5.4; $b^{(1)}$ depicts b in Case (1), whereas $b^{(2)}$ depicts b in Case (2).

We set $\beta_\rho := \top$ if $k_\rho \in (\mathsf{R}{+}\ell)$, and $\beta_\rho := \bot$ if $k_\rho \notin (\mathsf{R}{+}\ell)$. Putting the equations (10) and (11) together, the hnf-sentence β_ρ clearly satisfies statement (9).

Case (2): If $c_{n+1} \notin N^{\tau}_{2r+1}(\overline{c})$, then the sets $N^{\tau}_{r}(\overline{c})$ and $N^{\tau}_{r}(c_{n+1})$ are disjoint and there are no edges in the Gaifman graph of τ between the nodes from $N^{\tau}_{r}(\overline{c})$ and the nodes from $N^{\tau}_{r}(c_{n+1})$.

Case (2.a): If $\mathcal{N}^{\tau}_{r}(\overline{c})$ and $\mathcal{N}^{\rho}_{r}(\overline{e})$ are not isomorphic, then for every σ-structure $\mathcal{A}$ and every $\overline{a} \in A^n$ with $\mathcal{N}^{\mathcal{A}}_{4r}(\overline{a}) \cong \rho$ we have

$$|\{\, b \in A \,:\, \mathcal{N}^{\mathcal{A}}_{r}(\overline{a}, b) \cong \tau \,\}| \quad = \quad 0\,.$$

Hence, if $0 \in (\mathsf{R}{+}\ell)$ we let $\beta_\rho := \top$, and if $0 \notin (\mathsf{R}{+}\ell)$ we let $\beta_\rho := \bot$. Clearly, β_ρ is an hnf-sentence that satisfies statement (9).

Case (2.b): If $\mathcal{N}^{\tau}_{r}(\overline{c})$ and $\mathcal{N}^{\rho}_{r}(\overline{e})$ are isomorphic, then for every σ-structure $\mathcal{A}$ and every $\overline{a} \in A^n$ with $\mathcal{N}^{\mathcal{A}}_{4r}(\overline{a}) \cong \rho$ we have

$$\begin{aligned} & |\{\, b \in A \,:\, \mathcal{N}^{\mathcal{A}}_{r}(\overline{a}, b) \cong \tau \,\}| \\ = \quad & |\{\, b \in A \setminus N^{\mathcal{A}}_{2r+1}(\overline{a}) \,:\, \mathcal{N}^{\mathcal{A}}_{r}(b) \cong \mathcal{N}^{\tau}_{r}(c_{n+1}) \,\}|\,. \end{aligned} \tag{12}$$

On the other hand, since $2r{+}1{+}r \leqslant 4r$, we have

$$\begin{aligned} & |\{\, b \in N^{\mathcal{A}}_{2r+1}(\overline{a}) \,:\, \mathcal{N}^{\mathcal{A}}_{r}(b) \cong \mathcal{N}^{\tau}_{r}(c_{n+1}) \,\}| \\ = \quad & |\{\, f \in N^{\rho}_{2r+1}(\overline{e}) \,:\, \mathcal{N}^{\rho}_{r}(f) \cong \mathcal{N}^{\tau}_{r}(c_{n+1}) \,\}| \quad =: \ \ell_\rho \end{aligned} \tag{13}$$

Hence, putting together the equations (12) and (13), we obtain that

$$
\begin{aligned}
& |\{\, b \in A \,:\, \mathcal{N}_r^{\mathcal{A}}(\overline{a}, b) \cong \tau \,\}| \\
= \quad & |\{\, b \in A \,:\, \mathcal{N}_r^{\mathcal{A}}(b) \cong \mathcal{N}_r^{\tau}(c_{n+1}) \,\}| \;-\; \ell_\rho \,.
\end{aligned}
$$

Therefore, the hnf-sentence

$$\beta_\rho \;:=\; (\mathsf{R}{+}(\ell{+}\ell_\rho))\, y \; \mathrm{sph}_{\mathcal{N}_r^{\tau}(c_{n+1})}(y)$$

satisfies statement (9). Observe that $\ell_\rho \leqslant n \cdot \nu_d(2r{+}1)$ and hence, β_ρ has displacement $\leqslant \ell + n \cdot \nu_d(2r{+}1)$.

Observe that for each $\rho \in T_{4r}^{d,\sigma}(n)$ the hnf-sentence β_ρ has locality radius $\leqslant r$ and displacement $\leqslant \ell + n \cdot \nu_d(2r{+}1)$. Consequently, the hnf-formula $\beta(\overline{x})$ has locality radius $4r$ and displacement $\leqslant \ell + n \cdot \nu_d(2r{+}1)$.

Furthermore, for ultimately periodic R, it can easily be decided for a given number m if $m \in (\mathsf{R}{+}\ell)$. In Case (1) and Case (2.a), this is used for $m = k_\rho$ and $m = 0$, respectively, and leads to an algorithm which constructs the hnf-formula $\beta(\overline{x})$. ☐

Remark 5.5. *Note that the above proof also shows that the hnf-formula β is computable if R is decidable.*

Summarising the two above lemmas, we are now ready to transform a formula of the form $(\mathsf{Q}{+}k)\, y\, \psi'(\overline{x}, y)$ with ψ' an hnf-formula, into an hnf-formula.

Lemma 5.6. *Let $d \geqslant 2$ be a degree bound, σ a relational signature, and $\varphi(\overline{x}) = (\mathsf{Q}{+}k)\, y\, \psi'(\overline{x}, y)$ a formula from $\mathrm{FO}(\mathbf{D})[\sigma]$ with ψ' an hnf-formula. There exists a d-equivalent hnf-formula $\psi(\overline{x}) \in \mathrm{FO}(\mathbf{D})[\sigma]$.*

Let the locality radius of ψ' be at most r, for $r \geqslant 1$, and let $K \in \mathbb{N}$ be the displacement of ψ'. Let $p \geqslant 1$ be the period of the quantifier $\mathsf{Q} \in \mathbf{D} \cup \{\exists\}$ and let $n = |\overline{x}|$ the number of free variables of φ. The locality radius of $\psi(\overline{x})$ is $\leqslant 4r$ and its displacement is $\leqslant \max\{K, \max\{k, p\} + n \cdot \nu_d(2r{+}1)\}$.

Furthermore, from d, σ, and $\varphi(\overline{x})$, one can compute $\psi(\overline{x})$ in time

$$\|\psi'\| \cdot \max\{2k, 2p\}^{((n+1)\nu_d(4r))^{\mathsf{b}}}$$

for some $\mathsf{b} > 0$ of size $\mathcal{O}(\|\sigma\|)$.

Proof of the first two claims. [7]

Applying the algorithm from Lemma 5.3 to the formula φ, we obtain a d-equivalent Boolean combination ψ_1 of counting-formulas of locality radius $\leqslant r$.

If $n = 0$, i.e., if φ is a sentence, then ψ_1 is an hnf-sentence with displacement $\leqslant \max\{K, k, p\}$ and thus, the construction of ψ is finished by choosing $\psi := \psi_1$.

[7] See Section 6 for the proof of the third claim.

Otherwise, i.e., if $n \geqslant 1$, then ψ_1 is a Boolean combination of hnf-sentences with displacement $\leqslant K$ and counting-formulas with free variables and displacement $\leqslant \{k, p\}$. Applying the algorithm from Lemma 5.4 to each of these counting-formulas, we obtain a d-equivalent hnf-formula $\psi(\overline{x})$ of locality radius $\leqslant 4r$ and displacement $\leqslant \max\{K, \max\{k, p\} + n \cdot \nu_d(2r{+}1)\}$. $\square$

5.2.3. Putting things together: construction of hnf-formulas in FO(**D**)

Recall that by Nurmonen's result [22], any set $\{\mathsf{D}_p\}$ permits Hanf normal forms. It follows that also the set of all modulo-counting quantifiers **D** permits Hanf normal forms. The following proposition provides us with an algorithm that allows to compute such an hnf-formula.

Proposition 5.7. *For all degree bounds $d \geqslant 2$, relational signatures σ, and formulas $\varphi \in \mathrm{FO}(\mathbf{D})[\sigma]$, there exists a d-equivalent hnf-formula $\psi \in \mathrm{FO}(\mathbf{D})[\sigma]$.*

Let $q \geqslant 0$ and $K \geqslant 0$ be the generalised quantifier rank and the displacement of φ, respectively, and let $P \geqslant 1$ be an upper bound on the periods of the quantifiers occurring in φ. Then ψ has locality radius $\leqslant 4^q$ and displacement $\leqslant \max\{K, P\} + \|\varphi\| \cdot \nu_d(4^q)$.

Furthermore, ψ can be computed from d, σ, and φ in time

$$\max\{2K, 2P\}^{(\|\varphi\| \cdot \nu_d(4^q))^{\mathsf{c}}}$$

for some $\mathsf{c} > 0$ of size $\mathcal{O}(\|\sigma\|)$.

Proof of the first two claims. [8]

We describe the algorithm on input of a degree bound $d \geqslant 2$, a relational signature σ, and an $\mathrm{FO}(\mathbf{D})[\sigma]$-formula $\varphi(\overline{x})$.

Let $n := |\overline{x}|$ be the number of free variables of φ.

The algorithm proceeds by induction on the shape of $\varphi(\overline{x})$. We will show that for the hnf-formula $\psi(\overline{x})$ from $\mathrm{FO}(\mathbf{D})[\sigma]$, which the algorithm constructs, the following *inductive invariant* holds:

(a) $\psi(\overline{x})$ has locality radius $\leqslant 4^q$.

(b) $\psi(\overline{x})$ is d-equivalent to $\varphi(\overline{x})$.

(c) $\psi(\overline{x})$ has displacement $\leqslant \max\{K, P\} + \|\varphi\| \cdot \nu_d(4^q)$.

If φ is quantifier-free (i.e., $q = 0$), we employ the algorithm from Lemma 5.2 to construct an hnf-formula $\psi(\overline{x})$ which satisfies the inductive invariant.

The case of φ being a Boolean combination of formulas with generalised quantifier rank $q \geqslant 1$ is also easy: If $\varphi = \neg\varphi'$, we compute an hnf-formula ψ' that is d-equivalent to φ' and let $\psi := \neg\psi'$. If $\varphi = (\varphi' \vee \varphi'')$, we compute hnf-formulas ψ' and ψ'' that are d-equivalent to φ' and φ'', respectively, and let $\psi := (\psi' \vee \psi'')$. In both cases, the inductive invariant is obviously satisfied.

[8]See Section 6 for the proof of the third claim.

Finally, consider the case that $\varphi = (\mathsf{Q}{+}k)\, y\, \varphi'$ with $\mathsf{Q} \in \mathbf{D} \cup \{\exists\}$ of period $p \in [1, P]$ and with $k \in [0, K]$: We first compute an hnf-formula ψ' that is d-equivalent to φ'. By the inductive invariant, the locality radius of ψ' is $\leqslant 4^{q-1}$ and the displacement is $\leqslant \max\{K, P\} + \|\varphi'\| \cdot \nu_d(4^{q-1})$. Then we apply Lemma 5.6 to the formula $(\mathsf{Q}{+}k)\, y\, \psi'$. This results in a d-equivalent hnf-formula ψ of locality radius $\leqslant 4 \cdot 4^{q-1} = 4^q$ and displacement

$$\begin{aligned}
&\leqslant \quad \max\{\max\{K, P\} + \|\varphi'\| \cdot \nu_d(4^{q-1}),\ \max\{K, P\} + n \cdot \nu_d(2 \cdot 4^{q-1}{+}1)\} \\
&= \quad \max\{K, P\} + \max\{\|\varphi'\| \cdot \nu_d(4^{q-1}),\ n \cdot \nu_d(2 \cdot 4^{q-1}{+}1)\} \\
&\leqslant \quad \max\{K, P\} + \|\varphi\| \cdot \nu_d(4^q).
\end{aligned}$$

This finishes the inductive proof. □

5.3. Carrying out Step (3): From FO($\mathbf{D}$) *to* FO($\mathbf{Q}$)

It remains to transform an hnf-formula from FO($\mathbf{D}$) into a d-equivalent hnf-formula from FO($\mathbf{Q}$). All that needs to be done is to transform counting-sentences of the form $(\mathsf{D}_p{+}k)\, y\, \mathrm{sph}_\tau(y)$ into hnf-sentences from FO($\mathbf{Q}$). For this, we will use the following lemma.

Lemma 5.8. *Let $\alpha, \beta, \pi \in \{0,1\}^*$ such that π is primitive and of length $p \geqslant 1$, and β is a prefix of $\alpha\pi^\omega$ of length at least $|\alpha|$. Then*

$$|\alpha| \equiv |\beta| \mod p \qquad \Longleftrightarrow \qquad \beta\pi \text{ is a prefix of } \alpha\pi^\omega\,.$$

Proof. We first prove the implication "$\Longrightarrow$". So assume that $|\beta| - |\alpha|$ is a multiple of $p = |\pi|$. Since β is a prefix of $\alpha\pi^\omega$, there exists $i \geqslant 0$ with $\beta = \alpha\pi^i$. Hence $\beta\pi = \alpha\pi^{i+1}$ is a prefix of $\alpha\pi^\omega$.

To prove the converse implication "$\Longleftarrow$", we assume that $\beta\pi$ is a prefix of $\alpha\pi^\omega$. Let $i \in \mathbb{N}$ be maximal such that $\alpha\pi^i$ is a prefix of $\beta\pi$ (such an $i \geqslant 0$ exists since $\alpha = \alpha\pi^0$ is a prefix of β). Then there exists $u \in \{0,1\}^*$ of length $< |\pi|$ with $\alpha\pi^i u = \beta$. In the following, we show that $|u| = 0$ and thus $|\beta| - |\alpha|$ is a multiple of $p = |\pi|$.

Towards a contradiction, assume that $|u| \geqslant 1$. Since $\alpha\pi^{i+1}$ and $\beta\pi$ both are prefixes of $\alpha\pi^\omega$ and since $|\beta\pi| < |\alpha\pi^{i+1}|$, there exists $v \in \{0,1\}^*$ of length $\geqslant 1$ with $\alpha\pi^{i+1} = \beta\pi v$. Hence $\alpha\pi^{i+1} = \beta\pi v = \alpha\pi^i uv$, implying that $uv = \pi$.

Furthermore, $\beta\pi vu = \alpha\pi^i uvu = \alpha\pi^{i+1}u$ is a prefix of $\alpha\pi^\omega$. Note that $(\alpha\pi^\omega)[n] = (\alpha\pi^\omega)[n+p]$ for all $n \geqslant |\alpha|$ and therefore in particular for all $n \geqslant |\beta|$. Since $\beta\pi$ is a prefix of $\alpha\pi^\omega$, this implies that $\alpha\pi^\omega = \beta\pi^\omega$, i.e., $\beta\pi vu$ is a prefix of $\beta\pi^\omega$. Hence vu is a prefix of π^ω of length $|vu| = |uv| = |\pi|$, i.e., $vu = \pi$.

Thus, we have $uv = \pi = vu$. By a standard result on *word combinatorics* (see [19, Proposition 1.3.2]), there exists $w \in \{0,1\}^*$ with $u, v \in w^*$ and therefore $\pi \in w^*$. Since π is primitive, this implies that $w = \pi$. Since $|u| < |\pi|$ and $u \in \pi^*$, we obtain $u = \varepsilon$, which is a contradiction to the assumption that $|u| \geqslant 1$. □

To translate an hnf-formula from FO($\mathbf{D}$)[σ] into an equivalent hnf-formula from FO($\mathbf{Q}$)[σ], we prove the following slightly more general result.

Proposition 5.9. *For all $p \geqslant 1$, all ultimately periodic quantifiers Q of period p, and all formulas $(\mathsf{D}_p{+}k)\, y\ \varphi \in \mathrm{FO}(\mathbf{D})$, there exists an equivalent Boolean combination ψ of formulas of the form $(\mathsf{Q}{+}\ell)\, y\ \varphi$ and $(\exists{+}\ell)\, y\ \varphi$ with $0 \leqslant \ell < k + \|\mathsf{Q}\| + p$.*

Furthermore, ψ can be computed from p, rep(Q), and $(\mathsf{D}_p{+}k)\, y\ \varphi$ in time

$$\mathcal{O}((\|\mathsf{Q}\|{+}k)^3 \cdot \|\varphi\|).$$

Proof. Let p be the period of Q and let n_0 be the minimal offset of Q w.r.t. p. Let furthermore α be the shortest prefix of χ_{Q} of length $\geqslant \max\{n_0, k\}$ with $|\alpha| \equiv k \bmod p$. Then there exists a primitive word π of length p with $\chi_{\mathsf{Q}} = \alpha\pi^\omega$.

From Lemma 5.8 we obtain for all $n \in \mathbb{N}$ with $n \geqslant |\alpha\pi|$ that

$$n \ \in\ (\mathsf{D}_p{+}k) \quad\iff\quad \chi_{\mathsf{Q}}[n{-}p, n) \ =\ \pi\,. \tag{14}$$

Here, the implication "$\Longrightarrow$" holds since $\chi_{\mathsf{Q}} = \alpha\pi^\omega$, $|\pi| = p$, and $k \equiv |\alpha| \bmod p$. To obtain the implication "$\Longleftarrow$", let $\beta := \chi_{\mathsf{Q}}[0, n{-}p)$. Then, $\beta\pi = \chi_{\mathsf{Q}}[0, n)$ is a prefix of $\chi_{\mathsf{Q}} = \alpha\pi^\omega$. By Lemma 5.8, $|\alpha| \equiv |\beta| \bmod p$. Hence, $n = |\beta\pi| = |\beta| + p \equiv |\alpha| \equiv k \bmod p$, and therefore, $n \in (\mathsf{D}_p{+}k)$. This proves the equivalence (14).

To apply this equivalence, let $\pi_1, \pi_2, \ldots, \pi_p \in \{0, 1\}$ such that $\pi = \pi_p \cdots \pi_2\pi_1$. It is straightforward to see that for all $n \in \mathbb{N}$ with $n \geqslant |\alpha\pi|$,

$$\begin{aligned} & \chi_{\mathsf{Q}}[n{-}p, n) \ =\ \pi \\ \iff\quad & n - i \ \in\ \mathsf{Q} \quad \text{for each } i \in [1,p] \text{ with } \pi_i = 1, \text{ and} \\ & n - j \ \notin\ \mathsf{Q} \quad \text{for each } j \in [1,p] \text{ with } \pi_j = 0 \\ \iff\quad & n \ \in\ (\mathsf{Q}{+}i) \quad \text{for each } i \in [1,p] \text{ with } \pi_i = 1, \text{ and} \\ & n \ \notin\ (\mathsf{Q}{+}j) \quad \text{for each } j \in [1,p] \text{ with } \pi_j = 0. \end{aligned}$$

Consequently, the formula $(\mathsf{D}_p{+}k)\, y\ \varphi$ is equivalent to the formula

$$\begin{aligned} & \bigvee_{\ell \in S} \exists^{=\ell} y\ \varphi \\ \vee\ & \Bigg(\exists^{\geqslant|\alpha\pi|} y\ \varphi \quad\wedge \bigwedge_{\substack{i \in [1,p]:\\ \pi_i = 1}} (\mathsf{Q}{+}i)\, y\ \varphi \quad\wedge \bigwedge_{\substack{j \in [1,p]:\\ \pi_j = 0}} \neg\,(\mathsf{Q}{+}j)\, y\ \varphi \Bigg), \end{aligned} \tag{15}$$

where S is the set of all $n \in (\mathsf{D}_p{+}k)$ with $n < |\alpha\pi|$. Since $|\alpha\pi| \leqslant \max\{n_0, k\}{+}2p$, each of the quantifiers that explicitly occur in the formula (15) has displacement at most

$$n_0 + k + 2p \quad\leqslant\quad 2\|\mathsf{Q}\| + k. \tag{16}$$

Using this, it is straightforward to see that the formula (15) has size

$$\mathcal{O}((\|\mathsf{Q}\|{+}k)^3 \cdot \|\varphi\|)$$

and can be computed within the same time bound. □

5.4. Summary: Hanf normal forms for FO(**Q**)

We can now prove that any set of ultimately periodic quantifiers permits Hanf normal forms.

Proof of Theorem 3.3, "if"-direction. Let $\mathbf{Q}$ be a set of ultimately periodic quantifiers, let $d \in \mathbb{N}$ be a degree bound, and let $\varphi \in \mathrm{FO}(\mathbf{Q})[\sigma]$ be a formula. By Proposition 5.1, there exists a d-equivalent formula $\varphi_1 \in \mathrm{FO}(\mathbf{D})[\sigma]$ where

$$\mathbf{D} \;=\; \{\mathsf{D}_p \,:\, p \geqslant 1 \text{ and there is some } \mathsf{Q} \in \mathbf{Q} \text{ with period } p\}\,.$$

From Proposition 5.7 we obtain a d-equivalent hnf-formula ψ_1 from $\mathrm{FO}(\mathbf{D})[\sigma]$. Finally, Proposition 5.9 allows to translate this hnf-formula into a d-equivalent hnf-formula ψ from $\mathrm{FO}(\mathbf{Q})[\sigma]$. □

6. Runtime analysis

In this section, we prove the upper bounds on the runtime of the algorithms described in Section 5. In the proofs of this section, we use the same notation as in the description of the corresponding algorithms.

For the remainder of this section, $d \geqslant 2$ is a degree bound and σ denotes a finite relational signature.

Two basic tasks that are repeatedly used within the algorithms presented in this article are (a) to test two d-bounded r-types with n centres for isomorphism and (b) to compute the set $T_r^{d,\sigma}(n)$. Both can be accomplished by brute-force algorithms:

(a) *Testing two types for isomorphism:* Recall that the universe of each d-bounded r-type with n centres has size at most $N := n \cdot \nu_d(r)$. For every bijection between the elements of two such r-types, it can be checked in time $\mathcal{O}(n + N^{\|\sigma\|}) \leqslant (N{+}1)^{\mathcal{O}(\|\sigma\|)}$ whether the bijection is indeed an isomorphism. Since there are at most N^N bijections, the test whether the two r-types are isomorphic can be performed in at most

$$2^{\mathsf{a}_0 \cdot \|\sigma\| \cdot N^2} \tag{17}$$

time steps, where $\mathsf{a}_0 \geqslant 1$ is a suitable number of size $\mathcal{O}(\|\sigma\|)$.

(b) *Constructing $T_r^{d,\sigma}(n)$:* For computing $T_r^{d,\sigma}(n)$, we can assume that the universe of each element in $T_r^{d,\sigma}(n)$ is a subset of $\{1, \ldots, N\}$. Hence, the relations of an arbitrary r-type in $T_r^{d,\sigma}(n)$ can be represented by a bitstring of length at most $\|\sigma\|$ if $N = 1$, and $N^{\|\sigma\|}$ otherwise. Furthermore, each of the r-type's n centres is interpreted by an element from $\{1, \ldots, N\}$. Therefore, there are at most

$$N^n \cdot 2^{\max\{\|\sigma\|, N^{\|\sigma\|}\}} \quad \leqslant \quad 2^{N^2 + \max\{\|\sigma\|, N^{\|\sigma\|}\}} \quad \leqslant \quad 2^{\max\{2,N\}^{\|\sigma\|+2}}$$

candidates for types in $T_r^{d,\sigma}(n)$.

For any such candidate, we can test in at most $\max\{2,N\}^{\mathtt{b}_0}$ time steps, for a suitable number $\mathtt{b}_0 \geqslant 1$ of size $\mathcal{O}(\|\sigma\|)$, whether the candidate is indeed a d-bounded r-type. Furthermore, for any two such candidates, we use the above described algorithm for testing whether they are isomorphic. In summary, $T_r^{d,\sigma}(n)$ can be constructed in at most

$$\begin{aligned} & 2^{\max\{2,N\}^{\|\sigma\|+2}} \cdot \max\{2,N\}^{\mathtt{b}_0} + \left(2^{\max\{2,N\}^{\|\sigma\|+2}}\right)^2 \cdot 2^{\mathtt{a}_0 \cdot \|\sigma\| \cdot N^2} \\ \leqslant \; & 2^{\max\{2,N\}^{\mathtt{a}_1}} \end{aligned} \tag{18}$$

time steps, where $\mathtt{a}_1 > \max\{\mathtt{a}_0, \mathtt{b}_0\}$ is a suitable number of size $\mathcal{O}(\|\sigma\|)$.

6.1. Step (1): From FO(**Q**) *to* FO(**D**)

Proof of the third claim of Proposition 5.1. Recall that for an FO(**Q**)-formula φ we let $w(\varphi) := 1$ if φ is quantifier-free; and otherwise,

$$w(\varphi) \;:=\; \max\{\|\mathsf{Q}\| + k \,:\, (\mathsf{Q}{+}k) \text{ occurs in } \varphi\}\,.$$

The algorithm proceeds by induction on the shape of φ.

The cases of φ being quantifier-free or a Boolean combination of formulas with generalised quantifier rank $\geqslant 1$ are trivial.

Assume that φ is of shape $(\mathsf{Q}{+}k)\,y\,\varphi'$ for a quantifier $\mathsf{Q} \in \mathbf{Q} \cup \{\exists\}$ and a $k \geqslant 0$, and with generalised quantifier rank $q \geqslant 1$. In this case, the algorithm proceeds as follows:

Step (1): The algorithm is called recursively to compute an FO(**D**)[σ]-formula ψ' that is equivalent to φ' and that has displacement $\leqslant w(\varphi')$ and generalised quantifier rank $q{-}1$.

Step (2): *If* Q *is the existential quantifier*, the algorithm outputs $\psi := (\exists{+}k)\,y\,\psi'$.

Otherwise, $\mathsf{Q} \in \mathbf{Q}$, and the algorithm makes use of the formula (5) (using ψ' instead of φ) to obtain an FO(**D**)-formula ψ that is equivalent to $(\mathsf{Q}{+}k)\,y\,\psi'$.

Since the size of the formula $\exists^{=\ell}\,y\,\psi'$ is $\mathcal{O}(\ell^2 + \|\psi'\|)$, the formula ψ has size in

$$\mathcal{O}((n_0{+}k{+}p)^3 \cdot \|\psi'\|) \quad\leqslant\quad \mathcal{O}(w(\varphi)^3 \cdot \|\psi'\|)$$

where p is the period and n_0 is the smallest offset of Q (the latter inequality holds by definition of $w(\varphi)$, since $n_0 + k + p < \|\mathsf{Q}\| + k \leqslant w(\varphi)$).

Since the only step which increases the formula's size is the one for formulas of the shape $(\mathsf{Q}{+}k)\,y\,\varphi$, we get

$$\|\psi\| \quad\in\quad \|\varphi\| \cdot w(\varphi)^{\mathcal{O}(q)}\,.$$

It is easy to see that ψ can also be constructed within the same time bound. □

6.2. *Step (2): Hanf normal form for* FO($\mathbf{D}$)

6.2.1. *The base case*

Proof of the third claim of Lemma 5.2. We denote the steps of the computation in the same way as in the description of the algorithm:

Step (i): According to (18), the set $T_0^{d,\sigma}(n)$ can be computed in time

$$2^{\max\{2,n\cdot\nu_d(0)\}^{a_1}} \quad \leqslant \quad 2^{\max\{2,n\}^{\mathtt{a}_1}} \quad < \quad 2^{\|\varphi\|^{\mathtt{a}_1}},$$

for a number $\mathtt{a}_1 \geqslant 1$ of size $\mathcal{O}(\|\sigma\|)$.

Let $\tau = (\mathcal{T}, \overline{c}) \in T_0^{d,\sigma}(n)$. Since $\varphi(\overline{x})$ is quantifier-free, we can check if $\mathcal{T} \models \varphi[\overline{c}]$ in time $\mathcal{O}(\|\varphi\|)$. Since $|T_0^{d,\sigma}(n)| \leqslant 2^{\|\varphi\|^{\mathtt{a}_1}}$, the set T can therefore be computed in time $2^{\|\varphi\|^{\mathtt{a}_2}}$, where $\mathtt{a}_2 \geqslant \mathtt{a}_1$ is of size $\mathcal{O}(\|\sigma\|)$

Step (ii): From $\tau \in T_0^{d,\sigma}(n)$, one can compute the formula $\mathrm{sph}_\tau(\overline{x})$ in time $\mathcal{O}(\|\sigma\|)$ if $n = 1$, and otherwise in time $n^{\mathcal{O}(\|\sigma\|)}$.

In summary, there is a number $\mathtt{a} \geqslant \mathtt{a}_1$ of size $\mathcal{O}(\|\sigma\|)$ such that $\psi(\overline{x})$ can be constructed in $\leqslant 2^{\|\varphi\|^{\mathtt{a}}}$ time steps. □

6.2.2. *The inductive step*

Proof of the third claim of Lemma 5.3. In the following, let n be the number of free variables of φ and set $N = (n+1) \cdot \nu_d(r)$. Note that $N \geqslant 2$, since $r \geqslant 1$.

Step (i): According to (18), the set $T_r^{d,\sigma}(n+1)$ can be constructed in at most

$$2^{((n+1)\cdot\nu_d(r))^{\mathtt{a}_1}} \quad = \quad 2^{N^{\mathtt{a}_1}}$$

time steps.

Step (ii): To estimate the time needed to construct the formula (6) (given the set $T_r^{d,\sigma}(n+1)$), we need upper bounds on the number of possible values for τ, on the time used for constructing the formula $\mathrm{sph}_\tau(\overline{x}, y)$, and on the time used for constructing the sentence φ'_τ.

By (18), the number of types $\tau \in T_r^{d,\sigma}(n+1)$ is at most $2^{N^{\mathtt{a}_1}}$. Now let $\tau = (\mathcal{T}, \overline{c}, c_{n+1}) \in T_r^{d,\sigma}(n+1)$. Since the universe of τ contains at most $(n+1) \cdot \nu_d(r) = N$ elements, the formula $\mathrm{sph}_\tau(\overline{x}, y)$ can be constructed in time

$$((n+1) \cdot \nu_d(r))^{\mathcal{O}(\|\sigma\|)} \quad \leqslant \quad N^{\mathtt{a}_2}$$

for some number $\mathtt{a}_2 \geqslant 1$ of size $\mathcal{O}(\|\sigma\|)$.

To compute the hnf-sentence ψ'_τ, we have to decide for each sphere-formula $\mathrm{sph}_\rho(\overline{z})$ in $\psi'(\overline{x}, y)$ whether to replace it by $\top$ or by $\bot$. This is done by an isomorphism test which, according to (17), takes at most

$$2^{((n+1)\cdot\nu_d(r))^{\mathtt{a}_0}} \quad = \quad 2^{N^{\mathtt{a}_0}}.$$

time steps. Note that ψ' contains at most $\|\psi'\|$ occurrences of sphere-formulas. Hence, the construction of ψ'_τ can be carried out in

$$\|\psi'\| \cdot 2^{N^{\mathsf{a}_0}}$$

time steps.

In summary, the formula (6) can be obtained in time

$$2^{N^{\mathsf{a}_1}} \;+\; 2^{N^{\mathsf{a}_1}} \cdot \Big(N^{\mathsf{a}_2} + \|\psi'\| \cdot 2^{N^{\mathsf{a}_0}} \Big) \quad \leqslant \quad \|\psi'\| \cdot 2^{N^{\mathsf{a}_3}}$$

for a number $\mathsf{a}_3 \geqslant \mathsf{a}_0 + \mathsf{a}_1 + \mathsf{a}_2$ of size $\mathcal{O}(\|\sigma\|)$.

Step (iii): The construction of the formula $\xi(\overline{x})$ from the formula (7) is done recursively. Since, in every recursive call, the set M is subdivided into two sets of almost equal size, the recursion depth is $\mathcal{O}(\log(|T_r^{d,\sigma}(n{+}1)|) = \mathcal{O}(N^{\mathsf{a}_1})$. In each recursive call, the formula $(\mathsf{D}_p{+}\ell)\, y \bigvee_{\tau\in M} \alpha_\tau(\overline{x}, y)$ (with $0 \leqslant \ell \leqslant \max\{k,p\}$) is replaced by a Boolean combination of $2(\max\{\ell,p\}{+}1)^2 \leqslant 2(\max\{k,p\}{+}1)^2$ many formulas of the form $(\mathsf{D}_p{+}\ell)\, y \bigvee_{\tau\in M'} \alpha_\tau(\overline{x}, y)$ with $0 \leqslant \ell \leqslant \max\{k,p\}$. In total, the formula $\xi(\overline{x})$ is a Boolean combination of at most $(2(\max\{k,p\}{+}1)^2)^{\mathcal{O}(N^{\mathsf{a}_1})} \leqslant (2 \cdot \max\{k,p\})^{N^{\mathsf{a}_4}}$ formulas of the form (8) (for some $\mathsf{a}_4 \geqslant \mathsf{a}_1$ of size $\mathcal{O}(\|\sigma\|)$) and can be computed in time $\|\psi'\| \cdot (2 \cdot \max\{k,p\})^{N^{\mathsf{a}_4}}$.

Step (iv): The computation of the formula $\psi(\overline{x})$ from $\xi(\overline{x})$ can be carried out in time $\|\xi(\overline{x})\| \leqslant \|\psi'\| \cdot (2 \cdot \max\{k,p\})^{N^{\mathsf{a}_4}}$

In summary, the construction of ψ from $\varphi = (\mathsf{Q}{+}k)\, y\, \psi'$ can be carried out in time

$$2^{N^{\mathsf{a}_1}} + \|\psi'\| \cdot 2^{N^{\mathsf{a}_3}} + \|\psi'\| \cdot \max\{2k, 2p\}^{N^{\mathsf{a}_4}} + \|\psi'\| \cdot \max\{2k, 2p\}^{N^{\mathsf{a}_4}}$$

which is at most

$$\|\psi'\| \cdot \max\{2k, 2p\}^{N^{\mathsf{a}_5}}$$

for some a_5 of size $\mathcal{O}(\|\sigma\|)$. □

Proof of the third claim of Lemma 5.4. Let $N := n \cdot \nu_d(4r)$. Note that $N \geqslant 2$.

According to (18) the set $T_{4r}^{d,\sigma}(n)$ can be constructed in time $2^{N^{\mathsf{a}_1}}$, for a suitable number $\mathsf{a}_1 \geqslant 1$ of size $\mathcal{O}(\|\sigma\|)$.

For each $\rho \in T_{4r}^{d,\sigma}(n)$, it takes time $N^{\mathcal{O}(\|\sigma\|)}$ to construct the formula $\mathrm{sph}_\rho(\overline{x})$. For the construction of β_ρ, we have to make a case distinction depending on whether $c_{n+1} \in N_{2r+1}^{\tau}(\overline{c})$. To determine which of the two cases actually applies for the given τ, recall that the universe of τ has size at most $(n{+}1) \cdot \nu_d(r)$. Thus, the time needed to decide whether $c_{n+1} \in N_{2r+1}^{\tau}(\overline{c})$ is in $N^{\mathcal{O}(\|\sigma\|)}$.

Case (1): *If* $c_{n+1} \in N^{\tau}_{2r+1}(\overline{c})$, we compute the number k_ρ defined in (11). This requires us to check for at most $n \cdot \nu_d(2r{+}1) \leqslant N$ d-bounded r-types with $(n{+}1)$ centres whether they are isomorphic to τ. Since each of these types has a universe of size $\leqslant (n{+}1) \cdot \nu_d(r) \leqslant N$, the number k_ρ can be computed in time $2^{\mathcal{O}(\|\sigma\| \cdot N^2)}$ by using the brute-force isomorphism test described at the beginning of Section 6. For ultimately periodic R, testing whether $k_\rho \in (\mathsf{R}{+}\ell)$ is done in time $\mathcal{O}(k_\rho) \subseteq \mathcal{O}(N)$ since $k_\rho \leqslant N$. Hence, in Case (1), the algorithm uses time $2^{\mathcal{O}(\|\sigma\| \cdot N^2)}$ to construct the formula β_ρ.

Case (2): *If* $c_{n+1} \notin N^{\tau}_{2r+1}(\overline{c})$, we have to check whether the two r-types $\mathcal{N}^{\tau}_{r}(\overline{c})$ and $\mathcal{N}^{\rho}_{r}(\overline{e})$, which each have a universe of size $\leqslant N$, are isomorphic.

Case (2.a): *If they are not isomorphic*, we need time $\mathcal{O}(1)$ to check whether $0 \in (\mathsf{R}{+}\ell)$ (which is the case iff $\ell = 0$ and $0 \in \mathsf{R}$) and construct β_ρ accordingly in time $\mathcal{O}(1)$.

Case (2.b): *If they are isomorphic*, we have to compute the number ℓ_ρ defined in (13). This requires us to check for at most $n \cdot \nu_d(2r{+}1) \leqslant N$ d-bounded r-types with one centre, each with a universe of at most $\nu_d(r) \leqslant N$ elements, whether they are isomorphic to $\mathcal{N}^{\tau}_{r}(c_{n+1})$. By using the brute-force isomorphism test described at the beginning of Section 6, the number ℓ_ρ can be computed in time $2^{\mathcal{O}(\|\sigma\| \cdot N^2)}$. Afterwards, we construct the formula

$$\beta_\rho \ := \quad (\mathsf{R}{+}(\ell{+}\ell_\rho))\, y\ \mathrm{sph}_{\mathcal{N}^{\tau}_{r}(c_{n+1})}(y)\,.$$

For this, we use time $N^{\mathcal{O}(\|\sigma\|)}$ to construct the formula $\mathrm{sph}_{\mathcal{N}^{\tau}_{r}(c_{n+1})}(y)$. Resolving the quantification $(\mathsf{R}{+}(\ell{+}\ell_\rho))$ via the quantifiers R and $\exists$, as described in Section 2.3.2, needs time in

$$\begin{aligned} & \mathcal{O}((\ell{+}\ell_\rho)^2 + \|\mathsf{R}\, y\ \mathrm{sph}_{\mathcal{N}^{\tau}_{r}(c_{n+1})}(y)\|) \\ \leqslant \quad & (\ell^2{+}\|\mathsf{R}\|) \cdot N^{\mathcal{O}(\|\sigma\|)}. \end{aligned}$$

Altogether, we obtain that in Case (2) the algorithm uses time at most $(\ell^2{+}\|\mathsf{R}\|) \cdot 2^{N^{\mathcal{O}(\|\sigma\|)}}$ to construct the formula β_ρ.

Finally, the algorithm outputs the formula $\beta(\overline{x})$, which is the disjunction of the formulas $(\mathrm{sph}_\rho(\overline{x}) \wedge \beta_\rho)$, for all $\rho \in T^{d,\sigma}_{4r}(n)$. The overall time used for constructing this formula is

$$2^{N^{\mathsf{a}_1}} \ + \ 2^{N^{\mathsf{a}_1}} \cdot \big(N^{\mathcal{O}(\|\sigma\|)} + (\ell^2{+}\|\mathsf{R}\|) \cdot 2^{N^{\mathcal{O}(\|\sigma\|)}}\big) \quad \leqslant \quad (\ell^2{+}\|\mathsf{R}\|) \cdot 2^{N^{\mathcal{O}(\|\sigma\|)}}.$$

This completes the proof of Lemma 5.4. ☐

Proof of the third claim of Lemma 5.6. According to Lemma 5.3, the formula ψ_1 can be computed in time

$$\|\psi'\| \cdot \max\{2k, 2p\}^{((n+1)\nu_d(r))^{\mathcal{O}(\|\sigma\|)}}$$

(which also bounds the size of ψ_1) and its displacement is bounded by $\max\{K, k, p\}$.

If $n = 0$, then the algorithm is finished with the construction of $\psi := \psi_1$.

Otherwise, the algorithm from Lemma 5.4 is called at most $\|\psi_1\|$ times. Each call applies to a counting-formula of locality radius $\leqslant r$, displacement $\leqslant \max\{k, p\}$, and with a quantifier $\mathsf{R} \in \mathbf{D} \cup \{\exists\}$ of period $p \geqslant 1$ whose encoding *rep*(R) has length $\leqslant p{+}2$. Therefore, the call is completed in time

$$(\max\{k, p\}^2 {+} \|\mathsf{R}\|) \cdot 2^{(n \cdot \nu_d(4r))^{\mathcal{O}(\|\sigma\|)}} \quad \leqslant \quad \max\{k, p\}^2 \cdot 2^{(n \cdot \nu_d(4r))^{\mathcal{O}(\|\sigma\|)}}.$$

In summary, the number of time steps required for the algorithm is at most

$$\begin{aligned} & \|\psi'\| \cdot \max\{2k, 2p\}^{((n+1)\nu_d(r))^{\mathcal{O}(\|\sigma\|)}} \cdot \max\{k, p\}^2 \cdot 2^{(n \cdot \nu_d(4r))^{\mathcal{O}(\|\sigma\|)}} \\ \leqslant \quad & \|\psi'\| \cdot \max\{2k, 2p\}^{((n+1)\nu_d(4r))^{\mathcal{O}(\|\sigma\|)}}. \end{aligned}$$

□

6.2.3. Putting things together: construction of hnf-formulas in FO($\mathbf{D}$)

Proof of the third claim of Proposition 5.7. Let $\mathsf{a}, \mathsf{b} > 0$ be the numbers from the Lemmas 5.2 and 5.6, respectively, and set $\mathsf{c} = \max\{\mathsf{a}, \mathsf{b} + 1\}$. Note that $\mathsf{c} > 1$ is of size $\mathcal{O}(\|\sigma\|)$.

We extend the inductive invariant stated in the proof of the first two claims of Proposition 5.7 (in Section 5.2.3) by

(d) *The algorithm terminates after at most* $\max\{2K, 2P\}^{(\|\varphi\| \cdot \nu_d(4^q))^{\mathsf{c}}}$ *time steps.*

The cases of φ being quantifier-free or a Boolean combination of formulas with generalised quantifier rank $\geqslant 1$ are trivial.

Assume that φ is of the form $(\mathsf{Q}{+}k)\, y\, \varphi'$ for a quantifier $\mathsf{Q} \in \mathbf{D} \cup \{\exists\}$ and $k \geqslant 0$, and with generalised quantifier rank $q \geqslant 1$. By the induction hypothesis (d), the formula ψ' is computed in time

$$\max\{2K, 2P\}^{(\|\varphi'\| \cdot \nu_d(4^{q-1}))^{\mathsf{c}}}$$

which is also an upper bound on the size $\|\psi'\|$ of ψ'. Furthermore, ψ' has locality radius $\leqslant 4^{q-1}$. We apply the algorithm from Lemma 5.6 to the formula $(\mathsf{Q}{+}k)\, y\, \psi'$. This algorithm requires at most

$$\begin{aligned} & \|\psi'\| \cdot \max\{2k, 2p\}^{((n+1) \cdot \nu_d(4 \cdot 4^{q-1}))^{\mathsf{b}}} \\ \leqslant \quad & \max\{2K, 2P\}^{(\|\varphi'\| \cdot \nu_d(4^{q-1}))^{\mathsf{c}} + ((n+1) \cdot \nu_d(4^q))^{\mathsf{b}}} \end{aligned}$$

time steps. Let $N := \|\varphi\| \cdot \nu_d(4^q)$. Then we have

$$\begin{aligned} & (\|\varphi'\| \cdot \nu_d(4^{q-1}))^{\mathsf{c}} \ + \ ((n{+}1) \cdot \nu_d(4^q))^{\mathsf{b}} \\ \leqslant \quad & (\|\varphi\| \cdot \nu_d(4^q) - 1)^{\mathsf{c}} \ + \ (\|\varphi\| \cdot \nu_d(4^q))^{\mathsf{c}-1} \qquad = \quad (N-1)^{\mathsf{c}} + N^{\mathsf{c}-1} \\ = \quad & N^{\mathsf{c}} \cdot \Big(\big(\tfrac{N-1}{N}\big)^{\mathsf{c}} + \tfrac{1}{N} \Big) \\ \leqslant \quad & N^{\mathsf{c}} \cdot \big(\tfrac{N-1}{N} + \tfrac{1}{N} \big) \qquad = \quad (\|\varphi\| \cdot \nu_d(4^q))^{\mathsf{c}}. \end{aligned}$$

Consequently, the algorithm terminates after at most

$$\max\{2K, 2P\}^{(\|\varphi\|\cdot\nu_d(4^q))^c}$$

time steps. □

6.3. Summary: Proof of Theorem 3.4

We can now prove the upper bound on the runtime of the algorithm for the construction of Hanf normal forms for FO($\mathbf{Q}$), for sets $\mathbf{Q}$ of ultimately periodic quantifiers, as stated in Theorem 3.4.

Proof of Theorem 3.4. We analyse the algorithm's runtime when receiving as input a degree bound $d \geqslant 2$, a signature σ, and a formula $\varphi \in \mathrm{FO}(\mathbf{U})[\sigma]$ with $\|\sigma\| \leqslant \|\varphi\|$. Let $q \geqslant 0$ be the quantifier rank of φ. Note that the generalised quantifier rank of φ is $\leqslant q$. Let $\mathbf{Q} \subseteq \mathbf{U}$ be the set of unary counting quantifiers that occur in φ and let $\mathbf{D} = \{\mathsf{D}_p : p \geqslant 1 \text{ and there is a } \mathsf{Q} \in \mathbf{Q} \text{ of period } p\}$. Finally, let $P = \max\{p : p = 1 \text{ or } \mathsf{D}_p \in \mathbf{D}\}$. Note that $P \leqslant \|\varphi\|$.

Step (1): The algorithm of Proposition 5.1 computes an FO($\mathbf{D}$)[σ]-formula φ_1 that is equivalent to φ and that has displacement $K \leqslant w(\varphi)$ and generalised quantifier rank q. This takes time at most

$$\|\varphi\| \cdot w(\varphi)^{\mathcal{O}(q)} \quad \leqslant \quad 2^{\mathrm{poly}(\|\varphi\|)}$$

For the latter inequality recall that $q \leqslant \|\varphi\|$ and $w(\varphi) \leqslant \|\varphi\|$.

Step (2): The algorithm of Proposition 5.7 computes an hnf-formula $\psi_1 \in \mathrm{FO}(\mathbf{D})[\sigma]$ that is d-equivalent to φ_1 and of locality radius $\leqslant 4^q$. Since $\|\sigma\|, K, P, q \leqslant \|\varphi\|$ and $\|\varphi_1\| \leqslant 2^{\mathrm{poly}(\|\varphi\|)}$ this takes time at most

$$\max\{2K, 2P\}^{(\|\varphi_1\|\cdot\nu_d(4^q))^{\mathcal{O}(\|\sigma\|)}} \quad \leqslant \quad 2^{\left(2^{\mathrm{poly}(\|\varphi\|)}\cdot\nu_d(4^{\|\varphi\|})\right)^{\mathcal{O}(\|\varphi\|)}}.$$

For the same reason, the displacement of ψ_1 is bounded by

$$\max\{K, P\} + \|\varphi_1\| \cdot \nu_d(4^q) \quad \leqslant \quad 2^{\mathrm{poly}(\|\varphi\|)} \cdot \nu_d(4^{\|\varphi\|}).$$

We now consider the special cases $d = 2$ and $d \geqslant 3$.

For $d = 2$, we have $\nu_d(r) \leqslant 2r+1$ for every $r \geqslant 0$. Hence, the formula ψ_1 can be computed in time

$$2^{\left(2^{\mathrm{poly}(\|\varphi\|)}\cdot(2\cdot 4^{\|\varphi\|}+1)\right)^{\mathcal{O}(\|\varphi\|)}} \quad \leqslant \quad 2^{2^{\mathrm{poly}(\|\varphi\|)}}$$

and has displacement in

$$2^{\mathrm{poly}(\|\varphi\|)} \cdot (2 \cdot 4^{\|\varphi\|} + 1) \quad \leqslant \quad 2^{\mathrm{poly}(\|\varphi\|)}.$$

For $d \geqslant 3$, we have $\nu_d(r) \leqslant d^{r+1}$ for every $r \geqslant 0$. Hence, the formula ψ_1 can be computed in time

$$2^{\left(2^{\text{poly}(\|\varphi\|)} \cdot d^{4^{\|\varphi\|}+1}\right)^{\mathcal{O}(\|\varphi\|)}} \quad \leqslant \quad 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}}.$$

and has displacement in

$$2^{\text{poly}(\|\varphi\|)} \cdot d^{4^{\|\varphi\|}+1} \quad \leqslant \quad d^{2^{\mathcal{O}(\|\varphi\|)}}.$$

Step (3): Finally, the algorithm of Proposition 5.9 computes for each counting-sentence $(\mathsf{D}_p{+}k)\, y\, \gamma$ in ψ_1 an equivalent hnf-sentence from $\text{FO}(\mathbf{Q})[\sigma]$. The displacement of the constructed hnf-formula is at most

$$k + \|\mathsf{Q}\| + p \quad \leqslant \quad 2^{\text{poly}(\|\varphi\|)} \cdot \nu_d(4^{\|\varphi\|})$$

where $\mathsf{Q} \in \mathbf{Q}$ is ultimately periodic with period p. Furthermore, this takes time at most

$$\mathcal{O}((\|\mathsf{Q}\|{+}k)^3 \cdot \|\gamma\|) \quad \leqslant \quad \mathcal{O}((\|\varphi\|{+}k)^3 \cdot \|\psi_1\|)\,.$$

Hence, the construction of the hnf-formula ψ from ψ_1 takes time

$$\begin{aligned} & \mathcal{O}(\|\psi_1\| \cdot (\|\varphi\|{+}2^{\text{poly}(\|\varphi\|)} \cdot \nu_d(4^{\|\varphi\|}))^3 \cdot \|\psi_1\|) \\ = \quad & \mathcal{O}(\|\psi_1\|^2 \cdot 2^{\text{poly}(\|\varphi\|)} \cdot \nu_d(4^{\|\varphi\|})^3)\,. \end{aligned}$$

Note that ψ has the same locality radius $\leqslant 4^q$ as ψ_1 and displacement

$$2^{\text{poly}(\|\varphi\|)} \cdot \nu_d(4^{\|\varphi\|})\,.$$

Again, we distinguish the cases $d = 2$ and $d \geqslant 3$.

For $d = 2$, $\nu_d(r) \leqslant 2r{+}1$ (for all $r \geqslant 0$) and hence ψ has displacement $\leqslant 2^{\text{poly}(\|\varphi\|)}$ and is constructed in time

$$2^{2^{\text{poly}(\|\varphi\|)}}.$$

For $d \geqslant 3$, $\nu_d(r) \leqslant d^{r+1}$ (for all $r \geqslant 0$) and hence ψ has displacement $\leqslant d^{2^{\mathcal{O}(\|\varphi\|)}}$ and is constructed in time

$$2^{d^{2^{\mathcal{O}(\|\varphi\|)}}}.$$

Altogether, for a degree bound $d \geqslant 2$, ψ is an hnf-formula from $\text{FO}(\mathbf{Q})[\sigma]$ that is d-equivalent to φ and can be computed in time

$$2^{2^{\text{poly}(\|\varphi\|)}} \quad \text{for } d = 2, \qquad \text{and} \qquad 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}} \quad \text{for } d \geqslant 3.$$

This completes the proof of Theorem 3.4. □

7. Model-checking

This section is devoted to the proof of Theorem 3.5.

Proof of Theorem 3.5.
Let $\varphi(\overline{x})$, $\mathcal{A}$, and $\overline{a}$ be the algorithm's input, where σ is the relational signature that consists of precisely the relation symbols occurring in φ, and $\mathcal{A}$ is a σ-structure. Let $\mathbf{Q} \subseteq \mathbf{U}$ be the set of (ultimately periodic) unary counting quantifiers that occur in φ.

For checking whether $\mathcal{A} \models \varphi[\overline{a}]$, the algorithm proceeds as follows:

Step (1): Compute an upper bound $d \geqslant 2$ on the degree of $\mathcal{A}$.

Step (2): Use the algorithm from Theorem 3.4 to transform $\varphi(\overline{x})$ into a d-equivalent $\mathrm{FO}(\mathbf{Q})[\sigma]$-formula $\psi(\overline{x})$ in Hanf normal form.

Step (3): For each sphere-formula α that occurs in ψ, check if $\mathcal{A} \models \alpha[\overline{a}]$, and replace each occurrence of α in ψ with the Boolean constant **1** if $\mathcal{A} \models \alpha[\overline{a}]$, and with the Boolean constant **0** otherwise.

Step (4): For each counting-sentence χ that occurs in ψ, check if $\mathcal{A} \models \chi$, and replace each occurrence of χ in ψ with the Boolean constant **1** if $\mathcal{A} \models \chi$, and with the Boolean constant **0** otherwise.

Step (5): After having performed the steps (1)–(4), ψ is a Boolean combination of the Boolean constants **0** and **1**. Evaluate this Boolean combination and output "yes" if the result is **1**, and output "no" if the result is **0**.

Obviously, the algorithm's output is "yes" if, and only if, $\mathcal{A} \models \varphi[\overline{a}]$. For analysing the algorithm's runtime, we use the same conventions as in [7]. I.e., we use random-access machines with a uniform cost measure, and the input structure $\mathcal{A}$ is given by an adjacency list representation of size linear in

$$\|\mathcal{A}\| \quad := \quad \|\sigma\| \;+\; |A| \;+\; \sum_{R \in \sigma} |R^{\mathcal{A}}| \cdot \mathrm{ar}(R).$$

We let $n := |\overline{x}|$ and let $\overline{x} = (x_1, \ldots, x_n)$ and $\overline{a} = (a_1, \ldots, a_n)$. For the following runtime analysis of each of the steps (1)–(5), note that n, $\|\sigma\|$, and $\|\mathsf{Q}\|$ are smaller than $\|\varphi\|$, for each $\mathsf{Q} \in \mathbf{Q}$.

Step (1): To compute d, compute an adjacency list representation of $\mathcal{A}$'s Gaifman-graph $G_{\mathcal{A}}$: For each $R \in \sigma$ and each occurrence of an element of A in a tuple of $R^{\mathcal{A}}$, we have to add at most $\mathrm{ar}(R) \leqslant \|\sigma\|$ edges to $G_{\mathcal{A}}$. For each edge, this takes time $\mathcal{O}(d)$, since $\mathcal{A}$ is d-bounded. In summary, computing $G_{\mathcal{A}}$ and d takes time at most $\mathcal{O}(\|\mathcal{A}\| \cdot \|\sigma\| \cdot d)$.

Step (2): According to Theorem 3.4, this takes time at most

$$2^{2^{\mathrm{poly}(\|\varphi\|)}} \quad \text{for } d = 2, \quad \text{and} \quad 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}} \quad \text{for } d \geqslant 3.$$

Recall that $\psi(\overline{x})$ is a Boolean combination of sphere-formulas of the form $\mathrm{sph}_\rho(\overline{x}')$ and counting-sentences $(\mathsf{Q}{+}k)\,y\ \mathrm{sph}_\tau(y)$, where $\overline{x}'$ is a tuple of free variables from $\overline{x}$, ρ is a d-bounded type with radius $\leqslant 4^{\mathrm{qr}(\varphi)}$ and $|\overline{x}'|$ centres, $\mathsf{Q} \in \mathbf{Q}$, $k \in \mathbb{N}$, and τ is a d-bounded type with radius $\leqslant 4^{\mathrm{qr}(\varphi)}$ and one centre.

Step (3): Consider a sphere-formula $\alpha(\overline{x}') := \mathrm{sph}_\rho(\overline{x}')$, where $\overline{x}' = (x_{i_1}, \ldots, x_{i_m})$ for $1 \leqslant m \leqslant n$ and $i_1, \ldots, i_m \in [1, n]$, and where ρ is a d-bounded type with radius $r \leqslant 4^{\mathrm{qr}(\varphi)}$ and m centres.

To decide whether $\mathcal{A} \models \alpha[\overline{a}']$, for $\overline{a}' = (a_{i_1}, \ldots, a_{i_m})$, we have to check whether $\mathcal{N}_r^{\mathcal{A}}(\overline{a}') \cong \rho$. Note that the universe of $\mathcal{N}_r^{\mathcal{A}}(\overline{a}')$ has size at most

$$N \;:=\; m \cdot \nu_d(r) \;\leqslant\; \|\varphi\| \cdot \nu_d(4^{\mathrm{qr}(\varphi)}),$$

i.e., size at most

$$2^{\mathcal{O}(\|\varphi\|)} \quad \text{for } d = 2, \quad \text{and} \quad d^{2^{\mathcal{O}(\|\varphi\|)}} \quad \text{for } d \geqslant 3.$$

Furthermore, the structure $\mathcal{N}_r^{\mathcal{A}}(\overline{a}')$ can be computed within the same time bound, which we will denote by $t_1(\|\varphi\|)$ in the following. As explained at the beginning of Section 6, checking if $\mathcal{N}_r^{\mathcal{A}}(\overline{a}') \cong \rho$ can be done in time $t_2(\|\varphi\|) := 2^{\mathcal{O}(\|\sigma\| \cdot N^2)}$, which is at most

$$2^{2^{\mathcal{O}(\|\varphi\|)}} \quad \text{for } d = 2, \quad \text{and} \quad 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}} \quad \text{for } d \geqslant 3.$$

Since there are at most $\|\psi\|$ sphere-formulas in ψ, the entire Step (3) takes time at most $\|\psi\| \cdot (t_1(\|\varphi\|) + t_2(\|\varphi\|))$, i.e., time at most

$$2^{2^{\mathrm{poly}(\|\varphi\|)}} \quad \text{for } d = 2, \quad \text{and} \quad 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}} \quad \text{for } d \geqslant 3.$$

Step (4): Consider a counting-sentence $\chi := (\mathsf{Q}{+}k)\,y\ \mathrm{sph}_\tau(y)$ that occurs in ψ. In particular, τ is a d-bounded type with radius $r \leqslant 4^{\mathrm{qr}(\varphi)}$ and with one centre. To decide whether $\mathcal{A} \models \chi$, we first compute the number k_τ of elements $a \in A$ with $\mathcal{N}_r^{\mathcal{A}}(a) \cong \tau$, and then we check if $k_\tau \in (\mathsf{Q}{+}k)$. The latter can be done easily, as Q is ultimately periodic. To compute k_τ, we consider every $a \in A$, compute $\mathcal{N}_r^{\mathcal{A}}(a)$, and check whether $\mathcal{N}_r^{\mathcal{A}}(a) \cong \tau$. From Step (3) we know that for each $a \in A$ this can be done in time $t_1(\|\varphi\|) + t_2(\|\varphi\|)$. Since there are at most $\|\psi\|$ counting-sentences in ψ, the entire Step (4) takes time at most

$$\|\psi\| \cdot |A| \cdot (t_1(\|\varphi\|) + t_2(\|\varphi\|)),$$

i.e., time at most

$$2^{2^{\mathrm{poly}(\|\varphi\|)}} \cdot |A| \quad \text{for } d = 2, \quad \text{and} \quad 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}} \cdot |A| \quad \text{for } d \geqslant 3.$$

Step (5): Evaluating the resulting variable-free Boolean expression takes time polynomial in the length of this expression, i.e., time polynomial in $\|\psi\|$.

In summary, the total running time of the algorithm is

$$2^{2^{\mathrm{poly}(\|\varphi\|)}} \cdot \|\mathcal{A}\| \quad \text{for } d=2, \quad \text{and} \qquad 2^{d^{2^{\mathcal{O}(\|\varphi\|)}}} \cdot \|\mathcal{A}\| \quad \text{for } d \geqslant 3.$$

This completes the proof of Theorem 3.5. □

8. Lower bound

In this section we show that the algorithm of Theorem 3.4 is worst-case optimal for each $d \geqslant 2$, even for the special case of an input formula from plain first-order logic. Our proofs follow the basic idea of the lower bound from [2], which covers the case $d = 3$ for output formulas from plain first-order logic. This section's main result reads as follows.

Theorem 8.1. *Let $\sigma = \{E, P\}$ be a signature consisting of a binary relation symbol E and a unary relation symbol P. Let $d \in \mathbb{N}$ with $d \geqslant 2$.*

There is no algorithm that computes, upon input of an $\mathrm{FO}[\sigma]$*-sentence φ, a d-equivalent hnf-sentence ψ in* $\mathrm{FO}(\mathbf{U})[\sigma]$ *in time at most*

$$2^{2^{o(\|\varphi\|)}} \quad \textit{if} \;\; d=2, \qquad \textit{and} \quad 2^{d^{2^{o(\|\varphi\|)}}} \quad \textit{if} \;\; d \geqslant 3\,.$$

As an intermediate step for proving the theorem, we will consider signatures σ_d that depend on the degree bound d. For the remainder of this section, let P be a unary relation symbol and let $S_1, S_2, \ldots$ be pairwise distinct binary relation symbols. For each $d \geqslant 2$ let $\sigma_d := \{S_1, \ldots, S_{d-1}, P\}$. We will construct suitable sequences of "small" $\mathrm{FO}[\sigma_d]$-formulas for which we can show lower bounds on the size of d-equivalent hnf-sentences. These formulas are described in the proof of the following Theorem 8.2, for which we need some more notation.

Let $d \in \mathbb{N}$ with $d \geqslant 2$. For a σ_d-structure $\mathcal{A}$ and elements $a, b \in A$ with $(a, b) \in S_i^{\mathcal{A}}$ for some $i \in [1, d)$ we say that b is an *S_i-successor* of a. A *directed path* (of *length* $\ell \in \mathbb{N}$) in $\mathcal{A}$ (from node a_0 to node a_ℓ) is a sequence $(a_0, \ldots, a_\ell) \in A^{\ell+1}$ such that for each $j \in [0, \ell)$ there is an $i \in [1, d)$ with $(a_j, a_{j+1}) \in S_i^{\mathcal{A}}$. The structure $\mathcal{A}$ is called *acyclic* if it contains no node $a \in A$ such there is a directed path from a to a of length $\ell \geqslant 1$ in $\mathcal{A}$. For $a, b \in A$ we say that b is *reachable* from a in $\mathcal{A}$ if there is a directed path (of some length $\ell \geqslant 0$) from a to b in $\mathcal{A}$. We write $\mathrm{Reach}_a^{\mathcal{A}}$ to denote the set of all nodes $b \in A$ that are reachable from a in $\mathcal{A}$.

A *labelled and ordered $(d-1)$-ary tree* is a σ_d-structure $\mathcal{A}$ with the following properties:

1. $\mathcal{A}$ is acyclic,
2. there is a unique node $a_0 \in A$ (called the *root* of $\mathcal{A}$) such that every node $b \in A$ is reachable from a_0 in $\mathcal{A}$,
3. for every $b \in A$ with $b \neq a_0$ there is a unique $i \in [1, d)$ and a unique $a \in A$ such that $(a, b) \in S_i^{\mathcal{A}}$, and

4. for every $a \in A$ and every $i \in [1,d)$ there exists at most one $b \in A$ with $(a,b) \in S_i^{\mathcal{A}}$.

The *height* of a node a in a labelled and ordered $(d{-}1)$-ary tree $\mathcal{A}$ is the length of the unique path from the root to a. A *leaf* is a node without any successor. A node is called *full* if it has an S_i-successor for each $i \in [1,d)$. $\mathcal{A}$ is said to be *complete* with height $h \in \mathbb{N}$ if each of its nodes is either full or a leaf, and all leaves have height h.

For each $h \geqslant 0$, we let $\mathfrak{T}_{d,h}$ be the set of all (up to isomorphism) *complete* labelled and ordered $(d{-}1)$-ary trees with height *precisely* h. Clearly, every structure in $\mathfrak{T}_{d,h}$ has degree at most d.

Since $\mathfrak{T}_{2,h}$ is the set of all (up to isomorphism) labelled paths of length h, the cardinality of $\mathfrak{T}_{2,h}$ grows exponentially with h. More precisely, the universe of each structure in $\mathfrak{T}_{2,h}$ consists of $n_{2,h} := h{+}1$ elements, and there are exactly $2^{n_{2,h}}$ pairwise non-isomorphic complete labelled and ordered 1-ary trees with height h. Therefore,

$$|\mathfrak{T}_{2,h}| \quad = \quad 2^{h+1}\,. \tag{19}$$

For each $d \geqslant 3$, the cardinality of $\mathfrak{T}_{d,h}$ grows 2-fold exponentially with h. More precisely, the universe of each structure in $\mathfrak{T}_{d,h}$ consists of $n_{d,h} := \sum_{i=0}^{h}(d{-}1)^i$ elements, and there are exactly $2^{n_{d,h}}$ pairwise non-isomorphic complete labelled and ordered $(d{-}1)$-ary trees with height h. Therefore, $|\mathfrak{T}_{d,h}| = 2^{n_{d,h}}$, and since $(d{-}1)^h \leqslant n_{d,h} < (d{-}1)^{h+1}$, we have

$$2^{(d-1)^h} \quad \leqslant \quad |\mathfrak{T}_{d,h}| \quad \leqslant \quad 2^{(d-1)^{h+1}} \qquad \text{for each } \ d \geqslant 3\,. \tag{20}$$

This section's technical main result, from which we will infer Theorem 8.1, is the following theorem.

Theorem 8.2. *Let $d \in \mathbb{N}$ with $d \geqslant 2$. There is a number $c_d \geqslant 1$ of size $\mathcal{O}(d^2)$ and a sequence $(\varphi_{d,m})_{m \geqslant 1}$ of $\mathrm{FO}[\sigma_d]$-sentences such that the following holds for every $m \geqslant 1$:*

(1) $\varphi_{d,m}$ has size at most $c_d \cdot m$, and

(2) every hnf-sentence in $\mathrm{FO}(\mathbf{U})[\sigma_d]$ that is d-equivalent to $\varphi_{d,m}$ has size at least $|\mathfrak{T}_{d,2^m}|$.

Note that this immediately implies the statement of Theorem 8.1 for $d = 2$ (with $E := S_1$). To also obtain the statement of Theorem 8.1 for $d \geqslant 3$, we will interpret σ_d-structures in structures of signature $\sigma = \{E, P\}$ and modify the sentences provided by Theorem 8.2 accordingly.

The rest of this section is organised as follows. In Section 8.1 we prove a lemma that provides the "combinatorial essence", which is used for proving Theorem 8.2 in Section 8.2. Section 8.3 uses a standard interpretation argument to transfer the statement of Theorem 8.2 to a fixed signature $\sigma = \{E, P\}$ that is independent of d. The proof of Theorem 8.1 is then presented in Section 8.4.

8.1. The combinatorial essence used for proving Theorem 8.2

Let σ be a relational signature. We call a set $\mathfrak{C}$ of σ-structures *substructure-free* if for every $\mathcal{A}$ in $\mathfrak{C}$ and every $B \subseteq A$ with $\emptyset \neq B \neq A$, the induced substructure $\mathcal{A}[B]$ is not isomorphic to any structure in $\mathfrak{C}$. Observe that the set $\mathfrak{T}_{d,h}$ is substructure-free for each $d \geqslant 2$ and each $h \geqslant 0$.

A class $\mathfrak{C}$ of σ-structures is said to be *closed under induced substructures* if for all structures $\mathcal{A}$ in $\mathfrak{C}$ and all $\emptyset \neq B \subseteq A$, $\mathfrak{C}$ also contains a structure that is isomorphic to the induced substructure $\mathcal{A}[B]$. The *disjoint union* of two σ-structures $\mathcal{A}_1$ and $\mathcal{A}_2$ is defined as the σ-structure $\mathcal{B} := \mathcal{A}_1 \oplus \mathcal{A}_2$ with universe $B = (A_1 \times \{1\}) \cup (A_2 \times \{2\})$ and where $R^{\mathcal{B}} = \{\big((a_1,1),\ldots,(a_r,1)\big) : (a_1,\ldots,a_r) \in R^{\mathcal{A}_1}\} \cup \{\big((a_1,2),\ldots,(a_r,2)\big) : (a_1,\ldots,a_r) \in R^{\mathcal{A}_2}\}$ for every $R \in \sigma$ and $r := \operatorname{ar}(R)$. A class $\mathfrak{C}$ of σ-structures is said to be *closed under disjoint unions* if for all σ-structures $\mathcal{A}_1, \mathcal{A}_2$ in $\mathfrak{C}$, $\mathfrak{C}$ also contains a structure that is isomorphic to $\mathcal{A}_1 \oplus \mathcal{A}_2$. For a positive integer m and a σ-structure $\mathcal{A}$ we write $m\mathcal{A}$ to denote the disjoint union of m copies of $\mathcal{A}$.

A σ-structure is called *connected* if its Gaifman graph is connected. Note that for every (finite) σ-structure $\mathcal{A}$ there exists a finite set S of pairwise non-isomorphic *connected* σ-structures and a mapping $m : S \to \mathbb{N}_{\geqslant 1}$ such that $\mathcal{A}$ is isomorphic to the disjoint union of the structures $m(\mathcal{B})\mathcal{B}$ for all $\mathcal{B} \in S$. The set S and the mapping m are unique up to isomorphism. In the following, we will write $S_{\mathcal{A}}$ and $m_{\mathcal{A}}$ to denote the set S and the mapping m associated with a σ-structure $\mathcal{A}$, and we will identify $\mathcal{A}$ with the structure $\oplus_{\mathcal{B} \in S_{\mathcal{A}}} m_{\mathcal{A}}(\mathcal{B})\mathcal{B}$.

For a σ-structure $\mathcal{A}$ and a connected σ-structure $\mathcal{C}$, the *number of disjoint copies of $\mathcal{C}$ in $\mathcal{A}$* is defined to be 0 if $S_{\mathcal{A}}$ contains no structure that is isomorphic to $\mathcal{C}$, and $m_{\mathcal{A}}(\mathcal{C}')$ if $S_{\mathcal{A}}$ contains a structure $\mathcal{C}'$ that is isomorphic to $\mathcal{C}$. For an integer i we say that $\mathcal{A}$ *contains at most (at least, exactly) i disjoint copies of $\mathcal{C}$* if the number of disjoint copies of $\mathcal{C}$ in $\mathcal{A}$ is at most (at least, exactly) i.

Lemma 8.3. *Let σ be a finite relational signature and let $\mathfrak{D}$ be a class of σ-structures that is closed under disjoint unions and induced substructures. Let $\mathfrak{C}$ be a finite substructure-free subset of $\mathfrak{D}$ such that each structure in $\mathfrak{C}$ is connected and has a universe of size at least 2, and the structures in $\mathfrak{C}$ are pairwise non-isomorphic.*

Suppose that ψ is an hnf-sentence in $\mathrm{FO}(\mathbf{U})[\sigma]$ *such that for every $\mathcal{A} \in \mathfrak{D}$,*

$$\mathcal{A} \models \psi \iff \mathcal{A} \textit{ contains at most one disjoint copy of each structure from } \mathfrak{C}.$$

Then, ψ contains at least $|\mathfrak{C}|$ counting-sentences and hence, $\|\psi\| > |\mathfrak{C}|$.

Proof. Let ψ be such an hnf-sentence in $\mathrm{FO}(\mathbf{U})[\sigma]$ and suppose for contradiction that ψ contains less than $|\mathfrak{C}|$ counting-sentences. Let $(\mathsf{Q}_1{+}k_1)\, y\, \mathrm{sph}_{\tau_1}(y), \ldots, (\mathsf{Q}_s{+}k_s)\, y\, \mathrm{sph}_{\tau_s}(y)$ for $s < |\mathfrak{C}|$ be a list of all the counting-sentences that occur in ψ.

Claim 8.4. *There is a $\mathcal{C} \in \mathfrak{C}$ such that $(\mathcal{C}, c) \ncong \tau_i$, for all $c \in C$ and $i \in [s]$.*

Proof. Assume for contradiction that for each $\mathcal{C} \in \mathfrak{C}$ there is a $c_{\mathcal{C}} \in C$ and an $i_{\mathcal{C}} \in [s]$ such that $(\mathcal{C}, c_{\mathcal{C}}) \cong \tau_{i_{\mathcal{C}}}$. Since $s < |\mathfrak{C}|$, there exist two distinct $\mathcal{C}$ and $\mathcal{C}'$

in $\mathfrak{C}$ such that $i_{\mathcal{C}} = i_{\mathcal{C}'}$. But then, $(\mathcal{C}, c_{\mathcal{C}}) \cong \tau_{i_{\mathcal{C}}} \cong (\mathcal{C}', c_{\mathcal{C}'})$, which implies that $\mathcal{C}$ is isomorphic to $\mathcal{C}'$ and contradicts the assumption that the structures in $\mathfrak{C}$ are pairwise non-isomorphic. This completes the proof of Claim 8.4. $\square$

For the remainder of this proof let $\mathcal{C}$ be a fixed structure provided by Claim 8.4.

For each $i \in [s]$ let p_i be the period of the ultimately periodic set $(\mathsf{Q}_i{+}k_i)$, let n_i be an offset of $(\mathsf{Q}_i{+}k_i)$, and let r_i be the radius of the sphere-formula $\mathrm{sph}_{\tau_i}(y)$. Let $R := \max\{r_1, \ldots, r_s\}$, let $K := \max\{n_1, \ldots, n_s\}$, and let P be the least common multiple of the numbers $p_1, \ldots, p_s$.

Let $\mathcal{A}$ be the σ-structure obtained as the disjoint union of K copies of all *proper* induced substructures of $\mathcal{C}$ of the form $\mathcal{C}[N_r^{\mathcal{C}}(c)]$ with $c \in C$ and $r \leqslant R$ (such structures exist since the universe of $\mathcal{C}$ has at least two elements). Furthermore, let $\mathcal{B}$ be the σ-structure obtained as the disjoint union of $\mathcal{A}$ and of $2P$ copies of $\mathcal{C}$.

Claim 8.5. *$\mathcal{A}, \mathcal{B} \in \mathfrak{D}$, $\mathcal{A}$ contains 0 disjoint copies of each structure from $\mathfrak{C}$, and $\mathcal{B}$ contains $2P$ disjoint copies of $\mathcal{C}$ and 0 disjoint copies of each structure from $\mathfrak{C} \setminus \{\mathcal{C}\}$.*

Proof. The first statement is true since $\mathcal{C} \in \mathfrak{C} \subseteq \mathfrak{D}$ and $\mathfrak{D}$ is closed under disjoint unions and induced substructures. The second statement is true since $\mathfrak{C}$ is substructure-free and $\mathcal{A}$ only contains disjoint copies of proper induced substructures of the structure $\mathcal{C} \in \mathfrak{C}$. The third statement is true since $\mathcal{B}$ is the disjoint union of $\mathcal{A}$ and $2P$ copies of $\mathcal{C}$. This completes the proof of Claim 8.5 $\square$

Recall that by our assumption, a structure from $\mathfrak{D}$ satisfies the sentence ψ if and only if it contains at most one disjoint copy of each structure from $\mathfrak{C}$. Thus, by Claim 8.5 we have

$$\mathcal{A} \models \psi \qquad \text{and} \qquad \mathcal{B} \not\models \psi\,. \tag{21}$$

Recall that ψ is a Boolean combination of the counting-sentences $(\mathsf{Q}_i{+}k_i)\, y\, \mathrm{sph}_{\tau_i}(y)$ for $i \in [s]$. Consider an arbitrary $\tau_i \in \{\tau_1, \ldots, \tau_s\}$. By the choice of the structures $\mathcal{A}$ and $\mathcal{B}$ we have

$$\begin{aligned} n^{\mathcal{B}}_{\tau_i} \;&:=\; |\{\, b \in B \,:\, \mathcal{N}^{\mathcal{B}}_r(b) \cong \tau_i \,\}| \\ &=\; |\{\, a \in A \,:\, \mathcal{N}^{\mathcal{A}}_r(a) \cong \tau_i \,\}| \;+\; 2P \cdot |\{c \in C \,:\, \mathcal{N}^{\mathcal{C}}_r(c) \cong \tau_i\}| \\ &=\; n^{\mathcal{A}}_{\tau_i} \;+\; 2P \cdot n^{\mathcal{C}}_{\tau_i} \end{aligned}$$

for $n^{\mathcal{A}}_{\tau_i} := |\{\, a \in A \,:\, \mathcal{N}^{\mathcal{A}}_r(a) \cong \tau_i \,\}|$ and $n^{\mathcal{C}}_{\tau_i} := |\{\, c \in C \,:\, \mathcal{N}^{\mathcal{C}}_r(c) \cong \tau_i \,\}|$. Suppose $n^{\mathcal{A}}_{\tau_i} \neq n^{\mathcal{B}}_{\tau_i}$. Then $n^{\mathcal{B}}_{\tau_i} > n^{\mathcal{A}}_{\tau_i}$ and there is some $c \in C$ with $\mathcal{N}^{\mathcal{C}}_r(c) \cong \tau_i$. From Claim 8.4 we know that $(\mathcal{C}, c) \not\cong \tau_i$, and hence $\mathcal{C}[N^{\mathcal{C}}_r(c)]$ is a *proper* induced substructure of $\mathcal{C}$. Thus, $\mathcal{A}$ contains K disjoint copies of $\mathcal{C}[N^{\mathcal{C}}_r(c)]$. Therefore, if $n^{\mathcal{C}}_{\tau_i} \geqslant 1$, then $n^{\mathcal{A}}_{\tau_i} \geqslant K$. In summary, we thus have

$$\text{either} \qquad n^{\mathcal{B}}_{\tau_i} = n^{\mathcal{A}}_{\tau_i} \qquad \text{or} \qquad n^{\mathcal{B}}_{\tau_i} > n^{\mathcal{A}}_{\tau_i} \geqslant K \quad \text{and} \quad n^{\mathcal{B}}_{\tau_i} \equiv n^{\mathcal{A}}_{\tau_i} \mod P. \tag{22}$$

Recall that $(\mathsf{Q}_i{+}k_i)$ has period p_i and P is a multiple of p_i. Furthermore, $(\mathsf{Q}_i{+}k_i)$ has offset $n_i \leqslant K$. Thus, from (22) we obtain that $n_{\tau_i}^{\mathcal{A}} \in (\mathsf{Q}_i{+}k_i) \iff n_{\tau_i}^{\mathcal{B}} \in (\mathsf{Q}_i{+}k_i)$. Hence,

$$\mathcal{A} \models (\mathsf{Q}_i{+}k_i)\, y\, \mathrm{sph}_{\tau_i}(y) \quad \iff \quad \mathcal{B} \models (\mathsf{Q}_i{+}k_i)\, y\, \mathrm{sph}_{\tau_i}(y)\,,$$

and this holds for all $i \in [s]$. Since ψ is a Boolean combination of the sentences $(\mathsf{Q}_i{+}k_i)\, y\, \mathrm{sph}_{\tau_i}(y)$ for $i \in [s]$, this implies that

$$\mathcal{A} \models \psi \quad \iff \quad \mathcal{B} \models \psi\,.$$

This contradicts (21) and completes the proof of Lemma 8.3. □

8.2. Proof of Theorem 8.2

For the proof of Theorem 8.2, we will use FO$[\sigma_d]$-formulas which identify the roots of trees from $\mathfrak{T}_{d,2^m}$ in σ_d-structures, and which recognise isomorphic copies of trees from $\mathfrak{T}_{d,2^m}$. These formulas are provided by the following lemma, whose proof is implicit in [9].

Lemma 8.6. *For each $d \in \mathbb{N}$ with $d \geqslant 2$ there is a number $c_d \in \mathbb{N}_{\geqslant 1}$ of size $\mathcal{O}(d^2)$, and for each $m \in \mathbb{N}_{\geqslant 1}$ there are* FO$[\sigma_d]$*-formulas* $\mathrm{tree}_{d,2^m}(x)$ *and* $\mathrm{iso}_{d,2^m}(x,x')$ *of size at most $c_d \cdot m$, such that the following holds for every σ_d-structure $\mathcal{A}$.*

(a) For every $a \in A$,

$$\mathcal{A} \models \mathrm{tree}_{d,2^m}[a] \quad \iff \quad \mathcal{A}[\mathrm{Reach}_a^{\mathcal{A}}] \textit{ is isomorphic to a tree in } \mathfrak{T}_{d,2^m}.$$

(b) For all $a, a' \in A$ with $\mathcal{A} \models \mathrm{tree}_{d,2^m}[a]$ and $\mathcal{A} \models \mathrm{tree}_{d,2^m}[a']$,

$$\mathcal{A} \models \mathrm{iso}_{d,2^m}[a,a'] \quad \iff \quad \mathcal{A}[\mathrm{Reach}_a^{\mathcal{A}}] \cong \mathcal{A}[\mathrm{Reach}_{a'}^{\mathcal{A}}]\,.$$

Proof. In the same way as in the proof of [9, Lemma 25] we define for each $\ell \in \mathbb{N}$ an FO$[\sigma_d]$-formula $\delta_{d,\leqslant\ell}(x,y,x',y')$ with the following property: If $\mathcal{A}$ is a σ_d-structure and $a,b,a',b' \in A$, then $\mathcal{A} \models \delta_{d,\leqslant\ell}[a,b,a',b']$ if, and only if, there are directed paths $(c_0,\dots,c_n)$ and $(c'_0,\dots,c'_n)$ in $\mathcal{A}$ of the same length $n \leqslant \ell$ and with $a = c_0$, $b = c_n$, $a' = c'_0$, $b' = c'_n$, such that for every $j \in [0,n)$ there is an $i \in [1,d)$ such that (c_j, c_{j+1}) and (c'_j, c'_{j+1}) belong to $S_i^{\mathcal{A}}$.

For $\ell = 0$ we let

$$\delta_{d,\leqslant 0}(x,y,x',y') \;:=\; x{=}y \;\wedge\; x'{=}y'\,,$$

and for $\ell = 1$ we let

$$\delta_{d,\leqslant 1}(x,y,x',y') \;:=\; \delta_{d,\leqslant 0}(x,y,x',y') \;\vee \bigvee_{j\in[1,d)} \big(S_j(x,y) \wedge S_j(x',y')\big)\,.$$

For all $\ell \geqslant 1$ we let

$$\delta_{d,\leqslant 2\ell}(x,y,x',y') \;:=\; \exists z \exists z' \forall u \forall v \forall u' \forall v' \Big(\big((u{=}x \wedge u'{=}x' \wedge v{=}z \wedge v'{=}z') \vee (u{=}z \wedge u'{=}z' \wedge v{=}y \wedge v'{=}y') \big) \rightarrow \delta_{d,\leqslant \ell}(u,v,u',v') \Big)$$

and

$$\delta_{d,\leqslant 2\ell+1}(x,y,x',y') \;:=\; \exists z \exists z' \big(\delta_{d,\leqslant 1}(x,z,x',z') \;\wedge\; \delta_{d,\leqslant 2\ell}(z,y,z',y') \big) .$$

Clearly, there is a number c'_d of size $\mathcal{O}(d)$ such that the formula $\delta_{d,\leqslant \ell}$ has size at most $c'_d \cdot \log \ell$.

Note that the formula $\pi_{d,\leqslant \ell}(x,y) := \delta_{d,\leqslant \ell}(x,y,x,y)$ expresses that there is a directed path of length at most ℓ from x to y. We let $\text{tree}_{d,2^m}(x)$ be the conjunction of the following formulas:

- $\forall y \; \big(\, \pi_{d,\leqslant 2^m+1}(x,y) \;\rightarrow\; \pi_{d,\leqslant 2^m}(x,y) \, \big)$
- $\forall y \, \forall y' \; \big(\, \big(\pi_{d,\leqslant 2^m}(x,y) \wedge \pi_{d,\leqslant 2^m}(x,y') \wedge \pi_{d,\leqslant 2^m}(y,y') \wedge \pi_{d,\leqslant 2^m}(y',y) \big) \;\rightarrow\; y'{=}y \, \big)$
- $\forall y \; \big(\, \pi_{d,\leqslant 2^m-1}(x,y) \;\rightarrow\; \bigwedge_{i\in[1,d)} \exists z \, \big(\, S_i(y,z) \;\wedge\; \forall z' \, (\, S_i(y,z') \rightarrow z'{=}z \,) \, \big) \, \big)$
- $\bigwedge_{i\in[1,d)} \forall y \, \forall z \; \Big(\, \big(\, S_i(y,z) \;\wedge\; \pi_{d,\leqslant 2^m}(x,y) \;\wedge\; \pi_{d,\leqslant 2^m}(x,z) \, \big) \;\rightarrow\; \neg \, \exists y' \; \big(\, \pi_{d,\leqslant 2^m}(x,y') \;\wedge\; \bigvee_{j\neq i} S_j(y',z) \, \big) \, \Big)$
- $\bigwedge_{i\in[1,d)} \forall y \, \forall z \; \Big(\, \big(\, S_i(y,z) \;\wedge\; \pi_{d,\leqslant 2^m}(x,y) \;\wedge\; \pi_{d,\leqslant 2^m}(x,z) \, \big) \;\rightarrow\; \neg \, \exists y' \; \big(\, \pi_{d,\leqslant 2^m}(x,y') \;\wedge\; S_i(y',z) \;\wedge\; \neg y'{=}y \, \big) \, \Big)$

Clearly, there is a number c_d of size $\mathcal{O}(d^2)$ such that the formula $\text{tree}_{d,2^m}(x)$ has size at most $c_d \cdot m$. Furthermore, it is not difficult to verify that the formula $\text{tree}_{d,2^m}(x)$ has the property stated in part (a) of Lemma 8.6. To prove part (b) of Lemma 8.6, we can choose

$$\text{iso}_{d,2^m}(x,x') \;:=\; \forall y \, \forall y' \; \big(\, \delta_{d,\leqslant 2^m}(x,y,x',y') \;\rightarrow\; \big(\, P(y) \leftrightarrow P(y') \, \big) \, \big) .$$

This completes the proof of Lemma 8.6. □

The proof of Theorem 8.2 is now obtained by combining Lemma 8.6 with Lemma 8.3.

Proof of Theorem 8.2. Let $d \in \mathbb{N}$ with $d \geqslant 2$. For $h \in \mathbb{N}$ we let $\mathfrak{F}_{d,h}$ be the closure of $\mathfrak{T}_{d,h}$ under disjoint unions and induced substructures.

For each $m \in \mathbb{N}_{\geqslant 1}$ we use the formulas $\mathrm{tree}_{d,2^m}(x)$ and $\mathrm{iso}_{d,2^m}(x,x')$ from Lemma 8.6 and let

$$\mathrm{root}_{d,2^m}(x) \ := \quad \mathrm{tree}_{d,2^m}(x) \ \wedge \ \neg \exists y \bigvee_{i \in [1,d)} S_i(y,x)$$

and

$$\varphi_{d,m} \ := \quad \forall x \forall x' \Big(\big(\mathrm{root}_{d,2^m}(x) \wedge \mathrm{root}_{d,2^m}(x') \wedge \neg\, x{=}x' \big) \ \rightarrow \ \neg\, \mathrm{iso}_{d,2^m}(x,x') \Big).$$

From Lemma 8.6 we obtain a number $c_d \geqslant 1$ of size $\mathcal{O}(d^2)$ such that for each $m \geqslant 1$ the formula $\varphi_{d,m}$ has size at most $c_d \cdot m$. Moreover, for every σ_d-structure $\mathcal{A}$ in $\mathfrak{F}_{d,2^m}$ we have

$\mathcal{A} \models \varphi_{d,m} \iff \mathcal{A}$ contains at most one disjoint copy of each structure in $\mathfrak{T}_{d,2^m}$.

From Lemma 8.3 we obtain that every hnf-sentence ψ in $\mathrm{FO}(\mathbf{U})[\sigma_d]$ that is equivalent to $\varphi_{d,m}$ on $\mathfrak{F}_{d,2^m}$ has size at least $|\mathfrak{T}_{d,2^m}|$. Since all structures in $\mathfrak{F}_{d,2^m}$ have degree at most d, the proof of Theorem 8.2 is complete. □

8.3. Transferring the statement of Theorem 8.2 to the signature $\sigma = \{E, P\}$

By a standard interpretation argument, we obtain the following corollary.

Corollary 8.7. *Let $\sigma = \{E, P\}$ be the signature consisting of a binary relation symbol E and a unary relation symbol P. Let $d \in \mathbb{N}$ with $d \geqslant 2$. There is a number $c_d \in \mathbb{N}_{\geqslant 1}$ of size $\mathcal{O}(d^4)$ and a sequence $(\varphi'_{d,m})_{m \geqslant 1}$ of $\mathrm{FO}[\sigma]$-sentences such that the following holds for every $m \geqslant 1$:*

(1) $\varphi'_{d,m}$ has size at most $c_d \cdot m$, and

(2) every hnf-sentence in $\mathrm{FO}(\mathbf{U})[\sigma]$ that is d-equivalent to $\varphi'_{d,m}$ has size at least $|\mathfrak{T}_{d,2^m}|$.

Proof. For $d = 2$, the statement is obtained from Theorem 8.2 by letting $E := S_1$. Let $d \geqslant 3$ and recall that $\sigma_d = \{P, S_1, \ldots, S_{d-1}\}$. With every σ_d-structure $\mathcal{A}$ we associate a σ-structure $\mathcal{B}$ as follows. Every element $a \in A$ is represented by the element $a \in B$, and we have $(a,a) \notin E^{\mathcal{B}}$. We let $P^{\mathcal{B}} := P^{\mathcal{A}}$. For every $i \in [1,d)$, every S_i-edge $e = (a,a') \in S_i^{\mathcal{A}}$ is represented by a directed path $(a, e_1, \ldots, e_i, a')$ in $\mathcal{B}$ where $e_1, \ldots, e_i$ are pairwise distinct new nodes satisfying the following: for every $j \in [1,i]$ we have $(e_j, e_j) \in E^{\mathcal{B}}$, and there exists exactly one element $b \in B$ with $b \neq e_j$ and $(b, e_j) \in E^{\mathcal{B}}$, and there exists exactly one element $b \in B$ with $b \neq e_j$ and $(e_j, b) \in E^{\mathcal{B}}$. We write $\widetilde{\mathcal{A}}$ to denote this σ-structure $\mathcal{B}$ that represents the σ_d-structure $\mathcal{A}$.

For every $\mathrm{FO}(\mathbf{U})[\sigma_d]$-formula φ, let $\widetilde{\varphi}$ be the $\mathrm{FO}(\mathbf{U})[\sigma]$-formula obtained from φ as follows:

1. relativise every quantifier to elements x that satisfy $\neg E(x,x)$

2. replace every atomic subformula of the form $S_i(x,y)$ by the FO[σ]-formula

$$\mathit{edge}_{S_i}(x,y) \;:=\; \exists z_1 \cdots \exists z_i \left(\begin{array}{l} \bigwedge\limits_{1\leqslant j\leqslant i} E(z_j,z_j) \;\;\wedge\;\; \bigwedge\limits_{1\leqslant j<j'\leqslant i} \neg\, z_j{=}z_{j'} \;\;\wedge \\ E(x,z_1) \;\wedge\; \bigwedge\limits_{1\leqslant j<i} E(z_j,z_{j+1}) \;\wedge\; E(z_i,y) \end{array} \right).$$

Clearly, $\widetilde{\varphi}$ is of size $\mathcal{O}(d^2)\cdot\|\varphi\|$, and if φ is a sentence, then for all σ_d-structures $\mathcal{A}$ we have $\mathcal{A}\models\varphi \iff \widetilde{\mathcal{A}}\models\widetilde{\varphi}$.

Now, revisiting the proof of Theorem 8.2, let $\mathfrak{T}'_{d,h} := \{\widetilde{\mathcal{A}} : \mathcal{A}\in\mathfrak{T}_{d,h}\}$, let $\mathfrak{F}'_{d,h}$ be the closure of $\mathfrak{T}'_{d,h}$ under disjoint unions and induced substructures, and note that every structure in $\mathfrak{F}'_{d,h}$ has degree at most d. Note that $\mathfrak{T}'_{d,h}$ is a finite substructure-free subset of $\mathfrak{F}'_{d,h}$ such that each structure in $\mathfrak{T}'_{d,h}$ is connected and has a universe of size at least 2, and the structures in $\mathfrak{T}'_{d,h}$ are pairwise non-isomorphic. Furthermore, $|\mathfrak{T}'_{d,h}| = |\mathfrak{T}_{d,h}|$. From Lemma 8.3 we obtain that if ψ'_h is an hnf-sentence in FO($\mathbf{U}$)[σ] such that for every structure $\mathcal{B}$ in $\mathfrak{F}'_{d,h}$ we have

$$\mathcal{B}\models\psi'_h \iff \mathcal{B} \text{ contains at most one disjoint copy of each structure from } \mathfrak{T}'_{d,h}$$

then $\|\psi'_h\| > |\mathfrak{T}_{d,h}|$.

Recall the FO[σ_d]-sentence $\varphi_{d,m}$ from the proof of Theorem 8.2 and consider the associated FO[σ]-sentence $\widetilde{\varphi}_{d,m}$. Then, $\widetilde{\varphi}_{d,m}$ has size at most $\tilde{c}_d\cdot m$ for a number $\tilde{c}_d \geqslant 1$ of size $\mathcal{O}(d^4)$. Moreover, for every σ_d-structure $\mathcal{A}$ in $\mathfrak{F}_{d,2^m}$ we have

$$\begin{array}{rl} & \widetilde{\mathcal{A}} \models \widetilde{\varphi}_{d,m} \\ \iff & \mathcal{A} \models \varphi_{d,m} \\ \iff & \mathcal{A} \text{ contains at most one disjoint copy of each structure in } \mathfrak{T}_{d,2^m} \\ \iff & \widetilde{\mathcal{A}} \text{ contains at most one disjoint copy of each structure in } \mathfrak{T}'_{d,2^m}. \end{array}$$

To obtain the analogous statement not only for structures in $\{\widetilde{\mathcal{A}} : \mathcal{A}\in\mathfrak{F}_{d,2^m}\}$, but for *all* structures $\mathcal{B}$ in $\mathfrak{F}'_{d,2^m}$, we slightly modify the formula $\widetilde{\varphi}_{d,m}$ by choosing

$$\mathrm{root}'_{d,2^m}(x) \;:=\; \widetilde{\mathrm{root}}_{d,2^m}(x) \;\wedge\; \neg\exists y\, E(y,x)$$

and

$$\varphi'_{d,m} \;:=\; \forall x\forall x'\Big(\big(\mathrm{root}'_{d,2^m}(x)\wedge \mathrm{root}'_{d,2^m}(x')\wedge\neg\, x{=}x'\big) \;\to\; \neg\,\widetilde{\mathrm{iso}}_{d,2^m}(x,x')\Big).$$

Clearly, $\varphi'_{d,m}$ has size at most $c_d\cdot m$ for a number $c_d\geqslant 1$ of size $\mathcal{O}(d^4)$. Furthermore, it is straightforward to verify that for all $\mathcal{B}$ in $\mathfrak{F}'_{d,2^m}$ we have

$$\begin{array}{rl} & \mathcal{B} \models \varphi'_{d,m} \\ \iff & \mathcal{B} \text{ contains at most one disjoint copy of each structure in } \mathfrak{T}'_{d,2^m}. \end{array}$$

From Lemma 8.3 we conclude that every hnf-sentence in FO($\mathbf{U}$)[σ] that is d-equivalent to $\varphi'_{d,m}$ has size at least $|\mathfrak{T}'_{d,2^m}| = |\mathfrak{T}_{d,2^m}|$. This completes the proof of Corollary 8.7. □

8.4. Proof of Theorem 8.1

Theorem 8.1 follows easily by combining Corollary 8.7 with the following lemma.

Lemma 8.8. *Let σ be a relational signature and let $d \in \mathbb{N}$ with $d \geqslant 2$. Suppose that there is a number $c_d \in \mathbb{N}_{\geqslant 1}$, a sequence $(\varphi_{d,m})_{m \geqslant 1}$ of* $\mathrm{FO}[\sigma]$*-sentences, and a strictly increasing function $f_d \colon \mathbb{N} \to \mathbb{N}_{\geqslant 1}$ such that for every $m \in \mathbb{N}$ the following holds:*

(1) $\varphi_{d,m}$ has size at most $c_d \cdot m$, and

(2) every hnf-sentence in $\mathrm{FO}(\mathbf{U})[\sigma]$ that is d-equivalent to $\varphi_{d,m}$ has size at least $f_d(m{-}1)$.

Then, there is no algorithm which upon input of an $\mathrm{FO}[\sigma]$-sentence φ computes in time $f_d(o(\|\varphi\|))$ an hnf-sentence in $\mathrm{FO}(\mathbf{U})[\sigma]$ that is d-equivalent to φ.

Proof. For contradiction, assume that there are a monotonically increasing function $g(n) \in o(n)$ and an algorithm which upon input of an $\mathrm{FO}[\sigma]$-sentence φ computes in time $f_d(g(\|\varphi\|))$ an hnf-sentence ψ in $\mathrm{FO}(\mathbf{U})[\sigma]$ that is d-equivalent to φ. Then, in particular, $\|\psi\| \leqslant f_d(g(\|\varphi\|))$.

For each $m \in \mathbb{N}_{\geqslant 1}$ let $\psi_{d,m}$ be the hnf-sentence computed by the algorithm upon input of $\varphi_{d,m}$. Then, $\psi_{d,m}$ has size

$$f_d(m{-}1) \quad \leqslant \quad \|\psi_{d,m}\| \quad \leqslant \quad f_d\big(g(\|\varphi_{d,m}\|)\big) \quad \leqslant \quad f_d\big(g(c_d \cdot m)\big)\,.$$

Hence, for each $m \in \mathbb{N}_{\geqslant 1}$ we have

$$m{-}1 \quad \leqslant \quad g(c_d \cdot m)\,. \tag{23}$$

But by assumption we have $g(n) \in o(n)$. Therefore, in particular for $\varepsilon := 1/(2c_d)$ there is an n_0 such that $g(n) < \varepsilon \cdot n$ for all $n \geqslant n_0$. This implies that for all $m \in \mathbb{N}_{\geqslant 1}$ with $m \geqslant \max\{2, (n_0/c_d)\}$ we have $g(c_d \cdot m) < \varepsilon \cdot c_d \cdot m = m/2 \leqslant m{-}1$. This is a contradiction to (23) and completes the proof of Lemma 8.8. ☐

We are now ready for the proof of Theorem 8.1.

Proof of Theorem 8.1.
First, consider the case that $d = 2$. Let $E := S_1$ such that $\sigma_2 = \sigma$. Recall from (19) that $|\mathfrak{T}_{2,2^m}| = 2^{2^m+1}$. In particular, $|\mathfrak{T}_{2,2^m}| \geqslant f_2(m{-}1)$ for the function $f_2 : \mathbb{N} \to \mathbb{N}_{\geqslant 1}$ with $f_2(n) := 2^{2^n}$. For the sequence $(\varphi_{2,m})_{m \geqslant 1}$ of $\mathrm{FO}[\sigma_2]$-sentences provided by Theorem 8.2, we know that there is a number $c_2 \in \mathbb{N}_{\geqslant 1}$ such that for each $m \in \mathbb{N}_{\geqslant 1}$, $\|\varphi_{2,m}\| \leqslant c_2 \cdot m$ and, moreover, every hnf-sentence in $\mathrm{FO}(\mathbf{U})[\sigma_2]$ that is 2-equivalent to $\varphi_{2,m}$ has size $\geqslant f_2(m{-}1)$. From Lemma 8.8 we obtain that there is no algorithm which upon input of an $\mathrm{FO}[\sigma]$-sentence φ computes in time $f_2(o(\|\varphi\|)) = 2^{2^{o(\|\varphi\|)}}$ an hnf-sentence in $\mathrm{FO}(\mathbf{U})[\sigma]$ that is 2-equivalent to φ. This proves the statement of Theorem 8.1 for $d = 2$.

Now, consider the case that $d \in \mathbb{N}$ with $d \geqslant 3$. Recall from (20) that $|\mathfrak{T}_{d,2^m}| \geqslant 2^{(d-1)^{2^m}}$. Note that for all $d \geqslant 3$ we have $(d-1)^2 \geqslant d$. Hence, for all $h \geqslant 1$ we have $(d-1)^{2h} \geqslant d^h$, and therefore,

$$|\mathfrak{T}_{d,2^m}| \quad \geqslant \quad 2^{(d-1)^{2^m}} \quad \geqslant \quad 2^{d^{2^{m-1}}}.$$

Consider the function $f_d \colon \mathbb{N} \to \mathbb{N}_{\geqslant 1}$ with $f_d(n) = 2^{d^{2^n}}$ for all $n \in \mathbb{N}$. For the sequence $(\varphi'_{d,m})_{m \geqslant 1}$ of FO[σ]-sentences provided by Corollary 8.7, we know that there is a number $c_d \in \mathbb{N}_{\geqslant 1}$ such that for each $m \in \mathbb{N}_{\geqslant 1}$, $\|\varphi'_{d,m}\| \leqslant c_d \cdot m$ and, moreover, every hnf-sentence in FO(**U**)[σ] that is d-equivalent to $\varphi'_{d,m}$ has size $\geqslant |\mathfrak{T}_{d,2^m}| \geqslant f_d(m-1)$. From Lemma 8.8 we obtain that there is no algorithm which upon input of an FO[σ]-sentence φ computes in time $f_d(o(\|\varphi\|)) = 2^{d^{2^{(o(\|\varphi\|))}}}$ an hnf-sentence in FO(**U**)[σ] that is d-equivalent to φ. This completes the proof of Theorem 8.1. □

9. Conclusion

We have generalised the notion of Hanf normal forms from first-order logic FO to first-order logic with unary counting quantifiers FO(**Q**). Our first main result (Theorem 3.3) precisely characterises those sets **Q** of unary counting quantifiers that permit Hanf normal forms: the logic FO(**Q**) permits Hanf normal forms if, and only if, all sets in **Q** are ultimately periodic.

Our second main result (Theorem 3.4) provides an algorithm which, for any set **Q** of ultimately periodic sets and any degree bound $d \in \mathbb{N}$, transforms an input FO(**Q**)-formula φ into an FO(**Q**)-formula in Hanf normal form that is equivalent to φ on all (finite) structures of degree at most d. We provided a runtime analysis which showed that for $d \geqslant 3$ this algorithm uses time at most 3-fold exponential in the size of φ, and for $d = 2$ it uses time at most 2-fold exponential in the size of φ. We generalised a lower bound of [2] to show that for all $d \geqslant 3$ and plain first-order logic FO, our 3-fold exponential upper bound is worst-case optimal. Furthermore, we adapted this lower bound to show that also for $d = 2$, our 2-fold exponential upper bound is worst-case optimal (Theorem 8.1).

As an easy application of our algorithm, we obtained that for ultimately periodic sets **Q**, model-checking of FO(**Q**)-sentences against structures of degree at most d can be done in time

$$2^{2^{\mathrm{poly}(k)}} \cdot n \quad \text{for } d = 2, \qquad \text{and} \qquad 2^{d^{2^{\mathcal{O}(k)}}} \cdot n \quad \text{for } d \geqslant 3\,,$$

where k is the size of the formula and n is the size of the structure (see Theorem 3.5). For both cases, lower bounds of [9] show that our algorithm is worst-case optimal already for plain first-order logic.

In [1], our algorithm for transforming FO(**Q**)-formulas with ultimately periodic counting quantifiers **Q** into Hanf normal form was used to obtain efficient algorithms for evaluating FO(**Q**)-queries under updates on bounded degree databases. The recent article [16] presented a relaxed variant of a Hanf normal form that exists for all FO($\mathbf{Q}$)-formulas for arbitrary sets $\mathbf{Q}$ of unary counting quantifiers and, more generally, for an extension of first-order logic with counting terms and numerical predicates.

Regarding future work, it would be interesting to investigate to which extent our results can be extended to infinite structures.

References

[1] C. Berkholz, J. Keppeler, and N. Schweikardt. Answering FO+MOD queries under updates on bounded degree databases. In *Proceedings of the 20th International Conference on Database Theory, ICDT'17, March 21–24, 2017, Venice, Italy*, LIPIcs, pages 8:1–8:18. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017. Preprint available at `https://arxiv.org/abs/1702.08764`.

[2] B. Bollig and D. Kuske. An optimal construction of Hanf sentences. *J. Applied Logic*, 10(2):179–186, 2012.

[3] A. Dawar, M. Grohe, S. Kreutzer, and N. Schweikardt. Model Theory Makes Formulas Large. In *Proc. ICALP'07*, pages 1076–1088, 2007. Full version available as preprint NI07003-LAA, Isaac Newton Institute of Mathematical Sciences (2007).

[4] A. Durand and E. Grandjean. First-order queries on structures of bounded degree are computable with constant delay. *ACM Trans. Comput. Log.*, 8(4), 2007.

[5] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory.* Springer, 1999.

[6] R. Fagin, L. Stockmeyer, and M. Vardi. On monadic NP vs. monadic co-NP. *Inf. and Comp.*, 120(1):78–92, 1995.

[7] J. Flum and M. Grohe. *Parameterized Complexity Theory.* Springer, 2006.

[8] M. Frick and M. Grohe. Deciding first-order properties of locally tree-decomposable structures. *J. ACM*, 48(6):1184–1206, 2001.

[9] M. Frick and M. Grohe. The complexity of first-order and monadic second-order logic revisited. *Ann. Pure Appl. Logic*, 130(1-3):3–31, 2004.

[10] H. Gaifman. On local and non-local properties. In J. Stern, editor, *Proceedings of the Herbrand Symposium, Logic Colloquium '81*, pages 105–135. North Holland, 1982.

[11] W. Hanf. Model-theoretic methods in the study of elementary logic. In J. Addison, L. Henkin, and A. Tarski, editors, *The Theory of Models*, pages 132–145. North Holland, 1965.

[12] 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 '16, New York, NY, USA, July 5-8, 2016*, pages 277–286, 2016.

[13] 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.

[14] W. Kazana and L. Segoufin. First-order query evaluation on structures of bounded degree. *Logical Methods in Computer Science*, 7(2), 2011.

[15] S. Kreutzer. Algorithmic meta-theorems. In *Finite and Algorithmic Model Theory*. Cambridge University Press, 2011. London Mathematical Society Lecture Notes, No. 379.

[16] D. Kuske and N. Schweikardt. First-order logic with counting. In *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017*, pages 1–12, 2017. Preprint available at http://arxiv.org/abs/1703.01122.

[17] L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.

[18] S. Lindell. A normal form for first-order logic over doubly-linked data structures. *Int. J. Found. Comp. Sci.*, 19(1):205–217, 2008.

[19] M. Lothaire. *Combinatorics on words*. Cambridge Univ. Press, 1984.

[20] A. B. Matos. Periodic sets of integers. *Theor. Comput. Sci.*, 127(2):287–312, 1994.

[21] F. Neven, N. Schweikardt, F. Servais, and T. Tan. Distributed streaming with finite memory. In *Proc. 18th International Conference on Database Theory (ICDT 2015)*, pages 324–341, 2015.

[22] J. Nurmonen. Counting modulo quantifiers on finite structures. *Inf. Comput.*, 160(1-2):62–87, 2000.

[23] T. Schwentick and K. Bartelmann. Local normal forms for first-order logic with applications to games and automata. *Discrete Mathematics and Computer Science*, 3:109–124, 1999.

[24] D. Seese. Linear time computable problems and first-order descriptions. *Math. Struc. in Comp. Sci.*, 6(6):505–526, 1996.

[25] L. Segoufin. A glimpse on constant delay enumeration (invited talk). In *Proc. 31st International Symposium on Theoretical Aspects of Computer Science (STACS 2014)*, pages 13–27, 2014.