Notions of locality and their logical characterizations over finite models*

Lauri Hella
Department of Mathematics
P.O. Box 4 (Yliopistonkatu 5)
00014 University of Helsinki, Finland
Email: lauri.hella@helsinki.fi

Leonid Libkin
Bell Laboratories
600 Mountain Avenue
Murray Hill, NJ 07974, USA
Email: libkin@bell-labs.com

Juha Nurmonen†
Department of Mathematics
and Computer Science
University of Leicester
University Road
Leicester LE1 7RH, UK
Email: j.nurmonen@mcs.le.ac.uk

Abstract

Many known tools for proving expressibility bounds for first-order logic are based on one of several locality properties. In this paper we characterize the relationship between those notions of locality. We note that Gaifman's locality theorem gives rise to two notions: one deals with sentences and one with open formulae. We prove that the former implies Hanf's notion of locality, which in turn implies Gaifman's locality for open formulae. Each of these implies the bounded degree property, which is one of the easiest tools for proving expressibility bounds. These results apply beyond the first-order case. We use them to derive expressibility bounds for first-order logic with unary quantifiers and counting. We also characterize the notions of locality on structures of small degree.

1 Introduction

It is well known that first-order logic has limited expressive power. Typically, inexpressibility proofs are based on either a compactness argument, or Ehrenfeucht-Fraïssé games. In recent years, the expressive power of logics over *finite* models has been studied extensively. This increased interest is mostly due to a number of applications in computer science. For example, most database query languages have well known logical counterparts: traditional relational calculus has precisely the power of first-order logic, the language Datalog, with added negation and evaluated inflationary, corresponds to the least-fixpoint logic, and the query language with while loops is equivalent to the partial-fixpoint

*Some of these results appeared in preliminary form in the Proceedings of the 12th IEEE Symposium on Logic in Computer Science, in a paper by the second author.

†Supported by EPSRC grant GR/K 96564.

logic, cf. [1]. Another area of application is descriptive complexity. It turns out that familiar logics capture complexity classes over classes of (ordered) finite structures, cf. [8, 18].

Since compactness fails in restriction to finite structures [15], to prove results about the limits of expressiveness of first-order logic, one has to use Ehrenfeucht-Fraïssé games. Moreover, Ehrenfeucht-Fraïssé games are often used as the basic step in other, more sophisticated games for different logics. For example, playing the Ehrenfeucht-Fraïssé game is one of the steps in the Ajtai-Fagin game for monadic Σ^1_1 [2]. Since playing the game often involves an intricate combinatorial argument, it was suggested by Fagin, Stockmeyer and Vardi in [11] to build a library of winning strategies for those games. Or, more generally, one would like to have a collection of versatile and easily applicable tools for proving expressibility bounds for first-order logic.

A number of results proving expressibility bounds explain the nature of the limitations of first-order logic by saying that it can only express *local* properties. Intuitively, one cannot grasp the whole structure; instead, to answer a first-order query, one only looks at small portions of the input.

Several proposals have been made to formalize the notion of locality. Gaifman [12] proved that the outcome of a first-order definable query depends only on the isomorphism types of neighborhoods of a fixed radius. Fagin, Stockmeyer and Vardi [11], modifying a result by Hanf [16] for the finite case, proved that if a certain criterion relating the numbers of small neighborhoods in two structures holds, then these structures agree on sentences whose quantifier rank is determined by the size of those neighborhoods. Libkin and Wong [22] showed that if a first-order query operates on graphs, then the number of different in- and out-degrees in the output is below a bound given by the query and the maximal degree in the input graph. This property, called the *bounded degree property*, was generalized to first-order queries on arbitrary finite structures by Dong, Libkin and Wong [6].

Typically, inexpressibility proofs based on the bounded degree property are very easy (see, e.g., [22]). Proofs based on Hanf's theorem, while often easier than playing a game directly (compare, for example, the proofs that connectivity is not monadic Σ^1_1 in [2] and [11]) may still involve somewhat nontrivial combinatorial argument (see, e.g., [5]). On the other hand, Hanf's theorem being close to game characterization of logics, its extensions have been proved for several extensions of first-order logics [9, 24, 26]. Thus, it would be desirable to understand the relationship between various locality notions for first-order logic and its extensions.

This constitutes the main goal of the paper. We isolate the locality notions underlying Gaifman's and Hanf's theorems, and prove a chain of implications among them. In particular, we show that the bounded degree property and an analog of Gaifman's theorem hold in several counting extensions of first-order logic.

Organization and summary In Section 2, we introduce the notation and describe the basic notions of locality. We start by reviewing Gaifman's theorem, and note that it leads to two properties, called the *Gaifman-locality* and the *strong Gaifman-locality*. The result of [12] then says that first-order logic has both of these properties. We review the modification of Hanf's technique [16] for the finite case [11], and define the notion of *Hanf-locality*. We review the bounded degree property of [6, 22] which is implied by the Gaifman-locality [6].

In Section 3 we review the extensions of first-order logic we consider in this paper. These are fragments of infinitary logic, logics with unary quantifiers and first-order logic with second-sort counting. We then establish that all these logics have the Hanf-locality property. In Section 4 we give the main technical machinery which is used in the rest of the paper. Mainly, we examine Hanf's technique more closely and among other things give a new simple proof that the extension of first-order logic by all unary generalized quantifiers has the Hanf-locality property.

In Section 5 relationships between the notions of locality are considered. In Section 5.1, we show that the Hanf-locality implies the Gaifman-locality. We use this to derive a number of expressibility bounds for various logics; we also touch on some applications in descriptive complexity. This implies the bounded degree property for any logic that possesses the Hanf-locality property. In Section 5.2, we show that the strong Gaifman-locality implies the Hanf-locality. We do not yet know of any extension of first-order that is strongly Gaifman-local, so the main implication of this result is a very simple and intuitive proof that first-order logic is Hanf-local.

In Section 6, we give complete characterizations of the three main notions of locality on structures of small degree. We show that, in order to check whether a query has the bounded degree property, it is enough to check whether it is definable in a certain logic on structures of bounded degree. These results may also be helpful in proving expressibility bounds over finite models, as many counterexamples that are constructed in finite model theory turn out to be structures of small degree.

2 Notions of locality

Notation

Unless explicitly stated otherwise, all structures are assumed to be *finite*.

A relational signature σ is a set of relation symbols $\{R_1, ..., R_l\}$, with an associated arity function. In what follows, $p_i(> 0)$ denotes the arity of R_i. We write σ_n for σ extended with n new constant symbols. The signature of graphs (that is, one binary predicate R) is denoted by σ_{gr}.

A σ-structure is $\mathcal{A} = \langle A, R_1^{\mathcal{A}}, \ldots, R_l^{\mathcal{A}} \rangle$, where A is a finite set, and $R_i^{\mathcal{A}} \subseteq A^{p_i}$ interprets R_i. The class of finite σ-structures is denoted by $\text{STRUCT}[\sigma]$. When there is no confusion, we may write R_i in place of $R_i^{\mathcal{A}}$. Isomorphism of structures is denoted by $\cong$. We shall adopt the convention that the carrier of a structure $\mathcal{A}$ is always denoted by A and the carrier of $\mathcal{B}$ is denoted by B.

To make our results applicable to a number of logics, we state below the condition that is necessary for the proofs. Let $\mathcal{L}$ be a logic. Assume a vocabulary σ, and let $U_1, \ldots, U_m$ be relational symbols not in σ. Let $\sigma' = \sigma \cup \{U_1, \ldots, U_m\}$. Then, for every σ formula $\varphi(\vec{x})$ in $\mathcal{L}$, we can form a σ' sentence $\Phi = \forall \vec{x}(\gamma(\vec{x}) \to \varphi(\vec{x}))$ in $\mathcal{L}$, where γ is a Boolean combination of atomic formulas in $\{U_1, \ldots, U_m\}$ using variables from $\vec{x}$. That is, a σ'-structure $\mathcal{A}$ satisfies Φ iff for every $\vec{a}$ from A such that $\gamma(\vec{a})$ holds, it is the case that $\mathcal{A}_\sigma \models \varphi(\vec{a})$, where $\mathcal{A}_\sigma$ is the σ-reduct of $\mathcal{A}$. This condition can be formulated along the lines of [7, 20] for abstract logics. However, as all the logics we consider here are extensions of first-order that trivially satisfy this condition, we will not go into more detail. In what follows, whenever we speak of a logic closed under first-order operations, we mean that the condition above is satisfied.

With each formula $\psi(x_1, \ldots, x_m)$ in a logical language whose symbols are in σ, we associate a *query* that maps a σ-structure $\mathcal{A}$ to an m-ary relation $\psi^{\mathcal{A}} = \{(a_1, \ldots, a_m) \in A^m | \ \mathcal{A} \models \psi(a_1, \ldots, a_m)\}$; we denote the corresponding structure with universe A by $\psi[\mathcal{A}] = \langle A, \psi^{\mathcal{A}} \rangle$.

Given a structure $\mathcal{A}$, its *Gaifman graph* [11, 12, 8] $\mathcal{G}(\mathcal{A})$ is defined as $\langle A, E \rangle$ where (a, b) is in E iff there is a tuple $\vec{t} \in R_i^{\mathcal{A}}$ for some i such that both a and b are in $\vec{t}$. The distance $d(a, b)$ is defined as the length of the shortest path from a to b in $\mathcal{G}(\mathcal{A})$; we assume $d(a,a) = 0$. Given $a \in A$, its *r-sphere* $S_r^{\mathcal{A}}(a)$ is $\{b \in A \mid d(a,b) \leq r\}$. For a tuple $\vec{t}$, define $S_r^{\mathcal{A}}(\vec{t})$ as $\bigcup_{a \in \vec{t}} S_r^{\mathcal{A}}(a)$.

Given a tuple $\vec{t} = (t_1, \ldots, t_n)$, its *$r$-neighborhood* $N_r^{\mathcal{A}}(\vec{t})$ is defined as a σ_n structure

$$\mathcal{A} \restriction S_r^{\mathcal{A}}(\vec{t}) = \langle S_r^{\mathcal{A}}(\vec{t}), R_1^{\mathcal{A}} \cap S_r^{\mathcal{A}}(\vec{t})^{p_1}, \ldots, R_k^{\mathcal{A}} \cap S_r^{\mathcal{A}}(\vec{t})^{p_k}, t_1, \ldots, t_n \rangle$$

That is, the carrier of $N_r^{\mathcal{A}}(\vec{t})$ is $S_r^{\mathcal{A}}(\vec{t})$, the interpretation of the σ-relations is obtained by restricting them from $\mathcal{A}$ to the carrier, and the n extra constants are the elements of $\vec{t}$. If the structure $\mathcal{A}$ is clear from the context, we shall write $S_r(\vec{t})$ and $N_r(\vec{t})$.

The quantifier rank of a first-order formula ψ, $\mathsf{qr}(\psi)$, is defined as the maximum depth of quantifier nesting in ψ; that is, $\mathsf{qr}(\psi) = 0$ for atomic formulas ψ, $\mathsf{qr}(\psi) = \max\{\mathsf{qr}(\varphi), \mathsf{qr}(\eta)\}$ if ψ is $\varphi \vee \eta$, $\mathsf{qr}(\psi) = \mathsf{qr}(\varphi)$ if ψ is $\neg\varphi$, and $\mathsf{qr}(\psi) = \mathsf{qr}(\varphi) + 1$ if ψ is of the form $\exists x \varphi$ or $\forall x \varphi$.

Gaifman-locality

Before presenting Gaifman's theorem, note that for any σ-structure $\mathcal{A}$, there is a first order formula $\gamma_\sigma(x,y)$ such that $\mathcal{A} \models \gamma_\sigma(a,b)$ iff $(a,b) \in \mathcal{G}(\mathcal{A})$. Thus, for every fixed k, there are first order formulae $d_{<k}(x,y)$, $d_k(x,y)$ and $d_{>k}(x,y)$ such that $\mathcal{A} \models d_{<k}(a,b)$ iff $d(a,b) < k$, and similarly for d_k and $d_{>k}$. This means that bounded quantification of the form $\forall x \in S_k(\vec{y})$ and $\exists x \in S_k(\vec{y})$ is expressible for every constant k. If every quantifier in a formula is of this form, where $\vec{y}$ are among its free variables, and $k \leq r$, we call the formula *r-local*.

Theorem 2.1 (Gaifman [12]) *For every first-order formula $\psi(x_1, \ldots, x_n)$ there exist t and r such that ψ is equivalent to a Boolean combination of t-local formulae $\chi(x_{i_1}, \ldots, x_{i_s})$ and sentences of the form*

$$\exists y_1 \ldots \exists y_m \Big(\bigwedge_{i=1}^{m} \varphi(y_i) \wedge \bigwedge_{i,j \leq m, i \neq j} d_{>2r}(y_i, y_j)\Big) \tag{1}$$

where φ is r-local. Furthermore, we can choose $r \leq 7^{\mathsf{qr}(\psi)-1}$, $t \leq (7^{\mathsf{qr}(\psi)-1} - 1)/2$, $m \leq n + \mathsf{qr}(\psi)$, and, if ψ is a sentence, only sentences (1) occur in the Boolean combination. □

Note that this theorem holds both on infinite and finite structures. To abstract the notion of being local and extend it to other logics, we introduce the following definitions. For σ-structures $\mathcal{A}$ and $\mathcal{B}$, and two tuples $\vec{a}$ from $\mathcal{A}$ and $\vec{b}$ from $\mathcal{B}$ of the same length, we write $\vec{a} \approx_r^{\mathcal{A},\mathcal{B}} \vec{b}$ if $N_r^{\mathcal{A}}(\vec{a}) \cong N_r^{\mathcal{B}}(\vec{b})$. If $\vec{a}$ and $\vec{b}$ are both tuples of elements of $\mathcal{A}$, we abbreviate this as $\vec{a} \approx_r^{\mathcal{A}} \vec{b}$. Again, $\mathcal{A}$ and $\mathcal{B}$ are omitted if they are clear from the context.

Definition 2.2

- *A formula $\psi(x_1, \ldots, x_m)$, is* Gaifman-local *if there exists $r > 0$ such that, for every $\mathcal{A} \in \mathrm{STRUCT}[\sigma]$ and for every two m-tuples $\vec{a}$, $\vec{b}$ of elements of A, $\vec{a} \approx_r \vec{b}$ implies $\mathcal{A} \models \psi(\vec{a})$ iff $\mathcal{A} \models \psi(\vec{b})$. The minimum r for which this holds is called the* locality rank *of ψ, and is denoted by $\mathsf{lr}(\psi)$.*
- *A formula $\psi(x_1, \ldots, x_m)$, is* strongly Gaifman-local *if there exists $r > 0$ such that, for every $\mathcal{A}, \mathcal{B} \in \mathrm{STRUCT}[\sigma]$ and for every two m-tuples $\vec{a}$, $\vec{b}$ of elements of A and B respectively, $\vec{a} \approx_r \vec{b}$ implies $\mathcal{A} \models \psi(\vec{a})$ iff $\mathcal{B} \models \psi(\vec{b})$.*

- *A sentence* Ψ *is* strongly Gaifman-local *if it is equivalent to a Boolean combination of sentences of the form* $\exists\vec{y}\psi(\vec{y})$, *where* $\psi(\vec{y})$ *is a strongly Gaifman-local formula.*

Now we immediately see:

Proposition 2.3 *Every first-order formula is Gaifman-local, and every first-order sentence is strongly Gaifman-local. Moreover, for every* $\psi(\vec{x})$ *of quantifier rank* n, $\mathsf{lr}(\psi) \leq (7^n - 1)/2$.

Proof: Suppose $\psi(\vec{x})$ is a first-order formula. Then it is equivalent to a Boolean combination of formulae $\gamma_i(\vec{x})$ and sentences Γ_j, where each γ_i is r_i-local. Let $r = \max r_i$. Then $\mathsf{lr}(\psi) \leq r$. Indeed, take a structure $\mathcal{A}$ and let $\vec{a} \approx_r \vec{b}$. Since $N_r(\vec{a}) \cong N_r(\vec{b})$, we have $\mathcal{A} \models \gamma_i(\vec{a}) \leftrightarrow \gamma_i(\vec{b})$, which gives us $\mathcal{A} \models \psi(\vec{a}) \leftrightarrow \psi(\vec{b})$, since all Γ_js are sentences.

To prove strong Gaifman-locality, note that any formula of the form

$$\psi(\vec{y}) \;=\; Q_1 z_1 \in S_{r_1}(\vec{y}) \ldots Q_k z_k \in S_{r_m}(\vec{y}) \gamma(\vec{y}, \vec{z}), \tag{2}$$

where the Q_is are quantifiers and γ is quantifier-free, is strongly Gaifman-local, since $\max_{i=1}^{k} r_i$ witnesses strong locality. The formula $\bigwedge_{i=1,\ldots,m} \varphi(y_i) \wedge \bigwedge_{i,j \leq m, i \neq j} d_{>2r}(y_i, y_j)$, where φ is r-local, can be represented in the form (2) with $r_i \leq 2r+1$ for each $i = 1, \ldots, k$. This implies strong Gaifman-locality. $\Box$

Note that not every first-order formula is strongly Gaifman-local. Consider $\psi(x) \equiv (\forall y \neg R(y, x)) \wedge \exists z \forall y \neg R(z, y)$. Assume that it is strongly local, fix r as in the definition and consider two graphs: G_1 is a chain of length $r+1$, and G_2 is obtained from G_1 by adding a loop on the end-node of G_1. Let a_i be the start node of G_i. Then $a_1 \approx_r^{G_1, G_2} a_2$, but $G_1 \models \psi(a_1)$ and $G_2 \models \neg\psi(a_2)$.

Hanf-locality

Let τ be an isomorphism type of a structure in the language σ_1 (σ extended with one constant). A point a in a structure $\mathcal{A}$ d-realizes τ, written as $\tau_d(\mathcal{A}, a) = \tau$, if $N_d^{\mathcal{A}}(a)$ is of isomorphism type τ.

By $\#_d[\mathcal{A}, \tau]$ we denote the number of elements of A which d-realize τ, that is, the cardinality of $\{a \in A \mid \tau_d(\mathcal{A}, a) = \tau\}$.

We say that $\mathcal{A}, \mathcal{B} \in \mathrm{STRUCT}[\sigma]$ are *d-equivalent*, if for every isomorphism type τ of a σ_1-structure we have $\#_d[\mathcal{A}, \tau] = \#_d[\mathcal{B}, \tau]$. This is denoted by $\mathcal{A} \leftrightarrows_d \mathcal{B}$. If $d > d'$, then $\mathcal{A} \leftrightarrows_d \mathcal{B}$ implies $\mathcal{A} \leftrightarrows_{d'} \mathcal{B}$ [11]. Note that d-equivalence can also be defined by letting $\mathcal{A} \leftrightarrows_d \mathcal{B}$ iff there exists a bijection $f : A \to B$ such that $a \approx_d^{\mathcal{A},\mathcal{B}} f(a)$ for every $a \in A$.

It was shown by Hanf [16] that two (finite or infinite) models are elementary equivalent if their spheres of finite radius are finite and, for each d and each type τ, either $\#_d[\mathcal{A}, \tau] = \#_d[\mathcal{B}, \tau] < \omega$, or both $\#_d[\mathcal{A}, \tau]$ and $\#_d[\mathcal{B}, \tau]$ are infinite. This was recently modified for the finite case as follows.

Theorem 2.4 (Fagin-Stockmeyer-Vardi [11]) *Let* $n > 0$. *Then there exists an integer* $d > 0$ *such that whenever* $\mathcal{A} \leftrightarrows_d \mathcal{B}$, *then* $\mathcal{A}$ *and* $\mathcal{B}$ *agree on all first-order sentences* φ *with* $\mathsf{qr}(\varphi) \leq n$. $\Box$

It follows from the proof in [11] that d can be taken to be 3^{n-1}, see also [10]. This leads to the following definition.

Definition 2.5 *A sentence* Ψ *is* Hanf-local *if there exists a number* d *such that any two* d*-equivalent structures agree on* Ψ*. The minimum* d *for which this holds is called the Hanf locality rank of* Ψ*, and is denoted by* $\mathsf{hlr}(\Psi)$*.*

Thus, Fagin-Stockmeyer-Vardi's theorem says that every first order sentence Ψ is Hanf-local, and $\mathsf{hlr}(\Psi) \leq 3^{\mathsf{qr}(\Psi)-1}$. Several extensions of this theorem are known. We consider them in the next section.

Bounded degree property

We define the notions of *degrees* in the usual way. For a graph G, its *degree set* $deg\_set(G)$ is the set of all possible in- and out-degrees that are realized in G, and $deg\_count(G)$ is the cardinality of $deg\_set(G)$. These notions generalize to arbitrary σ-structures: Given a relation $R_i^{\mathcal{A}}$ in $\mathcal{A}$, $degree_j(R_i, a)$ is the number of tuples in $R_i^{\mathcal{A}}$ whose jth component is a. Then $deg\_set(\mathcal{A})$ is the set $\{degree_j(R_i, a) \mid R_i \in \sigma, a \in A, j \leq p_i\}$, and $deg\_count(\mathcal{A})$ is its cardinality. The class of σ-structures $\mathcal{A}$ with $deg\_set(\mathcal{A}) \subseteq \{0, 1, \ldots, k\}$ is denoted by $\mathrm{STRUCT}_k[\sigma]$.

Given a formula $\psi(x_1, \ldots, x_m)$ and a structure $\mathcal{A}$, one can apply these concepts to the output structure $\psi[\mathcal{A}] = \langle A, \psi^A \rangle$. The bounded degree property says that there is an upper bound on $deg\_count(\varphi[\mathcal{A}])$ that depends only on ψ and the maximal value in $deg\_set(\mathcal{A})$. More precisely,

Definition 2.6 (see [6]) *A formula* $\psi(x_1, \ldots, x_m)$ *has the* bounded degree property (BDP)*, if there is a function* $f : \mathbb{N} \to \mathbb{N}$ *such that* $deg\_count(\psi[\mathcal{A}]) \leq f(k)$ *for any* $\mathcal{A} \in \mathrm{STRUCT}_k[\sigma]$. □

The BDP was introduced and proved for first-order queries from graphs to graphs (that is, formulae $\psi(x, y)$ in the language σ_{gr}) in [22]. It was also shown there that the BDP proves many inexpressibility results effortlessly. For example, to prove that (deterministic) transitive closure [8, 18] is not first-order, consider the following $C_n \in \mathrm{STRUCT}_1[\sigma_{\mathrm{gr}}]$:

where n is the number of nodes. Since the degree-set of its (deterministic) transitive closure has n elements, it violates the BDP and thus is not first-order definable. Another example in [22] is testing for balanced binary trees (that is, all paths from the root to the leaves are of the same length; note that this involves both recursive computation and counting). Assume this test is definable, and assume G is an input graph. For every two nodes a, b in G, having two successors each, a_1, a_2 and b_1, b_2, we define a new graph $G_{a,b}$ by making b_1, b_2 the successors of a and a_1, a_2 the successors of b. If G were a balanced binary tree, then $G_{a,b}$ is a balanced binary tree iff a and b have the same distance to the root. Thus, we see that there is a first-order query that, when its input is a balanced binary tree $G \in \mathrm{STRUCT}_2[\sigma_{\mathrm{gr}}]$ of length n, returns the set of cliques of elements at the same distance from the root, that is, a graph with an $n+1$-element degree-set. This again violates the BDP.

Theorem 2.7 (Dong-Libkin-Wong [6]) *Every Gaifman-local formula has the bounded degree property.* □

Thus, from Gaifman's theorem, we obtain:

Corollary 2.8 *Every first-order formula has the bounded degree property.* □

We saw that simple forms of recursion (deterministic transitive closure) violate the BDP. So does the simplest form of second-order quantification: monadic Σ_1^1 is not local. The BDP was introduced in connection with studying expressive power of database languages with aggregation [14, 22], where it was asked if such languages have it. The positive answer given recently [6] also implies that first-order logic with Rescher and Härtig quantifiers (see below for a definition of these quantifiers) has the BDP, but it was not known (although conjectured) if any of these is Gaifman-local.

3 Extensions of first-order logic

In this section we introduce the extensions of first-order logic that are considered in this paper. These are extensions with unary quantifiers and counting, and a fragment of infinitary logic with unary quantifiers.

First of all, recall that infinitary logic $\mathcal{L}_{\infty\omega}$ is the extension of first-order logic where infinite disjunctions and conjunctions of formulas are also allowed. It is well known that any (isomorphism closed) class $\mathcal{C} \subseteq \mathrm{STRUCT}[\sigma]$ can be defined in $\mathcal{L}_{\infty\omega}$. Interest of this logic comes from its fragments which have weaker expressive power. One such fragment is $\mathcal{L}^k_{\infty\omega}$ where only k distinct variables, free or bound, are allowed. The *finite variable logic* $\mathcal{L}^\omega_{\infty\omega}$ is then the union of $\mathcal{L}^k_{\infty\omega}$ over all natural numbers k. For an extensive study of this logic, see e.g. [8].

Another fragment of $\mathcal{L}_{\infty\omega}$ we are interested in, is the one where the quantifier rank of each formula is allowed to be at most r. From now on, we use the notation $(\mathcal{L}_{\infty\omega})^r$ for this fragment and $(\mathcal{L}_{\infty\omega})^\omega$ for the union of the logics $(\mathcal{L}_{\infty\omega})^r$ over all natural numbers r. More generally, if $\mathcal{L}$ is any logic with a well-defined notion of quantifier rank, then the fragment of $\mathcal{L}$ consisting of all formulas with quantifier rank at most r is denoted by $\mathcal{L}^r$, $\mathcal{L}^r = \{\varphi \mid \varphi \in \mathcal{L}, \mathsf{qr}(\varphi) \leq r\}$, and their union over all natural numbers r is denoted by $\mathcal{L}^\omega$. (The parenthesis notation $(\mathcal{L}_{\infty\omega})^r$ is used to avoid confusion with the finite variable logic.)

Unary quantifiers

Suppose $\mathcal{L}$ is a logic. Let σ_k^{unary} be a signature of k unary symbols, and let $\mathcal{K}$ be a class of σ_k^{unary}-structures which is closed under isomorphisms. Then $\mathcal{L}(Q_\mathcal{K})$ extends the set of formulae of $\mathcal{L}$ with the following additional rule:

> if $\psi_1(x_1, \vec{y}_1), \ldots, \psi_k(x_k, \vec{y}_k)$ are formulae, then $Q_\mathcal{K} x_1 \ldots x_k(\psi_1(x_1, \vec{y}_1), \ldots, \psi_k(x_k, \vec{y}_k))$ is a formula.

Here $Q_\mathcal{K}$ binds x_i in the ith formula, for each $i = 1, \ldots, k$. A free occurrence of a variable y in $\psi_i(x_i, \vec{y}_i)$ remains free in this new formula unless $y = x_i$. The semantics of $Q_\mathcal{K}$ is defined as follows:

> $\mathcal{A} \models Q_\mathcal{K} x_1 \ldots x_k(\psi_1(x_1, \vec{a}_1), \ldots, \psi_k(x_k, \vec{a}_k))$ iff $(A, \psi_1[\mathcal{A}, \vec{a}_1], \ldots, \psi_k[\mathcal{A}, \vec{a}_k]) \in \mathcal{K}$, where $\psi_i[\mathcal{A}, \vec{a}_i] = \{a \in A \mid \mathcal{A} \models \psi_i(a, \vec{a}_i)\}$.

In this definition, $\vec{a}_i$ is a tuple of parameters that gives the interpretation for those free variables of $\psi_i(x_i, \vec{y}_i)$ which are not equal to x_i. The logic $\mathcal{L}(\mathbf{Q})$ for a set $\mathbf{Q}$ of unary generalized quantifiers is defined similarly with the corresponding rule for each quantifier $Q_\mathcal{K} \in \mathbf{Q}$. The quantifier rank $\mathsf{qr}(\varphi)$ of an $\mathcal{L}(\mathbf{Q})$ formula φ is defined as the quantifier rank for the logic $\mathcal{L}$ with the following additional rule for each $Q_\mathcal{K} \in \mathbf{Q}$:

$$\mathsf{qr}(Q_\mathcal{K} x_1, \ldots, x_k(\psi_1(x_1, \vec{y}_1), \ldots, \psi_k(x_k, \vec{y}_k))) = \max\{\mathsf{qr}(\psi_i(x_i, \vec{y}_i)) \mid i \leq k\} + 1.$$

Examples of unary quantifiers include the usual $\exists$ and $\forall$, as well as the Rescher (bigger cardinality) and the Härtig (equicardinality) quantifiers. The Rescher quantifier $Q_\mathcal{R}$ and the Härtig quantifier $Q_\mathcal{I}$ are classes of σ_2^{unary}-structures; the Rescher quantifier $Q_\mathcal{R} x_1 x_2(\psi_1(x_1, \vec{a}_1), \psi_2(x_2, \vec{a}_2))$ is true if and only if there are at most as many points a that satisfy $\psi_1(a, \vec{a}_1)$ as there are points b that satisfy $\psi_2(b, \vec{a}_2)$. The Härtig quantifier $Q_\mathcal{I} x_1 x_2(\psi_1(x_1, \vec{a}_1), \psi_2(x_2, \vec{a}_2))$ in turn is true if and only if there are equally many points a that satisfy $\psi_1(a, \vec{a}_1)$ as there are points b that satisfy $\psi_2(b, \vec{a}_2)$. We use the notation $\mathcal{L}(\mathbf{Q}_u)$ for $\mathcal{L}$ extended with *all* unary quantifiers.

Next we give a game characterization for queries definable in $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$, that is, the fragment of infinitary logic with unary quantifiers that consists of formulas of finite quantifier rank. For this, we recall the definition of *bijective Ehrenfeucht-Fraïssé game* [17]. There are two players in this game, called the spoiler and the duplicator. Furthermore, the number of rounds, say n, and two structures $\mathcal{A}, \mathcal{B} \in \mathrm{STRUCT}[\sigma]$ are given. In each round $i = 1, \ldots, n$, the duplicator selects a bijection $f_i : A \to B$, and the spoiler selects a point $a_i \in A$ (if $card(A) \neq card(B)$, then the spoiler wins). The duplicator wins the game, if after the last round the relation $\{(a_i, f_i(a_i)) \mid 1 \leq i \leq n\}$ is a partial isomorphism $\mathcal{A} \to \mathcal{B}$; otherwise the spoiler wins. From the results in [17] it follows that the bijective Ehrenfeucht-Fraïssé game characterizes the expressive power of $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$. To see this, we first recall the following result.

Theorem 3.1 ([17]) *Let $\mathcal{A}, \mathcal{B} \in \mathrm{STRUCT}[\sigma]$. Then $\mathcal{A}$ and $\mathcal{B}$ agree on all sentences of $\mathcal{FO}(\mathbf{Q}_u)$ of quantifier rank up to n if and only if the duplicator has a winning strategy in the n-round bijective Ehrenfeucht-Fraïssé game over $\mathcal{A}$ and $\mathcal{B}$.* □

The proof of this in [17] actually shows that if the duplicator has a winning strategy in the n-round bijective Ehrenfeucht-Fraïssé game, then the structures agree on all sentences of $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^n$. This yields the following characterization.

Proposition 3.2 *A class $\mathcal{C} \subseteq \mathrm{STRUCT}[\sigma]$ is definable in $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ iff there is n such that the spoiler has a winning strategy in the n-round bijective Ehrenfeucht-Fraïssé game over $\mathcal{A}$ and $\mathcal{B}$ whenever $\mathcal{A} \in \mathcal{C}$ and $\mathcal{B} \notin \mathcal{C}$.*

Proof: Suppose $\mathcal{C}$ is definable by a sentence φ of $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^n$, for some n. Then for every $\mathcal{A}, \mathcal{B} \in \mathrm{STRUCT}[\sigma]$ such that $\mathcal{A}$ and $\mathcal{B}$ satisfy the same sentences of $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)$ of quantifier rank up to n, we have $\mathcal{A} \in \mathcal{C}$ if and only if $\mathcal{B} \in \mathcal{C}$. Hence there cannot be structures $\mathcal{A}$ and $\mathcal{B}$ as in the claim.

Assume then that the spoiler has a winning strategy in the n-round bijective Ehrenfeucht-Fraïssé over all $\mathcal{A}$ and $\mathcal{B}$ where $\mathcal{A} \in \mathcal{C}$ and $\mathcal{B} \notin \mathcal{C}$. For every such pair $\mathcal{A}$ and $\mathcal{B}$, by Theorem 3.1 there is a sentence $\varphi_{\mathcal{A},\mathcal{B}}$ of $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^n$ such that $\mathcal{A} \models \varphi_{\mathcal{A},\mathcal{B}}$ but $\mathcal{B} \not\models \varphi_{\mathcal{A},\mathcal{B}}$. But now $\mathcal{C}$ is defined by the $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^n$-sentence

$$\bigvee_{\mathcal{A}\in\mathcal{C}} \left(\bigwedge_{\mathcal{B}\notin\mathcal{C}} \varphi_{\mathcal{A},\mathcal{B}} \right).$$

Note that the infinite conjunctions and disjunction above can be restricted to range over sets, since there are only countably many non-isomorphic structures in STRUCT[σ]. $\square$

Let us remark that $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ is strictly stronger in expressive power than $\mathcal{FO}(\mathbf{Q}_u)^\omega$. This follows because the second vectorization of Härtig quantifier cannot be defined in $\mathcal{FO}(\mathbf{Q}_u)^\omega$, as was shown by Luosto [23] (using Ramsey theory). On the other hand, each vectorization of a unary quantifier can be defined in $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ [20].

As mentioned in the previous section, several extensions of Theorem 2.4 are known. One such extension can be given for $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$. This is because in [24] it was shown that d-equivalence, for large enough d, guarantees a winning strategy for the duplicator in the n-round bijective Ehrenfeucht-Fraïssé game.

Theorem 3.3 (see [24, 26]) *Every $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ sentence Ψ is Hanf-local. Moreover,* $\mathsf{hlr}(\Psi) \leq 3^{\mathsf{qr}(\Psi)}$. $\square$

In the next section we give a new simple proof for this fact.

We also consider first-order logic with counting $\mathcal{FO} + COUNT$ [19]. We present it here following [9]. $\mathcal{FO} + COUNT$ is defined as a two sorted logic, with second sort being the sort of natural numbers. More precisely, in this approach a structure $\mathcal{A}$ is of the form

$$\mathcal{A} = \langle \{1, \ldots, n\}, \{v_1, \ldots, v_n\}, <, \mathrm{BIT}, 1, \max, R_1^{\mathcal{A}}, \ldots, R_l^{\mathcal{A}} \rangle.$$

Here the relations $R_i^{\mathcal{A}}$ are defined on the domain $\{v_1, \ldots, v_n\}$, while on the numerical domain $\{1, \ldots, n\}$ one has $1, \max, <$ and the BIT predicate available ($\mathrm{BIT}(i,j)$ iff the ith bit in the binary representation of j is one). It also has counting quantifiers $\exists ix\varphi(x)$, meaning that φ has at least i satisfiers; here i refers to the numerical domain and x to the domain $\{v_1, \ldots, v_n\}$. These quantifiers bind x but not i. Etessami noticed that the technique used in a proof of [24] (which is based on bijective Ehrenfeucht-Fraïssé games [17]) applies to $\mathcal{FO} + COUNT$:

Theorem 3.4 (see [9]) *Every $\mathcal{FO} + COUNT$ sentence is Hanf-local. Moreover,* $\mathsf{hlr}(\Psi) \leq 3^{\mathsf{qr}(\Psi)}$. $\square$

4 Technical machinery

In this section we give the technical machinery used repeatedly in the paper in examining the relationships of the notions of locality, and characterizations of these notions on structures of bounded degree.

We start with a lemma which is one of our main technical tools and we apply it several times in this section. The idea of the proof given below is similar to all the earlier applications of Hanf's technique mentioned before [11, 16, 24, 25], but we believe this proof is simpler.

First, we need two obvious facts stated previously in [6].

Claim 4.1 *Assume that $\mathcal{A} \in$ STRUCT[σ] and $h : N_r^{\mathcal{A}}(\vec{a}) \to N_r^{\mathcal{A}}(\vec{b})$ is an isomorphism. Let $d \leq r$. Then h restricted to $S_d^{\mathcal{A}}(\vec{a})$ is an isomorphism between $N_d^{\mathcal{A}}(\vec{a})$ and $N_d^{\mathcal{A}}(\vec{b})$.* $\square$

Claim 4.2 *Assume that* $\mathcal{A} \in \mathrm{STRUCT}[\sigma]$ *and* $h : N_r^{\mathcal{A}}(\vec{a}) \to N_r^{\mathcal{A}}(\vec{b})$ *is an isomorphism. Let* $d + l \le r$ *and* $\vec{x}$ *be a tuple from* $S_l(\vec{a})$*. Then* $h(S_d(\vec{x})) = S_d(h(\vec{x}))$*, and* $N_d(\vec{x})$ *and* $N_d(h(\vec{x}))$ *are isomorphic.* $\Box$

The next claim generalizes a result from [6].

Claim 4.3 *Let* $\mathcal{A}, \mathcal{B} \in \mathrm{STRUCT}[\sigma]$ *and let* $\vec{a}_1 \in A^n, \vec{b}_1 \in B^n$ *for* $n \ge 1$*, and* $\vec{a}_2 \in A^m, \vec{b}_2 \in B^m$ *for* $m \ge 1$*. Assume that* $\vec{a}_1 \approx_r \vec{b}_1$ *and* $\vec{a}_2 \approx_r \vec{b}_2$*. Let* $\vec{a}$ *be* $\vec{a}_1$ *followed by* $\vec{a}_2$ *and* $\vec{b}$ *be* $\vec{b}_1$ *followed by* $\vec{b}_2$*. Furthermore, assume that for any components* a^1, a^2 *of* $\vec{a}_1$ *and* $\vec{a}_2$ *respectively we have* $d(a^1, a^2) > 2r + 1$*, and similarly for any components* b^1, b^2 *of* $\vec{b}_1$ *and* $\vec{b}_2$ *respectively,* $d(b^1, b^2) > 2r + 1$*. Then* $\vec{a} \approx_r \vec{b}$*.*

Proof: Since the distance between any two components of $\vec{a}_1$ and $\vec{a}_2$ is at least $2r+1$, any tuple in any σ-relation in $N_r^{\mathcal{A}}(\vec{a})$ either has all its components in $S_r^{\mathcal{A}}(\vec{a}_1)$, or it has all its components in $S_r^{\mathcal{A}}(\vec{a}_2)$. Similarly, any tuple in a σ-relation in $N_r^{\mathcal{B}}(\vec{b})$ either has all its components in $S_r^{\mathcal{B}}(\vec{b}_1)$, or in $S_r^{\mathcal{B}}(\vec{b}_2)$. Thus, the isomorphism between $N_r^{\mathcal{A}}(\vec{a})$ and $N_r^{\mathcal{B}}(\vec{b})$ can be defined componentwise. $\Box$

Let $\vec{a}$ be an n-tuple. By $\vec{a}x$ we denote the $n+1$-tuple whose first n components are those of $\vec{a}$ and the last one is x.

We define d-equivalence for structures with parameters. Let $\mathcal{A}, \mathcal{B} \in \mathrm{STRUCT}[\sigma]$ and let $\vec{a}$ and $\vec{b}$ be two tuples of the same length of $\mathcal{A}$ and $\mathcal{B}$, respectively. We say that $(\mathcal{A}, \vec{a})$ and $(\mathcal{B}, \vec{b})$ are *d-equivalent* if

$$\begin{array}{rl} \mathit{card}(\{x \in A \mid N_d^{\mathcal{A}}(\vec{a}x) \text{ is of isomorphism type } \tau\}) & = \\ \mathit{card}(\{y \in B \mid N_d^{\mathcal{B}}(\vec{b}y) \text{ is of isomorphism type } \tau\}) & \end{array}$$

for every isomorphism type τ. That is, there is a bijection $f : A \to B$ such that $\vec{a}x \approx_d \vec{b}f(x)$ for every $x \in A$. This is denoted by $(\mathcal{A}, \vec{a}) \leftrightarrows_d (\mathcal{B}, \vec{b})$.

Lemma 4.4 *If* $\mathcal{A} \leftrightarrows_d \mathcal{B}$ *and* $\vec{a} \approx_{3d+1}^{\mathcal{A},\mathcal{B}} \vec{b}$*, then* $(\mathcal{A}, \vec{a}) \leftrightarrows_d (\mathcal{B}, \vec{b})$*.*

Proof: We need to define a bijection $f : A \to B$ such that $\vec{a}x \approx_d^{\mathcal{A},\mathcal{B}} \vec{b}f(x)$ for every $x \in A$. Since $\vec{a} \approx_{3d+1}^{\mathcal{A},\mathcal{B}} \vec{b}$, there is an isomorphism $h : N_{3d+1}^{\mathcal{A}}(\vec{a}) \to N_{3d+1}^{\mathcal{A}}(\vec{b})$. Then the restriction of h to $S_{2d+1}^{\mathcal{A}}(\vec{a})$ is an isomorphism between $N_{2d+1}^{\mathcal{A}}(\vec{a})$ and $N_{2d+1}^{\mathcal{A}}(\vec{b})$ (Claim 4.1), and thus

$$\mathit{card}(A - S_{2d+1}^{\mathcal{A}}(\vec{a})) \quad = \quad \mathit{card}(B - S_{2d+1}^{\mathcal{B}}(\vec{b})).$$

Now consider an arbitrary type τ of a d-neighborhood of a single point. Assume that $a \in S_{2d+1}^{\mathcal{A}}(\vec{a})$ realizes τ in $\mathcal{A}$. Since h is an isomorphism of $3d+1$-neighborhoods, we see that $S_d^{\mathcal{A}}(a) \subseteq S_{3d+1}^{\mathcal{A}}(\vec{a})$ and thus $h(a) \in S_{2d+1}^{\mathcal{B}}(\vec{b})$ realizes τ. Thus, the number of elements in $S_{2d+1}^{\mathcal{A}}(\vec{a})$ and $S_{2d+1}^{\mathcal{B}}(\vec{b})$ that realize τ is the same.

Since $\mathcal{A} \leftrightarrows_d \mathcal{B}$ (that is, $\#_d[\mathcal{A}, \tau] = \#_d[\mathcal{B}, \tau]$), the observation above implies that

$$\mathit{card}(\{a \in A - S_{2d+1}^{\mathcal{A}}(\vec{a}) \mid \tau_d(\mathcal{A}, a) = \tau\}) \quad = \quad \mathit{card}(\{b \in B - S_{2d+1}^{\mathcal{B}}(\vec{b}) \mid \tau_d(\mathcal{B}, b) = \tau\})$$

for any τ. Thus, we can find a bijection $g : A - S_{2d+1}^{\mathcal{A}}(\vec{a}) \to B - S_{2d+1}^{\mathcal{B}}(\vec{b})$ such that $a \approx_d g(a)$ for any $a \in A - S_{2d+1}^{\mathcal{A}}(\vec{a})$.

We now define f by

$$f(x) \quad = \quad \begin{cases} h(x) & \text{if } x \in S^{\mathcal{A}}_{2d+1}(\vec{a}) \\ g(x) & \text{if } x \notin S^{\mathcal{A}}_{2d+1}(\vec{a}) \end{cases}$$

It is clear that f is a bijection $A \to B$.

We now claim that $\vec{a}x \approx_d \vec{b}f(x)$ for every $x \in A$. If $x \in S^{\mathcal{A}}_{2d+1}(\vec{a})$, then $S^{\mathcal{A}}_d(x) \subseteq S^{\mathcal{A}}_{3d+1}(\vec{a})$, and $\vec{a}x \approx_d \vec{b}h(x)$ because h is an isomorphism. If $x \notin S^{\mathcal{A}}_{2d+1}(\vec{a})$, then $f(x) = g(x) \notin S^{\mathcal{B}}_{2d+1}(\vec{b})$, and $x \approx_d g(x)$. Hence, by Claim 4.3, $\vec{a}x \approx_d \vec{b}g(x)$. □

Lemma 4.4 says that d-equivalent structures are d-equivalent even with parameters if large enough neighborhoods of these parameters are isomorphic. A similar idea was used in [24] to show that d-equivalence, for large enough d, guarantees a win for duplicator in the r-round bijective Ehrenfeucht-Fraïssé game. Lemma 4.4 allows us to simplify the proof of Theorem 3.3 given in [24].

First we give an easy consequence of Lemma 4.4.

Corollary 4.5 *If $(\mathcal{A}, \vec{a}) \leftrightarrows_{3d+1} (\mathcal{B}, \vec{b})$ then there exists a bijection $f : A \to B$ such that $(\mathcal{A}, \vec{a}x) \leftrightarrows_d (\mathcal{B}, \vec{b}f(x))$ for every $x \in A$.*

Proof: If $(\mathcal{A}, \vec{a}) \leftrightarrows_{3d+1} (\mathcal{B}, \vec{b})$ then there exists a bijection $f : A \to B$ such that $\vec{a}x \approx_{3d+1} \vec{b}f(x)$ for every $x \in A$. On the other hand, $(\mathcal{A}, \vec{a}) \leftrightarrows_{3d+1} (\mathcal{B}, \vec{b})$ implies $\mathcal{A} \leftrightarrows_d \mathcal{B}$. Then by Lemma 4.4, $(\mathcal{A}, \vec{a}x) \leftrightarrows_d (\mathcal{B}, \vec{b}f(x))$ for every $x \in A$. □

This provides a winning strategy in the bijective Ehrenfeucht-Fraïssé game on $\mathcal{A}$ and $\mathcal{B}$, if their d-equivalence for large enough d can be guaranteed. That is:

Proposition 4.6 *Let $n \geq 1$ and $d = (3^{n-1} - 1)/2$. Assume that $\mathcal{A} \leftrightarrows_d \mathcal{B}$. Then the duplicator has a winning strategy in the n-round bijective Ehrenfeucht-Fraïssé game on $\mathcal{A}$ and $\mathcal{B}$.*

Proof: Let $d_0 = 0$, $d_1 = 3d_0 + 1, d_2 = 3d_1 + 1, \ldots, d_k = 3d_{k-1} + 1, \ldots$. Note that $d = d_{n-1}$. Suppose that $\mathcal{A} \leftrightarrows_d \mathcal{B}$. Assume that after a round $i < n$ the spoiler has chosen points $a_1, \ldots, a_i$ and the duplicator has chosen bijections $f_1, \ldots, f_i$ and the equivalence

$$(\mathcal{A}, (a_1, \ldots, a_i)) \leftrightarrows_{d_{n-i}} (\mathcal{B}, (f_1(a_1), \ldots, f_i(a_i)))$$

holds. By Corollary 4.5 the duplicator can choose a bijection f_{i+1} such that

$$(\mathcal{A}, (a_1, \ldots, a_{i+1})) \leftrightarrows_{d_{n-i-1}} (\mathcal{B}, (f_1(a_1), \ldots, f_{i+1}(a_{i+1})))$$

for all a_{i+1}. In particular, after the last round

$$(\mathcal{A}, (a_1, \ldots, a_n)) \leftrightarrows_0 (\mathcal{B}, (f_1(a_1), \ldots, f_n(a_n))),$$

which guarantees that $\{(a_i, f_i(a_i)) \mid 1 \leq i \leq n\}$ is a partial isomorphism. □

In particular, it follows that every $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ sentence is Hanf-local (Theorem 3.3).

4.1 (n, d)-equivalence

We extend the notion of d-equivalence from points to tuples. Let τ^n be an isomorphism type of a structure in the language σ_n (σ extended with n constants). An n-tuple $\vec{a}$ of a structure $\mathcal{A}$ d-realizes τ^n, written as $\tau_d(\mathcal{A}, \vec{a}) = \tau^n$, if $N_d(\vec{a})$ is of isomorphism type τ^n.

We denote the cardinality of $\{\vec{a} \in A^n \mid \tau_d(\mathcal{A}, \vec{a}) = \tau^n\}$ by $\#_d[\mathcal{A}, \tau^n]$, that is, the number of n-tuples of $\mathcal{A}$ which d-realize τ^n.

We say that structures $\mathcal{A}, \mathcal{B} \in \mathrm{STRUCT}[\sigma]$ are *(n, d)-equivalent,* $\mathcal{A} \leftrightarrows_{n,d} \mathcal{B}$, if for every isomorphism type τ^n we have $\#_d[\mathcal{A}, \tau^n] = \#_d[\mathcal{B}, \tau^n]$, i.e., there are equally many n-tuples in $\mathcal{A}$ and $\mathcal{B}$ whose d-neighborhoods realize τ^n. Obviously, $(1, d)$-equivalence is the same as d-equivalence.

We start by analyzing this notion of equivalence. First observe that $\mathcal{A} \leftrightarrows_{n,d} \mathcal{B}$ implies $\mathcal{A} \leftrightarrows_d \mathcal{B}$. Indeed, let $\vec{a} \in A^n$ be an n-tuple whose all components are $a \in A$. Since $\mathcal{A} \leftrightarrows_{n,d} \mathcal{B}$ there is $\vec{b} \in B^n$ such that $\vec{a} \approx_d \vec{b}$, and this isomorphism proves that all components of $\vec{b}$ are the same, say $b \in B$. Thus $a \approx_d b$, which shows that $\mathcal{A} \leftrightarrows_d \mathcal{B}$. Recall from a remark preceding Theorem 2.4 that this implies $\mathcal{A} \leftrightarrows_{d'} \mathcal{B}$ for every $d' \leq d$ [11].

Our main result in this section is that r-equivalence of $n+1$-tuples can be guaranteed by d-equivalence of n-tuples for sufficiently large d that depends on r only.

First, we give a simple criterion for (n, d)-equivalence.

Proposition 4.7 *$\mathcal{A} \leftrightarrows_{n,d} \mathcal{B}$ iff there is a bijection $\pi : A^n \to B^n$ such that for any $\vec{a} \in A^n$,*

$$\vec{a} \ \approx_d \ \pi(\vec{a}).$$

Proof: Let $\tau_1^n, \ldots, \tau_s^n$ be the collection of all isomorphism types of d-neighborhoods of n-tuples realized in $\mathcal{A}$ or $\mathcal{B}$. Let $A_i = \{\vec{a} \in A^n \mid \tau_d(\mathcal{A}, \vec{a}) = \tau_i^n\}$ and $B_i = \{\vec{b} \in B^n \mid \tau_d(\mathcal{B}, \vec{b}) = \tau_i^n\}$. Then $\{A_i\}_{i=1,\ldots,s}$ and $\{B_i\}_{i=1,\ldots,s}$ form partitions of A^n and B^n respectively. Assume $\mathcal{A} \leftrightarrows_{n,d} \mathcal{B}$. Then $card(A_i) = card(B_i)$ for every $i = 1, \ldots, s$, and the required π is defined as the union of bijective maps between A_i and B_i for all i. Conversely, if π satisfying $\vec{x} \approx_d \pi(\vec{x})$ exists, let τ^n be an isomorphism type and let $\vec{a}_1, \ldots, \vec{a}_k$ be the elements of A^n such that $\tau_d(\mathcal{A}, \vec{a}_i) = \tau^n$. Then $\tau_d(\mathcal{B}, \pi(\vec{a}_i)) = \tau^n$, and $\#_d[\mathcal{A}, \tau^n] \leq \#_d[\mathcal{B}, \tau^n]$. A symmetric argument shows the reverse inequality. $\square$

The proposition below provides the main technical tool for Section 5.

Proposition 4.8 *Let $n > 0$ and $d \geq 0$. Then $\mathcal{A} \leftrightarrows_{n,3d+1} \mathcal{B}$ implies $\mathcal{A} \leftrightarrows_{n+1,d} \mathcal{B}$.* $\square$

Proof: Suppose $\mathcal{A} \leftrightarrows_{n,3d+1} \mathcal{B}$. Then there exists a bijection $\mu : A^n \to B^n$ such that $\vec{a} \approx_{3d+1} \mu(\vec{a})$ for every $\vec{a} \in A^n$. As observed above, we also know that $\mathcal{A} \leftrightarrows_d \mathcal{B}$. Thus by Lemma 4.4, $(\mathcal{A}, \vec{a}) \leftrightarrows_d (\mathcal{B}, \mu(\vec{a}))$ for every $\vec{a} \in A^n$. By definition, for every $\vec{a} \in A^n$ there is a bijection $f_{\vec{a}} : A \to B$ such that $\vec{a}x \approx_d \mu(\vec{a})f_{\vec{a}}(x)$ for every $x \in A$. Now the bijection $\pi(\vec{a}x) = \mu(\vec{a})f_{\vec{a}}(x)$ proves the claim $\mathcal{A} \leftrightarrows_{n+1,d} \mathcal{B}$. $\square$

As an immediate consequence we state the following.

Corollary 4.9 *For any $r > 0$ and any $n \geq 1$ there exists a number d such that $\mathcal{A} \leftrightarrows_d \mathcal{B}$ implies $\mathcal{A} \leftrightarrows_{n,r} \mathcal{B}$. In fact, d can be taken to be $3^{n-1}r + (3^{n-2} - 1)/2$ for $n > 1$ and $d = r$ for $n = 1$.* $\square$

5 Relationships between the notions of locality

This section is dedicated to the study of the relationships between the notions of locality. We show that the Hanf-locality implies the Gaifman-locality, and the strong Gaifman-locality implies the Hanf-locality. We then see that each of these notions of locality implies the bounded degree property.

First, we extend Definition 2.5 from sentences to formulas.

Definition 5.1 *A formula $\psi(x_1, \ldots, x_n)$ is* Hanf-local *if there exists a number d such that for every $\mathcal{A}, \mathcal{B} \in \mathrm{STRUCT}[\sigma]$ and for every two n-tuples $\vec{a}$ and $\vec{b}$ of elements of $\mathcal{A}$ and $\mathcal{B}$ respectively, $(\mathcal{A}, \vec{a}) \leftrightarrows_d (\mathcal{B}, \vec{b})$ implies $\mathcal{A} \models \psi(\vec{a})$ iff $\mathcal{B} \models \psi(\vec{b})$. The minimum d for which this holds is called the Hanf locality rank of ψ, and is denoted by $\mathsf{hlr}(\psi)$.*

We start with a simple observation, which shows that in the study of Hanf-locality it is enough to consider just sentences.

Given a signature σ, by $\sigma^{(n)}$ we denote σ extended with n new unary symbols $U_1, \ldots, U_n$. Given a structure $\mathcal{A}$ and an n-tuple $\vec{a}$, by $\mathcal{A}[\vec{a}]$ we denote the $\sigma^{(n)}$ structure that extends $\mathcal{A}$ by interpreting the U_is as singletons containing the corresponding components of $\vec{a}$.

Let $\psi(x_1, \ldots, x_n)$ be a formula with n free variables. By $\Psi^{(n)}$ we denote a sentence in a logic $\mathcal{L}$ that is equivalent to $\forall x_1 \ldots \forall x_n \, ((U_1(x_1) \wedge \ldots \wedge U_n(x_n)) \to \psi(x_1, \ldots, x_n))$; it exists if $\mathcal{L}$ is closed under first-order operations. Obviously, for any $\mathcal{A}$ and any n-tuple $\vec{a}$, $\mathcal{A} \models \psi(\vec{a})$ if and only if $\mathcal{A}[\vec{a}] \models \Psi^{(n)}$.

Proposition 5.2 *Let $\mathcal{L}$ be a logic that is closed under first-order operations. If every sentence in $\mathcal{L}$ is Hanf-local, then every formula in $\mathcal{L}$ is Hanf-local.*

Proof. Let $\psi(x_1, \ldots, x_n)$ be a formula of $\mathcal{L}$, and let $\mathsf{hlr}(\Psi^{(n)}) = d$. Suppose $(\mathcal{A}, \vec{a}) \leftrightarrows_d (\mathcal{B}, \vec{b})$; then also $\mathcal{A}[\vec{a}] \leftrightarrows_d \mathcal{B}[\vec{b}]$. By the observation above we then have

$$\mathcal{A} \models \psi(\vec{a}) \ \text{ iff } \ \mathcal{A}[\vec{a}] \models \Psi^{(n)} \ \text{ iff } \ \mathcal{B}[\vec{b}] \models \Psi^{(n)} \ \text{ iff } \ \mathcal{B} \models \psi(\vec{b})\,.$$

Hence $\psi(\vec{x})$ is Hanf-local and $\mathsf{hlr}(\psi) \leq d$. □

5.1 Hanf-locality implies Gaifman-locality

The first main result of this section is:

Theorem 5.3 *Every Hanf-local formula is Gaifman-local.*

Proof: Let a formula $\psi(x_1, \ldots, x_n)$ be Hanf-local and suppose $\mathsf{hlr}(\psi) = d$. Let $\vec{a}$ and $\vec{b}$ be n-tuples of $\mathcal{A}$ such that $\vec{a} \approx_{3d+1} \vec{b}$. Since $\mathcal{A} \leftrightarrows_d \mathcal{A}$, by Lemma 4.4, $(\mathcal{A}, \vec{a}) \leftrightarrows_d (\mathcal{A}, \vec{b})$. From Hanf-locality of ψ we see that $\mathcal{A} \models \psi(\vec{a})$ iff $\mathcal{A} \models \psi(\vec{b})$. Thus ψ is Gaifman-local and $\mathsf{lr}(\psi) \leq 3d + 1$. □

By Proposition 5.2, the following holds as well.

Corollary 5.4 *Let $\mathcal{L}$ be a logic that is closed under first-order operations. Assume that every sentence in $\mathcal{L}$ is Hanf-local. Then every formula in $\mathcal{L}$ is Gaifman-local.* □

The proof above also shows that $\mathsf{lr}(\psi) \leq 3 \cdot \mathsf{hlr}(\Psi^{(n)}) + 1$. In the case ψ is a first-order formula, $\Psi^{(n)}$ is of quantifier rank $\mathsf{qr}(\psi) + n$, and hence we obtain a new bound that improves Gaifman's $(7^{\mathsf{qr}(\psi)} - 1)/2$.

Corollary 5.5 *Let $\psi(x_1, \ldots, x_n)$ be a first-order formula. Then $\mathsf{lr}(\psi) \leq 3^{\mathsf{qr}(\psi)+n} + 1$.* □

Note that this improves the locality rank implied by Gaifman's theorem, not the bound on the size of neighborhood in an *explicitly* constructed formula used in Gaifman's proof.

We now list some corollaries of Theorem 5.3. We immediately obtain

Corollary 5.6 *Every Hanf-local formula has the bounded degree property. If $\mathcal{L}$ is a logic closed under first-order operations and such that every $\mathcal{L}$-sentence is Hanf-local, then $\mathcal{L}$ has the bounded degree property.* □

Corollary 5.7 *$\mathcal{FO}(\mathbf{Q}_u)$ and $\mathcal{FO} + COUNT$ are Gaifman-local and have the bounded degree property.* □

More precisely, every $\mathcal{FO} + COUNT$ formula without free variables over the numerical domain is Gaifman-local and has the BDP. This generalizes a number of known results. For example, the bounded degree property of first-order logic with Härtig and Rescher quantifiers (proved in [6] by a lengthy and quite involved argument) follows straightforwardly. We also obtain a theorem by Etessami [9] that deterministic transitive closure is not definable in $\mathcal{FO} + COUNT$ in the presence of a successor relation. Note that this can be viewed as a small step towards separating TC^0 from DLOGSPACE, because $\mathcal{FO} + COUNT$ captures uniform TC^0 on linearly ordered structures [3] and $\mathcal{FO}$ with deterministic transitive closure captures DLOGSPACE with built-in successor relation [8, 18].

Corollary 5.7 allows us to make the next incremental step. First, recall from Section 3 that with the counting quantifiers $\exists i x \varphi(x)$ in $\mathcal{FO} + COUNT$ we can use the built-in relations (like $<$ and BIT) of the numerical second sort. Let then $k \in \mathbb{N}$ and let $\mathcal{S}_k$ be any family of built-in relations on the *non-numerical* domain whose degrees do not exceed k.

Corollary 5.8 *Deterministic transitive closure is not definable in $\mathcal{FO} + COUNT$ with the built-in relations $\mathcal{S}_k$.* [1] □

Furthermore, using locality, we can extend the above results to more complex auxiliary data. Consider a class of structures $\mathcal{C} \subseteq \mathrm{STRUCT}[\sigma']$ for some relational vocabulary σ'. Define a function $s_{\mathcal{C}} : \mathbb{N} \to \mathbb{N}$ by letting $s_{\mathcal{C}}(n)$ be the maximal possible degree in some n-element structure $\mathcal{A} \in \mathcal{C}$. We say that $\mathcal{C}$ is of *moderate degree* (see [11]) if $s_{\mathcal{C}}(n) \leq \log^{o(1)} n$. That is, there is a function $\delta : \mathbb{N} \to \mathbb{N}$ such that $\lim_{n\to\infty} \delta(n) = 0$ and $s_{\mathcal{C}}(n) \leq \log^{\delta(n)} n$.

The following was shown in [6].

Proposition 5.9 (see [6]) *Let ψ be a local graph query, of locality rank r. Then for any structure $\mathcal{A}$, the number of distinct in-degrees in the graph $\psi[\mathcal{A}]$ is at most the number of non-isomorphic $3r + 1$-neighborhoods realized in $\mathcal{A}$. The same is true for out-degrees.* □

[1] In fact, if $\mathcal{T}$ is any set of built-in relations defined on the *numerical* domain, then deterministic transitive closure still is not definable in $\mathcal{FO} + COUNT$ with the built-in relations $\mathcal{S}_k$ and the built-in relations $\mathcal{T}$.

Now one can use this proposition and calculate that, for structures of moderate degree, one cannot construct a graph that has n distinct in-degrees, where n is the number of nodes. This, and locality of $\mathcal{FO} + \mathit{COUNT}$, gives us

Corollary 5.10 *Transitive closure and deterministic transitive closure are* not *definable in* $\mathcal{FO} + \mathit{COUNT}$ *in the presence of built-in relations of moderate degree.* $\square$

However, the order relation adds all degrees from 0 to the cardinality of the input. Thus, we cannot generalize Corollary 5.8 to the case of built-in linear order.

5.2 Strong Gaifman-locality implies Hanf-locality

The next main result of the section is:

Theorem 5.11 *Every strongly Gaifman-local sentence is Hanf-local.*

From this and Gaifman's theorem, the theorem by Fagin, Stockmeyer and Vardi follows immediately (though not the bound produced by the proof in [11]). We also believe that the proof below is simpler than that in [11] and shows clearly why this result is indeed a form of locality, as claimed in [11].

Proof of Theorem 5.11. It is enough to consider a sentence Ψ which is equivalent to $\exists x_1 \ldots \exists x_n \psi(x_1, \ldots, x_n)$, where $\psi(\vec{x})$ is strongly Gaifman-local. Assume that $r > 0$ witnesses strong locality of ψ: that is, $\vec{a} \approx_r^{\mathcal{A},\mathcal{B}} \vec{b}$ implies $\mathcal{A} \models \psi(\vec{a})$ iff $\mathcal{B} \models \psi(\vec{b})$. Let d be given by Corollary 4.9; then $\mathcal{A} \leftrightarrows_d \mathcal{B}$ implies $\mathcal{A} \leftrightarrows_{n,r} \mathcal{B}$. We claim that $\mathsf{hlr}(\Psi) \leq d$. Indeed, assume $\mathcal{A} \leftrightarrows_d \mathcal{B}$. Let $\mathcal{A} \models \Psi$. Then $\mathcal{A} \models \psi(\vec{a})$ for some $\vec{a} \in A^n$. By Corollary 4.9, $\mathcal{A} \leftrightarrows_{n,r} \mathcal{B}$, and thus we find $\vec{b} \in B^n$ such that $\vec{b} \approx_r \vec{a}$. From strong Gaifman-locality of ψ we see $\mathcal{B} \models \psi(\vec{b})$ and thus $\mathcal{B} \models \Psi$. The converse (that is, $\mathcal{B} \models \Psi$ implies $\mathcal{A} \models \Psi$) is similar. Hence, $\mathsf{hlr}(\Psi) \leq d$, which completes the proof. $\square$

From Proposition 5.2 we get the following corollary.

Corollary 5.12 *Let $\mathcal{L}$ be a logic that is closed under first-order operations. Assume that every sentence in $\mathcal{L}$ is strongly Gaifman-local. Then every formula in $\mathcal{L}$ is Hanf-local.*

Combining the proof above with Gaifman's theorem, we see that for an arbitrary *first-order* sentence Ψ, we have the bound $\mathsf{hlr}(\Psi) \leq 2 \cdot 3^{\mathsf{qr}(\Psi)} \cdot 7^{\mathsf{qr}(\Psi)-1}$, which is much worse than $3^{\mathsf{qr}(\Psi)}$ that is given by [11]. However, it is not the bound itself, but its existence that is used in most applications. Also, the above proof reveals the close connection between Gaifman's and Hanf's theorems.

Another corollary of Theorem 5.11 is that the two parts of Gaifman's theorem are not independent:

Corollary 5.13 *Let $\mathcal{L}$ be a logic that is closed under first-order operations. Assume that every sentence in $\mathcal{L}$ is strongly Gaifman-local. Then every formula in $\mathcal{L}$ is Gaifman-local.* $\square$

6 Locality and structures of small degree

In this section we give characterizations of the notions of locality on structures of bounded degree. We start with a simple observation:

Lemma 6.1 *For any signature σ, there exist functions $f_\sigma, F_\sigma : \mathbb{N} \times \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ such that for any $\mathcal{A} \in \mathrm{STRUCT}_k[\sigma]$,*

$$card(\{\tau^n \mid \exists \vec{a} \in A^n : \tau_r(\mathcal{A}, \vec{a}) = \tau^n\}) \leq f_\sigma(k, r, n) \text{ and}$$
$$\forall \vec{a} \in A^n : \quad card(S_r^{\mathcal{A}}(\vec{a})) \leq F_\sigma(k, r, n).$$

□

The next two lemmas show us that on structures of bounded degree the relations $\approx_r$ and $\leftrightarrows_r$ are definable by formulas of certain logics.

Recall from Section 2 that for every fixed r there is a first-order formula $d_{\leq r}(z, \vec{x})$ which expresses that $d(z, \vec{x}) \leq r$, i.e., $d(z, x) \leq r$ for some component x of $\vec{x}$. A proof of the first lemma is essentially given already in [8, Section 1], and we only sketch the proof below.

Lemma 6.2
- *For every $\mathcal{A}$, $\vec{a} \in A^n$ and a positive integer r, there exists a first-order formula $\varphi^r_{\mathcal{A},\vec{a}}(\vec{x})$ such that for every $\mathcal{B}$ and $\vec{b} \in B^n$, $\mathcal{B} \models \varphi^r_{\mathcal{A},\vec{a}}(\vec{b})$ iff $\vec{a} \approx_r \vec{b}$.*
- *For every $\mathcal{A}$ and positive integers r and n, there exists a first-order sentence $\theta^{r,n}_{\mathcal{A}}$ such that for every $\mathcal{B}$, $\mathcal{B} \models \theta^{r,n}_{\mathcal{A}}$ iff exactly the same isomorphism types of n-tuples are r-realized in $\mathcal{A}$ and $\mathcal{B}$.*

Proof: We define $\varphi^r_{\mathcal{A},\vec{a}}(\vec{x})$ to be a first-order formula which says that $\vec{x}$ realizes the isomorphism type of $N_r^{\mathcal{A}}(\vec{a})$. For this, let $\vec{a} \in A^n$ and let $a_1, \ldots, a_n, b_1, \ldots, b_m$ be the elements of $S_r^{\mathcal{A}}(\vec{a})$. Let $\rho(x_1, \ldots, x_n, y_1, \ldots, \ldots, y_m)$ be the diagram of $N_r^{\mathcal{A}}(\vec{a})$, that is, the conjunction of atomic and negated atomic formulas realized in $N_r^{\mathcal{A}}(\vec{a})$. Then $\varphi^r_{\mathcal{A},\vec{a}}(\vec{x})$ can be defined as

$$\exists y_1, \ldots \exists y_m \forall z \left(\rho(x_1, \ldots, x_n, y_1, \ldots, y_m) \wedge \left(d_{\leq r}(z, \vec{x}) \leftrightarrow (\bigvee_{i=1}^{n} z = x_i \vee \bigvee_{i=1}^{m} z = y_i) \right) \right).$$

Suppose that for $\mathcal{B}$ and $\vec{b} \in B^n$ we have $\mathcal{B} \models \varphi^r_{\mathcal{A},\vec{a}}(\vec{b})$. Then $N_r^{\mathcal{A}}(\vec{a})$ and $N_r^{\mathcal{B}}(\vec{b})$ satisfy the same atomic formulas and the second part of the definition of $\varphi^r_{\mathcal{A},\vec{a}}(\vec{x})$ says that there are no other points in $S_r^{\mathcal{B}}(\vec{b})$. Hence $\vec{a} \approx_r \vec{b}$. Suppose then $\vec{a} \approx_r \vec{b}$, that is, $N_r^{\mathcal{A}}(\vec{a}) \cong N_r^{\mathcal{B}}(\vec{b})$. By construction we have $\mathcal{A} \models \varphi^r_{\mathcal{A},\vec{a}}(\vec{a})$ and hence also $\mathcal{B} \models \varphi^r_{\mathcal{A},\vec{a}}(\vec{b})$.

For the second claim, suppose $\vec{a} \in A^n$ and consider a first-order formula $\varphi^r_{\mathcal{A},\vec{a}}(\vec{x})$ which says that $\vec{x}$ r-realizes the isomorphism type of $N_r^{\mathcal{A}}(\vec{a})$. We define the first-order sentence $\theta^{r,n}_{\mathcal{A}}$ as

$$\bigvee_{\vec{a} \in A^n} \exists \vec{x} \varphi^r_{\mathcal{A},\vec{a}}(\vec{x}) \wedge \forall \vec{x} \bigvee_{\vec{a} \in A^n} \varphi^r_{\mathcal{A},\vec{a}}(\vec{x}).$$

If the same isomorphism types of n-tuples are realized in $\mathcal{A}$ and $\mathcal{B}$, then obviously $\mathcal{B} \models \theta^{r,n}_{\mathcal{A}}$. Suppose that $\mathcal{B} \models \theta^{r,n}_{\mathcal{A}}$. Then the first part of $\theta^{r,n}_{\mathcal{A}}$ implies that every isomorphism type of an n-tuple which is r-realized in $\mathcal{A}$, is also r-realized in $\mathcal{B}$. The second part of the formula in turn says that no other isomorphism types are realized in $\mathcal{B}$. Thus exactly the same isomorphism types are realized in $\mathcal{A}$ and $\mathcal{B}$. □

The second lemma below gives a formula which defines the relation $\leftrightarrows_r$. It shows us that on structures of bounded degree unary quantifiers allow us to keep quantifier rank bounded.

Lemma 6.3 *For every $\mathcal{A} \in \mathrm{STRUCT}_k[\sigma]$, $\vec{a} \in A^n$ and a positive integer r, there exists a positive integer m, which is independent of $\mathcal{A}$, and an $\mathcal{FO}(\mathbf{Q}_u)^m$-formula $\eta^r_{\mathcal{A},\vec{a}}(\vec{x})$ such that, for every $\mathcal{B}$ and $\vec{b} \in B^n$, $\mathcal{B} \models \eta^r_{\mathcal{A},\vec{a}}(\vec{b})$ iff $(\mathcal{A},\vec{a}) \leftrightarrows_r (\mathcal{B},\vec{b})$. Here m can be taken to be $F_\sigma(k,r,n) - n + 1$.*

Proof: Suppose $\mathcal{A} \in \mathrm{STRUCT}_k[\sigma]$ and $\vec{a} \in A^n$. For every $b \in A$ let $\varphi^r_{\mathcal{A},\vec{a}b}(\vec{x}y)$ be the first-order formula given by Lemma 6.2 which describes the isomorphism type of $N^{\mathcal{A}}_r(\vec{a}b)$. By Lemma 6.1 we see that $\mathsf{qr}(\varphi^r_{\mathcal{A},\vec{a}b}) \leq F_\sigma(k,r,n) - n$. Then for every isomorphism type $N^{\mathcal{A}}_r(\vec{a}b)$ we can express the number of points b which realize this isomorphism type by $\exists^{=j} y \varphi^r_{\mathcal{A},\vec{a}b}(\vec{x}y)$. Here $\exists^{=j} y$ is the unary quantifier given by the class of structures $\langle A, U \rangle$ where U is a j-element subset of A.

From Lemma 6.1 we see that there exists a bound M on the number of different isomorphism types $N^{\mathcal{C}}_r(\vec{d}e)$ in structures from $\mathrm{STRUCT}_k[\sigma]$. Let these isomorphism types be $\tau_1, \ldots, \tau_M$ and let $\beta_i(\vec{x}y)$ be the first-order formula given by Lemma 6.2 which describes the isomorphism type τ_i. Denote $n_i = \#_r[(\mathcal{A},\vec{a}), \tau_i]$. Then the formula $\eta^r_{\mathcal{A},\vec{a}}(\vec{x})$ can be defined as

$$\bigwedge_{1 \leq i \leq M} \exists^{=n_i} y \beta_i(\vec{x}y).$$

If $(\mathcal{A},\vec{a}) \leftrightarrows_r (\mathcal{B},\vec{b})$ then obviously $\mathcal{B} \models \eta^r_{\mathcal{A},\vec{a}}(\vec{b})$. By the definition of $\eta^r_{\mathcal{A},\vec{a}}(\vec{x})$, if $\mathcal{B} \models \eta^r_{\mathcal{A},\vec{a}}(\vec{b})$, then there are exactly the same number of points in $\mathcal{B}$ with the same isomorphism type as $N^{\mathcal{A}}_r(\vec{a}b)$ for every $b \in A$. Hence there exists a bijection $A \to B$ which shows that $(\mathcal{A},\vec{a}) \leftrightarrows_r (\mathcal{B},\vec{b})$. $\square$

We first give a characterization for strongly Gaifman-local queries. We say that a Boolean query Ψ is strongly Gaifman-local on $\mathrm{STRUCT}_k[\sigma]$ if in Definition 2.2 $\mathrm{STRUCT}[\sigma]$ is replaced by $\mathrm{STRUCT}_k[\sigma]$, i.e., we restrict the consideration to structures where each point has degree at most k. The idea of the proof given below is similar to the one in [25], where a characterization for Boolean queries definable in $\mathcal{FO}$ (and in $\mathcal{FO}$ with modular counting quantifiers) on structures of bounded degree, was given.

Proposition 6.4 *Let Ψ be a Boolean query and k a natural number. Then Ψ is strongly Gaifman-local on $\mathrm{STRUCT}_k[\sigma]$ iff Ψ is definable in first-order logic.*

Proof: The implication from right to left is already established in Proposition 2.3. For the converse, let Ψ be strongly Gaifman-local, let r witness strong Gaifman-locality. That is, Ψ is a Boolean combination of sentences of the form $\exists \vec{x} \psi(\vec{x})$ where each $\vec{a} \approx^{\mathcal{A},\mathcal{B}}_r \vec{b}$ implies $\mathcal{A} \models \psi(\vec{a})$ iff $\mathcal{B} \models \psi(\vec{b})$. Let n be the maximum length of the tuples $\vec{x}$ in these formulas ψ. Consider then the sentence Φ:

$$\bigvee_{\substack{\mathcal{A} \models \Psi \\ \mathcal{A} \in \mathrm{STRUCT}_k[\sigma]}} \theta^{r,n}_{\mathcal{A}}$$

where the sentences $\theta^{r,n}_{\mathcal{A}}$ are given by Lemma 6.2. Since there are only 2^m sentences $\theta^{r,n}_{\mathcal{A}}$ when $\mathcal{A} \in \mathrm{STRUCT}_k[\sigma]$ (up to logical equivalence), where $m = f_\sigma(k,r,n)$, this disjunction is finite and hence Φ is a first-order sentence. Intuitively, Φ describes the isomorphism types of n-tuples which are r-realized in the structures satisfying Ψ.

We claim that Φ is equivalent to Ψ on structures from $\mathrm{STRUCT}_k[\sigma]$. Suppose that $\mathcal{B} \in \mathrm{STRUCT}_k[\sigma]$ and $\mathcal{B} \models \Psi$. Then $\mathcal{B} \models \theta^{r,n}_{\mathcal{B}}$ and thus $\mathcal{B} \models \Phi$. If $\mathcal{B} \models \Phi$ then $\mathcal{B} \models \theta^{r,n}_{\mathcal{A}}$ for some $\mathcal{A} \in \mathrm{STRUCT}_k[\sigma]$ for which $\mathcal{A} \models \Psi$. Then exactly the same isomorphism types of n-tuples are r-realized in $\mathcal{A}$ and $\mathcal{B}$. Thus, $\mathcal{A}$ and $\mathcal{B}$ agree on every sentence of the form $\exists \vec{x} \psi(\vec{x})$, where r witnesses strong Gaifman-locality of ψ. This implies $\mathcal{B} \models \Psi$. $\Box$

Next we give a characterization for Hanf-local queries in terms of logical expressibility.

Proposition 6.5 *Let k be a natural number, and let $\psi(\vec{x})$ be a query on* $\mathrm{STRUCT}_k[\sigma]$. *Then ψ is Hanf-local iff ψ is definable in $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$.*

Proof: The implication from right to left follows from Theorem 3.3 and Proposition 5.2. For the other direction, let $\psi(x_1, \ldots, x_n)$ be Hanf-local and $\mathsf{hlr}(\psi) = r$. Consider the formula $\varphi(\vec{x})$ defined as

$$\bigvee_{\substack{\mathcal{A} \models \psi(\vec{a}) \\ \mathcal{A} \in \mathrm{STRUCT}_k[\sigma]}} \eta^r_{\mathcal{A},\vec{a}}(\vec{x})$$

where the formulas $\eta^r_{\mathcal{A},\vec{a}}(\vec{x})$ are given by Lemma 6.3. Since each $\eta^r_{\mathcal{A},\vec{a}}(\vec{x})$ is an $\mathcal{FO}(\mathbf{Q}_u)^m$-formula, $\varphi(\vec{x})$ is an $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^m$-formula.

We show that $\varphi(\vec{x})$ and $\psi(\vec{x})$ are equivalent on $\mathrm{STRUCT}_k[\sigma]$. Let $\mathcal{B} \in \mathrm{STRUCT}_k[\sigma]$ and $\vec{b} \in B^n$. If $\mathcal{B} \models \psi(\vec{b})$ then $\mathcal{B} \models \eta^r_{\mathcal{B},\vec{b}}(\vec{b})$ and hence $\mathcal{B} \models \varphi(\vec{b})$. Suppose then that $\mathcal{B} \models \varphi(\vec{b})$. Now $\mathcal{B} \models \eta^r_{\mathcal{A},\vec{a}}(\vec{b})$ for some $\mathcal{A}$ and $\vec{a}$ which satisfy $\mathcal{A} \models \psi(\vec{a})$. Then $(\mathcal{A}, \vec{a}) \leftrightarrows_r (\mathcal{B}, \vec{b})$ by Lemma 6.3 and since ψ has Hanf locality rank r, we have $\mathcal{B} \models \psi(\vec{b})$. $\Box$

In particular, when studying the bounded degree property, Hanf-locality can be replaced by definability in $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$. Using this, we can now give a very simple alternative proof for Corollary 5.6: Let $\psi(x_1, \ldots, x_n)$ be a formula of $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^r$. Then for each j and l there is a formula $\varphi_{j,l}(x)$ of $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^{r+n-1}$ which says that $\mathit{degree}_j(\psi, x)$, the jth degree of x in the output of ψ, is exactly l. Hence, if for any two points a and b we have $a \approx_d b$, where $d = 3^{r+n-1}$, then a and b have the same degrees with respect to ψ. But by Lemma 6.1 the number of different isomorphism types of 3^{r+n-1}-neighborhoods realized in structures from $\mathrm{STRUCT}_k[\sigma]$ is bounded, and thus ψ has the bounded degree property.

To describe Gaifman-local formulae on structures of small degree, we need the following definition.

Definition 6.6 *A formula $\psi(x_1, \ldots, x_n)$ in a language σ is given by a* first-order definition by cases *on a class $\mathcal{C} \subseteq$* $\mathrm{STRUCT}[\sigma]$ *if there exists a partition $\mathcal{C} = \mathcal{C}_1 \cup \ldots \cup \mathcal{C}_m$ and first-order formulae $\alpha_1(x_1, \ldots, x_n), \ldots, \alpha_m(x_1, \ldots, x_n)$ in the language σ, such that on all structures $\mathcal{A} \in \mathcal{C}_i$, ψ is equivalent to α_i. That is, for any $1 \leq i \leq m$ and any $\mathcal{A} \in \mathcal{C}_i$,*

$$\mathcal{A} \models \forall \vec{x} (\psi(\vec{x}) \leftrightarrow \alpha_i(\vec{x})).$$

This is reminiscent of the familiar `case` statement (or equivalently, a nested `if-then-else` statement) in many programming languages.

Proposition 6.7 *Let $\psi(x_1, \ldots, x_n)$ be a formula and k a natural number. Then ψ is Gaifman-local on structures from* $\mathrm{STRUCT}_k[\sigma]$ *iff ψ is given by a first-order definition by cases on* $\mathrm{STRUCT}_k[\sigma]$. *Furthermore, if ψ is given by a first-order definition by cases on* $\mathrm{STRUCT}_k[\sigma]$, *and ψ is in a logic $\mathcal{L}$*

that is closed under first-order operations, then each class $\mathcal{C}_i$ of the corresponding partition is definable by a sentence in $\mathcal{L}$.

Proof: That a formula given by a first-order definition by cases is Gaifman-local follows from locality of first-order formulae. Indeed, $\mathsf{lr}(\psi) \leq \max_{i=1}^{m}\{\mathsf{lr}(\alpha_i)\}$. For the converse, assume that ψ is of locality rank r. We know that there exists a bound, $M = f_\sigma(k,r,n)$, on the number of different isomorphism types of r-neighborhoods of n-tuples in structures from $\mathrm{STRUCT}_k[\sigma]$, see Lemma 6.1. Let $\tau_1, \ldots, \tau_M$ be an enumeration of those isomorphism types, and let $\beta_i(\vec{x})$ be the first-order formula given by Lemma 6.2 saying that $\vec{x}$ r-realizes τ_i. Let Φ_i be the sentence $\exists \vec{y}(\beta_i(\vec{y}) \wedge \psi(\vec{y}))$ (note that Φ_i is not necessarily in $\mathcal{L}$, unless $\mathcal{L}$ is closed under first-order operations). We now claim that $\psi(\vec{x})$ is equivalent to

$$\varphi(\vec{x}) \quad = \quad \bigvee_{i=1}^{M} (\beta_i(\vec{x}) \wedge \Phi_i)$$

on structures from $\mathrm{STRUCT}_k[\sigma]$. Indeed, if $\mathcal{A} \in \mathrm{STRUCT}_k[\sigma]$ and $\mathcal{A} \models \psi(\vec{a})$, assuming that $\vec{a}$ realizes τ_i in $\mathcal{A}$, we see that $\mathcal{A} \models \Phi_i$ and thus $\mathcal{A} \models \varphi(\vec{a})$. Conversely, let $\mathcal{A} \models \varphi(\vec{a})$; that is, $\mathcal{A} \models \beta_i(\vec{a})$ and $\mathcal{A} \models \Phi_i$ for some i. In particular, there exists an n-tuple $\vec{b}$ such that $\mathcal{A} \models \beta_i(\vec{b}) \wedge \psi(\vec{b})$. Since $\mathcal{A} \models \beta_i(\vec{b})$, we obtain $\vec{a} \approx_r^{\mathcal{A}} \vec{b}$, and thus $\mathcal{A} \models \psi(\vec{a})$ by locality, which proves the claim.

Finally, for each subset $I \subseteq \{1, \ldots, M\}$, let $\Xi_I = (\bigwedge_{i \in I} \Phi_i) \wedge (\bigwedge_{i \notin I} \neg \Phi_i)$, and let $\alpha_I(\vec{x})$ be a first-order formula equivalent to $\bigvee_{i \in I} \beta_i(\vec{x})$. Then we still have

$$\psi(\vec{x}) \quad \leftrightarrow \quad \bigvee_{I \subseteq \{1,\ldots,M\}} (\alpha_I(\vec{x}) \wedge \Xi_I).$$

Since the classes $\mathcal{C}_I = \{\mathcal{A} \in \mathrm{STRUCT}_k[\sigma] \mid \mathcal{A} \models \Xi_I\}$ form a partition of $\mathrm{STRUCT}_k[\sigma]$, the above gives the desired first-order definition by cases. Furthermore, the second claim follows from the proof. $\Box$

This proposition gives us yet another proof of the bounded degree property of arbitrary local formulae (assuming the BDP of first-order formulae). Indeed, for each k, the upper bound on *deg_count*$(\psi[\mathcal{A}])$ for $\mathcal{A} \in \mathrm{STRUCT}_k[\sigma]$ can be calculated as the maximum of $f_{\alpha_i}(k)$, where α_is come from the first-order definition by cases of ψ, and f_{α_i} is the function giving the bound for first-order formula α_i.

As another corollary of the above characterization of locality, we have the following Ramsey-style property, similar to those studied in [4].

Corollary 6.8 *Let $\mathcal{C}$ be an infinite class of structures in $\mathrm{STRUCT}_k[\sigma]$. Let $\psi(\vec{x})$ be a Gaifman-local formula. Then there exists an infinite subclass $\mathcal{C}' \subseteq \mathcal{C}$ and a first-order formula $\varphi(\vec{x})$ that is equivalent to ψ on $\mathcal{C}'$. That is, $\mathcal{A} \models \forall \vec{x}(\psi(\vec{x}) \leftrightarrow \varphi(\vec{x}))$ for all $\mathcal{A} \in \mathcal{C}'$.* $\Box$

Studying the bounded degree property on structures of small degree appears to be of little interest, since any query producing a result that has a small number of in- and out-degrees will have the BDP. This leads to an easy characterization of formulae that have the BDP on structures of small degree. Namely, let R be an n-ary relational symbol, and χ_m a formula in the language $\{R\}$ such that $\mathcal{A} \models \chi_m(\vec{a})$ iff $\vec{a} \in R^{\mathcal{A}}$ and there are at most m distinct vectors of the form *deg_count*(b) under *deg_count*(a_i), $a_i \in \vec{a}$, in the lexicographic order. Here by *deg_count*(b) we mean the vector $(\mathit{degree}_1(R^{\mathcal{A}}, b), \ldots, \mathit{degree}_n(R^{\mathcal{A}}, b))$. Note that χ_m is easily definable in $\mathcal{L}^{\omega}_{\infty\omega}$. Thus, formulae having the bounded degree property on structures of small degree can be characterized as those of the form $\chi_m[\varphi(\vec{x})/R(\vec{x})]$, where φ is an arbitrary $\mathcal{L}^{\omega}_{\infty\omega}$ formula (that is, each occurrence of $R(\vec{x})$ is replaced by $\varphi(\vec{x})$).

7 Conclusion

We examined the main notions of locality of first-order formulae, and proved that these notions are closely related. We showed that the Hanf-locality implies the Gaifman-locality for open formulae, and the strong Gaifman-locality for sentences implies the Hanf-locality. Each of these implies the bounded degree property, which is one of the easiest tools for proving inexpressibility results. Our presentation goes beyond the first-order case, and thus allows us to infer new results for logics with unary quantifiers and counting.

We believe that the most challenging problem is to extend these techniques to the ordered setting. Note that on ordered structures all queries are trivially strongly Gaifman-local; this is simply because the distance between any two elements in the Gaifman graph of an ordered structure is at most 1. Thus, to obtain interesting results one should restrict the attention to *order-independent* queries, such as the transitive closure. More precisely, we are interested in structures of the form $\langle A, R_1, \ldots, R_l, < \rangle$, where, in addition to σ-relations $R_1, \ldots, R_l$, we have a binary relation $<$ which is interpreted as a linear order. When we talk about neighborhoods, we mean neighborhoods in the σ-reduct $\langle A, R_1, \ldots, R_l \rangle$ – this gives us the definitions of all the notions of locality. Finally, if we have a σ-structure $\mathcal{A}$, and an ordering $<$ on A, by $\mathcal{A}_<$ we denote the corresponding ordered structure. The order-independence of a query $\psi(\vec{x})$ in a language that includes both σ and $<$ means that for any $\mathcal{A}$, and any two orderings $<_1$ and $<_2$, we have $\mathcal{A}_{<_1} \models \psi(\vec{a})$ iff $\mathcal{A}_{<_2} \models \psi(\vec{a})$ for every $\vec{a}$.

The first natural question is whether the locality properties of first-order logic could be extended to the order-independent setting. A positive answer to this question was recently obtained by Grohe and Schwentick [13], who proved that all order-independent queries in $\mathcal{FO}$ are Gaifman-local. On the other hand, the corresponding problem for Hanf-locality is still open.

It would be tempting to conjecture that the Gaifman-locality of $\mathcal{FO} + COUNT$ could also be extended to the case of order-independent queries. Indeed, this would imply that deterministic transitive closure is not in TC^0, which in turn would imply the separation of TC^0 and DLOGSPACE. However, the following counterexample shows that this conjecture, made in [21], is false.

Proposition 7.1 *There is an order-independent query ψ in $\mathcal{FO} + COUNT$ which does not have the bounded degree property, and hence is not Gaifman-local.*

Proof: Consider structures of the type $\mathcal{A} = \langle A, P, E \rangle$, where $P \subseteq A$ and $\langle A, E \rangle$ is a directed graph such that $E \subseteq P^2$ is the graph of a successor relation. Let $\theta(x, y)$ be a formula of $\mathcal{FO} + COUNT$ in the extended signature $\{P, E, <\}$ saying that $x \in P$ and $card(\{a \in P \mid a < x\}) = card(\{b \in A \mid b < y\})$. Thus, for each ordering $<$ of A, θ defines in $\mathcal{A}_<$ a bijection from P to an initial segment P' of $<$. Let E' be the image of the relation E under this bijection. Clearly E' is definable in $\mathcal{A}_<$ by a formula of $\mathcal{FO} + COUNT$, and the graph $\langle P', E' \rangle$ is an isomorphic copy of $\langle P, E \rangle$.

It is known that the BIT predicate corresponding to an ordering $<$ is definable in $\mathcal{FO} + COUNT$ (see [3]). The BIT predicate in turn can be used for encoding subsets of P' by elements of A as follows: for each $a \in A$, let $S_a = \{b \in P' \mid \mathrm{BIT}(b, a)\}$. If the initial segment P' is of length at most $\log(card(A))$, then all subsets of P' are encoded by at least one element, i.e., for every $S \subseteq P'$ there is $a \in A$ such that $S = S_a$. Hence, assuming that

$$(*) \qquad card(P) \leq \log(card(A))$$

we can simulate monadic second-order quantifiers over P' by first-order quantifiers over A. In particular, there is a formula $\varphi'(x, y)$ of $\mathcal{FO} + COUNT$ such that if $(*)$ holds, then for all $a, b \in P'$,

$\langle A, P', E', < \rangle \models \varphi'(a, b)$ iff there is a directed E'-path from a to b.

Let $\varphi(x, y)$ be the formula $\exists u \exists v(\theta(x, u) \wedge \theta(y, v) \wedge \varphi'(u, v))$. Thus, $\varphi(x, y)$ says that $\varphi'(u, v)$ holds for the images u, v of x, y under the bijection defined by θ, which, assuming condition $(*)$, is equivalent to saying that there is a directed E-path from x to y. Let $\psi(x, y)$ be the conjunction of $\varphi(x, y)$ with a sentence of $\mathcal{FO} + COUNT$ expressing the condition $(*)$. Obviously ψ is order-independent, and it is easy to see that it does not have the bounded degree property, as it defines the transitive closure on a subset of nodes of size $\log(card(A))$. $\square$

Note that the restriction $card(P) \leq \log(card(A))$ is crucial in the proof above; the simulation of monadic second-order quantifiers over P' by first-order quantifiers over A is not possible without this assumption. Thus, it is still meaningful to ask whether some reasonable weak version of locality holds for order-independent queries in $\mathcal{FO} + COUNT$. For example, we may ask whether for each order-independent query $\psi(\vec{x})$ in $\mathcal{FO} + COUNT$ there is a sublinear function $f : \mathbb{N} \to \mathbb{N}$ such that if $\vec{a} \approx_r^{\mathcal{A}} \vec{b}$ for $r = f(card(A))$, then $\mathcal{A} \models \psi(\vec{a})$ iff $\mathcal{A} \models \psi(\vec{b})$.

Acknowledgement We would like to thank Alexei Stolboushkin, Michael Taitslin and anonymous referees for their helpful comments.

References

[1] S. Abiteboul, R. Hull, V. Vianu, *Foundations of Databases*, Addison Wesley, 1995.

[2] M. Ajtai and R. Fagin. Reachability is harder for directed than for undirected graphs. *Journal of Symbolic Logic*, 55(1):113–150, March 1990.

[3] D.A. Barrington, N. Immerman, H. Straubing. On uniformity within NC^1. *J. Comput. and Syst. Sci.*, 41:274–306,1990.

[4] M. Benedikt, L. Libkin. On the structure of queries in constraint query languages. In *Proc. 11th IEEE Symp. on Logic in Computer Science*, New Brunswick, NJ, July 1996, pages 25–34.

[5] M. Benedikt, T. Griffin and L. Libkin. Verifiable properties of database transactions. *Information and Computation*, 145 (1998), 57–88.

[6] G. Dong, L. Libkin, L. Wong. Local properties of query languages. *Theoretical Computer Science*, to appear. Extended abstract in *Proc. Int. Conf. on Database Theory*, Springer LNCS 1186, 1997, pages 140–154.

[7] H.-D. Ebbinghaus. Extended logics: the general framework. In J. Barwise and S. Feferman, editors, *Model-Theoretic Logics*, Springer-Verlag, 1985, pages 25–76.

[8] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory.* Springer Verlag, 1995.

[9] K. Etessami. Counting quantifiers, successor relations, and logarithmic space, *Journal of Computer and System Sciences*, 54 (1997), 400–411.

[10] R. Fagin. Easier ways to win logical games. In *Descriptive Complexity and Finite Models*, N. Immerman and Ph. Kolaitis, editors, AMS, 1997, pages 1–32.

[11] R. Fagin, L. Stockmeyer, M. Vardi, On monadic NP vs monadic co-NP, *Information and Computation*, 120 (1994), 78–92.

[12] H. Gaifman, On local and non-local properties, *in* "Proceedings of the Herbrand Symposium, Logic Colloquium '81," North Holland, 1982.

[13] M. Grohe and T. Schwentick. Locality of order-invariant first-order formulas. In *Proc. Conference on Mathematical Foundations of Computer Science (MFCS'98)*, Springer LNCS vol. 1450, Springer Verlag, 1998, pages 437–445.

[14] S. Grumbach, L. Libkin, T. Milo and L. Wong. Query languages for bags: expressive power and complexity. *SIGACT News*, 27 (1996), 30–37.

[15] Y. Gurevich. Logic and the challenge of computer science. In E. Börger, ed., *Current trends in theoretical computer science*, Computer Science Press, 1988, pages 1–57.

[16] W. Hanf. Model-theoretic methods in the study of elementary logic. In J.W. Addison et al, eds, *The Theory of Models*, North Holland, 1965, pages 132–145.

[17] L. Hella. Logical hierarchies in PTIME. *Information and Computation*, 129 (1996), 1–19.

[18] N. Immerman. Descriptive complexity: A logician's approach to computation. *Notices of the AMS* 42 (1995), 1127–1133.

[19] N. Immerman and E. Lander. Describing graphs: A first order approach to graph canonization. In *"Complexity Theory Retrospective"*, Springer Verlag, Berlin, 1990.

[20] Ph. Kolaitis and J. Väänänen. Generalized quantifiers and pebble games on finite structures. *Annals of Pure and Applied Logic* 74 (1995), 23–75.

[21] L. Libkin. On the forms of locality over finite models. In *Proc. 12th IEEE Symp. on Logic in Computer Science*, Warsaw, Poland, June–July 1996, pages 204–215.

[22] L. Libkin, L. Wong. Query languages for bags and aggregate functions. *Journal of Computer and System Sciences*, 55 (1997), 241–272.

[23] K. Luosto. Hierarchies of monadic generalized quantifiers. *Journal of Symbolic Logic*, to appear.

[24] J. Nurmonen. On winning strategies with unary quantifiers. *J. Logic and Computation*, 6 (1996), 779–798.

[25] J. Nurmonen. Counting modulo quantifiers on finite structures. *Information and Computation*, to appear. Extended abstract in *Proc. 11th IEEE Symp. on Logic in Computer Science*, New Brunswick, NJ, July 1996, pages 484–493.

[26] J. Nurmonen. Unary quantifiers and finite structures. PhD Thesis, University of Helsinki, 1996.

[27] T. Schwentick. On winning Ehrenfeucht games and monadic NP. *Annals of Pure and Applied Logic*, 79 (1996), 61–92.