

Counting and Locality over Finite Structures

A Survey

Leonid Libkin¹

Juha Nurmonen^{*2}

¹ Bell Laboratories

600 Mountain Avenue

Murray Hill, NJ 07974, USA

Email: libkin@bell-labs.com

² Department of Mathematics and Computer Science

University of Leicester

University Road

Leicester LE1 7RH, UK

Email: j.nurmonen@mcs.le.ac.uk

Abstract. We survey recent results on logics with counting and their local properties. We first consider game-theoretic characterizations of first-order logic and its counting extensions provided by unary generalized quantifiers. We then study Gaifman's and Hanf's locality theorems, their connection with game characterizations, and examples of their usage in proving expressivity bounds for first-order logic and its extensions. We review the abstract notions of Gaifman's and Hanf's locality, and show how they are related. We also consider a closely related bounded degree property, and demonstrate its usefulness in proving expressivity bounds. We discuss two applications. One deals with proving lower bounds for the complexity class TC^0 . In particular, we use logical characterization of TC^0 and locality theorems for first-order with counting quantifiers to provide lower bounds. We then explain how the notions of locality are used in database theory to prove that extensions of relational calculus with aggregate functions and grouping still lack the power to express fixpoint computation.

1 Introduction

Finite model theory is an active area of research, mostly due to its connections to theoretical computer science, in particular, database theory [1] and complexity theory [13, 28]. Several important complexity classes have nice logical characterizations. For instance, the star-free languages are exactly the ones definable in first-order logic FO, least fixpoint logic LFP captures PTIME, and partial fixpoint logic PFP captures PSPACE on ordered finite structures (see [13]). These logics also have their counterparts in the theory of database query languages.

* Supported by EPSRC grant GR/K 96564.

For example, relational calculus, that underlies majority of practical query languages, has exactly the power of first-order logic. The language Datalog with negation corresponds to LFP, and the extension of relational calculus with while loops corresponds to PFP (see [1]).

Several counting properties have also been considered in complexity and database theory. We give a few examples here. A logic can be extended with a linearly ordered second sort universe of numbers, and *counting quantifiers* ‘there are at least n elements.’ A very strong result in [9] shows that least fixpoint logic with such second sort counting fails to capture the complexity class PTIME. In circuit complexity theory, *counting modulo quantifiers* have been studied. Allowing gates which count inputs modulo a constant p , for every p , in the definition of AC^0 , one obtains the class ACC. Again, this class has a logical characterization, and the problem whether the containment $ACC \subseteq NC^1$ is strict, is still open. In [5], it was shown that the class TC^0 , that extends AC^0 with threshold gates, can be captured on ordered structures by first-order logic with the second sort counting, and simple arithmetic predicates on numbers.

In database theory, one often extends traditional first-order based languages with aggregate functions, such as summing up all values in a column in a relation, or finding the average value. One also extends those languages with grouping, that permits queries such as finding the average salary in each department. While the expressive power of relational calculus and many other query languages is well understood, much less is known about their aggregate extensions.

These applications have motivated a systematic study of the expressive power of counting properties from the viewpoint of finite model theory. In finite model theory, one normally uses games to prove expressivity bounds. For example, Ehrenfeucht-Fraïssé games are used for first-order logic, and bijective Ehrenfeucht-Fraïssé games are used for FO extended with unary generalized quantifiers. In Section 3 we review rules of such games for FO and some of its extensions with counting.

Playing a game often involves a complicated combinatorial argument. Many results on the expressive power of first-order logic and its extensions with counting, in particular, those obtained as an attempt to avoid game arguments, give us the intuition that these logics can express only local properties, and lack a mechanism for fixpoint computation. In Section 4, we review several results of this kind. We consider Gaifman’s theorem [18] for first-order logic, which shows that every first-order formula is equivalent to a local one, in the sense that only a small part of a structure is relevant for evaluating the query given by a formula. We also study modifications of Hanf’s result [22]. In this approach one counts the number of isomorphism types of fixed radius neighborhoods of points. If the result of this counting satisfies certain criteria, then the structures considered are guaranteed to be elementary equivalent in a certain logic. This technique has been modified for first-order logic [17], first-order logic with counting modulo quantifiers [41] and first-order logic extended by all unary generalized quantifiers [40], for the

case of finite structures. Proofs of applicability of Hanf’s technique typically are not very difficult [17, 15, 40, 42]. We will see some examples in Section 4.

The above results have motivated a study of general notions of locality [32, 24]. We review this line of work in Section 5. We show that Gaifman’s theorem gives rise to two general notions, one for sentences and one for open formulas. We formulate an abstract notion of locality that captures Hanf’s condition, and study the relationship between the notions of locality. We also consider the bounded degree property, which is implied by all other notions of locality, and is particularly well suited for proving expressivity bounds, especially for properties involving fixpoint computation.

In Section 6, we discuss applications of the concept of locality in complexity theory. In particular, we study the relationship between the circuit complexity class TC^0 and other complexity classes such as L and NL. In Section 7, we consider applications in database theory. We review results on expressive power of relational languages that resemble commercial languages such as SQL. We show that queries such as transitive closure, are inexpressible in a theoretical language that has the power of core SQL.

2 Preliminaries

A relational signature σ is a finite set of relation symbols $\{R_1, \dots, R_l\}$, each of which has an arity $p_i > 0$. We write σ_n for σ extended with n new constant symbols. A σ -structure is $\mathcal{A} = (A, \overline{R}_1, \dots, \overline{R}_l)$, where A is a non-empty set and $\overline{R}_i \subseteq A^{p_i}$ interprets R_i . If the *universe* A is finite, the structure \mathcal{A} is called finite. Unless mentioned otherwise, all structures considered here are assumed to be finite. When the notation is clear from the context, we write R_i in place of \overline{R}_i . The class of finite σ -structures is denoted by $\text{STRUCT}[\sigma]$. For instance, a graph $\mathcal{A} = (A, E)$ is a structure over a signature which consists of a single binary relation symbol E . If E is required to be symmetric, then \mathcal{A} is an undirected graph.

If $X \subseteq A$, by $\mathcal{A} \upharpoonright X$ we mean the structure with universe X , where the interpretation of each R_i is restricted to X . An *isomorphism* $f : \mathcal{A} \rightarrow \mathcal{B}$ is a bijection $A \rightarrow B$ such that $(a_1, \dots, a_{p_i}) \in \overline{R}_i$ holds in \mathcal{A} if and only if $(f(a_1), \dots, f(a_{p_i})) \in \overline{R}_i$ holds in \mathcal{B} . If $a_1, \dots, a_n \in A$ and $b_1, \dots, b_n \in B$, we say that $\{(a_i, b_i) \mid 1 \leq i \leq n\}$ is a *partial isomorphism* $\mathcal{A} \rightarrow \mathcal{B}$ if it is an isomorphism $\mathcal{A} \upharpoonright \{a_1, \dots, a_n\} \rightarrow \mathcal{B} \upharpoonright \{b_1, \dots, b_n\}$.

Every formula $\psi(x_1, \dots, x_m)$ with free variables x_1, \dots, x_m defines a *query* which maps a σ -structure \mathcal{A} to an m -ary relation $q_\psi(\mathcal{A}) = \{\mathbf{a} \in A^m \mid \mathcal{A} \models \psi(\mathbf{a})\}$. We denote the corresponding structure by $\psi[\mathcal{A}] = (A, q_\psi(\mathcal{A}))$. An m -ary query q is definable in a logic \mathcal{L} if there is a formula $\varphi(\mathbf{x})$ of that logic such that for every $\mathcal{A} \in \text{STRUCT}[\sigma]$ we have $q(\mathcal{A}) = \{\mathbf{a} \in A^m \mid \mathcal{A} \models \varphi(\mathbf{a})\}$. As an example, consider the *transitive closure* query $TRCL$. Suppose we are given a graph $\mathcal{A} = (A, E)$. Then $TRCL$ consists of all pairs $(a, b) \in A^2$ such that there

is an E -path from a to b , that is,

$$qTRCL = \left\{ (a, b) \in A^2 \mid \begin{array}{l} \exists a_1, \dots, a_k \in A \text{ such that } a = a_1 \text{ and } b = a_k \text{ and} \\ E(a_i, a_{i+1}) \text{ for all } i \end{array} \right\}.$$

This query is definable in many fixpoint logics, and in second-order logic, but we shall see that it is not definable in FO and its counting extensions.

The *Gaifman graph* $\mathcal{G}(\mathcal{A})$ of a σ -structure \mathcal{A} is the undirected graph (A, E) where $(a, b) \in E$ if and only if there is a tuple $\mathbf{t} \in \overline{R}_i$ for some $R_i \in \sigma$ such that $a, b \in \mathbf{t}$. Note that if \mathcal{A} is an undirected graph, then $\mathcal{A} = \mathcal{G}(\mathcal{A})$. The *degree* of a point is its degree in the Gaifman graph. The *distance* $d(a, b)$ is the length of the shortest path from a to b in $\mathcal{G}(\mathcal{A})$. For $a \in A$, its r -*sphere* is $S_r^{\mathcal{A}}(a) = \{b \in A \mid d(b, a) \leq r\}$. For an n -tuple \mathbf{t} we define $S_r^{\mathcal{A}}(\mathbf{t}) = \cup_{a \in \mathbf{t}} S_r^{\mathcal{A}}(a)$. The r -*neighborhood* of a tuple \mathbf{t} is the σ_n -structure $N_r^{\mathcal{A}}(\mathbf{t}) = (\mathcal{A} \upharpoonright S_r^{\mathcal{A}}(\mathbf{t}), \mathbf{t})$. That is, it is the restriction of \mathcal{A} to $S_r^{\mathcal{A}}(\mathbf{t})$, with \mathbf{t} being n distinguished constants. For instance, if \mathcal{A} is a chordless cycle of length at least $2r + 2$ and $a \in A$, then $N_r^{\mathcal{A}}(a)$ is the chordless path of length $2r + 1$ with one distinguished point a . We denote the *isomorphism type* of $N_r^{\mathcal{A}}(\mathbf{t})$ by $tp_r^{\mathcal{A}}(\mathbf{t})$. We emphasize here that if $h : N_r^{\mathcal{A}}(\mathbf{a}) \rightarrow N_r^{\mathcal{A}}(\mathbf{b})$ is an isomorphism, where $\mathbf{a} = (a_1, \dots, a_n)$ and $\mathbf{b} = (b_1, \dots, b_n)$, then $h(a_1) = b_1, \dots, h(a_n) = b_n$, as neighborhoods $N_r^{\mathcal{A}}(\mathbf{a})$ and $N_r^{\mathcal{A}}(\mathbf{b})$ are σ_n -structures.

If the structure \mathcal{A} is understood, we omit it from the notations when convenient.

3 Logics and games

In this section we introduce the logics that are considered in this paper. These logics are first-order logic FO and its extensions with various generalized quantifiers. All logics we consider are closed under Boolean connectives and first-order quantification, and are regular in the sense of [12, 31] (we do not go into details in this survey). We also review game-theoretic characterizations for elementary equivalence in these logics, and show how the games can be used to prove expressivity bounds.

We use the standard definitions for formulas and semantics of first-order logic FO (see e.g. [13]). Equality is treated as a special relation symbol which is not a member of the signature. The *quantifier rank* of a formula φ is defined to be the depth of quantifier nesting in φ , and is denoted by $qr(\varphi)$.

The rules of the *first-order Ehrenfeucht-Fraïssé game* are as follows. There are two players, *the spoiler* and *the duplicator*. Two σ -structures \mathcal{A} and \mathcal{B} and the number of rounds, say n , are given. In each round the spoiler first selects a point of one of the structures and the duplicator selects a point of the other structure. Let a_1, \dots, a_n and b_1, \dots, b_n be the points selected after the last round from \mathcal{A} and \mathcal{B} , respectively. The duplicator is declared the winner if $\{(a_i, b_i) \mid 1 \leq i \leq n\}$ is a partial isomorphism $\mathcal{A} \rightarrow \mathcal{B}$; otherwise the spoiler wins. We say that a player

has a *winning strategy* if he can guarantee a win, no matter how the other player plays. This game is interesting because of the following result.

Theorem 1 (cf. [13]). *Let $\mathcal{A}, \mathcal{B} \in \text{STRUCT}[\sigma]$. Then the duplicator has a winning strategy in the n -round Ehrenfeucht-Fraïssé game if and only if \mathcal{A} and \mathcal{B} agree on all first-order sentences of quantifier rank up to n .* \square

We use the notation $\mathcal{A} \equiv_{\text{FO}}^n \mathcal{B}$ if the duplicator has a winning strategy in the n -round game on \mathcal{A} and \mathcal{B} . The above theorem can be used to provide the following tool for proving expressivity bounds.

Corollary 1 (cf. [13]). *A class $\mathcal{C} \subseteq \text{STRUCT}[\sigma]$ is not definable in first-order logic if and only if for every n there are σ -structures $\mathcal{A} \in \mathcal{C}$ and $\mathcal{B} \notin \mathcal{C}$ such that the duplicator has a winning strategy in the n -round Ehrenfeucht-Fraïssé game on \mathcal{A} and \mathcal{B} .* \square

In other words, if for every n we can find one σ -structure from the class \mathcal{C} and another from the complement of \mathcal{C} , such that the duplicator can maintain a partial isomorphism for n rounds, then \mathcal{C} is not definable in FO.

We now give some examples. Note that in those examples, we do not spell out every single detail of the game argument – this may require more space than this entire section. We shall offer much simpler proofs of the existence of winning strategies in the next section, after we have introduced the ideas of locality.

First, one can use Ehrenfeucht-Fraïssé games to show that connectivity of finite graphs cannot be expressed in first-order logic. Assume that connectivity is definable, and take as a counterexample, for each finite n , \mathcal{A} to be a chordless cycle of length 2^n and \mathcal{B} to be a disjoint union of two chordless cycles of length 2^n . The duplicator's strategy in the n -round Ehrenfeucht-Fraïssé game is to preserve, in each round $j \leq n$, distances up to 2^{n-j} . The only way the spoiler could win is to show that in \mathcal{B} there are two points with no path between them, whereas in \mathcal{A} there always is a path between any two points. But if the spoiler cannot build a path between two points in n rounds, it does not matter how far these points are, or whether there is a path between them at all. This informal reasoning can be formalized to show that the duplicator has a winning strategy in the n -round Ehrenfeucht-Fraïssé game over \mathcal{A} and \mathcal{B} . This result holds also for ordered structures, see [13]. Note that this shows that the transitive closure query *TRCL* cannot be defined in FO. Indeed, assume that a first-order formula $\psi(x, y)$ defines *TRCL*. Then $\forall x \forall y \psi(x, y)$ would be a first-order sentence defining connectivity.

Next, assume that we are given two distinguished points a and b of a graph \mathcal{A} . Then we cannot define, in first-order logic, the property that a and b have equally many neighbors. As a counterexample we can take, for every n , a to be a point which has $n+2$ neighbors and b to be a point with $n+1$ neighbors in some graph \mathcal{A} . Since there are only n rounds available, the spoiler cannot demonstrate that b has fewer neighbors.

3.1 Unary quantifiers

A commonly used way to increase the expressive power of first-order logic is to extend it with *generalized quantifiers*, cf. [31, 48]. The basic idea is that we are given a class of structures, and we can check whether a substructure defined by a given family of formulas belongs to this class. In computational complexity theory, generalized quantifiers are often considered as oracles. In this section, however, we concentrate more on logical aspects of generalized quantifiers. We now review this method in detail in the case of *unary* generalized quantifiers.

Let σ_k^{unary} be a signature of k unary symbols. Suppose \mathcal{K} is an isomorphism closed class of σ_k^{unary} -structures. Then $\text{FO}(Q_{\mathcal{K}})$ is the extension of FO by a new formula formation rule:

if $\psi_i(x_i, \mathbf{y}_i)$ is a formula of $\text{FO}(Q_{\mathcal{K}})$ for $i = 1, \dots, k$, then
 $Q_{\mathcal{K}}x_1, \dots, x_k(\psi_1(x_1, \mathbf{y}_1), \dots, \psi_k(x_k, \mathbf{y}_k))$ is a formula of $\text{FO}(Q_{\mathcal{K}})$.

The corresponding semantic rule is:

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

Here the tuple \mathbf{a}_i gives the interpretation for those free variables in $\psi_i(x_i, \mathbf{y}_i)$ which are not equal to x_i . The extension $\text{FO}(\mathbf{Q})$ by a set \mathbf{Q} of unary quantifiers is defined similarly by adding to FO the above rule for each $Q_{\mathcal{K}} \in \mathbf{Q}$. The quantifier rank of a formula of $\text{FO}(\mathbf{Q})$ is defined as usually by the maximum depth of nesting of quantifiers (counting both first-order and generalized quantifiers). We also write $\text{FO}(\mathbf{Q}_u)$ for FO extended by *all* (continuum many) unary quantifiers. Note that the same definition can be used with other ambient logics, not just FO. In particular, we shall use $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)$, the infinitary logic extended with unary generalized quantifiers.

We now list some well-known examples.

- Example 1. (1)* The existential quantifier \exists corresponds to the class of structures $\{(A, P) \mid \emptyset \neq P \subseteq A\}$. Similarly, the universal quantifier \forall can be identified with the unary quantifier which is defined by the class $\{(A, P) \mid P = A\}$.
- (2) Counting quantifier $\exists^{\geq k}$* can be defined by the class $\{(A, P) \mid \text{card}(P) \geq k\}$. Note that each $\exists^{\geq k}$ can be defined in first-order logic; however, this requires k quantifiers \exists , and increases the quantifier rank by k . In contrast, the counting quantifier $\exists^{\geq k}$ increases the quantifier rank by one. We denote the set of all counting quantifiers $\exists^{\geq k}$ by \mathbf{C} , i.e., $\mathbf{C} = \{\exists^{\geq k} \mid k \geq 1\}$.
- (3) Counting modulo m quantifier D_m* is defined by the class $\{(A, P) \mid \text{card}(P) \equiv 0 \pmod{m}\}$. An easy first-order Ehrenfeucht-Fraïssé game argument shows that D_m is not definable in FO whenever $m > 1$. Thus

- $\text{FO}(D_m)$ is strictly more expressive than FO. For instance, the sentence $D_2x(x = x)$ of $\text{FO}(D_2)$ says that the number of points in a structure is even; it is well-known that this property is not definable in first-order logic alone.
- (4) *Majority quantifier* MAJ, which is defined by the class $\{(A, P) \mid \text{card}(P) \geq \frac{1}{2}\text{card}(A)\}$, is not definable in FO, either. For example, in $\text{FO}(\text{MAJ})$ we can say that there is a node in a graph \mathcal{A} that is connected to at least half of the nodes of \mathcal{A} : the defining sentence is $\exists y \text{MAJ}x E(y, x)$. This quantifier is also interesting in connection with capturing complexity classes, see Section 6.
 - (5) Extending FO with *Rescher* (bigger cardinality) or *Härtig* (equicardinality) quantifiers also increases the expressive power. Rescher quantifier R is defined by the class $\{(A, P, S) \mid \text{card}(P) \leq \text{card}(S)\}$ and Härtig quantifier H by the class $\{(A, P, S) \mid \text{card}(P) = \text{card}(S)\}$. Thus, for instance, given two points a and b in an undirected graph, $Rx, y(E(a, x), E(b, y))$ says that a has at most as many neighbors as b . Similarly, $Hx, y(E(a, x), E(b, y))$ says that a and b have equally many neighbors.

For each $Q_{\mathcal{K}}$, there is a natural Ehrenfeucht-Fraïssé style game-theoretic characterization for elementary equivalence in $\text{FO}(Q_{\mathcal{K}})$. We formulate the rules of this game for the counting modulo m quantifier. The rules of the game are as for first-order Ehrenfeucht-Fraïssé game except that now the spoiler may also choose a subset from one of the structures, say $X \subseteq A$. The duplicator has to respond by choosing a subset of the other structure, $Y \subseteq B$, which has modulo m equal cardinality to the spoiler's choice, that is, $\text{card}(Y) \equiv \text{card}(X) \pmod{m}$. The spoiler then challenges the duplicator's choice by selecting a point from the duplicator's structure, $b \in B$, and the duplicator has to choose a point from the other structure, $a \in A$, such that $a \in X$ if and only if $b \in Y$. Again, if a_1, \dots, a_n are the points chosen from \mathcal{A} and b_1, \dots, b_n are the points chosen from \mathcal{B} during n rounds, the duplicator wins if and only if $\{(a_i, b_i) \mid 1 \leq i \leq n\}$ is a partial isomorphism $\mathcal{A} \rightarrow \mathcal{B}$. We call this game the *counting modulo m Ehrenfeucht-Fraïssé game*.

Note that the first-order Ehrenfeucht-Fraïssé game can be seen as a special case of the counting modulo m Ehrenfeucht-Fraïssé game: the spoiler can choose the empty subset from one of the structures, and, in order to win, the duplicator has to respond with the empty set. The spoiler and the duplicator then choose their points from the complement of the empty set, that is, without any restrictions, just as in the regular Ehrenfeucht-Fraïssé game.

The following theorem shows that the counting modulo m Ehrenfeucht-Fraïssé game indeed gives us the game-theoretic characterization we were looking for. We use the notation $\mathcal{A} \equiv_{\text{FO}(D_m)}^n \mathcal{B}$ when the duplicator has a winning strategy in the n -round counting modulo m Ehrenfeucht-Fraïssé game.

Theorem 2 (see [31]). *Let $\mathcal{A}, \mathcal{B} \in \text{STRUCT}[\sigma]$. Then $\mathcal{A} \equiv_{\text{FO}(D_m)}^n \mathcal{B}$ if and only if \mathcal{A} and \mathcal{B} agree on all $\text{FO}(D_m)$ sentences of quantifier rank up to n . \square*

Corollary 2. *A class $\mathcal{C} \subseteq \text{STRUCT}[\sigma]$ is not definable in $\text{FO}(D_m)$ if and only if for every n there are σ -structures $\mathcal{A} \in \mathcal{C}$ and $\mathcal{B} \notin \mathcal{C}$ such that the duplicator has a winning strategy in the n -round counting modulo m Ehrenfeucht-Fraïssé game on \mathcal{A} and \mathcal{B} .*

The intuition behind this theorem is that if all subsets of \mathcal{A} and \mathcal{B} that the spoiler and the duplicator can use in the game look similar, and there are modulo m equally many of them, then $\text{FO}(D_m)$ can distinguish between \mathcal{A} and \mathcal{B} no more than FO can. In the next section we give a precise formulation for this intuition.

Theorem 2 can be used to show limits of expressive bounds of counting modulo quantifiers. For instance, we can show that connectivity of finite graphs is not definable in $\text{FO}(D_m)$, for any m . The construction is similar to the first-order case, but we also have to require that there are modulo m equally many points in both structures. Then we can proceed very much like in the proof for first-order. Similarly, we can show that the majority quantifier MAJ (or Rescher or Härtig quantifiers) cannot be defined in $\text{FO}(D_m)$. To see this, we can, for instance, take, for each n , $\mathcal{A} = (A, U)$ to be a set with a unary relation U such that $\text{card}(A) = 3nm$ and $\text{card}(U) = 2nm$. Take $\mathcal{B} = (B, U)$ which satisfies $\text{card}(B) = 3nm$ but $\text{card}(U) = nm$. It is not difficult to show that the duplicator has a winning strategy in the n -round counting modulo m Ehrenfeucht-Fraïssé game over \mathcal{A} and \mathcal{B} . But obviously $\mathcal{A} \models \text{MAJ}xU(x)$ and $\mathcal{B} \not\models \text{MAJ}xU(x)$. These results can be extended to the ordered case, see [41].

If we want to give a game-theoretic method to prove expressive bounds for first-order logic with *all* unary quantifiers, different techniques must be used. The method we employ here is based on *bijjective Ehrenfeucht-Fraïssé* games. The rules of the game are the following. As before, the players are the spoiler and the duplicator. There are two σ -structures \mathcal{A} and \mathcal{B} and the number of rounds, say n , given. In each round i the duplicator first selects a bijection $f_i : A \rightarrow B$ (if $\text{card}(A) \neq \text{card}(B)$, then the duplicator loses), and then the spoiler selects a point $a_i \in A$. The duplicator has to select the point $f_i(a_i)$ from B . This continues for n rounds. After the last round, the duplicator is declared the winner if and only if $\{(a_i, f_i(a_i)) \mid 1 \leq i \leq n\}$ is a partial isomorphism from \mathcal{A} to \mathcal{B} . We use the notation $\mathcal{A} \equiv_{bij}^n \mathcal{B}$ if the duplicator has a winning strategy in the n -move bijjective game on \mathcal{A} and \mathcal{B} .

It turns out that this game characterizes elementary equivalence in a logic that is stronger than $\text{FO}(\mathbf{Q}_u)$. This logic is obtained from $\text{FO}(\mathbf{Q}_u)$ by allowing infinite disjunctions and conjunctions, but by keeping quantifier rank bounded. More precisely, let $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^k$ be the extension of $\text{FO}(\mathbf{Q}_u)$ where infinite disjunctions and conjunctions are allowed but quantifier rank of each formula is at most k . The union of all these logics $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^k$ over all natural numbers $k < \omega$ is denoted by $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ (that is, the depth of nesting of quantifiers in each formula is finite). Methods used in [23] give us the following result (a proof can be found in [24]).

Theorem 3 (see [23, 24]). *Let $\mathcal{A}, \mathcal{B} \in \text{STRUCT}[\sigma]$. Then $\mathcal{A} \equiv_{bij}^n \mathcal{B}$ if and only if \mathcal{A} and \mathcal{B} agree on all $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ sentences of quantifier rank up to n . \square*

Corollary 3. *A class $\mathcal{C} \subseteq \text{STRUCT}[\sigma]$ is not definable in $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ if and only if for every n there are $\mathcal{A} \in \mathcal{C}$ and $\mathcal{B} \notin \mathcal{C}$ such that the duplicator has a winning strategy in the n -round bijective Ehrenfeucht-Fraïssé game on \mathcal{A} and \mathcal{B} . \square*

Note that the expressive power of $\text{FO}(\mathbf{Q}_u)$ is strictly weaker than that of $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$. In $\text{FO}(\mathbf{Q}_u)$, it is not possible to express the second vectorization of Härtig quantifier (that is, the equicardinality quantifier for pairs) [39], while techniques used in [31] show that every vectorization of every unary quantifier can be defined in $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$. It also follows from [31] that $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ is as strong in expressive power as $\mathcal{L}_{\infty\omega}(\mathbf{C})^\omega$. (In [31], this was shown for finite-variable logics, but the same proof technique works for $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ [23].)

Although $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ has strong counting power, the game characterization can be used to show that connectivity and transitive closure are not definable in it. The idea is the same as before: For each n , we can take \mathcal{A} to be a chordless cycle of length 2^{n+1} and \mathcal{B} a disjoint union of two chordless cycles of length 2^n . Now the duplicator's strategy is to choose in round j a bijection that preserves the distances up to 2^{n-j} between the points next to the already chosen ones. A combinatorial proof can be given that shows the existence of such a strategy [27]; however, we shall see a much easier way to establish this in the next section.

We remark that bijective games are useless in the presence of a linear order. In order to win, the duplicator has to follow the linear order when he chooses bijections (otherwise the spoiler wins in the next two rounds). Thus, there is essentially only one bijection the duplicator can choose. This gives us an example of difficulties that arise when one attempts to prove expressivity bounds in the ordered setting. In the subsequent sections, we shall face similar problems several times.

Another interesting counting logic is $\text{FO} + \text{COUNT}$, cf. [15, 30]. This is a two-sorted logic with the second sort being the sort of natural numbers. More precisely, in this approach a structure is of the form

$$\mathcal{A} = (\{v_1, \dots, v_n\}, \{1, \dots, n\}, \overline{R}_1, \dots, \overline{R}_l, \leq, BIT, 1, \max).$$

Here relations R_i apply to the non-numerical domain $\{v_1, \dots, v_n\}$, while the linear order \leq , the BIT predicate and the constants 1 and \max (interpreted as 1 and n) refer to the numerical domain $\{1, \dots, n\}$. Here $BIT(i, j)$ holds if and only if the i th bit in the binary representation of j is one. These two disjoint domains are connected by allowing formulas of the form $\exists ix\varphi(x)$ with the semantics that at least i elements satisfy φ , i.e., $\text{card}(\{a \in A \mid \mathcal{A} \models \varphi(a)\}) \geq i$. Here i refers to the numerical domain and x refers to the non-numerical domain; the quantifier $\exists ix$ binds x but not i . As an example, consider the sentence

$$\exists i \exists j [(j + j = i) \wedge \exists ix\varphi(x) \wedge \forall k (\exists kx\varphi(x) \rightarrow k \leq i)].$$

This sentence tests if the cardinality of $\{a \mid \varphi(a)\}$ is even. Indeed, $\exists ix\varphi(x) \wedge \forall k(\exists kx\varphi(x) \rightarrow k \leq i)$ holds iff exactly i elements satisfy φ , and i is even since it is of the form $2j$. Note that we used the fact that $+$ and $*$ are definable as ternary predicates in the presence of the *BIT* relation, cf. [15, 28].

Remark 1. L. Hella made the following observation. While first-order logic extended by Rescher quantifier R is in general strictly weaker than $\text{FO} + \text{COUNT}$, in the presence of a built-out linear order, $\text{FO}(R)$ and $\text{FO} + \text{COUNT}$ have the same expressive power. Clearly, $Rx, y(\varphi(x), \psi(y))$ can be written in $\text{FO} + \text{COUNT}$. On the other hand, $\exists ix\varphi(x)$ can be expressed by $Rj, x(j < i, \varphi(x))$.

A game-theoretic characterization for elementary equivalence in $\text{FO} + \text{COUNT}$ was introduced in [30] and used subsequently in [14]. However, we do not go into detail here, mainly due to the fact that the counting games of [30] are subsumed by the bijective games. The logic $\text{FO} + \text{COUNT}$ has a number of applications in computer science, in particular, in complexity theory. This will be discussed in Section 6.

Finally, we refer the reader to [31, 48] for a more detailed overview of results on generalized quantifiers in finite-model theory.

4 Gaifman's and Hanf's conditions

The game-theoretic characterizations for elementary equivalence of logics considered in the previous section gave us a vague intuition that these logics can only express local properties. This intuition will be formalized in this section. We review theorems by Gaifman and Hanf, and their modifications.

4.1 Gaifman's theorem

We start with Gaifman's theorem [18]. Let \mathcal{A} be a σ -structure. Recall that the distance $d(a, b)$ is the length of the shortest path from a to b in the Gaifman graph of \mathcal{A} , and $S_r(\mathbf{a}) = \{b \mid d(b, a) \leq r, a \in \mathbf{a}\}$. For each fixed k , there are first-order formulas that define the relations $d(a, b) > k$, $d(a, b) = k$ and $d(a, b) < k$ (see [13]). Hence, bounded quantifications of the form $\forall x \in S_k(\mathbf{y})$ and $\exists x \in S_k(\mathbf{y})$ are expressible in first-order logic for every fixed k . A formula $\varphi^{(r)}(\mathbf{y})$ is called *r-local around \mathbf{y}* if every quantifier in it is of the form $\forall x \in S_k(\mathbf{y})$ or $\exists x \in S_k(\mathbf{y})$ with $k \leq r$. A sentence ψ is called *basic r-local* if it is of the form

$$\exists x_1 \dots \exists x_m \left(\bigwedge_{1 \leq i \leq m} \varphi^{(r)}(x_i) \wedge \bigwedge_{1 \leq i < j \leq m} d(x_i, x_j) > 2r \right)$$

where $\varphi^{(r)}(x)$ is an r -local formula around x .

Theorem 4 (Gaifman [18]). *Every first-order sentence is equivalent to a Boolean combination of basic r -local sentences, and every first-order formula $\varphi(x_1, \dots, x_n)$ is equivalent to a Boolean combination of t -local formulas around x_1, \dots, x_n and basic r -local sentences. Furthermore, $r \leq 7^{qr(\varphi)-1}$, $t \leq (7^{qr(\varphi)-1})/2$ and $m \leq n + qr(\varphi)$. \square*

Gaifman's theorem tells us that every first-order formula can see only a bounded number of local neighborhoods in a structure, i.e., only a small part of the input. This is indeed a formalization of the informal statement that first-order logic can express only local properties.




Fig. 1. Formula $\varphi(x, y)$ cannot distinguish (a, b) from (b, a) .

Example 2. We show that first-order logic cannot express the transitive closure of a directed graph. Assume, to the contrary, that there is a first-order formula $\varphi(x, y)$ in the language of a single binary relation E , such that $\varphi[\mathcal{A}] = (A, \{(a, b) \in A \mid \mathcal{A} \models \varphi(a, b)\})$ is the transitive closure of \mathcal{A} . Apply Gaifman's theorem to it, and find t and r . Now consider the graph shown in Figure 1. It is a successor relation, on which we select two points, a and b . Assume that $d(a, b) > 2t$, and the distances from a and b to the start and the end node of the graph are at least $t + 1$. Then the t -neighborhoods of (a, b) and (b, a) are isomorphic, and by Gaifman's theorem, φ cannot distinguish (a, b) from (b, a) . Thus, φ cannot define the transitive closure, since $(a, b) \in TRCL(\mathcal{A})$, but $(b, a) \notin TRCL(\mathcal{A})$.

4.2 Hanf's theorem and its modifications

While Gaifman's theorem helps prove expressivity bounds for FO directly, without resorting to establishing a winning strategy for the duplicator¹, Hanf's theorem [22] and its numerous modifications [17, 24, 40, 41] provide criteria for the existence of a strategy for the duplicator that is based on counting of small neighborhoods in two structures.

Hanf's theorem was originally proved for infinite structures. It was observed by Fagin, Stockmeyer and Vardi [17] that the technique can be modified to be

¹ New winning conditions for the duplicator based on Gaifman's theorem were presented recently in [47].

applicable to finite structures. The extensions of Hanf's technique [24, 40, 41] follow the ideas of [17].

Let \mathcal{A} be a σ -structure and $a \in A$. Recall that the isomorphism type of $N_d^{\mathcal{A}}(a)$ is denoted by $tp_d^{\mathcal{A}}(a)$. Let τ be an isomorphism type of a σ_1 -structure (σ extended with one constant). We denote the number of points $a \in A$ whose d -neighborhoods realize τ by $n_d(\mathcal{A}, \tau)$. That is,

$$n_d(\mathcal{A}, \tau) = \text{card}(\{a \in A \mid tp_d^{\mathcal{A}}(a) = \tau\}).$$

For example, if \mathcal{A} is a chordless undirected cycle of length at least $2d + 2$ then there is only one isomorphism type τ of a d -neighborhood of a point occurring in \mathcal{A} : the chordless path of length $2d + 1$. In this case $n_d(\mathcal{A}, \tau) = \text{card}(A)$.

We call structures \mathcal{A} and \mathcal{B} (d, m) -equivalent if for every isomorphism type τ they have exactly the same number of points whose d -neighborhoods realize τ , or both structures have at least m such points, that is,

$$\min(n_d(\mathcal{A}, \tau), m) = \min(n_d(\mathcal{B}, \tau), m).$$

The modification of Hanf's theorem for the finite case is the following.

Theorem 5 ([17]). *Let n and f be positive integers. There are positive integers d and m such that whenever \mathcal{A} and \mathcal{B} are (d, m) -equivalent structures where every point has degree at most f , then $\mathcal{A} \equiv_{\text{FO}}^n \mathcal{B}$, that is, \mathcal{A} and \mathcal{B} satisfy the same sentences of FO of quantifier rank up to n . \square*

Note that since we consider finite models, for any \mathcal{A} there is a number f that exceeds degrees of all points of \mathcal{A} . This leads to the following. We say that \mathcal{A} and \mathcal{B} are d -equivalent, written as $\mathcal{A} \simeq_d \mathcal{B}$, if for every type of a d -neighborhood of a point, τ , equally many points realize it in \mathcal{A} and \mathcal{B} . That is,

$$\mathcal{A} \simeq_d \mathcal{B} \quad \text{iff} \quad n_d(\mathcal{A}, \tau) = n_d(\mathcal{B}, \tau) \quad \text{for every } \tau.$$

Corollary 4 ([17]). *Let n be a positive integer. There there exists a positive integer d such that $\mathcal{A} \simeq_d \mathcal{B}$ implies $\mathcal{A} \equiv_{\text{FO}}^n \mathcal{B}$. \square*

This result makes precise the intuition that counting power of first-order logic is rather limited. It also shows that only local neighborhoods are relevant for elementary equivalence in FO. Most importantly, the result above yields much simpler proofs of expressivity bounds than those based on games. Below we give a canonical example of applicability of Hanf's technique.

Example 3. We show that connectivity of finite graphs is not definable in FO. Assume, to the contrary, that it is definable by a FO sentence Φ of quantifier rank n . Apply Corollary 4 to find $d > 0$ such that $\mathcal{A} \simeq_d \mathcal{B}$ would imply $\mathcal{A} \models \Phi$ iff $\mathcal{B} \models \Phi$. Now let \mathcal{A} be a (chordless) cycle which has length $4d + 4$, and let \mathcal{B} be a disjoint union of two chordless cycles of length $2d + 2$, see Figure 2.




Fig. 2. Hanf's technique proves that connectivity is not first-order

As we noticed before, there is only one type of a d -neighborhood that these graphs realize, namely a chain on $2d + 1$ points. Thus, $\mathcal{A} \equiv_d \mathcal{B}$, since they have the same number of vertices. At the same time, \mathcal{A} is connected, but \mathcal{B} is not, proving that connectivity cannot be expressed by Φ .

Other examples, proved previously with games, can be shown to be derivable from Hanf's technique. For example, with the concept of (d, m) -equivalence, it is now easy to show that the majority quantifier (or Rescher and Härtig quantifiers) are not definable in first-order logic. \square

Before we outline the proof of Theorem 5 in [17], we state Hanf's theorem from [22]. It says that two arbitrary structures \mathcal{A} and \mathcal{B} are elementary equivalent if $S_r^{\mathcal{A}}(a)$ and $S_r^{\mathcal{B}}(b)$ are finite, for every finite r and every $a \in A$ and $b \in B$, and, for each $r > 0$ and each type τ of an r -neighborhood of a point, either $n_r(\mathcal{A}, \tau) = n_r(\mathcal{B}, \tau) < \omega$, or both \mathcal{A} and \mathcal{B} have infinitely many realizers of type τ .

Now we explain how Theorem 5 is proved. First, d is taken to be $\frac{3^n - 1}{2}$, and m is taken to exceed the size of any d -neighborhood of a point in a structure whose degrees are bounded by f . Define $d_j = \frac{3^{n+1-j} - 1}{2}$ for $j \leq n$; in particular, $d_j = 3 \cdot d_{j+1} + 1$, $d_1 = d$ and $d_n = 1$. It is then shown that the duplicator can play in the first-order Ehrenfeucht-Fraïssé game in such a way that after each round $j \leq n$, if $\mathbf{a} \in A^j$ and $\mathbf{b} \in B^j$ are the points chosen so far, then $tp_{d_j}^{\mathcal{A}}(\mathbf{a}) = tp_{d_j}^{\mathcal{B}}(\mathbf{b})$. This suffices, since $tp_1^{\mathcal{A}}(a_1, \dots, a_n) = tp_1^{\mathcal{B}}(b_1, \dots, b_n)$ implies that the mapping $a_i \mapsto b_i$ is a partial isomorphism. The condition that \mathcal{A} and \mathcal{B} are (d, m) -equivalent allows the duplicator to make the first move by choosing a point with a given type of its d -neighborhood. Suppose that the duplicator

maintained the condition successfully for the first j rounds. That is, we have $tp_{d_j}^A(\mathbf{a}) = tp_{d_j}^B(\mathbf{b})$. Assume that $a \in A$ is spoiler's next choice in the game. If $a \in S_{2d_{j+1}+1}^A(\mathbf{a})$, then the duplicator selects the point $b \in S_{2d_{j+1}+1}^B$ which is given by the isomorphism between $N_{d_j}^A(\mathbf{a})$ and $N_{d_j}^B(\mathbf{b})$. Then one easily checks that $tp_{d_{j+1}}^A(\mathbf{aa}) = tp_{d_{j+1}}^B(\mathbf{bb})$. If $a \notin S_{2d_{j+1}+1}^A(\mathbf{a})$ then the duplicator can choose any $b \notin S_{2d_{j+1}+1}^B(\mathbf{b})$ such that $tp_{d_{j+1}}^A(a) = tp_{d_{j+1}}^B(b)$. If m is chosen as above, such a point b always exists. Then again we can see that $tp_{d_{j+1}}^A(\mathbf{aa}) = tp_{d_{j+1}}^B(\mathbf{bb})$ holds, since no points in $S_{d_{j+1}}^A(\mathbf{a})$ and $S_{d_{j+1}}^B(\mathbf{b})$ belong to the same tuple of an A -relation, and likewise for B .

Having established Hanf's condition for FO, we turn to counting extensions of FO. In [41], the proof of Theorem 5 was modified for $\text{FO}(D_k)$, $k > 1$. Two structures \mathcal{A} and \mathcal{B} are called (d, m, D_k) -equivalent, if for each isomorphism type τ they have the same number of points whose d -neighborhoods realize τ , or in both structures there are at least m such points but modulo k equally many, that is,

$$\min(n_d(\mathcal{A}, \tau), m) = \min(n_d(\mathcal{B}, \tau), m) \quad \text{and} \quad n_d(\mathcal{A}, \tau) \equiv n_d(\mathcal{B}, \tau) \pmod{k}.$$

Note that (d, m, D_1) -equivalence is just the (d, m) -equivalence. Thus, the following result extends Theorem 5.

Theorem 6 ([41]). *Let n, f and k be positive integers. There are positive integers d and m such that whenever \mathcal{A} and \mathcal{B} are (d, m, D_k) -equivalent structures where every point has degree at most f , then $\mathcal{A} \equiv_{\text{FO}(D_k)}^n \mathcal{B}$. That is, \mathcal{A} and \mathcal{B} satisfy the same sentences of $\text{FO}(D_k)$ of quantifier rank up to n . \square*

The intuition behind this theorem is that the counting power of $\text{FO}(D_k)$ is rather limited, and that $\text{FO}(D_k)$ is not much stronger than FO. For instance, we can use Theorem 6 to show that connectivity is not definable in $\text{FO}(D_k)$ for any k , by using the same argument as in Example 3.

The proof of Theorem 6 in [41] follows the same idea than the proof of Theorem 5. One could ask whether this proof technique can be easily modified to prove elementary equivalence for $\text{FO}(Q)$, where Q is an arbitrary unary quantifier. It turns out that in this technique additivity property of the quantifier is crucial, and all unary quantifiers which satisfy this property are essentially counting modulo quantifiers. However, we now show a more general result that describes a winning condition in a *bijective* game.

Theorem 7 ([24, 40]). *Let n be a positive integer. There is a positive integer d such that $\mathcal{A} \sqsubseteq_d \mathcal{B}$ implies $\mathcal{A} \equiv_{bij}^n \mathcal{B}$. In particular, if $\mathcal{A} \sqsubseteq_d \mathcal{B}$, then \mathcal{A} and \mathcal{B} satisfy the same sentences of $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)$ of quantifier rank up to n . \square*

Before we explain how this result is proved, we give the following alternative definition of d -equivalence. Two structures \mathcal{A} and \mathcal{B} are d -equivalent if there

exists a bijection $f : A \rightarrow B$ such that for every $a \in A$,

$$tp_d^A(a) = tp_d^B(f(a)).$$

Now the proof of Theorem 7 is very similar to the proof of Hanf's theorem for first-order logic. We again let d_j be $\frac{3^{n+1-j}-1}{2}$, and take d to be d_1 . The duplicator's strategy is to play so that after each round j in the bijective Ehrenfeucht-Fraïssé game we have $tp_{d_j}^A(\mathbf{a}) = tp_{d_j}^B(\mathbf{b})$, if $\mathbf{a} \in A^j$ and $\mathbf{b} \in B^j$ are chosen during those j rounds. The first round bijection is given by the reformulation of d -equivalence above. Assume j rounds have been played, and we have $tp_{d_j}^A(\mathbf{a}) = tp_{d_j}^B(\mathbf{b})$, where $d_j = 3 \cdot d_{j+1} + 1$. Assume that $tp_{d_{j+1}}^A(a) = \tau$, for $a \in S_{2d_{j+1}+1}^A(\mathbf{a})$. Then $tp_{d_{j+1}}^B(h(a)) = \tau$, and $h(a) \in S_{2d_{j+1}+1}^B(\mathbf{b})$, where $h : N_{d_j}^A(\mathbf{a}) \rightarrow N_{d_j}^B(\mathbf{b})$ is an isomorphism. Thus, the number of points realizing τ is the same in complements of $S_{2d_{j+1}+1}^A(\mathbf{a})$ and $S_{2d_{j+1}+1}^B(\mathbf{b})$, and hence we have a bijection $g : A - S_{2d_{j+1}+1}^A(\mathbf{a}) \rightarrow B - S_{2d_{j+1}+1}^B(\mathbf{b})$ with the property $tp_{d_{j+1}}^A(x) = tp_{d_{j+1}}^B(g(x))$. We now define a bijection f_{j+1} for the round $j+1$ to be h on $S_{2d_{j+1}+1}^A(\mathbf{a})$, and g on its complement. It is routine to verify that f_{j+1} is a bijection, and that $tp_{d_{j+1}}^A(\mathbf{a}a) = tp_{d_{j+1}}^B(\mathbf{b}f_{j+1}(a))$ for every $a \in A$. After the last round we have vectors with isomorphic 1-neighborhoods; hence they define a partial isomorphism.

Theorem 7 shows the limits of expressive power provided by all unary quantifiers. It also significantly simplifies proofs of expressivity bounds, as applying bijective games is usually not a very easy task. For instance, we can use Example 3 to show that connectivity of finite graphs is not definable in $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$, thus avoiding all the tedious combinatorics involved in an argument based on bijective games.

As another example, we consider classes of undirected graphs which are *closed under stretching*. More precisely, let \mathcal{A} be a graph and let \mathcal{A}_d be the graph obtained by replacing every edge of \mathcal{A} by a path of length $2d+1$. Then each d -neighborhood of a point in \mathcal{A}_d contains at most one node whose degree is greater than two. We say that a class \mathcal{C} of graphs is closed under stretching if $\mathcal{A}_d \in \mathcal{C}$ for every $\mathcal{A} \in \mathcal{C}$ and for every positive integer d . Now it is easy to see that if \mathcal{A} and \mathcal{B} are 1-equivalent graphs, then \mathcal{A}_d and \mathcal{B}_d are d -equivalent. In other words, if there are the same number of points in both graphs \mathcal{A} and \mathcal{B} of each degree, then \mathcal{A}_d and \mathcal{B}_d are d -equivalent. It can then be proved that if a class \mathcal{C} and its complement are closed under stretching, and there are $\mathcal{A} \in \mathcal{C}$ and $\mathcal{B} \notin \mathcal{C}$ as above, then \mathcal{C} is not definable in $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ [26]. This argument shows that graph properties such as planarity and 3-colorability are not definable in $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$.

Hanf's technique was also used by Etessami [15] (although a preliminary conference version [14] had a proof based on counting games of [30]). It was shown in [15] that a linear order cannot be defined in FO + COUNT from its underlying successor relation. The proof relies on Hanf's technique for FO + COUNT. Its applicability follows from Theorem 7, since every FO + COUNT sentence is definable in $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$.

Summing up, the combinatorial arguments in this section allow us to simplify many proofs that were originally shown by using difficult game-theoretic arguments. Furthermore, we can also prove, often quite easily, new nondefinability results. Note, however, that in the presence of a linear ordering, none of these results can be applied to derive inexpressibility results. This is because in the presence of an order, every point is a neighbor of any other point, and thus $S_1^A(a)$ contains all elements of \mathcal{A} . Many results proving expressive bounds of these logics on ordered structures still use games (see [13] for first-order case and [41] for $\text{FO}(D_m)$). However, bijective Ehrenfeucht-Fraïssé games cannot be applied at all in the ordered case. In fact, this is closely connected to some deep problems in circuit complexity, as will be explained in Section 6.

5 Abstract notions of locality

On the surface, Gaifman’s theorem and various forms of Hanf’s theorem appear to be quite unrelated. Nevertheless, we shall see soon that there is a very close relationship between these results. To make statements like this, we study abstract concepts behind locality theorems. This approach can be seen as the study of the essential ideas behind the proofs of locality theorems, rather than using the statements of the theorems for proving expressivity bounds. From these theorems, we extract abstract notions of locality, and show how they are related. We also discuss a new form of locality, the bounded degree property, and show how it is related to other forms. This property turns out to be particularly simple to use in proving expressivity bounds.

5.1 Gaifman’s locality

We start by analyzing Gaifman’s theorem. This theorem says that only local neighborhoods are important for elementary equivalence in first-order logic. This is captured by the following definition.

Definition 1 ([32, 24]). *A formula $\psi(x_1, \dots, x_m)$ is Gaifman-local if there exists $r > 0$ such that for every $\mathcal{A} \in \text{STRUCT}[\sigma]$ and for every two m -ary vectors $\mathbf{a}, \mathbf{b} \in A^m$,*

$$tp_r^{\mathcal{A}}(\mathbf{a}) = tp_r^{\mathcal{A}}(\mathbf{b}) \quad \text{implies} \quad \mathcal{A} \models \psi(\mathbf{a}) \text{ if and only if } \mathcal{A} \models \psi(\mathbf{b}).$$

The minimum r for which this holds is called the locality rank of ψ and is denoted by $\text{lr}(\psi)$.

This definition formulates that if a logic is Gaifman-local, i.e., every formula of a logic is Gaifman-local, then indeed only small parts of the input are relevant for elementary equivalence in the logic. The part of Gaifman’s theorem that deals with sentences, gives rise to the following notion.

Definition 2 ([32, 24]).

- A formula $\psi(x_1, \dots, x_m)$, $m \geq 1$, is strongly Gaifman-local if there exists $r > 0$ such that for every $\mathcal{A}, \mathcal{B} \in \text{STRUCT}[\sigma]$ and for every two m -ary vectors $\mathbf{a} \in A^m$ and $\mathbf{b} \in B^m$, $tp_r^{\mathcal{A}}(\mathbf{a}) = tp_r^{\mathcal{B}}(\mathbf{b})$ implies $\mathcal{A} \models \psi(\mathbf{a})$ if and only if $\mathcal{B} \models \psi(\mathbf{b})$.
- A sentence Ψ is strongly Gaifman-local if it is equivalent to a Boolean combination of sentences of the form $\exists \mathbf{y} \psi(\mathbf{y})$, where $\psi(\mathbf{y})$ is a strongly Gaifman-local formula.

Before going further we notice that not every first-order formula is *strongly* Gaifman local. Consider the class of directed graphs. Let $\psi(x)$ be the formula $\forall y \neg E(y, x) \wedge \exists z \forall y \neg E(z, y)$. Then $\psi(x)$ says that x has not a predecessor and there is a point which has no a successor. Suppose that ψ is strongly Gaifman-local with locality rank r . Let G_1 be a chain graph of length $r + 1$ and let G_2 be a chain graph of the same length with a loop on the end node. Denote the start node of G_i by a_i . Then $tp_r^{G_1}(a_1) = tp_r^{G_2}(a_2)$ but obviously $G_1 \models \psi(a_1)$ whereas $G_2 \not\models \psi(a_2)$.

From Gaifman's theorem, we can derive the following.

Corollary 5 ([18, 32, 24]). *Every first-order formula is Gaifman-local and every first-order sentence is strongly Gaifman-local. Moreover, for every $\psi(\mathbf{x})$, $\text{lr}(\psi) \leq (7^{r(\psi)} - 1)/2$. \square*

5.2 Hanf's locality

When we studied Hanf's condition and proofs of Theorems 5, 6 and 7, we noticed that the essential part in all the proofs was the concept of d -equivalence. We now consider this concept in more detail and define the notion of Hanf's locality. Before giving the definition of the abstract notion we need more notation and definitions.

We extend the notion of d -equivalence (see Section 4) for structures with parameters. Let $\mathcal{A}, \mathcal{B} \in \text{STRUCT}[\sigma]$ and $\mathbf{a} \in A^m$, $\mathbf{b} \in B^m$. Then $(\mathcal{A}, \mathbf{a})$ and $(\mathcal{B}, \mathbf{b})$ are d -equivalent, $(\mathcal{A}, \mathbf{a}) \simeq_d (\mathcal{B}, \mathbf{b})$, if for every isomorphism type τ ,

$$n_d((\mathcal{A}, \mathbf{a}), \tau) = n_d((\mathcal{B}, \mathbf{b}), \tau).$$

In other words, there are as many points $a \in A$ and $b \in B$ such that $tp_d^{\mathcal{A}}(aa) = tp_d^{\mathcal{B}}(bb)$; or equivalently, there is a bijection $f : A \rightarrow B$ such that $tp_d^{\mathcal{A}}(\mathbf{a}x) = tp_d^{\mathcal{B}}(\mathbf{b}f(x))$ for all $x \in A$.

Definition 3 ([32, 24]). *A formula $\psi(x_1, \dots, x_m)$ is Hanf-local if there exists a number d such that for every $\mathcal{A}, \mathcal{B} \in \text{STRUCT}[\sigma]$ and for every two m -ary vectors \mathbf{a} and \mathbf{b} of elements of \mathcal{A} and \mathcal{B} respectively,*

$(\mathcal{A}, \mathbf{a}) \sqsubseteq_d (\mathcal{B}, \mathbf{b})$ implies $\mathcal{A} \models \psi(\mathbf{a})$ if and only if $\mathcal{B} \models \psi(\mathbf{b})$.

The minimum d for which this holds is called the Hanf locality rank of ψ , and is denoted by $\text{hlf}(\psi)$.

Thus, a sentence φ is Hanf-local, if there exists a number d such that $\mathcal{A} \sqsubseteq_d \mathcal{B}$ implies $\mathcal{A} \models \varphi$ iff $\mathcal{B} \models \varphi$. From results of the previous section, the following theorem is immediate.

Theorem 8 ([32, 24]). *Every sentence φ of $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$, or FO + COUNT, or FO is Hanf-local, and $\text{hlf}(\varphi) \leq 3^{qr(\varphi)}$. \square*

This result confirms that although these logics have substantial counting power, they can only recognize properties of small neighborhoods, and cannot grasp a structure as a whole. The definitions above extract the essential notions that were crucial for proving theorems in the previous section. Now these notions allow us to *compare* different locality results.

5.3 Relationship between the notions of locality

The result below is the main technical lemma that establishes the relationship between strong Gaifman-locality, Gaifman-locality and Hanf-locality. It states that d -equivalence of structures extends to d -equivalence of structures with parameters, if large enough neighborhoods of the parameters are isomorphic.

Lemma 1 ([24]). *If $\mathcal{A} \sqsubseteq_d \mathcal{B}$ and $tp_{3d+1}^{\mathcal{A}}(\mathbf{a}) = tp_{3d+1}^{\mathcal{B}}(\mathbf{b})$, then $(\mathcal{A}, \mathbf{a}) \sqsubseteq_d (\mathcal{B}, \mathbf{b})$. \square*

Using this, we prove the following.

Theorem 9 ([32, 24]). *Every Hanf-local formula is Gaifman-local.*

Proof. Suppose $\psi(x_1, \dots, x_m)$ is a Hanf-local formula with $\text{hlf}(\psi) = d$. We show that ψ is Gaifman-local. Take any two m -vectors \mathbf{a} and \mathbf{b} of a structure \mathcal{A} such that $tp_r(\mathbf{a}) = tp_r(\mathbf{b})$, where $r = 3d + 1$. Since $\mathcal{A} \sqsubseteq_d \mathcal{A}$, by Lemma 1 we obtain $(\mathcal{A}, \mathbf{a}) \sqsubseteq_d (\mathcal{A}, \mathbf{b})$. Thus, $\mathcal{A} \models \psi(\mathbf{a})$ if and only if $\mathcal{A} \models \psi(\mathbf{b})$. Hence ψ is Gaifman-local and $\text{lr}(\psi) \leq 3d + 1$. \square

We now consider the relationship between Hanf's locality and strong Gaifman's locality. As a technical tool, we need to extend the notion of d -equivalence to tuples. The number of different m -tuples whose d -neighborhoods realize an isomorphism type τ^m of a σ_m -structure \mathcal{A} , is denoted by $n_d(\mathcal{A}, \tau^m)$. We write $\mathcal{A} \sqsubseteq_{m,d} \mathcal{B}$, if for every isomorphism type τ^m ,

$$n_d(\mathcal{A}, \tau^m) = n_d(\mathcal{B}, \tau^m).$$

Equivalently, $\mathcal{A} \simeq_{m,d} \mathcal{B}$ if and only if there is a bijection $f : A^m \rightarrow B^m$ such that $tp_d^{\mathcal{A}}(\mathbf{a}) = tp_d^{\mathcal{B}}(f(\mathbf{a}))$ for every $\mathbf{a} \in A^m$. By considering m -tuples whose components are the same, we see that for all $m > 0$, $\mathcal{A} \simeq_{m,d} \mathcal{B}$ implies $\mathcal{A} \simeq_d \mathcal{B}$.

Our crucial lemma is that r -equivalence of $(m + 1)$ -tuples can be guaranteed by d -equivalence of tuples for large enough d that depends only on r . This can be shown by considering parametrized versions of these structures and applying Lemma 1.

Proposition 1 ([32, 24]). *Let $m > 0$ and $d \geq 0$. Then $\mathcal{A} \simeq_{m,3d+1} \mathcal{B}$ implies $\mathcal{A} \simeq_{m+1,d} \mathcal{B}$. In particular, for every $r > 0$ and $m \geq 1$ there is d such that $\mathcal{A} \simeq_d \mathcal{B}$ implies $\mathcal{A} \simeq_{m,r} \mathcal{B}$. \square*

This can be used to show the following.

Theorem 10 ([32, 24]). *Every strongly Gaifman-local sentence is Hanf-local.*

Consider a sentence Ψ which is equivalent to $\exists x_1 \dots \exists x_m \psi(x_1, \dots, x_m)$, where $\psi(\mathbf{x})$ is strongly Gaifman-local. Let $r > 0$ witness strong Gaifman's locality of ψ . Take d given by Proposition 1. Then $\text{hlf}(\Psi) \leq d$. Indeed, let $\mathcal{A} \simeq_d \mathcal{B}$ and $\mathcal{A} \models \Psi$. Then $\mathcal{A} \models \psi(\mathbf{a})$ for some $\mathbf{a} \in A^m$. By Proposition 1 we know that $\mathcal{A} \simeq_{m,r} \mathcal{B}$, and thus we find $\mathbf{b} \in B^m$ such that $tp_r^{\mathcal{A}}(\mathbf{a}) = tp_r^{\mathcal{B}}(\mathbf{b})$. Since ψ is strongly Gaifman-local, $\mathcal{B} \models \psi(\mathbf{b})$ and thus $\mathcal{B} \models \Psi$. Hence, $\text{hlf}(\Psi) \leq d$. \square

This implies that the two parts of Gaifman's theorem (those dealing with sentences and open formulas) are not independent. In fact, for any logic satisfying some regularity properties, strong Gaifman-locality of its sentences implies Gaifman-locality of its open formulae. See [24, 32] for details.

5.4 Bounded degree property

One of the easiest ways to prove expressivity bounds is the *bounded degree property*. It was first introduced for graph queries in studying limits of expressive power of database query languages [36]. Later it was generalized to arbitrary (finite) structures in [11]. We now review this concept, show its usefulness in proving expressivity bounds, and relate it to other notions of locality.

For a relation \overline{R}_i in \mathcal{A} , we define $\text{degree}_j(R_i, a)$ to be the number of tuples in \overline{R}_i whose j th component is a . For directed graphs, this gives us the familiar notions of in- and out-degree. The set

$$\{\text{degree}_j(R_i, a) \mid R_i \in \sigma, a \in A, j \leq p_i\}$$

of all degrees realized in \mathcal{A} is denoted by $\text{deg_set}(\mathcal{A})$. We use $\text{deg_count}(\mathcal{A})$ for $\text{card}(\text{deg_set}(\mathcal{A}))$. The class of σ -structures \mathcal{A} for which $\text{deg_set}(\mathcal{A}) \subseteq \{0, \dots, k\}$ is denoted by $\text{STRUCT}_k[\sigma]$.

Informally, a query has the bounded degree property if an upper bound on the degrees in an input structure implies an upper bound on the number of degrees realized in the output structure produced by the query. Recall that the output of $\psi(x_1, \dots, x_m)$ on \mathcal{A} , $\psi[\mathcal{A}]$, is the structure with one m -ary relation $(\mathcal{A}, \{\mathbf{a} \in A^m \mid \mathcal{A} \models \psi(\mathbf{a})\})$.

Definition 4 ([11]). *A formula $\psi(x_1, \dots, x_m)$ has the bounded degree property (BDP), if there is a function $f_\psi : \mathbb{N} \rightarrow \mathbb{N}$ such that*

$$\text{deg_count}(\psi[\mathcal{A}]) \leq f_\psi(k) \text{ for any } \mathcal{A} \in \text{STRUCT}_k[\sigma].$$

The bounded degree property is a very useful tool in proving inexpressibility results of recursive properties, i.e., for those queries that require fixpoint computation. As a simple example, we show that the transitive closure query violates the BDP. Assume that $TRCL$ does have the BDP; that is, there is a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that $\text{deg_count}(TRCL(\mathcal{A})) \leq f(k)$ if all in- and out-degrees in \mathcal{A} do not exceed k . Let $N = f(1) + 1$. Consider \mathcal{A} which is a successor relation on N points (see Figure 1). Since all in- and out-degrees in \mathcal{A} are at most 1, we get $\text{deg_count}(TRCL(\mathcal{A})) \leq f(1) < N$, but one can easily see that $\text{deg_count}(TRCL(\mathcal{A})) = N$. This contradiction shows that $TRCL$ does not have the bounded degree property. This proof also shows that *deterministic* transitive closure violates the BDP. (Deterministic transitive closure is defined just as transitive closure, except that one only considers paths where each node other than the last one has outdegree 1, see [29].)

What makes the BDP particularly interesting, is the following result.

Theorem 11 ([11]). *Every Gaifman-local formula has the bounded degree property.* \square

From results in the previous subsection, we conclude that first-order logic and various counting logics we considered have the bounded degree property. This confirms our intuition that these logics lack mechanisms for expressing recursive (fixpoint) computation.

Corollary 6 ([32, 24]). *Every Hanf-local formula has the bounded degree property. In particular, $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ formulae, FO + COUNT formulae (without free numerical variables), and FO formulae have the bounded degree property.* \square

Since deterministic transitive closure does not have the BDP, we obtain the following result.

Corollary 7 ([32, 24]). *Deterministic transitive closure is not definable in $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ nor in FO + COUNT.* \square

This follows immediately from the BDP, and avoids all the combinatorial arguments in Section 4, and especially the ones in Section 3, that are based on game-theoretic techniques.

We now give another example that shows how the BDP can be applied to prove inexpressibility results.

Example 4. A *balanced binary tree* is a (directed) binary tree in which all paths from the root to the leaves are of the same length. Can this property be tested in FO, or perhaps in more expressive logics such as $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$? We now use the BDP to give the negative answer.

Suppose that we have a sentence Φ that tests if a given graph is balanced binary tree. We next define a query $\varphi(x, y)$ as follows. It first defines a new graph, by interchanging the immediate successors of x, x' and x'' , and the immediate successors of y, y' and y'' , as shown in Figure 3 below, and then it tests if the resulting graph is a balanced binary tree. If either x or y fails to have exactly two immediate successors, then $\varphi(x, y)$ will evaluate to *false*. Assuming Φ is in the logic, so is $\varphi(x, y)$, for logics like FO and $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$.




Fig. 3. Changing successors of nodes in a balanced binary tree

We now show that $\varphi(x, y)$ violates the BDP. Assume it does have the BDP, and let $N = f_\varphi(2) + 1$. Let \mathcal{A} be a balanced binary tree where each path from the root to a leaf has length N . Since degrees in \mathcal{A} do not exceed 2, $\text{deg_count}(\varphi[\mathcal{A}]) < N$ by the BDP. We can see that $\mathcal{A} \models \varphi(a, b)$, for two nodes a, b , if and only if a and b are at the same level in \mathcal{A} . Thus, $\varphi[\mathcal{A}]$ is a disjoint union of N cliques of different sizes, and hence $\text{deg_count}(\varphi[\mathcal{A}]) = N$. This contradiction shows that φ does not have the BDP, and hence cannot be defined in $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$. Consequently, testing for balanced binary trees is not $\mathcal{L}_{\infty\omega}(\mathbf{Q}_u)^\omega$ -definable.

One may have noticed that there is a certain asymmetry in the definition of the BDP. In the assumption, we deal with $\text{deg_set}(\mathcal{A})$, but the conclusion puts a

bound on $\text{deg_count}(\varphi[\mathcal{A}])$. Can the definition be made symmetric? To formalize this, define the *strong bounded degree property* of $\varphi(\mathbf{x})$ as follows: there exists a function $f_\varphi : \mathbb{N} \rightarrow \mathbb{N}$ such that $\text{deg_count}(\varphi[\mathcal{A}]) \leq f_\varphi(\text{deg_count}(\mathcal{A}))$.

Proposition 2 ([11]). *There are first-order definable graph queries that violate the strong bounded degree property.* \square

In fact, [11] shows that even a weaker property is violated by some first-order queries. Define the *interval bounded degree property* of a query $\varphi(\mathbf{x})$ as the existence of a function $f_\varphi : \mathbb{N} \rightarrow \mathbb{N}$ such that $\text{deg_count}(\varphi[\mathcal{A}]) \leq f_\varphi(k)$ whenever $\text{deg_set}(\mathcal{A}) \subseteq \{n, n+1, \dots, n+k\}$ for some number n . Then there exist first-order definable queries on graphs that violate this property.

To summarize, we have seen four different locality conditions: strong Gaifman-locality, Hanf-locality, Gaifman-locality, and the bounded degree property. The relationship between them is shown in Figure 5.4. While Hanf-locality is closely tied to a game and relatively easy to show for FO and some of its extensions, Gaifman-locality and the bounded degree property are very easy to use in expressibility proofs. Fortunately, they are implied by Hanf-locality of a logic.




Fig. 4. The relationship between the notions of locality

6 Applications in complexity theory

Fagin’s theorem, that equates existential second-order logic and complexity class NP, started a new line of research in complexity theory. In the past 20 years, many complexity classes have been characterized in logical terms, see [13, 28] for an overview. For example, polynomial time and space can be characterized by least- and partial fixpoint logics, respectively. Essential for many characterizations is the presence of a linear order on the input. The intuition behind having an order is simulating the order in which elements of the input appear on the tape of a Turing machine. While for order-invariant queries, the exact order does not affect the output, its presence is required for a logic to simulate the computation of a machine. In fact, it remains an open problem whether there is a logic for polynomial time properties of unordered graphs, for example.

In this section, we deal with a circuit complexity class TC^0 . This class is defined via Boolean circuits. Consider a family of circuits $C = \{c_1, c_2, \dots, c_n, \dots\}$, where the circuit c_n has n inputs and one output. Given a Boolean string x , we say that C accepts x if the output of c_n on x is 1, whenever x is of length n .

The class AC^0 is defined as the class of languages accepted by circuits C where each gate is either an AND, or an OR, or a NOT gate, with AND and OR gates having unbounded fan-in (no restriction on the number of inputs). The number of gates in c_n is polynomial in n , and the depth of circuits c_n is constant. (More generally, for AC^k , the depth of c_n is allowed to be $O(\log^k n)$.) The class TC^0 is defined as AC^0 , except that majority gates MAJ are also allowed. Assume such a gate has k inputs. Then its single output is 1 iff at least $\lfloor k \rfloor + 1$ of its inputs are 1.

The class TC^0 is not an idle creation of complexity theory; in fact, it is of special importance in computer science. It characterizes the complexity of such important operations as integer multiplication, division, and sorting, and serves as a computational model for neural nets [44]. We refer the reader to survey [4] for additional information on circuit complexity.

Despite its importance, not much is known about the relationship between TC^0 and other complexity classes. We do know that $AC^0 \subset TC^0$, as the parity language (strings with even number of 1s) is in TC^0 , but not in AC^0 [8]. We also know (see [4], for example) that

$$TC^0 \subseteq NC^1 \subseteq DLOG \subseteq NLOG \subseteq NC \subseteq PTIME \subseteq NP,$$

but we do not know if *any* of the inclusions is proper! In fact, [45] showed that there is inherent difficulty in separating TC^0 from NP, at least using conventional techniques of circuit complexity. A general notion of *natural proof* was formulated in [45]; this notion subsumes most of the existing lower bound proofs. Then [45] showed that the existence of a natural proof separating these two classes would imply that no good pseudo-random number generators are computable in TC^0 . Putting it in the language of [4], it would imply that no cryptographically-secure functions can be computed in TC^0 , even though cryptographers believe that such functions do exist. As the notion of natural proof is quite different from results on logical expressibility in finite model theory, one might attempt to avoid obstacles of [45] by using a logical characterization of TC^0 . Below, we survey results in this direction.

Notice that in the definition of classes AC^0 and TC^0 , we did not say anything about the relationship between circuits $c_n \in C$ when n varies; in fact, they can compute completely “different” things for different n . However, in most applications, those circuits compute the same property, like parity. Capturing this intuition leads to the notion of *uniformity*. The weakest notion of uniformity is PTIME-uniformity, meaning that the mapping $n \mapsto c_n$ is computable in polynomial time. Similarly, one can define logspace-uniformity (see [3] for using these notions with the class TC^0). However, the most widely used notion of uniformity is DLOGTIME-uniformity. We spare the reader the more technical definition, that can be found in [5], and instead give the characterization theorem.

Theorem 12 ([5]). *DLOGTIME-uniform* $TC^0 = FO + COUNT + <$. \square

From now on, when we speak of TC^0 , we mean its DLOGTIME-uniform version; that is, $\text{FO} + \text{COUNT} + <$. The latter is the class of problems definable by $\text{FO} + \text{COUNT}$ formulae in the presence of an order relation $<$. We will mostly deal with order-independent properties. The notion of order-independence is defined as follows. Suppose we have a $\text{FO} + \text{COUNT}$ formula φ in the language of σ and $<$, and suppose \mathcal{A} is a σ -structure. Then, for any two orderings $<_1$ and $<_2$ on \mathcal{A} , and for any \mathbf{a} , $\mathcal{A}_{<_1} \models \varphi(\mathbf{a})$ iff $\mathcal{A}_{<_2} \models \varphi(\mathbf{a})$, where $\mathcal{A}_{<}$ denotes the extension of \mathcal{A} with the order $<$.

Even though we restrict our attention to order-independent properties, the mere presence of an order relation does increase the expressive power:

Proposition 3 ([6]). *There are order-independent properties definable in $\text{FO} + \text{COUNT} + <$ but not in $\text{FO} + \text{COUNT}$.* \square

The example of a separating query (not the proof!) is quite simple. Let $\sigma = (E, U)$, where E is binary and U is unary. Consider the following property: If the interpretation of E is an equivalence relation, then the number of distinct sizes of equivalence classes of E equals the cardinality of U . This query is not definable in $\text{FO} + \text{COUNT}$ [6], but, as shown in [38], can easily be defined with order, since all elements whose equivalence classes have the same size, can be canonically represented by the $<$ -minimal such element. Then one just checks if the number of those elements equals the cardinality of U .

Theorem 12 reduces the problem of separating TC^0 from classes above it to the problem of logical expressibility; for example, to show $\text{TC}^0 \neq \text{NLOG}$, it would suffice to show that transitive closure, an NLOG-complete problem, is not definable in $\text{FO} + \text{COUNT} + <$. Since locality gives us an easy proof that transitive closure is not in $\text{FO} + \text{COUNT}$, one might try to push the ideas of locality into the ordered setting.

We do not know whether the above expressivity bound on $\text{FO} + \text{COUNT} + <$ is true, although we conjecture that it is. Below, we survey some of the partial results confirming the intuition. We state the results for the NLOG-complete problem of computing the transitive closure, but they also hold for *deterministic* transitive closure, which is complete for DLOG.

Assume that instead of an order relation, we have a *successor* relation SUCC . Since it realizes only degrees 0 and 1, as an immediate consequence of the bounded degree property of $\text{FO} + \text{COUNT}$, we obtain

Corollary 8 ([14]). *Transitive closure is not definable in $\text{FO} + \text{COUNT} + \text{SUCC}$.* \square

Note that FO plus transitive closure TRCL plus successor relation capture NLOG (cf. [13, 28]); hence, $\text{FO} + \text{COUNT} + \text{SUCC} \subset \text{NLOG} = \text{FO} + \text{TRCL} + \text{SUCC}$. This result was first shown in [14], via a rather complex argument based on games of [30]. Later, using the results of [40], the journal version of [14] (see

[15]) gave a much simpler proof based on Hanf's condition. Finally, using the bounded degree property, we gave a completely elementary proof.

The use of bounded degree property allows us to substitute any auxiliary relation for SUCC, as long as its degrees are bounded by a constant. For example, we could use balanced binary trees (note that using such a structure would most certainly make a game-based proof unmanageable). The next question is: How can we lift the results for FO + COUNT from the constant world to that where degrees are allowed to depend on the size of a structure?

First such result was given in [32], and it used the notion of *moderate* degree of [17]. Let \mathcal{C} be a class of structures. Let $\text{maxdeg}_{\mathcal{C}}(n)$ denote the maximal degree of a structure in \mathcal{C} , whose cardinality is n . Then we say that \mathcal{C} is a class of relations of moderate degree if $\text{maxdeg}_{\mathcal{C}}(n) \leq \log^{o(1)} n$. That is, for some function $\delta : \mathbb{N} \rightarrow \mathbb{N}$ with $\lim_{n \rightarrow \infty} \delta(n) = 0$, we have $\text{maxdeg}_{\mathcal{C}}(n) \leq \log^{\delta(n)} n$. Combination of results from [11] and [32] led to the following.

Proposition 4 ([32]). *Transitive closure is not definable in FO + COUNT in the presence of relations of moderate degree.* \square

A linear order on an n -element set realizes n different degrees, from 0 to $n - 1$. Thus, we need to lift the results from relations of small (constant or moderate) degree to relations of large (comparable with the size of the input) degree. The concept of moderate degree was introduced in [17] to show that connectivity is not definable in monadic Σ_1^1 in the presence of those relations. Later, [46] extended this to linear orders. Thus, one may ask if a similar avenue of attack on the separation problem can be pursued in the case of FO + COUNT.

A partial result in this direction exists. Let \mathcal{O}_k stand for the class of relations which are pre-orders $\langle A, \prec \rangle$ (i.e., \prec is reflexive and transitive), and each equivalence class of the relation $x \sim y \equiv (x \prec y) \wedge (y \prec x)$ has at most k elements. In particular, \mathcal{O}_1 is the class of linear orders.

Theorem 13 ([38]). *Transitive closure is not definable in FO + COUNT + \mathcal{O}_k , for any $k > 1$.* \square

In fact, the result of [38] is even stronger as it allows relations from \mathcal{O}_2 to be of special form: those that have at most $g(n)$ elements in equivalence classes of size 2, and the rest, $n - g(n)$ elements, form a linear order. The function $g(n)$ can be chosen to be arbitrarily small, e.g., $\log \log \dots \log n$, but it cannot be bounded by a constant. The version of Theorem 13, proved in [38] for deterministic transitive closure, implies $\text{DLOG} \not\subseteq \text{FO} + \text{COUNT} + \mathcal{O}_k$ for any $k > 1$. But this still falls short of resolving the most important case of $k = 1$.

One of the main goals of pursuing this line of research was to avoid the obstacles on the path towards separating TC^0 from other classes posed by the natural proofs of [45]. At the first glance, our expressivity bounds look nothing like natural proofs, although we must admit that there was no systematic study

conducted on the relationship between classical lower bounds proofs in circuit complexity and logical expressivity bounds for $\text{FO} + \text{COUNT}$ and the likes.

However, while avoiding one kind of problems, we encountered (perhaps even created) different ones. For example, the proof of Theorem 13 is based on the following technique. One first shows that if for a class \mathcal{C} of auxiliary relations, there exists a class of graphs \mathcal{G} such that the pair $(\mathcal{C}, \mathcal{G})$ satisfies two properties, P1 and P2, then transitive closure of graphs in \mathcal{G} is not definable in $\text{FO} + \text{COUNT} + \mathcal{C}$. Then, [38] showed how to construct \mathcal{G} for \mathcal{C} being \mathcal{O}_k , with $k > 1$, thus proving inexpressibility of transitive closure in $\text{FO} + \text{COUNT} + \mathcal{O}_k$. Unfortunately, [38] also showed that if $\mathcal{C} = \mathcal{O}_1$, the class of linear orders, then there is *no* class of graphs \mathcal{G} such that $(\mathcal{C}, \mathcal{G})$ satisfies P1 and P2. That is, the approach is inherently limited for showing results in the ordered setting.

Another large obstacle to using the ideals of locality for proving the separation is a recent result:

Theorem 14 ([24]). *There exist order-invariant formulae in $\text{FO} + \text{COUNT} + <$ that do not have the bounded degree property. Consequently, order-invariant $\text{FO} + \text{COUNT} + <$ is neither Gaifman-local nor Hanf-local.* \square

Nevertheless, we believe it was a useful exercise to study the expressivity of $\text{FO} + \text{COUNT}$ with relations of large degree, viewing it as a reasonable approximation of uniform TC^0 . The bounds of [14], obtained less than 4 years ago with a great deal of effort, were shown here by a simple application of the bounded degree property. The best lower bound given by locality arguments (Theorem 13) shows that one can get very close to a linear order, and thus very close to uniform TC^0 , without achieving the power of logspace computation. In a way, results of this section can be viewed as separating “very uniform” versions of TC^0 from DLOG and classes above it.

7 Applications in database theory

The theory of database query languages is firmly grounded in finite-model theory, and also provides a major motivation for finite-model theory research. Traditional databases query languages, such as relational algebra and calculus, have precisely the power of first-order logic. An important subclass, called conjunctive queries, is simply the $\{\exists, \wedge\}$ -fragment of first-order logic. Various extensions of relational calculus, such as Datalog, Datalog with negation, and the language of *while* loops, correspond to various fixpoint extensions of first-order logic, see [1].

In the relational model, data is stored in relations. For example, a database may have two relations, for storing information about employees (called **Emp**) and departments (called **Dept**). Assume that **Emp** stores triples containing employee name, department, and salary, and **Dept** stores triples containing department name, manager, and manager’s salary. Below is an example of a database:

EName	EDept	ESalary
John	A1	50
Ann	A1	60
Jim	B2	75

DName	Manager	MSalary
A1	Bob	80
B2	Steve	85
C3	Mary	80

An example of a query is “For each employee, find his or her manager.” This can be written in first-order logic as:

$$q(\text{emp}, \text{manager}) \equiv \exists \text{dept}, \text{esal}, \text{msal} \text{Emp}(\text{emp}, \text{dept}, \text{esal}) \wedge \text{Dept}(\text{dept}, \text{manager}, \text{msal}).$$

In real life, of course, programmers do not write first-order formulae; instead they write queries in the language called SQL, which is the *lingua franca* of the commercial database world. The above query, in SQL, will look like:

```
SELECT E.EName, D.Manager
FROM Emp E, Dept D
WHERE E.Edept = D.Dname
```

The best way to read such statements is as set-theoretic comprehensions: the above becomes

$$\left\{ (EName, D.Manager) \left| \begin{array}{l} (EName, EDept, ESalary) \in \text{Emp}, \\ (DName, Manager, MSalary) \in \text{Dept}, \\ EDept = DName \end{array} \right. \right\}$$

The basis of SQL is the select-from-where statement, with the addition union and difference, and features such as view creation, which make the language compositional. This basis has precisely the power of the first-order logic. However, all practical implementations of SQL come equipped with two additional features: *arithmetic operations*, and *aggregate functions*. For example, consider the following query: “Find all departments that have more than 5 employees, together with the name of the manager and the average salary of the employees.”

In SQL, this will be written as

```
SELECT D.DName, D.Manager, AVG(E.ESalary)
FROM Emp E, Dept D
WHERE E.Edept = D.Dname
GROUPBY D.DName
HAVING COUNT(E.EName) > 5
```

There are two key new features in this query. *Grouping* is given by the clauses GROUPBY and HAVING: for each department manager, we group together all the employees in his/her department, provided there are more than five of them. The other feature is *aggregate function*: these are AVG, for computing the average salary, and COUNT, for counting the number of employees.

Let us now see why adding these features is indeed an extension of first-order logic. We consider a query that we know is inexpressible in first-order logic alone:

given a graph G , find the set of nodes x with $\text{in-deg}(x) = \text{out-deg}(x)$. We let graphs be represented as a binary relation `Edges` with two attributes, `From` and `To`. As the first step, we create two new relations, one storing nodes together with their in-degrees, and the other storing nodes with their out-degrees. Such intermediate relations are called *views* in SQL, and are created as

```
CREATE VIEW INDEGV(Node, Indeg) AS
  SELECT E.To, COUNT(E.From)
  FROM   Edges E
  GROUPBY E.To
```

```
CREATE VIEW OUTDEGV(Node, Outdeg) AS
  SELECT E.From, COUNT(E.To)
  FROM   Edges E
  GROUPBY E.From
```

Using these views, one computes the answer to the query as

```
SELECT INDEGV.Node
FROM   INDEGV, OUTDEGV
WHERE  (INDEGV.Node = OUTDEGV.Node)
      AND (INDEGV.Indeg = OUTDEGV.Outdeg)
```

By now, the reader must be convince that Härtig and Rescher quantifiers can be expressed in SQL. Thus, it is more powerful than first-order logic. The question now is:

How expressive is SQL?

More precisely, there seems to be a “folk result” saying that SQL cannot compute recursive queries, such as transitive closure. That is, it lacks a mechanism for recursive computation. The question now is: What kind of formal statement can one prove to confirm this intuition?

The approach of some textbooks is to restrict SQL to its subset which is essentially first-order logic, and use expressivity bounds for the latter. However, as we have just seen, this is not satisfactory. The difficulty with answering the question above is that there are dozens of different version of SQL (see [43] for an overview of standards and dialects), and they often support different sets of operators. For example, some versions even add the transitive closure operator. Thus, we restrict our attention to features that are common to *all* versions of SQL; that is, grouping and aggregation (cf. [1, 43]). The first result on the expressive power of such a language was based on yet unproven assumption from complexity theory, and two observations: SQL queries can be evaluated in deterministic logspace, and transitive closure is complete for nondeterministic logspace. Thus,

Proposition 5 ([10]). *Assume that $DLOGSPACE \neq NLOGSPACE$. Then transitive closure cannot be expressed in SQL.* \square

Can we get rid of the unproven assumption? The problem we face is that SQL per se is quite inconvenient to work with – its syntax is quite awkward, and in fact it has been an object of persistent criticism. SQL combines sets and multisets in order to evaluate aggregates: for example, computing the average salary, one *cannot* first project out the salary attribute and then compute its average, as the elimination of duplicates will produce an incorrect result. Following [20, 35], a language that deals correctly with multiset and set semantics was proposed in [36]. As this language can model the main features of SQL, and extends first-order logic, it was suggested to use it as a rational reconstruction of SQL.

Let us now give an informal introduction into this language, which we call here $AGGR_{\mathbb{Q}}$. It deals with objects which can freely combine rational numbers, elements of the domain of atomic values, D , tuples, and sets. In particular, it permits sets of sets. The language is *statically typed* (cf. [21, 36]), but we shall not go into detail here. The full description can be found in [21, 36]; here we highlight the salient features. The language allows one to apply a function to each element of a set, that is, obtain a set $\{f(x) \mid x \in X\}$ from X . If all elements of a set \mathcal{X} are sets themselves, their union can be taken: $\bigcup_{X \in \mathcal{X}} X$. Basic arithmetic operations ($+$, $-$, $*$, \div , $<$) are available on \mathbb{Q} . On the domain of basic values D , only equality test is available (notice the absence of order). Finally, if f is a function into \mathbb{Q} , its values on a set X can be added up; that is, one can compute $\sum_{x \in X} f(x)$. Note that if f is identically 1, then the above is the cardinality of X .

It was proved in [36] that these features model the main features of SQL. Furthermore, extending a result from [49], a conservativity property was shown in [36]. It says that nesting of sets is in a sense superfluous: every query from relational databases to relational databases in $AGGR_{\mathbb{Q}}$ can be written in a way that does not use sets of sets. Nesting is essential for modeling grouping, and thus the (nontrivial) conservativity result gives us a rather pleasant language (without higher-order features) to model *all* the features of SQL. In [36], the following was proved:

Proposition 6 ([36]). *Transitive closure cannot be expressed in $AGGR_{\mathbb{Q}}$.* \square

While this does provide useful bounds, there are two problems with the proof in [36]. First, it is very syntactic. It proceeds by establishing a normal form result for queries on a special class of inputs; the property depends both on the class of inputs, and the properties of the chosen syntax of the language. Thus, making a minor change in the syntax that does not affect expressiveness would mean that the proof must be redone from scratch. Even more unpleasantly, the proof in [36] establishes inexpressibility of transitive closure, but fails to establish a general property that will give us expressivity bounds.

In [21,36] it was conjectured that *relational queries* in $\text{AGGR}_{\mathbb{Q}}$ have the BDP. By relational queries we mean those whose inputs and outputs only contain elements of D , but no numbers, although rational numbers can be used in the process of evaluating a query. An SQL example of finding nodes with equal in- and out-degrees shows that there are relational queries definable in SQL but not in first-order. Clearly, proving the above conjecture would resolve the problem for queries such as transitive closure or deterministic transitive closure. By extending the normal form result $\text{AGGR}_{\mathbb{Q}}$, the following was proved.

Proposition 7 ([11]). *Every relational query in $\text{AGGR}_{\mathbb{Q}}$ has the bounded degree property.* \square

It is still unpleasant that the proof depends on a particular syntax for the language. Also, it would be interesting to know if relational queries are Gaifman-local. These two questions were considered in [32]. The approach taken by [32] was the following. Restrict the language to *natural numbers* only. That is, one can compute COUNT but not AVG. Let us call this language $\text{AGGR}_{\mathbb{N}}$. With rationals out of the way, try to embed it into $\text{FO} + \text{COUNT}$ to prove locality.

Embedding into $\text{FO} + \text{COUNT}$ turns out to be problematic, as in $\text{FO} + \text{COUNT}$ the sizes of the first-sort and second-sort universes are the same. In contrast, in SQL, one can create numbers much bigger than the size of the database. A simple example is this: consider an n -element set X , and compute $\sum_{x \in X} f(x)$ where $f(x) = 2$ for every x . The result is $2n$. The solution to the problem is to modify a query, essentially by putting a huge linear order “on the side” and having all arithmetic done on that linear order. This technique led to the following result.

Proposition 8 ([32]). *For any relational query Q in $\text{AGGR}_{\mathbb{N}}$, another query Q' can be found such that (a) Q is Gaifman-local iff Q' is, and (b) Q' can be defined in $\text{FO} + \text{COUNT}$. Consequently, every relational query in $\text{AGGR}_{\mathbb{N}}$ is Gaifman-local.* \square

Using this technique, [37] returned to the main question: what is the expressive power of SQL? It defined a new language, called $\text{AGGR}_{\mathbb{N}}^{\text{flat}}$, which first restricts $\text{AGGR}_{\mathbb{N}}$ to objects that do not contain sets of sets, and then adds new arithmetic operation, and *product* over a set. That is, given a set X and a function f into \mathbb{N} definable in $\text{AGGR}_{\mathbb{N}}^{\text{flat}}$, one can compute $\prod_{x \in X} f(x)$ in $\text{AGGR}_{\mathbb{N}}^{\text{flat}}$. Then the following sequence of results was proved.

Theorem 15 ([37]). *1) Every relational $\text{AGGR}_{\mathbb{Q}}$ query can be expressed in $\text{AGGR}_{\mathbb{N}}^{\text{flat}}$.*

2) For any relational query Q in $\text{AGGR}_{\mathbb{N}}^{\text{flat}}$, another query Q' in $\text{AGGR}_{\mathbb{N}}^{\text{flat}}$ can be found such that (a) Q is Gaifman-local iff Q' is, and (b) Q' can be defined in $\text{FO} + \text{COUNT}$. \square

The bulk of the proof of this theorem is in showing (b) of part 2). Combining these results with locality of $\text{FO} + \text{COUNT}$, we obtain:

Corollary 9. *Every relational query in $\text{AGGR}_{\mathbb{Q}}$ is Gaifman-local.* □

Consequently, plain SQL queries are Gaifman-local and have the bounded degree property. Therefore, queries that need a recursion mechanism (transitive closure, deterministic transitive closure, connectivity test, etc.) *cannot* be computed in SQL.

Thus, locality helped us answer important questions about expressive power of real world database query languages.

We conclude this section by a remark about the set of basic operations on the domain of atomic values, D . In all the above results, we assume that it is only possible to test if two elements of D are equal. In many applications, there is a meaningful linear order on such domains (for examples, salaries can be compared). What happens if D is ordered? As it often happens, things become a lot more complicated in the ordered setting. In a way, we saw it in the last section. For SQL, one can show that it is possible to express any query from $\text{FO} + \text{COUNT} + <$ if D is ordered. Thus, every uniform TC^0 property becomes definable in SQL, and hence the problems of expressivity in of recursive queries (such as transitive closure) hinges on the separation of complexity classes, thus confirming the original intuition of [10]. Note that there are other well-known examples of close connection between separation of complexity classes and expressivity bounds of query languages, see [2].

One can go further and add more operations to D . For example, one can assume that D is the field of real numbers with the usual operations $+$, $*$, $-$, or perhaps more complex such as e^x . The results of [7] show that for the class of *generic* [1] queries (those that commute with permutations of D), these extra operations do not add expressive power, beyond a linear order. Most examples of queries we consider – transitive closure, connectivity test, etc. – are generic. Thus, adding extra operations beyond $<$ does not lead to an increase in power of languages such as first-order logic or SQL. This assertion, in the case of an interpreted domain D , depends on what one means by quantification $\exists x$. In most application, quantification in queries assumes the finite database; however, it is conceivable (and in some applications, important) to quantify over D . We refer the reader to [7] for the discussion on this topic.

8 Conclusion

In this paper, we reviewed some results that were developed for proving lower bounds for logical expressibility. We considered first-order logic and its extensions with several kinds of counting mechanisms. We presented the usual game-theoretic characterizations of those logics, as well as Gaifman's and Hanf's theorems, and general notions behind these results. We also studied the relationship

between these notions and the bounded degree property. We reviewed applications of these notions in descriptive complexity theory and database theory.

Note. Several results on locality of logics appeared recently. As it is too late to include their detailed descriptions in this survey, we just give pointers to relevant papers. Most importantly, [19] partly extended locality techniques to the ordered setting. Namely, [19] proves that every order-invariant query definable in $\text{FO} + <$ is Gaifman-local, where the concept of neighborhood does not take into account the order relation. An elementary proof of Gaifman-locality of a very powerful counting logic was given in [33]; that paper also used a refined argument for proving locality to find tight bounds on the radii of neighborhoods required in locality theorems. Here, we typically used the radii of the order $O(3^{qr(\varphi)})$; [33] improves this to $O(2^{qr(\varphi)})$. Hanf-locality of the counting logic from [33] was proved in [25]. It was also shown in [25] that Hanf-locality remains valid under the addition of aggregate functions, which provided a transparent proof of Corollary 9. Finally, [34] improved Theorem 13 by showing that invariant queries in $\text{FO} + \text{COUNT} + \mathcal{O}_k$, $k > 1$, have the bounded degree property. This remains true for more powerful logics, and a more general class of auxiliary relations.

Acknowledgement We thank Jouko Väänänen for the invitation to write this survey. We are grateful to Lauri Hella and Limsoon Wong for their comments.

References

1. S. Abiteboul, R. Hull and V. Vianu. *Foundations of Databases*, Addison Wesley, 1995.
2. S. Abiteboul and V. Vianu. Computing with first-order logic. *Journal of Computer and System Sciences* 50 (1995), 309–335.
3. M. Agrawal, E. Allender and S. Datta. On TC^0 , AC^0 , and arithmetic circuits. In *Proc. 12th IEEE Conf. on Computational Complexity*, 1997.
4. E. Allender. Circuit complexity before the dawn of the new millennium. In *Proc. 16th Conf. on Foundations of Software Technology and Theoretical Computer Science (FST&TCS'96)*, Springer LNCS vol. 1180, 1996, 1–18.
5. D.A. Barrington, N. Immerman and H. Straubing. On uniformity within NC^1 . *J. Comput. and Syst. Sci.*, 41:274–306, 1990.
6. M. Benedikt, H.J. Keisler. On expressive power of unary counters. *Proc. Intl. Conf. on Database Theory (ICDT'97)*, Springer LNCS 1186, January 1997.
7. M. Benedikt and L. Libkin. Relational queries over interpreted structures. *J. ACM*, to appear. Extended abstract in *PODS'97*, pages 87–98.
8. R.B. Boppana and M. Sipser. The Complexity of Finite Functions. In *“Handbook of Theoretical Computer Science,”* Volume A, Chapter 14, pages 759–804, North Holland, 1990.
9. J. Cai, M. Fürer and N. Immerman. On optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12 (1992), 389–410.
10. M. Consens and A. Mendelzon. Low complexity aggregation in GraphLog and Datalog, *Theoretical Computer Science* 116 (1993), 95–116. Extended abstract in *ICDT'90*.

11. G. Dong, L. Libkin and L. Wong. Local properties of query languages. *Theoretical Computer Science*, to appear. Extended abstract in *ICDT'97*, LNCS vol. 1186, pages 140–154.
12. H.-D. Ebbinghaus. Extended logics: the general framework. In J. Barwise and S. Feferman, editors, *Model-Theoretic Logics*, Springer-Verlag, 1985, pages 25–76.
13. H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer Verlag, 1995.
14. K. Etessami. Counting quantifiers, successor relations, and logarithmic space, In *Proc. Structure in Complexity Theory*, 1995.
15. K. Etessami. Counting quantifiers, successor relations, and logarithmic space, *Journal of Computer and System Sciences*, 54 (1997), 400–411.
16. R. Fagin. Easier ways to win logical games. In *Proc. DIMACS Workshop on Descriptive Complexity and Finite Models*, AMS 1997.
17. R. Fagin, L. Stockmeyer and M. Vardi, On monadic NP vs monadic co-NP, *Information and Computation*, 120 (1994), 78–92.
18. H. Gaifman. On local and non-local properties, *Proceedings of the Herbrand Symposium, Logic Colloquium '81*, North Holland, 1982.
19. M. Grohe and T. Schwentick. Locality of order-invariant first-order formulas. In *MFCS'98*, pages 437–445.
20. S. Grumbach and T. Milo. Towards tractable algebras for bags. *Journal of Computer and System Sciences*, 52 (1996), 570–588.
21. S. Grumbach, L. Libkin, T. Milo and L. Wong. Query languages for bags: expressive power and complexity. *SIGACT News*, 27 (1996), 30–37.
22. 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.
23. L. Hella. Logical hierarchies in PTIME. *Information and Computation*, 129 (1996), 1–19.
24. L. Hella, L. Libkin and J. Nurmonen. Notions of locality and their logical characterizations over finite models. *Journal of Symbolic Logic*, to appear.
25. L. Hella, L. Libkin, J. Nurmonen and L. Wong. Logics with aggregate operators. In *Proc. 14th IEEE Symp. on Logic in Computer Science (LICS'99)*, Trento, Italy, July 1999, pages 35–44.
26. L. Hella and J. Nurmonen. Vectorization hierarchies of some graph quantifiers. *Archive for Mathematical Logic*, to appear.
27. L. Hella and G. Sandu. Partially ordered connectives and finite graphs. In M. Krynicki, M. Mostowski and L. Szczerba, eds, *Quantifiers: Logics, Models and Computation II*, Kluwer Academic Publishers, 1995, pages 79–88.
28. N. Immerman. *Descriptive Complexity*. Springer-Verlag, 1999.
29. N. Immerman. Languages that capture complexity classes. *SIAM J. Comput.* 16 (1987), 760–778.
30. N. Immerman and E. Lander. Describing graphs: A first order approach to graph canonization. In *“Complexity Theory Retrospective”*, Springer Verlag, Berlin, 1990.
31. Ph. Kolaitis and J. Väänänen. Generalized quantifiers and pebble games on finite structures. *Annals of Pure and Applied Logic*, 74 (1995), 23–75.
32. L. Libkin. On the forms of locality over finite models. In *Proc. 12th IEEE Symp. on Logic in Computer Science (LICS'97)*, Warsaw, Poland, June–July 1997, pages 204–215.
33. L. Libkin. On counting logics and local properties. In *Proc. 13th IEEE Symp. on Logic in Computer Science (LICS'98)*, Indianapolis, June 1998, pages 501–512.
34. L. Libkin. Logics with counting, auxiliary relations, and lower bounds for invariant queries. In *Proc. 14th IEEE Symp. on Logic in Computer Science (LICS'99)*, Trento, Italy, July 1999, pages 316–325.

35. L. Libkin and L. Wong. Some properties of query languages for bags. In *Proc. Database Programming Languages 1993*, Springer Verlag, 1994, pages 97–114.
36. L. Libkin and L. Wong. Query languages for bags and aggregate functions. *Journal of Computer and System Sciences* 55 (1997), 241–272. Extended abstract in *PODS'94*, pages 155–166.
37. L. Libkin and L. Wong. On the power of aggregation in relational query languages. In *Proc. Database Programming Languages 1997*, Springer LNCS 1369, pages 260–280.
38. L. Libkin and L. Wong. Unary quantifiers, transitive closure, and relations of large degree. In *Proc. 15th Symp. on Theoretical Aspects of Computer Science (STACS'98)*, Springer LNCS 1373, pages 183–193.
39. K. Luosto. Hierarchies of monadic generalized quantifiers. *Journal of Symbolic Logic*, to appear.
40. J. Nurmonen. On winning strategies with unary quantifiers. *J. Logic and Computation*, 6 (1996), 779–798.
41. 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 (LICS'96)*, New Brunswick, NJ, July 1996, pages 484–493.
42. J. Nurmonen. Unary quantifiers and finite structures. PhD Thesis, University of Helsinki, 1996.
43. P. O'Neil. *Database*. Morgan Kaufmann, 1994.
44. I. Parberry and G. Schmitger. Parallel computation and threshold functions. *Journal of Computer and System Sciences* 36 (1988), 278–302.
45. A. Razborov and S. Rudich. Natural proofs. *Journal of Computer and System Sciences* 55 (1997), 24–35.
46. T. Schwentick. On winning Ehrenfeucht games and monadic NP. *Annals of Pure and Applied Logic*, 79 (1996), 61–92.
47. T. Schwentick and K. Barthelmann. Local normal forms for first-order logic with applications to games and automata. In *Proc. 15th Symp. on Theoretical Aspects of Computer Science (STACS'98)*, Springer LNCS 1373, pages 444–454.
48. J. Väänänen. Generalized quantifiers. *Bulletin of the European Association for Theoretical Computer Science*, 62 (1997), 115–136.
49. L. Wong. Normal forms and conservative properties for query languages over collection types. *Journal of Computer and System Sciences* 52 (1996), 495–505. Extended abstract in *PODS 93*.