



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Locality of Queries and Transformations

Citation for published version:

Libkin, L 2006, 'Locality of Queries and Transformations', *Electronic Notes in Theoretical Computer Science*, vol. 143, pp. 115-127. <https://doi.org/10.1016/j.entcs.2005.04.041>

Digital Object Identifier (DOI):

[10.1016/j.entcs.2005.04.041](https://doi.org/10.1016/j.entcs.2005.04.041)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Electronic Notes in Theoretical Computer Science

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.





Locality of Queries and Transformations

Leonid Libkin¹

*Department of Computer Science
University of Toronto
Toronto, Canada*

Abstract

Locality is a standard notion of finite model theory. There are two well known flavors of it, based on Hanf's and Gaifman's theorems. Essentially they say that structures that locally look alike cannot be distinguished by first-order sentences. Very recently these standard notions have been generalized in two ways. The first extension makes the notion of "looking alike" depend on logical indistinguishability, rather than isomorphism, of local neighborhoods. The second extension considers transformations defined by FO formulae, and requires that small neighborhoods be preserved by those transformations. In this survey we explain these new notions – as well as the standard ones – and show how they behave with respect to Hanf's and Gaifman's conditions.

Keywords: First-order logic, finite models, locality, Hanf's theorem, Gaifman's theorem, local consistency, locality under logical equivalence

1 Introduction

Locality is a property of logics that finds its origins in the work by Hanf [12] and Gaifman [10], and that was shown to be very useful in the context of finite model theory. Locality is useful for proving inexpressibility results, and for establishing normal forms for logical formulae. It has found applications in various areas of computer science including: complexity theory (e.g., the study of monadic NP and monadic coNP [8], or circuit complexity classes [5,20]), databases (e.g., expressiveness of aggregate query languages [14], and – more recently – query answering and rewriting in data exchange [6,1]), formal languages (e.g., locally threshold-testable languages [24]), algorithm design

¹ Email: libkin@cs.toronto.edu

(e.g. low-complexity model-checking algorithms [9,11,23]), automata theory (e.g., the normal form of [22]), etc.

There are two closely related ways of stating locality of logical formulae. One, originating in Hanf's work [12], says that if two structures \mathfrak{A} and \mathfrak{B} realize the same multiset of types of neighborhoods of radius d , then they agree on a given sentence Φ . Here d depends only on Φ .

The notion of locality inspired by Gaifman's theorem [10] says that if the d -neighborhoods of two tuples \bar{a}_1 and \bar{a}_2 in a structure \mathfrak{A} are isomorphic, then $\mathfrak{A} \models \varphi(\bar{a}_1) \leftrightarrow \varphi(\bar{a}_2)$. Again, d depends on φ , and not on \mathfrak{A} .

Recently these notions have been extended in two ways. The first way has to deal with the underlying assumptions on locality. Namely, the standard notions of locality refer to isomorphism of neighborhoods, which is a strong property that is often not expressible in a logic that satisfies one of the locality properties. Intuitively, it seems that instead of isomorphism, one can require that neighborhoods be indistinguishable in a logic whose locality we want to establish. This direction was pursued in [2], where it was shown that this simple intuition works for some logics of interest but fails for others.

A different direction has to do with viewing logical formulae as defining not queries (where we are simply interested in knowing whether a tuple \bar{a} is in the result) but rather *transformations*, where we are interested in a whole structure that results by applying a formula. This was largely motivated by database applications, specifically by *data exchange* [6]. Data exchange is the problem of transforming an instance of a database schema (that can be viewed as a finite relational structure) into an instance of another schema. Rules specifying such transformations are often given by means of logical formulae.

While the basic locality notions are now found in texts (see, e.g., [4,15,19]), the extensions mentioned above are quite recent. Here I survey them, concentrating on the case of first-order logic (FO).

2 Notations

In this paper we work with finite structures, whose vocabularies are relational, that is, finite sequences of relation symbols $\sigma = \langle R_1, \dots, R_l \rangle$. A σ -structure \mathfrak{A} consists of a finite universe A and an interpretation of each p_i -ary relation symbol R_i in σ as $R_i^{\mathfrak{A}} \subseteq A^{p_i}$. We adopt the convention that the universe of a structure is denoted by the corresponding Roman letter, that is, the universe of \mathfrak{A} is A , the universe of \mathfrak{B} is B , etc. Isomorphism of structures will be denoted by \cong . We shall use the notation σ_n for σ expanded with n constant symbols.

Given a structure \mathfrak{A} , its *Gaifman graph* [19,10,8] $\mathcal{G}(\mathfrak{A})$ is defined as $\langle A, E \rangle$

where (a, b) is in E iff there is a tuple $\bar{c} \in R_i^{\mathfrak{A}}$ for some i such that both a and b are in \bar{c} . The distance $d(a, b)$ is defined as the length of the shortest path from a to b in $\mathcal{G}(\mathfrak{A})$; we assume $d(a, a) = 0$. If $\bar{a} = (a_1, \dots, a_n)$, then $d(\bar{a}, b) = \min_i d(a_i, b)$. Given \bar{a} over A , its r -ball $B_r^{\mathfrak{A}}(\bar{a})$ is $\{b \in A \mid d(\bar{a}, b) \leq r\}$. If $|\bar{a}| = n$, its r -neighborhood $N_r^{\mathfrak{A}}(\bar{a})$ is defined as a σ_n structure

$$\langle B_r^{\mathfrak{A}}(\bar{a}), R_1^{\mathfrak{A}} \cap B_r^{\mathfrak{A}}(\bar{a})^{p_1}, \dots, R_l^{\mathfrak{A}} \cap B_r^{\mathfrak{A}}(\bar{a})^{p_l}, a_1, \dots, a_n \rangle.$$

That is, the carrier of $N_r^{\mathfrak{A}}(\bar{a})$ is $B_r^{\mathfrak{A}}(\bar{a})$, the interpretation of the σ -relations is inherited from \mathfrak{A} , and the n extra constants are the elements of \bar{a} .

If $N_r^{\mathfrak{A}}(\bar{a})$ and $N_r^{\mathfrak{B}}(\bar{b})$ are isomorphic, we write $N_r^{\mathfrak{A}}(\bar{a}) \cong N_r^{\mathfrak{B}}(\bar{b})$. Note that for any isomorphism $h : N_r^{\mathfrak{A}}(\bar{a}) \rightarrow N_r^{\mathfrak{B}}(\bar{b})$ it must be the case that $h(\bar{a}) = \bar{b}$.

Given a tuple $\bar{a} = (a_1, \dots, a_n)$ and an element c , we write $\bar{a}c$ for the tuple (a_1, \dots, a_n, c) .

An m -ary query is a mapping $Q : \mathfrak{A} \rightarrow A^m$ that is closed under isomorphism (that is, if $h : \mathfrak{A} \rightarrow \mathfrak{B}$ is an isomorphism, then $Q(\mathfrak{B}) = h(Q(\mathfrak{A}))$). A logical formula $\varphi(\bar{x})$, with $|\bar{x}| = m$, defines an m -ary query Q by $\bar{a} \in Q(\mathfrak{A}) \Leftrightarrow \mathfrak{A} \models \varphi(\bar{a})$. A query also gives us a transformation of structures that sends \mathfrak{A} into $\langle A, Q(\mathfrak{A}) \rangle$.

We write $\mathfrak{A} \equiv_k \mathfrak{B}$ if \mathfrak{A} and \mathfrak{B} agree on all FO sentences of quantifier rank up to k , and $(\mathfrak{A}, \bar{a}) \equiv_k (\mathfrak{B}, \bar{b})$ if $\mathfrak{A} \models \varphi(\bar{a}) \Leftrightarrow \mathfrak{B} \models \varphi(\bar{b})$ for every FO formula $\varphi(\bar{x})$ of quantifier rank up to k . It is well known (see [4,15,19]) that $\mathfrak{A} \equiv_k \mathfrak{B}$ iff the duplicator has a winning strategy in the k -round Ehrenfeucht-Fraïssé game on \mathfrak{A} and \mathfrak{B} , and $(\mathfrak{A}, \bar{a}) \equiv_k (\mathfrak{B}, \bar{b})$ iff the duplicator has a winning strategy in the k -round Ehrenfeucht-Fraïssé game on \mathfrak{A} and \mathfrak{B} starting in position (\bar{a}, \bar{b}) .

3 Locality: Hanf and Gaifman conditions

Hanf’s locality condition essentially says that if two structures realize the same multiset of neighborhood types of points, then they cannot be distinguished in a logic. It was presented in [12] for infinite structures and modified for the finite case in [8]. A slight generalization presented in [13] allows one to deal with free variables, and this is the definition we use here.

If \mathfrak{A} and \mathfrak{B} are σ -structures, and $\bar{a} \in A^n$, $\bar{b} \in B^n$, then we write

$$(\mathfrak{A}, \bar{a}) \rightleftharpoons_d (\mathfrak{B}, \bar{b})$$

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

$$N_d^{\mathfrak{A}}(\bar{a}c) \cong N_d^{\mathfrak{B}}(\bar{b}f(c)).$$




Fig. 1. Hanf-locality

In the case of $n = 0$, $\mathfrak{A} \sqsubseteq_d \mathfrak{B}$ means that $N_d^{\mathfrak{A}}(c) \cong N_d^{\mathfrak{B}}(f(c))$ for some bijection $f : A \rightarrow B$.

The \sqsubseteq_d relation says, in a sense, that locally two structures look the same, with respect to a certain bijection f ; that is, f sends each element c into $f(c)$ that has the same neighborhood.

Definition 3.1 [Hanf-locality] An m -ary query Q , for $m \geq 0$, on σ -structures is *Hanf-local* if there exists a number $d \geq 0$ such that for two every σ -structures \mathfrak{A} and \mathfrak{B} , and for $\bar{a} \in A^m, \bar{b} \in B^m$,

$$(\mathfrak{A}, \bar{a}) \sqsubseteq_d (\mathfrak{B}, \bar{b}) \text{ implies } (\bar{a} \in Q(\mathfrak{A}) \Leftrightarrow \bar{b} \in Q(\mathfrak{B})).$$

This is illustrated in Fig. 1 for Boolean (0-ary) query Q .

The “canonical” example of using Hanf-locality is proving that connectivity is not expressible in a logic that is Hanf-local. This is illustrated in Fig. 2: suppose connectivity were Hanf-local, with d witnessing Hanf-locality. Then pick two graphs: G_1 is a union of two cycles of length $d' > 2d + 1$ and G_2 is one cycle of length $2d'$. Then every d -neighborhood is a chain of length of $2d$, and thus *any* bijection $f : G_1 \rightarrow G_2$ witnesses $G_1 \sqsubseteq_d G_2$, and yet G_1 and G_2 disagree on the connectivity query.

So the question is which logics define Hanf-local queries. This was answered in [12,8]:

Theorem 3.2 *Every FO-definable query is Hanf-local.*

While Hanf-locality is a useful criterion, it is often easier to work with one structure than with two. This is achieved by using a different locality notion.

Definition 3.3 [Gaifman-locality] An m -ary query Q , for $m > 0$, on σ -structures, is called *Gaifman-local* if there exists a number $d \geq 0$ such that for every σ -structure \mathfrak{A} and every $\bar{a}_1, \bar{a}_2 \in A^m$,

$$N_d^{\mathfrak{A}}(\bar{a}_1) \cong N_d^{\mathfrak{A}}(\bar{a}_2) \text{ implies } (\bar{a}_1 \in Q(\mathfrak{A}) \Leftrightarrow \bar{a}_2 \in Q(\mathfrak{A})).$$

This notion is illustrated in Fig. 3.




Fig. 2. Connectivity is not Hanf-local




Fig. 3. Gaifman-locality




Fig. 4. A local formula cannot distinguish (a, b) from (b, a)

The “canonical” example of using Gaifman-locality is showing that any logic that defines only Gaifman-local queries is incapable of expressing the transitive closure of a graph. This is illustrated in Fig. 4: in a long enough chain, we can find two points a and b such that $N_d(a, b) \cong N_d(b, a)$ and yet the transitive closure query distinguishes these pairs.

It follows immediately from [10] that all FO-definable queries are Gaifman-local. In fact there is a close connection between the two notions of locality:

Theorem 3.4 (see [13]) *Every Hanf-local query is Gaifman-local.*

4 Locality based on logical equivalence

Intuitively, the notion of isomorphism of neighborhoods seems to be too strong an assumption for concluding that formulae of a logic cannot distinguish certain tuples, if a logic itself cannot describe neighborhoods up to isomorphism. So it seems natural to weaken the requirement that neighborhoods be isomorphic and instead require that they be indistinguishable in a logic. That is, instead of requiring $N_d(\bar{a}) \cong N_d(\bar{b})$, we now require that $N_d(\bar{a}) \equiv_k N_d(\bar{b})$ for some $k \geq 0$. Before defining these notions formally, we need a new version of

the \Leftrightarrow_d relation. We write

$$(\mathfrak{A}, \bar{a}) \Leftrightarrow_{d,k}^{\equiv} (\mathfrak{B}, \bar{b})$$

iff there is a bijection $f : A \rightarrow B$ such that $N_d^{\mathfrak{A}}(\bar{a}c) \equiv_k N_d^{\mathfrak{B}}(\bar{b}f(c))$ for all $c \in A$.

Definition 4.1 [Locality under logical equivalence] a) An m -ary query Q , $m \geq 0$, on σ -structures, is called *Hanf-local under logical equivalence* if there exists numbers $d, k \geq 0$ such that for every two σ -structures \mathfrak{A} and \mathfrak{B} , and $\bar{a} \in A^m$, $\bar{b} \in B^m$, if $(\mathfrak{A}, \bar{a}) \Leftrightarrow_{d,k}^{\equiv} (\mathfrak{B}, \bar{b})$ then $\bar{a} \in Q(\mathfrak{A})$ iff $\bar{b} \in Q(\mathfrak{B})$.

b) An m -ary query Q , $m > 0$, on σ -structures, is called *Gaifman-local under logical equivalence* if there exists numbers $d, k \geq 0$ such that for every σ -structure \mathfrak{A} and every $\bar{a}_1, \bar{a}_2 \in A^m$,

$$N_d^{\mathfrak{A}}(\bar{a}_1) \equiv_k N_d^{\mathfrak{A}}(\bar{a}_2) \quad \text{implies} \quad (\bar{a}_1 \in Q(\mathfrak{A}) \Leftrightarrow \bar{a}_2 \in Q(\mathfrak{A})).$$

What can be said about these notions? First, from [10] we get:

Theorem 4.2 *Every FO-definable query is Gaifman-local under logical equivalence.*

However, Hanf-locality is lost with this new notion, as noticed in [22]. Indeed, consider an FO sentence φ in the vocabulary of one binary predicate E saying that E is a total relation (that is, $\forall x \forall y E(x, y)$). Suppose this φ is Hanf-local under logical equivalence, with d and k witnessing it. Consider two graphs: G_1 is a clique with $2k$ vertices, and G_2 is a disjoint union of two cliques having k vertices each. Then for an arbitrary bijection $f : G_1 \rightarrow G_2$ we have $N_d^{G_1}(c) = N_1^{G_1}(c) \equiv_k N_1^{G_2}(f(c)) = N_d^{G_2}(f(c))$ for all c and $d \geq 1$, and yet $G_1 \models \varphi$ and $G_2 \models \neg\varphi$.

Furthermore, the implication Hanf-locality \Rightarrow Gaifman-locality fails for the notions of locality based on logical equivalence. Of course a counterexample query Q cannot be FO-definable, by Theorem 4.2. The query Q works on graphs G with the edge relation E and additional binary relation R and unary relation C . If G is a union of two E -connected components of different sizes, each containing one element in C from which there is an E -edge to all other elements of the component, and R -edges between the two elements in C , then Q selects the C -element of the largest component. Otherwise, the output of Q is empty. It is not hard to see that $d = k = 1$ witness Hanf-locality under logical equivalence for Q . On the other hand, if both connected components are large enough, they cannot be distinguished in FO, and yet only in one of them an element is selected.




Fig. 5. Locally consistent transformation

5 Locally consistent transformation

We now move from queries – where we are interested in knowing whether $\bar{a} \in Q(\mathfrak{A})$ – to transformations that associate to a structure \mathfrak{A} another structure $\langle A, Q(\mathfrak{A}) \rangle$. A natural extension of the notion of locality is to require that isomorphic neighborhoods be sent to isomorphic neighborhoods. We first make it precise in the setting of one structure, following the definition of [1].

Definition 5.1 [Locally consistent transformation] A transformation $Q : \mathfrak{A} \rightarrow \langle A, Q(\mathfrak{A}) \rangle$ is *Gaifman-locally consistent* if for every $d \geq 0$, one can find $r \geq 0$ such that for every \mathfrak{A} and every $a, b \in A$ we have

$$N_r^{\mathfrak{A}}(a) \cong N_r^{\mathfrak{A}}(b) \Rightarrow N_d^{\langle A, Q(\mathfrak{A}) \rangle}(a) \cong N_d^{\langle A, Q(\mathfrak{A}) \rangle}(b).$$

This notion is illustrated in Fig. 5.

A natural question is whether all FO-definable transformations are locally consistent. The answer is negative in general.

Proposition 5.2 *There are transformations definable by quantifier-free FO formulae that are not Gaifman-locally consistent.*

An example is given in Figure 6. Suppose we have a graph with edge relation E whose nodes may be colored C_0 and C_1 , and an FO formula

$$\varphi(x, y) \equiv E(x, y) \vee (C_0(x) \wedge C_1(y)). \tag{1}$$

Then, given the graph G in Figure 6, with a, b being the nodes in C_0 , if G' is the graph obtained by applying φ to G , then $N_1^{G'}(a) \not\cong N_1^{G'}(b)$ as long as the two C_1 -chains are of different length (note that nodes to which a and b are connected belong to neither C_0 nor C_1). But for any fixed d we can pick them to be long enough so that $N_d^G(a) \cong N_d^G(b)$.

But there are cases when local consistency can be ensured for FO-definable transformations. The problem with the previous example is that in $\varphi(G)$, there are edges between elements that are very far apart in G . We exclude




Fig. 6. Failure of local consistency

this type of situations by means of the following definition. We say that an FO formula $\varphi(\bar{x})$ is *d-bounded* if, for every structure \mathfrak{A} and every tuple \bar{a} such that $\mathfrak{A} \models \varphi(\bar{a})$, the distance between any two components of \bar{a} in \mathfrak{A} itself does not exceed d .

Proposition 5.3 (see [1]) *If φ is a d-bounded FO formula for $d \geq 0$, then it defines a Gaifman-locally consistent transformation.*

While for general FO formulae the notion of *d-boundedness* is of course undecidable, there are many useful examples of *d-bounded* formulae. Atomic formulae are, of course, 1-bounded (and this has been used in data exchange applications, where they correspond to an important subclass of data exchange problems arising under the local-as-view scenario [6,1]). Conjunctive queries whose graph is connected are also *d-bounded* for some d that depends on the query.

So far we looked at transformations that define a structure with one relation. But we could have considered transformations given by several queries Q_1, \dots, Q_m : then we have a transformation $[Q_1, \dots, Q_m] : \mathfrak{A} \mapsto \langle A, Q_1(\mathfrak{A}), \dots, Q_m(\mathfrak{A}) \rangle$. It turns out that even if all Q_i 's are locally consistent, then $[Q_1, \dots, Q_m]$ need not be. The example is essentially the same as formula (1): both $E(x, y)$ and $(C_0(x) \wedge C_1(y))$ define locally consistent transformations, and yet taken together they do not produce one.

We now move to the Hanf-based notion. We say that a transformation given by Q is *Hanf-locally consistent* if for every $d \geq 0$, there exists $r \geq 0$ such that

$$\mathfrak{A} \stackrel{r}{\simeq} \mathfrak{B} \Rightarrow \langle A, Q(\mathfrak{A}) \rangle \stackrel{d}{\simeq} \langle B, Q(\mathfrak{B}) \rangle.$$

Then the same example as in (1) shows that this notion need not hold even for transformations definable by quantifier-free FO formulae. But just as in the case of Gaifman-based notion of local consistency, we can recover this for *d-bounded* transformations.

Proposition 5.4 *If φ is a d-bounded FO formula for $d \geq 0$, then it defines a Hanf-locally consistent transformation.*

The proof combines the proof of Proposition 5.3 and Gaifman's theorem [10]. Results in this section can also be extended to the case of neighborhoods of tuples rather than single elements [1].

6 Locally consistent transformations and logical equivalence

Just as we changed the notion of locality from isomorphism-based to logical equivalence-based, we can modify the definitions of locally consistent transformations by making weaker assumptions on neighborhoods.

For a query Q , we say that a transformation given by Q is *Hanf-locally consistent under logical equivalence* if for every $d, k \geq 0$, there exist $r, \ell \geq 0$ such that

$$\mathfrak{A} \stackrel{\equiv}{\leftarrow}_{r,\ell} \mathfrak{B} \Rightarrow \langle A, Q(\mathfrak{A}) \rangle \stackrel{\equiv}{\leftarrow}_{d,k} \langle B, Q(\mathfrak{B}) \rangle.$$

Then a slight modification of the example showing that Hanf-locality of queries under logical equivalence fails proves that this new notion fails too even for very simple FO sentences.

But unlike the case of isomorphism-based local consistency, moving to logical equivalence allows us to recover the Gaifman condition.

Definition 6.1 A transformation $Q : \mathfrak{A} \rightarrow \langle A, Q(\mathfrak{A}) \rangle$ is *Gaifman-locally consistent under logical equivalence* if for every $d, k \geq 0$, one can find $r, \ell \geq 0$ such that for every \mathfrak{A} and every $a, b \in A$ we have

$$N_r^{\mathfrak{A}}(a) \equiv_{\ell} N_r^{\mathfrak{A}}(b) \Rightarrow N_d^{\langle A, Q(\mathfrak{A}) \rangle}(a) \equiv_k N_d^{\langle A, Q(\mathfrak{A}) \rangle}(b).$$

Since the rank- k type of a d -neighborhood can be defined by a formula whose quantifier rank depends on d and k only, Gaifman's theorem gives us the following.

Proposition 6.2 *Every transformation given by an FO formula is Gaifman-locally consistent under logical equivalence.*

7 More expressive logics

While the Hanf notion of locality does not withstand more general definitions, the Gaifman notion appears more robust and can be recovered under logical equivalence for both queries and transformations.

We have looked at FO so far, so the next natural question is whether more expressive logics possess similar notions of locality. For isomorphism-based notions the answer is known. First, many extensions with counting or unary quantifiers are local [17,20,14]. Maximal in terms of expressiveness logics that are Gaifman- or Hanf-local (under isomorphism) have been characterized [18].

So it is natural to ask whether in logics that are known to be Gaifman-local under isomorphism, Gaifman-locality can be recovered under the notion of

logical equivalence, where, of course, equivalence with respect to FO formulae is replaced by equivalence with respect to formulae of the logic.

That is, assume that if we have a logic \mathcal{L} with a notion of $\equiv_k^{\mathcal{L}}$ similar to that of \equiv_k (that is, $\mathfrak{A} \equiv_k^{\mathcal{L}} \mathfrak{B}$ if \mathfrak{A} and \mathfrak{B} agree on all sentences of quantifier rank – appropriately defined for \mathcal{L} – up to k , or if there is a game for \mathcal{L} and the duplicator has a winning strategy in k rounds). Then locality under logical equivalence for \mathcal{L} states that for every formula $\varphi(\bar{x})$, there are numbers $d, k \geq 0$ such that $N_d^{\mathfrak{A}}(\bar{a}) \equiv_k^{\mathcal{L}} N_d^{\mathfrak{B}}(\bar{b})$ implies $\mathfrak{A} \models \varphi(\bar{a}) \leftrightarrow \varphi(\bar{b})$.

While it seems natural that such extensions should be true for logics that are Gaifman-local under isomorphism, the situation is more complex than expected. Here we consider extensions of FO with *simple unary quantifiers* [16,25]. A simple unary quantifier is a class \mathcal{K} of structures $\langle A, U \rangle$, where U is unary, closed under isomorphism. The logic $\text{FO}(\mathcal{Q}_{\mathcal{K}})$ extends FO by means of the following rule: if $\varphi(\bar{x}, y)$ is a formula, then $\psi(\bar{x}) = \mathcal{Q}_{\mathcal{K}}y\varphi(\bar{x}, y)$ is a formula. The semantics is as follows: $\mathfrak{A} \models \psi(\bar{a})$ iff the structure $\langle A, \{c \in A \mid \mathfrak{A} \models \varphi(\bar{a}, c)\} \rangle$ is in \mathcal{K} . The most commonly used unary quantifiers (besides the usual \exists and \forall) are *modulo quantifiers* [21,25] based on classes \mathcal{K}_p of structures $\langle A, U \rangle$ where $|U| = 0 \pmod{p}$. In this case we write \mathcal{Q}_p instead of $\mathcal{Q}_{\mathcal{K}_p}$.

We also consider the quantifier $\mathcal{Q}_{\text{prime}} = \mathcal{Q}_{\mathcal{K}_{\text{prime}}}$ where $\mathcal{K}_{\text{prime}}$ is the class of structures $\langle A, U \rangle$ where $|U|$ is prime.

It is known that extensions of FO by means of an arbitrary collection of simple unary generalized quantifiers are both Hanf- and Gaifman-local (under isomorphism) [13]. The following shows that the behavior of the logical equivalence-based notion of Gaifman-locality is not nearly as uniform.

Theorem 7.1 (see [2]) *a) For an arbitrary set of natural numbers p_1, \dots, p_m , the extension $\text{FO}(\mathcal{Q}_{p_1}, \dots, \mathcal{Q}_{p_m})$ of FO with $\mathcal{Q}_{p_1}, \dots, \mathcal{Q}_{p_m}$ is Gaifman-local under logical equivalence.*

b) $\text{FO}(\mathcal{Q}_{\text{prime}})$ is not Gaifman-local under logical equivalence.

8 Transformations that invent new values

So far we looked at transformations that sent a structure \mathfrak{A} to $\langle A, Q(\mathfrak{A}) \rangle$. But in many applications, especially in data exchange [6,7], transformations need to invent new values. Consider, for example, a database relation that stores data about employee ids, names, and salaries, and suppose one needs to create another relation that stores employee ids and departments. In this case one copies data about ids and creates new and distinct values, that are not present elsewhere in the database, for department names.

We assume that elements of finite universes come from some set D , and we also have a countable set V of values (often called variables or nulls in the

database context) that will be used as new invented values. A transformation of σ -structures to σ' -structures is given by rules of the form

$$(r) \quad \psi(\bar{x}, \bar{z}) \quad :- \quad \varphi(\bar{x}, \bar{y}),$$

where ψ is a conjunction of σ' -atoms, and φ is an FO formula of vocabulary σ . The semantics of this rule (r) as follows: if $\varphi(\bar{a}, \bar{b})$ is true in a σ -structure \mathfrak{A} , then one picks a tuple \bar{v} with $|\bar{v}| = |\bar{z}|$ of fresh elements from V , and adds tuples specified by $\psi(\bar{a}, \bar{v})$ to a σ' -structure. For example, if (r) is the rule $R_1(x_1, z), R_2(z, z', x_2) \quad :- \quad \varphi(x_1, x_2, y)$, and $\varphi(a, b, c)$ is true in an input structure \mathfrak{A} , then we pick two new elements v, v' from V and add a tuple (a, v) to R_1 and a tuple (v, v', b) to R_2 .

The result of applying (r) to \mathfrak{A} will be denoted by $\mathcal{F}_r(\mathfrak{A})$. Notice that $\mathcal{F}_r(\mathfrak{A})$ is unique up to renaming elements from V . Notice also that if (r) is of the form $R(\bar{x}) \quad :- \quad \varphi(\bar{x})$, then $\mathcal{F}_r(\mathfrak{A})$ is simply $\varphi(\mathfrak{A}) = \langle A, \{\bar{a} \mid \mathfrak{A} \models \varphi(\bar{a})\} \rangle$.

One can reformulate the previous notions of Hanf- and Gaifman-local consistency for transformations of the form \mathcal{F}_r . It turns out they behave in the same way as transformations that do not invent new values.

Theorem 8.1 (see [1]) *The transformation \mathcal{F}_r is Gaifman-locally consistent under logical equivalence. Moreover, if φ in (r) is d -bounded for some $d \geq 0$, then \mathcal{F}_r is Gaifman-locally consistent.*

In fact Theorem 8.1 is true for more general transformations that consist of a sequence rules, with each rule referring in its body to relations mentioned in the heads of rules appearing earlier in the sequence [1]. These rule-based transformations are also closely related to FO transductions [3], but they appear to be more convenient for presenting transformations and their locality properties.

A central problem in data exchange is query answering, since queries are formulated over σ' -relations, but then input is normally a σ -structure \mathfrak{A} . Several ways of answering queries have been proposed [6,7]. The most common one seems to be the following. If we have an input \mathfrak{A} and a formula $\alpha(\bar{x})$ over σ' , then $\text{answer}_{(r)}(\alpha, \mathfrak{A})$ is the set $\{\bar{a} \in A^{|\bar{a}|} \mid \mathcal{F}_r(\mathfrak{A}) \models \alpha(\bar{a})\}$ (one only keeps elements of A since $\mathcal{F}_r(\mathfrak{A})$ is specified up to renaming of elements of V) [6].

Theorem 8.1 gives us a useful locality property for $\text{answer}_{(r)}(\alpha, \mathfrak{A})$:

Corollary 8.2 *Given a rule (r) as above and a formula α over σ' , there is a number d such that $N_d^{\mathfrak{A}}(\bar{a}_1) \cong N_d^{\mathfrak{A}}(\bar{a}_2)$ implies $\bar{a}_1 \in \text{answer}_{(r)}(\alpha, \mathfrak{A}) \Leftrightarrow \bar{a}_2 \in \text{answer}_{(r)}(\alpha, \mathfrak{A})$, for every σ -structure \mathfrak{A} .*

This result can easily be used to show that certain queries are *not* answerable in the data exchange scenario [1].

	FO queries: locality		FO transformations: local consistency	
	Hanf	Gaifman	Hanf	Gaifman
under isomorphism	<i>yes</i>	<i>yes</i>	<i>no</i> in general <i>yes</i> for d -bounded	<i>no</i> in general <i>yes</i> for d -bounded
under logical equivalence	<i>no</i>	<i>yes</i>	<i>no</i>	<i>yes</i>

Fig. 7. Summary of main results

We conclude with a summary of the main results surveyed here, presented in Figure 7.

Acknowledgement

The notion of local consistency, investigated in [1], was originally formulated by Ron Fagin in 2003. I am grateful to Ron, and to Marcelo Arenas and Pablo Barceló, for collaboration on locality results upon which this paper is based.

References

- [1] M. Arenas, P. Barceló, R. Fagin, L. Libkin. Locally consistent transformations and query answering in data exchange. In *PODS'04*, pages 229–240.
- [2] M. Arenas, P. Barceló, L. Libkin. Game-based notions of locality over finite models. In *CSL'04*, pages 175–189.
- [3] B. Courcelle. Graph operations, graph transformations and monadic second-order logic: a survey. *WoLLIC'01*, Electr. Notes Theor. Comput. Sci. 51: (2001).
- [4] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer Verlag, 1995.
- [5] K. Etessami. Counting quantifiers, successor relations, and logarithmic space, *JCSS* 54 (1997), 400–411.
- [6] R. Fagin, Ph. Kolaitis, R. Miller, L. Popa. Data exchange: semantics and query answering. *TCS*, to appear. Extended abstract in *ICDT'03*, pages 207–224.
- [7] R. Fagin, Ph. Kolaitis, L. Popa. Data exchange: getting to the core. In *PODS'03*, pages 90–101.
- [8] R. Fagin, L. Stockmeyer and M. Vardi, On monadic NP vs monadic co-NP, *Inf.&Comput.* 120 (1994), 78–92.
- [9] J. Flum, M. Grohe. Fixed-parameter tractability, definability, and model-checking. *SIAM J. Comput.* 31(1): 113–145 (2001).
- [10] H. Gaifman. On local and non-local properties, *Logic Colloquium '81*, North Holland, 1982.
- [11] M. Grohe, S. Wöhrle. An existential locality theorem. In *CSL 2001*, pages 99–114.
- [12] 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.

- [13] L. Hella, L. Libkin and J. Nurmonen. Notions of locality and their logical characterizations over finite models. *JSL* 64 (1999), 1751–1773.
- [14] L. Hella, L. Libkin, J. Nurmonen and L. Wong. Logics with aggregate operators. *Journal of the ACM* 48 (2001), 880–907.
- [15] N. Immerman. *Descriptive Complexity*. Springer Verlag, 1998.
- [16] Ph. Kolaitis and J. Väänänen. Generalized quantifiers and pebble games on finite structures. *Annals of Pure and Applied Logic*, 74 (1995), 23–75.
- [17] L. Libkin. On counting logics and local properties. *ACM Trans. on Comput. Logic*, 1 (1) (2000), 33–59.
- [18] L. Libkin. Logics capturing local properties. *ACM Trans. on Comput. Logic*, 2 (1) (2001), 135–153.
- [19] L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
- [20] J. Nurmonen. On winning strategies with unary quantifiers. *J. Logic Comput.*, 6 (1996), 779–798.
- [21] J. Nurmonen. Counting modulo quantifiers on finite structures. *Inf. & Comput.*, 160 (2000), 62–87.
- [22] T. Schwentick and K. Barthelmann. Local normal forms for first-order logic with applications to games and automata. *STACS'98*, pages 444–454.
- [23] D. Seese. Linear time computable problems and first-order descriptions. *Math. Str. Comput. Sci.* 6(6): 505–526 (1996).
- [24] W. Thomas. Languages, automata, and logic. In *Handbook of Formal Languages, Vol. 3*, Springer, 1997, pages 389–455.
- [25] J. Väänänen. Unary quantifiers on finite models. *J. Logic, Language and Information* 6 (1997), 275–304.