

ON THE LOCALITY OF ARB-INVARIANT FIRST-ORDER FORMULAS WITH MODULO COUNTING QUANTIFIERS*

FREDERIK HARWATH^a AND NICOLE SCHWEIKARDT^b

^a Institut für Informatik, Goethe-Universität Frankfurt, Germany
e-mail address: harwath@cs.uni-frankfurt.de

^b Institut für Informatik, Humboldt-Universität zu Berlin, Germany
e-mail address: schweikn@informatik.hu-berlin.de

ABSTRACT. We study Gaifman locality and Hanf locality of an extension of first-order logic with modulo p counting quantifiers (FO+MOD $_p$, for short) with arbitrary numerical predicates. We require that the validity of formulas is independent of the particular interpretation of the numerical predicates and refer to such formulas as arb-invariant formulas. This paper gives a detailed picture of locality and non-locality properties of arb-invariant FO+MOD $_p$. For example, on the class of all finite structures, for any $p \geq 2$, arb-invariant FO+MOD $_p$ is neither Hanf nor Gaifman local with respect to a sublinear locality radius. However, in case that p is an odd prime power, it is *weakly* Gaifman local with a polylogarithmic locality radius. And when restricting attention to the class of string structures, for odd prime powers p , arb-invariant FO+MOD $_p$ is both Hanf and Gaifman local with a polylogarithmic locality radius. Our negative results build on examples of order-invariant FO+MOD $_p$ formulas presented in Niemistö's PhD thesis. Our positive results make use of the close connection between FO+MOD $_p$ and Boolean circuits built from NOT-gates and AND-, OR-, and MOD $_p$ - gates of arbitrary fan-in.

1. INTRODUCTION

Expressivity of logics over finite structures plays an important role in various areas of computer science. In descriptive complexity, logics are used to characterise complexity classes, and concerning databases, common query languages have well-known logical equivalents. These applications have motivated a systematic study of the expressive powers of logics on finite structures. The classical inexpressibility arguments for logics over finite structures (i.e., back-and-forth systems or Ehrenfeucht-Fraïssé games; cf. [12]) often involve nontrivial combinatorics. Notions of *locality* have been proposed as an alternative that allows to contain much of the hard combinatorial work in generic results.

2012 ACM CCS: [Theory of computation]: Computational complexity and cryptography—Circuit complexity / Complexity classes; Logic—Finite Model Theory; Theory and algorithms for application domains—Database theory—Database query languages (principles);

Key words and phrases: finite model theory, Gaifman and Hanf locality, first-order logic with modulo counting quantifiers, order-invariant and arb-invariant formulas, lower bounds in circuit complexity .

* This is the full version of the conference contribution [10].

The two best known notions of locality are *Gaifman locality* and *Hanf locality*, introduced in [11, 8]. A k -ary query is called *Gaifman local with locality radius $f(n)$* if in a structure of cardinality n , the question whether a given tuple satisfies the query only depends on the isomorphism type of the tuple's neighbourhood of radius $f(n)$. A Boolean query is *Hanf local with locality radius $f(n)$* if the question whether a structure of size n satisfies the query only depends on the number of occurrences of isomorphism types of neighbourhoods of radius $f(n)$. If a given logic is capable of defining only Gaifman or Hanf local queries with a sublinear locality radius, then this logic cannot express “non-local” queries such as, e.g., the query asking whether two nodes of a graph are connected by a path, or the query asking whether a graph is acyclic (cf., e.g., the textbook [12]). It is well-known that first-order logic FO, as well as extensions of FO by various kinds of counting quantifiers, are Gaifman local and Hanf local with a constant locality radius [11, 8]. Also, locality properties of extensions of FO by invariant uses of order and arithmetic have been considered [9, 1].

Order-invariant and arb-invariant formulas were introduced to capture the data independence principle in databases: An implementation of a database query may exploit the order in which the database elements are stored in memory, and thus identify the elements with natural numbers on which arithmetic can be performed. But the use of order and arithmetic should be restricted in such a way that the result of the query does not depend on the particular order in which the data is stored. In addition to the relations of a given structure, *arb-invariant* formulas are allowed to use a linear order $<$ and arithmetic predicates such as $+$ or \times induced by the order. But the truth of such a formula in a structure must be independent of the particular linear order which is chosen for $<$. Arb-invariant formulas that only use the linear order, but no further arithmetic predicates, are called *order-invariant*. In [9] it was shown that order-invariant FO can express only queries that are Gaifman local with a constant locality radius, and from [1] we know that arb-invariant FO can express only queries that are Gaifman local with a polylogarithmic locality radius. The proof of [1] relies on a reduction using strong lower bound results from circuit complexity, concerning AC^0 -circuits. Similar lower bounds are known also for the extension of AC^0 -circuits by modulo p counting gates, for a prime power p [17]. This naturally raises the question whether the locality results from [1] can be generalised to the extension of FO by modulo p counting quantifiers (FO+MOD $_p$, for short), which precisely corresponds to AC^0 -circuits with modulo p counting gates [3]. This question was the starting point for the investigations carried out in the present paper.

Our results give a detailed picture of the locality and non-locality properties of order-invariant and arb-invariant FO+MOD $_p$: For every natural number $p \geq 2$, order-invariant FO+MOD $_p$ is neither Hanf nor Gaifman local with a sublinear locality radius (see Section 4 and Proposition 3.2). For *even* numbers $p \geq 2$, order-invariant FO+MOD $_p$ is not even *weakly* Gaifman local with a sublinear locality radius (Proposition 3.4). Here, *weak* Gaifman locality is a relaxed notion of Gaifman locality referring only to tuples with disjoint neighbourhoods (cf., [12]). However, for *odd prime powers* p we can show that arb-invariant FO+MOD $_p$ is weakly Gaifman local with a polylogarithmic locality radius (Theorem 3.5). For showing the latter result, we introduce a new locality notion called *shift locality*, for which we can prove for all prime powers p that arb-invariant FO+MOD $_p$ is shift local with a polylogarithmic locality radius (Theorem 3.7). Our proof relies on Smolensky's circuit lower bound [17]. Generalising our result from prime powers p to arbitrary numbers p can be expected to be difficult, since it would solve long-standing open questions in circuit complexity (see Remark 3.13).

When restricting attention to the class of *string structures*, we obtain for odd prime powers p , that arb-invariant $\text{FO}+\text{MOD}_p$ is both Hanf and Gaifman local with a polylogarithmic locality radius (Theorem 4.6 and Corollary 4.7). On the other hand, for even numbers $p \geq 2$, order-invariant $\text{FO}+\text{MOD}_p$ on string structures is neither Gaifman nor Hanf local with a sublinear locality radius (Proposition 3.4 and Section 4). This, in particular, implies that order-invariant $\text{FO}+\text{MOD}_p$ is strictly more expressive on strings than $\text{FO}+\text{MOD}_p$, refuting a conjecture of Benedikt and Segoufin [6].

Outline. Section 2 fixes the basic notation, introduces order-invariance and arb-invariance, and recalls two examples of order-invariant $\text{FO}+\text{MOD}_p$ -formulas from Niemistö's PhD-thesis [14]. Section 3 presents our results concerning Gaifman locality, weak Gaifman locality, and shift locality. Section 4 deals with Gaifman and Hanf locality on string structures. Section 5 concludes the paper.

Acknowledgements. We want to thank the anonymous reviewers for their valuable suggestions.

2. PRELIMINARIES

We assume that the reader is familiar with the basic concepts and notations concerning first-order logic and extensions thereof which can be found in most modern introductory books on mathematical logic. In particular, we use the standard semantics of first-order logic. The exposition given in the textbooks [12, 7] is close to our concerns. These textbooks also provide the reader with the necessary background on the relation of first-order logic to circuit complexity, on notions of locality, and on order-invariant first-order formulas.¹

Basic notation. We write \mathbb{N} for the set of non-negative integers and let $\mathbb{N}_{\geq 1} := \mathbb{N} \setminus \{0\}$. For $n \in \mathbb{N}_{\geq 1}$ we write $[n]$ for the set $\{i \in \mathbb{N} : 0 \leq i < n\}$, i.e., $[n] = \{0, \dots, n-1\}$. For integers i, j, p with $p \geq 2$, we write $i \equiv j \pmod p$ (and say that i is congruent j modulo p) iff there exists an integer k such that $i = j + kp$. For integers i, i' , the term $(i+i' \pmod p)$ denotes the number $j \in [p]$ such that $i+i' \equiv j \pmod p$. A number p is called a *prime power* if $p = \hat{p}^i$ for a prime \hat{p} and an integer $i \geq 1$, and p is called an *odd prime power* if p 's prime factor is different from 2 (i.e., p is odd). We write $\log n$ to denote the logarithm of a number n with respect to base 2, and we often simply write $\log n$ instead of $\lfloor \log n \rfloor$.

By 2^A we denote the power set of A , i.e., the set $\{Y : Y \subseteq A\}$. The set of all non-empty finite strings built from symbols in A is denoted A^+ . We write $|w|$ for the length of a string $w \in A^+$. For an $a \in A$ we write $|w|_a$ for the number of occurrences of the letter a in the string w .

Structures. A *signature* σ is a set of relation symbols R , each of them associated with a fixed arity $\text{ar}(R) \in \mathbb{N}_{\geq 1}$. Throughout this paper, σ will usually denote a fixed finite signature.

A σ -*structure* \mathcal{A} consists of a non-empty set A called the *universe* of \mathcal{A} , and a relation $R^{\mathcal{A}} \subseteq A^{\text{ar}(R)}$ for each relation symbol $R \in \sigma$. The *cardinality* of a σ -structure \mathcal{A} is the cardinality of its universe. *Finite* σ -structures are σ -structures of finite cardinality. For σ -structures \mathcal{A} and \mathcal{B} and tuples $\bar{a} = (a_1, \dots, a_k) \in A^k$ and $\bar{b} = (b_1, \dots, b_k) \in B^k$ we write $(\mathcal{A}, \bar{a}) \cong (\mathcal{B}, \bar{b})$ to indicate that there is an isomorphism π from \mathcal{A} to \mathcal{B} that maps \bar{a} to \bar{b} (i.e., $\pi(a_i) = b_i$ for each $i \leq k$).

¹These topics are more emphasised in [12] which, in particular, contains several examples for the different notions of locality.

We represent strings over a finite alphabet Σ by successor-based structures as follows: We choose $\sigma_\Sigma := \{E\} \cup \{P_a : a \in \Sigma\}$, where E is a binary relation symbol and P_a is a unary relation symbol, for each $a \in \Sigma$. We represent a non-empty string $w \in \Sigma^+$ by the σ_Σ -structure \mathcal{S}_w , where the universe of \mathcal{S}_w is the set $\{1, \dots, |w|\}$ of positions of w , the edge relation $E^{\mathcal{S}_w}$ is the successor relation, i.e., $E^{\mathcal{S}_w} = \{(i, i+1) : 1 \leq i < |w|\}$, and for each $a \in \Sigma$, the set $P_a^{\mathcal{S}_w}$ consists of all positions of w that carry the letter a . Structures of the form \mathcal{S}_w (for a string w) are called *string structures*. Note that in the literature, strings are also often represented by linearly ordered structures, whereas our representation contains only the successor relation, but not the linear order.

In this paper, we restrict attention to *finite* structures. All *classes* \mathfrak{C} of σ -structures will be closed under isomorphism, i.e., if \mathcal{A} and \mathcal{B} are isomorphic σ -structures, then $\mathcal{A} \in \mathfrak{C}$ iff $\mathcal{B} \in \mathfrak{C}$. We will write Σ -*strings* to denote the class of all σ_Σ -structures that represent strings in Σ^+ (i.e., Σ -*strings* is the closure under isomorphisms of the set $\{\mathcal{S}_w : w \in \Sigma^+\}$).

First-order logic with modulo counting quantifiers. By $free(\varphi)$ we denote the set of all free variables of a formula φ . A *sentence* is a formula φ with $free(\varphi) = \emptyset$. We often write $\varphi(\bar{x})$, for $\bar{x} = (x_1, \dots, x_k)$, to indicate that $free(\varphi) = \{x_1, \dots, x_k\}$. If \mathcal{A} is a σ -structure and $\bar{a} = (a_1, \dots, a_k) \in A^k$, we write $\mathcal{A} \models \varphi[\bar{a}]$ to indicate that the formula $\varphi(\bar{x})$ is satisfied in \mathcal{A} when interpreting the free occurrences of the variables x_1, \dots, x_k with the elements a_1, \dots, a_k .

We write $FO(\sigma)$ to denote the class of all first-order formulas of signature σ . In this paper, we consider the extension of $FO(\sigma)$ by *modulo counting quantifiers*, defined as follows: Let p be a natural number with $p \geq 2$. A *modulo p counting quantifier* is of the form $\exists^{i \bmod p}$ for some $i \in [p]$. A formula of the form $\exists^{i \bmod p} x \varphi(x, \bar{y})$ is satisfied by a finite σ -structure \mathcal{A} and an interpretation $\bar{b} \in A^k$ of the variables \bar{y} iff the number of elements $a \in A$ such that $\mathcal{A} \models \varphi[a, \bar{b}]$ is congruent i modulo p .

For a fixed natural number $p \geq 2$ we write $FO+MOD_p(\sigma)$ to denote the extension of $FO(\sigma)$ by modulo p counting quantifiers. That is, $FO+MOD_p(\sigma)$ is built from atomic formulas of the form $x_1=x_2$ and $R(x_1, \dots, x_{ar(R)})$, for $R \in \sigma$ and variables $x_1, x_2, \dots, x_{ar(R)}$, and closed under Boolean connectives \wedge, \vee, \neg , existential and universal first-order quantifiers \exists, \forall , and modulo p counting quantifiers $\exists^{i \bmod p}$, for $i \in [p]$. This logic has been studied in depth, see e.g., [18, 11, 3]. Note that if m is a multiple of p , then $FO+MOD_m$ can express modulo p counting quantifiers, since $\exists^{i \bmod p} x \varphi(x, \bar{y})$ is equivalent to $\bigvee_{0 \leq j < m/p} \exists^{jp+i \bmod m} x \varphi(x, \bar{y})$.

Arb-invariant formulas. We can extend the expressive power of a logic by allowing formulas to use, apart from the relation symbols present in the signature σ , also a linear order $<$, arithmetic predicates such as $+$ or \times , or arbitrary numerical predicates. We define an r -ary *numerical predicate* $P^{\mathbb{N}}$ as an r -ary relation on \mathbb{N} (i.e., $P^{\mathbb{N}} \subseteq \mathbb{N}^r$). Note that, compared to the definition of numerical predicates used, e.g., in [18], this definition simplifies some details, while it is essentially equivalent. That is, each r -ary numerical predicate according to the definition of [18] can be represented by an $(r+1)$ -ary numerical predicate according to our definition. Two examples of numerical predicates are the linear order $<^{\mathbb{N}}$ consisting of all tuples $(a, b) \in \mathbb{N}^2$ with $a < b$, and the addition predicate $+^{\mathbb{N}}$ consisting of all triples $(a, b, c) \in \mathbb{N}^3$ with $a + b = c$.

To allow formulas to use numerical predicates, we fix the following notation: For every $r \in \mathbb{N}_{\geq 1}$ and every r -ary numerical predicate $P^{\mathbb{N}}$, let P be a new relation symbol of arity r (“new” meaning that P does not belong to σ). We write η_{arb} to denote the set of all the

relation symbols P obtained this way, and let $\sigma_{\text{arb}} := \sigma \cup \eta_{\text{arb}}$ (the subscript “arb” stands for “arbitrary numerical predicates”).

Next, we would like to allow $\text{FO}+\text{MOD}_p(\sigma_{\text{arb}})$ -formulas to make meaningful statements about finite σ -structures. To this end, for a finite σ -structure \mathcal{A} , we consider embeddings ι of the universe of \mathcal{A} into the initial segment of \mathbb{N} of size $n = |\mathcal{A}|$, i.e., the set $[n] = \{0, \dots, n-1\}$.

Definition 2.1 (Embedding). Let \mathcal{A} be a finite σ -structure, and let $n := |\mathcal{A}|$. An *embedding* ι of \mathcal{A} is a bijection $\iota : A \rightarrow [n]$.

Given a finite σ -structure \mathcal{A} and an embedding ι of \mathcal{A} , we can translate r -ary numerical predicates $P^{\mathbb{N}}$ into r -ary predicates on A as follows: $P^{\mathbb{N}}$ induces the r -ary predicate P^ι on A , consisting of all r -tuples $\bar{a} = (a_1, \dots, a_r) \in A^r$ where $\iota(\bar{a}) = (\iota(a_1), \dots, \iota(a_r)) \in P^{\mathbb{N}}$. In particular, the linear order $<^{\mathbb{N}}$ induces the linear order $<^\iota$ on A where for all $a, b \in A$ we have $a <^\iota b$ iff $\iota(a) <^{\mathbb{N}} \iota(b)$.

The σ_{arb} -structure \mathcal{A}^ι associated with \mathcal{A} and ι is the expansion of \mathcal{A} by the predicates P^ι for all $P \in \eta_{\text{arb}}$. That is, \mathcal{A}^ι has the same universe as \mathcal{A} , all relation symbols $R \in \sigma$ are interpreted in \mathcal{A}^ι in the same way as in \mathcal{A} , and every numerical symbol $P \in \eta_{\text{arb}}$ is interpreted by the relation P^ι .

To ensure that an $\text{FO}+\text{MOD}_p(\sigma_{\text{arb}})$ -formula φ makes a meaningful statement about a σ -structure \mathcal{A} , we evaluate φ in \mathcal{A}^ι , and we restrict attention to those formulas whose truth value is independent of the particular choice of the embedding ι . This is formalised by the following notion.

Definition 2.2 (Arb-invariance). Let $\varphi(\bar{x})$ be an $\text{FO}+\text{MOD}_p(\sigma_{\text{arb}})$ -formula with k free variables, and let \mathcal{A} be a finite σ -structure. The formula $\varphi(\bar{x})$ is *arb-invariant on \mathcal{A}* if for all embeddings ι_1 and ι_2 of \mathcal{A} and for all tuples $\bar{a} \in A^k$ we have: $\mathcal{A}^{\iota_1} \models \varphi[\bar{a}] \iff \mathcal{A}^{\iota_2} \models \varphi[\bar{a}]$.

Let $\varphi(\bar{x})$ be arb-invariant on \mathcal{A} . We write $\mathcal{A} \models \varphi[\bar{a}]$, if $\mathcal{A}^\iota \models \varphi[\bar{a}]$ for some (and hence every) embedding ι of \mathcal{A} .

Definition 2.3 (arb-inv- $\text{FO}+\text{MOD}_p$). An $\text{FO}+\text{MOD}_p(\sigma_{\text{arb}})$ -formula $\varphi(\bar{x})$ is *arb-invariant on a class \mathfrak{C}* of finite σ -structures, if $\varphi(\bar{x})$ is arb-invariant on every $\mathcal{A} \in \mathfrak{C}$. We denote the set of all $\text{FO}+\text{MOD}_p(\sigma_{\text{arb}})$ -formulas that are arb-invariant on \mathfrak{C} by $\text{arb-inv-FO}+\text{MOD}_p^{\mathfrak{C}}(\sigma)$.

A formula $\varphi(\bar{x})$ is called *arb-invariant* if it is arb-invariant on the class of all finite σ -structures. We write $\text{arb-inv-FO}+\text{MOD}_p(\sigma)$ to denote the set of all arb-invariant $\text{FO}+\text{MOD}_p(\sigma_{\text{arb}})$ -formulas.

Definition 2.4 (Order-invariance and $<$ -inv- $\text{FO}+\text{MOD}_p$).

An arb-invariant formula that only uses the numerical predicate $<^{\mathbb{N}}$ is called *order-invariant*. By $<$ -inv- $\text{FO}+\text{MOD}_p(\sigma)$ we denote the set of all arb-invariant $\text{FO}+\text{MOD}_p(\sigma \cup \{<\})$ -formulas.

Next, we present two examples of $<$ -inv- $\text{FO}+\text{MOD}_p(\sigma)$ -sentences that were developed by Niemistö in [14] and that will be used later on in this paper as examples for the locality and non-locality properties of $\text{arb-inv-FO}+\text{MOD}_p(\sigma)$ -sentences.

Example 2.5 (Niemistö (Proposition 6.22 in [14])). Let $\sigma = \{E\}$ be the signature consisting of a binary relation symbol E . This example presents an $<$ -inv- $\text{FO}+\text{MOD}_2(\sigma)$ -sentence $\varphi_{\text{even cycles}}$ that is satisfied by exactly those finite σ -structures \mathcal{A} that are disjoint unions of directed cycles where the number of cycles of even length is even.

Clearly, a finite σ -structure \mathcal{A} is a disjoint union of directed cycles iff every element $a \in A$ has in-degree 1 and out-degree 1 — and this can easily be expressed by an $\text{FO}(\sigma)$ -sentence

φ_{cycles} .

Each \mathcal{A} with $\mathcal{A} \models \varphi_{cycles}$ can be identified with the permutation $\pi_{\mathcal{A}}$ of A where, for every $a \in A$, $\pi_{\mathcal{A}}(a) = b$ for the unique element $b \in A$ with $(a, b) \in E^{\mathcal{A}}$. Note that the cycles of \mathcal{A} precisely correspond to the cycle decomposition of the permutation $\pi_{\mathcal{A}}$.

Now let $n := |A|$, let ι be an arbitrary embedding of \mathcal{A} into $[n]$, and let $\pi_{\mathcal{A}}^{\iota}$ be defined via $\pi_{\mathcal{A}}^{\iota}(j) := \iota(\pi_{\mathcal{A}}(\iota^{-1}(j)))$, for every $j \in [n]$. Clearly, $\pi_{\mathcal{A}}^{\iota}$ is a permutation of $[n]$, and the cycle decomposition of $\pi_{\mathcal{A}}^{\iota}$ is obtained from the cycle decomposition of $\pi_{\mathcal{A}}$ by replacing every $a \in A$ with the number $\iota(a)$.

It is a well-known fact concerning permutation groups (see [15]) that the cycle decomposition of $\pi_{\mathcal{A}}^{\iota}$ has an even number of cycles of even length if, and only if, the number of *inversions*, i.e., pairs $(x, y) \in [n]^2$ with $x < y$ and $\pi_{\mathcal{A}}^{\iota}(y) < \pi_{\mathcal{A}}^{\iota}(x)$, is even. The latter can easily be expressed in $\text{FO}+\text{MOD}_2(\sigma \cup \{<\})$ by the formula

$$\varphi_{inversions} := \exists^{0 \bmod 2} x \exists^{1 \bmod 2} y (x < y \wedge \pi(y) < \pi(x)),$$

where $\pi(y) < \pi(x)$ is an abbreviation for the formula $\exists x' \exists y' (E(x, x') \wedge E(y, y') \wedge y' < x')$.

In summary, the formula $\varphi_{even\ cycles} := (\varphi_{cycles} \wedge \varphi_{inversions})$ is an $<$ -inv- $\text{FO}+\text{MOD}_2(\sigma)$ -sentence that is satisfied by exactly those finite σ -structures that are disjoint unions of cycles where the number of cycles of even length is even. \square

Example 2.6 (Niemistö (Proposition 6.20 in [14])). Let $\sigma = \{E_1, E_2\}$ be the signature consisting of two binary relation symbols E_1 and E_2 . Let $h \in \mathbb{N}$. A *torus* is a σ -structure \mathcal{T} with universe $[h] \times [w]$, for some $w \geq 2$, and relations

$$\begin{aligned} E_1^{\mathcal{T}} &:= \{((i, j), (i+1 \bmod h, j)) : i \in [h], j \in [w]\} \\ E_2^{\mathcal{T}} &:= \{((i, j), (i, j+1)) : i \in [h], j \in [w-1]\} \cup \\ &\quad \{((i, w-1), (i+k \bmod h, 0)) : i \in [h]\}, \text{ for some } k \in [h]. \end{aligned}$$

The number w is the *width* of \mathcal{T} . The set $[h] \times \{j\}$ is the j -th *column* of \mathcal{T} , for each $j \in [w]$. The number h is the *height* of \mathcal{T} which will be fixed throughout this example. The number k is the *twist* of \mathcal{T} which we also denote by $twist(\mathcal{T})$. See Figure 1 for an illustration of two tori with different twist.





Figure 1: Tori of height 3 and width 4 with twist 0 (left) and twist 1 (right). The E_1 - and E_2 -edges are depicted by solid arcs and dotted arcs, respectively.

We consider disjoint unions of tori of height h . The *twist* of a disjoint union \mathcal{A} of tori $\mathcal{T}_1, \dots, \mathcal{T}_\ell$ is defined as $twist(\mathcal{A}) := twist(\mathcal{T}_1) + \dots + twist(\mathcal{T}_\ell) \bmod h$. This example presents an $<$ -inv- $\text{FO}+\text{MOD}_h(\sigma)$ -sentence φ_h which defines the class of disjoint unions of tori of height h and twist 0 on the class of finite σ -structures.

Using the Hanf-locality of $\text{FO}+\text{MOD}_h$ (cf. [11]), it can be shown that plain $\text{FO}+\text{MOD}_h$ cannot distinguish between tori with twist 0 and tori with twist 1 if their width is sufficiently

large, depending on the sentence. Here, we show that $\langle\text{-inv-FO+MOD}_h$ can distinguish between tori with twist 0 and tori with twist 1. To this end, we use the order to choose a distinguished element from each column of a torus.

A *list of representatives* for a torus \mathcal{T} of width w and height h is a tuple $R = (a_0, \dots, a_{w-1})$ of nodes of \mathcal{T} where, for each $j \in [w]$, a_j belongs to the j -th column of \mathcal{T} . The j -th *turn path* with respect to R is the directed $E_1^{\mathcal{T}}$ -path from a_j to the $E_2^{\mathcal{T}}$ -successor of $a_{(j-1) \bmod w}$. The length of the j -turn path w.r.t R is called the R -turn of the j -th column of \mathcal{T} , in symbols: $\text{turn}_j(\mathcal{T}, R)$. Note that $\text{turn}_j(\mathcal{T}, R) \in [h]$. The R -turn of \mathcal{T} is defined as $\text{turn}(\mathcal{T}, R) := \text{turn}_0(\mathcal{T}, R) + \dots + \text{turn}_{w-1}(\mathcal{T}, R)$. See the picture to the right for an illustration of a torus with a set of representatives R with R -turn 0; the turn paths are drawn thick and red.



Claim. $\text{turn}(\mathcal{T}, R) \equiv \text{twist}(\mathcal{T}) \pmod{h}$, for every list R of representatives for \mathcal{T} .

Proof. Let $k := \text{twist}(\mathcal{T})$. First note that if $R = (a_0, \dots, a_{w-1})$ is such that $a_j = (0, j)$, for every $j \in [w]$, then the following is true: $\text{turn}_j(\mathcal{T}, R) = 0$ for every $j \in \{1, \dots, w-1\}$ (as a_j is the $E_2^{\mathcal{T}}$ -successor of a_{j-1}), and $\text{turn}_0(\mathcal{T}, R) = k$ (as $a_0 = (0, 0)$ and the $E_2^{\mathcal{T}}$ -successor of $(0, w-1)$ is $(k, 0)$). Thus, $\text{turn}(\mathcal{T}, R) = k$.

Now consider arbitrary lists $R = (a_0, \dots, a_{w-1})$ of representatives. By induction on the number n_R of elements a_j such that $a_j \neq (0, j)$, we show that $\text{turn}(\mathcal{T}, R) \equiv k \pmod{h}$.

The induction base where $n_R = 0$ has been shown already. For the induction step, let j be such that $a_j = (p, j)$ for some $p \in \{1, \dots, h-1\}$. Let R' be obtained from R by replacing a_j with the element $(0, j)$. By induction we know that $\text{turn}(\mathcal{T}, R') \equiv k \pmod{h}$. By a straightforward case distinction, one can verify that the following is true: $\text{turn}_j(\mathcal{T}, R') \equiv \text{turn}_j(\mathcal{T}, R) + p \pmod{h}$ and $\text{turn}_{j+1 \bmod w}(\mathcal{T}, R') \equiv \text{turn}_{j+1 \bmod w}(\mathcal{T}, R) - p \pmod{h}$. Furthermore, $\text{turn}_{j'}(\mathcal{T}, R') = \text{turn}_{j'}(\mathcal{T}, R)$ for all $j' \in [w] \setminus \{j, j+1 \bmod w\}$. Hence, $\text{turn}(\mathcal{T}, R) \equiv \text{turn}(\mathcal{T}, R') \pmod{h}$. This completes the proof of the claim. \square

Now consider a disjoint union \mathcal{A} of tori $\mathcal{T}_1, \dots, \mathcal{T}_\ell$ of height h . From the previous claim, we obtain

$$\begin{aligned} \text{twist}(\mathcal{A}) &\equiv \text{twist}(\mathcal{T}_1) + \dots + \text{twist}(\mathcal{T}_\ell) \pmod{h} \\ &\equiv \text{turn}(\mathcal{T}_1, R_1) + \dots + \text{turn}(\mathcal{T}_\ell, R_\ell) \pmod{h}, \end{aligned}$$

for all lists of representatives R_1, \dots, R_ℓ of $\mathcal{T}_1, \dots, \mathcal{T}_\ell$.

We proceed with the construction of an $\text{FO+MOD}_h(\sigma \cup \{\langle\})$ -sentence ψ_h which, when evaluated in a disjoint union of tori $\mathcal{T}_1, \dots, \mathcal{T}_\ell$ of height h and width w_1, \dots, w_ℓ , computes $\text{turn}(\mathcal{T}_1, R_1) + \dots + \text{turn}(\mathcal{T}_\ell, R_\ell) \pmod{h}$, for lists of representatives R_1, \dots, R_ℓ of $\mathcal{T}_1, \dots, \mathcal{T}_\ell$ which depend on \langle . The sentence ψ_h is satisfied iff this sum modulo h is 0, i.e., iff $\text{twist}(\mathcal{A}) = 0$. The sentence uses the order \langle only to choose a particular list of representatives $R = (a_0, \dots, a_{w-1})$ for each torus \mathcal{T} of the disjoint union \mathcal{A} by letting a_j be the smallest element of the j -th column. Hence, ψ_h is order-invariant on the class of all disjoint unions of tori of height h . To compute $\text{twist}(\mathcal{A})$, the formula uses a formula $\psi_{\text{on-turn-path}}(x)$ which is satisfied by exactly those elements b in the j -th column, for each $j \in [w]$, of a torus \mathcal{T} that are different from a_j and which lie on the j -turn path w.r.t R . Then, $\text{twist}(\mathcal{T}) = \text{turn}(\mathcal{T}, R)$ is exactly the number, modulo h , of all nodes b of \mathcal{T} for which $\mathcal{T} \models \psi_{\text{on-turn-path}}[b]$. Consequently,

$\text{twist}(\mathcal{A})$ is exactly the number, modulo h , of all nodes b of \mathcal{A} for which $\mathcal{A} \models \psi_{\text{on-turn-path}}[b]$. Thus, we can choose

$$\psi_h := \exists^{0 \bmod h} x \psi_{\text{on-turn-path}}(x).$$

It remains to define the formula $\psi_{\text{on-turn-path}}(x)$. Since the height h is fixed, it is easy to see that there are FO(σ)-formulas $\varphi_{\text{col}}(x, y)$ and $\varphi_{\text{between}}(x, y, z)$, and an FO($\sigma \cup \{<\}$)-formula $\varphi_R(x)$, such that for each torus \mathcal{T} and all elements $a, b, c \in T$ the following holds:

- $\mathcal{T} \models \varphi_{\text{col}}[a, b]$ iff a and b belong to the same column of \mathcal{T} .
- $\mathcal{T} \models \varphi_{\text{between}}[a, b, c]$ iff a, b, c belong to the same column, $b \neq a$, and b lies on the directed $E_1^{\mathcal{T}}$ -path of length at most $h-1$ from a to c .
- $\mathcal{T} \models \varphi_R[a]$ iff a is the smallest element with respect to $<$ in its column.

Now, we can choose

$$\psi_{\text{on-turn-path}}(x) := \exists y \exists z \exists z' (\varphi_R(y) \wedge \varphi_R(z) \wedge E_2(z, z') \wedge \varphi_{\text{between}}(y, x, z')).$$

Finally, we want to construct the sentence φ_h which is order-invariant on the class of all finite σ -structures and which defines the class of disjoint unions of tori with twist 0. It is not difficult to see that a finite σ -structure \mathcal{A} is isomorphic to a disjoint union of tori iff it satisfies the following three properties:

- The relations $E_1^{\mathcal{A}}$ and $E_2^{\mathcal{A}}$ are graphs of permutations π_1 and π_2 of A (i.e. disjoint unions of cycles).
- The cycles of $E_1^{\mathcal{A}}$, which we call *columns*, all have length h .
- The cycles of $E_2^{\mathcal{A}}$ all have length at least 2.
- The π_2 -image of each column \mathcal{C} is a column \mathcal{C}' and the restriction of π_2 to \mathcal{C} is an isomorphism of \mathcal{C} and \mathcal{C}' .

Note that these properties can be expressed by an FO(σ)-sentence θ_h . For the last property, this can be done easily since the length h of the columns is fixed. We let $\varphi_h := \theta_h \wedge \psi_h$. Since ψ_h is order-invariant on disjoint unions of tori and θ_h defines the class of all disjoint unions of tori, φ_h is order-invariant on the class of all finite σ -structures.

3. LOCALITY OF QUERIES

A k -ary query q is a mapping that associates with every finite σ -structure \mathcal{A} a k -ary relation $q(\mathcal{A}) \subseteq A^k$, which is invariant under isomorphisms, i.e., if π is an isomorphism from a σ -structure \mathcal{A} to a σ -structure \mathcal{B} , then for all $\bar{a} = (a_1, \dots, a_k) \in A^k$ we have $\bar{a} \in q(\mathcal{A})$ iff $\pi(\bar{a}) = (\pi(a_1), \dots, \pi(a_k)) \in q(\mathcal{B})$. If \mathfrak{C} is a class of finite σ -structures, then every arb-inv-FO+MOD $_p^{\mathfrak{C}}(\sigma)$ -formula $\varphi(\bar{x})$ with k free variables defines a k -ary query q_φ on \mathfrak{C} via $q_\varphi(\mathcal{A}) = \{\bar{a} \in A^k : \mathcal{A} \models \varphi[\bar{a}]\}$, for every σ -structure $\mathcal{A} \in \mathfrak{C}$.

The *Gaiifman graph* of a σ -structure \mathcal{A} is the undirected graph $\mathcal{G}(\mathcal{A})$ with vertex set A , where for any $a, b \in A$ with $a \neq b$ there is an undirected edge between a and b iff there is an $R \in \sigma$ and a tuple $(a_1, \dots, a_{\text{ar}(R)}) \in R^{\mathcal{A}}$ such that $a, b \in \{a_1, \dots, a_{\text{ar}(R)}\}$. The *distance* $\text{dist}^{\mathcal{A}}(a, b)$ between two elements $a, b \in A$ is the length of a shortest path between a and b in $\mathcal{G}(\mathcal{A})$; in case that there is no path between a and b in $\mathcal{G}(\mathcal{A})$, we let $\text{dist}^{\mathcal{A}}(a, b) := \infty$. The distance $\text{dist}^{\mathcal{A}}(b, \bar{a})$ between an element $b \in A$ and a tuple $\bar{a} = (a_1, \dots, a_k) \in A^k$ is the minimum of $\text{dist}^{\mathcal{A}}(b, a_i)$ for all $i \in \{1, \dots, k\}$. For every $r \in \mathbb{N}$, the r -ball $N_r^{\mathcal{A}}(\bar{a})$ around a tuple $\bar{a} \in A^k$ is the set of all elements b with $\text{dist}^{\mathcal{A}}(b, \bar{a}) \leq r$. The r -neighbourhood of \bar{a} is the induced substructure $\mathcal{N}_r^{\mathcal{A}}(\bar{a})$ of \mathcal{A} on $N_r^{\mathcal{A}}(\bar{a})$.

3.1. Gaifman locality. The notion of *Gaifman locality* provides a standard tool for showing that particular queries are not definable in certain logics (cf., e.g., the textbook [12] for an overview).

Definition 3.1 (Gaifman locality). Let \mathfrak{C} be a class of finite σ -structures, $k \in \mathbb{N}_{\geq 1}$ and $f : \mathbb{N} \rightarrow \mathbb{N}$. A k -ary query q is *Gaifman $f(n)$ -local on \mathfrak{C}* if there is an $n_0 \in \mathbb{N}$ such that for every $n \in \mathbb{N}$ with $n \geq n_0$ and every σ -structure $\mathcal{A} \in \mathfrak{C}$ with $|A| = n$, the following is true for all k -tuples $\bar{a}, \bar{b} \in A^k$ with $(\mathcal{N}_{f(n)}^{\mathcal{A}}(\bar{a}), \bar{a}) \cong (\mathcal{N}_{f(n)}^{\mathcal{A}}(\bar{b}), \bar{b})$: $\bar{a} \in q(\mathcal{A}) \iff \bar{b} \in q(\mathcal{A})$.

The query q is *Gaifman $f(n)$ -local* if it is Gaifman $f(n)$ -local on the class of all finite σ -structures.

That is, in a σ -structure of cardinality n , a query that is Gaifman $f(n)$ -local cannot distinguish between k -tuples of nodes whose neighbourhoods of radius $f(n)$ are isomorphic. The function $f(n)$ is called the *locality radius* of the query. As a simple example, consider the unary query which maps an undirected graph to the set of all vertices which are the central vertex of a star with an even number of rays. Here, a *star* is a connected graph where all but the unique central vertex are rays, i.e. vertices of degree one. This query is Gaifman 2-local on the class of all undirected graphs. As an example of a query which is not Gaifman $f(n)$ -local for any sublinear function f , consider the binary reachability query which maps a graph to all pairs of vertices such that there exists a path from the first vertex to the second vertex.

It is well-known that queries definable in FO or FO+MOD $_p$ (for any $p \geq 2$) are Gaifman local with a constant locality radius [11]. The articles [9] and [1] generalised this to order-invariant FO (for constant locality radius) and arb-invariant FO (for polylogarithmic locality radius) in the following sense: Let $k \in \mathbb{N}_{\geq 1}$, and let q be a k -ary query. If q is definable in $<$ -inv-FO(σ), then there is a $c \in \mathbb{N}$ such that q is Gaifman c -local. If q is definable in arb-inv-FO(σ), then there is a $c \in \mathbb{N}$ such that q is Gaifman $(\log n)^c$ -local. However, for every $d \in \mathbb{N}$ there is a unary query q_d that is definable in arb-inv-FO($\{E\}$) and that is not Gaifman $(\log n)^d$ -local.

Somewhat surprisingly, using Example 2.6 one obtains that the Gaifman locality result cannot be generalised to order- or arb-invariant FO+MOD $_p$. In fact, $<$ -inv-FO+MOD $_p$ can define queries that are not even Gaifman local with locality radius as big as $(\frac{n}{h}-2)$, for the smallest prime divisor h of p :

Proposition 3.2. *Let $h \in \mathbb{N}$ with $h \geq 2$, and let $\sigma = \{R, E_1, E_2\}$ be a signature consisting of a unary relation symbol R and two binary relation symbols E_1, E_2 . There exists a unary query q that is not Gaifman $(\frac{n}{h}-2)$ -local, but definable in $<$ -inv-FO+MOD $_p(\sigma)$, for every multiple $p \geq 2$ of h .*

Proof. Recall the $\{E_1, E_2\}$ -structures from Example 2.6 which we called tori. From a torus \mathcal{T} of height h and width w , we obtain a σ -structure \mathcal{T} by deleting all E_2 -edges from the last to the first column and marking the least element of the last column with a unary relation. We call the structure \mathcal{T} a *hose*. More precisely, the universe of \mathcal{T} is $H := T = [h] \times [w]$ and the relations of \mathcal{T} are

$$\begin{aligned} E_1^{\mathcal{T}} &:= E_1^{\mathcal{T}}, \\ E_2^{\mathcal{T}} &:= E_2^{\mathcal{T}} \setminus \{((i, w-1), (i + \text{twist}(\mathcal{T}) \bmod h, 0)) : i \in [h]\}, \\ R^{\mathcal{T}} &:= \{(0, w-1)\}. \end{aligned}$$




Figure 2: A torus \mathcal{T} (left) and the hose \mathcal{T} (right) obtained from \mathcal{T} . The unique node in the relation $R^{\mathcal{T}}$ of the hose is the rightmost node in the top row, depicted by a red star.

From Example 2.6 we obtain an $\langle\text{-inv-FO+MOD}_h(\{E_1, E_2, \langle\})\rangle$ -sentence $\varphi_{h\text{-torus}}$ which defines the class of all tori \mathcal{T} of height h and $\text{twist}(\mathcal{T}) = 0$ on the class of all finite $\{E_1, E_2\}$ -structures. We modify $\varphi_{h\text{-torus}}$ in such a way that we obtain an $\langle\text{-inv-FO+MOD}_h(\sigma \cup \{\langle\})\rangle$ -formula $\psi(x)$ which, when evaluated in a hose \mathcal{T} with x interpreted as the element $a := (0, 0)$ or $b := (1, 0)$, simulates $\varphi_{h\text{-torus}}$ evaluated on a torus of height h with twist 0 or twist 1, respectively. To this end, we let $\psi(x)$ state that each of the following is satisfied:

- There is a unique element y_0 satisfying $R(y_0)$,
- there are elements y_1, \dots, y_{h-1} such that $E_1(y_i, y_{i+1 \bmod h})$ is true for all $i \in [h]$,
- there are elements x_0, \dots, x_{h-1} such that $x_0 = x$ and $E_1(x_i, x_{i+1 \bmod h})$ is true for all $i \in [h]$,
- the formula φ' is satisfied, where φ' is obtained from $\varphi_{h\text{-torus}}$ by replacing every atom of the form $E_2(u, v)$ by the formula $(E_2(u, v) \vee \bigvee_{0 \leq i < h} (u = y_i \wedge v = x_i))$.

Clearly, $\mathcal{T} \models \psi[a]$ (since $\mathcal{T} \models \varphi_{h\text{-torus}}$ if $\text{twist}(\mathcal{T}) = 0$), and $\mathcal{T} \not\models \psi[b]$ (since $\mathcal{T} \not\models \varphi_{h\text{-torus}}$ if $\text{twist}(\mathcal{T}) = 1$). Thus, $a \in q_\psi(\mathcal{T})$ and $b \notin q_\psi(\mathcal{T})$. Note that the $(w-2)$ -neighbourhoods of a and b in the hose \mathcal{T} are isomorphic, i.e., $(\mathcal{N}_{w-2}^{\mathcal{T}}(a), a) \cong (\mathcal{N}_{w-2}^{\mathcal{T}}(b), b)$. The cardinality of \mathcal{T} is $n := hw$, and hence $w-2 = \frac{n}{h}-2$. Thus, the query defined by $\psi(x)$ is not Gaifman $(\frac{n}{h}-2)$ -local.

By Example 2.6, $\varphi_{h\text{-torus}}$ is order-invariant on the class of all finite $\{E_1, E_2\}$ -structures. Therefore, the formula $\psi(x)$ is order-invariant on the class of all finite σ -structures. Furthermore, $\varphi_{h\text{-torus}}$ uses only modulo counting quantifiers with modulus h and the construction of $\psi(x)$ adds no new modulo counting quantifiers. As we observed in Section 2, if p is a multiple of h , each modulo counting quantifier with modulus h can also be expressed using quantifiers with modulus p . \square

3.2. Weak Gaifman locality. *Weak Gaifman locality* (cf., [12]) is a relaxed notion of Gaifman locality where “ $\bar{a} \in q(\mathcal{A}) \iff \bar{b} \in q(\mathcal{A})$ ” needs to be true only for those tuples \bar{a} and \bar{b} whose $f(n)$ -neighbourhoods are disjoint.

Definition 3.3 (Weak Gaifman locality). Let \mathfrak{C} be a class of finite σ -structures, $k \in \mathbb{N}_{\geq 1}$ and $f : \mathbb{N} \rightarrow \mathbb{N}$. A k -ary query q is *weakly Gaifman $f(n)$ -local on \mathfrak{C}* if there is an $n_0 \in \mathbb{N}$ such that for every $n \in \mathbb{N}$ with $n \geq n_0$ and every σ -structure $\mathcal{A} \in \mathfrak{C}$ with $|\mathcal{A}| = n$, the following is true for all k -tuples $\bar{a}, \bar{b} \in A^k$ with $(\mathcal{N}_{f(n)}^{\mathcal{A}}(\bar{a}), \bar{a}) \cong (\mathcal{N}_{f(n)}^{\mathcal{A}}(\bar{b}), \bar{b})$ and $N_{f(n)}^{\mathcal{A}}(\bar{a}) \cap N_{f(n)}^{\mathcal{A}}(\bar{b}) = \emptyset$: $\bar{a} \in q(\mathcal{A}) \iff \bar{b} \in q(\mathcal{A})$. The query q is *weakly Gaifman $f(n)$ -local* if it is weakly Gaifman $f(n)$ -local on the class of all finite σ -structures.

Note that the example presented in the proof of Proposition 3.2 does not provide a counter-example to *weak Gaifman locality*, since the elements a and b considered in the proof of Proposition 3.2 are of distance 1, and thus their $f(n)$ -neighbourhoods are not disjoint. However, using Example 2.5, one obtains a counter-example to *weak Gaifman locality* for




Figure 3: A string structure \mathcal{S}_{w_ℓ} with $x = a_\ell$ (top) and $x = b_\ell$ (bottom). The formula $\psi(x)$ simulates $\varphi_{\text{even cycles}}$ on the structure where the dashed edges are added and the light nodes and edges are removed.

$<$ -inv-FO+MOD $_p$ for *even* numbers p ; see Proposition 6.23 in [14]. Here, we present a refinement of Niemistö's proof which provides a counter-example to weak Gaifman locality already for the restricted case of string structures.

Proposition 3.4. *Let $\Sigma := \{0, 1\}$, and let $\sigma_\Sigma = \{E, P_0, P_1\}$ be the signature used for representing strings over Σ . There exists a unary query q that is not weakly Gaifman $(\frac{n}{4}-1)$ -local on Σ -strings, but definable in $<$ -inv-FO+MOD $_p(\sigma_\Sigma)$, for every even number $p \geq 2$.*

Proof. For every $\ell \in \mathbb{N}_{\geq 1}$, let \mathcal{A}_ℓ and \mathcal{B}_ℓ be $\{E\}$ -structures whose universe consists of 2ℓ vertices, the edge relation of \mathcal{A}_ℓ consists of two directed cycles of length ℓ , and the edge relation of \mathcal{B}_ℓ consists of a single directed cycle of length 2ℓ . Furthermore, we choose w_ℓ to be the string $1^\ell 0^\ell 1^\ell 0^\ell$, and we let $a_\ell := \ell$ be the rightmost position of the first block of 1s, and $b_\ell := 3\ell$ the rightmost position of the second block of 1s.

From Example 2.5 we obtain an $<$ -inv-FO+MOD $_2(\{E\})$ -sentence $\varphi_{\text{even cycles}}$ that is satisfied by a finite $\{E\}$ -structure \mathcal{A} iff \mathcal{A} is a disjoint union of directed cycles where the number of cycles of even length is even. Thus, for every $\ell \in \mathbb{N}_{\geq 1}$ we have: $\mathcal{A}_\ell \models \varphi_{\text{even cycles}}$ and $\mathcal{B}_\ell \not\models \varphi_{\text{even cycles}}$. We modify the formula $\varphi_{\text{even cycles}}$ in such a way that we obtain an $<$ -inv-FO+MOD $_2(\sigma_\Sigma)$ -formula $\psi(x)$ which, when evaluated in the σ_Σ -structure \mathcal{S}_{w_ℓ} representing the string w_ℓ with x interpreted as the position a_ℓ or b_ℓ simulates $\varphi_{\text{even cycles}}$ evaluated on \mathcal{A}_ℓ or \mathcal{B}_ℓ , respectively. To this end, we let $\psi(x)$ be a formula stating that each of the following is satisfied:

- There is a unique position $x' \neq x$ that carries the letter 1 such that the position directly to the right of x' carries the letter 0.
- There is a unique position y of in-degree 0, and this position carries the letter 1. Furthermore, there is a unique position y' that carries the letter 1, such that the position directly to the left of y' carries the letter 0.
- The formula φ' is satisfied, where φ' is obtained from $\varphi_{\text{even cycles}}$ by relativisation of all quantifiers to positions that carry the letter 1, and by replacing every atom of the form $E(u, v)$ by the formula $(E(u, v) \vee (u=x \wedge v=y) \vee (u=x' \wedge v=y'))$.

Clearly, for every $\ell \in \mathbb{N}_{\geq 1}$ we have: $\mathcal{S}_{w_\ell} \models \psi[a_\ell]$ (since $\mathcal{A}_\ell \models \varphi_{\text{even cycles}}$), and $\mathcal{S}_{w_\ell} \not\models \psi[b_\ell]$ (since $\mathcal{B}_\ell \not\models \varphi_{\text{even cycles}}$). Thus, $a_\ell \in q_\psi(\mathcal{S}_{w_\ell})$ and $b_\ell \notin q_\psi(\mathcal{S}_{w_\ell})$. Note that the $(\ell-1)$ -neighbourhoods of a_ℓ and b_ℓ in \mathcal{S}_{w_ℓ} are disjoint and isomorphic. The cardinality of \mathcal{S}_{w_ℓ} is

$n := 4\ell$, and hence $\ell - 1 = \frac{n}{4} - 1$. Thus, the query defined by $\psi(x)$ is not weakly Gaifman $(\frac{n}{4} - 1)$ -local. Since $\varphi_{\text{even cycles}}$ is order-invariant on all finite $\{E\}$ -structures, the formula $\psi(x)$ is order-invariant on the class of all finite σ_Σ -structures. Note that $\psi(x)$ is definable in $<$ -inv-FO+MOD $_2(\sigma_\Sigma)$, and hence also in $<$ -inv-FO+MOD $_p(\sigma_\Sigma)$, for every multiple p of 2. \square

In light of Proposition 3.4 it is somewhat surprising that for *odd* numbers p , unary queries definable in $<$ -inv-FO+MOD $_p$ are weakly Gaifman local with constant locality radius — this is a result obtained by Niemistö (see Corollary 6.37 in [14]). For *odd prime powers* p we can generalise this to k -ary queries definable in arb-inv-FO+MOD $_p$, when allowing polylogarithmic locality radius. Note that we cannot hope for a smaller locality radius, since [1] provides, for every $d \in \mathbb{N}$, a unary query definable in arb-inv-FO($\{E\}$) that is not weakly Gaifman $(\log n)^d$ -local. Precisely, we will show the following.

Theorem 3.5. *Let \mathfrak{C} be a class of finite σ -structures. Let $k \in \mathbb{N}_{\geq 1}$, let q be a k -ary query, and let p be an odd prime power. If q is definable in arb-inv-FO+MOD $_p^{\mathfrak{C}}(\sigma)$ on \mathfrak{C} , then there is a $c \in \mathbb{N}$ such that q is weakly Gaifman $(\log n)^c$ -local on \mathfrak{C} .*

The proof of this theorem will be given in the next subsection, as an easy consequence of Theorem 3.7 below. A generalisation of Theorem 3.5 from odd prime powers to arbitrary odd numbers p would lead to new separations concerning circuit complexity classes and can therefore be expected to be rather difficult (see Remark 3.13).

3.3. Shift locality.

The following notion of *shift locality* generalises the notion of *alternating Gaifman locality* introduced by Niemistö in [14]. In some sense discussed below, it unifies this notion and the notion of weak Gaifman locality.

Definition 3.6 (Shift locality). Let \mathfrak{C} be a class of finite σ -structures. Let $k, t \in \mathbb{N}_{\geq 1}$ with $t \geq 2$, and let $f : \mathbb{N} \rightarrow \mathbb{N}$. A kt -ary query q is *shift $f(n)$ -local w.r.t. t* on \mathfrak{C} if there is an $n_0 \in \mathbb{N}$ such that for every $n \in \mathbb{N}$ with $n \geq n_0$ and every σ -structure $\mathcal{A} \in \mathfrak{C}$ with $|\mathcal{A}| = n$, the following is true for all k -tuples $\bar{a}_0, \dots, \bar{a}_{t-1} \in A^k$ with $(\mathcal{N}_{f(n)}^{\mathcal{A}}(\bar{a}_i), \bar{a}_i) \cong (\mathcal{N}_{f(n)}^{\mathcal{A}}(\bar{a}_j), \bar{a}_j)$ and $N_{f(n)}^{\mathcal{A}}(\bar{a}_i) \cap N_{f(n)}^{\mathcal{A}}(\bar{a}_j) = \emptyset$ for all $i, j \in [t]$ with $i \neq j$: $(\bar{a}_0, \bar{a}_1, \dots, \bar{a}_{t-1}) \in q(\mathcal{A}) \iff (\bar{a}_1, \dots, \bar{a}_{t-1}, \bar{a}_0) \in q(\mathcal{A})$.

Query q is *shift $f(n)$ -local w.r.t. t* if it is shift $f(n)$ -local w.r.t. t on the class of all finite σ -structures.

The case of $k = 1$ and $t = 3$ for a constant function f yields Niemistö's notion of alternating Gaifman locality [14]. From query q of Proposition 3.4, we obtain an example of an alternatingly Gaifman local query. For all even numbers $p \geq 2$, this query q is definable by an $<$ -inv-FO+MOD $_p$ -formula $\psi(x)$. The query defined by the formula $\varphi(x, y, z) := \psi(x) \wedge y = y \wedge z = z$ is an alternatingly Gaifman local query. This can be seen directly, but it is also a consequence of the result of [14] that queries which are definable by formulas of $<$ -inv-FO+MOD $_p$ for even numbers p are alternatingly Gaifman local. Several examples of non-shift locality will be given in Section 3.4. To understand the relation between shift locality and weak Gaifman locality, consider a k -ary query q and the $2k$ -ary query \tilde{q} with $\tilde{q}(\mathcal{A}) := \{\bar{a}_1 \bar{a}_2 : \bar{a}_1 \in q(\mathcal{A}), \bar{a}_2 \in A^k\}$. Then q is weakly Gaifman $f(n)$ -local iff \tilde{q} is shift $f(n)$ -local w.r.t. 2. The notion of shift locality helps to discuss both kinds of locality in a uniform way.

In a technical lemma (Lemma 6.36 in [14]), Niemistö showed that for $k = 1$ and $p, t \in \mathbb{N}$ with $p, t \geq 2$ and p and t coprime, for every t -ary query q definable in $<$ -inv-FO+MOD $_p(\sigma)$, there is a $c \in \mathbb{N}$ such that q is shift c -local w.r.t. t . Our next result deals with the general case of shift locality and the more expressive logic arb-inv-FO+MOD $_p(\sigma)$, when allowing polylogarithmic locality radius.

Theorem 3.7. *Let \mathfrak{C} be a class of finite σ -structures. Let $k, t \in \mathbb{N}_{\geq 1}$ with $t \geq 2$, let q be a kt -ary query, and let p be a prime power such that p and t are coprime. If q is definable in arb-inv-FO+MOD $_p^{\mathfrak{C}}(\sigma)$ on \mathfrak{C} , then there is a $c \in \mathbb{N}$ such that q is shift $(\log n)^c$ -local w.r.t. t on \mathfrak{C} .*

Our proof of Theorem 3.7 relies on lower bounds achieved in circuit complexity. A generalisation of Theorem 3.7 from prime powers to arbitrary numbers $p \geq 2$ would lead to new separations of circuit complexity classes and can therefore be expected to be rather difficult (see Remark 3.13). Before giving the proof of Theorem 3.7, let us first point out that it immediately implies Theorem 3.5.

Proof of Theorem 3.5 (using Theorem 3.7). Let $\varphi(\bar{x})$ be an arb-inv-FO+MOD $_p^{\mathfrak{C}}(\sigma)$ -formula with k free variables $\bar{x} = (x_1, \dots, x_k)$, defining a k -ary query q_φ on \mathfrak{C} . Let $\bar{y} = (y_1, \dots, y_k)$ be k variables different from the variables in \bar{x} . Then, $\psi(\bar{x}, \bar{y}) := (\varphi(\bar{x}) \wedge \bigwedge_{1 \leq i \leq k} y_i = x_i)$ is an arb-inv-FO+MOD $_p^{\mathfrak{C}}(\sigma)$ -formula that defines a $2k$ -ary query q_ψ . By Theorem 3.7, there exists a $c \in \mathbb{N}$ such that q_ψ is shift $(\log n)^c$ -local w.r.t. $t := 2$ on \mathfrak{C} . It is straightforward to see that the shift $(\log n)^c$ -locality of q_ψ w.r.t. $t = 2$ implies that the query q_φ is weakly Gaifman $(\log n)^c$ -local. \square

The remainder of this subsection is devoted to the proof of Theorem 3.7. We follow the overall method of [1] for the case of disjoint neighbourhoods (see [16] for an overview) and make use of the connection between arb-inv-FO+MOD $_p$ and MOD $_p$ -circuits [3], along with a circuit lower bound by Smolensky [17].

We assume that the reader is familiar with basic notions and results in circuit complexity (cf., e.g., the textbook [2]). A MOD $_p$ -gate returns the value 1 iff the number of ones at its inputs is congruent 0 modulo p . We consider Boolean MOD $_p$ -circuits consisting of AND-, OR-, and MOD $_p$ -gates of unbounded fan-in, input gates, negated input gates, and constant gates $\mathbf{0}$ and $\mathbf{1}$. More precisely, a MOD $_p$ -circuit with m input bits is a directed acyclic graph whose vertices without ingoing edges are called *input gates* and are labelled with either $\mathbf{0}$, $\mathbf{1}$, w_ν , or $\neg w_\nu$ for $\nu \in \{1, \dots, m\}$, whose internal nodes are called *gates* and are labelled either AND or OR or MOD $_p$, and which has a distinguished vertex without outgoing edges called the *output gate*. A MOD $_p$ -circuit C with m input bits naturally defines a function from $\{0, 1\}^m$ to $\{0, 1\}$. For an input string $w \in \{0, 1\}^m$ we say that C *accepts* w if $C(w) = 1$. Accordingly, C *rejects* w if $C(w) = 0$. The *size* of a circuit is the number of gates it contains, and the *depth* is the length of the longest path from any of the input gates to the output gate.

Our proof of Theorem 3.7 relies on Smolensky’s following circuit lower bound.

Theorem 3.8 (Smolensky [17] (see also [18])). *Let p be a prime power and let r be a number which has a prime factor which is different from the prime factor of p .*

There exist numbers $\varepsilon, \ell > 0$ such that for every $d \in \mathbb{N}_{\geq 1}$ there is an $m_d \in \mathbb{N}_{\geq 1}$ such that for every $m \in \mathbb{N}$ with $m \geq m_d$ the following is true: No MOD $_p$ -circuit of depth d and size at most $2^{\varepsilon \ell \sqrt{m}}$ accepts exactly those bitstrings $w \in \{0, 1\}^m$ that contain a number of ones congruent 0 modulo r .

In the literature, Smolensky's theorem is usually stated only for primes p . Note, however, that (for each fixed $k \in \mathbb{N}_{\geq 1}$) MOD_p^k -gates can easily be simulated by MOD_p -circuits of constant depth and polynomial size (cf., [18]), and hence Smolensky's theorem also holds for prime powers p , as stated in Theorem 3.8. It is still open whether an analogous result also holds for numbers p composed of more than one prime factor (see Chapter VIII of [18] and Chapter 14.4 of [2] for discussions on this).

To establish the connection between MOD_p -circuits and $\text{arb-inv-FO+MOD}_p(\sigma)$, we need to represent σ -structures \mathcal{A} and K -tuples $\bar{a} \in A^K$ (for $K \in \mathbb{N}$) by bitstrings. This is done in a straightforward way: Let $\sigma = \{R_1, \dots, R_{|\sigma|}\}$ and let $r_i := ar(R_i)$ for each $i \leq |\sigma|$. Consider a finite σ -structure \mathcal{A} with $|A| = n$. Let ι be an embedding of \mathcal{A} into $[n]$. For each $R_i \in \sigma$ we let $\text{Rep}^\iota(R_i^{\mathcal{A}})$ be the bitstring of length n^{r_i} whose j -th bit is 1 iff the j -th smallest element in A^{r_i} w.r.t. the lexicographic order associated with $<^\iota$ belongs to the relation $R_i^{\mathcal{A}}$. Similarly, for each component a_i of a K -tuple $\bar{a} = (a_1, \dots, a_K) \in A^K$ we let $\text{Rep}^\iota(a_i)$ be the bitstring of length n whose j -th bit is 1 iff a_i is the j -th smallest element of A w.r.t. $<^\iota$. Finally, we let

$$\text{Rep}^\iota(\mathcal{A}, \bar{a}) := \text{Rep}^\iota(R_1^{\mathcal{A}}) \cdots \text{Rep}^\iota(R_{|\sigma|}^{\mathcal{A}}) \text{Rep}^\iota(a_1) \cdots \text{Rep}^\iota(a_K)$$

be the *binary representation of (\mathcal{A}, \bar{a}) w.r.t. ι* . Note that, independently of ι , the length of the bitstring $\text{Rep}^\iota(\mathcal{A}, \bar{a})$ is $\lambda_K^\sigma(n) := \sum_{i=1}^{|\sigma|} n^{r_i} + Kn$.

The connection between $\text{FO+MOD}_p(\sigma_{\text{arb}})$ and MOD_p -circuits is obtained by the following result.

Theorem 3.9 (implicit in [3] (see also [18])). *Let σ be a finite relational signature, let $K \in \mathbb{N}$, and let $p \in \mathbb{N}$ with $p \geq 2$. For every $\text{FO+MOD}_p(\sigma_{\text{arb}})$ -formula $\varphi(\bar{x})$ with K free variables there exist numbers $d, s \in \mathbb{N}$ such that for every $n \in \mathbb{N}_{\geq 1}$ there is a MOD_p -circuit C_n with $\lambda_K^\sigma(n)$ input bits, depth d , and size n^s such that the following is true for all σ -structures \mathcal{A} with $|A| = n$, all $\bar{a} \in A^K$, and all embeddings ι of \mathcal{A} into $[n]$: C_n accepts $\text{Rep}^\iota(\mathcal{A}, \bar{a}) \iff \mathcal{A}^\iota \models \varphi[\bar{a}]$.*

The proof of Theorem 3.7 can be outlined informally as follows. Assume that, for some prime power p and a number t which is coprime with p , there is an $\text{arb-inv-FO+MOD}_p(\sigma)$ -definable kt -ary query q which is not shift $(\log n)^c$ -local w.r.t t for each c . We wish to obtain a contradiction to Smolensky's theorem (Theorem 3.8). From the defining arb-inv -invariant sentence for q , we obtain a MOD_p -circuit-family $(C_n)_{n \in \mathbb{N}}$ which accepts exactly the bitstrings which are representations of structures (\mathcal{A}, \bar{a}) such that $\bar{a} \in q(\mathcal{A})$ (cf. Theorem 3.9). Since q is not shift $(\log n)^c$ -local w.r.t t , for each c , we obtain an infinite family of counter examples to the shift $(\log n)^c$ -locality of q . That is, for each n_0 there is a structure \mathcal{A} on $n \geq n_0$ elements which contains tuples $\bar{a}_0, \dots, \bar{a}_{t-1}$ of k elements each such that the tuples have pairwise isomorphic and disjoint $(\log n)^c$ -neighbourhoods, but $(\bar{a}_0, \dots, \bar{a}_{t-1}) \in q(\mathcal{A})$ and $(\bar{a}_1, \dots, \bar{a}_{t-1}, \bar{a}_0) \notin q(\mathcal{A})$, i.e. C_n accepts all representations of $(\mathcal{A}, \bar{a}_0, \dots, \bar{a}_{t-1})$ and rejects all representations of $(\mathcal{A}, \bar{a}_1, \dots, \bar{a}_{t-1}, \bar{a}_0)$.

We transform each C_n into a circuit \tilde{C}_m (cf. Lemma 3.10 below) of exactly the same size which accepts all bitstrings where the number of ones is 0 modulo t and rejects all bitstrings where this number is 1 modulo t ; furthermore, it does not distinguish between strings with the same number of ones modulo t . To this end, on input of a word w , the circuit \tilde{C}_m simulates the circuit C_n on the representation of a structure which is obtained from \mathcal{A} as follows. For each position i of w which is labelled by 1, the relations between elements at distance i and $i+1$ from the tuple $(\bar{a}_0, \dots, \bar{a}_{t-1})$ are changed in such a way

that the resulting structure looks like \mathcal{A} where the tuple $(\bar{a}_0, \dots, \bar{a}_{t-1})$ has been shifted. See Figure 4 for an example where the structure \mathcal{A} is a graph. This will be done in a way which ensures that \tilde{C}_m accepts w if the number of ones is 0 modulo t — in this case, the simulated structure looks like $(\mathcal{A}, \bar{a}_0, \dots, \bar{a}_{t-1})$ and hence its representation is accepted by C_n — and that it rejects w if the number of ones is 1 modulo t — in this case, the simulated structure looks like $(\mathcal{A}, \bar{a}_1, \dots, \bar{a}_{t-1}, \bar{a}_0)$ and hence it is rejected by C_n .

This construction is not yet sufficient for a contradiction to Smolenky's theorem. For this, we need to show (cf. Lemma 3.11 below) that the circuit \tilde{C}_m can be transformed into a circuit \hat{C}_m of roughly the same size and depth which accepts *exactly* the bitstrings with 0 ones modulo r , for some factor $r \geq 2$ of t . To this end, we show that there exists a factor $r \geq 2$ of t such that, basically, it is possible to determine if the number of ones in a bitstring w is 0 modulo r by computing, for each $j \in [t]$, whether \tilde{C}_m accepts w after replacing the first j zeros of w by ones. Having achieved all this, if we fix c and n_0 appropriately in terms of Smolenky's theorem, we obtain the desired contradiction.

We proceed with the formal proof of Theorem 3.7. To simplify notation, we define $\vec{a}^{(0)} := (\bar{a}_0, \bar{a}_1, \dots, \bar{a}_{t-1})$ and $\vec{a}^{(i)} := (\bar{a}_i, \bar{a}_{i+1}, \dots, \bar{a}_{t-1}, \bar{a}_0, \bar{a}_1, \dots, \bar{a}_{i-1})$, for all $i \in [t]$ with $i \leq 1$.

Lemma 3.10. *Let $m, k, t \in \mathbb{N}_{\geq 1}$ with $t \geq 2$. Let \mathcal{A} be a finite σ -structure with $n := |\mathcal{A}|$. For each $i \in [t]$ let $\bar{a}_i \in A^k$ such that for all $i, j \in [t]$ with $i \neq j$ we have $(\mathcal{N}_m^{\mathcal{A}}(\bar{a}_i), \bar{a}_i) \cong (\mathcal{N}_m^{\mathcal{A}}(\bar{a}_j), \bar{a}_j)$ and $N_m^{\mathcal{A}}(\bar{a}_i) \cap N_m^{\mathcal{A}}(\bar{a}_j) = \emptyset$. Let $p \in \mathbb{N}$ with $p \geq 2$. Let C be a MOD_p -circuit with $\lambda_{kt}^{\sigma}(n)$ input bits such that:*

- (1) C accepts $\text{Rep}^{\iota_1}(\mathcal{A}, \vec{a}^{(i)})$ iff it accepts $\text{Rep}^{\iota_2}(\mathcal{A}, \vec{a}^{(i)})$, for all embeddings ι_1 and ι_2 of \mathcal{A} and for every $i \in [t]$, and
- (2) C accepts $\text{Rep}^{\iota}(\mathcal{A}, \vec{a}^{(0)})$ and rejects $\text{Rep}^{\iota}(\mathcal{A}, \vec{a}^{(1)})$, for every embedding ι of \mathcal{A} .

There exists a MOD_p -circuit \tilde{C} with m input bits, such that:

- (a) \tilde{C} has the same depth and size as C ,
- (b) for all $w, w' \in \{0, 1\}^m$ with $|w|_1 \equiv |w'|_1 \pmod{t}$, \tilde{C} accepts w iff it accepts w' , and
- (c) \tilde{C} accepts all $w \in \{0, 1\}^m$ with $|w|_1 \equiv 0 \pmod{t}$ and rejects all $w \in \{0, 1\}^m$ with $|w|_1 \equiv 1 \pmod{t}$.

Proof. Let $I \subset [t]$ be the set containing $i \in [t]$ iff C accepts $\text{Rep}^{\iota_1}(\mathcal{A}, \vec{a}^{(i)})$ for some (i.e., due to property (1) of C , every) embedding ι_1 of \mathcal{A} . By property (2) of C , we know that $0 \in I$ and $1 \notin I$.

For the remainder of this proof, fix an embedding ι of \mathcal{A} into $[n]$. Note that ι is also an embedding of any other σ -structure that has the same universe as \mathcal{A} . For every $w \in \{0, 1\}^m$, we will define a σ -structure \mathcal{A}_w with the same universe as \mathcal{A} , which has the following property for every $i \in [t]$:

$$\text{If } |w|_1 \equiv i \pmod{t}, \text{ then } (\mathcal{A}_w, \vec{a}^{(0)}) \cong (\mathcal{A}, \vec{a}^{(i)}). \quad (3.1)$$

Note that if $(\mathcal{A}_w, \vec{a}^{(0)}) \cong (\mathcal{A}, \vec{a}^{(i)})$, then there is an embedding ι_1 such that $\text{Rep}^{\iota_1}(\mathcal{A}, \vec{a}^{(i)}) = \text{Rep}^{\iota}(\mathcal{A}_w, \vec{a}^{(0)})$. Hence, due to property (1), C accepts $\text{Rep}^{\iota}(\mathcal{A}_w, \vec{a}^{(0)})$ iff it accepts $\text{Rep}^{\iota}(\mathcal{A}, \vec{a}^{(i)})$.

The circuit \tilde{C} will be constructed so that on input $w \in \{0, 1\}^m$ it does the same as circuit C does on input $\text{Rep}^{\iota}(\mathcal{A}_w, \vec{a}^{(0)})$. Thus, the following will be true for every $w \in \{0, 1\}^m$ and the particular number $i \in [t]$ such that $|w|_1 \equiv i \pmod{t}$:

$$\tilde{C} \text{ accepts } w \iff C \text{ accepts } \text{Rep}^{\iota}(\mathcal{A}_w, \vec{a}^{(0)}) \iff C \text{ accepts } \text{Rep}^{\iota}(\mathcal{A}, \vec{a}^{(i)}) \iff i \in I.$$




Figure 4: Illustration of a structure \mathcal{A} (left) and structures \mathcal{A}_w for two different bitstrings w .

This immediately implies that \tilde{C} satisfies property (b); and since $0 \in I$ and $1 \notin I$, the circuit \tilde{C} also satisfies property (c).

Definition of \mathcal{A}_w : For each $j \in [t]$, we partition $N_m^{\mathcal{A}}(\bar{a}_j)$ into shells $S_\nu(\bar{a}_j) := \{x \in A : \text{dist}^{\mathcal{A}}(x, \bar{a}_j) = \nu\}$, for all $\nu \in \{0, \dots, m\}$. We write S_ν for the set $S_\nu(\bar{a}_0) \cup \dots \cup S_\nu(\bar{a}_{t-1})$. For each $j \in [t]$ let π_j be an isomorphism from $(\mathcal{N}_m^{\mathcal{A}}(\bar{a}_j), \bar{a}_j)$ to $(\mathcal{N}_m^{\mathcal{A}}(\bar{a}_{(j+1 \bmod t)}), \bar{a}_{(j+1 \bmod t)})$. Note that $\pi_j(S_\nu(\bar{a}_j)) = S_\nu(\bar{a}_{(j+1 \bmod t)})$ for each $j \in [t]$ and each $\nu \leq m$.

For a bitstring $w = w_1 \dots w_m \in \{0, 1\}^m$ the structure \mathcal{A}_w has the same universe as \mathcal{A} . For each $R \in \sigma$ of arity r , the relation $R^{\mathcal{A}_w}$ is obtained from $R^{\mathcal{A}}$ as follows: We start with $R^{\mathcal{A}_w} := \emptyset$, and then for each tuple $\bar{c} \in R^{\mathcal{A}}$ we insert the tuple \bar{c}_w into $R^{\mathcal{A}_w}$, where \bar{c}_w is defined as follows:

- (i) If $\bar{c} \notin (S_{\nu-1} \cup S_\nu)^r$ for any $\nu \leq m$, or $\bar{c} \in S_\nu^r$ for some $\nu \leq m$, then $\bar{c}_w := \bar{c}$.
- (ii) Otherwise, if $\bar{c} \in (S_{\nu-1} \cup S_\nu)^r$ for some $\nu \leq m$, then note that (since $\bar{c} \in R^{\mathcal{A}}$), there is a unique $j \in [t]$ such that $\bar{c} \in (S_{\nu-1}(\bar{a}_j) \cup S_\nu(\bar{a}_j))^r$ (since $N_m^{\mathcal{A}}(\bar{a}_j) \cap N_m^{\mathcal{A}}(\bar{a}_{j'}) = \emptyset$, for all $j, j' \in [t]$ with $j \neq j'$). To keep the notation simple, assume that $\bar{c} = (\bar{c}_{\nu-1}, \bar{c}_\nu)$, where all elements of $\bar{c}_{\nu-1}$ belong to $S_{\nu-1}(\bar{a}_j)$ and all elements of \bar{c}_ν belong to $S_\nu(\bar{a}_j)$. We define \bar{c}_w depending on the ν -th bit w_ν of w : If $w_\nu = 0$, then $\bar{c}_w := \bar{c}$. If $w_\nu = 1$, then $\bar{c}_w := (\bar{c}_{\nu-1}, \pi_j(\bar{c}_\nu))$.

Note that for every $\nu \in \{1, \dots, m\}$ with $w_\nu = 1$, this construction enforces that the role that was formerly played by shell $S_\nu(\bar{a}_j)$ is afterwards played by shell $S_\nu(\bar{a}_{(j+1 \bmod t)})$; see Figure 4 for an illustration. This shows that \mathcal{A}_w satisfies the property (3.1). A more precise argument will be given at the end of the proof.

Construction of \tilde{C} : On input of $w \in \{0, 1\}^m$ the circuit \tilde{C} will simulate circuit C on input $\text{Rep}^t(\mathcal{A}_w, \vec{a}^{(0)})$. We construct \tilde{C} in a way which mirrors the construction of \mathcal{A}_w . To this end, all inputs of C corresponding, in $\text{Rep}^t(\mathcal{A}, \vec{a}^{(0)})$, to tuples in relations of \mathcal{A} that are unchanged by the construction of \mathcal{A}_w (i.e. tuples to which case (i) of the definition of \mathcal{A}_w applies) are fixed to constant values. Inputs of C corresponding to tuples to which case (ii) of the definition of \mathcal{A}_w applies (i.e. tuples $\bar{c} \in (S_{\nu-1} \cup S_\nu)^r$ for some $\nu \leq m$) are changed according to the ν -th bit in the input w of \tilde{C} in a way that mirrors case (ii).

More precisely, for each relation symbol $R \in \sigma$ of arity r and each tuple $\bar{c} \in R^{\mathcal{A}}$, we proceed as follows. Let g and $\neg g$ be the non-negated and the negated input gate of C that correspond to the bit b in $\text{Rep}^t(\mathcal{A}, \vec{a}^{(0)})$ which represents the information whether \bar{c} belongs to $R^{\mathcal{A}}$.

- (i) If $\bar{c} \notin (S_{\nu-1} \cup S_\nu)^r$ for all $\nu \leq m$, or $\bar{c} \in S_\nu^r$ for some $\nu \leq m$, the gate g is replaced by the constant gate $\mathbf{1}$, and the negated input gate $\neg g$ is replaced by the constant gate $\mathbf{0}$.
- (ii) Otherwise, if $\bar{c} \in (S_{\nu-1}(\bar{a}_j) \cup S_\nu(\bar{a}_j))^r$ for some $\nu \leq m$ and $j \in [t]$, assume, for simplicity, that $\bar{c} := (\bar{c}_{\nu-1}, \bar{c}_\nu)$, where all elements of $\bar{c}_{\nu-1}$ belong to $S_{\nu-1}(\bar{a}_j)$ and

all elements of \bar{c}_ν belong to $S_\nu(\bar{a}_j)$. Let g' and $\neg g'$ be the non-negated and the negated input gate of C that correspond to the bit b' in $\text{Rep}'(\mathcal{A}, \bar{a}^{(0)})$ representing the information whether $(\bar{c}_{\nu-1}, \pi_j(\bar{c}_\nu))$ belongs to $R^{\mathcal{A}}$.

- We replace gate g by the new input gate $\neg w_\nu$, and gate g' by the input gate w_ν .
- Accordingly, the negated input gates $\neg g$ and $\neg g'$ are replaced by the input gates w_ν and $\neg w_\nu$.

Afterwards, we consider all non-negated input gates g of C that have not yet been replaced, and we let b be the bit of $\text{Rep}'(\mathcal{A}, \bar{a}^{(0)})$ that is inserted at input gate g . We replace g by the constant gate $\mathbf{1}$ if $b = 1$, and by the constant gate $\mathbf{0}$ if $b = 0$. Accordingly, the negated input gate $\neg g$ is replaced by $\mathbf{0}$ if $b = 1$, and by $\mathbf{1}$ if $b = 0$.

It is easy to see that the resulting circuit \tilde{C} does the same on input $w \in \{0, 1\}^m$ as circuit C does on input $\text{Rep}'(\mathcal{A}_w, \bar{a}^{(0)})$. Furthermore, \tilde{C} obviously has the same depth as C , and the size of \tilde{C} is smaller than or equal to the size of C .

To finish the proof of Lemma 3.10, it remains to show that \mathcal{A}_w satisfies the property (3.1). For each $\nu \leq m$, we let $i_\nu := |w_1 \dots w_\nu|_1 \bmod t$. In particular, $i_0 = 0$. For each map f and each subset X of its domain, we let $f \upharpoonright X$ denote the restriction of f to X . By id_X we denote the identity map on a set X . For each $\nu \leq m$, we let $\text{id}_\nu := \text{id}_{S_\nu^{\mathcal{A}_w}(\bar{a}^{(0)})}$. Here we have extended our notation for shells which was defined above with respect to \mathcal{A} to \mathcal{A}_w . We write $g \circ f$ for the composition of maps $f : X \rightarrow Y$ and $g : Y \rightarrow Z$. Recall the maps π_1, \dots, π_t from the definition of \mathcal{A}_w and observe that, since these maps have disjoint domains and images, their union π is a well-defined map, and that this map is an isomorphism from $(\mathcal{N}_m^{\mathcal{A}}(\bar{a}^{(i)}), \bar{a}^{(i)})$ to $(\mathcal{N}_m^{\mathcal{A}}(\bar{a}^{(i+1 \bmod t)}), \bar{a}^{(i+1 \bmod t)})$, for each $i \in [t]$. For each $\nu \leq m$, we let $\gamma_\nu := \pi \upharpoonright N_\nu^{\mathcal{A}}(\bar{a}^{(0)})$ (which is the same as $\pi \upharpoonright N_\nu^{\mathcal{A}}(\bar{a}^{(i)})$, for each $i \in [t]$).

By induction on $\nu \leq m$, we prove that $N_\nu^{\mathcal{A}_w}(\bar{a}^{(0)}) = N_\nu^{\mathcal{A}}(\bar{a}^{(i_\nu)})$ and that there exists a map ρ_ν such that

$$\rho_\nu : (\mathcal{N}_\nu^{\mathcal{A}_w}(\bar{a}^{(0)}), \bar{a}^{(0)}) \cong (\mathcal{N}_\nu^{\mathcal{A}}(\bar{a}^{(i_\nu)}), \bar{a}^{(i_\nu)}). \quad (3.2)$$

For $\nu = 0$, we let $\rho_0 := \text{id}_0$. If $1 \leq \nu \leq m$, we let

$$\rho_\nu := \begin{cases} \rho_{\nu-1} \cup \text{id}_\nu & \text{if } w_\nu = 0, \\ (\gamma_{\nu-1} \circ \rho_{\nu-1}) \cup \text{id}_\nu & \text{if } w_\nu = 1. \end{cases}$$

Before we show that ρ_ν satisfies (3.2), for each $\nu \leq m$, we show how to obtain the isomorphism ρ from $(\mathcal{A}_w, \bar{a}^{(0)})$ to $(\mathcal{A}, \bar{a}^{(i_m)})$ required by (3.1) if ρ_ν satisfies (3.2). We let $\rho := \rho_m \cup \text{id}_{A \setminus N_m^{\mathcal{A}_w}(\bar{a}^{(0)})}$, which is a bijection since ρ_m is a bijection from $N_m^{\mathcal{A}_w}(\bar{a}^{(0)})$ to $N_m^{\mathcal{A}}(\bar{a}^{(i_m)}) = N_m^{\mathcal{A}_w}(\bar{a}^{(0)})$. According to (3.2), the restriction of ρ to $N_m^{\mathcal{A}_w}(\bar{a}^{(0)})$ is an isomorphism. The restriction of ρ to $A \setminus N_m^{\mathcal{A}_w}(\bar{a}^{(0)})$ is an isomorphism since the definition of \mathcal{A}_w (according to its case (i)) includes a tuple of elements which all have distance greater than m to $\bar{a}^{(0)}$ in a relation of \mathcal{A}_w iff it is included in the corresponding relation of \mathcal{A} . To verify that ρ is an isomorphism from $(\mathcal{A}_w, \bar{a}^{(0)})$ to $(\mathcal{A}, \bar{a}^{(i_m)})$, it suffices to show that it preserves the relations between elements from $S_m^{\mathcal{A}_w}$ and $S_{m+1}^{\mathcal{A}_w}$. This follows since, restricted to $S_m^{\mathcal{A}_w} \cup S_{m+1}^{\mathcal{A}_w}$, the map ρ is the identity.

Now we prove the construction of ρ_ν correct. For $\nu = 0$, we have $i_0 = 0$ and from the construction of \mathcal{A}_w , we see that $(\mathcal{N}_0^{\mathcal{A}_w}(\bar{a}^{(0)}), \bar{a}^{(0)}) = (\mathcal{N}_0^{\mathcal{A}}(\bar{a}^{(0)}), \bar{a}^{(0)})$. Hence, $\text{id}_0 : (\mathcal{N}_0^{\mathcal{A}_w}(\bar{a}^{(0)}), \bar{a}^{(0)}) \cong (\mathcal{N}_0^{\mathcal{A}}(\bar{a}^{(0)}), \bar{a}^{(0)})$. Suppose now that $\nu \geq 1$. First, we consider the case that $w_\nu = 0$ and hence $i_\nu = i_{\nu-1}$. By induction, $N_{\nu-1}^{\mathcal{A}_w}(\bar{a}^{(0)}) = N_{\nu-1}^{\mathcal{A}}(\bar{a}^{(i_{\nu-1})})$ and

$\rho_{\nu-1} : (\mathcal{N}_{\nu-1}^{\mathcal{A}_w}(\vec{a}^{(0)}), \vec{a}^{(0)}) \cong (\mathcal{N}_{\nu-1}^{\mathcal{A}}(\vec{a}^{(i_\nu)}), \vec{a}^{(i_\nu)})$. If $w_\nu = 0$, case (ii) of the definition of \mathcal{A}_w includes a tuple \vec{c} with elements from $S_{\nu-1}^{\mathcal{A}_w} \cup S_\nu^{\mathcal{A}_w}$ in a relation of \mathcal{A}_w iff it is included in the corresponding relation in \mathcal{A} . Hence, we obtain $N_\nu^{\mathcal{A}_w}(\vec{a}^{(0)}) = N_\nu^{\mathcal{A}}(\vec{a}^{(i_\nu)})$. By definition of ρ_ν for $w_\nu = 0$, we have $\rho_\nu \upharpoonright N_{\nu-1}^{\mathcal{A}_w}(\vec{a}^{(0)}) = \rho_{\nu-1}$ and hence $\rho_\nu \upharpoonright S_{\nu-1}^{\mathcal{A}_w}(\vec{a}^{(0)}) = \text{id}_{\nu-1}$. Since also $\rho_\nu \upharpoonright S_\nu^{\mathcal{A}_w}(\vec{a}^{(0)}) = \text{id}_\nu$, we obtain $\rho_\nu : (\mathcal{N}_\nu^{\mathcal{A}_w}(\vec{a}^{(0)}), \vec{a}^{(0)}) \cong (\mathcal{N}_\nu^{\mathcal{A}}(\vec{a}^{(i_\nu)}), \vec{a}^{(i_\nu)})$.

Now consider the case that $w_\nu = 1$ and hence $i_\nu \equiv i_{\nu-1} + 1 \pmod t$. Since, by induction, (3.2) holds for $\nu - 1$, we can actually apply $\gamma_{\nu-1}$ after $\rho_{\nu-1}$, and hence ρ_ν is a well-defined bijection of $N_\nu^{\mathcal{A}_w}(\vec{a}^{(0)})$ and $N_\nu^{\mathcal{A}}(\vec{a}^{(i_\nu)})$. Furthermore, we have $\rho_{\nu-1}(\bar{a}_j) = \bar{a}_{(j+i_{\nu-1} \bmod t)}$, for each $j \in [t]$. Since $\gamma_{\nu-1}$ is an isomorphism from $(\mathcal{N}_{\nu-1}^{\mathcal{A}}(\vec{a}^{(i_{\nu-1})}), \vec{a}^{(i_{\nu-1})})$ to $(\mathcal{N}_{\nu-1}^{\mathcal{A}}(\vec{a}^{(i_\nu)}), \vec{a}^{(i_\nu)})$, we have $\gamma_{\nu-1}(\bar{a}_{(j+i_{\nu-1} \bmod t)}) = \bar{a}_{(j+i_\nu \bmod t)}$. Altogether, we obtain that ρ_ν maps each element of the tuple $\vec{a}^{(0)}$ to the element in the same position of the tuple $\vec{a}^{(i_\nu)}$, as required. In order to show that (3.2) holds for ν , it remains to show for each relation $R \in \sigma$ of arity r and each tuple $\vec{c} \in A^r$ that $\vec{c} \in R^{\mathcal{A}_w}$ iff $\rho_\nu(\vec{c}) \in R^{\mathcal{A}}$. If \vec{c} contains only elements from $N_{\nu-1}^{\mathcal{A}_w}(\vec{a}^{(0)})$, this follows since $\gamma_{\nu-1} \circ \rho_{\nu-1}$ is an isomorphism. It remains to consider tuples \vec{c} containing elements from $S_\nu^{\mathcal{A}_w}(\vec{a}^{(0)})$. That is, $\vec{c} \in (S_{\nu-1}^{\mathcal{A}_w} \cup S_\nu^{\mathcal{A}_w})^r$. As in the definition of \mathcal{A}_w , we assume for notational simplicity that $\vec{c} = (\bar{c}_{\nu-1}, \bar{c}_\nu)$ where the elements of $\bar{c}_{\nu-1}$ and \bar{c}_ν belong to $S_{\nu-1}^{\mathcal{A}_w}$ and $S_\nu^{\mathcal{A}_w}$, respectively. Here, $\bar{c}_{\nu-1}$ could be the empty tuple. By case (ii) of the definition of \mathcal{A}_w , we see that $\vec{c} \in R^{\mathcal{A}_w}$ iff $(\bar{c}_{\nu-1}, \pi^{-1}(\bar{c}_\nu)) \in R^{\mathcal{A}}$. Since π is, in particular, a partial isomorphism from \mathcal{A} to \mathcal{A} which is defined on all elements of $(\bar{c}_{\nu-1}, \pi^{-1}(\bar{c}_\nu))$, the previous statement holds iff $\pi((\bar{c}_{\nu-1}, \pi^{-1}(\bar{c}_\nu))) = (\pi(\bar{c}_{\nu-1}), \bar{c}_\nu) \in R^{\mathcal{A}}$. Since, as an isomorphism, π preserves distances, the elements of $\pi(\bar{c}_{\nu-1})$ belong to $S_{\nu-1}^{\mathcal{A}}(\vec{a}^{(i)})$, for each $i \in [t]$. Since $\rho_{\nu-1} \upharpoonright S_{\nu-1}^{\mathcal{A}_w}(\vec{a}^{(0)}) = \text{id}_{\nu-1}$ and $\rho_\nu \upharpoonright S_\nu^{\mathcal{A}_w}(\vec{a}^{(0)}) = \text{id}_\nu$, we have $\pi(\bar{c}_{\nu-1}) = \gamma_{\nu-1}(\rho_{\nu-1}(\bar{c}_{\nu-1})) = \rho_\nu(\bar{c}_{\nu-1})$ and $\bar{c}_\nu = \rho_\nu(\bar{c}_\nu)$. Hence, $(\pi(\bar{c}_{\nu-1}), \bar{c}_\nu) = (\rho_\nu(\bar{c}_{\nu-1}), \rho_\nu(\bar{c}_\nu)) = \rho_\nu(\vec{c})$ and we are done. \square

For the proof of Theorem 3.7, we want to convert the circuit \tilde{C} of the previous lemma, without increasing its size too much, to a circuit which accepts exactly the bitstrings that contain a number of ones which is divisible by r , where $r \geq 2$ is some factor of the number t from the previous lemma.

Lemma 3.11. *Let $m, d, M, t, p \in \mathbb{N}_{\geq 1}$ with $m > 9$ and $p, t \geq 2$ such that p and t are coprime. Let \tilde{C} be a MOD_p -circuit of depth d and size M which has the property that for all words $w, w' \in \{0, 1\}^m$ with $|w|_1 \equiv |w'|_1 \pmod t$, it accepts w iff it accepts w' . Furthermore, let \tilde{C} accept all $w \in \{0, 1\}^m$ with $|w|_1 \equiv 0 \pmod t$, and reject all $w \in \{0, 1\}^m$ with $|w|_1 \equiv 1 \pmod t$.*

There is a MOD_p -circuit \hat{C} of depth $(d+6)$ and size $(tM+2m^t)$ which, for some factor $r \geq 2$ of t , accepts exactly those bitstrings $w \in \{0, 1\}^m$ where $|w|_1 \equiv 0 \pmod r$.

Proof. We let $b = b_0 b_1 \cdots b_{t-1}$ be the bitstring of length t where, for every $j \in [t]$ we have $b_j = 1$ iff \tilde{C} accepts bitstrings $w \in \{0, 1\}^m$ with $|w|_1 \equiv j \pmod t$.

For a bitstring $w \in \{0, 1\}^m$ with $|w|_0 \geq t-1$, we let $\text{pattern}(w) = a_0 a_1 \cdots a_{t-1} \in \{0, 1\}^t$ with $a_j = 1$ iff \tilde{C} accepts the bitstring obtained from w by replacing the first j zeros with ones. Note that if $|w|_1 \equiv i \pmod t$, then $\text{pattern}(w) = b_i b_{i+1} \cdots b_{t-1} b_0 \cdots b_{i-1}$.

Claim. *There is a factor $r \geq 2$ of t such that for all $w \in \{0, 1\}^m$ with $|w|_0 \geq t-1$ we have: $\text{pattern}(w) = b \iff |w|_1 \equiv 0 \pmod r$.*

Proof. Obviously, $pattern(w) = b$ is true for all $w \in \{0, 1\}^m$ with $|w|_0 \geq t - 1$ and $|w|_1 \equiv 0 \pmod t$. In case that for all $w \in \{0, 1\}^m$ with $|w|_0 \geq t - 1$ we have

$$pattern(w) = b \iff |w|_1 \equiv 0 \pmod t,$$

we are done by choosing $r := t$.

In case that there is a $w \in \{0, 1\}^m$ with $|w|_0 \geq t - 1$, $pattern(w) = b$, and $|w|_1 \equiv i \pmod t$ for an $i \in \{1, \dots, t-1\}$, we know that $b_0 b_1 \cdots b_{t-1} = b_i b_{i+1} \cdots b_{t-1} b_0 \cdots b_{i-1}$. Thus, for $x := b_0 \cdots b_{i-1}$ and $y := b_i \cdots b_{t-1}$ we have $b = xy = yx$, and $x, y \in \{0, 1\}^+$.

A basic result in word combinatorics (see Proposition 1.3.2 in [13]) states that two words $x, y \in \{0, 1\}^+$ commute (i.e., $xy = yx$) iff they are powers of the same word (i.e., there is a $z \in \{0, 1\}^+$ and $\nu, \mu \in \mathbb{N}_{\geq 1}$ such that $x = z^\nu$ and $y = z^\mu$). We choose $z \in \{0, 1\}^+$ of minimal length such that $b = z^s$ for some $s \in \mathbb{N}$. Clearly, $|z| \geq 2$, since by assumption we have $b_0 b_1 = 10$.

Since z is of minimal length, it is straightforward to see that for every $w \in \{0, 1\}^m$ with $|w|_0 \geq t-1$ we have: $pattern(w) = z^s \iff |w|_1 \equiv 0 \pmod{|z|}$. \square

We choose r according to the claim. Obviously, the following is true for every $w \in \{0, 1\}^m$:

$$|w|_1 \equiv 0 \pmod r \iff \begin{cases} (1) & |w|_0 \geq t-1 \text{ and } pattern(w) = b, \text{ or} \\ (2) & \text{there is a } j \in [t-1] \text{ with } m-j \equiv 0 \pmod r \text{ such that } |w|_0 = j. \end{cases}$$

To complete the proof of Lemma 3.11, it therefore suffices to construct circuits $C_{(1)}$ and $C_{(2)}$ testing for (1) and (2), respectively, and to let \hat{C} be the disjunction of $C_{(1)}$ and $C_{(2)}$. For constructing these circuits, note the following:

- For each $j \leq m$, a circuit C_j of depth 2 and size $\leq m^j$ which accepts an input $w \in \{0, 1\}^m$ iff $|w|_0 = j$ can be realised via $\bigvee_{\substack{I \subseteq \{1, \dots, m\} \\ |I|=j}} (\bigwedge_{i \in I} \neg w_i \wedge \bigwedge_{i \in \{1, \dots, m\} \setminus I} w_i)$.
- The circuit $C_{(2)}$ can be realised via $\bigvee_{\substack{j \in [t-1] \\ j \equiv 0 \pmod r}} C_j$. In particular, $C_{(2)}$ has depth 2 and size $< m^t$.
- A circuit $C_{\geq t-1}$ of depth 2 and size $\leq m^{t-1}$ which accepts an input $w \in \{0, 1\}^m$ iff $|w|_0 \geq t-1$ can be realised via $\bigvee_{\substack{I \subseteq \{1, \dots, m\} \\ |I|=t-1}} \bigwedge_{i \in I} \neg w_i$.
- It is not difficult to construct, for each $j \in [t]$ a circuit $C^{(j)}$ with $2m$ output gates $w_1^{(j)}, \neg w_1^{(j)}, \dots, w_m^{(j)}, \neg w_m^{(j)}$, such that on input of a $w \in \{0, 1\}^m$ with $|w|_0 \geq t-1$, this circuit produces at its output gates $w_1^{(j)}, \dots, w_m^{(j)}$ the string $w^{(j)} = w_1^{(j)} \cdots w_m^{(j)}$ obtained from w by replacing the first j zeros with ones. For $j \geq 1$ note that $w_i^{(j)}$ can be expressed as

$$w_i \vee \left(\neg w_i \wedge \bigvee_{\substack{I \subseteq \{1, \dots, i-1\} \\ |I| < j}} \left(\bigwedge_{i' \in I} \neg w_{i'} \wedge \bigwedge_{i' \in \{1, \dots, i-1\} \setminus I} w_{i'} \right) \right).$$

Thus, the circuit $C^{(j)}$ has depth 4 and size $2m(3 + m^{j-1}) \leq 8m^j$.

- For each $j \in [s]$ let $\tilde{C}^{(j)}$ be the concatenation of $C^{(j)}$ and \tilde{C} . Note that on input of a bitstring $w \in \{0, 1\}^m$ with $|w|_0 \geq t-1$, the circuit $\tilde{C}^{(j)}$ computes, at its output gate, the letter a_j of $pattern(w) = a_0 \cdots a_{t-1}$. Furthermore, $\tilde{C}^{(j)}$ has depth $(4+d)$ and size $\leq 8m^j + M$.

- The circuit $C_{(1)}$ can now be realised via

$$C_{\geq t-1} \wedge \bigwedge_{\substack{j \in [t] \\ b_j=1}} \tilde{C}^{(j)} \wedge \bigwedge_{\substack{j \in [t] \\ b_j=0}} \neg \tilde{C}^{(j)}.$$

This circuit has depth $(d+5)$ and size $\leq 1 + m^{t-1} + \sum_{j=0}^{t-1} (8m^j + M) \leq tM + m^t$ (for $m > 9$).

- Finally, the circuit \hat{C} is realised via $C_{(1)} \vee C_{(2)}$. Thus, \hat{C} has depth $(d+6)$ and size $\leq tM + 2m^t$.

□

We are now ready for the proof of Theorem 3.7.

Proof of Theorem 3.7. Let q be a kt -ary query defined on \mathfrak{C} by an arb-inv-FO+MOD $_p^{\mathfrak{C}}(\sigma)$ -formula $\varphi(\bar{x}_0, \dots, \bar{x}_{t-1})$, where \bar{x}_i is a k -tuple of variables, for each $i \in [t]$. By Theorem 3.9, there exist numbers $d, s \in \mathbb{N}$ such that for every $n \in \mathbb{N}_{\geq 1}$ there is a MOD $_p$ -circuit C_n with $\lambda_{kt}^{\sigma}(n)$ input bits, depth d , and size n^s such that the following is true for all σ -structures $\mathcal{A} \in \mathfrak{C}$ with $|A| = n$, all k -tuples $\bar{a}_0, \dots, \bar{a}_{t-1} \in A^k$, and all embeddings ι of \mathcal{A} into $[n]$:

$$C_n \text{ accepts } \text{Rep}^{\iota}(\mathcal{A}, \bar{a}_0, \dots, \bar{a}_{t-1}) \iff \mathcal{A}^{\iota} \models \varphi[\bar{a}_0, \dots, \bar{a}_{t-1}] \iff \mathcal{A} \models \varphi[\bar{a}_0, \dots, \bar{a}_{t-1}].$$

For contradiction, assume that for every $c \in \mathbb{N}$ the query q is *not* shift $(\log n)^c$ -local w.r.t. t on \mathfrak{C} . Thus, in particular for $c := 2\ell(d+6)$ (with ℓ chosen as in Theorem 3.8), we obtain that for all $n_0 \in \mathbb{N}$ there is an $n \geq n_0$, and a σ -structure $\mathcal{A} \in \mathfrak{C}$ with $|A| = n$, and k -tuples $\bar{a}_0, \dots, \bar{a}_{t-1} \in A^k$ such that for $m := (\log n)^c = (\log n)^{2\ell(d+6)}$ we have:

- $(\mathcal{N}_m^{\mathcal{A}}(\bar{a}_i), \bar{a}_i) \cong (\mathcal{N}_m^{\mathcal{A}}(\bar{a}_j), \bar{a}_j)$ and $N_m^{\mathcal{A}}(\bar{a}_i) \cap N_m^{\mathcal{A}}(\bar{a}_j) = \emptyset$ for all $i, j \in [t]$ with $i \neq j$, and
- $\mathcal{A} \models \varphi[\bar{a}_0, \bar{a}_1, \dots, \bar{a}_{t-1}]$ and $\mathcal{A} \not\models \varphi[\bar{a}_1, \dots, \bar{a}_{t-1}, \bar{a}_0]$.

We fix $n \in \mathbb{N}$ sufficiently large such that, for $\hat{d} := (d+6)$ and ε and $m_{\hat{d}}$ chosen as in Theorem 3.8, we have for $m = (\log n)^c$ that $m > 9$, $m \geq m_{\hat{d}}$, and $n^{\varepsilon \log n} > tn^s + 2(\log n)^{ct}$.

Clearly, C_n is a MOD $_p$ -circuit with $\lambda_{kt}^{\sigma}(n)$ input bits which, for every $i \in [t]$ and all embeddings ι_1 and ι_2 of \mathcal{A} accepts $\text{Rep}^{\iota_1}(\mathcal{A}, \bar{a}^{(i)})$ iff it accepts $\text{Rep}^{\iota_2}(\mathcal{A}, \bar{a}^{(i)})$. Furthermore, C_n accepts $\text{Rep}^{\iota}(\mathcal{A}, \bar{a}^{(0)})$ and rejects $\text{Rep}^{\iota}(\mathcal{A}, \bar{a}^{(1)})$, for every embedding ι of \mathcal{A} . Thus, from Lemma 3.10 we obtain a MOD $_p$ -circuit \tilde{C} on m input bits, with depth d and size n^s , such that \tilde{C} has the property that for all words $w, w' \in \{0, 1\}^m$ with $|w|_1 \equiv |w'|_1 \pmod{t}$, it accepts w iff it accepts w' . Furthermore, \tilde{C} accepts all $w \in \{0, 1\}^m$ with $|w| \equiv 0 \pmod{t}$ and rejects all $w \in \{0, 1\}^m$ with $|w| \equiv 1 \pmod{t}$. From Lemma 3.11, we therefore obtain a MOD $_p$ -circuit \hat{C} of depth $\hat{d} := (d+6)$ and size $(tn^s + 2m^t) = (tn^s + 2(\log n)^{ct})$ which, for some factor $r \geq 2$ of t , accepts exactly those bitstrings $w \in \{0, 1\}^m$ where $|w|_1 \equiv 0 \pmod{r}$.

Since p and t are coprime by assumption, and $r \geq 2$ is a factor of t , we know that r has a prime factor different from p 's prime factor. Thus, from Theorem 3.8 (for $\varepsilon, \ell, m_{\hat{d}}$ chosen as in Theorem 3.8, and for $m \geq m_{\hat{d}}$) we know that the size $(tn^s + 2(\log n)^{ct})$ of \hat{C} must be bigger than $2^{\varepsilon \ell \hat{d} \overline{m}}$. However, we chose $m = (\log n)^c = (\log n)^{2\ell \hat{d}}$, and hence $2^{\varepsilon \ell \hat{d} \overline{m}} = 2^{\varepsilon (\log n)^2} = n^{\varepsilon \log n} > tn^s + 2(\log n)^{ct}$ for all sufficiently large n — a contradiction! Thus, the proof of Theorem 3.7 is complete. □

3.4. Applications. In the same way as Gaifman locality (cf., e.g., [12]), also shift locality can be used for showing that certain queries are not expressible in particular logics. The first example query we consider here is the reachability query *reach* which associates, with every finite directed graph $\mathcal{A} = (A, E^{\mathcal{A}})$, the relation $\text{reach}(\mathcal{A}) := \{(a, b) : \mathcal{A} \text{ contains a directed path from node } a \text{ to node } b\}$.

Proposition 3.12. *Let $\sigma = \{E\}$ consist of a binary relation symbol E . Let $p, t \in \mathbb{N}$ with $p, t \geq 2$ be such that every t -ary query q definable in $\text{arb-inv-FO+MOD}_p(\sigma)$ is shift $f_q(n)$ -local w.r.t. t , for a function $f_q : \mathbb{N} \rightarrow \mathbb{N}$ where $f_q(n) \leq (\frac{n}{2t} - \frac{1}{2})$ for all sufficiently large n . Then, the reachability query is not definable in $\text{arb-inv-FO+MOD}_p(\sigma)$.*

Proof. Assume, for contradiction, that *reach* is definable by an $\text{arb-inv-FO+MOD}_p(\sigma)$ -formula $\varrho(x, y)$. Then, $\psi(x_0, \dots, x_{t-1}) := \varrho(x_0, x_1) \wedge \varrho(x_1, x_2) \wedge \dots \wedge \varrho(x_{t-2}, x_{t-1})$ is an $\text{arb-inv-FO+MOD}_p(\sigma)$ -formula expressing in a finite directed graph \mathcal{A} , that there is a directed path from node x_i to node x_{i+1} , for every $i \in [t-1]$. Let q be the t -ary query defined by $\psi(x_0, \dots, x_{t-1})$. By assumption, this query is shift $f_q(n)$ -local w.r.t. t , for a function f_q with $f_q(n) \leq \frac{n}{2t} - \frac{1}{2}$ for all sufficiently large n .

Now, consider for each $\ell \in \mathbb{N}_{\geq 1}$ the graph \mathcal{A}_ℓ consisting of a single directed path $v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_{t(2\ell+1)}$ on $t \cdot (2\ell+1)$ nodes. For each $i \in [t]$ let $a_i := v_{i(2\ell+1) + (\ell+1)}$. Then, the ℓ -neighbourhoods of the a_i , for $i \in [t]$, are pairwise disjoint and isomorphic. The cardinality of \mathcal{A}_ℓ is $n := t \cdot (2\ell+1)$, and thus $\ell = \frac{n}{2t} - \frac{1}{2} \geq f_q(n)$. Since q is shift $f_q(n)$ -local w.r.t. t , we obtain that $\mathcal{A}_\ell \models \psi[a_0, a_1, \dots, a_{t-1}] \iff \mathcal{A}_\ell \models \psi[a_1, \dots, a_{t-1}, a_0]$. But in \mathcal{A}_ℓ there is a directed path from a_i to a_{i+1} for every $i \in [t-1]$, but there is no directed path from a_{t-1} to a_0 . According to the choice of ψ , we have that $\mathcal{A}_\ell \models \psi[a_0, a_1, \dots, a_{t-1}]$ but $\mathcal{A}_\ell \not\models \psi[a_1, \dots, a_{t-1}, a_0]$ — a contradiction! \square

As an immediate consequence of Proposition 3.12 and Theorem 3.7 we obtain (the known fact) that the reachability query is not definable in $\text{arb-inv-FO+MOD}_p(\{E\})$, for any prime power p . Proposition 3.12 also shows why it can be expected to be difficult to generalise Theorem 3.5 and Theorem 3.7 from prime powers p to arbitrary numbers $p \geq 2$:

Remark 3.13. Assume, we could generalise Theorem 3.7 from prime powers p to arbitrary numbers $p \geq 2$. By Proposition 3.12, we would then obtain that the reachability query is not definable in $\text{arb-inv-FO+MOD}_p(\{E\})$, for any $p \in \mathbb{N}$ with $p \geq 2$. The “opposite direction” of Theorem 3.9, obtained in [3], would then tell us that the reachability query is not computable in ACC^0 . Here, $\text{ACC}^0 = \bigcup_{p \geq 2} \text{AC}^0[p]$, where $\text{AC}^0[p]$ is the class of all problems computable by a family of constant depth, polynomial size MOD_p -circuits. Since the reachability query can be computed in nondeterministic logarithmic space, we would thus obtain that $\text{NLOGSPACE} \not\subseteq \text{ACC}^0$. This would constitute a major breakthrough in computational complexity: The current state-of-the-art (see [19] for a recent survey) states that $\text{NEXP} \not\subseteq \text{ACC}^0$, but does not know a problem in PTIME that provably does not belong to ACC^0 .

Similarly, a generalisation of Theorem 3.5 to all odd numbers p would imply that the reachability query is not definable in $\text{AC}^0[p]$, for any odd number p . Also this is currently not known.

The fact that FO-definable queries are Gaifman local simplifies and unifies many non-expressibility results for FO. In fact, to someone acquainted with Gaifman locality, the proof of a non-expressibility result by a locality argument can often be communicated in

a simple visual way, while the complicated combinatorial arguments in game based non-expressibility results are hidden within the proof of Gaifman locality for FO. We give some further examples of well-known queries which have already been used in a similar context in the literature on locality, and we show that shift locality can also be used to obtain non-expressibility of these queries in $\text{arb-inv-FO+MOD}_p(\{E\})$, for prime powers p . For none of these examples, the fact that they are not expressible in $\text{arb-inv-FO+MOD}_p(\{E\})$, for any prime power p , is new. That is, all examples could be proved directly using Smolensky's theorem. But we believe that our examples show that the notion of shift locality serves a similar purpose as Gaifman locality. That is, the locality argument simplifies and unifies the proofs, i.e. all details of the necessary reductions are done in Theorem 3.7.

Using similar constructions as in the proof of Proposition 3.12, we show that none of the following queries is definable in $\text{arb-inv-FO+MOD}_p(\{E\})$, for any prime power p :

- $\text{cycle}(\mathcal{A}) := \{a \in A : a \text{ is a node that lies on a cycle of the graph } \mathcal{A} = (A, E^{\mathcal{A}})\}$,
- $\text{triangle-reach}(\mathcal{A}) := \{a \in A : a \text{ is reachable from a triangle in the graph } \mathcal{A} = (A, E^{\mathcal{A}})\}$,
- $\text{same-distance}(\mathcal{A}) := \{(a, b, c) \in A^3 : \text{dist}^{\mathcal{A}}(a, c) = \text{dist}^{\mathcal{A}}(b, c)\}$.

In the literature, variants of these queries have previously served as examples of queries which are not Gaifman local. In the proofs below, we will use the following notion. Given a directed graph $\mathcal{A} = (A, E^{\mathcal{A}})$ and an edge $e = (u, v) \in E^{\mathcal{A}}$, the ℓ -fold subdivision of (u, v) replaces the edge (u, v) by a path $u \rightarrow u_{e,1} \rightarrow \dots \rightarrow u_{e,\ell} \rightarrow v$, where $u_{e,1}, \dots, u_{e,\ell} \notin A$ are new nodes introduced for subdividing e . This notion extends naturally to edge sets instead of single edges: The ℓ -fold subdivision of a set E of directed edges is obtained by replacing each edge $(u, v) \in E$ by a path $u \rightarrow u_{e,1} \rightarrow \dots \rightarrow u_{e,\ell} \rightarrow v$, where $u_{e,1}, \dots, u_{e,\ell}$ are new nodes introduced for subdividing the edge e .

Proposition 3.14.

For any prime power p , the cycle query is not definable in $\text{arb-inv-FO+MOD}_p(\{E\})$.

Proof. Let $t \in \mathbb{N}$ with $t \geq 2$ such that t and p are coprime (e.g., $t = 2$ if p is odd, and $t = 3$ if p is even). Assume, for contradiction, that the cycle query is definable by an $\text{arb-inv-FO+MOD}_p(\{E\})$ -formula $\varrho(x)$. Let x_0, \dots, x_{t-1} be first-order variables that do not occur in $\varrho(x)$ and let

$$\psi(x_0, \dots, x_{t-1}) := \varrho(x_0) \wedge \bigwedge_{1 \leq i < t} x_i = x_i.$$

Let q be the t -ary query defined by $\psi(x_0, \dots, x_{t-1})$. By Theorem 3.7, the query q is shift $(\log n)^c$ -local, for some $c \in \mathbb{N}$.

Let \mathcal{A}_0 be a directed graph which is a disjoint union of a directed cycle $u_0 \rightarrow u_1 \rightarrow \dots \rightarrow u_0$ and a directed path $v_0 \rightarrow v_1 \rightarrow \dots \rightarrow v_t$.

For each $\ell \in \mathbb{N}_{\geq 1}$, let \mathcal{A}_ℓ be obtained by the ℓ -fold subdivision of the edge set of \mathcal{A}_0 . In particular, \mathcal{A}_ℓ has cardinality $n_\ell = t+3 + (t+2)\ell$.

Let $a_0 := u_0$, and for each $i \in \{1, \dots, t-1\}$, let $a_i := v_i$. For sufficiently large values of ℓ , the $(\log n_\ell)^c$ -neighbourhoods of the nodes a_0, \dots, a_{t-1} in \mathcal{A}_ℓ are pairwise disjoint and isomorphic, each of them being isomorphic to a directed path of length $2(\log n_\ell)^c + 1$ with node a_i in the middle.

Since q is shift $(\log n)^c$ -local w.r.t. t , we obtain for all sufficiently large ℓ that

$$\mathcal{A}_\ell \models \psi[a_0, a_1, \dots, a_{t-1}] \iff \mathcal{A}_\ell \models \psi[a_1, \dots, a_{t-1}, a_0].$$

However, $a_0 \in \text{cycle}(\mathcal{A}_\ell)$, because a_0 lies on a cycle, and $a_1 \notin \text{cycle}(\mathcal{A}_\ell)$, because a_1 does not lie on a cycle. That is, $\mathcal{A}_\ell \models \varrho[a_0]$ and $\mathcal{A}_\ell \not\models \varrho[a_1]$. Thus, according to the choice of the formula ψ , we have that $\mathcal{A}_\ell \models \psi[a_0, a_1, \dots, a_{t-1}]$, but $\mathcal{A}_\ell \not\models \psi[a_1, \dots, a_{t-1}, a_0]$, which is a contradiction. \square

Proposition 3.15.



For any prime power p , the triangle-reach query is not definable in $\text{arb-inv-FO+MOD}_p(\{E\})$.

Proof. Let $t \in \mathbb{N}$ with $t \geq 2$ such that t and p are coprime. Assume, for contradiction, that the triangle-reach query is definable by an $\text{arb-inv-FO+MOD}_p(\{E\})$ -formula $\varrho(x)$. Let x_0, \dots, x_{t-1} be first-order variables that do not occur in $\varrho(x)$ and let

$$\psi(x_0, \dots, x_{t-1}) := \varrho(x_{t-1}) \wedge \bigwedge_{0 \leq i < t-1} x_i = x_i.$$

Let q be the t -ary query defined by $\psi(x_0, \dots, x_{t-1})$. By Theorem 3.7, the query q is shift $(\log n)^c$ -local, for some $c \in \mathbb{N}$.

Let \mathcal{A}_0 be the graph consisting of the directed path $v_0 \rightarrow \dots \rightarrow v_t$, the directed triangle $v_t \rightarrow v_{t+1} \rightarrow v_{t+2} \rightarrow v_t$, and the directed path $v_{t+2} \rightarrow v_{t+3} \rightarrow v_{t+4}$.



For each $\ell \in \mathbb{N}_{\geq 1}$, let \mathcal{A}_ℓ be obtained by the ℓ -fold subdivision of all but the triangle's edges of \mathcal{A}_0 . In particular, \mathcal{A}_ℓ has cardinality $n_\ell = t+5 + (t+2)\ell$. For each $i \in [t-1]$, let $a_i := v_{i+1}$, and let $a_{t-1} := v_{t+3}$. For sufficiently large values of ℓ , the $(\log n_\ell)^c$ -neighbourhoods of the nodes a_0, \dots, a_{t-1} in \mathcal{A}_ℓ are pairwise disjoint and isomorphic, each of them being isomorphic to a directed path of length $2(\log n_\ell)^c + 1$ with node a_i in the middle.

Since q is shift $(\log n)^c$ -local w.r.t. t , we obtain for all sufficiently large ℓ that

$$\mathcal{A}_\ell \models \psi[a_0, a_1, \dots, a_{t-1}] \iff \mathcal{A}_\ell \models \psi[a_1, \dots, a_{t-1}, a_0].$$

However, $a_{t-1} \in \text{triangle-reach}(\mathcal{A}_\ell)$, because a_{t-1} is reachable from the triangle, and $a_0 \notin \text{triangle-reach}(\mathcal{A}_\ell)$, because there is no directed path from any node of the triangle to a_0 . That is, $\mathcal{A}_\ell \models \varrho[a_{t-1}]$ and $\mathcal{A}_\ell \not\models \varrho[a_0]$. Thus, due to the choice of the formula ψ , we have that $\mathcal{A}_\ell \models \psi[a_0, a_1, \dots, a_{t-1}]$, but $\mathcal{A}_\ell \not\models \psi[a_1, \dots, a_{t-1}, a_0]$, which is a contradiction. \square

Proposition 3.16.

For any prime power p , the same-distance query is not definable in $\text{arb-inv-FO+MOD}_p(\{E\})$.

Proof. Let $t \in \mathbb{N}$ with $t \geq 3$ such that t and p are coprime. Assume, for contradiction, that the same-distance query is definable by an $\text{arb-inv-FO+MOD}_p(\{E\})$ -formula $\varrho(x, y, z)$.


Let x_0, \dots, x_{t-1} be first-order variables that do not occur in $\varrho(x, y, z)$ and let


$$\psi(x_0, \dots, x_{t-1}) := \varrho(x_{t-3}, x_{t-2}, x_{t-1}) \wedge \bigwedge_{0 \leq i < t-3} x_i = x_i.$$

Let q be the t -ary query defined by $\psi(x_0, \dots, x_{t-1})$. By Theorem 3.7, the query q is shift $(\log n)^c$ -local, for some $c \in \mathbb{N}$.

Let \mathcal{A}_0 be the graph consisting of the edges $u \rightarrow v_i$ for each $i \in \{0, \dots, t-2\}$, and a directed path $u \rightarrow v_{t-1} \rightarrow v_t$.

For each $\ell \in \mathbb{N}_{\geq 1}$, let \mathcal{A}_ℓ be the graph obtained by the ℓ -fold subdivision of the edges $u \rightarrow v_i$ in \mathcal{A}_0 , for each $i \in [t]$.





In particular, \mathcal{A}_ℓ has cardinality $n_\ell = t+2 + t\ell$. For each $i \in [t-1]$ let $a_i := v_i$, and let $a_{t-1} := v_t$. For sufficiently large values of ℓ , the $(\log n_\ell)^c$ -neighbourhoods of nodes a_0, \dots, a_{t-1} in \mathcal{A}_ℓ are pairwise disjoint and isomorphic, each of them being isomorphic to a directed path of length $(\log n_\ell)^c + 1$ with a_i at the end of the path.

Since q is shift $(\log n)^c$ -local w.r.t. t , we obtain for all sufficiently large ℓ that

$$\mathcal{A}_\ell \models \psi[a_0, a_1, \dots, a_{t-1}] \iff \mathcal{A}_\ell \models \psi[a_1, \dots, a_{t-1}, a_0].$$

However, $(a_{t-3}, a_{t-2}, a_{t-1}) \in \text{same-distance}(\mathcal{A}_\ell)$, because a_{t-3} and a_{t-2} both have distance $2(\ell+1)+1$ to a_{t-1} in \mathcal{A}_ℓ , and $(a_{t-2}, a_{t-1}, a_0) \notin \text{same-distance}(\mathcal{A}_\ell)$, because a_{t-2} has distance $2(\ell+1)$ to a_0 , but a_{t-1} has distance $2(\ell+1)+1$ to a_0 . That is, $\mathcal{A}_\ell \models \varrho[a_{t-3}, a_{t-2}, a_{t-1}]$ and $\mathcal{A}_\ell \not\models \varrho[a_{t-2}, a_{t-1}, a_0]$. Thus, due to the choice of the formula ψ , we have that $\mathcal{A}_\ell \models \psi[a_0, a_1, \dots, a_{t-1}]$, but $\mathcal{A}_\ell \not\models \psi[a_1, \dots, a_{t-1}, a_0]$, which is a contradiction. \square

4. HANF LOCALITY AND LOCALITY ON STRING STRUCTURES

For giving the precise definition of Hanf locality, we need the following notation: As in [12], for σ -structures \mathcal{A} and \mathcal{B} , for k -tuples $\bar{a} \in A^k$ and $\bar{b} \in B^k$, and for an $r \in \mathbb{N}$, we write $(\mathcal{A}, \bar{a}) \simeq_r (\mathcal{B}, \bar{b})$ (or simply $\mathcal{A} \simeq_r \mathcal{B}$ in case that $k=0$) if there is a bijection $\beta : A \rightarrow B$ such that $(\mathcal{N}_r^{\mathcal{A}}(\bar{a}c), \bar{a}c) \cong (\mathcal{N}_r^{\mathcal{B}}(\bar{b}\beta(c)), \bar{b}\beta(c))$ is true for every $c \in A$.

Definition 4.1 (Hanf locality). Let \mathfrak{C} be a class of finite σ -structures, $k \in \mathbb{N}$, and $f : \mathbb{N} \rightarrow \mathbb{N}$. A k -ary query q is *Hanf $f(n)$ -local on \mathfrak{C}* if there is an $n_0 \in \mathbb{N}$ such that for every $n \in \mathbb{N}$ with $n \geq n_0$ and all σ -structures $\mathcal{A}, \mathcal{B} \in \mathfrak{C}$ with $|A| = |B| = n$, the following is true for all k -tuples $\bar{a} \in A^k$ and $\bar{b} \in B^k$ with $(\mathcal{A}, \bar{a}) \simeq_{f(n)} (\mathcal{B}, \bar{b})$: $\bar{a} \in q(\mathcal{A}) \iff \bar{b} \in q(\mathcal{B})$.

The query q is called *Hanf $f(n)$ -local* if it is Hanf $f(n)$ -local on the class of all finite σ -structures.

Consider a 0-ary query which maps a structure, containing a unary relation P , to the relation containing the empty tuple if the number of elements contained in P is even, and to the empty set, otherwise. This is an example of a Hanf 0-local query. The 0-ary query which maps a graph to the relation containing only the empty tuple if it is connected, and to the empty set otherwise, is an example of a query which is not Hanf $f(n)$ -local for any sublinear function f .

Hanf locality is an even stronger locality notion than Gaifman locality:

Theorem 4.2 (Hella, Libkin, Nurmonen [11]). *Let \mathfrak{C} be a class of finite σ -structures and let $f : \mathbb{N} \rightarrow \mathbb{N}$. Let $k \in \mathbb{N}_{\geq 1}$ and let q be a k -ary query. If q is Hanf $f(n)$ -local on \mathfrak{C} , then q is Gaifman $(3f(n)+1)$ -local on \mathfrak{C} .*

It is well-known that queries definable in FO or FO+MOD $_p$ (for any $p \geq 2$) are Hanf local with a constant locality radius [8, 11]. For order-invariant or arb-invariant FO it is still open whether they are Hanf local with respect to any sublinear locality radius. As an immediate consequence of Proposition 3.2 and Theorem 4.2 one obtains for every $p \in \mathbb{N}$ with $p \geq 2$ that order-invariant FO+MOD $_p$ is *not* Hanf local with respect to any sublinear locality radius.

An arb-inv-FO+MOD $_p(\sigma_\Sigma)$ -sentence φ defines the string-language $L_\varphi := \{w \in \Sigma^+ : S_w \models \varphi\}$. A string-language $L \subseteq \Sigma^+$ corresponds to a 0-ary query q_L on Σ -strings defined,

for each $w \in \Sigma^+$, via

$$q_L(\mathcal{S}_w) := \begin{cases} \{()\} & \text{if } w \in L \\ \emptyset & \text{if } w \notin L, \end{cases}$$

where $()$ denotes the unique tuple of arity 0. A string-language $L \subseteq \Sigma^+$ is called *Hanf $f(n)$ -local* iff the 0-ary query q_L is Hanf $f(n)$ -local on Σ -strings, i.e., there is an $n_0 \in \mathbb{N}$ such that for every $n \in \mathbb{N}$ with $n \geq n_0$ and all strings $u, v \in \Sigma^+$ of length n with $\mathcal{S}_u \stackrel{f(n)}{\leftrightarrow} \mathcal{S}_v$, we have: $u \in L \iff v \in L$.

For the restricted case of string structures, Benedikt and Segoufin [6] have shown that on Σ -strings order-invariant FO has the same expressive power as FO and thus is Hanf local with constant locality radius (in fact, [6] obtains the same result also for finite labelled ranked trees). In [1] it was shown that every query definable in arb-invariant FO on Σ -strings is Hanf local with polylogarithmic locality radius, and that in the worst case the locality radius can indeed be of polylogarithmic size. As an immediate consequence of Proposition 3.4 and Theorem 4.2 we obtain that for $\Sigma := \{0, 1\}$ there is a unary query q that is *not* Hanf $(\frac{n-8}{12})$ -local on Σ -strings, but definable in $\langle\text{-inv-FO}+\text{MOD}_p(\sigma_\Sigma)\rangle$ for every *even* number $p \geq 2$. From this observation, we can also obtain an example of a language which is not Hanf $(\frac{n-8}{12})$ -local on strings over an extended alphabet. The existence of this language can be obtained using a general principle which allows to convert a k -ary query q over strings to a 0-ary query, i.e. a language A_q over strings over an extended alphabet. The language A_q inherits the relevant definability and locality properties for our purposes from the query q . This principle is stated in Lemma 4.10 below, since we will need it for the proof of Theorem 4.6 below. Using this approach here, however, it becomes rather hard to describe the shape of the language which is not Hanf $(\frac{n-8}{12})$ -local concretely. To this end, we show directly how this language can be obtained by simple modification of the proof of Proposition 3.4. Consider the languages

$$\begin{aligned} L_{\text{left}} &:= 1^+ 2 0^+ 1^+ 0^+ \\ L_{\text{right}} &:= 1^+ 0^+ 1^+ 2 0^+ \end{aligned}$$

over the alphabet $\Sigma := \{0, 1, 2\}$. Note that the definitions of L_{left} and L_{right} are very similar, the only difference being the position of the unique 2 occurring in strings from L_{left} and L_{right} . We define

$$\begin{aligned} L_{\text{even}} &:= \{w \in L_{\text{right}} : |w|_1 \bmod 2 \equiv 0\} \\ L_{\text{odd}} &:= \{w \in L_{\text{left}} : |w|_1 \bmod 2 \equiv 1\} \end{aligned}$$

Proposition 4.3. *Let $L := L_{\text{even}} \cup L_{\text{odd}}$.*

- (a) *The language L is not Hanf $(\frac{n-1}{8})$ -local.*
- (b) *L is definable by a sentence φ in $\langle\text{-inv-FO}+\text{MOD}_p(\sigma_\Sigma)\rangle$, for every even number $p \geq 2$.*

Proof. (a): For every $\ell \in \mathbb{N}_{\geq 1}$ let

$$\begin{aligned} u_\ell &:= 1^\ell 1^\ell 2 0^\ell 0^\ell 1^\ell 1^\ell 0^\ell 0^\ell = \mathbf{x} \mathbf{y}_1 \mathbf{y}_2 \mathbf{y}_3 \mathbf{z}, \\ v_\ell &:= 1^\ell 1^\ell 0^\ell 0^\ell 1^\ell 1^\ell 2 0^\ell 0^\ell = \mathbf{x} \mathbf{y}_3 \mathbf{y}_2 \mathbf{y}_1 \mathbf{z}, \end{aligned}$$

for $\mathbf{x} := 1^\ell$, $\mathbf{y}_1 := 1^\ell 2 0^\ell$, $\mathbf{y}_2 := 0^\ell 1^\ell$, $\mathbf{y}_3 := 1^\ell 0^\ell$, $\mathbf{z} := 0^\ell$.

It is not difficult to see that $\mathcal{S}_{u_\ell} \stackrel{\ell}{\leftrightarrow} \mathcal{S}_{v_\ell}$: the bijection β , for which

$$(\mathcal{N}_\ell^{\mathcal{S}_{u_\ell}}(c), c) \cong (\mathcal{N}_\ell^{\mathcal{S}_{v_\ell}}(\beta(c)), \beta(c)) \quad \text{for every } c \in \{1, \dots, |u_\ell|\},$$

can be chosen as follows. It maps each position of

- \mathbf{x} in u_ℓ onto the according position of \mathbf{x} in v_ℓ ,
- \mathbf{y}_s (for $s \in \{1, 2, 3\}$) in u_ℓ onto the according position of \mathbf{y}_s in v_ℓ ,
- \mathbf{z} in u_ℓ onto the according position of \mathbf{z} in v_ℓ .

It is straightforward to verify that this bijection β indeed witnesses that $\mathcal{S}_{u_\ell} \stackrel{\beta}{\leftrightarrow}_\ell \mathcal{S}_{v_\ell}$. Furthermore, $v_\ell \in L$ and $u_\ell \notin L$. The length of u_ℓ and v_ℓ is $n := 8\ell + 1$, thus $\ell = \frac{n-1}{8}$. Hence, the language L is not Hanf $(\frac{n-1}{8})$ -local. This completes the proof of (a).

(b): First, note that the language

$$M := L_{left} \cup L_{right}$$

can be defined by an $\text{FO}(\sigma_\Sigma)$ -sentence φ_M which states the following:

- The first position of the string carries the letter 1. The last position of the string carries the letter 0.
- For each position x that carries the letter 1, the position directly to the right of x carries one of the letters 0, 1, 2. Furthermore, there is exactly one position x that carries the letter 1, such that the position directly to the right of x carries the letter 0. And there is exactly one position x that carries the letter 1, such that the position directly to the right of x carries the letter 2.
- For each position y that carries the letter 0, the position directly to the right of y carries one of the letters 0, 1. Furthermore, there is exactly one position y that carries the letter 0, such that the position directly to the right of y carries the letter 1.
- There is exactly one position z that carries the letter 2. The position directly to the right of z carries the letter 0.

From Example 2.5 we obtain an $\text{<-inv-FO+MOD}_2(\{E\})$ -sentence $\varphi_{\text{even cycles}}$ that is satisfied by a finite $\{E\}$ -structure \mathcal{A} iff \mathcal{A} is a disjoint union of directed cycles where the number of cycles of even length is even. We choose

$$\varphi := \varphi_M \wedge \varphi',$$

where φ' is the formula obtained from $\varphi_{\text{even cycles}}$ by relativisation of all quantifiers to the positions that carry the letters 1 or 2, and by replacing every atom of the form $E(\mu, \nu)$ (for first-order variables μ and ν) by a formula stating that

- $E(\mu, \nu)$ is true, or
- position μ carries the letter 2, and ν is the leftmost position of the string, or
- the positions μ and ν both carry the letter 1, and the positions directly to the right of μ and directly to the left of ν both carry the letter 0.

Clearly, the obtained formula is order-invariant on the class of all finite σ_Σ -structures, since the formula $\varphi_{\text{even cycles}}$ is order-invariant on the class of all finite $\{E\}$ -structures.

It is straightforward to see that, when evaluated in a string $w \in L_{left}$ of the form $1^i 2^j 1^{i'} 0^{j'}$, the formula φ' simulates the evaluation of the formula $\varphi_{\text{even cycles}}$ in a graph that consists of the disjoint union of two cycles of lengths $i+1$ and i' ; and hence φ' is satisfied iff $i+1$ and i' are either both even or both odd — and this is equivalent to the statement that $|w|_1 = i+i'$ is odd. Similarly, when evaluated in a string $w \in L_{right}$ of the form $1^i 0^j 1^{i'} 2^j$, the formula φ' simulates the evaluation of the formula $\varphi_{\text{even cycles}}$ in a graph that consists of a single cycle of length $i+i'+1$; and hence φ' is satisfied iff $i+i'+1$ is odd, i.e., $|w|_1 = i+i'$ is even.

In summary, we obtain that φ is an $\langle\text{-inv-FO+MOD}_2(\sigma_\Sigma)\text{-sentence}$ that defines the language L . Since modulo 2 counting quantifiers can be simulated by modulo p counting quantifiers, for every *even* number $p \geq 2$, the proof of (b) is complete. \square

Proposition 4.4. *Let $L := L_{\text{even}} \cup L_{\text{odd}}$.*

- (a) *The language L is not Hanf $(\frac{n-1}{8})$ -local.*
 (b) *L is definable by a sentence φ in $\langle\text{-inv-FO+MOD}_p(\sigma_\Sigma)$, for every even number $p \geq 2$.*

Proof.

(a): For every $\ell \in \mathbb{N}_{\geq 1}$ let

$$\begin{aligned} u_\ell &:= 1^\ell 1^\ell 20^\ell 0^\ell 1^\ell 1^\ell 0^\ell 0^\ell = \mathbf{x} \mathbf{y}_1 \mathbf{y}_2 \mathbf{y}_3 \mathbf{z}, \\ v_\ell &:= 1^\ell 1^\ell 0^\ell 0^\ell 1^\ell 1^\ell 20^\ell 0^\ell = \mathbf{x} \mathbf{y}_3 \mathbf{y}_2 \mathbf{y}_1 \mathbf{z}, \end{aligned}$$

for $\mathbf{x} := 1^\ell$, $\mathbf{y}_1 := 1^\ell 20^\ell$, $\mathbf{y}_2 := 0^\ell 1^\ell$, $\mathbf{y}_3 := 1^\ell 0^\ell$, $\mathbf{z} := 0^\ell$.

It is not difficult to see that $\mathcal{S}_{u_\ell} \leftrightarrow_\ell \mathcal{S}_{v_\ell}$: the bijection β , for which

$$(\mathcal{N}_\ell^{\mathcal{S}_{u_\ell}}(c), c) \cong (\mathcal{N}_\ell^{\mathcal{S}_{v_\ell}}(\beta(c)), \beta(c)) \quad \text{for every } c \in \{1, \dots, |u_\ell|\},$$

can be chosen as follows. It maps each position of

- \mathbf{x} in u_ℓ onto the according position of \mathbf{x} in v_ℓ ,
- \mathbf{y}_s (for $s \in \{1, 2, 3\}$) in u_ℓ onto the according position of \mathbf{y}_s in v_ℓ ,
- \mathbf{z} in u_ℓ onto the according position of \mathbf{z} in v_ℓ .

It is straightforward to verify that this bijection β indeed witnesses that $\mathcal{S}_{u_\ell} \leftrightarrow_\ell \mathcal{S}_{v_\ell}$. Furthermore, $v_\ell \in L$ and $u_\ell \notin L$. The length of u_ℓ and v_ℓ is $n := 8\ell + 1$, thus $\ell = \frac{n-1}{8}$. Hence, the language L is not Hanf $(\frac{n-1}{8})$ -local. This completes the proof of (a).

(b): First, note that the language

$$M := L_{\text{left}} \cup L_{\text{right}}$$

can be defined by an $\text{FO}(\sigma_\Sigma)$ -sentence φ_M which states the following:

- The first position of the string carries the letter 1. The last position of the string carries the letter 0.
- For each position x that carries the letter 1, the position directly to the right of x carries one of the letters 0, 1, 2. Furthermore, there is exactly one position x that carries the letter 1, such that the position directly to the right of x carries the letter 0. And there is exactly one position x that carries the letter 1, such that the position directly to the right of x carries the letter 2.
- For each position y that carries the letter 0, the position directly to the right of y carries one of the letters 0, 1. Furthermore, there is exactly one position y that carries the letter 0, such that the position directly to the right of y carries the letter 1.
- There is exactly one position z that carries the letter 2. The position directly to the right of z carries the letter 0.

From Example 2.5 we obtain an $\langle\text{-inv-FO+MOD}_2(\{E\})\text{-sentence}$ $\varphi_{\text{even cycles}}$ that is satisfied by a finite $\{E\}$ -structure \mathcal{A} iff \mathcal{A} is a disjoint union of directed cycles where the number of cycles of even length is even. We choose

$$\varphi := \varphi_M \wedge \varphi',$$

where φ' is the formula obtained from $\varphi_{\text{even cycles}}$ by relativisation of all quantifiers to the positions that carry the letters 1 or 2, and by replacing every atom of the form $E(\mu, \nu)$ (for first-order variables μ and ν) by a formula stating that

- $E(\mu, \nu)$ is true, or
- position μ carries the letter 2, and ν is the leftmost position of the string, or
- the positions μ and ν both carry the letter 1, and the positions directly to the right of μ and directly to the left of ν both carry the letter 0.

Clearly, the obtained formula is order-invariant on the class of all finite σ_Σ -structures, since the formula $\varphi_{\text{even cycles}}$ is order-invariant on the class of all finite $\{E\}$ -structures.

It is straightforward to see that, when evaluated in a string $w \in L_{\text{left}}$ of the form $1^i 2 0^j 1^{i'} 0^{j'}$, the formula φ' simulates the evaluation of the formula $\varphi_{\text{even cycles}}$ in a graph that consists of the disjoint union of two cycles of lengths $i+1$ and i' ; and hence φ' is satisfied iff $i+1$ and i' are either both even or both odd — and this is equivalent to the statement that $|w|_1 = i+i'$ is odd. Similarly, when evaluated in a string $w \in L_{\text{right}}$ of the form $1^i 0^j 1^{i'} 2 0^{j'}$, the formula φ' simulates the evaluation of the formula $\varphi_{\text{even cycles}}$ in a graph that consists of a single cycle of length $i+i'+1$; and hence φ' is satisfied iff $i+i'+1$ is odd, i.e., $|w|_1 = i+i'$ is even.

In summary, we obtain that φ is an $<$ -inv-FO+MOD₂(σ_Σ)-sentence that defines the language L . Since modulo 2 counting quantifiers can be simulated by modulo p counting quantifiers, for every *even* number $p \geq 2$, the proof of (b) is complete. \square

Remark 4.5. Benedikt and Segoufin [6] conjectured that order-invariant FO+MOD (i.e. formulas with arbitrary modulo counting quantifiers) has the same expressive power as FO+MOD on trees. The previous proposition shows that, for even numbers $p \geq 2$, there exist order-invariant FO+MOD _{p} -definable languages which are not FO+MOD-definable, since FO+MOD is Hanf local. This refutes the conjecture even for strings instead of trees.

From Niemistö's Corollary 7.25 in [14] it follows that for *odd* numbers p , order-invariant FO+MOD _{p} (σ_Σ) on Σ -strings has exactly the same expressive power as FO+MOD_{PFC(p)}(σ_Σ), where PFC(p) is the set of all numbers whose prime factors are prime factors of p , and FO+MOD_{PFC(p)} is first-order logic with modulo m counting quantifiers for all $m \in \text{PFC}(p)$.

The present section's main result shows that for *odd prime powers* p , the Hanf locality result of [1] can be generalised from arb-invariant FO to arb-invariant FO+MOD _{p} on Σ -strings:

Theorem 4.6. *Let Σ be a finite alphabet. Let $k \in \mathbb{N}$, let q be a k -ary query, and let p be an odd prime power. If q is definable in arb-inv-FO+MOD _{p} (σ_Σ) on Σ -strings, then there is a $c \in \mathbb{N}$ such that q is Hanf $(\log n)^c$ -local on Σ -strings.*

Together with Theorem 4.2 this implies (general instead of weak) Gaifman locality on Σ -strings:

Corollary 4.7. *Let Σ be a finite alphabet. Let $k \in \mathbb{N}_{\geq 1}$, let q be a k -ary query, and let p be an odd prime power. If q is definable in arb-inv-FO+MOD _{p} (σ_Σ) on Σ -strings, then there is a $c \in \mathbb{N}$ such that q is Gaifman $(\log n)^c$ -local on Σ -strings.*

Note that this corollary does not contradict the non-locality result of Proposition 3.2, as the counter-example given in the proof of that proposition is not a string structure.

The remainder of this section is devoted to the proof of Theorem 4.6. We follow the overall approach of [1]. The crucial step is to prove Theorem 4.6 for queries q of arity $k = 0$.




Figure 5: The picture shows how the disjoint r -swap operation turns the edge relation of the string structure of a string $w = xuyvz$ (top) into the edge relation of the string structure of $w' = xvyuz$ (middle and bottom).

The case for queries of arity $k \geq 1$ can then easily be reduced to the case for queries of arity 0 by adding k extra symbols to the alphabet; see below.

Note that a 0-ary query q defines the string-language $L_q := \{w \in \Sigma^+ : () \in q(\mathcal{S}_w)\}$, where $()$ denotes the unique tuple of arity 0. The language L_q is called *Hanf $f(n)$ -local* iff q is Hanf $f(n)$ -local on Σ -strings. For proving Theorem 4.6 for the case $k = 0$, we consider the following notion.

Definition 4.8 (Disjoint swaps [1]). Let $r \in \mathbb{N}$ and let $w \in \Sigma^+$ be a string over a finite alphabet Σ . A string $w' \in \Sigma^+$ is obtained from w by a *disjoint r -swap operation* if there exist strings x, u, y, v, z such that $w = xuyvz$ and $w' = xvyuz$, and for the positions i, j, i', j' of w just before u, y, v, z the following is true: The neighbourhoods $N_r^{\mathcal{S}_w}(i), N_r^{\mathcal{S}_w}(j), N_r^{\mathcal{S}_w}(i'), N_r^{\mathcal{S}_w}(j')$ are pairwise disjoint, and $(\mathcal{N}_r^{\mathcal{S}_w}(i), i) \cong (\mathcal{N}_r^{\mathcal{S}_w}(i'), i')$ and $(\mathcal{N}_r^{\mathcal{S}_w}(j), j) \cong (\mathcal{N}_r^{\mathcal{S}_w}(j'), j')$.

Let $f : \mathbb{N} \rightarrow \mathbb{N}$. A string-language $L \subseteq \Sigma^+$ is *closed under disjoint $f(n)$ -swaps* if there exists an $n_0 \in \mathbb{N}$ such that for every $n \in \mathbb{N}_{\geq 1}$ with $n \geq n_0$, all strings $w \in \Sigma^+$ of length n , and all strings w' obtained from w by a disjoint $f(n)$ -swap operation, we have: $w \in L \iff w' \in L$.

It was shown in [1] (see Proposition 5.7, Lemma 5.2, and the proof of Theorem 5.1 in [1]) that if a language $L \subseteq \Sigma^+$ is closed under disjoint $(\log n)^d$ -swaps, for some $d \in \mathbb{N}$, then it is Hanf $(\log n)^c$ -local, for some $c > d$. Hence, the following lemma immediately implies Theorem 4.6 for the case $k = 0$.

Lemma 4.9. *Let Σ be a finite alphabet, let $L \subseteq \Sigma^+$, and let p be an odd prime power. If L is definable by an arb-inv-FO+MOD $_p(\sigma_\Sigma)$ -sentence, then there exists a constant $d \in \mathbb{N}$ such that L is closed under disjoint $(\log n)^d$ -swaps.*

Proof. We proceed in the same way as in the proof of Proposition 5.5 in [1], which obtained the analogue of Lemma 4.9 for arb-inv-FO(σ_Σ)-sentences. However, we cannot just copy the proof from there, since that proof relies on (general) Gaifman locality (with polylogarithmic locality radius) of queries definable in arb-inv-FO(σ_Σ), while in the present case we have

available (from Theorem 3.5) only *weak* Gaifman locality (with polylogarithmic locality radius) of queries definable in $\text{arb-inv-FO+MOD}_p(\sigma_\Sigma)$.

Let φ be an $\text{arb-inv-FO+MOD}_p(\sigma_\Sigma)$ -sentence which defines a string language $L_\varphi = \{w \in \Sigma^+ : \mathcal{S}_w \models \varphi\}$. For contradiction, assume that there is no $d \in \mathbb{N}$ such that L_φ is closed under disjoint $(\log n)^d$ -swaps. For each fixed $d, n_0 \in \mathbb{N}$ let $n \geq n_0$ and let w, w' be strings of length n which witness the violation of the ‘‘closure under disjoint $(\log n)^d$ -swaps’’ property. That is, $w \in L_\varphi, w' \notin L_\varphi$, and w' is obtained from w by a disjoint $(\log n)^d$ -swap operation. Thus, there exist strings $\mathbf{x}, \mathbf{u}, \mathbf{y}, \mathbf{v}, \mathbf{z}$ such that $w = \mathbf{x}\mathbf{y}\mathbf{v}\mathbf{z}$ and $w' = \mathbf{x}\mathbf{v}\mathbf{y}\mathbf{z}$, and for the positions i, j, i', j' of w just before $\mathbf{u}, \mathbf{y}, \mathbf{v}, \mathbf{z}$ the following is true for $r := (\log n)^d$: the neighbourhoods $N_r^{\mathcal{S}_w}(i), N_r^{\mathcal{S}_w}(j), N_r^{\mathcal{S}_w}(i'), N_r^{\mathcal{S}_w}(j')$ are pairwise disjoint, and $(\mathcal{N}_r^{\mathcal{S}_w}(i), i) \cong (\mathcal{N}_r^{\mathcal{S}_w}(i'), i')$ and $(\mathcal{N}_r^{\mathcal{S}_w}(j), j) \cong (\mathcal{N}_r^{\mathcal{S}_w}(j'), j')$.

The overall proof idea is as follows:

- (1) Choose an appropriate extension $\tilde{\sigma}$ of the signature σ_Σ ,
- (2) modify the formula φ into a suitable $\text{FO+MOD}_p(\tilde{\sigma}_{\text{arb}})$ -formula $\psi(x_1)$ with one free variable, and
- (3) define for each string w and all tuples $\bar{p} := (i, i', j, j')$ of positions of w which satisfy $0 \leq i < i' < j < j' < |w|$ a $\tilde{\sigma}$ -structure $\mathcal{A} := \mathcal{A}_w^{\bar{p}}$ with the same universe as \mathcal{S}_w , such that the positions $a := i+1$ and $a' := i'+1$ have disjoint and isomorphic $((\log n)^d - 1)$ -neighbourhoods in \mathcal{A} if the $(\log n)^d$ -neighbourhoods of i and i' and those of j and j' in \mathcal{S}_w are isomorphic and all these $(\log n)^d$ -neighbourhoods are pairwise disjoint

such that the following is satisfied:

- (4) $\psi(x_1)$ is arb-invariant on \mathcal{A} ,
- (5) $\mathcal{A} \models \psi[a] \iff \mathcal{S}_w \models \varphi$,
- (6) $\mathcal{A} \models \psi[a'] \iff \mathcal{S}_{w'} \models \varphi$.

Note that $w \in L_\varphi$ and $w' \notin L_\varphi$ imply that $\mathcal{A} \models \psi[a]$ and $\mathcal{A} \not\models \psi[a']$. In combination with (3) this shows that the unary query defined by $\psi(x_1)$ is *not* weakly Gaifman $((\log n)^d - 1)$ -local on the class \mathfrak{C} containing all structures $\mathcal{A}_w^{\bar{p}}$ which we defined above. Property (4) is true for each choice of the string w and positions \bar{p} . Hence, $\psi(x_1)$ is arb-invariant on \mathfrak{C} . Since d can be chosen arbitrarily large, and our choice of the formula $\psi(x_1)$ will not depend on d , this contradicts the fact that the formula $\psi(x_1)$ is weakly Gaifman local on \mathfrak{C} with polylogarithmic locality radius according to Theorem 3.5.

The details described in items (1)–(4) are carried out as follows.

ad (1): Let $\tilde{\sigma} := \sigma_\Sigma \cup \{F, X, Y_1, Y_2, Z\}$, where F is a binary relation symbol and X, Y_1, Y_2, Z are unary relation symbols. Thus, $\tilde{\sigma} = \{E, F, X, Y_1, Y_2, Z\} \cup \{P_a : a \in \Sigma\}$.

ad (3): Let $\mathcal{A} := \mathcal{A}_w^{\bar{p}}$ be the $\tilde{\sigma}$ -structure defined as follows (an illustration can be found in Figure 6):

- \mathcal{A} has the same universe as \mathcal{S}_w , i.e., $A = \{1, \dots, |w|\}$.
- $E^{\mathcal{A}}$ is obtained from $E^{\mathcal{S}_w}$ by removing all edges between the strings $\mathbf{x}, \mathbf{u}, \mathbf{y}, \mathbf{v}, \mathbf{z}$, i.e. $E^{\mathcal{A}} := E^{\mathcal{S}_w} \setminus \{(i, i+1), (j, j+1), (i', i'+1), (j', j'+1)\}$.
- $F^{\mathcal{A}}$ relates the first and last position of \mathbf{u} , and the first and last position of \mathbf{v} , i.e. $F^{\mathcal{A}} := \{(i+1, j), (i'+1, j')\}$.
- $X^{\mathcal{A}}$ marks the last position of \mathbf{x} , $Y_1^{\mathcal{A}}$ and $Y_2^{\mathcal{A}}$ mark the first and the last position of \mathbf{y} , and $Z^{\mathcal{A}}$ marks the first position of \mathbf{z} , i.e. $X^{\mathcal{A}} := \{i\}, Y_1^{\mathcal{A}} := \{j+1\}, Y_2^{\mathcal{A}} := \{i'\}, Z^{\mathcal{A}} := \{j'+1\}$.

- for each $a \in \Sigma$, $P_a^{\mathcal{A}}$ is identical to $P_a^{\mathcal{S}_w}$, i.e. $P_a^{\mathcal{A}} = P_a^{\mathcal{S}_w}$.

It is easy to see that, if in \mathcal{S}_w the $(\log n)^d$ -neighbourhoods of i and i' are isomorphic and the $(\log n)^d$ -neighbourhoods of j and j' are isomorphic and the $(\log n)^d$ -neighbourhoods of i, i', j, j' are pairwise disjoint, then in \mathcal{A} the $((\log n)^d - 1)$ -neighbourhoods of $a := i+1$ and $a' := i'+1$ are disjoint and isomorphic.




Figure 6: Structure \mathcal{A} (top) of Lemma 4.9 and the edge relations simulated by the formula $\psi(x_1)$ if x_1 is assigned the value $a := i+1$ (middle) or the value $a' := i'+1$ (bottom). F -edges in \mathcal{A} are depicted by dotted arcs; the E -edges in \mathcal{A} are given as successor relations within each of the framed boxes.

ad (2) and (4), (5), (6): We define $\psi(x_1)$ in such a way that, when evaluated in one of the structures $\mathcal{A} := \mathcal{A}_w^{\bar{p}}$, it does the following:

- if the variable x_1 is assigned the value $a := i+1$, then $\psi(x_1)$ simulates φ on \mathcal{S}_w .
- if the variable x_1 is assigned the value $a' := i'+1$, then $\psi(x_1)$ simulates φ on $\mathcal{S}_{w'}$.
- if x_1 is assigned to a value different from a, a' , then $\psi(x_1)$ is not satisfied in \mathcal{A} .

The first two items imply items (5) and (6) above. Together with the last item and the arb-invariance of φ on \mathcal{S}_w and $\mathcal{S}_{w'}$, this implies that item (4) is satisfied.

We let $\psi(x_1) := \exists \tilde{x}_1 \exists x_2 \exists \tilde{x}_2 \psi'$, where ψ' is a conjunction of formulas stating that:

- $F(x_1, \tilde{x}_1)$ holds. (This ensures that $\psi(x_1)$ is satisfied in \mathcal{A} only if x_1 is assigned one of the values $a := i+1$ or $a' := i'+1$.)
- $x_2 \neq x_1$ and $F(x_2, \tilde{x}_2)$ holds.
- The formula $\varphi'(x_1, \tilde{x}_1, x_2, \tilde{x}_2)$ is satisfied, where φ' is obtained from φ by replacing each atom of the form $E(\mu, \nu)$ with the formula $\chi_E(x_1, \tilde{x}_1, x_2, \tilde{x}_2, \mu, \nu)$ defined as follows:
 χ_E is the disjunction of formulas stating that
 - $E(\mu, \nu)$ holds (i.e., there already exists an E -edge from μ to ν in \mathcal{A}),
 - $X(\mu)$ holds and $\nu = x_1$ (i.e., a new E -edge from the last position of the string x to the position assigned to the variable x_1 is introduced)
 - $\mu = \tilde{x}_1$ and $Y_1(\nu)$ holds (i.e., a new E -edge from the position assigned to \tilde{x}_1 to the first position of the string y is introduced)
 - $Y_2(\mu)$ and $\nu = x_2$ (i.e., a new E -edge from the last position of the string y to the position assigned to the variable x_2 is introduced)

- $\mu = \tilde{x}_2$ and $Z(\nu)$ (i.e., a new E -edge from the position assigned to the variable \tilde{x}_2 to the first position of the string \mathbf{z} is introduced).

It is not difficult to see that χ_E simulates the edge relation of \mathcal{S}_w in \mathcal{A} if the variable x_1 is assigned the value $a := i+1$; and it simulates the edge relation of $\mathcal{S}_{w'}$ if the variable x_1 is assigned the value $a' := i'+1$ (see Figure 6 for an illustration).

This finishes the construction for items (2) and (4), (5), (6). In summary, the proof of Lemma 4.9 is complete. \square

Now that we have proved Theorem 4.6 for the case $k = 0$, we explain how the general case can be obtained from this. To this end, we convert a k -ary query on strings to a language over an extended alphabet with the same relevant definability and locality properties. This can be done using a standard technique which encodes variable assignments in an extended alphabet.

For each alphabet Σ and each $k \geq 1$, we let $\Sigma_{\text{var}(k)} := \Sigma \times 2^{[k]}$. The subsets of $[k]$ are used to encode an assignment of k variables to the positions of a string over the alphabet Σ . We let $L_{\text{assign}(k)}$ denote the language of strings $w \in \Sigma_{\text{var}(k)}^+$ where for each number $i \in [k]$ there is a unique position $\text{occ}_i(w) \in \{1, \dots, |w|\}$ such that the label of $\text{occ}_i(w)$ is $(a, X) \in \Sigma_{\text{var}(k)}$ with $i \in X$. Note that this language is clearly $\text{FO}[\sigma_{\Sigma_{\text{var}(k)}}]$ -definable. For each string $w \in L_{\text{assign}(k)}$, we let $\tilde{w} \in \Sigma^+$ denote the string where each symbol (a, X) is replaced by a . Furthermore, we let $\overline{\text{occ}}_k(w) := (\text{occ}_0(w), \dots, \text{occ}_{k-1}(w))$.

Lemma 4.10. *Let Σ be a finite alphabet and let $k \in \mathbb{N}_{\geq 1}$. With each k -ary query q on Σ -strings we associate the language*

$$A_q := \{w \in L_{\text{assign}(k)} : \overline{\text{occ}}_k(w) \in q(S_{\tilde{w}})\}.$$

Then

- (1) For each $p \in \mathbb{N}_{\geq 1}$, if q is definable in $\langle\text{-inv-FO+MOD}_p(\sigma_\Sigma)$ on Σ -strings, then A_q is definable in $\langle\text{-inv-FO+MOD}_p(\sigma_{\Sigma_{\text{var}(k)}})$.
- (2) The query q is Hanf $f(n)$ -local on Σ -strings iff the language A_q is Hanf $f(n)$ -local.

Proof.

Ad (1): Let $\varphi(\bar{x})$ be a formula of $\langle\text{-inv-FO+MOD}_p(\sigma_\Sigma)$ defining q , where $\bar{x} := (x_0, \dots, x_{k-1})$. For each $i \in [k]$, let $\psi_{\text{occ},i}(x) := \bigvee_{a \in \Sigma, X \in 2^{[k]}, i \in X} P_{(a,X)}(x)$. That is, the formula states that i occurs in the set in the second component of the label at position x . Let

$$\hat{\varphi} := \exists x_0 \cdots \exists x_{k-1} \bigwedge_{i \in [k]} \psi_{\text{occ},i}(x_i) \wedge \varphi',$$

where φ' is obtained from φ by replacing each occurrence of a relation symbol $P_a(x)$, for $a \in \Sigma$, by the formula $\bigvee_{X \in 2^{[k]}} P_{(a,X)}(x)$ which states for a position x of \mathcal{S}_w , for each string $w \in \Sigma_{\text{var}(k)}^+$, that the label at position x in \tilde{w} is a . Let ψ be an FO-sentence which defines $L_{\text{assign}(k)}$ on $\Sigma_{\text{var}(k)}$ -strings. It is straightforward to verify that $\hat{\varphi} \wedge \psi$ is a formula of $\langle\text{-inv-FO+MOD}_p(\sigma_{\Sigma_{\text{var}(k)}})$ which defines the language A_q .

Ad (2): Note that $\mathcal{S}_u \stackrel{f(n)}{\leftrightarrow} \mathcal{S}_v$ iff $(\mathcal{S}_{\tilde{u}}, \overline{\text{occ}}_k(u)) \stackrel{f(n)}{\leftrightarrow} (\mathcal{S}_{\tilde{v}}, \overline{\text{occ}}_k(v))$, for all $u, v \in L_{\text{assign}(k)}$ of length n . More concretely, a bijection β witnesses the first statement iff it witnesses the second statement. This follows from the fact that β preserves the second component of the labelling of \mathcal{S}_u (i.e. the label of $\beta(x)$ is the same as the

label of x , for all elements $x \in \{1, \dots, n\}$ iff $\beta(\text{occ}_i(u)) = \text{occ}_i(v)$, for all $i \in [k]$, and that the label of each position of \tilde{u} and \tilde{v} is the first component of the labelling of that position in u and v .

Note also that $L_{\text{assign}(k)}$ is 0-local since its definition depends only on the labelling.

First, we show that locality of q implies locality of A_q . Suppose that $\mathcal{S}_u \xleftrightarrow{f(n)} \mathcal{S}_v$, for $u, v \in \Sigma_{\text{var}(k)}^n$. We have $u \in L_{\text{assign}(k)}$ iff $v \in L_{\text{assign}(k)}$. If $u, v \notin L_{\text{assign}(k)}$, then $u, v \notin A_q$. Suppose that $u, v \in L_{\text{assign}(k)}$. Then $(\mathcal{S}_{\tilde{u}}, \overline{\text{occ}}_k(u)) \xleftrightarrow{f(n)} (\mathcal{S}_{\tilde{v}}, \overline{\text{occ}}_k(v))$ and we obtain that $\overline{\text{occ}}_k(u) \in q(\mathcal{S}_{\tilde{u}})$ iff $\overline{\text{occ}}_k(v) \in q(\mathcal{S}_{\tilde{v}})$, by $f(n)$ -locality of q , and so $u \in A_q$ iff $v \in A_q$.

Now we show that locality of A_q implies locality of q . Let $u, v \in \Sigma^n$ and $\bar{a} = (a_0, \dots, a_{k-1}), \bar{b} := (b_0, \dots, b_{k-1}) \in \{1, \dots, n\}^k$ be such that $(\mathcal{S}_u, \bar{a}) \xleftrightarrow{f(n)} (\mathcal{S}_v, \bar{b})$. There exist $u', v' \in L_{\text{assign}(k)}$ such that $u = \tilde{u}'$ and $v = \tilde{v}'$ and such that $a_i = \text{occ}_i(u')$ and $b_i = \text{occ}_i(v')$, for each $i \in [k]$. Then $(\mathcal{S}_u, \bar{a}) = (\mathcal{S}_{\tilde{u}'}, \overline{\text{occ}}_k(u'))$ and $(\mathcal{S}_v, \bar{b}) = (\mathcal{S}_{\tilde{v}'}, \overline{\text{occ}}_k(v'))$. Since $(\mathcal{S}_u, \bar{a}) \xleftrightarrow{f(n)} (\mathcal{S}_v, \bar{b})$, we also have $\mathcal{S}_{u'} \xleftrightarrow{f(n)} \mathcal{S}_{v'}$. By Hanf $f(n)$ -locality of A_q , this means that $u' \in A_q$ iff $v' \in A_q$. Hence, $\overline{\text{occ}}_k(u') \in q(\mathcal{S}_{\tilde{u}'})$ iff $\overline{\text{occ}}_k(v') \in q(\mathcal{S}_{\tilde{v}'})$. That is, $\bar{a} \in q(\mathcal{S}_u)$ iff $\bar{b} \in q(\mathcal{S}_v)$ and hence q is Hanf $f(n)$ -local on Σ -strings. \square

The proof of Theorem 4.6 can now be finished as follows. Suppose that there is a k -ary query q , for some $k > 0$, which is definable in $\text{arb-inv-FO+MOD}_p(\sigma_\Sigma)$ on Σ -strings, for some odd prime power p and some finite alphabet Σ , such that there is no $c \in \mathbb{N}$ for which q is Hanf $(\log n)^c$ -local on Σ -strings. By the previous lemma, we obtain a language A_q which is definable in $\text{arb-inv-FO+MOD}_p(\sigma_{\Sigma_{\text{var}(k)}})$ and which is also not Hanf $(\log n)^c$ -local. But we have already proved that this is impossible.

5. CONCLUSION

We have introduced a new notion of locality called *shift locality* which generalises Niemistö's notion of alternating locality with constant locality radius [14] and Libkin's notion of weak Gaifman locality [12]. We have presented a comprehensive picture of the locality of arb-inv-FO+MOD_p for prime powers p .

We have also investigated notions of locality on string structures. Here some natural questions remain open. We have shown that there is an <-inv-FO+MOD_2 -definable language which is not Hanf local and hence not FO+MOD_q -definable for any modulus q . It would be interesting to understand the expressive power of <-inv-FO+MOD_2 on strings in more detail. Is there a decidable algebraic characterisation of the <-inv-FO+MOD_2 -definable languages? Is there a logic with the same expressive power as <-inv-FO+MOD_2 but with an effective syntax?

We have derived Hanf locality on string structures from weak Gaifman locality. To accomplish this, we have used the characterisation of Hanf locality in terms of disjoint swap operations. The origin of these swap operations goes back to [5] where similar operations are, among other things, used to obtain a new proof of an algebraic characterisation of first-order logic on string structures from [4]. As a first step towards proving an algebraic characterisation of <-inv-FO+MOD_2 , we believe that one could define a variant of Hanf locality whose relation to shift locality is similar to the relation of Hanf locality and weak Gaifman locality that we established. Then one could try to characterise this notion of locality by a variant of the disjoint swap operations.

REFERENCES

- [1] M. Anderson, D. van Melkebeek, N. Schweikardt, and L. Segoufin. Locality from circuit lower bounds. *SIAM Journal on Computing*, 41(6):1481–1523, 2012.
- [2] S. Arora and B. Barak. *Computational Complexity: A Modern Approach*. Cambridge Univ. Press, 2009.
- [3] D. A. M. Barrington, N. Immerman, and H. Straubing. On uniformity within NC^1 . *J. Comput. Syst. Sci.*, 41(3):274–306, 1990.
- [4] D. Beauquier and J.-E. Pin. Factors of words. In *Proc. ICALP*, pages 63–79, 1989.
- [5] M. Benedikt and L. Segoufin. Regular tree languages definable in FO and FOMod. *ACM Transactions on Computational Logic*, 11(1), 2009.
- [6] M. Benedikt and L. Segoufin. Towards a characterization of order-invariant queries over tame graphs. *Journal of Symbolic Logic*, 74(1):pp. 168–186, 2009.
- [7] H.-D. Ebbinghaus and J. Flum. *Finite model theory*. Springer, 1999.
- [8] R. Fagin, L. J. Stockmeyer, and M. Y. Vardi. On Monadic NP vs. Monadic co-NP. *Information and Computation*, 120(1):78–92, 1995.
- [9] M. Grohe and T. Schwentick. Locality of order-invariant first-order formulas. *ACM Transactions on Computational Logic*, 1(1):112–130, 2000.
- [10] F. Harwath and N. Schweikardt. On the locality of arb-invariant first-order logic with modulo counting quantifiers. In *Computer Science Logic 2013, CSL 2013*, volume 23 of *LIPICs*, pages 363–379. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013.
- [11] L. Hella, L. Libkin, and J. Nurmonen. Notions of locality and their logical characterizations over finite models. *Journal of Symbolic Logic*, 64(4):1751–1773, 1999.
- [12] L. Libkin. *Elements of Finite Model Theory*. Springer, 2004.
- [13] M. Lothaire. *Combinatorics on words*. Cambridge University Press, 1984.
- [14] H. Niemistö. *Locality and Order-Invariant Logics*. PhD thesis, Department of Mathematics and Statistics, University of Helsinki, 2007.
- [15] J. J. Rotman. *An Introduction to the Theory of Groups*. Springer, 4th edition, 1994.
- [16] N. Schweikardt. A short tutorial on order-invariant first-order logic. In *Computer Science - Theory and Applications - 8th International Computer Science Symposium in Russia, CSR 2013, Ekaterinburg, Russia, June 25-29, 2013. Proceedings*, pages 112–126, 2013.
- [17] R. Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, STOC '87, pages 77–82, New York, NY, USA, 1987. ACM.
- [18] H. Straubing. *Finite Automata, Formal Logic, and Circuit Complexity*. Birkhäuser, 1994.
- [19] R. Williams. Guest column: A casual tour around a circuit complexity bound. *SIGACT News*, 42(3):54–76, 2011.