A Short Tutorial on Order-Invariant First-Order Logic

Nicole Schweikardt

Goethe-Universität Frankfurt am Main, Germany
schweika@informatik.uni-frankfurt.de

Abstract. This paper gives a short introduction to order-invariant first-order logic and arb-invariant first-order logic. We present separating examples demonstrating the expressive power, as well as tools for proving certain expressive weaknesses of these logics.

1 Introduction

Expressibility 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. Order-invariant and arb-invariant logics 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.

Arb-invariant queries are queries that can make use of an order predicate $<$ and of arithmetic predicates such as $+$ or $\times$, but only in such a way that the answer is independent of the particular interpretation of $<, +, \times$. Queries that only use the linear order, but no further arithmetic predicates, are called *order-invariant*.

It is known that order-invariant least fixed-point logic LFP precisely captures the polynomial time computable queries [12,26], while arb-invariant LFP and arb-invariant first-order logic capture the queries computable in $\mathrm{P}_{/\mathrm{poly}}$ [15] and AC^0 [13], respectively. Order-invariant queries and arb-invariant queries have been studied in depth, cf. e.g. [1,5,25,15,8,14,17,22,16,19,20,7,4,24,9,2]. A short overview of the state-of-the-art concerning these logics can be found in [23].

The aim of this paper is to give a short tutorial on order-invariant and arb-invariant first-order logic FO. In Section 2 we fix the basic notation. Section 3 gives the precise definition of order-invariant and arb-invariant FO, along with a few easy examples. Section 4 presents examples that separate order-invariant FO from plain FO. Section 5 shows how to prove that certain queries are not definable in order- or arb-invariant FO. Section 6 gives a list of open research questions.

2 Preliminaries

Basic Notation. We write $\mathbb{N}$ for the set of non-negative integers, and we let $\mathbb{N}_{\geq 1} := \mathbb{N} \setminus \{0\}$. For $n \in \mathbb{N}_{\geq 1}$ we write $[n]$ to denote the set $\{i \in \mathbb{N} : 0 \leq i < n\}$, i.e.,

A.A. Bulatov and A.M. Shur (Eds.): CSR 2013, LNCS 7913, pp. 112–126, 2013.

$[n] = \{0, \ldots, n-1\}$. For a positive real number r, the logarithm of r with respect to base 2 is denoted $\log r$.

For a finite set A we write $|A|$ to denote the cardinality of A. 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 words built from symbols in A is denoted A^+. We write $|w|$ for the length of a word $w \in A^+$.

Structures. A *signature* σ is a set of relation symbols R, each of them associated with a fixed arity $ar(R) \in \mathbb{N}_{\geq 1}$. A *$\sigma$-structure* $\mathcal{A}$ consists of a non-empty set A called the *universe* of $\mathcal{A}$, and a relation $R^{\mathcal{A}} \subseteq A^{ar(R)}$ 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 $\overline{a} = (a_1, \ldots, a_k) \in A^k$ and $\overline{b} = (b_1, \ldots, b_k) \in B^k$ we write $(\mathcal{A}, \overline{a}) \cong (\mathcal{B}, \overline{b})$ to indicate that there is an isomorphism π from $\mathcal{A}$ to $\mathcal{B}$ that maps $\overline{a}$ to $\overline{b}$ (i.e., $\pi(a_i) = b_i$ for each $i \leq k$).

First-Order Logic. We assume that the reader is familiar with basic concepts and notations concerning first-order logic (cf., e.g., the textbooks [14,6]). We write FO(σ) to denote the class of all first-order formulas of signature σ. By *free*(φ) we denote the set of all free variables of an FO(σ)-formula φ. A *sentence* is a formula φ with *free*(φ) $= \emptyset$. We often write $\varphi(\overline{x})$, for $\overline{x} = (x_1, \ldots, x_k)$, to indicate that *free*(φ) $= \{x_1, \ldots, x_k\}$.

If $\mathcal{A}$ is a σ-structure and $\overline{a} = (a_1, \ldots, a_k) \in A^k$, we write $\mathcal{A} \models \varphi[\overline{a}]$ to indicate that the formula $\varphi(\overline{x})$ is satisfied in $\mathcal{A}$ when interpreting the free occurrences of the variables $x_1, \ldots, x_k$ with the elements $a_1, \ldots, a_k$.

If $\mathcal{A}$ and $\mathcal{B}$ are σ-structures and r is a natural number, we write $\mathcal{A} \equiv_r \mathcal{B}$ to indicate that $\mathcal{A}$ and $\mathcal{B}$ satisfy exactly the same FO(σ)-sentences of quantifier rank r.

Throughout the remainder of this paper, we will assume that σ is a fixed finite signature.

3 Order-Invariant Logic and Arb-Invariant Logic

The idea: 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.

Definition 3.1 (Numerical predicate)
For $r \in \mathbb{N}_{\geq 1}$, an r-ary *numerical predicate* is an r-ary relation on $\mathbb{N}$.

Three particular numerical predicates that will often be used in this paper are

- the linear order $<^{\mathbb{N}}$ consisting of all tuples $(a, b) \in \mathbb{N}^2$ with $a < b$,
- the addition predicate $+^{\mathbb{N}}$ consisting of all triples $(a, b, c) \in \mathbb{N}^3$ with $a + b = c$, and
- the multiplication predicate $\times^{\mathbb{N}}$ consisting of all triples $(a, b, c) \in \mathbb{N}^3$ with $a \times b = c$.

To allow logical formulas to use numerical predicates, we fix the following notation: For every $r \in \mathbb{N}_{\geq 1}$ and for every r-ary numerical predicate $P^{\mathbb{N}}$ (i.e., $P^{\mathbb{N}} \subseteq \mathbb{N}^r$), we let P be a new relation symbol of arity r — here, "new" means 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 we let σ_{arb} be the disjoint union of σ and η_{arb} (the subscript "arb" stands for "arbitrary numerical predicates").

Next, we would like to allow FO(σ_{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 $n = |A|$, i.e., the set $[n] = \{0, \ldots, n{-}1\}$.

Definition 3.2 (Embedding). Let $\mathcal{A}$ be a finite σ-structure, and let $n := |A|$.
An *embedding* ι of $\mathcal{A}$ is a bijection $\iota : A \to [n]$.

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: The linear order $<^{\mathbb{N}}$ induces a linear order $<^{\iota}$ on A where for all $a, b \in A$ we have $a <^{\iota} b$ iff $\iota(a) < \iota(b)$. The addition predicate $+^{\mathbb{N}}$ induces an addition predicate $+^{\iota}$ on A where for all $a, b, c \in A$ we have $(a, b, c) \in +^{\iota}$ iff $\iota(a) + \iota(b) = \iota(c)$. In general, an arbitrary r-ary numerical predicate $P^{\mathbb{N}}$ induces the r-ary predicate P^{ι} on A, consisting of all r-tuples $\overline{a} = (a_1, \ldots, a_r) \in A^r$ where $\iota(\overline{a}) = (\iota(a_1), \ldots, \iota(a_r)) \in P^{\mathbb{N}}$.

The σ_{arb}-structure $\mathcal{A}^{\iota}$ associated with $\mathcal{A}$ and ι is the expansion of $\mathcal{A}$ by the predicates P^{ι} for all $P \in \eta_{\text{arb}}$. I.e., $\mathcal{A}^{\iota}$ has the same universe as $\mathcal{A}$, all relation symbols $R \in \sigma$ are interpreted in $\mathcal{A}^{\iota}$ 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 FO(σ_{arb})-formula φ makes a meaningful statement about a σ-structure $\mathcal{A}$, we evaluate φ in $\mathcal{A}^{\iota}$, 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 3.3 (Arb-invariance and arb-inv-FO)
Let $\varphi(\overline{x})$ be an FO(σ_{arb})-formula with k free variables, and let $\mathcal{A}$ be a finite σ-structure.

(a) The formula $\varphi(\overline{x})$ is *arb-invariant on* $\mathcal{A}$ if for *all* embeddings ι_1 and ι_2 of $\mathcal{A}$ and for all tuples $\overline{a} \in A^k$ we have: $\mathcal{A}^{\iota_1} \models \varphi[\overline{a}] \iff \mathcal{A}^{\iota_2} \models \varphi[\overline{a}]$.
(b) Let $\varphi(\overline{x})$ be arb-invariant on $\mathcal{A}$.
We write $\mathcal{A} \models \varphi[\overline{a}]$, if $\mathcal{A}^{\iota} \models \varphi[\overline{a}]$ for some (i.e., every) embedding ι of $\mathcal{A}$.
(c) $\varphi(\overline{x})$ is called *arb-invariant* if it is arb-invariant on every finite σ-structure $\mathcal{A}$.
(d) We write arb-inv-FO(σ) to denote the set of all arb-invariant FO(σ_{arb})-formulas.

Example 3.4. We present an arb-invariant FO(σ_{arb})-sentence $\varphi_{\textit{even}}$ which is satisfied by exactly those finite σ-structures that have even cardinality. The formula is chosen as follows:

$$\varphi_{\textit{even}} \ := \ \exists x \, \forall y \, \big((\, y{<}x \ \lor \ y{=}x \,) \ \land \ \textit{Odd}(x) \big),$$

where $\textit{Odd}^{\mathbb{N}}$ is the unary numerical predicate consisting of all odd numbers.

Let us consider a finite σ-structure $\mathcal{A}$ of size $n = |A|$ and an embedding ι of $\mathcal{A}$ into the set $[n] = \{0, \ldots, n{-}1\}$. Obviously, $\mathcal{A}^{\iota} \models \varphi_{\textit{even}}$ iff the maximum element in $[n]$, i.e., the number $n{-}1$, is odd, i.e., the number n is even. Thus, the formula $\varphi_{\textit{even}}$ is arb-invariant on $\mathcal{A}$, and it expresses that $\mathcal{A}$ is of even cardinality.

Note that $<^{\mathbb{N}}$ and $\textit{Odd}^{\mathbb{N}}$ are the only numerical predicates used by the formula $\varphi_{\textit{even}}$. Both predicates can be replaced by uses of the addition predicate $+^{\mathbb{N}}$, since $\textit{Odd}(x)$ is equivalent to $\neg \exists z \, z{+}z{=}x$, and $y{<}x$ is equivalent to $(\, \neg \, y{=}x \ \land \ \exists z \, y{+}z{=}x \,)$.

Thus, "even cardinality" of finite σ-structures can also be expressed by an arb-invariant FO(σ_{arb})-sentence that only uses the numerical predicate $+^{\mathbb{N}}$. Recall that it is well-known that "even cardinality" can neither be expressed by FO(σ), nor by arb-invariant FO(σ_{arb})-sentences that only use the numerical predicate $<^{\mathbb{N}}$ (cf., e.g., the textbooks [14,6], where it is shown that "even cardinality of linear orders" is not definable in first-order logic).

Definition 3.5 (Order-invariance and addition-invariance)

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

Example 3.4 shows that $<$-inv-FO(σ) is less expressive than $+$-inv-FO(σ).

It is known that for any signature σ that contains at least one symbol of arity ≥ 2, there is no algorithm that decides whether an input FO($\sigma \cup \{<\}$)-sentence is order-invariant (this can be shown by an easy reduction using Trakhtenbrot's theorem, see e.g. [14]). However, if σ contains only unary relation symbols, order-invariance of an input sentence is decidable, since commutativity of regular languages is decidable (via checking if the language's syntactic monoid is commutative).

Definition 3.6. An arb-invariant formula that only uses numerical predicates that belong to a subset S of η_{arb} is called S-*invariant*. By S-inv-FO(σ) we denote the set of all S-invariant FO($\sigma \cup S$)-formulas.

The next two examples show that $+$-inv-FO(σ) is less expressive than $\{+, \times\}$-inv-FO(σ) which, in turn, is less expressive than arb-inv-FO(σ).

Example 3.7. Consider the formula φ_{even} from Example 3.4. Let φ_{square} be the formula obtained from φ_{even} by replacing the atom $Odd(x)$ by $Square'(x)$, where $Square'^{\mathbb{N}} := \{i^2{-}1 \,:\, i \in \mathbb{N}_{\geq 1}\}$. Obviously, φ_{square} is an arb-invariant sentence satisfied by exactly those finite σ-structures whose cardinality is a square number.

Note that $i^2{-}1 = (i{-}1)^2 + 2(i{-}1)$. Thus, $Square'(x)$ is equivalent to

$$\exists y\, \exists z_1\, \exists z_2\, \big(y{\times}y{=}z_1 \;\wedge\; y{+}y{=}z_2 \;\wedge\; z_1{+}z_2{=}x \big).$$

Therefore, "square number cardinality" can be expressed in $\{+, \times\}$-inv-FO(σ). It is well-known that this cannot be expressed in $+$-inv-FO(σ) (since, by the theorem of Ginsburg and Spanier, FO($+$)-definable subsets of $\mathbb{N}$ are semi-linear; see e.g. [21] for an overview).

Example 3.8. It is straightforward to see that $\{+, \times\}$-inv-FO(σ)-sentences can only define *decidable* properties of finite σ-structures. However, arb-inv-FO(σ) can define also undecidable properties. For example, let $U^{\mathbb{N}}$ be an undecidable subset of $\mathbb{N}$ (such a set exists, since there are uncountably many subsets of $\mathbb{N}$, but only a countable number of decidable sets). Now let φ_U be the formula obtained from φ_{even} by replacing the atom $Odd(x)$ by $U'(x)$, where $U'^{\mathbb{N}} := \{i{-}1 \,:\, i \in U \text{ and } i \neq 0\}$. Clearly, φ_U is an arb-invariant sentence satisfied by exactly those finite σ-structures whose cardinality belongs to $U^{\mathbb{N}}$. As $U^{\mathbb{N}}$ is undecidable, also the class of finite σ-structures satisfying φ_U is undecidable.

The examples seen so far are simple in the sense that they only refer to the cardinality of structures and do not make use of the relation symbols present in σ. Showing that $<$-inv-FO(σ) is more expressive than plain FO(σ) requires much more sophisticated constructions and depends on the particular choice of the signature σ. In fact, if σ is the *empty* signature $\emptyset$ (as could have been chosen for the examples above), it is straightforward to see that $<$-inv-FO($\emptyset$) has exactly the same expressive power as FO($\emptyset$).

4 Three Examples Showing That Order-Invariant FO Is More Expressive Than FO

In the literature, basically only three examples are known that separate $<$-inv-FO(σ) from plain FO(σ), for various signatures σ. These examples go back to Gurevich (who did not publish this example; but it can be found in the textbooks [1,14]), Potthoff [18], and Otto [17].

Gurevich's Example. For a finite set X let $\mathcal{B}_X := (2^X, \subseteq)$ be the Boolean algebra over X. Thus, $\mathcal{B}_X$ is a finite σ-structure, where $\sigma := \{\subseteq\}$ is the signature consisting of a single binary relation symbol $\subseteq$.

Theorem 4.1 (Gurevich). *There is an order-invariant* FO($\sigma \cup \{<\}$)*-sentence* $\varphi_{Gurevich}$*, but no* FO(σ)*-sentence, such that for every finite set X we have: $\mathcal{B}_X \models \varphi_{Gurevich} \iff |X|$ is even.*

Proof sketch. Part 1: Construction of the $<$-inv-FO(σ)*-sentence* $\varphi_{Gurevich}$. An element $y \in 2^X$ is called an *atom* if it is a singleton set. Thus, $|X|$ is the number of atoms in 2^X. Obviously, the following FO(σ)-formula *atom*(x) expresses that x is an atom:

$$atom(x) \;:=\; \big(\neg emptyset(x) \;\wedge\; \forall y \;\big(\, y \subseteq x \;\rightarrow\; (\, y{=}x \;\vee\; emptyset(y)\,)\,\big)\big),$$

where $emptyset(y) := \forall z\; y \subseteq z$.

The order-invariant FO($\sigma \cup \{<\}$)-sentence $\varphi_{Gurevich}$ states that

(1) the underlying σ-structure is indeed a Boolean algebra $(2^X, \subseteq)$, and
(2) there exists a set $z \in 2^X$ that contains the first (w.r.t. $<$) atom of 2^X and every other atom (w.r.t. $<$) of 2^X, and that has the property that the last (w.r.t. $<$) atom of 2^X does not belong to z.

Note that the statements (1) and (2) ensure that $\varphi_{Gurevich}$ is order-invariant on the class of all finite σ-structures, and that a finite Boolean algebra $\mathcal{B}_X$ satisfies $\varphi_{Gurevich}$ if and only if $|X|$ is even. It is an easy exercise to express statement (1) by an FO(σ)-sentence and statement (2) by an FO($\sigma \cup \{<\}$)-sentence.

Part 2: Proof of the non-expressibility in FO(σ). By using a standard Ehrenfeucht-Fraïssé game argument one can show that $\mathcal{B}_{X_1} \equiv_r \mathcal{B}_{X_2}$ is true for all $r \in \mathbb{N}$ and all finite sets X_1 and X_2 of size at least 2^r. Thus, for every quantifier rank $r \in \mathbb{N}$, we can find sufficiently large finite sets X_1 and X_2 of odd and even cardinality, respectively, that cannot be distinguished by FO(σ)-sentences of quantifier rank r.

A detailed exposition of Gurevich's proof can be found in the textbook [14]. □

Potthoff's Example. We consider unordered finite binary trees T where every node is either a leaf or has exactly two children. The *height* of a leaf x of T is the length of the path from the root to x. The height of T is the largest height of a leaf of T. A tree T is *full* if all leaves are of the same height.

Let $\sigma := \{E, D\}$ be the signature consisting of two binary relation symobls E and D. We represent an unordered finite binary tree T by a σ-structure $\mathcal{A}_T$ whose universe is the set of nodes of T, and where E is the directed edge relation connecting every non-leaf node with its two children, and D is the descendant relation, i.e., the transitive closure of E.

Theorem 4.2 (Potthoff [18]). *There exists an order-invariant* FO($\sigma \cup \{<\}$)*-sentence* $\varphi_{\textit{Potthoff}}$*, but no* FO($\sigma$)*-sentence, such that for every full unordered finite binary tree* T *we have:* $\mathcal{A}_T \models \varphi_{\textit{Potthoff}} \iff T$ *is of even height.*

Proof sketch. Part 1: Construction of the $<$-inv-FO(σ)*-sentence* $\varphi_{\textit{Potthoff}}$. To keep the description of the formula simple, we here present a sentence $\varphi_{\textit{Potthoff}}$ that is order-invariant only on the class of all *full* unordered finite binary trees. A more sophisticated sentence that is order-invariant on all finite σ-structures is outlined below, after Lemma 4.3.

For constructing $\varphi_{\textit{Potthoff}}$ let us consider a full binary tree T of height h. We use the linear order $<$ to order the children of each node a of T: If b_1 and b_2 are a's children and $b_1 < b_2$, then b_1 is called *the 1-child*, and b_2 is called *the 2-child* of a. Now, we consider the *zig-zag-path* which starts in the root, visits the root's 1-child, that node's 2-child, that nodes 1-child, etc. I.e., the zig-zag-path is the path $(x_0, x_1, x_2, \ldots, x_h)$ where x_0 is the root, x_h is a leaf, and for odd $i \geq 1$, x_i is the 1-child of x_{i-1}, whereas for even $i \geq 1$, x_i is the 2-child of x_{i-1}.

As T is a *full* binary tree, the height h of T is even if and only if the last node of the zig-zag-path is a 2-child — and this is exactly the statement made by the formula $\varphi_{\textit{Potthoff}}$. Note that a formula making this statement will be order-invariant on the structure $\mathcal{A}_T$, for all *full* binary trees T.

The statement "the last node of the zig-zag-path is a 2-child" can be formalised by an FO($\sigma \cup \{<\}$)-sentence $\varphi_{\textit{Potthoff}}$, which states the following:

(1) There exists a node x_0 which is the root, and there exists a node x_h which is a leaf, such that
(2) x_h is the 2-child (w.r.t. $<$) of its parent,
(3) the node x_1 which satisfies $\big(E(x_0, x_1) \wedge D(x_1, x_h)\big)$ is the 1-child (w.r.t. $<$) of its parent, and
(4) for any three nodes u, v, w such that $E(u, v)$ and $E(v, w)$ and $\big(w{=}x_h \vee D(w, x_h)\big)$ we have that v is the 1-child of its parent iff w is the 2-child of its parent.

Part 2: Proof of the non-expressibility in FO(σ). By using a standard Ehrenfeucht-Fraïssé game argument, one can show that $\mathcal{A}_{T_1} \equiv_r \mathcal{A}_{T_2}$ is true for all $r \in \mathbb{N}$ and all full unordered finite binary trees T_1 and T_2 of height $\geq 2^{r+1}$ (the duplicator's winning strategy in the Ehrenfeucht-Fraïssé game is a straightforward generalisation of the winning strategy in the game on two linear orders, cf. e.g. [14]). Thus, for every $r \in \mathbb{N}$, we can find sufficiently big full unordered finite binary trees of odd and even height, respectively, that cannot be distinguished by FO(σ)-sentences of quantifier rank r. □

For completeness, let us give the precise statement of Potthoff's result. Instead of constructing an order-invariant FO(σ)-formula, Potthoff constructs an FO(σ')-formula for the signature $\sigma' = \sigma \cup \{C_1, C_2\}$, where C_1 and C_2 are unary relation symbols. An *ordered* finite binary tree T is represented by the σ'-structure $\mathcal{B}_T$ which is the expansion of the structure $\mathcal{A}_T$ by unary relations C_1 and C_2, where C_1 consists of all nodes which are the first child of their parent, and C_2 consists of all nodes which are the second child of their parent.

Lemma 4.3 (Lemma 5.1.8 in [18]). *There is an* FO(σ')*-sentence* $\psi_{\textit{Potthoff}}$ *such that for every ordered finite binary tree* T *we have:* $\mathcal{B}_T \models \psi_{\textit{Potthoff}} \iff$ *every leaf of* T *is of even height.*

The order-invariant sentence $\varphi_{\textit{Potthoff}}$ whose existence is claimed in Theorem 4.2 is now obtained as the conjunction of

- a straightforward FO(σ)-axiomatisation of unordered binary trees, and
- the formula obtained from $\psi_{\textit{Potthoff}}$ by replacing atoms of the form $C_i(x)$ (for $i \in \{1,2\}$) with an FO($E,<$)-formula stating that x is the i-child (w.r.t. $<$) of its parent.

Otto's Example. For every $n \in \mathbb{N}_{\geq 1}$ and every undirected graph G on $2n$ vertices, we consider a σ-structure $\mathcal{S}_{2n}(G)$ into which G is embedded. The signature $\sigma = \{E, \sim, \in, V, V', P'\}$ consists of three binary relation symbols $E, \sim, \in$ and three unary relation symbols V, V', P'.

We let $\sigma' := \sigma \setminus \{E\}$ and define, for each $n \in \mathbb{N}_{\geq 1}$, the σ'-structure $\mathcal{S}_{2n}$ as follows: The universe of $\mathcal{S}_{2n}$ is partitioned into three disjoint sets V, V', P', where $V = \{v_0, \ldots, v_{2n-1}\}$, $V' = \{v'_0, \ldots, v'_{2n-1}\}$, and $P' = 2^{V'}$. The relation $\in$ is the "element"-relation between V' and $2^{V'}$, connecting for each X in $2^{V'}$ every node $v' \in X$ with X. The relation $\sim$ is the equivalence relation on $V \cup V'$ whose equivalence classes are $\{v_i, v'_i, v_{n+i}, v'_{n+i}\}$ for all $i < n$. For every graph $G = (V, E)$, the σ-structure $\mathcal{S}_{2n}(G)$ is the expansion of the structure $\mathcal{S}_{2n}$ with the graph's edge relation E. An illustration can be found in Figure 1.

Theorem 4.4 (Otto [17]). *There is an order-invariant* FO($\sigma \cup \{<\}$)*-sentence* $\varphi_{\textit{Otto}}$*, but no* FO(σ)*-sentence, such that for every* $n \in \mathbb{N}_{\geq 1}$ *and every graph* G *on* $2n$ *nodes we have:* $\mathcal{S}_{2n}(G) \models \varphi_{\textit{Otto}} \iff G$ *is connected.*

Proof sketch. Part 1: Construction of the $<$-inv-FO(σ)*-sentence* $\varphi_{\textit{Otto}}$. Otto's proof shows a stronger result, namely that every monadic second-order sentence Φ of signature $\{E\}$ can be translated into an order-invariant FO($\sigma \cup \{<\}$)-sentence φ_{Φ}, such that

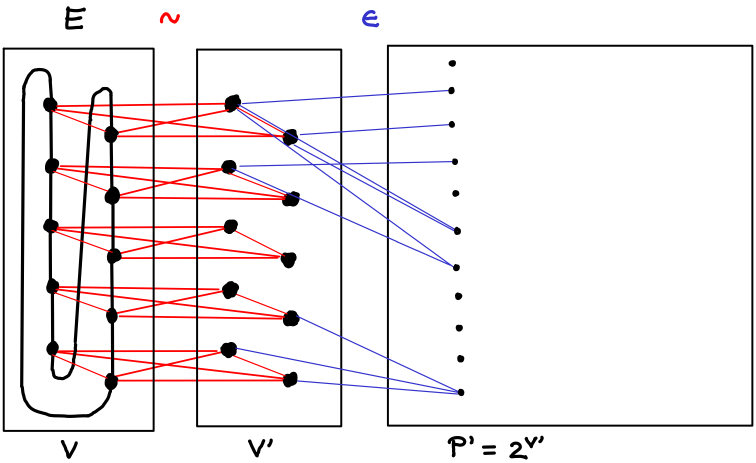

Fig. 1. The σ-structure $S_{2n}(G)$ where $G = G^1_{2n}$ is a cycle on $2n$ nodes $V = \{v_0, v_1, \ldots, v_{2n-1}\}$. G is represented in the leftmost box of the picture. The box in the middle contains the set $V' = \{v' : v \in V\}$. The 4-cliques between V and V' represent the equivalence relation $\sim$. The box on the right contains a node X for each element X in $P' = 2^{V'}$. The edges between the box in the middle and the box on the right represent the $\in$-relation connecting, for each X in $2^{V'}$, every node $v' \in X$ with the node X.

for every $n \in \mathbb{N}_{\geq 1}$ and every graph G on $2n$ nodes we have: $\mathcal{S}_{2n}(G) \models \varphi_\Phi \iff G \models \Phi$.

The claimed formula φ_{Otto} can then be chosen as φ_Φ where Φ is a monadic second-order formalisation of graph connectivity.
For constructing φ_Φ, the following observations are crucial:

(1) We can use the linear order $<$ to define a bijection $\beta_<$ from V to V' such that $v \sim \beta_<(v)$ for every $v \in V$. This bijection can be described by an FO$(V, V', \sim, <)$-formula.
(2) Using this bijection, we can identify V with V'. And using P' and the $\in$-relation between V' and P', we can simulate monadic second-order quantification over V by first-order quantification of elements in P'. Utilising this, it is straightforward to translate Φ into an FO$(\sigma \cup \{<\})$-sentence ψ_Φ.
(3) Finally, the formula φ_Φ is chosen as the conjunction of ψ_Φ with an FO(V, E)-sentence stating that E is a subset of $V \times V$, and an FO(σ')-axiomatisation of σ'-structures isomorphic to $\mathcal{S}_{2n}$ for some $n \in \mathbb{N}_{\geq 1}$.

The resulting formula φ_Φ is satisfied by $\mathcal{A}^\iota$, for a finite σ-structure $\mathcal{A}$ and an embedding ι of $\mathcal{A}$ if, and only if, $\mathcal{A}$ is isomorphic to $\mathcal{S}_{2n}(G)$ for some $n \in \mathbb{N}_{\geq 1}$ and some graph G on $2n$ vertices satisfying Φ. This also shows that φ_Φ is order-invariant on all finite σ-structures.

Part 2: Proof of the non-expressibility in FO(σ). For each $n \in \mathbb{N}_{\geq 1}$ let G^1_{2n} be the cycle $(v_0, v_1, \ldots, v_{2n-1}, v_0)$, and let G^2_{2n} be the disjoint union of the two cycles $(v_0, v_1, \ldots, v_{n-1}, v_0)$ and $(v_n, v_{n+1}, \ldots, v_{2n-1}, v_n)$. An illustration of G^1_{2n} is given in the leftmost box of Figure 1. It suffices to show that for all $r \in \mathbb{N}$ and all sufficiently large n, the structures $\mathcal{S}_{2n}(G^1_{2n})$ and $\mathcal{S}_{2n}(G^2_{2n})$ cannot be distinguished by FO(σ)-sentences of quantifier rank r.

For each $b \in \{1, 2\}$ let $\widehat{G}^b_{2n}$ be the expansion of G^b_{2n} where each node is labeled by its equivalence class with respect to $\sim$ in $\mathcal{S}_{2n}(G^b_{2n})$. An easy Hanf-locality argument (cf., [14]) shows that for every $r \in \mathbb{N}$ and all sufficiently large n, the structures $\widehat{G}^1_{2n}$ and $\widehat{G}^2_{2n}$ cannot be distinguished by first-order sentences of quantifier rank r.

A closer inspection of the structures $\mathcal{S}_{2n}(G^1_{2n})$ and $\mathcal{S}_{2n}(G^2_{2n})$ shows that the duplicator's winning strategy in the r-round Ehrenfeucht-Fraïssé game on $\widehat{G}^1_{2n}$ and $\widehat{G}^2_{2n}$ can be translated into a winning strategy on $\mathcal{S}_{2n}(G^1_{2n})$ and $\mathcal{S}_{2n}(G^2_{2n})$. Thus, the latter two structures cannot be distinguished by FO(σ)-sentences of quantifier rank r. □

5 Limitations of the Expressive Power of Arb-Invariant FO

The results stated in this section hold for arbitrary signatures σ. For simplicity of presentation, however, we let $\sigma = \{E\}$ be the signature consisting of a single binary relation symbol E. Thus, finite σ-structures are finite directed graphs.

Connections between arb-inv-FO(σ) and Circuit Complexity. For proving non-expressibility results for arb-inv-FO(σ), tools from circuit complexity are of major use. We assume that the reader is familiar with basic notions and results in circuit complexity (cf., e.g., the textbook [3]). We consider Boolean circuits consisting of AND- and OR-gates of unbounded fan-in, NOT-gates, input gates, and constant gates **0** and **1**. The size of a circuit is the number of its gates, and the depth is the length of the longest path from an input gate to the output gate.

Let C_m be a circuit with $m \in \mathbb{N}_{\geq 1}$ input gates, and let $w \in \{0, 1\}^m$ be a bitstring of length m. We say that C_m *accepts* w if C_m evaluates to 1 when for every $i \leq m$ the i-th input gate of C_m is assigned the i-th symbol of w.

The non-expressibility proofs for arb-inv-FO(σ) presented in this section rely on Håstad's following well-known circuit lower bound.

Theorem 5.1 (Håstad [10]). *There exist numbers $\ell, m_0 > 0$ such that for every $d \in \mathbb{N}$ with $d \geq 2$ and every $m \in \mathbb{N}$ with $m \geq m_0$ the following is true: No circuit of depth d and size at most $2^{\ell \cdot \sqrt[d-1]{m}}$ accepts exactly those bitstrings $w \in \{0, 1\}^m$ that contain an even number of ones.*

To establish the connection between circuits and arb-inv-FO(σ), we need to represent graphs by bitstrings. This is done in a straightforward way: Consider a directed graph $G = (V, E)$ on $|V| = n$ nodes. Let ι be an embedding of G into $[n]$, and let $(a_{i,j})_{0 \leq i,j < n}$ be the adjacency matrix of G with respect to ι, i.e., $a_{i,j} = 1$ if $(\iota^{-1}(i), \iota^{-1}(j)) \in E$, and $a_{i,j} = 0$ otherwise. The bitstring representation $\text{Rep}^\iota(G)$ of G w.r.t. ι is then chosen as $\text{Rep}^\iota(G) := a_{0,0} \cdots a_{0,n-1} \cdots a_{n-1,0} \cdots a_{n-1,n-1}$. I.e., $\text{Rep}^\iota(G)$ is the concatenation of all rows of the adjacency matrix $(a_{i,j})_{0 \leq i,j < n}$. The connection between FO(σ_{arb}) and Boolean circuits is obtained by the following result.

Theorem 5.2 (Immerman [13]). *For every* FO(σ_{arb})*-sentence* φ *there exist numbers* $d, s \in \mathbb{N}$ *(with* $d \geq 2$*) such that for every* $n \in \mathbb{N}_{\geq 1}$ *there is a circuit* C_{n^2} *with* n^2 *input gates, depth* d*, and size* n^s *such that the following is true for all graphs* $G = (V, E)$ *with* $|V| = n$ *and all embeddings* ι *of* V *into* $[n]$*:* C_{n^2} *accepts* $\mathrm{Rep}^\iota(G) \iff G^\iota \models \varphi$.

Proof sketch. For every fixed n, we translate φ into a Boolean formula with n^2 Boolean variables:

(1) Replace every existential quantification "$\exists x$" of φ into a big disjunction $\bigvee_{0 \leq x < n}$,
(2) replace every universal quantification "$\forall x$" of φ into a big conjunction $\bigwedge_{0 \leq x < n}$.

After these two tranformation steps, the "atomic formulas" remaining in φ are either of the form $E(x, y)$ for $x, y \in [n]$, where E is the edge relation of the graph, or of the form $P(x_1, \ldots, x_r)$ for $x_1, \ldots, x_r \in [n]$, where P is a symbol for a numerical predicate $P^{\mathbb{N}}$ of arity r (here, equality of the form "$x_1 = x_2$" is also viewed as a numerical predicate).

(3) Replace every "atom" of the form $E(x, y)$ for $x, y \in [n]$ with the Boolean variable $a_{x,y}$ representing the edge from $\iota^{-1}(x)$ to $\iota^{-1}(y)$ in G^ι, and
(4) replace every "atom" of the form $P(x_1, \ldots, x_r)$ for $x_1, \ldots, x_r \in [n]$, where P is a symbol for a numerical predicate $P^{\mathbb{N}}$ by the constant **1** if $(x_1, \ldots, x_r) \in P^{\mathbb{N}}$, and by the constant **0** otherwise.

The result of this transformation is a Boolean formula with Boolean variables $a_{x,y}$ for $x, y \in [n]$. This Boolean formula can easily be turned into the desired circuit C_{n^2}. □

As an immediate consequence of the Theorems 5.2 and 5.1 one obtains the following.

Corollary 5.3. *There is no* arb-inv-FO(σ)*-sentence* φ *that is satisfied by exactly those finite directed graphs that consist of an even number of edges.*

Proof. For contradiction, assume that φ is an arb-inv-FO(σ)-sentence satisfied by exactly those finite directed graphs that consist of an even number of edges.

Let d, s and C_{n^2} (for every $n \in \mathbb{N}_{\geq 1}$) be chosen according to Theorem 5.2. It can easily be seen that the circuit C_{n^2} accepts exactly those bitstrings w of length n^2 that contain an even number of ones: Every $w \in \{0, 1\}^{n^2}$ can be viewed as the bitstring representation $\mathrm{Rep}^\iota(G)$ of some graph $G = (V, E)$ on n nodes. Clearly, $G^\iota \models \varphi$ iff $G \models \varphi$ iff G contains an even number of edges. By Theorem 5.2 we furthermore know that $G^\iota \models \varphi$ iff $w = \mathrm{Rep}^\iota(G)$ is accepted by C_{n^2}.

Thus, for $m := n^2$, C_m is a circuit of depth d and size $n^s = m^{s/2}$ that accepts exactly those bitstrings $w \in \{0, 1\}^m$ that contain an even number of ones. However, for any fixed ℓ and all sufficiently large m we have $m^{s/2} < 2^{\ell \cdot \sqrt[d-1]{m}}$, contradicting Theorem 5.1. □

Gaifman Locality of Arb-Invariant FO. A *k-ary query* q is a mapping that associates with every finite directed graph $G = (V, E)$ a k-ary relation $q(G) \subseteq V^k$, which is invariant under isomorphisms, i.e., if π is an isomorphism from a graph G to a graph H, then for all $\overline{a} = (a_1, \ldots, a_k) \in A^k$ we have $\overline{a} \in q(G)$ iff $\pi(\overline{a}) = (\pi(a_1), \ldots, \pi(a_k)) \in q(H)$. Every arb-inv-FO($\sigma$)-formula $\varphi(\overline{x})$ with k free variables defines a k-ary query q_φ via $q_\varphi(G) = \{\overline{a} \in V^k : G \models \varphi[\overline{a}]\}$.

The notion of *Gaifman locality* is a standard tool for showing that particular queries are not definable in certain logics (cf., e.g., the textbook [14] for an overview). For presenting the precise definition of Gaifman locality, we need the following notation.

The *Gaifman graph* of a directed graph $G = (V, E)$ is the undirected graph $\mathcal{G}(G)$ with the same vertex set as G, where for any $a, b \in V$ with $a \neq b$ there is an undirected edge between a and b iff $(a, b) \in E$ or $(b, a) \in E$. The *distance* $dist^G(a, b)$ between two nodes a, b of G is the length of the shortest path between a and b in $\mathcal{G}(G)$.

For every $r \in \mathbb{N}$, the *r-ball* $N_r^G(a)$ around a node a is the set of all nodes b with $dist^G(a, b) \leq r$. The r-ball $N_r^G(\overline{a})$ around a tuple $\overline{a} = (a_1, \ldots, a_k) \in V^k$ is the union of the r-balls around the nodes $a_1, \ldots, a_k$. The *r-neighborhood* of $\overline{a}$ is the induced subgraph $\mathcal{N}_r^G(\overline{a})$ of G on $N_r^G(\overline{a})$.

Definition 5.4 (Gaifman locality). Let $k \in \mathbb{N}_{\geq 1}$ and $f : \mathbb{N} \to \mathbb{N}$. A k-ary query q is *Gaifman $f(n)$-local* if there is an $n_0 \in \mathbb{N}$ such that for every $n \in \mathbb{N}$ with $n \geq n_0$ and every directed graph $G = (V, E)$ on n nodes, the following is true for all k-tuples $\overline{a}, \overline{b} \in V^k$ with $(\mathcal{N}_{f(n)}^G(\overline{a}), \overline{a}) \cong (\mathcal{N}_{f(n)}^G(\overline{b}), \overline{b})$: $\quad \overline{a} \in q(G) \iff \overline{b} \in q(G)$.

I.e., in a graph of size n, a query that is Gaifman $f(n)$-local cannot distinguish between k-tuples of nodes whose neighborhoods of radius $f(n)$ are isomorphic. Gaifman locality is a powerful tool for showing that certain queries cannot be defined by formulas of particular logics.

Example 5.5. Let F be a class of formulas such that every query q definable by a formula in F is Gaifman $f_q(n)$-local for a function $f_q : \mathbb{N} \to \mathbb{N}$ where $f_q(n) \leq n/5$ for all sufficiently large n. Then, none of the following queries is definable in F:

- *reach*$(G) := \{(a, b) : G$ contains a directed path from node a to node $b\}$,
- *cycle*$(G) := \{a : a$ is a node that lies on a cycle of $G\}$,
- *triangle-reach*$(G) := \{a : a$ is reachable from a triangle in $G\}$,
- *same-distance*$(G) := \{(a, b, c) : dist^G(a, c) = dist^G(b, c)\}$.

Assume, for contradiction, that *reach* is definable in F. By assumption, $f_{reach}(n) \leq n/5$ for all sufficiently large n. Now, consider for each n the graph G_n consisting of two disjoint directed paths of length $n/2$, and let a be the first node of the first path, let b be the last node of the first path, and let b' be the last node of the second path. Then, $\mathcal{N}_{n/5}^{G_n}(a, b)$ consists of two disjoint paths of length $n/5$, where a is the first node of the first path and b is the last node of the second path. Obviously, $(\mathcal{N}_{n/5}^{G_n}(a, b), (a, b)) \cong (\mathcal{N}_{n/5}^{G_n}(a, b'), (a, b'))$. Thus, due to the assumed Gaifman $f_{reach}(n)$-locality of the query *reach*, we have $(a, b) \in$ *reach*(G_n) iff $(a, b') \in$ *reach*(G_n). However, in G_n there is a directed path from a to b, but no directed path from a to b' — a contradiction.
Similar constructions can be used to show that none of the queries *cycle*, *triangle-reach*, *same-distance* is definable in F.

It is well-known that FO(σ)-definable queries are Gaifman local with constant locality radius, i.e., for every FO(σ)-definable query q there is a constant c such that q is Gaifman c-local [11]. This can be generalised to order-invariant FO:

Theorem 5.6 (Grohe, Schwentick [8]). *Order-invariant* FO *is Gaifman local with constant locality radius. I.e., for every* $<$-inv-FO(σ)*-definable query* q *there is a constant* c *such that* q *is Gaifman* c*-local.*

The result for *constant* locality radius (independent of the size of the graph) cannot be lifted to arb-invariant FO: In [2] it was shown that for every $d \in \mathbb{N}$ there is an $\{+, \times\}$-inv-FO(σ)-definable unary query q_d that is not Gaifman $(\log n)^d$-local. But still, for arb-invariant FO we get a Gaifman locality result for neighborhoods whose radius is bounded polylogarithmically in the size of the underlying graphs:

Theorem 5.7 (Anderson, Melkebeek, Schweikardt, Segoufin [2]). *Arb-invariant* FO *is Gaifman local with polylogarithmic locality radius. I.e., for every query* q *definable in* arb-inv-FO(σ) *there is a constant* c *such that* q *is Gaifman* $(\log n)^c$*-local.*

Note that this suffices to conclude that none of the queries mentioned in Example 5.5 is definable in arb-inv-FO(σ).

The proof of Theorem 5.6 relies on a sophisticated construction using Ehrenfeucht-Fraïssé games. A simplified proof a weaker version of Theorem 5.6 can be found in the textbook [14]. The proof of Theorem 5.7 exploits the connection between arb-invariant FO and Boolean circuits. In the following, we present the proof of a weaker version of Theorem 5.7 for the particular case of *unary* queries and the notion of *weak Gaifman locality* [14], where "$\overline{a} \in q(G) \iff \overline{b} \in q(G)$" needs to be true only for those tuples $\overline{a}$ and $\overline{b}$ whose $f(n)$-neighborhoods are disjoint.

Definition 5.8 (Weak Gaifman locality). Let $f : \mathbb{N} \to \mathbb{N}$. A unary query q is called *weakly Gaifman* $f(n)$*-local* if there is an $n_0 \in \mathbb{N}$ such that for every $n \in \mathbb{N}$ with $n \geq n_0$ and every directed graph $G = (V, E)$ on n nodes, the following is true for all nodes $a, b \in V$ with $(\mathcal{N}^G_{f(n)}(a), a) \cong (\mathcal{N}^G_{f(n)}(b), b)$ and $N^G_{f(n)}(a) \cap N^G_{f(n)}(b) = \emptyset$: $a \in q(G) \iff b \in q(G)$.

We give a proof of the following weaker version of Theorem 5.7:

Proposition 5.9. *For every unary query* q *definable in* arb-inv-FO(σ) *there is a constant* c *such that* q *is weakly Gaifman* $(\log n)^c$*-local.*

Proof. Let q be a unary query expressed by an arb-inv-FO(σ)-formula $\varphi(x)$. By using a variation of Theorem 5.2, there exist numbers $d, s \in \mathbb{N}$ such that for every $n \in \mathbb{N}_{\geq 1}$ there is a circuit C_{n^2+n} with n^2+n input gates, depth d, and size n^s such that the following is true for all graphs $G = (V, E)$ with $|V| = n$, for all nodes $a \in V$, and all embeddings ι of V into $[n]$:

$$C_{n^2+n} \text{ accepts } \mathrm{Rep}^\iota(G, a) \iff G \models \varphi[a]. \tag{1}$$

Here, $\mathrm{Rep}^\iota(G, a) = \mathrm{Rep}^\iota(G)\mathrm{Rep}^\iota(a)$ is the bitstring representation of (G, a), where $\mathrm{Rep}^\iota(a)$ is the bitstring $b_0 \cdots b_{n-1}$ with $b_{\iota(a)} = 1$ and $b_j = 0$ for all $j \neq \iota(a)$.

For contradiction, let us now assume that for every $c \in \mathbb{N}$ the query q defined by $\varphi(x)$ is *not* weakly Gaifman $(\log n)^c$-local. Thus, in particular for $c := 2(d-1)$ we obtain that for all $n_0 \in \mathbb{N}$ there exists an $n \geq n_0$, and

$(*)$: a graph $G = (V, E)$ on n nodes, and nodes $a, b \in V$ such that
for $m := (\log n)^c = (\log n)^{2(d-1)}$ we have:

$$(\mathcal{N}_m^G(a), a) \cong (\mathcal{N}_m^G(b), b), \quad N_m^G(a) \cap N_m^G(b) = \emptyset, \quad G \models \varphi[a], \quad G \not\models \varphi[b].$$

Claim. *The circuit C_{n^2+n} can be transformed into a circuit $\tilde{C}_m$ on m input bits, such that $\tilde{C}_m$ has the same depth and size as C_{n^2+n} and accepts exactly those bitstrings $w \in \{0,1\}^m$ that contain an even number of ones.*

Before proving this claim, let us point out how it can be used to conclude the proof of Proposition 5.9. According to the claim, $\tilde{C}_m$ is a circuit of depth d and size n^s, which accepts exactly those bitstrings $w \in \{0,1\}^m$ that contain an even number of ones.

From Theorem 5.1 we know that the size n^s of $\tilde{C}_m$ must be bigger than $2^{\ell \cdot \sqrt[d-1]{m}}$. However, we had chosen $m = (\log n)^{2(d-1)}$, and hence $2^{\ell \cdot \sqrt[d-1]{m}} = 2^{\ell \cdot (\log n)^2} = n^{\ell \cdot \log n} > n^s$ for all sufficiently large n — a contradiction! Thus, for concluding the proof of Proposition 5.9, it suffices to prove the claim.

Proof of the claim. Let $G = (V, E)$ be the graph chosen according to $(*)$, and let ι be an arbitrary embedding of V into $[n]$. The idea is to define, for every bitstring $w \in \{0,1\}^m$, a graph G_w such that

$$(**): \qquad (G_w, a) \cong \begin{cases} (G, a) & \text{if } w \text{ contains an even number of ones,} \\ (G, b) & \text{otherwise.} \end{cases}$$

The circuit $\tilde{C}_m$ is constructed in such a way that on input $w \in \{0,1\}^m$ it does the same as circuit C_{n^2+n} does on input $\mathrm{Rep}^\iota(G_w, a)$. From (1) we then obtain that

$$\tilde{C}_m \text{ accepts } w \iff C_{n^2+n} \text{ accepts } \mathrm{Rep}^\iota(G_w, a) \iff G_w \models \varphi[a].$$

If w contains an even number of ones, $(G_w, a) \cong (G, a)$. As we know from $(*)$ that $G \models \varphi[a]$, we therefore obtain that $\tilde{C}_m$ accepts w.
If w contains an odd number of ones, $(G_w, a) \cong (G, b)$. As we know from $(*)$ that $G \not\models \varphi[b]$, we therefore obtain that $\tilde{C}_m$ does not accept w.
Thus, circuit $\tilde{C}_m$ accepts exactly those $w \in \{0,1\}^m$ that contain an even number of 1s.

Definition of G_w: From $(*)$ we know that there exists an isomorphism π from $\mathcal{N}_m^G(a)$ to $\mathcal{N}_m^G(b)$ with $\pi(a) = b$. Furthermore, we know that $N_m^G(a) \cap N_m^G(b) = \emptyset$, and thus G contains no edges that link vertices of $N_{m-1}^G(a)$ with vertices of $N_{m-1}^G(b)$. For $x \in \{a, b\}$, we partition $N_m^G(x)$ into *shells* $S_i(x) := \{y \in V : \mathit{dist}^G(x, y) = i\}$, for all $i \leq m$. Note that $\pi(S_i(a)) = S_i(b)$.
In the following, we write S_i for the set $S_i(a) \cup S_i(b)$.
For a bitstring $w = w_1 \cdots w_m \in \{0,1\}^m$ the graph G_w is defined as follows:

- G_w has the same vertex set V as the graph G.
- All edges of G that do *not* link a vertex of shell S_{i-1} with a vertex of shell S_i, for some $i \leq m$, are copied into G_w.
- Edges of G that link a vertex of shell S_{i-1} with a vertex of shell S_i, for some $i \leq m$ are modified depending on the i-th bit w_i of the bitstring w:

– If $w_i = 0$, then edges of G that link a vertex of shell S_{i-1} with a vertex of shell S_i are copied into G_w.
– If $w_i = 1$, then for every edge of G that links a vertex u of shell $S_{i-1}(a)$ with a vertex v of shell $S_i(a)$, we insert into G_w an edge that links vertex u (in shell $S_{i-1}(a)$) with vertex $\pi(v)$ (in shell $S_i(b)$), and we also insert the according edge that links vertex $\pi(u)$ (in shell $S_{i-1}(b)$) with the vertex v (in shell $S_i(a)$).

Thus, for every i with $w_i = 1$, the roles of the shells $S_i(a)$ and $S_i(b)$ are swapped. It is straightforward to see that the resulting graph G_w satisfies $(**)$; see Figure 2 for an illustration.

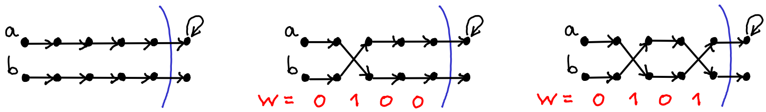

Fig. 2. Illustration of the graph G_w for neighborhoods of radius $m = 4$

Construction of $\tilde{C}_m$*:* Let us fix an embedding ι of V into $[n]$. The circuit $\tilde{C}_m$ is obtained from C_{n^2+n} by replacing the input gates of C_{n^2+n} as follows:

Let $u, v \in V$, and let $g_{\mu,\nu}$ for $\mu := \iota(u)$ and $\nu := \iota(v)$ be the input gate of C_{n^2+n} that corresponds to the entry $a_{\mu,\nu}$ of G's adjacency matrix w.r.t. ι (i.e., $a_{\mu,\nu} = 1$ iff $(u, v) \in E$).

In case that (u, v) does *not* belong to $(S_{i-1} \times S_i) \cup (S_i \times S_{i-1})$ for any $i \leq m$, the gate $g_{\mu,\nu}$ is replaced by the constant gate **1** if $(u, v) \in E$, and by the constant gate **0** if $(u, v) \notin E$.

In case that (u, v) belongs to $(S_{i-1}(a) \times S_i(a)) \cup (S_i(a) \times S_{i-1}(a))$ for some $i \leq m$, let $g := g_{\mu,\nu}$, and let $g', \tilde{g}, \tilde{g}'$ be the input gates of C_{n^2+n} corresponding to the potential edges $(\pi(u), \pi(v))$, $(u, \pi(v))$, and $(\pi(u), v)$, respectively.
If $(u, v) \notin E$, then $g, g', \tilde{g}, \tilde{g}'$ are replaced by the constant gate **0**.
If $(u, v) \in E$, then g and g' are replaced by $\neg w_i$, whereas $\tilde{g}$ and $\tilde{g}'$ are replaced by w_i, where w_i is the input gate for the i-th bit of the input bitstring of length m.

It is straightforward to see that on input $w \in \{0, 1\}^m$ the circuit $\tilde{C}_m$ does the same as circuit C_{n^2+n} does on input $\text{Rep}^{\iota}(G_w, a)$. This completes the proof of Proposition 5.9. □

6 Some Open Questions

We conclude with a list of open research questions:

(1) Is addition-invariant FO Gaifman local with *constant* locality radius? (Cf., [8,14].)
(2) Can addition-invariant FO define string-languages that are not regular? (See [24] for details.)
(3) Are there analogues of the Theorems 5.6 and 5.7 for the notion of *Hanf locality*? (Cf. [14] for the definition of Hanf locality.)
(4) Does order-invariant FO have a zero-one law? (See [6,14] for zero-one laws.)
(5) Are there *decidable* characterisations of order-invariant FO, addition-invariant FO, or $\{+, \times\}$-invariant FO? (See [4,24,9] for related results.)

References

1. Abiteboul, S., Hull, R., Vianu, V.: Foundations of Databases. Addison-Wesley (1995)
2. Anderson, M., van Melkebeek, D., Schweikardt, N., Segoufin, L.: Locality from circuit lower bounds. SIAM Journal on Computing 41(6), 1481–1523 (2012)
3. Arora, S., Barak, B.: Computational Complexity: A Modern Approach. Cambridge Univ. Press (2009)
4. Benedikt, M., Segoufin, L.: Towards a characterization of order-invariant queries over tame structures. Journal of Symbolic Logic 74(1), 168–186 (2009)
5. Courcelle, B.: The monadic second-order logic of graphs X: Linear orderings. Theoretical Computer Science 160(1&2), 87–143 (1996)
6. Ebbinghaus, H.-D., Flum, J.: Finite model theory. Springer (1999)
7. Ganzow, T., Rubin, S.: Order-invariant MSO is stronger than Counting MSO in the finite. In: Proc. STACS 2008, pp. 313–324 (2008)
8. Grohe, M., Schwentick, T.: Locality of order-invariant first-order formulas. ACM Transactions on Computational Logic 1(1), 112–130 (2000)
9. Harwath, F., Schweikardt, N.: Regular tree languages, cardinality predicates, and addition-invariant FO. In: Proc. STACS 2012, pp. 489–500 (2012)
10. Håstad, J.: Computational limitations for small-depth circuits. PhD thesis. MIT (1986)
11. Hella, L., Libkin, L., Nurmonen, J.: Notions of locality and their logical characterizations over finite models. Journal of Symbolic Logic 64(4), 1751–1773 (1999)
12. Immerman, N.: Relational queries computable in polynomial time. Information and Control 68(1-3), 86–104 (1986)
13. Immerman, N.: Languages that capture complexity classes. SIAM Journal on Computing 16(4), 760–778 (1987)
14. Libkin, L.: Elements of Finite Model Theory. Springer (2004)
15. Makowsky, J.A.: Invariant Definability and **P**/*poly*. In: Gottlob, G., Grandjean, E., Seyr, K. (eds.) CSL 1998. LNCS, vol. 1584, pp. 142–158. Springer, Heidelberg (1999)
16. Niemistö, H.: On locality and uniform reduction. In: Proc. LICS 2005, pp. 41–50 (2005)
17. Otto, M.: Epsilon-logic is more expressive than first-order logic over finite structures. Journal of Symbolic Logic 65(4), 1749–1757 (2000)
18. Potthoff, A.: Logische Klassifizierung regulärer Baumsprachen. PhD thesis, Christian-Albrechts-Universität Kiel (1994)
19. Rossman, B.: Successor-invariant first-order logic on finite structures. Journal of Symbolic Logic 72(2), 601–618 (2007)
20. Rossman, B.: On the constant-depth complexity of k-clique. In: Proc. STOC 2008, pp. 721–730 (2008)
21. Schweikardt, N.: Arithmetic, first-order logic, and counting quantifiers. ACM Transactions on Computational Logic 6(3), 634–671 (2005)
22. Schweikardt, N.: On the expressive power of monadic least fixed point logic. Theoretical Computer Science 350(2-3), 325–344 (2006)
23. Schweikardt, N.: On the expressive power of logics with invariant uses of arithmetic predicates. In: Ong, L., de Queiroz, R. (eds.) WoLLIC 2012. LNCS, vol. 7456, pp. 85–87. Springer, Heidelberg (2012)
24. Schweikardt, N., Segoufin, L.: Addition-invariant FO and regularity. In: Proc. LICS 2010, pp. 273–282 (2010)
25. Schwentick, T.: On Winning Ehrenfeucht Games and Monadic NP. Annals of Pure and Applied Logic 79(1), 61–92 (1996)
26. Vardi, M.: The complexity of relational query languages. In: Proc. STOC 1982, pp. 137–146 (1982)