Lower Bounds for Invariant Queries in Logics with Counting

Leonid Libkin [a,*,1] Limsoon Wong [b]

[a] *Department of Computer Science, University of Toronto, Toronto, Ontario M5S 3H5, Canada*

[b] *Kent Ridge Digital Labs, 21 Heng Mui Keng Terrace, Singapore 119613*

Abstract

We study the expressive power of counting logics in the presence of auxiliary relations such as orders and preorders. The simplest such logic is the first-order logic with counting. This logic captures the complexity class TC^0 over ordered structures. We also consider first-order logic with arbitrary unary quantifiers and with infinitary extensions.

We start by giving a simple direct proof that first-order logic with counting, in the presence of pre-orders that are almost-everywhere linear orders, cannot express the transitive closure of a binary relation. The proof is based on locality of formulae. We then show that the technique cannot be extended to linear orders. We further show that this result does not say anything about the power of invariant queries in first-order logic with counting vs. the class TC^0, in the presence of these preorders.

In the second part of the paper, we prove a separation result showing that, for all the counting logics above, a linear order is more powerful than a preorder that is a linear order almost everywhere. In fact, we prove that the expressive power of invariant queries in the presence of such preorders can be characterized by a property normally associated with first-order definability over *unordered* structures. We do this by using locality techniques from finite-model theory. However, as some standard notions of locality fail in this setting, we have to modify them to prove the main result.

* Corresponding author.

Email addresses: libkin@cs.toronto.edu (Leonid Libkin), limsoon@krdl.org.sg (Limsoon Wong).

[1] Research affiliation: Bell Labs.

Preprint submitted to Theoretical Computer Science

1 Introduction

The development of Descriptive Complexity suggests a very close connection between proving lower bounds in complexity theory and proving inexpressibility results in logic. The latter are of the form "a property P cannot be expressed in a logic $\mathcal{L}$ over a class of finite models." Developing tools for proving such expressivity bounds is one of the central problems in Finite-Model Theory. In this paper we show how tools based on locality of logics can be applied to the complexity class TC^0 and how they allow us to derive new expressivity bounds of counting extensions of first-order logic in the presence of complex auxiliary relations.

The class TC^0 is an important complexity class. Problems such as integer multiplication and division, and sorting belong to TC^0. This class has also been studied in connection with neural nets [29]. Despite serious efforts and a number of proved lower bounds [1], it is still not known if $\mathrm{TC}^0 \subsetneq \mathrm{NP}$. In fact, the results of [30] show that traditional approaches to circuit lower bounds are unlikely to succeed in proving this separation.

A starting point for our study is a result by Barrington, Immerman and Straubing [2] stating that:

$$\mathrm{FO}(\mathbf{C}) + < \quad = \quad \text{uniform } \mathrm{TC}^0$$

Here, TC^0 is the class of problems solvable by polynomial-size, constant-depth threshold circuits; and uniform means DLOGTIME-uniform; see [2] for more details. From now on, we write TC^0 whenever we mean the uniform class. $\mathrm{FO}(\mathbf{C})$ is the extension of first-order logic with counting quantifiers $\exists i$, where $\exists ix.\varphi(x)$ means that there are at least i elements x that satisfy φ. $\mathrm{FO}(\mathbf{C})+ <$ is $\mathrm{FO}(\mathbf{C})$ in the presence of a built-in order relation. We give full definitions later. At this point, we offer an example: $\exists i, j((j + j = i) \wedge \exists! ix.\varphi(x))$, where $\exists! i$ is a shorthand for "exists exactly i". This formula states that the number of x satisfying φ is even; this property is known to be inexpressible in first-order logic alone.

The problem of separation of uniform TC^0 from classes such as DLOGSPACE, NLOGSPACE, P, etc., is thus reduced to proving that their complete problems are inexpressible in $\mathrm{FO}(\mathbf{C})+ <$. However, it appears that the presence of an order relation is a major obstacle to proving such expressivity bounds for $\mathrm{FO}(\mathbf{C})$. Several partial results [8,21] show that there are problems complete for DLOGSPACE that cannot be defined by $\mathrm{FO}(\mathbf{C})$ in the presence of auxiliary relations whose degrees are bounded by a fixed constant k. If we talk about directed graphs, by *degrees* we mean in- and out-degrees of nodes. For example, in the graph of a successor relation, every node has in-and out-degree either 0 or 1. In contrast, in a linear order on an n-element set, all n different (in- and out-) degrees from 0 to $n-1$ are realized. Thus, in order to move closer to proving expressivity bounds in the presence of an order relation, one has to at least be able to lift the results from constant degrees to those that depend on the size of the input.

A result in this direction was proved in [21] using a definition of *moderate degree* by Fagin, Stockmeyer and Vardi [9]. We say that a class $\mathcal{C}$ of graphs (more generally, relational structures) is of moderate degree, if $degmax_{\mathcal{C}}(n)$, the maximal in- or out-degree of an n-element graph from $\mathcal{C}$, is at most $\log^{o(1)} n$. That is, for some function $\delta(n)$ such that $\lim_{n\to\infty}\delta(n) = 0$, we have $degmax_{\mathcal{C}}(n) \leq \log^{\delta(n)} n$. Then [21] proved that there is a DLOGSPACE-complete problem which is not definable in FO($\mathbf{C}$) in the presence of auxiliary relations of moderate degree.

In [9], auxiliary relations of moderate degree were shown to be of no help for expressing connectivity of graphs in monadic Σ^1_1. Starting from their result, Schwentick extended it to degrees $n^{o(1)}$ [31] and to a linear order [32]. So one may wonder if a similar program can be carried out for FO($\mathbf{C}$).

The intuition behind the introduction of a linear order is that it allows us to simulate encodings of structures on the tape of a Turing machine (or the order of inputs of a circuit). While for order-invariant properties it does not matter in which order elements appear on the tape (indeed, properties like connectivity of graphs do not depend on how graphs are represented), they do appear in *some* order, and one must be able to use this order in logical formulae. Even though the particular ordering does not change the truth value of an order-invariant formula, the mere presence of an order gives many logics extra power. For example, while FO+LFP and FO+PFP capture PTIME and PSPACE over ordered structures [15,33], they possess the 0-1 law over unordered structures [19], meaning that such a simple PTIME property as parity cannot be expressed. The lower bound of Cai, Fürer and Immerman [4] shows that there are PTIME properties of unordered structures not definable even in FO+LFP extended with counting quantifiers. A similar phenomenon is observed for other logics, e.g., FO and FO($\mathbf{C}$) [3,28].

Our main goal is to study the impact of auxiliary relations, such as orderings, on the expressive power of logics with counting. Our results apply to a variety of logics, starting with FO and FO($\mathbf{C}$), and ending with a logic $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$ proposed in [22]. This logic subsumes FO($\mathbf{C}$) and all other known pure counting extensions of FO. Note that when we speak of counting extensions of FO, we mean extensions that only add a counting mechanism, as opposed to those extensively studied in the literature [27] that add both counting and fixpoint.

We consider a class of relations which are extremely close to linear orderings.

These are preorders, with equivalence classes of size at most 2, that coincide with linear orders almost everywhere. See Section 2 for precise definition, and Figure 1 for a picture. We first prove, by a simple direct argument, that there are DLOGSPACE and NLOGSPACE-complete problems not definable in all the counting logics above, in the presence of such relations. This immediately leads to a question whether the expressivity of, say, FO($\mathbf{C}$) in the presence of such relations is the same as that of FO($\mathbf{C}$)+ $<$. In the second part of the paper, we prove a more involved result showing that this is *not* the case. In particular, logics such as FO($\mathbf{C}$) and $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$, in the presence of preorders that are almost everywhere linear orders, exhibit very tame behavior, normally associated with first-order definable properties. To prove the main result, we exploit the locality techniques in Finite-Model Theory.

The idea of locality in Finite-Model Theory was first introduced by Gaifman [10]. Informally, a logic is local if the result of any query or property definable in it can be determined by examining a "small neighborhoods" of its arguments. An interesting consequence of locality is the "bounded number of degrees property." Informally, a logic has the bounded number of degrees property if any graph definable in it in terms of a second graph has a small number of distinct in- and out-degrees that depends only on the defining formula and the maximum in- and out-degree of the second graph. These properties make it straightforward to infer many inexpressibility results. For example, if a logic is known to be local and thus has the bounded number of degrees property, then we can immediately conclude that it cannot define the transitive closure of a chain graph, as the number of distinct degrees in the transitive closure obviously depends on the length of the chain.

Organization In Section 2, we give formal definitions of various counting extensions of FO, notions of locality, and definability with auxiliary relations. We also give an example that shows how the presence of auxiliary relations affects expressiveness.

In Section 3, we give a direct proof that the transitive closure query is not expressible in FO($\mathbf{C}$) in the presence of almost-everywhere linear orders. We also explain that the technique of the proof does not straightforwardly generalize to proving separation results in the ordered case.

In Section 4, we address the question of whether it is possible to use the almost-everywhere linear orders to prove separation results in the ordered case. We give a negative answer. Indeed, for all counting logics we consider here, adding a linear order is strictly more expressive than adding a preorder, however close to a linear order that preorder might be. We state the result and some of its corollaries. In Section 5, we give the proof, where we first describe notions of weak locality and then combine them with bijective Ehrenfeucht-Fraïssé games.

Two extended abstracts with the results of this paper appeared in the Proceedings of 15th Symposium on Theoretical Aspects of Computer Science [25], and the Proceedings of the 14th IEEE Symposium on Logic in Computer Science [23].

2 Notations

Finite Structures and Logics All structures are assumed to be *finite*. A relational signature σ is a set of relation symbols $\{R_1, ..., R_l\}$, with associated arities $p_i > 0$. For directed graphs, the signature consists of one binary predicate. A σ-structure is $\mathcal{A} = \langle A, R_1^{\mathcal{A}}, \ldots, R_l^{\mathcal{A}} \rangle$, where A is a finite set, and $R_i^{\mathcal{A}} \subseteq A^{p_i}$ interprets R_i. We also allow some constants into a σ-structure where needed. The class of finite σ-structures is denoted by STRUCT[σ]. When there is no confusion, we write R_i in place of $R_i^{\mathcal{A}}$. Isomorphism is denoted by $\cong$. The carrier of a structure $\mathcal{A}$ is always denoted by A.

We abbreviate first-order logic by FO and omit the standard definitions. *FO with counting*, denoted by FO($\mathbf{C}$), is a two-sorted logic with second sort being interpreted as an initial segment of natural numbers. Here, a structure $\mathcal{A}$ is of the form

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

The relations $R_i^{\mathcal{A}}$ are defined on the domain $\{v_1, \ldots, v_n\}$. The constants $\underline{1}$ and $\underline{\max}$ are defined on the numerical domain $\{1, \ldots, n\}$ and are interpreted as 1 and n respectively. On the numerical domain the logic also has a linear order $<$ and the BIT predicate available, where BIT(i, j) iff the ith bit in the binary representation of j is one. This logic also has *counting quantifiers* $\exists ix.\varphi(x)$, meaning that there are at least i elements x that satisfy $\varphi(x)$; here i refers to the numerical domain and x to the domain $\{v_1, \ldots, v_n\}$. These quantifiers bind x but not i. Ternary predicates $+$ and $*$ are definable on the numerical domain [8]. The quantifier $\exists! ix$ meaning the existence of exactly i elements satisfying a formula is also definable. For example, $\exists i \exists j \ [(j+j) = i \wedge \exists! ix.\varphi(x)]$ tests if the number of x satisfying φ is even. Note that this example property is not definable in FO alone. We separate first-sort variables from second-sort variables by semicolon: $\varphi(\vec{x}; \vec{\jmath})$.

There are several counting extensions of FO that are more powerful than FO($\mathbf{C}$). Amongst them is FO($\mathbf{Q}_u$). This logic is FO extended with *all* unary quantifiers. We refer the reader to [12,18] for the definition of FO($\mathbf{Q}_u$) and its properties. Here, we mostly work with an even more powerful logic that is defined below.

We denote the infinitary logic by $\mathcal{L}_{\infty\omega}$. It extends FO by allowing infinite conjunctions $\bigwedge$ and disjunctions $\bigvee$. Then $\mathcal{L}_{\infty\omega}(\mathbf{C})$ is a two-sorted logic that extends the infinitary logic $\mathcal{L}_{\infty\omega}$. Its structures are of the form $(\mathcal{A}, \mathbb{N})$, where $\mathcal{A}$ is a finite relational structure and $\mathbb{N}$ is a copy of natural numbers. Assume that every constant $n \in \mathbb{N}$ is a second-sort term. To $\mathcal{L}_{\infty\omega}$, add counting quantifiers $\exists ix$ for every $i \in \mathbb{N}$, and *counting terms:* If φ is a formula and $\vec{x}$ is a tuple of free first-sort variables in φ, then $\#\vec{x}.\varphi$ is a term of the second sort, and its free variables are those in φ except $\vec{x}$. Its interpretation is the number of tuples $\vec{a}$ over the finite first-sort universe that satisfy φ. That is, given a structure $\mathcal{A}$, a formula $\varphi(\vec{x}, \vec{y}; \vec{\jmath})$, $\vec{b} \subseteq A$, and $\vec{\jmath}_0 \subset \mathbb{N}$, the value of the term $\#\vec{x}.\varphi(\vec{x}, \vec{b}; \vec{\jmath}_0)$ is the cardinality of the finite set $\{\vec{a} \subseteq A \mid \mathcal{A} \models \varphi(\vec{a}, \vec{b}; \vec{\jmath}_0)\}$. For example, the interpretation of $\#x.E(x, y)$ is the in-degree of node y in a graph with the edge-relation E.

As this logic expresses every property of finite structures, it is too powerful. We restrict it by means of the *rank* of formulae and terms, denoted by rk. It is defined as quantifier rank. That is, it is 0 for atomic formulae; $\mathsf{rk}(\bigvee_i \varphi_i) = \max_i \mathsf{rk}(\varphi_i)$; $\mathsf{rk}(\neg\varphi) = \mathsf{rk}(\varphi)$; and $\mathsf{rk}(\exists x\varphi) = \mathsf{rk}(\exists ix\varphi) = \mathsf{rk}(\varphi) + 1$ as usual. But it does not take into account quantification over $\mathbb{N}$: $\mathsf{rk}(\exists i\varphi) = \mathsf{rk}(\varphi)$. Furthermore, $\mathsf{rk}(\#\vec{x}.\psi) = \mathsf{rk}(\psi) + |\vec{x}|$.

Definition 1 (see [22]) *The logic $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$ is defined to be the restriction of $\mathcal{L}_{\infty\omega}(\mathbf{C})$ to terms and formulae of finite rank.*

It is known [22] that $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$ formulae are closed under Boolean connectives and all quantification, and that every predicate on $\mathbb{N} \times \ldots \times \mathbb{N}$ is definable by a $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$ formula of rank 0. Thus, we assume that $+, *, -, \leq$, and in fact *every* predicate on natural numbers is available. Known counting expansions of FO are contained in $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$. That is, for every FO, FO($\mathbf{C}$), or FO($Q_{\mathbf{u}}$) formula, there exists an equivalent $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$ formula of the same rank. A counting logic of [3] can also be embedded into $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$.

Definability with auxiliary relations An *m-ary query* on σ-structures, Q, is a mapping that associates to each $\mathcal{A} \in \mathrm{STRUCT}[\sigma]$ a structure $\langle A, S\rangle$, where $S \subseteq A^m$. We write $\vec{a} \in Q(\mathcal{A})$ if $\vec{a} \in S$, where $\langle A, S\rangle = Q(\mathcal{A})$. A query Q is definable in a logic $\mathcal{L}$ if there exists an $\mathcal{L}$ formula $\varphi(x_1, \ldots, x_m)$ such that $Q(\mathcal{A}) = \varphi[\mathcal{A}] \stackrel{\text{def}}{=} \langle A, \{\vec{a} \mid \mathcal{A} \models \varphi(\vec{a})\}\rangle$.

Let σ' be a relational signature disjoint from σ. If $\mathcal{A}$ is a σ-structure on a universe A, and $\mathcal{A}'$ is a σ'-structure on A, we use the notation $(\mathcal{A}, \mathcal{A}')$ for the $\sigma \cup \sigma'$-structure on A which inherits the interpretation of σ relational symbols from $\mathcal{A}$, and the interpretation of σ' symbols from $\mathcal{A}'$.

Let $\mathcal{C}$ be a class of σ'-structures, with σ and σ' being disjoint. Let $\mathcal{A} \in \text{STRUCT}[\sigma]$. A formula $\varphi(\vec{x})$ in the language of $\sigma \cup \sigma'$ is called *$\mathcal{C}$-invariant on $\mathcal{A}$* if for any two $\mathcal{C}$ structures $\mathcal{A}'$ and $\mathcal{A}''$ on A we have $\varphi[(\mathcal{A}, \mathcal{A}')] = \varphi[(\mathcal{A}, \mathcal{A}'')]$. Associated with such a formula is the following m-ary query (where $m = | \vec{x} |$):

$$Q^w_\varphi(\mathcal{A}) = \begin{cases} \varphi[(\mathcal{A}, \mathcal{A}')], & \varphi \text{ is } \mathcal{C}\text{-invariant on } \mathcal{A} \\ \emptyset, & \text{otherwise.} \end{cases}$$

where $\mathcal{A}'$ is any structure from $\mathcal{C}$ on A. We use the notation $(\mathcal{L}+\mathcal{C})_w$ to denote all queries defined in such a way when φ ranges over formulae of $\mathcal{L}$.

A formula φ is *$\mathcal{C}$-invariant* if it is $\mathcal{C}$-invariant on every structure. With such a φ, we associate a query Q_φ given by $Q_\varphi(\mathcal{A}) = \varphi[(\mathcal{A}, \mathcal{A}')]$ where $\mathcal{A}'$ is a structure from $\mathcal{C}$ on A. The class of all such queries is denoted by $\mathcal{L} + \mathcal{C}$. Clearly,

$$\mathcal{L} + \mathcal{C} \quad \subseteq \quad (\mathcal{L} + \mathcal{C})_w$$

We thus shall aim to establish expressivity bounds for $(\mathcal{L} + \mathcal{C})_w$.

When $\mathcal{C}$ is the class of order relations, we shall write $<$ instead of $\mathcal{C}$. The capture results for complexity classes deal with the classes of queries of the form $\mathcal{L}+ <$. For example, uniform TC^0 equals $\text{FO}(\mathbf{C})+ <$ [2]. While queries in $\mathcal{L}+ <$ are independent of any particular order relation used, the mere presence of such a relation can have an impact on the expressivity of a logic.

We give an example for $\text{FO}(\mathbf{C})$. Assume that σ has one binary and one unary relation, i.e. its structures are graphs with a selected subset of nodes. Let Q_0 be the following Boolean query [3]: given a structure $\langle A, E, X \rangle$, where $A \neq \emptyset$, $E \subseteq A^2$ and $X \subseteq A$, return true iff E is an equivalence relation, and the number of distinct sizes of E-classes equals $|X|$. It is known that Q_0 is not expressible in $\text{FO}(\mathbf{C})$ [3]. However, it *is* expressible in $\text{FO}(\mathbf{C})+ <$. Indeed, the equivalence relation $x\theta y$ iff the E-equivalence classes of x and y have the same cardinality is definable in $\text{FO}(\mathbf{C})$. Thus, in $\text{FO}(\mathbf{C})$ one defines the set of smallest (wrt $<$) elements of each such class, and then compares, in $\text{FO}(\mathbf{C})$, the size of this set to X. The two are the same iff the value of Q_0 is true. Note that *any* linear order suffices to express this query.

Thus, $\text{FO}(\mathbf{C}) \subsetneqq \text{FO}(\mathbf{C})+ <$. Since the latter captures uniform TC^0, this means that there are problems in TC^0 not definable in $\text{FO}(\mathbf{C})$ over unordered structures. It is also known that $\text{FO} \subsetneqq \text{FO}+ <$. We shall see later that this continues to be true for other counting logics.

Bounded number of degrees property (BNDP) If $\mathcal{A} \in \mathrm{STRUCT}[\sigma]$ and R_i is of arity p_i, then $degree_j(R_i^{\mathcal{A}}, a)$, for $1 \leq j \leq p_i$, is the number of tuples $\vec{a}$ in $R_i^{\mathcal{A}}$ having a in the jth position. In the case of directed graphs, this gives us the usual notions of in- and out-degree. By $deg\_set(\mathcal{A})$ we mean the set of all $degree_j(R_i^{\mathcal{A}}, a)$ realized in $\mathcal{A}$, and $deg\_count(\mathcal{A})$ stands for the cardinality of $deg\_set(\mathcal{A})$. We use the notation $\mathrm{STRUCT}_k[\sigma]$ for $\{\mathcal{A} \in \mathrm{STRUCT}[\sigma] \mid deg\_set(\mathcal{A}) \subseteq \{0, 1, \ldots, k\}\}$.

Definition 2 (see [24,5,21]) *An m-ary query Q, $m \geq 1$, is said to have the* bounded number of degrees property [2] *, or* BNDP*, if there exists a function $f_Q : \mathbb{N} \to \mathbb{N}$ such that $deg\_count(Q(\mathcal{A})) \leq f_Q(k)$ for every $\mathcal{A} \in \mathrm{STRUCT}_k[\sigma]$. QED*

The BNDP is very easy to use for proving expressivity bounds [24]. For example, it is very easy to verify that (deterministic) transitive closure violates the BNDP.

Locality All existing proofs of the BNDP establish first that a logic is *local*. We now define this concept. Given a structure $\mathcal{A}$, its *Gaifman graph* [7,10,9] $\mathcal{G}(\mathcal{A})$ is defined as $\langle A, E\rangle$ where (a, b) is in E iff there is a tuple $\vec{c} \in R_i^{\mathcal{A}}$ for some i such that both a and b are in $\vec{c}$. The distance $d(a, b)$ is defined as the length of the shortest path from a to b in $\mathcal{G}(\mathcal{A})$; we assume $d(a, a) = 0$. If $\vec{a} = (a_1, \ldots, a_n)$ and $\vec{b} = (b_1, \ldots, b_m)$, then $d(\vec{a}, \vec{b}) = \min_{ij} d(a_i, b_j)$. Given $\vec{a}$ over A, its *r-sphere* $S_r^{\mathcal{A}}(\vec{a})$ is $\{b \in A \mid d(\vec{a}, b) \leq r\}$. Its *$r$-neighborhood* $N_r^{\mathcal{A}}(\vec{a})$ is defined as a structure $N_r^{\mathcal{A}}(\vec{a})$

$$\langle S_r^{\mathcal{A}}(\vec{a}), R_1^{\mathcal{A}} \cap S_r^{\mathcal{A}}(\vec{a})^{p_1}, \ldots, R_k^{\mathcal{A}} \cap S_r^{\mathcal{A}}(\vec{a})^{p_l}, a_1, \ldots, a_n\rangle$$

in the signature that extends σ with n constant symbols. That is, the carrier of $N_r^{\mathcal{A}}(\vec{a})$ is $S_r^{\mathcal{A}}(\vec{a})$, the interpretation of the σ-relations is inherited from $\mathcal{A}$, and the n extra constants are the elements of $\vec{a}$. If $\mathcal{A}$ is understood, we write $S_r(\vec{a})$ and $N_r(\vec{a})$.

If $\mathcal{A}, \mathcal{B} \in \mathrm{STRUCT}[\sigma]$ and there is an isomorphism $N_r^{\mathcal{A}}(\vec{a}) \to N_r^{\mathcal{B}}(\vec{b})$ that sends $\vec{a}$ to $\vec{b}$, we write $\vec{a} \approx_r^{\mathcal{A},\mathcal{B}} \vec{b}$. If $\mathcal{A} = \mathcal{B}$, we write $\vec{a} \approx_r^{\mathcal{A}} \vec{b}$.

Given tuples $\vec{a} = (a_1, \ldots, a_n)$ and $\vec{b} = (b_1, \ldots, b_m)$ and an element c, we write $\vec{a}\vec{b}$ for the tuple $(a_1, \ldots, a_n, b_1, \ldots, b_m)$, and $\vec{a}c$ for $(a_1, \ldots, a_n, c)$.

[2] This property was formerly known as the bounded degree property, or the BDP, see [5,13,22,24,25, etc.] However, many found the name confusing, as the property refers to the *number* of degrees in the output being bounded, rather than the degrees themselves. Following a suggestion by Neil Immerman, we decided to change the name from BDP to BNDP.

Definition 3 (cf. [21]) *An m-ary query Q is called* local *if there exists a number $r \geq 0$ such that, for any structure $\mathcal{A}$ and any $\vec{a}, \vec{b} \in A^m$*

$$\vec{a} \approx_r^{\mathcal{A}} \vec{b} \quad \textit{implies} \quad \vec{a} \in Q(\mathcal{A}) \ \textit{ iff } \ \vec{b} \in Q(\mathcal{A}).$$

The minimum such r is called the locality rank *of Q, and is denoted by $\mathsf{lr}(Q)$. QED*

It follows from Gaifman's theorem [10] that every FO-definable query is local. Moreover, if Q is definable by a formula $\varphi(\vec{x})$, then $\mathsf{lr}(Q) \leq (7^{\mathsf{qr}(\varphi)} - 1)/2$. It was shown in [21,22] that every FO($Q_{\mathbf{u}}$), FO($\mathbf{C}$), and $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$-definable query is local. Furthermore, $\mathsf{lr}(Q) \leq 2^{\mathsf{rk}(\varphi)}$ [22].

Fact 1 (see [5]) *Every local query has the bounded number of degrees property. QED*

Thus, without auxiliary relations, queries such as transitive closure cannot be expressed in FO($\mathbf{C}$) and even in $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$.

Proviso: When we deal with queries in $\mathcal{L}+\mathcal{C}$ and $(\mathcal{L}+\mathcal{C})_w$, which are defined on structures $(\mathcal{A}, \mathcal{A}')$, $\mathcal{A}' \in \mathcal{C}$, all locality concepts (neighborhoods, degrees, etc) refer only to the σ-structure $\mathcal{A}$, and not to the auxiliary structure $\mathcal{A}'$ from $\mathcal{C}$.

Almost-linear-orders We next define the class of relations that we view as "almost linear orders." First, let $\precsim_k$ stand for the class of preorders R (reflexive transitive relations) in which every equivalence class of $R \cap R^{-1}$ (that is, $x \equiv y$ iff xRy and yRx) has size at most k. Note that $\precsim_1$ is the class of linear orders.

Let $g : \mathbb{N} \to \mathbb{R}$ be a nondecreasing function [3]. Define $<_{\restriction g}$ as the class of binary relations (A, R) such that there exists a partition $A = B \cup C$ with the following properties:

(1) The cardinality of B is at least $n - g(n)$.
(2) R restricted to B is a linear order.
(3) R restricted to C is a relation from $\precsim_2$, that is, a preorder where every equivalence class has at most two elements.
(4) For any $b \in B$ and $c \in C$, $(b, c) \in R$, and $(c, b) \notin R$.

See Figure 1 for a preorder from $<_{\restriction g}$. Actually, we show the associated successor relation in the figure. A relation from $<_{\restriction g}$ is really the transitive closure of the one shown in Figure 1. Intuitively, if g is very small (e.g., $\log\log\ldots\log n$),

[3] One can deal with functions $g : \mathbb{N} \to \mathbb{N}$ as well; however, as in many examples we use $\log_2$, we prefer to have $\mathbb{R}$ as the range.

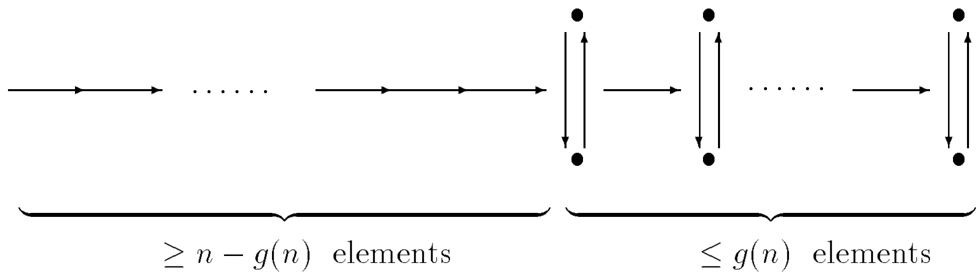

Fig. 1. A relation from $<\restriction_g$

then this can be viewed as the least possible "damage" that can be done to a linear ordering: We make a small subset at the end into a preorder, with its classes having no more than 2 elements.

3 Expressivity bounds for FO(C) and FO($Q_{\mathbf{u}}$) in the presence of relations of large degree

We start by giving a general technique for proving expressivity bounds for local logics. Then we apply it to FO(**C**) to prove our main result that DLOGSPACE-complete problems, such as deterministic transitive closure, cannot be expressed in it in the presence of relations that are very close to linear orderings. In particular, it will follow that DLOGSPACE $\not\subseteq$ FO(**C**)+ $\precsim_k$ for any $k > 1$.

Proving expressivity bounds in local logics Let Q be a query that takes structures from STRUCT[σ] as inputs and returns m-ary relations; eg., transitive closure takes graphs from STRUCT[$\sigma_{\rm gr}$] as inputs and returns graphs. Let $\mathcal{R}$ be a class of relations, and $\mathcal{L}$ a logic. Suppose we want to prove that $Q \notin (\mathcal{L} + \mathcal{R})_w$. For that purpose, we introduce two conditions.

Def$_{\mathcal{L}[\sigma]}[\mathcal{R}, \mathcal{C}]$ Assume $\mathcal{C} \subseteq$ STRUCT[σ]. Then there exists a number n and an $\mathcal{L}$ formula φ in the vocabulary σ such that $\varphi[\mathcal{A}] \in \mathcal{R}$ for every $\mathcal{A} \in \mathcal{C}$ with $|A| > n$.

That is, relations from $\mathcal{R}$ are definable by σ-formulae of $\mathcal{L}$ on large enough structures from $\mathcal{C}$.

Sep$_{\mathcal{L}[\sigma]}[Q, \mathcal{C}]$ For any two numbers $r, n > 0$, there exists $\mathcal{A} \in \mathcal{C}$ with $|A| > n$ and two m-ary vectors $\vec{a}$, $\vec{b}$ of elements of A such that $\vec{a} \approx_r^{\mathcal{A}} \vec{b}$, $\vec{a} \in Q(\mathcal{A})$ and $\vec{b} \notin Q(\mathcal{A})$.

That is, Q separates similarly looking (in a local neighborhood) tuples on arbitrarily large structures from $\mathcal{C}$.

Theorem 1 *Assume that $\mathcal{L}$ is* FO, FO(**C**), FO($Q_{\mathbf{u}}$), *or* $\mathcal{L}^*_{\infty\omega}$(**C**). *Suppose for a given query Q on σ-structures, one can find $\mathcal{C} \subseteq$* STRUCT[σ] *such that both*

$\mathbf{Def}_{\mathcal{L}[\sigma]}[\mathcal{R},\mathcal{C}]$ *and* $\mathbf{Sep}_{\mathcal{L}[\sigma]}[Q,\mathcal{C}]$ *hold. Then* $Q \notin (\mathcal{L}+\mathcal{R})_w$.

Proof: Assume that Q is in $(\mathcal{L}+\mathcal{R})_w$. Since $\mathbf{Sep}_{\mathcal{L}[\sigma]}[Q,\mathcal{C}]$ holds, Q returns nonempty results for arbitrarily large structures. Thus, we assume that it is definable by a formula ψ in the vocabulary that includes σ and a symbol R for the relation from $\mathcal{R}$. Let ψ' be obtained from ψ by replacing each occurrence of $R(\cdots)$ by $\varphi(\cdots)$, where φ is given by $\mathbf{Def}_{\mathcal{L}[\sigma]}[\mathcal{R},\mathcal{C}]$. Note that ψ' is a $\mathcal{L}$-formula in the vocabulary σ. By [21,22], ψ' is local. Let $r = \mathsf{lr}(\psi')$. For an arbitrary $N > n$, we find a structure $\mathcal{A}$ of cardinality at least N and $\vec{a} \approx_r^{\mathcal{A}} \vec{b}$ such that $\vec{a} \in Q(\mathcal{A})$ and $\vec{b} \notin Q(\mathcal{A})$. Since $Q \in (\mathcal{L}+\mathcal{R})_w$, we know that ψ is invariant on $\mathcal{A}$. Thus, $Q(\mathcal{A}) = \psi[(\mathcal{A},R)]$ where R is any interpretation of a relation from $\mathcal{R}$ on A. ¿From the invariance we obtain $\psi'[\mathcal{A}] = Q(\mathcal{A})$. Thus, for $\vec{a},\vec{b}$, we have $\mathcal{A} \models \neg(\psi'(\vec{a}) \leftrightarrow \psi'(\vec{b}))$, which contradicts the locality of ψ'. This proves the theorem. QED

Note that this theorem can be straightforwardly extended to the case of several built-in relations of possibly different arities by considering $\vec{\mathcal{R}}$ instead of $\mathcal{R}$, where $\vec{\mathcal{R}}$ is a tuple of classes of auxiliary relations. Then $\mathbf{Def}_{\mathcal{L}[\sigma]}[\vec{\mathcal{R}},\mathcal{C}]$ says that relations from each component of $\vec{\mathcal{R}}$ can be defined by a σ-formula of $\mathcal{L}$ on sufficiently large structures from $\mathcal{C}$.

Lower bounds for transitive closure Recall [16] that deterministic transitive closure of a graph is obtained by closing its deterministic paths. That is, if $G = \langle V,E\rangle$ is a directed graph, then $dtc(G) = \langle V,E'\rangle$ where $(a,b) \in E'$ iff either $(a,b) \in E$ or there exists a path $(a,a_1),(a_1,a_2),\ldots,(a_{n-1},a_n),(a_n,b) \in E$ such that a and each a_i, $i = 1,\ldots,n$ have outdegree 1. That is, the edge (a,a_1) is the only outgoing edge from a, etc.

We shall use tc to denote the transitive closure of a graph. According to [16], FO$+dtc+<$ captures DLOGSPACE and FO$+tc+<$ captures NLOGSPACE. Note that $<$ can be replaced by a successor relation, since its (deterministic) transitive closure is a linear order.

Theorem 2 *Let* $g : \mathbb{N} \to \mathbb{R}$ *be a nondecreasing function that is not bounded by a constant. Then (deterministic) transitive closure is not in* $(\mathcal{L}+ <_{\lceil g})_w$*, where* $\mathcal{L}$ *is* FO(**C**)*,* FO($Q_{\mathbf{u}}$)*, or* $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$*. In particular,* DLOGSPACE $\not\subseteq$ FO(**C**)$+ <_{\lceil g}$. *QED*

This can be compared with the results of [4] where it was shown that first-order logic with fixpoint and counting fails to express some polynomial-time problems even in the presence of relations from $\precsim_4$. Of course, first-order logic with fixpoint captures polynomial time in the presence of an order relation, cf. [7].

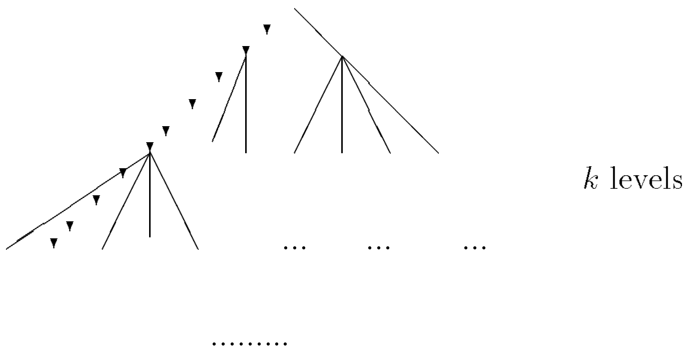

Fig. 2. Canonical k-bushy tree

In view of Theorem 1, to establish Theorem 2, it is enough to prove the following:

Proposition 1 *Let q be (deterministic) transitive closure, and $\mathcal{L}$ be* FO(**C**) *or* FO($Q_{\mathbf{u}}$). *Assume that $g : \mathbb{N} \to \mathbb{R}$ is a nondecreasing function that is not bounded by a constant. Then there exists a class $\mathcal{C}$ of graphs such that both* $\mathbf{Def}_{\mathcal{L}[\sigma_{\mathrm{gr}}]}(<\restriction_g, \mathcal{C})$ *and* $\mathbf{Sep}_{\mathcal{L}[\sigma_{\mathrm{gr}}]}(q, \mathcal{C})$ *hold.*

Bushy trees In what follows, trees are directed graphs with edges oriented from the root to the leaves.

A tree is called *bushy* if, for any two non-leaf nodes $x \neq y$, $\mathit{out\text{-}deg}(x) \neq \mathit{out\text{-}deg}(y)$. A *$k$-bushy* tree is a bushy tree in which every path from the root to a leaf has the same length k. A *canonical* k-bushy tree is obtained as follows. We start with the root of outdegree 2. Its first child has 3 children, the second child has 4 children. This completes level 2, and we now have 7 elements at level 3. They will have 5, 6, 7, 8, 9, 10 and 11 children respectively. This gives us 56 nodes at level 4, which will have 12(=11+1), 13, ..., 67(=11+56) children respectively. We continue until we fully filled all k levels. See the picture in Figure 2. We use B_k to denote the canonical k-bushy tree.

Proof of Proposition 1: We start by defining a family of graphs $G^0_{d,k}$, $d, k \in \mathbb{N}_+$, $d > k + 1$. Let s_k be the total number of nodes in the canonical k-bushy tree. The root of $G^0_{d,k}$ has $s_k + 1$ children. Two of them are roots of two copies of a canonical k-bushy tree, denoted here by B^1_k and B^2_k. The other $s_k - 1$ nodes at the second level, we give $s_k + 2, s_k + 3, \ldots, s_k + s_k = 2s_k$ children respectively. Now, to those nodes at the second level that do not belong to the two canonical k-bushy trees, we give $2s_k + 1$, $2s_k + 2, \ldots$ children, as before, increasing the number by one. We continue this process until we fully fill the $k + 1$st level.

Now that the $k + 1$st level is filled (i.e. we have a graph with all paths from root to leaves being of length $k + 1$), we look at the node at the level k with most children, say M of them, and start giving nodes at the $k+1$st level $M+1, M+2, M+3, \ldots$ children. We stop the process when we completely fill the dth level.

This is the graph $G^0_{d,k}$. Note that every two non-leaf nodes $x \neq y$ have different outdegrees, unless one of them is in B^1_k and the other is in B^2_k. We now define $G_{d,k}$ by adding an arbitrary linear ordering on the leaves. That is, if Z is the set of leaves of $G^0_{d,k}$, and L_Z is a linear order on Z, then the set of edges of $G_{d,k}$ is that of $G^0_{d,k}$ plus L_Z. When we speak of "leaf nodes" of $G_{d,k}$, we actually mean the leaf nodes of $G^0_{d,k}$.

Let B°_1 and B°_2 be the sets of non-leaf nodes in B^1_k and B^2_k. Then, for any two distinct nodes $x, y \notin B^\circ_1 \cup B^\circ_2$, it is the case that $(\textit{in-deg}(x), \textit{out-deg}(x)) \neq (\textit{in-deg}(y), \textit{out-deg}(y))$. Indeed, outdegrees are different for non-leaf nodes of $G^0_{d,k}$, except for those in $B^\circ_1 \cup B^\circ_2$); all in-degrees are different for the leaf nodes; and all in-degrees for the leaf nodes, except one, are different from those of the non-leaf nodes. The exceptional leaf node is the one with in-degree one; it is thus the smallest one in the linear order, and thus has an outdegree that exceeds that of all the internal nodes.

Define two binary relations on the set of nodes: $x \prec_0 y$ iff $\textit{in-deg}(x) < \textit{in-deg}(y)$ or $\textit{in-deg}(x) = \textit{in-deg}(y)$ and $\textit{out-deg}(x) < \textit{out-deg}(y)$. Let B° be $B^\circ_1 \cup B^\circ_2$. Define:

$$x \prec y \quad \text{iff} \quad \begin{cases} x \notin B^\circ, y \in B^\circ, & \text{or} \\ x, y \in B^\circ \text{ and } x \prec_0 y, & \text{or} \\ x, y \notin B^\circ \text{ and } x \prec_0 y \end{cases}$$

From this description, it is clear that $\prec$ is definable in FO($\mathbf{C}$) and thus in FO($Q_\mathbf{u}$) and $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$, because these logics can define the set B° as the set of nodes for which there exists another node with the same in- and out-degree). Let $\alpha_\prec$ denote a formula defining $\prec$.

Now, given k, let d_k be the smallest number $d > k+1$ such that $2s_k < g(n)$ for all $n \geq N_{d,k}$, where $N_{d,k}$ is the total number of nodes in $G_{d,k}$. Since for every fixed k, $N_{d,k}$ grows with d, and g is nondecreasing, d_k is well-defined and indeed depends only on k. Let $\mathcal{C}_g = \{G_{d,k} \mid d, k \in \mathbb{N}_+, d > d_k\}$. From the construction above, it follows that, for a nondecreasing g, $\alpha_\prec$ defines an element of $<\restriction_g$ on every $G_{d,k} \in \mathcal{C}_g$. Thus, $\mathbf{Def}_{\mathcal{L}[\sigma_{\text{gr}}]}[<\restriction_g, \mathcal{C}_g]$ holds.

We now show that $\mathbf{Sep}_{\mathcal{L}[\sigma_{\text{gr}}]}[tc, \mathcal{C}_g]$ holds. Fix two numbers $r, n > 0$. We must show that there exist a graph $G_{d,k} \in \mathcal{C}_g$ such that, for two pairs of nodes (a, b) and (a', b') with $(a, b) \approx_r (a', b')$, there is a path from a to b, but there is no path from a' to b'.

Let $k = 2r + 2$, and let $d > 4r + 6$ be such that $G_{d,k} \in \mathcal{C}_g$. Let a be a node at the middle ($(r+1)$th) level of B^1_k, and a' a node in B^2_k with the same out-degree as a. Then $a \approx_r a'$. We now consider a path from a to a leaf, say l, and choose a node b on this path such that $d(b,a) > 2r+1$ and $d(b,l) > r$; this is possible since $d > 4r+6$. It is clear that $d(b,a') > 2r+1$; hence we have $(a,b) \approx_r (a',b)$. Furthermore, $(a,b) \in tc(G_{d,k})$, but $(a',b) \notin tc(G_{d,k})$, since $G^0_{d,k}$ is a tree. This proves $\mathbf{Sep}_{\mathcal{L}[\sigma_{\text{gr}}]}[tc, \mathcal{C}_g]$.

To complete the proof for deterministic transitive closure, we just reverse all the edges of $G_{d,k}$. Since all paths not involving leaves now become deterministic, the above proof works for the deterministic case. QED

Corollary 1 *(Deterministic) transitive closure is not definable in* FO(**C**) *or* FO($Q_{\mathbf{u}}$) *in the presence of relations from* $\precsim_k$ *(preorders of width at most* k*) for any* $k > 1$*. In particular,* DLOGSPACE $\not\subseteq$ FO(**C**)+ $\precsim_k$*.* *QED*

Limitations of the technique To summarize what has been achieved so far, we know that FO(**C**)+ $<$ $=$ TC^0, and the above results show that for any $k > 1$, DLOGSPACE $\not\subseteq$ FO(**C**)+ $\precsim_k$. Furthermore, DLOGSPACE $\not\subseteq$ FO(**C**)+ $<_{\lceil g}$ for any nondecreasing function g that is not bounded by a constant. Thus, one may ask if the techniques can be pushed further to prove expressivity bounds for FO(**C**)+ $<$. A possible avenue for attacking the problem of expressivity with linear order seems to be the following: try to find a class of structures $\mathcal{C}$ so that both $\mathbf{Def}_{\mathcal{L}}[<, \mathcal{C}]$ and $\mathbf{Sep}_{\mathcal{L}}[Q, \mathcal{C}]$ would hold, where Q is tc, or dtc, or any other query we want to show to be outside of FO(**C**)+ $<$. If we were able to find such a class $\mathcal{C}$, it would show that $Q \notin$ FO(**C**)+ $<$. Unfortunately, as the following proposition shows, no such class exists.

Proposition 2 *Let* Q *be a query invariant under isomorphisms. Let* $\mathcal{L}$ *be* FO(**C**)*,* FO($Q_{\mathbf{u}}$)*, or* $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$*. Then there does not exist a class* $\mathcal{C}$ *of structures such that both* $\mathbf{Def}_{\mathcal{L}}[<, \mathcal{C}]$ *and* $\mathbf{Sep}_{\mathcal{L}}[Q, \mathcal{C}]$ *hold.*

Proof: Assume that a class $\mathcal{C}$ of structures satisfying both $\mathbf{Def}_{\mathcal{L}}[<, \mathcal{C}]$ and $\mathbf{Sep}_{\mathcal{L}}[Q, \mathcal{C}]$ exists. That is, a linear order is definable by an $\mathcal{L}$ formula $\alpha(x,y)$ on large enough structures (cardinality $> n$) in $\mathcal{C}$. Since every query definable in FO(**C**), FO($Q_{\mathbf{u}}$), or $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$ are local, we know that α is local. Let $r = \mathsf{lr}(\alpha)$, and let $d = 3r+1$. Now apply $\mathbf{Sep}_{\mathcal{L}}[Q, \mathcal{C}]$ to d and n to find a structure $\mathcal{A}$ of cardinality at least n such that for some $\vec{a}, \vec{b}$, we have $\vec{a} \approx_d \vec{b}$. Let a_0 be the first component of $\vec{a}$ and b_0 be the first component of $\vec{b}$. We then have $a_0 \approx_d b_0$.

Assume without loss of generality that $a_0 \prec b_0$ but $b_0 \not\prec a_0$ in the linear order $\prec$ defined by α on $\mathcal{A}$. It follows from [5,21] that there is a permutation μ on $S_{2r+1}(a_0, b_0)$ such that $N_r(a_0, x) \cong N_r(b_0, \mu(x))$ for every $x \in S_{2r+1}(a_0, b_0)$.

From the locality of α and $r = \mathsf{lr}(\alpha)$, we get that for every $x \in S_{2r+1}(a_0, b_0)$, $a_0 \prec x$ iff $b_0 \prec \mu(x)$. That is, μ maps $\{x \in S_{2r+1}(a_0, b_0) \mid a_0 \prec x\}$ into $\{x \in S_{2r+1}(a_0, b_0) \mid b_0 \prec x\}$. But we know that the latter has fewer elements than the former (since $a_0 \prec b_0$ but $b_0 \not\prec a_0$), so we have an injective map from a bigger finite set to a smaller finite set. This contradiction completes the proof. QED

4 Expressive power with preorders

While we showed that the technique of Theorem 2 cannot be straightforwardly extended to deal with linear orders, we have not answered the following question: Is $\mathrm{FO}(\mathbf{C})+ <_{\restriction g}$ properly contained in $\mathrm{FO}(\mathbf{C})+ <$? If the two were shown to be equal, the bounds of Theorem 2 would apply to prove that the transitive closure is not in TC^0. However, we will show that there is an enormous gap between $\mathcal{L}+ <_{\restriction g}$ and $\mathcal{L}+ <$, where $\mathcal{L}$ is one of the counting logics we consider here, and g is very small. Namely, our main result is the following.

Theorem 3 *Let $g : \mathbb{N} \to \mathbb{R}$ be a nondecreasing function that is not bounded by a constant. Then every query in $(\mathcal{L}^*_{\infty\omega}(\mathbf{C})+ <_{\restriction g})_w$ has the bounded number of degrees property.*

That is, with auxiliary structures arbitrarily close to linear orders, the most powerful of counting logics, $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$, still exhibits the very tame behavior typical for FO queries over unordered structures.

The proof of this result is somewhat involved, and will be given in the next section. Here we state some corollaries.

Corollaries With g as above, the (deterministic) transitive closure, and, more generally, problems complete for classes DLOGSPACE and above it under first-order reductions, are not definable in any of the counting logics we consider, even in the presence of relations from $<_{\restriction g}$. That is,

Corollary 2 *Let $g : \mathbb{N} \to \mathbb{R}$ be a nondecreasing function that is not bounded by a constant. Then every query in $(\mathrm{FO}(Q_{\mathbf{u}})+ <_{\restriction g})_w$, $(\mathrm{FO}(\mathbf{C})+ <_{\restriction g})_w$, $\mathcal{L}^*_{\infty\omega}(\mathbf{C})+ <_{\restriction g}$, FOQU$+ <_{\restriction g}$, or $\mathrm{FO}(\mathbf{C})+ <_{\restriction g}$ has the BNDP.*

The following corollaries demonstrate the enormous gain in expressiveness by going from "almost orders" to orders. By a *colored* graph we mean a structure of the signature $(E, U_1, \ldots, U_m)$ where E is binary, and U_is are unary. That is, it is a graph with a few selected subsets of nodes. A colored graph query is a binary query Q on colored graphs; that is, it returns graphs. The *hardness* of such a query is defined as the function $\mathcal{H}_Q : \mathbb{N} \to \mathbb{N}$ where $\mathcal{H}_Q(n)$ is $\max\{deg\_count(Q(\mathcal{A}))\}$ with $\mathcal{A}$ ranging over structures with $|A| = n$ and E being a successor relation.

Recall that $deg\_count(\cdot)$ is the cardinality of the set of all degrees realized in a structure. That is, the hardness shows how complex the output might look like if the input is a successor relation with a few colored subsets. Note that $0 \leq \mathcal{H}_Q(n) \leq n+1$. Since every property of ordered structures is definable in $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$ [22], we obtain the following dichotomy result:

Corollary 3 • *Let $g : \mathbb{N} \to \mathbb{R}$ be any nondecreasing function that is not bounded by a constant. Let Q be a colored graph query in $\mathcal{L}^*_{\infty\omega}(\mathbf{C}) + <_{\restriction g}$. Then there exists a constant C such that $\mathcal{H}_Q(n) < C$ for all n.*
- *For any function $f : \mathbb{N} \to \mathbb{N}$ such that $0 \leq f(n) \leq n+1$, there exists a colored graph query Q in $\mathcal{L}^*_{\infty\omega}(\mathbf{C}) + <$ such that $\mathcal{H}_Q = f$.*

Thus, dropping a tiny portion of linear order (e.g., $\log\log \ldots \log n$ elements) accounts for the increase in hardness from constant to arbitrary one!

FO($\mathbf{C}$) also admits this kind of dichotomy, as there exists a colored graph query Q definable in FO($\mathbf{C}$)+ $<$ such that $\mathcal{H}_Q(n) \geq \log n$ [13]. We thus obtain:

Corollary 4 *There are problems in uniform* TC^0 *that cannot be expressed in* FO($\mathbf{C}$)+ $<_{\restriction g}$. *QED*

Moreover, it is known that there are uniform AC^0 (equivalently, first order logic with the BIT predict and a linear order, FO(BIT)+ $<$) queries that violate the BNDP [6,11]. Hence, we obtain:

Corollary 5 $\mathrm{AC}^0 \not\subseteq (\mathcal{L}^*_{\infty\omega}(\mathbf{C}) + <_{\restriction g})_w$. *QED*

We thus can finally compare the expressive power of counting logics with linear orders vs. their expressiveness with preorders:

Corollary 6 *Let $g : \mathbb{N} \to \mathbb{R}$ be as in Theorem 3, and $\mathcal{L}$ be* FO($\mathbf{C}$), *or* FO($Q_{\mathbf{u}}$), *or* $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$. *Then $\mathcal{L} + <_{\restriction g} \neq \mathcal{L} + <$. Furthermore,* FO($\mathbf{C}$) $\subsetneqq$ FO($\mathbf{C}$)+ $<_{\restriction g}$.

Note that the presence of some form of counting is essential in these results. It was shown recently [11] that every invariant query in FO+ $<$ has the BNDP. That is, $\mathcal{H}_Q$ is bounded by a constant for colored graph queries in FO+ $<$.

5 Proof of the Main Theorem: Failure of Locality, Weak Locality, and Bijective Games

5.1 Failure of locality

All proofs of the BNDP that are currently known derive it from locality of queries. Unfortunately, we cannot use this method as queries in $(\mathcal{L}^*_{\infty\omega}(\mathbf{C}) + <_{\restriction g})_w$ need not be local.

Proposition 3 *Let $g(n) < \frac{\log n}{\log\log n}$ be nondecreasing, and not bounded by a constant. Then there exist nonlocal queries in $(\mathcal{L}^*_{\infty\omega}(\mathbf{C}) + <_{\restriction g})_w$.*

Proof: We construct a query Q definable by a formula $\varphi(x)$ and a sequence of structures $\mathcal{A}_n$, $n \in \mathbb{N}$, with an n-element universe, so that for each n large enough and for any $P \in <_{\restriction g}$, there are two points a, b in $\mathcal{A}_n$ with isomorphic r-neighborhoods, and $(\mathcal{A}_n, P) \models \varphi(a) \wedge \neg\varphi(b)$, where $r = O(\log\log n)$. This will prove that $(\mathcal{L}^*_{\infty\omega}(\mathbf{C}) + <_{\restriction g})_w$ is not local. By $\log n$ we mean $\lfloor \log_2(n+1) \rfloor$.

The signature σ consists of three unary relations U_1, U_2 and C, and one binary relation E. We use P for the auxiliary relation from $<_{\restriction g}$. Let $l(n) = \lfloor \frac{\log(n - \log n)}{g(n)+1} \rfloor$. This function is not bounded by a constant as n grows. In $\mathcal{A}_n$, whose universe is denoted by A, U_1 has cardinality $M_n = l(n)(g(n)+1) \leq \log(n - \log n)$, and U_2 is interpreted as $A - U_1^{\mathcal{A}_n}$. The unary relation C is interpreted as a two-element subset of U_2. Let E' be defined on U_1 to be a disjoint union of $g(n)+1$ successor relations of length $l(n)$ each. For each such successor relation E'_i, $i = 1, \ldots, g(n)+1$, let c_i be the node at the distance $\lfloor l(n)/3 \rfloor$ from the start node, and d_i be the node at the distance $\lfloor 2 \cdot l(n)/3 \rfloor$ from the start node. Let $C^{\mathcal{A}_n} = \{a, b\}$. We then define

$$E^{\mathcal{A}_n} \;\; = \;\; E' \; \cup \; \bigcup_{i=1}^{g(n)+1} \{(a, c_i), (b, d_i)\}.$$

Next, define $\alpha(x) \equiv \forall y.(P(x,y) \wedge P(y,x) \rightarrow y = x)$ saying that x is in the linear order part of P, which we shall denote by P_α. From the above, we obtain $M_n \leq \log(|P_\alpha|)$. We now show that there exists a formula $\beta(x,y)$ in FO($\mathbf{C}$) such that $\beta(x,y)$ implies $x, y \in C$ and $(\mathcal{A}_n, P) \models \beta(a,b)$ and $(\mathcal{A}_n, P) \models \neg\beta(b,a)$ for any interpretation of P as a relation from $<_{\restriction g}$ on A. This will clearly suffice to conclude the proof, since one then defines $\varphi(x) \equiv C(x) \wedge \exists y.(C(y) \wedge \beta(x,y) \wedge \neg(x = y))$ and notices $(\mathcal{A}_n, P) \models \varphi(a) \wedge \neg\varphi(b)$ while a and b have isomorphic neighborhoods of radius $O(l(n))$.

The formula $\beta(x,y)$ is defined as $C(x) \wedge C(y) \wedge \exists u, v.(E(x,u) \wedge E(y,v) \wedge \gamma(u,v))$ where $\gamma(u,v)$ holds iff there is an E-path from u to v all of whose nodes are in P_α. Since the number of successor relations in E is $g(n)+1$, at least one of them is totally contained in P_α, which shows that $(\mathcal{A}_n, P) \models \beta(a,b)$. Since there is no path between d_i and c_i for every i, we have $(\mathcal{A}_n, P) \models \neg\beta(a,b)$. Thus, we must show how to express γ in $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$ (in fact, one can express it in $\text{FO}(\mathbf{C})$).

To express γ, we follow the proof of the failure of the BNDP for $\text{FO}(\mathbf{C})+<$ given in [13]. Let $P^1_\alpha = U_1 \cap P_\alpha$. Since $|P^1_\alpha| \leq |U_1| = M_n \leq \log(|P_\alpha|)$, subsets of P^1_α can be coded by the elements of P_α: a set $S \subseteq P^1_\alpha$ is coded by $c_S \in P_\alpha$ such that

$$\{x \mid \text{BIT}(m_1, m_2), \text{ where } m_1 = |\{y \mid y < x\}|,\ m_2 = |\{y \mid y < c_S\}|\,\} \quad = \quad S$$

With this coding, we can simulate monadic second-order on P^1_α in $\text{FO}(\mathbf{C})$, which suffices to express γ. This concludes the proof. QED

Proposition 3 provides the first nontrivial example that separates the notion of locality from the BNDP. Now one needs a different technique to prove Theorem 3. We introduce this technique in two steps. In the next subsection, we consider two ways of weakening the notion of locality, and we show that one of them, weak semi-locality, implies the BNDP. In subsection 5.3, we show how the bijective games [12] can be used to prove weak semi-locality of $(\mathcal{L}^*_{\infty\omega}(\mathbf{C})+<_{\lceil g})_w$ queries.

5.2 Weak locality

To define locality of a query, we considered the equivalence relation $\vec{a} \approx^{\mathcal{A}}_r \vec{b}$ iff $N^{\mathcal{A}}_r(\vec{a}) \cong N^{\mathcal{A}}_r(\vec{b})$. We now consider two refinements that lead to weaker notions of locality. For the first refinement, we write $\vec{a} \leftrightsquigarrow^{\mathcal{A}}_r \vec{b}$ if $\vec{a} \approx^{\mathcal{A}}_r \vec{b}$ and $S^{\mathcal{A}}_r(\vec{a}) \cap S^{\mathcal{A}}_r(\vec{b}) = \emptyset$.

For the other refinement, consider a partition $\mathcal{I} = (I_1, I_2)$ of the set $\{1,\ldots,n\}$. Given $\vec{x} = (x_1,\ldots,x_n)$, we denote by $\vec{x}^{\mathcal{I}}_1$ and $\vec{x}^{\mathcal{I}}_2$ the subtuples of $\vec{x}$ that consist of those components whose indices belong to I_1 or I_2, respectively. For example, if $n=4$ and $\mathcal{I} = (\{1,3\},\{2,4\})$, then $\vec{x}^{\mathcal{I}}_1 = (x_1,x_3)$ and $\vec{x}^{\mathcal{I}}_2 = (x_2,x_4)$. We then write $\vec{a} \leftrightsquigarrow^{\mathcal{A}}_r \vec{b}$, for $\vec{a},\vec{b} \in A^n$, if there exists a partition $\mathcal{I} = (I_1, I_2)$ of $\{1,\ldots,n\}$ such that

- $\vec{a}^{\mathcal{I}}_1 \approx^{\mathcal{A}}_r \vec{b}^{\mathcal{I}}_1$;
- $\vec{a}^{\mathcal{I}}_2 = \vec{b}^{\mathcal{I}}_2$;
- $S^{\mathcal{A}}_r(\vec{a}^{\mathcal{I}}_1)$, $S^{\mathcal{A}}_r(\vec{a}^{\mathcal{I}}_2)$, $S^{\mathcal{A}}_r(\vec{b}^{\mathcal{I}}_1)$ are disjoint.

Clearly, $\vec{a} \stackrel{\leftrightsquigarrow}{\leftrightsquigarrow}{}_r^{\mathcal{A}} \vec{b}$ implies $\vec{a} \leftrightsquigarrow_r^{\mathcal{A}} \vec{b}$ (by taking I_2 to be empty), and $\vec{a} \leftrightsquigarrow_{r+1}^{\mathcal{A}} \vec{b}$ implies $\vec{a} \approx_r^{\mathcal{A}} \vec{b}$.

Definition 4 *An m-ary query Q on σ-structures is called* weakly local *if there exists a number $r \in \mathbb{N}$ such that for any $\mathcal{A} \in \mathrm{STRUCT}[\sigma]$ and any $\vec{a}, \vec{b} \in A^m$, $\vec{a} \stackrel{\leftrightsquigarrow}{\leftrightsquigarrow}{}_r^{\mathcal{A}} \vec{b}$ implies $\vec{a} \in Q(\mathcal{A})$ iff $\vec{b} \in Q(\mathcal{A})$.*

A query Q is said to be weakly semi-local *if there exists a number $r \in \mathbb{N}$ such that for any $\mathcal{A} \in \mathrm{STRUCT}[\sigma]$ and any $\vec{a}, \vec{b} \in A^m$, $\vec{a} \leftrightsquigarrow_r^{\mathcal{A}} \vec{b}$ implies $\vec{a} \in Q(\mathcal{A})$ iff $\vec{b} \in Q(\mathcal{A})$.*

Proposition 4 *Every local query is weakly semi-local, and every weakly semi-local query is weakly local. There exist queries that are weakly local but not weakly semi-local, and there exist queries that are weakly semi-local but not local. That is,*

$$\textsc{Local} \quad \subsetneq \quad \textsc{Weakly Semi-Local} \quad \subsetneq \quad \textsc{Weakly Local}.$$

Proof. The chain of implications local $\Rightarrow$ weakly semi-local $\Rightarrow$ weakly local is obvious. Next, consider the following query Q_0 on graphs. If for the input graph $\langle V, E\rangle$, with vertices V and edges E, $E = \{(x_1, x_2), (x_2, x_3), \ldots, (x_{n-1}, x_n)\} \cup \{(x_i, x_i), (x_j, x_j)\}$ with $i < j$, where $V = \{x_1, \ldots, x_n\}$, then the output of Q_0 is a graph $\langle V, \{(x_i, x_j)\}\rangle$. Otherwise, the output has no edges. Clearly, this query is not local: For any r, we consider a graph as above with $i > r$, $j < n - r$ and $j - i > 2r + 1$; then $N_r(x_i, x_j) \cong N_r(x_j, x_i)$, showing that $\mathsf{lr}(Q_0)$ cannot equal r. At the same time, Q_0 is weakly semi-local, with $r = 1$ witnessing weak semi-locality. Indeed, assume that in a graph G as above $(x_k, x_l) \leftrightsquigarrow_1^G (x_i, x_j)$, with $(x_k, x_l) \neq (x_i, x_j)$. Since there are only two nodes x_i and x_j with loops, we get that $k = j$ and $l = i$, but this contradicts $(x_k, x_l) \leftrightsquigarrow_1^G (x_i, x_j)$. Thus, whenever we have $(x_k, x_l) \leftrightsquigarrow_1^G (x_s, x_t)$ with $(x_k, x_l) \neq (x_s, x_t)$, none of the pairs is (x_i, x_j) and hence $(x_k, x_l) \notin Q_0(G)$ and $(x_s, x_t) \notin Q_0(G)$, proving weak semi-locality.

To separate weak locality from weak semi-locality, consider another graph query Q_1. If its input $G = \langle V, E\rangle$ is of the form $E = \{(x_1, x_2), (x_2, x_3), \ldots, (x_{n-1}, x_n)\} \cup \{(x_i, x_i), (x_j, x_j), (x_k, x_k)\}$ with $i < j < k$, where $V = \{x_1, \ldots, x_n\}$, then the output of Q_1 is the graph $\langle V, \{(x_i, x_j)\}\rangle$. Otherwise, the output has no edges. To show that it is not weakly semi-local, let G be as above, $r > 0$, and let i, j, k be such that $i > r, j > i + 2r + 1, k > j + 2r + 1, n - k > 2r + 1$. Then $(x_i, x_j) \leftrightsquigarrow_r^G (x_i, x_k)$ but $(x_i, x_j) \in Q_1(G)$ while $(x_i, x_k) \notin Q_1(G)$, and r cannot witness weak semi-locality. To show that Q_1 is weakly local, consider again G as above, and let $(x_i, x_j) \stackrel{\leftrightsquigarrow}{\leftrightsquigarrow}{}_1^G (x_s, x_t)$. Since x_s and x_t must then be distinct and have loops, it is impossible that $S_1^G(x_i, x_j) \cap S_1^G(x_s, x_t) = \emptyset$. Thus, whenever we have $(x_s, x_t) \stackrel{\leftrightsquigarrow}{\leftrightsquigarrow}{}_1^G (x_p, x_q)$, none of the pairs is (x_i, x_j) and hence $(x_s, x_t) \notin Q_1(G)$ and $(x_p, x_q) \notin Q_1(G)$. Therefore, Q_1 is weakly local. QED

We study these notions because they are easier to prove than the BNDP, and we will see that the BNDP can be derived from them. The notion of weak locality is particularly simple: the only difference between it and locality is the disjointness of neighborhoods. However, it only gives us a partial result:

Proposition 5 *a) Let Q be a binary weakly local query (i.e., the output is a graph). Then Q has the bounded number of degrees property.*

b) For every $m > 2$, there exists an m-ary weakly local query that does not have the bounded number of degrees property.

Proof. We first prove a). Fix the relational signature σ, and let $F_0(d, k)$ be an upper bound on the size of a d-neighborhood of a point in $\mathcal{A}$ where degrees are bounded by k, and $F_1(d, k)$ be an upper bound on the number of isomorphism types of d-neighborhoods of points in such structures; such bounds do exist and depend on σ, d and k only [21].

Let Q be a binary weakly local query, with d witnessing weak locality. We now calculate the number of different outdegrees in the graph $Q(\mathcal{A})$, where all degrees in the structure $\mathcal{A}$ are at most k. Whenever we say *out-deg*(x), we mean out-degree in $Q(\mathcal{A})$.

Claim 2 *There is a number $M_{d,k}$ that depends on d and k only, such that*

$$|out\text{-}deg(a) - out\text{-}deg(b)| \;\leq\; M_{d,k}$$

whenever $a \leftrightsquigarrow_{2d+1} b$.

Proof of the claim: We call an isomorphism type τ of a d-neighborhood of a point (a, b)-*good* if there exist three points, $c_1, c_2, c_3 \in A - S_{2d+1}(a, b)$ such that c_1, c_2, c_3 realize τ and $d(c_i, c_j) > 4d, i, j = 1, 2, 3, i \neq j$. We call τ (a, b)-*bad* otherwise.

Let c be a point of a (a, b)-good type τ. Let c_1, c_2, c_3 witness the goodness of the τ. Then at most one of them can belong to $S_{2d}(c)$ (otherwise the distance would be below $4d$). Thus, there are two points c', c'' of type τ such that $d(c, c'), d(c, c''), d(c', c'') > 2d$; that is, their d-neighborhoods are disjoint. We now obtain

$$\begin{aligned} & (a, c) \in Q(\mathcal{A}) \\ \Leftrightarrow\; & (b, c') \in Q(\mathcal{A}) \\ \Leftrightarrow\; & (a, c'') \in Q(\mathcal{A}) \\ \Leftrightarrow\; & (b, c) \in Q(\mathcal{A}) \end{aligned}$$

by weak locality.

Now, letting M_0 be the number of points realizing (a,b)-bad types, we see that the difference between *out-deg*(a) and *out-deg*(b) cannot exceed the size of $S_{2d+1}(a,b) + M_0$; that is,

$$|\textit{out-deg}(a) - \textit{out-deg}(b)| \;\leq\; M_0 + 2 \cdot F_0(2d+1, k)$$

It thus remains to show that M_0 is determined by d and k.

Fix an (a,b)-bad type τ, and let $x \notin S_{2d+1}(a,b)$ realize τ. Suppose $y \notin S_{2d+1}(a,b) \cup S_{4d}(x)$ realizes τ. Then every other point z realizing τ must be either in $S_{2d+1}(a,b)$ or in $S_{4d}(x,y)$, for otherwise x,y,z would witness (a,b)-goodness of τ. Thus, the number of points realizing τ is at most $2 \cdot F_0(2d+1,k) + 2 \cdot F_0(4d,k)$, and hence M_0 is bounded above by

$$F_1(d,k) \cdot (2 \cdot F_0(2d+1,k) + 2 \cdot F_0(4d,k))$$

finishing the proof of the claim. QED

Using this, we show the following.

Claim 3 *Let $m_0 = F_0(8d+4,k)+1$. Suppose $a \approx_{2d+1} b$, and suppose that there are at least m_0 realizers of the isomorphism type of the $2d+1$-neighborhood of a. Then*

$$|\textit{out-deg}(a) - \textit{out-deg}(b)| \;\leq\; 2 \cdot M_{d,k}$$

Proof of the claim: If $S_{2d+1}(a) \cap S_{2d+1}(b) = \emptyset$, this follows from the previous claim. Assume then $S_{2d+1}(a) \cap S_{2d+1}(b) \neq \emptyset$. We have $S_{2d+1}(a,b) \subseteq S_{6d+3}(a) = C$, and from the assumptions, we obtain there exists an element $c \notin S_{2d+1}(C)$ such that $a \approx_{2d+1} c$. In particular, in this case $S_{2d+1}(a) \cap S_{2d+1}(c) = \emptyset$ and $S_{2d+1}(b) \cap S_{2d+1}(c) = \emptyset$. Then by the previous claim we have $|\textit{out-deg}(a) - \textit{out-deg}(c)| \leq M_{d,k}$ and $|\textit{out-deg}(b) - \textit{out-deg}(c)| \leq M_{d,k}$, which proves the claim. QED

Let now m_0 be as in Claim 3. Suppose τ is an isomorphism type of a $2d+1$-neighborhood that has fewer than m_0 realizers. The total number of points of such types is at most $M_1 = F_1(2d+1,k) \cdot m_0$, and thus they give rise to at most M_1 different outdegrees in the output. For any point a of a type τ (of $2d+1$-neighborhood) that is realized at least m_0 times, the possible outdegrees belong to a $2(2M_{d_k})+1$ element interval, by Claim 3. Thus, the total number of outdegrees in $Q(\mathcal{A})$ is at most

$$M_1 + F_1(2d+1,k) \cdot (4M_{d,k} + 1)$$

and thus depends on the signature, d and k only. The proof for a bound on the number of indegrees is identical. This completes the proof of the BNDP for weakly local graph queries.

To show b), we consider graphs as inputs, and let $m = 3$; extension to $m > 3$ is straightforward. For a graph $G = \langle V, E \rangle$, with vertices V and edges E, $(a, b, c) \in Q(G)$ iff the following two conditions hold. First, the graph G is of the special form: there is an element $v \in V$ such that $(v, v) \in E$, $(v, v'), (v', v) \notin E$ for any $v' \neq v$, and G restricted to $V - \{v\}$ is a chain (i.e., the graph of a successor relation). Second, $c = v$, $a, b \neq v$, and (a, b) is in the transitive closure of G.

Clearly, Q violates the BNDP: for every $a \in V - \{v\}$, there are k_a tuples (a, b, c) in $Q(G)$, where k_a is the number of nodes reachable from a. Thus, $deg\_count(Q(G)) = O(|V|)$. On the other hand, Q is weakly local (in fact, $r = 1$). Indeed, suppose $(a, b, c) \in Q(G)$ and $(a', b', c') \notin Q(G)$. Then $c = v$, and since there is only one loop in G, for $(a, b, c) \approx_1 (a', b', c')$ to hold we must have $c = c' = v$. However, in this case $S_r(a, b, c) \cap S_r(a', b', c') \neq \emptyset$, which shows that $(a, b, c) \leftrightsquigarrow_1 (a', b', c')$ does not hold, and thus proves the weak locality of Q. This completes the proof. QED

Combined with the results of Section 5.3, that would be sufficient to derive Theorem 3 for queries that return graphs. However, for arbitrary queries, we need the more involved notion of weak semi-locality:

Theorem 4 *Every weakly semi-local query has the bounded number of degrees property.*

Proof: Let $\vec{x} = (x_1, \ldots, x_n)$ and let $\mathcal{I} = \{I_1, I_2, I_3\}$ be a partition of $\{1, \ldots, n\}$. Then by $\vec{x}_j^{\mathcal{I}}$, $j = 1, 2, 3$, we denote the subtuple of $\vec{x}$ which consists of the components of $\vec{x}$ whose indices are in I_j, appearing in the same order as in $\vec{x}$.

Given a σ-structure $\mathcal{A}$, two positive integers $d, r > 0$ and $a \leftrightsquigarrow_d^{\mathcal{A}} b$, we define a binary relation $\twoheadleftarrow\!\!\twoheadrightarrow_{r,l}^{(\mathcal{A},a,b)}$ on A^l as follows. Given $\vec{x}, \vec{y} \in A^l$, $a\vec{x} \twoheadleftarrow\!\!\twoheadrightarrow_{r,l}^{(\mathcal{A},a,b)} b\vec{y}$ iff there exists an isomorphism $h : N_d^{\mathcal{A}}(a) \to N_d^{\mathcal{A}}(b)$ and a partition $\mathcal{I}$ of $\{1, \ldots, l\}$ such that

- $S_r^{\mathcal{A}}(\vec{x}_1^{\mathcal{I}}) \subseteq S_d^{\mathcal{A}}(a)$ and $\vec{y}_1^{\mathcal{I}} = h(\vec{x}_1^{\mathcal{I}})$;
- $S_r^{\mathcal{A}}(\vec{x}_2^{\mathcal{I}}) \subseteq S_d^{\mathcal{A}}(b)$ and $\vec{y}_2^{\mathcal{I}} = h^{-1}(\vec{x}_2^{\mathcal{I}})$;
- $\vec{x}_3^{\mathcal{I}} = \vec{y}_3^{\mathcal{I}}$ and $S_r^{\mathcal{A}}(\vec{x}_3^{\mathcal{I}}) \cap (S_r^{\mathcal{A}}(a, b) \cup S_r^{\mathcal{A}}(\vec{x}_1^{\mathcal{I}}\vec{x}_2^{\mathcal{I}}\vec{y}_1^{\mathcal{I}}\vec{y}_2^{\mathcal{I}})) = \emptyset$.

Note that these conditions imply $S_r^{\mathcal{A}}(\vec{y}_1^{\mathcal{I}}) \subseteq S_d^{\mathcal{A}}(b)$ and $S_r^{\mathcal{A}}(\vec{y}_2^{\mathcal{I}}) \subseteq S_d^{\mathcal{A}}(a)$.

We now need the following lemma.

Lemma 1 *For any positive integers r and l, there exists a positive integer d such that for any relational vocabulary σ, any σ-structure $\mathcal{A}$ and $a \leftrightsquigarrow_d^{\mathcal{A}} b$,* *there exists a permutation* $\pi : A^l \to A^l$ *such that*

$$a\vec{x} \;\twoheadleftarrow\!\!\twoheadrightarrow^{(\mathcal{A},a,b)}_{(r,l)}\; b\pi(\vec{x})$$

for every $\vec{x} \in A^l$.

Proof of the lemma: Let $d_0 = r, d_1 = 3d_0 + 1, \ldots, d_l = 3d_{l-1} + 1$. We claim that $d = d_l$. The proof is by induction on l. Below h stands for the isomorphism given by $a \leftrightsquigarrow^{\mathcal{A}}_{d} b$. If $l = 1$, we define $\pi : A \to A$ as follows:

$$\pi(x) \;=\; \begin{cases} h(x), & \text{if } x \in S^{\mathcal{A}}_{2r+1}(a); \\ h^{-1}(x), & \text{if } x \in S^{\mathcal{A}}_{2r+1}(b); \\ x, & \text{otherwise.} \end{cases}$$

The partition of the set $\{1\}$ containing the unique index is determined as follows:

$$\mathcal{I} \;=\; \begin{cases} (\{1\}, \emptyset, \emptyset), & \text{if } x \in S^{\mathcal{A}}_{2r+1}(a); \\ (\emptyset, \{1\}, \emptyset), & \text{if } x \in S^{\mathcal{A}}_{2r+1}(b); \\ (\emptyset, \emptyset, \{1\}), & \text{otherwise.} \end{cases}$$

It is routine to verify that π is a permutation and $x \twoheadleftarrow\!\!\twoheadrightarrow^{(\mathcal{A},a,b)}_{(r,1)} \pi(x)$.

Now assume $a \leftrightsquigarrow^{\mathcal{A}}_{d} b$, where $d = d_l$, $l > 1$. By the hypothesis for $l-1$, we find a permutation $\mu : A^{l-1} \to A^{l-1}$ such that for any $\vec{x} \in A^{l-1}$, $a\vec{x} \twoheadleftarrow\!\!\twoheadrightarrow^{(\mathcal{A},a,b)}_{(d_1,l-1)} b\mu(\vec{x})$; that is,

$$a\vec{x} \twoheadleftarrow\!\!\twoheadrightarrow^{(\mathcal{A},a,b)}_{(3r+1,l-1)} b\mu(\vec{x}).$$

We now show that for every $\vec{x} \in A^{l-1}$, there exists a permutation $\eta_{\vec{x}} : A \to A$ such that $a\vec{x}z \twoheadleftarrow\!\!\twoheadrightarrow^{(\mathcal{A},a,b)}_{(r,l)} b\mu(\vec{x})\eta_{\vec{x}}(z)$; this will suffice to conclude the proof as the function given by $(x_1, \ldots, x_{l-1}, x_l) \mapsto \mu((x_1, \ldots, x_{l-1}))\eta_{(x_1,\ldots,x_{l-1})}(x_l)$ is a permutation.

Fix an isomorphism $h : N^{\mathcal{A}}_d(a) \to N^{\mathcal{A}}_d(b)$ and a partition $\mathcal{I} = (I_1, I_2, I_3)$ witnessing $a\vec{x} \twoheadleftarrow\!\!\twoheadrightarrow^{(\mathcal{A},a,b)}_{(3r+1,l-1)} b\mu(\vec{x})$. That is, $S^{\mathcal{A}}_{3r+1}(\vec{x}^{\mathcal{I}}_1) \subseteq S^{\mathcal{A}}_d(a)$, $S^{\mathcal{A}}_{3r+1}(\vec{x}^{\mathcal{I}}_2) \subseteq S^{\mathcal{A}}_d(b)$, $\mu(\vec{x})^{\mathcal{I}}_1 = h(\vec{x}^{\mathcal{I}}_1)$, $\mu(\vec{x})^{\mathcal{I}}_1 = h^{-1}(\vec{x}^{\mathcal{I}}_1)$, and $S^{\mathcal{A}}_{3r+1}(\vec{x}^{\mathcal{I}}_3)$ does not intersect the $3r+1$-spheres of $a, b, \vec{x}^{\mathcal{I}}_1, \vec{x}^{\mathcal{I}}_2, \mu(\vec{x})^{\mathcal{I}}_1$, and $\mu(\vec{x})^{\mathcal{I}}_2$. Below we show how to define $\eta_{\vec{x}}$ and a new partition $\mathcal{I}'$ that will witness $a\vec{x}z \twoheadleftarrow\!\!\twoheadrightarrow^{(\mathcal{A},a,b)}_{(r,l)} b\mu(\vec{x})\eta_{\vec{x}}(z)$; $\mathcal{I}'$ is obtained from $\mathcal{I}$ by adding the last index l to one of its blocks. The isomorphism h remains the same.

$$\eta_{\vec{x}}(z) \;=\; \begin{cases} h(z), & \text{if } z \in S^{\mathcal{A}}_{2r+1}(a) \cup S_{2r+1}(\vec{x}^{\mathcal{I}}_1) \cup S^{\mathcal{A}}_{2r+1}(\mu(\vec{x})^{\mathcal{I}}_2);\; (I'_1{:=}I_1 \cup \{l\}); \\ h^{-1}(z), & \text{if } z \in S^{\mathcal{A}}_{2r+1}(b) \cup S_{2r+1}(\mu(\vec{x})^{\mathcal{I}}_1) \cup S^{\mathcal{A}}_{2r+1}(\vec{x}^{\mathcal{I}}_2);\; (I'_2{:=}I_2 \cup \{l\}); \\ z, & \text{otherwise;} \qquad (I'_3{:=}I_3 \cup \{l\}). \end{cases}$$

Consider the first case. Let $z' = \eta_{\vec{x}}(z)$. Since $C = S^{\mathcal{A}}_{2r+1}(a) \cup S_{2r+1}(\vec{x}^{\mathcal{I}}_1) \cup S^{\mathcal{A}}_{2r+1}(\mu(\vec{x})^{\mathcal{I}}_2)$ has the property that $S^{\mathcal{A}}_r(C) \subseteq S^{\mathcal{A}}_d(a)$, we have $S^{\mathcal{A}}_r(\vec{x}^{\mathcal{I}'}_1) \subseteq S^{\mathcal{A}}_d(a)$, and since $z' = h(z)$, we obtain $S^{\mathcal{A}}_r(\mu(\vec{x})^{\mathcal{I}'}_1) \subseteq S^{\mathcal{A}}_d(b)$. Furthermore, from $S^{\mathcal{A}}_r(z) \subseteq S^{\mathcal{A}}_{3r+1}(a) \cup S_{3r+1}(\vec{x}^{\mathcal{I}}_1) \cup S^{\mathcal{A}}_{3r+1}(\mu(\vec{x})^{\mathcal{I}}_2)$, we obtain $S^{\mathcal{A}}_r(z) \cap S^{\mathcal{A}}_{3r+1}(\vec{x}^{\mathcal{I}}_3) = \emptyset$ and thus $S^{\mathcal{A}}_r(z) \cap S^{\mathcal{A}}_r(\vec{x}^{\mathcal{I}}_3) = \emptyset$. Similarly, $S^{\mathcal{A}}_r(z') \cap S^{\mathcal{A}}_r(\vec{x}^{\mathcal{I}}_3) = \emptyset$. The proof of correctness in the second case is identical. For the third case, as $z \notin S^{\mathcal{A}}_{2r+1}(a,b) \cup S^{\mathcal{A}}_{2r+1}(\vec{x}^{\mathcal{I}}_1\vec{x}^{\mathcal{I}}_2\mu(\vec{x})^{\mathcal{I}}_1\mu(\vec{x})^{\mathcal{I}}_2)$, we obtain $S^{\mathcal{A}}_r(z) \cap (S^{\mathcal{A}}_r(a,b) \cup S^{\mathcal{A}}_r(\vec{x}^{\mathcal{I}}_1\vec{x}^{\mathcal{I}}_2\mu(\vec{x})^{\mathcal{I}}_1\mu(\vec{x})^{\mathcal{I}}_2)) = \emptyset$, completing the proof that $a\vec{x}z \leftrightsquigarrow^{(\mathcal{A},a,b)}_{(r,l)} b\mu(\vec{x})\eta_{\vec{x}}(z)$. The lemma is proved. QED

We now prove the theorem using the lemma. Let Q be an m-ary weakly semi-local query, $m > 1$. Let r witness weak semi-locality; ie., $\vec{a} \leftrightsquigarrow^{\mathcal{A}}_r \vec{b}$ implies $\vec{a} \in Q(\mathcal{A})$ iff $\vec{b} \in Q(\mathcal{A})$. Let d be the positive integer given by Lemma 1 for r and $l = m - 1$. We now fix $k > 0$ and let $\mathcal{A} \in \mathrm{STRUCT}_k[\sigma]$. As before, let $F_0(d,k)$ be the maximum size of a d-neighborhood around a point in a σ-structure whose degrees do not exceed k; such a bound exists and is determined by d, k and σ [21,13]. Let τ be an isomorphism type of a d-neighborhood of a point in a σ-structure. We call it *$\mathcal{A}$-good* if there are more than $(m+1)F_0(2d+1,k)$ realizers of τ in $\mathcal{A}$, and *$\mathcal{A}$-bad* if there are at most $(m+1)F_0(2d+1,k)$ realizers.

In what follows, the structure $\mathcal{A} \in \mathrm{STRUCT}_k[\sigma]$ is fixed, and *degree* refers to the degree in the output $Q(\mathcal{A})$.

Claim 4 *Let $a \leftrightsquigarrow\!\!\!\!\leftrightsquigarrow_d b$, and assume that the isomorphism type of the d-neighborhood of a is $\mathcal{A}$-good. Then $degree_1(a) = degree_1(b)$.*

Proof of the claim: Assume that $a\vec{x} \leftrightsquigarrow^{(\mathcal{A},a,b)}_{(r,m-1)} b\vec{y}$. Let $\mathcal{I} = (I_1, I_2, I_3)$ be a partition of $\{1,\ldots,m-1\}$ and $h : N_d(a) \to N_d(b)$ an isomorphism witnessing that. Assume that $a\vec{x} \in Q(\mathcal{A})$. Let $C = S_{2d+1}(a,b) \cup S_{2d+1}(\vec{x}^{\mathcal{I}}_3)$. Since $\vec{x}^{\mathcal{I}}_3$ has at most $m-1$ elements, $|C| \leq (m+1)F_0(2d+1,k)$, and thus there exists $c \notin C$ such that $c \approx_d a \approx_d b$, since the type of a is $\mathcal{A}$-good. Note that $a \leftrightsquigarrow\!\!\!\!\leftrightsquigarrow_d c$, $b \leftrightsquigarrow\!\!\!\!\leftrightsquigarrow_d c$ and $S_d(c) \cap S_d(\vec{x}^{\mathcal{I}}_3) = \emptyset$. Let $h_a : N_d(a) \to N_d(c)$ be an isomorphism. Let $\vec{z}_1 = h_a(\vec{x}^{\mathcal{I}}_1)$. Note that $S_r(c\vec{z}_1) \subseteq S_d(c)$. We then obtain $a\vec{x}^{\mathcal{I}}_1\vec{x}^{\mathcal{I}}_2\vec{x}^{\mathcal{I}}_3 \leftrightsquigarrow_r c\vec{z}_1\vec{x}_2\vec{x}^{\mathcal{I}}_3$. By weak semi-locality of Q, we obtain $c\vec{z}_1\vec{x}^{\mathcal{I}}_2\vec{x}^{\mathcal{I}}_3 \in Q(\mathcal{A})$. Since $S_r(\vec{x}^{\mathcal{I}}_2) \cap S_r(\vec{y}^{\mathcal{I}}_2) = \emptyset$, we obtain $c\vec{z}_1\vec{x}^{\mathcal{I}}_2\vec{x}^{\mathcal{I}}_3 \leftrightsquigarrow_r c\vec{z}_1\vec{y}^{\mathcal{I}}_2\vec{y}_3$, and thus $c\vec{z}_1\vec{y}^{\mathcal{I}}_2\vec{y}^{\mathcal{I}}_3 \in Q(\mathcal{A})$. Next, notice that $h \circ h_a^{-1}$ maps $N_r(c\vec{z}_1)$ isomorphically onto $N_r(b\vec{y}^{\mathcal{I}}_1)$, which is disjoint from $S_r(\vec{y}^{\mathcal{I}}_2\vec{y}^{\mathcal{I}}_3)$; we thus conclude that $c\vec{z}_1\vec{y}^{\mathcal{I}}_2\vec{y}^{\mathcal{I}}_3 \leftrightsquigarrow_r b\vec{y}^{\mathcal{I}}_1\vec{y}^{\mathcal{I}}_2\vec{y}^{\mathcal{I}}_3 = b\vec{y}$, and $b\vec{y} \in Q(\mathcal{A})$ by weak semi-locality of Q.

An identical proof shows that $b\vec{y} \in Q(\mathcal{A})$ implies $a\vec{x} \in Q(\mathcal{A})$. Thus, by Lemma 1, we have a permutation π on A^{m-1} such that $a\vec{x} \in Q(\mathcal{A})$ iff $b\pi(\vec{x}) \in Q(\mathcal{A})$, which proves $degree_1(a) = degree_1(b)$. QED

Using this claim, we show that if $a \approx^{\mathcal{A}}_d b$, and the isomorphism type τ of $N^{\mathcal{A}}_d(b)$

is $\mathcal{A}$-good, then $degree_1(a) = degree_1(b)$. Indeed, if $S_d^{\mathcal{A}}(a) \cap S_d^{\mathcal{A}}(b) = \emptyset$, then $a \leftrightsquigarrow_d^{\mathcal{A}} b$, and we are done by Claim 4. If $S_d^{\mathcal{A}}(a) \cap S_d^{\mathcal{A}}(b) \neq \emptyset$, then, as the cardinality of $S_{2d+1}^{\mathcal{A}}(a,b)$ is at most $2F_0(2d+1,k)$, there is a point $c \notin S_{2d+1}^{\mathcal{A}}(a,b)$ realizing τ. We have now $a \leftrightsquigarrow_d^{\mathcal{A}} c$ and $b \leftrightsquigarrow_d^{\mathcal{A}} c$, and thus $degree_1(a) = degree_1(c)$, and $degree_1(b) = degree_1(c)$; hence $degree_1(a) = degree_1(b)$.

We now calculate the number of different values of $degree_1(\cdot)$. The cardinality of $\{degree_1(a) \mid \text{type of } N_d^{\mathcal{A}}(a) \text{ is } \mathcal{A} - \text{good}\}$ is at most $F_1(d,k)$, the maximum possible number of different isomorphism types of d-neighborhoods in a structure from $\text{STRUCT}_k[\sigma]$. This number depends on d, k and σ only. The cardinality of $\{degree_1(a) \mid \text{type of } N_d^{\mathcal{A}}(a) \text{ is } \mathcal{A} - \text{bad}\}$ is at most the number of points realizing $\mathcal{A}$-bad types. As each $\mathcal{A}$-bad type has at most $(m+1)F_0(2d+1,k)$ realizers, the number of realizers of $\mathcal{A}$-bad types is at most $(m+1) \cdot F_0(2d+1,k) \cdot F_1(d,k)$. We thus obtain that

$$|\{degree_1(a) \mid a \in A\}| \quad \leq \quad F_1(d,k) + (m+1) \cdot F_0(2d+1,k) \cdot F_1(d,k)$$

and thus depends only on d, k and σ. As analogous proofs work for all $degree_i$, we get $deg\_count(Q(\mathcal{A})) \leq m \cdot F_1(d,k) \cdot ((m+1)F_0(2d+1,k)+1)$, which finishes the proof of the BNDP. QED

To incorporate the information about the function g, we modify the definition as follows: $\vec{a} \leftrightsquigarrow_{g,r}^{\mathcal{A}} \vec{b}$ if $\vec{a} \leftrightsquigarrow_r^{\mathcal{A}} \vec{b}$ and $\left|S_r^{\mathcal{A}}(\vec{a}) \cup S_r(\vec{b})\right| \leq g(|A|)$. Then a query Q is *g-weakly semi-local* if there exists an $r \in \mathbb{N}$ such that $\vec{a} \leftrightsquigarrow_{g,r}^{\mathcal{A}} \vec{b}$ implies $\vec{a} \in Q(\mathcal{A})$ iff $\vec{b} \in Q(\mathcal{A})$. The following is easily derived from Theorem 4.

Corollary 7 *Let $g : \mathbb{N} \to \mathbb{R}$ be nondecreasing and not bounded by a constant. Then every g-weakly semi-local query has the BNDP.*

Proof. Let Q be g-weakly local m-ary query, with d witnessing weak locality. Let $N_{d,k}$ be the smallest number such that $g(N) > 2m \cdot F_0(d,k)$ for any $N > N_{d,k}$. Then, if $\mathcal{A} \in \text{STRUCT}_k[\sigma]$ and $|A| > N_{d,k}$, for $\vec{a}, \vec{b} \in A^m$, $\vec{a} \leftrightsquigarrow_{g,r}^{\mathcal{A}} \vec{b}$ implies $\vec{a} \leftrightsquigarrow_r \vec{b}$, since $S_d(\vec{a}) \cup S_r(\vec{b})$ has fewer than $g(|A|)$ elements. Then the proof of Theorem 4 applies verbatim to show that for some function $f_0 : \mathbb{N} \to \mathbb{N}$, $deg\_count(Q(\mathcal{A})) \leq f_0(k)$. Since $deg\_count(Q(\mathcal{A})) \leq N_{d,k} + 1$ if $|A| \leq N_{d,k}$, we obtain that $deg\_count(Q(\mathcal{A})) \leq \max\{N_{d,k}^{m-1} + 1, f_0(k)\}$, thus proving the BNDP. QED

5.3 Games and weak semi-locality

The goal of this section is to prove the g-weak semi-locality of queries in $(\mathcal{L}_{\infty\omega}^*(\mathbf{C}) + <_{\lceil g})_w$. We do this by using *bijective games* of [12].

The game is played by two players, called the spoiler and the duplicator, on two structures $\mathcal{A}, \mathcal{B} \in$ STRUCT$[\sigma]$. For the n-round game, in each round $i = 1, \ldots, n$, the duplicator selects a bijection $f_i : A \to B$, where B is the carrier of $\mathcal{B}$, and the spoiler selects a point $a_i \in A$. If $|A| \neq |B|$, then the spoiler immediately wins. The duplicator wins after n rounds if the relation $\{(a_i, f_i(a_i)) \mid 1 \leq i \leq n\}$ is an isomorphism from $\mathcal{A} \cap (\bigcup_{1 \leq i \leq n} a_i) \to \mathcal{B} \cap (\bigcup_{1 \leq i \leq n} f_i(a_i))$. Otherwise the spoiler wins. If the duplicator has a winning strategy in the n-move bijective game on $\mathcal{A}$ and $\mathcal{B}$, we write $\mathcal{A} \equiv_n^{bij} \mathcal{B}$. We write $(\mathcal{A}, \vec{a}) \equiv_n^{bij} (\mathcal{B}, \vec{b})$ if the duplicator has a winning strategy in the n-move bijective game that starts with the position $(\vec{a}, \vec{b})$; ie., each f_i sends $\vec{a}$ to $\vec{b}$. This condition implies that for a FO (or FO$(Q_{\mathbf{u}})$) formula $\varphi(\vec{x})$ of quantifier rank n, $\mathcal{A} \models \varphi(\vec{a})$ iff $\mathcal{B} \models \varphi(\vec{b})$ [12]. We extend this to $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$. Note that the lemma below follows from a slightly more general result of [14].

Lemma 2 *Let $\varphi(x_1, \ldots, x_m)$ be a $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$ formula in the language of σ, with all free variables of the first sort. Let $(\mathcal{A}, \vec{a}) \equiv^{bij}_{\mathsf{rk}(\varphi)} (\mathcal{B}, \vec{b})$, where $\vec{a} \in A^m, \vec{b} \in B^m$. Then $\mathcal{A} \models \varphi(\vec{a})$ iff $\mathcal{B} \models \varphi(\vec{b})$.* *QED*

The following is the key lemma, which is proved by a technique reminiscent of that in [32], extended to deal with bijective games.

Lemma 3 *Let $g : \mathbb{N} \to \mathbb{R}$ be nondecreasing and not bounded by a constant. For any $\mathcal{A}$, $m > 0$, $\vec{a}, \vec{b} \in A^m$, and $n > 0$, if $\vec{a} \leftrightsquigarrow^{\mathcal{A}}_{g,2^n} \vec{b}$, then there exists a preorder P on A such that $P \in <\restriction_g$ and*

$$(\mathcal{A}, P, \vec{a}) \quad \equiv_n^{bij} \quad (\mathcal{A}, P, \vec{b})$$

Proof: Let $r = 2^n$ and $\vec{a} \leftrightsquigarrow^{\mathcal{A}}_{g,r} \vec{b}$. Let $\mathcal{I} = (I_1, I_2)$ be a partition witnessing that. We assume without loss of generality that I_1 is nonempty and equals $\{1, \ldots, l\}$, $l \leq m$. Let $\vec{a}' = (a_1, \ldots, a_l)$, $\vec{b}' = (b_1, \ldots, b_l)$, and $\vec{c} = (a_{l+1}, \ldots, a_m) = (b_{l+1}, \ldots, b_m)$. Then $\vec{a}' \leftrightsquigarrow\!\!\!\leftrightsquigarrow^{\mathcal{A}}_r \vec{b}'$, $S^{\mathcal{A}}_r(\vec{a}'\vec{b}') \cap S^{\mathcal{A}}_r(\vec{c}) = \emptyset$, and $\left|S^{\mathcal{A}}_r(\vec{a}'\vec{b}'\vec{c})\right| \leq g(|A|)$.

We now construct P. Let A_0 be $S^{\mathcal{A}}_r(\vec{a}') - \{a_1, \ldots, a_l\}$. Pick any ordering $\prec_1$ on $S^{\mathcal{A}}_r(\vec{a}')$ such that $a_1 \prec_1 a_2 \prec_1 \ldots \prec_1 a_l$ and further, for any $a \in A_0$ we have $a_i \prec_1 a$, for each $i = 1, \ldots, l$, and for any $a', a'' \in A_0$,

$$d(a', \vec{a}') < d(a'', \vec{a}') \quad \Rightarrow \quad a' \prec_1 a''$$

Let h be an isomorphism of $N^{\mathcal{A}}_r(\vec{a})$ onto $N^{\mathcal{A}}_r(\vec{b})$. Define on $S^{\mathcal{A}}_r(\vec{b}')$ an ordering $\prec_2$ by letting $b' \prec_2 b''$ iff $h^{-1}(b') \prec_1 h^{-1}(b'')$. Clearly, the initial fragment of $\prec_2$ is $(b_1, \ldots, b_l)$, and it respects the distance to $\vec{b}'$: $d(b', \vec{b}') < d(b'', \vec{b}')$ implies $b' \prec_2 b''$.

Let P_0 be an arbitrary linear ordering on $A - S^{\mathcal{A}}_r(\vec{a}'\vec{b}')$. Intuitively, P is P_0

followed by a preorder obtained by putting together $\prec_1$ and $\prec_2$, and tying them by h. Formally,

$$(x,y) \in P \text{ iff } \begin{cases} x, y \notin S_r^{\mathcal{A}}(\vec{a}'\vec{b}') \text{ and } (x,y) \in P_0, \text{ or} \\ x \notin S_r^{\mathcal{A}}(\vec{a}'\vec{b}') \text{ and } y \in S_r^{\mathcal{A}}(\vec{a}'\vec{b}'), \text{ or} \\ x \in S_r^{\mathcal{A}}(\vec{a}'), y \in S_r^{\mathcal{A}}(\vec{a}') \text{ and } x \prec_1 y, \text{ or} \\ x \in S_r^{\mathcal{A}}(\vec{b}'), y \in S_r^{\mathcal{A}}(\vec{b}') \text{ and } x \prec_2 y, \text{ or} \\ x \in S_r^{\mathcal{A}}(\vec{a}'), y \in S_r^{\mathcal{A}}(\vec{b}') \text{ and } h(x) \prec_2 y, \text{ or} \\ x \in S_r^{\mathcal{A}}(\vec{b}'), y \in S_r^{\mathcal{A}}(\vec{a}') \text{ and } x \prec_2 h(y) \end{cases}$$

It easily follows from $\vec{a}' \leftrightsquigarrow^{\mathcal{A}}_{g,r} \vec{b}'$ that $P \in <\restriction_g$.

Our next claims give a winning strategy for the duplicator in the bijective game on $\mathcal{A}_{\vec{a}} = (\mathcal{A}, P, \vec{a})$ and $\mathcal{A}_{\vec{b}} = (\mathcal{A}, P, \vec{b})$. Note that the universe of both structures is the same, A, and in the game the spoiler selects points in A, and the duplicator select bijections $f : A \to A$.

Define a binary relation H on $S_r^{\mathcal{A}}(\vec{a}'\vec{b}')$ by letting $(x,y) \in H$ iff $x = h(y)$ or $y = h(x)$. We will show that the duplicator can play in such a way that, if $\vec{x} = (x_1, \ldots, x_n)$ and $\vec{y} = (y_1, \ldots, y_n)$ are points played on $\mathcal{A}_{\vec{a}}$ and $\mathcal{A}_{\vec{b}}$ respectively after n rounds, then there exists a set $J \subseteq \{1, \ldots, n\}$ with the following properties. (1) If $j \in J$, then $(x_j, y_j) \in H$. (2) If $j \notin J$, then $x_j = y_j$. (3) $\vec{a}'\vec{x}^J \approx_0^{\mathcal{A}} \vec{b}'\vec{y}^J$, where $\vec{x}^J$ is the subtuple of $\vec{x}$ that consists of the component of $\vec{x}$ whose indices are in J, and likewise for $\vec{y}^J$. (4) $d_{\mathcal{A}}(\vec{a}'\vec{x}^J, \vec{x}^{\overline{J}}) > 1$, and $d_{\mathcal{A}}(\vec{b}'\vec{y}^J, \vec{x}^{\overline{J}}) > 1$, where $d_{\mathcal{A}}$ is the distance in $\mathcal{G}(\mathcal{A})$, the Gaifman graph of $\mathcal{A}$, and $\vec{x}^{\overline{J}}$ consists of the components of $\vec{x}$ whose indices are not in J.

We first prove that this suffices to show that the duplicator wins. For this we need to establish $\vec{a}'\vec{c}\vec{x} \approx_0^{\mathcal{A}} \vec{b}'\vec{c}\vec{y}$, and furthermore, show that the mapping F induced by these two tuples preserves P. The latter is clear though as for any $v = F(u)$, either $u = v$ or $(u,v) \in H$, by construction, and thus P is preserved. To see that $\vec{a}'\vec{c}\vec{x} \approx_0^{\mathcal{A}} \vec{b}'\vec{c}\vec{y}$, notice that $\vec{a}'\vec{x}^J \approx_0 \vec{b}'\vec{y}^J$ by (3), and by (4) and the definition of $\vec{c}$, $d_{\mathcal{A}}(\vec{a}'\vec{x}^J, \vec{c}\vec{x}^{\overline{J}}) > 1$, and $d_{\mathcal{A}}(\vec{b}'\vec{y}^J, \vec{c}\vec{x}^{\overline{J}}) > 1$. Thus no σ-relation can have a tuple containing an element of $\vec{a}'\vec{x}^J$ and an element of $\vec{c}\vec{x}^{\overline{J}}$, or an element of $\vec{b}'\vec{y}^J$ and an element of $\vec{c}\vec{x}^{\overline{J}}$. This suffices to conclude that $\vec{a}'\vec{c}\vec{x} \approx_0^{\mathcal{A}} \vec{b}'\vec{c}\vec{y}$, and thus the duplicator wins the n-round game, provided (1)-(4) hold.

To prove that the duplicator can play as required, we show, by induction on the number of moves, that the duplicator can maintain these four conditions. The play is somewhat similar to the one used in [32] for ordinary (not bijective) games. We shall classify all moves into two types, *type 1* moves and *type 2* moves. We use the following notation. Let $\vec{x} = (x_1, \ldots, x_i)$ and $\vec{y} = (y_1, \ldots, y_i)$ be points played on $\mathcal{A}_{\vec{a}}$ and $\mathcal{A}_{\vec{b}}$ after i rounds. That is, in the jth round, $j \le i$, x_j is played in $\mathcal{A}_{\vec{a}}$, and $y_j = f_j(x_j)$ is the duplicator's response, where f_j is the bijection chosen by the duplicator for this round. By $\vec{x}^{\mathbf{1}}$ and $\vec{y}^{\mathbf{1}}$ we denote subtuples consisting of points played in type 1 moves, and by $\vec{x}^{\mathbf{2}}$ and $\vec{y}^{\mathbf{2}}$ we denote subtuples of points played in type 2 moves.

Let $d_i = 2^{m-i}$. The first two conditions are as follows.

(1) If jth move is a type 2 move, then $x_j = y_j$; that is, $\vec{x}^{\mathbf{2}} = \vec{y}^{\mathbf{2}}$.
(2) if jth move is a type 1 move, then $S^{\mathcal{A}}_{d_i}(x_j, y_j) \subseteq S^{\mathcal{A}}_r(\vec{a}'\vec{b}')$, and $(x_j, y_j) \in H$.

Suppose that these conditions hold, and j is a type 1 move. Since $S^{\mathcal{A}}_r(\vec{a}) \cap S^{\mathcal{A}}_r(\vec{b}) = \emptyset$, if $x_j \in S^{\mathcal{A}}_r(\vec{a}')$, then $y_j \in S^{\mathcal{A}}_r(\vec{b}')$, and vice versa. We use the notation $\vec{x}^{\mathbf{1}}_a$ for the subtuple of $\vec{x}^{\mathbf{1}}$ whose components are in $S^{\mathcal{A}}_r(\vec{a}')$, and $\vec{x}^{\mathbf{1}}_b$ for the remaining components, that is, those in $S^{\mathcal{A}}_r(\vec{b}')$. We similarly define $\vec{y}^{\mathbf{1}}_a$ and $\vec{y}^{\mathbf{1}}_b$. Notice that $\vec{y}^{\mathbf{1}}_b = h(\vec{x}^{\mathbf{1}}_a)$ and $\vec{x}^{\mathbf{1}}_b = h(\vec{y}^{\mathbf{1}}_a)$.

We can now formulate the next two requirements:

3. $h : N^{\mathcal{A}}_{d_i}(\vec{a}'\vec{x}^{\mathbf{1}}_a\vec{y}^{\mathbf{1}}_a) \cong N^{\mathcal{A}}_{d_i}(\vec{b}'\vec{y}^{\mathbf{1}}_b\vec{x}^{\mathbf{1}}_b)$ (that is, h is an isomorphism between these neighborhoods).
4. $d_{\mathcal{A}}(\vec{a}'\vec{x}^{\mathbf{1}}_a\vec{y}^{\mathbf{1}}_a, \vec{x}^{\mathbf{2}}) > d_i$ and $d_{\mathcal{A}}(\vec{b}'\vec{x}^{\mathbf{1}}_b\vec{y}^{\mathbf{1}}_b, \vec{x}^{\mathbf{2}}) > d_i$.

Proving these 4 conditions indeed suffices to conclude that the duplicator wins after n rounds, as then the conditions (1)-(4) are easily verified: we take J to be the set of type 1 moves. Note that permuting indices of moves has no effect on whether the resulting map is a partial isomorphism; thus we normally put subtuples of $\vec{x}$ and $\vec{y}$ in an order that keeps notation simple.

We prove, by induction on i, that the duplicator can maintain conditions 1–4. For the first move, the duplicator's bijection f is taken to be $f(x) = h(x)$ if $x \in S^{\mathcal{A}}_{d_1}(\vec{a})$, $f(x) = h^{-1}(x)$ if $x \in S^{\mathcal{A}}_{d_1}(\vec{b})$, and $f(x) = x$ in other cases. If the spoiler's move is in $S^{\mathcal{A}}_{d_1}(\vec{a}, \vec{b})$, then it is a type 1 move, otherwise it is a type 2 move. It is routine to verify that conditions 1-4 are satisfied.

For the inductive step, assume that i rounds have been played, and conditions 1–4 are satisfied. As $d_i = 2d_{i+1}$, we define f_{i+1}, duplicator's bijection for the round $i+1$, as follows:

$$f_{i+1}(x) \quad = \quad \begin{cases} h(x) & \text{if } x \in S^{\mathcal{A}}_{d_{i+1}}(\vec{a}'\vec{x}^{\mathbf{1}}_a\vec{y}^{\mathbf{1}}_a) \\ h^{-1}(x) & \text{if } x \in S^{\mathcal{A}}_{d_{i+1}}(\vec{b}'\vec{y}^{\mathbf{1}}_b\vec{x}^{\mathbf{1}}_b) \\ x & \text{otherwise.} \end{cases}$$

If a move is made according to the first or the second clause (that is, if the spoiler plays in $S^{\mathcal{A}}_{d_{i+1}}(\vec{a}'\vec{b}'\vec{x}^{\mathbf{1}}\vec{y}^{\mathbf{1}})$), then the move is a type 1 move; otherwise it is a type 2 move. Let x_{i+1}, y_{i+1} be the $i+1$st move. It remains to verify conditions 1–4.

Condition 1 is obvious, and so is the second part of condition 2 ($(x_{i+1}, y_{i+1}) \in H$). If $x_{i+1} \in S^{\mathcal{A}}_{d_{i+1}}(\vec{a}'\vec{x}^{\mathbf{1}}_a\vec{y}^{\mathbf{1}}_a)$, then $S^{\mathcal{A}}_{d_{i+1}}(x_{i+1}) \subseteq S^{\mathcal{A}}_{d_i}(\vec{a}'\vec{x}^{\mathbf{1}}_a\vec{y}^{\mathbf{1}}_a)$, as $d_i = 2d_{i+1}$. Hence $S^{\mathcal{A}}_{d_{i+1}}(x_{i+1}) \subseteq S^{\mathcal{A}}_r(\vec{a}')$, and $S^{\mathcal{A}}_{d_{i+1}}(y_{i+1}) \subseteq S^{\mathcal{A}}_r(\vec{b}')$, proving the second part of condition 2. The proof for the case $x_{i+1} \in S^{\mathcal{A}}_{d_{i+1}}(\vec{b}'\vec{y}^{\mathbf{1}}_b\vec{x}^{\mathbf{1}}_b)$ is similar.

For a type 2 move, conditions 3 and 4 follow immediately from the definition of f_{i+1} and the hypothesis. Assume that $x_{i+1} \in S^{\mathcal{A}}_{d_{i+1}}(\vec{a}'\vec{x}^{\mathbf{1}}_a\vec{y}^{\mathbf{1}}_a)$ (the case of $x_{i+1} \in S^{\mathcal{A}}_{d_{i+1}}(\vec{b}'\vec{y}^{\mathbf{1}}_b\vec{x}^{\mathbf{1}}_b)$ is symmetric). As $h : N^{\mathcal{A}}_{2d_{i+1}}(\vec{a}'\vec{x}^{\mathbf{1}}_a\vec{y}^{\mathbf{1}}_a) \to N^{\mathcal{A}}_{2d_{i+1}}(\vec{b}'\vec{y}^{\mathbf{1}}_b\vec{x}^{\mathbf{1}}_b)$ is an isomorphism, we obtain that h maps $N^{\mathcal{A}}_{d_{i+1}}(\vec{a}'\vec{x}^{\mathbf{1}}_a\vec{y}^{\mathbf{1}}_a x_{i+1})$ isomorphically onto $N^{\mathcal{A}}_{2d_{i+1}}(\vec{b}'\vec{y}^{\mathbf{1}}_b\vec{x}^{\mathbf{1}}_b y_{i+1})$, and thus condition 3 holds. Furthermore, any component of $\vec{x}^{\mathbf{2}}$ is outside of $S^{\mathcal{A}}_{2d_{i+1}}(\vec{a}'\vec{x}^{\mathbf{1}}_a\vec{y}^{\mathbf{1}}_a)$ by the hypothesis, and hence outside $S^{\mathcal{A}}_{d_{i+1}}(x_{i+1})$. Thus, $d_{\mathcal{A}}(x_{i+1}, \vec{x}^{\mathbf{2}}) > d_{i+1}$. As $y_{i+1} = h(x_{i+1}) \in S^{\mathcal{A}}_{d_{i+1}}(\vec{b}'\vec{y}^{\mathbf{1}}_b\vec{x}^{\mathbf{1}}_b)$, and every element of $\vec{x}^{\mathbf{2}}$ is at a distance exceeding $2d_{i+1}$ from $\vec{b}'\vec{y}^{\mathbf{1}}_b\vec{x}^{\mathbf{1}}_b$, we obtain $d_{\mathcal{A}}(y_{i+1}, \vec{x}^{\mathbf{2}}) > d_{i+1}$, thus proving condition 4. This completes the proof that the duplicator can play to maintain conditions 1–4. The lemma is proved. QED

We now put these two lemmas together to show

Theorem 5 *Let g be nondecreasing and not bounded by a constant, and let Q be an m-ary query in $(\mathcal{L}^*_{\infty\omega}(\mathbf{C})+ <_{\restriction g})_w$. Then Q is g-weakly semi-local.*

Proof: Let Q be definable by $\varphi(x_1, \ldots, x_m)$, where φ is a $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$ formula in the language of σ and an extra symbol S for the auxiliary preorder. Let $\mathcal{A}$ be a σ-structure, with $\vec{a}, \vec{b} \in A^m$ and $\vec{a} \leftrightsquigarrow^{\mathcal{A}}_{g,2^n} \vec{b}$, where $n = \mathsf{rk}(\varphi)$. Assume that φ is $<_{\restriction g}$-invariant on $\mathcal{A}$. Let P_0 be a preorder on A, such that $P_0 \in <_{\restriction g}$. Let $\vec{a} \in Q(\mathcal{A}) = \varphi[(\mathcal{A}, P_0)]$. Choose P to be the preorder given by Lemma 3. Due to the invariance of φ, $\vec{a} \in \varphi[(\mathcal{A}, P)]$; that is, $(\mathcal{A}, P) \models \varphi(\vec{a})$. By Lemmas 3 and 2, $(\mathcal{A}, P) \models \varphi(\vec{b})$, and again by invariance $(\mathcal{A}, P_0) \models \varphi(\vec{b})$. Thus, $\vec{b} \in \varphi[(\mathcal{A}, P_0)] = Q(\mathcal{A})$. This proves g-weak semi-locality of Q. QED

If $g = id$, we obtain

Corollary 8 *Let $\precsim_2$ be the class of preorders in which every equivalence class has size at most 2. Then every query definable in $\mathcal{L}^*_{\infty\omega}(\mathbf{C})+ \precsim_2$ is weakly semi-local, and has the BNDP.* *QED*

Proof of Theorem 3 Let Q be in $(\mathcal{L}^*_{\infty\omega}(\mathbf{C})+ <_{\restriction g})_w$. By Theorem 5, it is g-weakly semi-local. By Corollary 7, it has the BNDP. QED

6 Conclusion

We have shown that queries definable in counting logics $\mathrm{FO}(\mathbf{C})$, $\mathrm{FO}(Q_{\mathbf{u}})$ and $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$, in the presence of relations from the class $<\restriction_g$ have the bounded number of degrees property. In other words, even extremely powerful counting logics in the presence of relations which are almost-everywhere linear orders have a very tame behavior. The situation changes drastically when $<\restriction_g$ is replaced by a linear order. For example, $\mathcal{L}^*_{\infty\omega}(\mathbf{C})+<$ expresses every query on ordered structures. A similar phenomenon is observed for other logics, most notably, $\mathrm{FO}(\mathbf{C})$ which captures uniform TC^0 on ordered structures.

The techniques of this paper cannot be straightforwardly extended to prove separation results in the ordered case. The logic $\mathcal{L}^*_{\infty\omega}(\mathbf{C})$ is very powerful, as it expresses every property of natural numbers, and all other known counting extensions of FO can be embedded into it. We also relied on bijective games to prove the main result. However, bijective games characterize expressiveness of a logic which defines all queries on ordered finite structures. Thus, in the ordered case one cannot use the generic techniques from [12,21,22,26] that apply to a variety of counting logics.

It was shown in [8] that if there is a proof of inexpressibility of some property in $\mathrm{FO}(\mathbf{C})+<$, then there must be a proof of that based on the counting games of [17]. The counting game is weaker than the bijective game; on the other hand, it does not have the inherent limitations of the latter in the ordered case. Thus, a possible way of proving a separation result may be to modify the locality techniques to work with the counting, rather than bijective, games.

Another approach would be to modify the ordered conjecture of [20] to include counting. Namely, such a modified conjecture would say that there is no unbounded class of ordered structures on which $\mathrm{FO}(\mathbf{C})$ captures polynomial time. One reason to consider this is that there are strong indications that for FO the conjecture holds [20]. With counting, however, one has to be careful: by considering the class of linear orders and adding unary quantifiers which test for polynomial time properties of cardinalities, one obtains a counting logic for which the conjecture fails. However, $\mathrm{FO}(\mathbf{C})$ has rather limited arithmetic, and perhaps an attempt to understand why it fails to capture polynomial time on various classes of structures may lead to a better understanding of its structural properties which are not shared by other counting logics.

Acknowledgements We thank Martin Grohe, Lauri Hella, Moshe Vardi, and Mihalis Yannakakis for their comments on earlier drafts of this paper.

References

[1] E. Allender. Circuit complexity before the dawn of the new millennium. In *Proc. 16th Conf. on Foundations of Software Technology and Theoretical Computer Science*, Springer LNCS vol. 1180, 1996, 1–18.

[2] D.A. Barrington, N. Immerman, H. Straubing. On uniformity within NC^1. *JCSS*, 41:274–306,1990.

[3] M. Benedikt, H.J. Keisler. Expressive power of unary counters. *Proc. Int. Conf. on Database Theory (ICDT'97)*, Springer LNCS 1186, January 1997.

[4] J. Cai, M. Fürer and N. Immerman. On optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12 (1992), 389–410.

[5] G. Dong, L. Libkin and L. Wong. Local properties of query languages. *TCS*, 239 (2000), 277–308.

[6] A. Durand, C. Lautemann and M. More. Counting results in weak formalisms. Manuscript, 1998.

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

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

[9] R. Fagin, L. Stockmeyer and M. Vardi, On monadic NP vs monadic co-NP, *Information and Computation*, 120 (1995), 78–92.

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

[11] M. Grohe and T. Schwentick. Locality of order-invariant first-order formulas. In *MFCS'98*, pages 437–445.

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

[13] L. Hella, L. Libkin and J. Nurmonen. Notions of locality and their logical characterizations over finite models. *Journal of Symbolic Logic*, 64 (1999), 1751–1773.

[14] L. Hella, L. Libkin, J. Nurmonen and L. Wong. Logics with aggregate operators. In *LICS'99*, pages 35–44. Full version to appear in *J. ACM.*

[15] N. Immerman. Relational queries computable in polynomial time. *Information and Control*, 68 (1986), 86–104.

[16] N. Immerman. Languages that capture complexity classes. *SIAM Journal of Computing* 16 (1987), 760–778.

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

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

[19] Ph. Kolaitis, M. Vardi. Infinitary logic and 0-1 laws. *Inf.& Comput.*, 98 (1992), 258–294.

[20] Ph. Kolaitis and M. Vardi. Fixpoint logic vs. infinitary logic in finite-model theory. In *LICS'92*, pages 46–57.

[21] L. Libkin. On the forms of locality over finite models. In *Proc. 12th IEEE Symp. on Logic in Computer Science (LICS'97)*, Warsaw, Poland, June–July 1996, pages 204–215.

[22] L. Libkin. On counting logics and local properties. In *Proc. 12th IEEE Symp. on Logic in Computer Science (LICS'97)*, Warsaw, Poland, June–July 1996, pages 204–215.

[23] L. Libkin. Logics with counting, auxiliary relations, and lower bounds for invariant queries. In *LICS'99*, 316–325.

[24] L. Libkin, L. Wong. Query languages for bags and aggregate functions. *JCSS* 55 (1997), 241–272.

[25] L. Libkin and L. Wong. Unary quantifiers, transitive closure, and relations of large degree. In *Proc. 15th Symp. on Theoretical Aspects of Computer Science (STACS'98)*, Springer Verlag, 1998, to appear.

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

[27] M. Otto. *Bounded Variable Logics and Counting: A Study in Finite Models.* Springer Verlag, 1997.

[28] M. Otto. ϵ-Logic is more expressive than first-order on finite structures. *Journal of Symbolic Logic*, to appear.

[29] I. Parberry and G. Schnitger. Parallel computation and threshold functions. *JCSS* 36 (1988), 278–302.

[30] A. Razborov and S. Rudich. Natural proofs. *Journal of Computer and System Sciences* 55 (1997), 24–35.

[31] T. Schwentick. Graph connectivity, monadic NP and built-in relations of moderate degree. In *Proc. 22nd International Colloq. on Automata, Languages, and Programming*, Springer LNCS 944, 1995, pages 405–416.

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

[33] M. Vardi. The complexity of relational query languages. In *Proceedings, 14th ACM Symposium on Theory of Computing*, 1982, 137–146.