Unary Quantifiers, Transitive Closure, and Relations of Large Degree*

Leonid Libkin
Bell Laboratories
600 Mountain Avenue
Murray Hill, NJ 07974, USA
Email: libkin@bell-labs.com

Limsoon Wong
BioInformatics Center &
Institute of Systems Science
Singapore 119597
Email: limsoon@iss.nus.sg

Abstract

This paper studies expressivity bounds for extensions of first-order logic with counting and unary quantifiers in the presence of relations of large degree. There are several motivations for this work. First, it is known that first-order logic with counting quantifiers captures uniform TC^0 over *ordered* structures. Thus, proving expressivity bounds for first-order with counting can be seen as an attempt to show $\mathrm{TC}^0 \subsetneq \mathrm{DLOG}$ using techniques of descriptive complexity. Second, the presence of auxiliary built-in relations (e.g., order, successor) is known to make a big impact on expressivity results in finite-model theory and database theory (where logics with counting and unary quantifiers have recently been used to model query languages with aggregation). For those logics, our goal is to extend techniques from "pure" setting to that of auxiliary relations.

Until now, all known results on the limitations of expressive power of the counting and unary-quantifier extensions of first-order logic dealt with auxiliary relations of "small" degree. For example, it is known that these logics fail to express some DLOG-queries in the presence of a sucessor relation. Our main result is that these extensions cannot define the deterministic transitive closure (a DLOG-complete problem) in the presence of auxiliary relations of "large" degree, in particular, those which are "almost linear orders." They are obtained from linear orders by replacing them by "very thin" preorders on arbitrarily small number of elements. We also show that the technique of the proof (in a precise sense) cannot be extended to provide the proof of separation of TC^0 from DLOG. We also discuss a general impact of having built-in (pre)orders, and give some expressivity statements in the pure setting that would imply separation results for the ordered case. We also briefly discuss database applications.

Keywords: descriptive complexity, TC^0, logarithmic space, counting quantifiers, unary quantifiers, orders, expressive power.

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 logic $\mathcal{L}$ over the 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, more generally, how they allows us to derive new expressivity bounds in the presence of complex auxiliary relations.

*Part of this work was done when the first author was visiting Institute of Systems Science.

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, cf. [25]. Despite serious efforts and a number of proved lower bounds (see [2] for a survey), it is still not known if $\mathrm{TC}^0 \subsetneq \mathrm{NP}$, and the results of [26] show that traditional approaches to lower bounds are unlikely to succeed in proving this separation.

A starting point for our study is a result by Barrington, Immerman and Straubing [4] stating that:

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

Here, as usual, TC^0 is the class of problems solvable by polynomial-size, constant-depth threshold circuits, and uniform means DLOGTIME-uniform, see [4] for more details. From now on, whenever we write TC^0, we mean the uniform class.

By $\mathrm{FO}(\mathbf{C})$ we mean the extension of first-order logic with counting quantifiers $\exists i$, where $\exists ix.\varphi(x)$ means that φ has at least i satisfiers. We shall give a full definition later, and at this point offer an example: $\exists i, j((j + j = i) \wedge \exists! ix.\varphi(x))$ (where $\exists! i$ is a shorthand for "exists exactly i") states that the number of satisfiers of φ is even — this is known not to be expressible in first-order logic alone. By $\mathrm{FO}(\mathbf{C})+ <$ we mean $\mathrm{FO}(\mathbf{C})$ in the presence of a built-in order relation. Note that if we are interested in $\mathrm{FO}(\mathbf{C})+ <$ sentences, then it does not matter which linear order is used. However, it is known that the mere presence of an order relation increases expressiveness (cf. [3]), and this will be discussed later.

This result reduces the problem of separation of uniform TC^0 from classes such as DLOG, NLOG, P, etc, 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 these expressivity bounds. The first partial result was given by Etessami, using counting games of [17]:

Fact 1 ([11]) *There exist a problem complete for* DLOG *under first-order reductions that cannot be defined by* $\mathrm{FO}(\mathbf{C})$ *in the presence of a successor relation.* □

Etessami's result also shows that *dtc*, deterministic transitive closure, is not in $\mathrm{FO}(\mathbf{C}) + \mathit{succ}$, while $\mathrm{FO} + \mathit{dtc} + \mathit{succ}$ captures the class DLOG. This was extended in [19] as follows.

Fact 2 ([19]) *Deterministic transitive closure 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. (A more general definition can be given for arbitrary relational structures, cf. [8].) In the successor relation, every node has in- and out-degree either 0 or 1, so this does generalize Fact 1. In contrast to these two results, 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 [19], using a definition on *moderate degree* by Fagin, Stockemeyer and Vardi [12]. We say that a class $\mathcal{C}$ of graphs (more generally, relational structures) is of moderate degree, if $\mathit{degmax}_{\mathcal{C}}(n)$, 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 $\mathit{degmax}_{\mathcal{C}}(n) \leq \log^{\delta(n)} n$.

Fact 3 ([19]) *Deterministic transitive closure cannot be defined by* FO(**C**) *in the presence of auxiliary relations of moderate degree.* □

In [12], 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 $n^{o(1)}$ [28] and to a linear order [27]. So one may wonder if a similar program can be carried out for FO(**C**).

There is a significant difference between Facts 1, 2 and 3, and the desired separation for the ordered case: in those Facts, we only deal with auxiliary relations of *small* degrees – these are either constant, or very small compared to the size of the input structure. In contrast, a linear order realizes as many degrees as there are elements in the input. Hence, one needs techniques to lift the results for FO(**C**) from relations of *small* degrees to relations of *large* degrees, i.e. those comaparable with the size of the input.

Before outlining the results of the paper, let us mention two more motivations for our study. First, the effect of a built-in order relation has been studied extensively in finite-model theory. Many expressivity bounds proved for logics without a built-in order are known not to carry over to the ordered setting. For example, extensions of first-order logic with least, inflationary and partial fixpoint operators are known to possess a 0-1 law [18]; hence, they cannot express the parity test ("is the cardinality of the input even?"). However, with a built-in order relation, these logics capture complexity classes PTIME and PSPACE, and of course can define parity. It is also known that there exist order-invariant queries that one can express in FO+ $<$, but not in FO (an example, due to Y. Gurevich, can be found in [3]). And in fact extending the expressivity bounds to the ordered setting is generally a big open problem in finite-model theory. Thus, we would like to know to what extent we can recover known expressivity bounds if complex auxiliary relations (like a linear order) are present.

Another motivation is the study of relational database query languages with aggregate functions, cf. [7, 8, 20]. It is also closely connected with the study of power of languages that allow correct manipulation of duplicate records in databases, cf. [14]. These languages adequately model features present in commercial languages (like SQL) but often ignored in database theory. It was shown recently [19, 21] that (a theoretical reconstruction of) such languages can be simulated by FO(**C**). This, and known expressivity bounds for FO(**C**) [11, 22, 23, 19] answer some of the open problems [7, 14] about expressivity of such languages without order relation on elements of base types. However, typically base types, such as *int* and *string*, come equipped with an ordering, or at least a preorder (for example, elements of type *real* can be only preordered if some digits are meaningless). Thus, extending techniques for proving expressivity bounds to the setting where (pre)orders on base types are allowed will give us new insight into expressiveness of practical database languages.

Organization In Section 2 we introduce the notation, and in Section 3 we describe the technical machinery based on local properties of logics.

In Section 4 we describe a general approach to proving expressivity bounds for local logics in the presence of auxiliary relations. This approach is based on establishing two properties: $\mathbf{Def}_{\mathcal{L}}[\mathcal{R}, \mathcal{C}]$, which essentially says that a class of relations $\mathcal{R}$ is definable on structures from $\mathcal{C}$ by $\mathcal{L}$-formulae, and $\mathbf{Sep}_{\mathcal{L}}[\varphi, \mathcal{C}]$ essentially saying that a $\mathcal{L}$-formula φ can express properties on $\mathcal{C}$ that are not "local." We show that for logics like FO(**C**) these two imply inexpressibility of φ in $\mathcal{L}$ in the presence of relations from $\mathcal{R}$.

We then define the class of "almost linear orders" (shown in figure 1) and use the above technique to show that deterministic transitive closure (and thus ofther DLOG-complete problems) are not expressible in $\mathrm{FO}(\mathbf{C})$ in the presence of those relations. We also show that this holds for some other logics as well.

In Section 5, we show that, in a precise sense, this is the best partial result that can be obtained using locality techniques. That is, any structure "between" our almost linear orders and real linears provides the same gain in expressiveness as linear orders, and an attempt to extend techniques of Section 4 to those fails as the conditions $\mathbf{Def}_{\mathcal{L}}[\mathcal{R}, \mathcal{C}]$ and $\mathbf{Sep}_{\mathcal{L}}[\varphi, \mathcal{C}]$ become inconsistent.

In Section 6, we analyze expressivity of $\mathrm{FO}(\mathbf{C})$ in the pure case (without auxiliary relations) vs. built-in orders or preorders. We also describe problems whose inexpressibility in $\mathrm{FO}(\mathbf{C})$ (note the absence of an order relation!) would imply $\mathrm{TC}^0 \subsetneq \mathrm{DLOG}(\mathrm{NLOG})$. Concluding remarks are given in Section 7. No proofs are given in this extended abstract; they can found in the full version.

2 Notations

A relational signature σ is a set of relation symbols $\{R_1, ..., R_l\}$, with an associated arity function. In what follows, $p_i (> 0)$ denotes the arity of R_i. We write σ_n for σ extended with n new constant symbols. We use σ_{gr} for the signature of graphs (that is, one binary predicate E). 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. If $\mathcal{A}$ is understood, we will omit the superscript. The class of finite σ-structures is denoted by $\mathrm{STRUCT}[\sigma]$. Isomorphism of structures is denoted by $\cong$. The carrier of a structure $\mathcal{A}$ is always denoted by A.

We mostly deal with three logics: $\mathrm{FO}, \mathrm{FO}(\mathbf{C})$ and $\mathrm{FO}(\mathbf{Q}_u)$, the last one being first-order logic with unary quantifiers. First-order formulae are built-up from atomic formulae by using Boolean connectives and quantifiers $\exists$ and $\forall$. First-order logic with counting, $\mathrm{FO}(\mathbf{C})$, is defined as a two sorted logic, with second sort being the sort of natural numbers. More precisely, in this approach a structure $\mathcal{A}$ is of the form

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

Here the relations $R_i^{\mathcal{A}}$ are defined on the domain $\{v_1, \ldots, v_n\}$, while on the numerical domain $\{1, \ldots, n\}$ one has $1, \max, <$ and the BIT predicate available ($\mathrm{BIT}(i, j)$ iff the ith bit in the binary representation of j is one). It also has counting quantifiers $\exists ix.\varphi(x)$, meaning that φ has at least i satisfiers; here i refers to the numerical domain and x to the domain $\{v_1, \ldots, v_n\}$. These quantifiers bind x but not i.

Let σ_k^{unary} be a signature of k unary symbols, and let $\mathcal{K}$ be a class of σ_k^{unary}-structures which is closed under isomorphisms. Then $\mathrm{FO}(Q_{\mathcal{K}})$ extends the set of formulae of FO with the following additional rule: if $\psi_1(x_1, \vec{y}_1), \ldots, \psi_k(x_k, \vec{y}_k)$ are formulae, then $Q_{\mathcal{K}} x_1 \ldots x_k.(\psi_1(x_1, \vec{y}_1), \ldots, \psi_k(x_k, \vec{y}_k))$ is a formula. Here $Q_{\mathcal{K}}$ binds x_i in the ith formula, for each $i = 1, \ldots, k$. The semantics of $Q_{\mathcal{K}}$ is defined as follows:

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

In this definition, $\vec{a}_i$ is a tuple of parameters that gives the interpretation for those free variables of $\psi_i(x_i, \vec{y}_i)$ which are not equal to x_i. Examples of unary quantifiers include the usual $\exists$ and $\forall$, as well as Rescher (bigger cardinality) and Härtig (equicardinality) quantifiers. We use the notation $\mathrm{FO}(\mathbf{Q}_u)$ for FO extended with *all* unary quantifiers.

Every FO($\mathbf{C}$) sentence can be expressed in FO($\mathbf{Q}_u$), while there exist properties definable in FO($\mathbf{Q}_u$) but not in FO($\mathbf{C}$).

Most of the results in this paper refer to these three logics. They can also be extended to abstract logics in the sense of [9], which are regular (e.g., closed under first-order operations and substititions).

With each formula $\psi(x_1, \ldots, x_m)$ in the logical language whose symbols are in σ, we associate a *query* (semantic mapping) that maps a σ-structure $\mathcal{A}$ into a m-ary relation $\psi[\mathcal{A}] = \langle A, \{(a_1, \ldots, a_m) \in A^m | \ \mathcal{A} \models \psi(a_1, \ldots, a_m)\}\rangle$.

Given a relational signature σ and a class $\mathcal{R}$ of σ'-structures, where σ' is another relational signature, disjoint from σ, we say that a query q, producing an m-ary relation, is definable on σ-structures in the presence of $\mathcal{R}$-structures if there exists a $\sigma \cup \sigma'$-formula $\varphi(\vec{x})$ such that, for any σ-structure $\mathcal{A}$ with carrier A and for any structure $\mathcal{A}' \in \mathcal{R}$ on A, we have:

$$q(\mathcal{A}) = \{\vec{a} \in A^m \mid (\mathcal{A}, \mathcal{A}') \models \varphi(\vec{a})\}$$

where $(\mathcal{A}, \mathcal{A}')$ is the $\sigma \cup \sigma'$ structure obtained by putting $\mathcal{A}$ and $\mathcal{A}'$ together. We most often encounter the situation where $\mathcal{R}$ is the class of preorders (with special properties), or linear orders. Note that, according to this definition, a query definable in the presence of $\mathcal{R}$-structures is independent of a particular $\mathcal{R}$ structure being used. We use the notation $\mathcal{L} + \mathcal{R}$ for the class of queries definable in $\mathcal{L}$ the presence of relations from $\mathcal{R}$.

We shall use $\mathcal{O}_k$ for the class of preorders in which no equivalence class has more than k elements; these can be viewed as being very close to linear orders for small k. We also call them preorders of width k. In particular, $\mathcal{O}_1$ is the class of linear orders. We also write $\mathcal{L}+ <$ instead of $\mathcal{L} + \mathcal{O}_1$ for the class of queries definable in $\mathcal{L}$ in the presence of built-in order relation.

3 Local queries over finite models

A number of notions of locality have been introduced in finite-model theory in order to prove inexpressibility results, cf. [10, 13, 12, 8, 19]. Here we describe one of these notions, which will serve as a main technical tool.

Given a structure $\mathcal{A}$, its *Gaifman graph* [10, 13] $\mathcal{G}(\mathcal{A})$ is defined as $\langle A, E\rangle$ where (a, b) is in E iff there is a tuple $\vec{t} \in R_i^{\mathcal{A}}$ for some i such that both a and b are in $\vec{t}$. For example, if $\mathcal{A}$ is a graph itself, then $\mathcal{G}(\mathcal{A})$ is its reflexive-symmetric closure. 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 $d(a, a) = 0$. Given $a \in A$, its *r-sphere* $S_r^{\mathcal{A}}(a)$ is $\{b \in A \mid d(a, b) \le r\}$. For a tuple $\vec{t}$, define $S_r^{\mathcal{A}}(\vec{t})$ as $\bigcup_{a \in \vec{t}} S_r^{\mathcal{A}}(a)$.

Given a tuple $\vec{t} = (t_1, \ldots, t_n)$, its *$r$-neighborhood* $N_r^{\mathcal{A}}(\vec{t})$ is defined as a σ_n structure

$$\langle S_r^{\mathcal{A}}(\vec{t}), R_1^{\mathcal{A}} \cap S_r^{\mathcal{A}}(\vec{t})^{p_1}, \ldots, R_k^{\mathcal{A}} \cap S_r^{\mathcal{A}}(\vec{t})^{p_k}, t_1, \ldots, t_n\rangle$$

That is, the carrier of $N_r^{\mathcal{A}}(\vec{t})$ is $S_r^{\mathcal{A}}(\vec{t})$, the interpretation of the σ-relations is obtained by restricting them from $\mathcal{A}$ to the carrier, and the n extra constants are the elements of $\vec{t}$. If the structure $\mathcal{A}$ is understood, we shall write $S_r(\vec{t})$ and $N_r(\vec{t})$.

We use the notation $\vec{a} \approx_r^{\mathcal{A}} \vec{b}$, or $\vec{a} \approx_r \vec{b}$ if $\mathcal{A}$ is understood, if $N_r^{\mathcal{A}}(\vec{a})$ and $N_r^{\mathcal{A}}(\vec{b})$ are isomorphic. Note that an isomorphism between these maps ith component of $\vec{a}$ onto ith component of $\vec{b}$.

A formula $\psi(x_1,\ldots,x_m)$ in a logic $\mathcal{L}$ is called *local* [8, 19] if there exists $r>0$ such that, for every $\mathcal{A}\in \mathrm{STRUCT}[\sigma]$ and for every two m-ary vectors $\vec{a}$, $\vec{b}$ of elements of A, $N_r(\vec{a}) \cong_r N_r(\vec{b})$ implies $\mathcal{A}\models\psi(\vec{a})$ iff $\mathcal{A}\models\psi(\vec{b})$. The minimum r for which this holds is called the *locality rank* of ψ, and is denoted by $\mathsf{lr}(\psi)$. Based on earlier results by Hella [15] and Nurmonen [22], the following was shown in [19]:

Fact 4 *Every* $\mathrm{FO}(\mathbf{C})$ *formula without free second-sort variables is local, and every* $\mathrm{FO}(\mathbf{Q}_u)$ *formula is local.* $\square$

Note that a direct application of locality is of no help if a built-in (pre)order relation is avilable, because then $d(a,b)\le 1$ for any two elements, and $N_r^{\mathcal{A}}(\vec{a})\cong N_r^{\mathcal{A}}(\vec{b})$, for $r>0$, is simply a restatement of $(\mathcal{A},\vec{a})\cong(\mathcal{A},\vec{b})$, which would imply $\mathcal{A}\models(\varphi(\vec{a})\leftrightarrow\varphi(\vec{b}))$ for any reasonable logic.

4 Expressivity bounds for $\mathrm{FO}(\mathbf{C})$ and $\mathrm{FO}(\mathbf{Q}_u)$ 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 $\mathrm{FO}(\mathbf{C})$ to prove our main result that DLOG-complete problems (in particular, 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 $\mathrm{DLOG}\not\subseteq \mathrm{FO}(\mathbf{C})+\mathcal{O}_k$ for any $k>1$.

Proving expressivity bounds in local logics

Let q be a query that takes structures from $\mathrm{STRUCT}[\sigma]$ as inputs and returns m-ary relations (e.g, transitive closure takes graphs from $\mathrm{STRUCT}[\sigma_{\mathrm{gr}}]$ as inputs and returns graphs). Let $\mathcal{R}$ be a class of relations, and $\mathcal{L}$ a logic. Suppose we want to prove that $q\notin\mathcal{L}+\mathcal{R}$. For that purpose, we introduce two conditions.

$\mathbf{Def}_{\mathcal{L}[\sigma]}[\mathcal{R},\mathcal{C}]$ Assume $\mathcal{C}\subseteq\mathrm{STRUCT}[\sigma]$. Then there exists a number n and an $\mathcal{L}$ formula φ in the volcabulary 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 $\mathcal{C}$.

$\mathbf{Sep}_{\mathcal{L}[\sigma]}[q,\mathcal{C}]$ 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, *or* $\mathrm{FO}(\mathbf{C})$, *or* $\mathrm{FO}(\mathbf{Q}_u)$. *Suppose for a given query* q *on* σ-*structures, one can find* $\mathcal{C}\subseteq\mathrm{STRUCT}[\sigma]$ *such that both* $\mathbf{Def}_{\mathcal{L}[\sigma]}[\mathcal{R},\mathcal{C}]$ *and* $\mathbf{Sep}_{\mathcal{L}[\sigma]}[q,\mathcal{C}]$ *hold. Then* $q\notin\mathcal{L}+\mathcal{R}$. $\square$

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

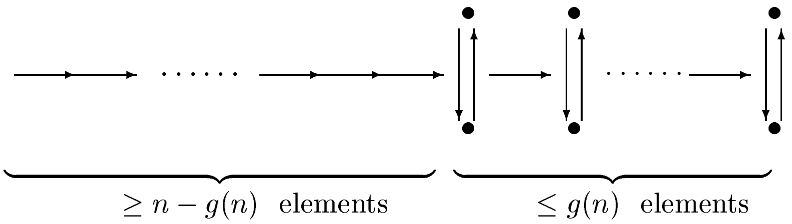

Figure 1: A relation from $\mathcal{P}_g$

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 $\mathcal{C}$.

Theorem 1 can also be extended to any local logic that is closed under first-order operations and allows a notion of substitution in a way that was used in the proof. All naturally occuring extensions of FO that are known to be local have these properties.

Lower bounds for (deterministic) transitive closure: FO(**C**) and FO($\mathbf{Q}_u$) with "thin" preorders

Recall [16] that deterministic transitive closure of a graph is obtained by closing its determistic 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 DLOG and FO + tc+ $<$ captures NLOG. Note that $<$ can be replaced by a successor relation, since its (deterministic) transitive closure is a linear order.

We next define the class of relations that we view as "almost linear orders." Let $g : \mathbb{N} \to \mathbb{R}$ be a nondecreasing function. Define $\mathcal{P}_g$ 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 $\mathcal{O}_2$, that is, a preorder where every equivalence class has at most two elements.
4. For any $b \in B$ and $c \in C$, $(b, c) \in R$.

See Figure 1 for a preorder from $\mathcal{P}_g$. Actually, we show the associated successor relation in the Figure. A relation from $\mathcal{P}_g$ is really the transitive closure of the one shown in Figure 1. Intuitively, if g is very small, then this is 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. In the result below, g can indeed be taken to be very small, for example, g could be $\log\log\ldots\log n$.

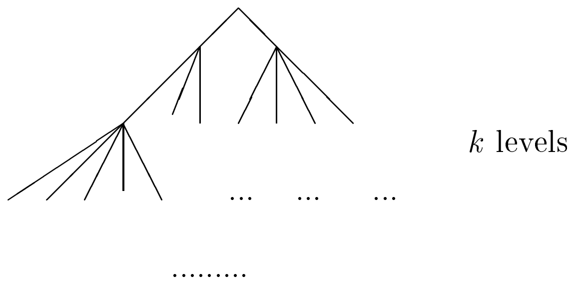

Figure 2: Canonical k-bushy tree

Theorem 2 *Let $g : \mathbb{N} \to \mathbb{R}$ be a nondcreasing function that is not bounded by a constant. Then (deterministic) transitive closure is not definable in* $\mathrm{FO}(\mathbf{C})$ *or* $\mathrm{FO}(\mathbf{Q}_u)$ *in the presence of relations from $\mathcal{P}_g$. In particular,* $\mathrm{DLOG} \not\subseteq \mathrm{FO}(\mathbf{C}) + \mathcal{P}_g$. $\square$

Corollary 1 *(Deterministic) transitive closure is not definable in* $\mathrm{FO}(\mathbf{C})$ *or* $\mathrm{FO}(\mathbf{Q}_u)$ *in the presence of relations from $\mathcal{O}_k$ for any $k > 1$. In particular,* $\mathrm{DLOG} \not\subseteq \mathrm{FO}(\mathbf{C}) + \mathcal{O}_k$. $\square$

This can be compared with the results of [6] where it was shown that first-order with fixpoint and counting fails to express some polynomial-time problems even in the presence of relations from $\mathcal{O}_4$ (of course first-order with fixpoint captures polynomial time in the presence of an order relation, cf. [10]).

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* $\mathrm{FO}(\mathbf{C})$ *or* $\mathrm{FO}(\mathbf{Q}_u)$. *Assume that $g : \mathbb{N} \to \mathbb{R}$ is a nondcreasing function that is not bounded by a constant. Then there exists a class $\mathcal{C}$ of graphs such that both $\mathbf{Def}_{\mathcal{L}[\sigma_{\mathrm{gr}}]}(\mathcal{P}_g, \mathcal{C})$ 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 $x \neq y$, $\textit{out-deg}(x) \neq \textit{out-deg}(y)$. 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, resp. 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 sketch 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. To other $s_k - 1$ nodes at the second level, we give $s_k+2, s_k+3, \ldots, s_k+(s_k-1) = 2s_k - 1$ children respectively. Now, to those nodes at the second level that do not belong to the two canonical k-bushy trees, we give $2s_k, 2s_k+1, \ldots$ children, as before, increasing the number by one. We continue this process until we fully fill the $k+1$st level. After that, 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 define $G_{d,k}$ by adding an arbitrary linear ordering on the leaves. 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 $(\textit{in-deg}(x), \textit{out-deg}(x)) \neq (\textit{in-deg}(y), \textit{out-deg}(y))$. Next, 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 $B^\circ_1 \cup B^\circ_2$. Then we let $x \prec y$ iff either $x \notin B^\circ, y \in B^\circ$, or $x, y \in B^\circ$ and $x \prec_0 y$, or $x, y \notin B^\circ$ and $x \prec_0 y$. This binary relation $\prec$ is definable in $\mathrm{FO}(\mathbf{C})$ and $\mathrm{FO}(\mathbf{Q}_u)$.

For a 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 depends only on k. Let $\mathcal{C}_g = \{G_{d,k} \mid d, k \in \mathbb{N}_+, d > d_k\}$. The rest of the proof is to verify that both $\mathbf{Def}_{\mathcal{L}[\sigma_{\mathrm{gr}}]}[\mathcal{P}_g, \mathcal{C}_g]$ and $\mathbf{Sep}_{\mathcal{L}[\sigma_{\mathrm{gr}}]}[tc, \mathcal{C}_g]$ hold. 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. □

5 Limitations of the technique

To summarize what has been achieved so far, we know that $\mathrm{FO}(\mathbf{C}) + \mathcal{O}_1 = \mathrm{TC}^0$, and the above results show that for any $k > 1$, $\mathrm{DLOG} \not\subseteq \mathrm{FO}(\mathbf{C}) + \mathcal{O}_k$. Furthermore, $\mathrm{DLOG} \not\subseteq \mathrm{FO}(\mathbf{C}) + \mathcal{P}_g$ 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 $\mathrm{FO}(\mathbf{C})+ <$.

The lemma below shows that removing the assumption that g is not bounded by a constant is essentially equivalent to having a linear order:

Lemma 1 *Let $g(n)$ be bounded by some constant M for all $n \in \mathbb{N}$. Assume that (deterministic) transitive closure is not in $\mathrm{FO}(\mathbf{C}) + \mathcal{P}_g$. Then it is not in $\mathrm{FO}(\mathbf{C})+ <$ either.*

Thus, 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 $\mathbf{Def}_{\mathcal{L}}[\mathcal{P}_M, \mathcal{C}]$ 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 $\mathrm{FO}(\mathbf{C})+ <$. Here we use $\mathcal{P}_M$ to denote $\mathcal{P}_g$ where $g(n) < M$ for all n.

If we were able to find such a class $\mathcal{C}$, it would show that $q \notin \mathrm{FO}(\mathbf{C})+ <$. Unfortunately, as the following theorem shows, no such class exists!

Theorem 3 *Let M be a constant, and let q be a query invariant under isomorphism. Let $\mathcal{L}$ be $\mathrm{FO}(\mathbf{C})$ or $\mathrm{FO}(\mathbf{Q}_u)$. Then there does not exist a class of structures $\mathcal{C}$ such that both $\mathbf{Def}_{\mathcal{L}}[\mathcal{P}_M, \mathcal{C}]$ and $\mathbf{Sep}_{\mathcal{L}}[q, \mathcal{C}]$ hold.* □

6 On the relative expressive power of auxiliary relations

In this section, we give a few comments about the murky area of expressivity with (pre)orders vs. expressivity without (pre)orders. As was mentioned before, it is known that $\mathrm{FO} \subsetneq \mathrm{FO}+ <$. Note that by

FO+ $<$ we mean the class of *order-independent* queries in FO+ $<$, so this is not a trivial oservation. A similar result can be shown for FO($\mathbf{C}$). In fact, it is implicit in [5].

Proposition 2 (Benedikt-Keisler) FO($\mathbf{C}$) $\subsetneq$ FO($\mathbf{C}$)+ $<$.

Proof skecth: Consider structures $\mathcal{A} = \langle A, R, U \rangle$ where R is binary and U is unary. The separating query q is the following: If R is an equivalence relation, does the number of distinct sizes of equivalence classes in R equal to the cardinality of U? It was shown to be inexperessible in FO($\mathbf{C}$) in [5]. However, we can show that it is definable in FO($\mathbf{C}$)+ $<$. $\Box$

However, this does not shed any light on what kind of examples might exist (if there are any!) in (FO($\mathbf{C}$)+ $<$) $-$ (FO($\mathbf{C}$) $+ \mathcal{O}_k$), $k > 1$:

Proposition 3 FO($\mathbf{C}$) $\subsetneq$ FO($\mathbf{C}$) $+ \mathcal{O}_k$. $\Box$

To find a separation from FO($\mathbf{C}$)+ $<$, we can only use a small class of preorders which, in a sense, do not have equivalence classes comparable to the size of the universe. Formally, let δ be a function from natural numbers to the reals, and let $\mathcal{O}_{\delta(n)}$ be the class of all preorder such the maximal equivalence class of a preorder on an n-element set is at most $n\delta(n)$. Let $\mathcal{O}_{o(n)}$ be the union of all $\mathcal{O}_{\delta(n)}$ where $\lim_{n\to\infty} \delta(n) = 0$. Then, modifying slightly an example shown to us by M. Otto [24], we can prove:

Proposition 4 FO($\mathbf{C}$) $+ \mathcal{O}_{o(n)} \subsetneq$ FO($\mathbf{C}$)+ $<$. $\Box$

While this gives us some partial results about expressivity of FO($\mathbf{C}$) with (pre)orders, it is still not clear how to prove bounds for FO($\mathbf{C}$)+ $<$ and ultimately for TC^0. We conclude by presenting a query whose inexpressibility in FO($\mathbf{C}$) (note the absence of order!) will imply bounds on TC^0.

Proposition 5 *a) If there is no* FO($\mathbf{C}$) *query that defines transitive closure on bushy trees, then* TC$^0 \subsetneq$ NLOG.
b) If there is no FO($\mathbf{C}$) *query that defines deterministic transitive closure on inverses of bushy trees, then* TC$^0 \subsetneq$ DLOG. $\Box$

While we do not know whether queries of Proposition 5 are definable in FO($\mathbf{C}$), we can give two partial results for *canonical* bushy trees.

Proposition 6 *a) Transitive closure of canonical bushy trees is* FO($\mathbf{Q}_u$)*-definable.*
b) Transitive closure of canonical bushy trees is not FO*-definable.* $\Box$

7 Conclusion

In this paper, we studied expressivity bounds for logics FO($\mathbf{C}$) and FO($\mathbf{Q}_u$) in the presence of auxiliary relations of large degree, i.e., those whose in- and out-degrees are comparable with the size of the input. Our main motivation was to extend results from the pure setting (where expressivity bounds have been proved recently) to the ordered setting (where FO($\mathbf{C}$) captures the DLOGTIME-uniform TC^0). We showed that in the presence of special auxiliary relations problems such as (deterministic) transitive closure are not expressible in these logics. The auxiliary relations we considered are obtained from linear orders by replacing them with a "thin" preorder on an arbitrarily small subset. By "thin" we mean that each equivalence class has the size at most 2. We also showed that, in a precise sense, the proof technique cannot be extened to the case of linear orders. Finally, we looked at the impact of auxiliary relations on the expressive power of logics with counting and unary quantifiers.

In the Introduction, we have mentioned that this work also has database motivation. In database theory, most query languages are based on first-order logic, while practical languages have features such as grouping and aggregate arithmetic function. A typical example would be a query in the language SQL (the de-facto standard of the commercial world) that asks for the average salary in each department, grouped by the department number. In several papers [8, 19, 20, 21], it was shown how to obtain expressivity bounds for such languages in the unordered setting. In [19, 21], it was done by producing an encoding of languages with grouping and aggregation in FO($\mathbf{C}$). However, in real languages some order (or at least preorder) is available at every base type. Thus, we can view the results of this paper as providing partial answers in the case when only a preorder is available at base types. One such case is the type *real*, where precision considerations often make it possible to talk about just a preorder relation (by simply disregarding meaningless digits). In such a case, our techniques imply that SQL-like languages cannot express the transitive closure of a graph. We shall provide more detail in the full version.

It still remains an open problem to prove expressivity bounds in the presence of an order relation. We believe that a descriptive complexity approach holds a promise, partly because it does not appear to fit the general scheme of natural proofs of Razborov and Rudich [26]. It is partly the case because we do not know how to translate expressivity bounds with various kinds of auxiliary relations into lower bounds for circuits (if indeed such a translation is possible). Another attempt to interpret such expressivity bounds in terms of circuit complexity is to find different notions of unifromity that will perhaps correspond to different auxiliary relations. We have not explored this yet.

An approach to proving lower bounds for TC^0 circuits, based on arithmetic circuits, was recently proposed in [1]. It may avoid the problems presented in [26]. In [1], the strongest results are obtained in the P-uniform and nonuniform setting, and the weakest for DLOGTIME-uniformity. Our results only apply to a fragment of DLOGTIME-uniform TC^0, so they appear to be of different nature than those in [1].

In conclusion, we believe that recent advances in the study of expressive power of logics with counting and unary quantifiers make it promising to use the tools of finite-model theory and descriptive complexity to attack some of the hard separation problems. The inverstigation carried out here produced the best possible partial results, but failed short of the final goal. Nevertheless, we plan to continue to work on refining the toolkit of expressive power of modest extensions of first-order logic in hope of proving the desired bounds.

References

[1] M. Agrawal, E. Allender and S. Datta. On TC^0, AC^0, and arithmetic circuits. In *Proc. 12th IEEE Conf. on Computational Complexity*, 1997.

[2] 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.

[3] S. Abiteboul, R. Hull, V. Vianu, *Foundations of Databases*, Addison Wesley, 1995.

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

[5] M. Benedikt, H.J. Keisler. On expressive power of unary counters. *Proc. Int. Conf. on Database Theory*, Springer LNCS 1186, 1997, pages 291–305.

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

[7] M. Consens and A. Mendelzon. Low complexity aggregation in GraphLog and Datalog, *Theoretical Computer Science* 116 (1993), 95–116.

[8] G. Dong, L. Libkin, L. Wong. Local properties of query languages. *Proc. Int. Conf. on Database Theory*, Springer LNCS 1186, 1997, pages 140–154.

[9] H.-D. Ebbinghaus. Extended logics: the general framework. In J. Barwise and S. Feferman, editors, *Model-Theoretic Logics*, Springer-Verlag, 1985, pages 25–76.

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

[11] K. Etessami. Counting quantifiers, successor relations, and logarithmic space, *JCSS*, 54 (1997), 400–411.

[12] R. Fagin, L. Stockmeyer, M. Vardi, On monadic NP vs monadic co-NP, *Information and Computation*, 120 (1994), 78–92.

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

[14] S. Grumbach, L. Libkin, T. Milo and L. Wong. Query languages for bags: expressive power and complexity. *SIGACT News*, 27 (1996), 30–37.

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

[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, M. Vardi. Infinitary logic and 0-1 laws. *Information and Computation*, 98 (1992), 258–294.

[19] L. Libkin. On the forms of locality over finite models. In *Proc. 12th IEEE Symp. on Logic in Computer Science*, 1997, pages 204–215. Full paper "Notions of locality and their logical characterization over finite models" by L. Hella, L. Libkin and J. Nurmonen is availabe as Bell Labs Technical Memo.

[20] L. Libkin, L. Wong. Query languages for bags and aggregate functions. *JCSS*, 55 (1997). Extended abstract in *PODS'94*, pages 155–166.

[21] L. Libkin and L. Wong. On the power of aggregattion in relational query languages. In *Proc. Database Programming Languages*, 1997, Springer LNCS, to appear.

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

[23] J. Nurmonen. Unary quantifiers and finite structures. PhD Thesis, University of Helsinki, 1996.

[24] M. Otto. Private communication. DIMACS, July 1997.

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

[26] A. Razborov and S. Rudich. Natural proofs. In *Proceedings of 26th Symposium on Theory of Computing*, pages 204–213, 1994.

[27] T. Schwentick. Graph connectivity and monadic NP. In *Proceedings of 35th Symposium on Foundations of Computer Science*, pages 614–622, 1994.

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