Logics with Rank Operators

Anuj Dawar* Martin Grohe† Bjarki Holm* Bastian Laubner†

*University of Cambridge
{anuj.dawar, bjarki.holm}@cl.cam.ac.uk

†Humboldt-Universität zu Berlin
{grohe, laubner}@informatik.hu-berlin.de

Abstract

*We introduce extensions of first-order logic (*FO*) and fixed-point logic (*FP*) with operators that compute the rank of a definable matrix. These operators are generalizations of the counting operations in* FP+C *(i.e. fixed-point logic with counting) that allow us to count the dimension of a definable vector space, rather than just count the cardinality of a definable set. The logics we define have data complexity contained in polynomial time and all known examples of polynomial time queries that are not definable in* FP+C *are definable in* FP+rk*, the extension of* FP *with rank operators. For each prime number p and each positive integer n, we have rank operators rk_p for determining the rank of a matrix over the finite field GF_p defined by a formula over n-tuples. We compare the expressive power of the logics obtained by varying the values p and n can take. In particular, we show that increasing the arity of the operators yields an infinite hierarchy of expressive power. The rank operators are surprisingly expressive, even in the absence of fixed-point operators. We show that* FO+rk$_p$ *can define deterministic and symmetric transitive closure. This allows us to show that, on ordered structures,* FO+rk$_p$ *captures the complexity class* $\mathrm{MOD}_p\,\mathrm{L}$*, for all prime values of p.*

1. Introduction

The question whether there is a logic capturing polynomial time was first raised by Chandra and Harel [6] in the context of database theory and later reformulated by Gurevich [15] in the form it is usually stated today. It asks for a logic, satisfying some minimal requirements to exclude pathological examples, in which precisely those properties of finite structures which are decidable in polynomial time are definable. The question is still wide open. It is considered to be one of the main open problems in finite model theory and in database theory.

A logic that has been intensively studied as a candidate for a logic capturing polynomial time is *fixed point logic with counting (*FP+C*)*. Proposed by Immerman [20] in the late 1980s, it has been shown to capture polynomial time on many natural classes of structures including planar graphs and structures of bounded tree width [12, 13, 14, 22]. Indeed, FP+C captures polynomial time on almost all structures in a precise technical sense [18]. However, Cai, Fürer, and Immerman [5] proved that FP+C does not capture polynomial time (on the class of all finite structures). The property they used to separate FP+C from polynomial time will play an important role in this paper; we call it the *CFI-property* henceforth.

It has been observed in recent years [1, 8] that all known examples of properties separating FP+C from polynomial time can be explained by the inability of the logic to express certain basic linear-algebraic properties such as the solvability of systems of linear equations; such properties are easily seen to be decidable in polynomial time by Gaussian elimination. Hence it is a natural idea to extend FP+C by operators that enable it to perform basic linear algebra. In this paper, we propose to extend fixed-point logic by operators defining the rank of definable matrices. The resulting logic FP+rk is at least as expressive as FP+C, because counting can be simulated by rank operators using the observation that the rank of a diagonal matrix is precisely the number of nonzero entries. It is easy to show that the CFI-property is definable in FP+rk. Hence FP+rk is strictly more expressive than FP+C. It is also not hard to see that all other known properties separating FP+C from polynomial time are definable in FP+rk. Still, the choice of adding a rank operator, rather than other possible linear-algebraic operators, to fixed-point logic may seem arbitrary at first. Having considered a number of other possibilities, we believe it is a good choice for a number of reasons. First of all, FP+rk still has polynomial time data complexity, that is to say that all FP+rk-definable properties are decidable in polynomial time. Secondly, as we have seen above, counting can directly be simulated with a rank operator and also the solvability of systems of linear equations can easily be expressed in terms of matrix rank (the system $A\mathbf{x} = \mathbf{b}$ is solvable if and only if the matrices A and $(A|\mathbf{b})$ have the same rank). This enables us, among other things, to define the CFI-property. Finally, the rank operator has a sim-

ple and natural formalisation in the familiar framework of two-sorted structures that is used to formalise the counting operators in FP+C. Indeed, it may be argued that the rank operator is no more than a generalised form of counting operator that counts the dimension of a vector space rather than the cardinality of a set. An alternative one might consider is an operator for computing the *determinant* of a matrix rather than its rank. However, as we show in Section 6, such an operator is already definable in FP+C, over all finite fields as well as the field of rationals and the ring of integers (building on earlier work of Rossman and Blass and Gurevich [2]). By similar techniques, we prove that even the rank of matrices over the field of rationals is expressible in FP+C. It is therefore seen that the additional expressive power of the logic FP+rk comes from the rank operators over finite fields.

While we do not have reason to believe that the logic FP+rk captures polynomial time, we believe that linear algebraic algorithms based on Gaussian elimination need to be understood in a logical context if progress is to be made on the existence of a logic capturing polynomial time. These would seem to be the most fundamental polynomial time algorithms that cannot be represented in FP+C. It turns out to be difficult to understand the interaction between logic and linear algebra, and we think it is most promising to study it in isolation first and hence consider a logic such as FP+rk before proceeding to more complicated logics that may potentially capture polynomial time.

Indeed, a considerable part of this paper is devoted to even simpler logics, the extension FO+rk of first-order logic by rank operators and, for every prime p, its fragment $\mathsf{FO{+}rk}_p$ that only has a rank operator over the field GF_p. We show that undirected graph reachability is expressible in $\mathsf{FO{+}rk}_p$ (for every prime p) and hence that $\mathsf{FO{+}rk}_p$ is strictly more expressive than the extensions DTC and STC of first-order logic by deterministic and symmetric transitive closure operators. Over ordered finite structures, we characterise the complexity of the logics by proving that $\mathsf{FO{+}rk}_p$ captures the complexity class $\mathrm{MOD}_p\,\mathrm{L}$. In particular, over ordered structures the logic $\mathsf{FO{+}rk}_2$ captures the complexity class $\oplus\mathrm{L} = \mathrm{MOD}_2\,\mathrm{L}$ (known as "parity L") consisting of all problems for which there is a nondeterministic logarithmic space Turing machine accepting an input if and only if an odd number of computation paths halt in an accepting state.

Our main result on the logic FP+rk is that it has a strict arity hierarchy. The row and column indices of definable matrices are tuples of elements of a structure, and the *arity* of a rank operator is the sum of the lengths of the tuples used as row and column indices of the matrices it applies to. The strictness of the arity hierarchy distinguishes the rank operators from counting operators: it is known that unary counting operators can simulate counting operators of all arities. Our result is based on a theorem due to Hella [17] stating that no extension of finite variable infinitary logic by Lindström quantifiers of bounded arity can define all polynomial time properties. Even though rank operators are not themselves Lindström quantifiers, it can be shown that the n-ary fragment of FP+rk can be embedded into the extension of finite variable infinitary logic by n-ary Lindström quantifiers.

The paper is organised as follows: After some preliminaries, we formally introduce the rank operators in Section 2. In Section 3, we present our results on the expressive power of first-order rank logic FO+rk. In Section 4 we prove that over ordered finite structures, $\mathsf{FO{+}rk}_p$ captures the complexity class $\mathrm{MOD}_p\,\mathrm{L}$ for prime p. Then in Section 5, we prove the strictness of the arity hierarchy. In Section 6 we prove that determinants over the integers, rationals and finite fields and matrix rank over the rationals are expressible in FP+C. Finally, in Section 7 we discuss several interesting directions for further research.

Due to space constraints, proofs of some of our results have been deferred to the full version of this paper.

2. Rank Operators

A *vocabulary* τ is a finite sequence of relation symbols $(R_1,\dots R_k)$ in which each R_i has a fixed arity n_i. A *τ-structure* $\mathscr{A} = \big(U(\mathscr{A}), R_1^{\mathscr{A}},\dots,R_k^{\mathscr{A}}\big)$ is a non-empty set $U(\mathscr{A})$, called the *universe* of $\mathscr{A}$, together with relations $R_i^{\mathscr{A}} \subseteq U(\mathscr{A})^{n_i}$ for $1 \le i \le k$. We let $|\mathscr{A}| := |U(\mathscr{A})|$.

We write $\mathbb{N}$ and $\mathbb{N}_0$ for the positive and non-negative integers, respectively. For $m,n \in \mathbb{N}_0$, let $[m,n] := \{\ell \in \mathbb{N}_0 \big| m \le \ell \le n\}$ and $[n] := [1,n]$. We often denote tuples $(v_1,\dots,v_k)$ by $\vec{v}$ and denote their length by $|\vec{v}|$.

Given a logic L and a vocabulary τ, $\mathsf{L}[\tau]$ denotes the set of τ-formulas of L. For a formula $\varphi \in \mathsf{L}[\tau]$ we write $\varphi(\vec{x})$ to indicate that all of φ's free variables are among $\vec{x}$. Given a τ-structure $\mathscr{A}$, we write $\mathscr{A} \models_{\mathsf{L}} \varphi[\vec{a}]$ if $\vec{a} \in U(\mathscr{A})^{|\vec{x}|}$ and $\mathscr{A}$ satisfies φ under the assignment of a_i to x_i for every $i \in [|\vec{x}|]$. When the logic L is clear from the context, we omit the subscript to the satisfaction relation.

A τ-formula $\varphi(\vec{x})$ with $|\vec{x}| = k$ defines a *k-ary query* that takes any τ-structure $\mathscr{A}$ to the set of k-tuples $\vec{a}$ from $U(\mathscr{A})$ for which $\mathscr{A} \models \varphi[\vec{a}]$. We say that a logic L_1 is (at least) as expressive as a logic L_2, and write $\mathsf{L}_2 \leqq \mathsf{L}_1$, if every query definable in L_2 is also definable in L_1. We write $L_1 \equiv L_2$ if $L_1 \leqq L_2$ and $L_2 \leqq L_1$.

Two-sorted structures. We equip structures with an additional integer sort. For a τ-structure $\mathscr{A} = \big(U(\mathscr{A}), (R^{\mathscr{A}})_{R\in\tau}\big)$ we define $\mathscr{A}^+$ to be the extension of $\mathscr{A}$ by the standard model of arithmetic. In other words, $\mathscr{A}^+$ is the two-sorted structure $\big(U(\mathscr{A}), \mathbb{N}_0, (R^{\mathscr{A}})_{R\in\tau}, +, \cdot, \le, 0, 1\big)$, where $+$ and $\cdot$ are binary functions denoting standard addition and multiplication over the integer sort, $\le$ is the linear order on the integers, and 0 and 1 are the usual constants from $\mathbb{N}_0$.

For logics over such two-sorted structures, we assume all variables to be typed, so each variable x ranges either over the universe $U(\mathscr{A})$ or over the numbers $\mathbb{N}_0$. We say that a tuple of variables $\vec{x}$ has type $k.\ell$ if $|\vec{x}| = k+\ell$, the first k variables of $\vec{x}$ range over $U(\mathscr{A})$ and the last ℓ variables range over $\mathbb{N}_0$. A *numeric term* is a term in the language of $\mathscr{A}^+$ that takes values in $\mathbb{N}_0$. If t is a numeric term, then $t^{\mathscr{A}}$ denotes the integer value t takes over $\mathscr{A}^+$. It will be convenient to view formulas as numeric terms taking the values 1, 0 corresponding to their truth values.

In order to avoid undecidability, quantification over number variables has to be bounded. Thus, if x is a number variable, its binding quantifier must appear in the form $\forall x \le t\ \varphi$ or $\exists x \le t\ \varphi$ for a numeric term t and a formula φ. Let us denote first-order logic over the two-sorted extension of structures with bounded quantification over the numerical sort by FO^+. We usually write $\mathscr{A} \models \varphi$ to mean $\mathscr{A}^+ \models_{\mathsf{FO}^+} \varphi$.

Rank operators. We are now ready to define rank operators. Consider a numeric term η and variables $\vec{x}$ and $\vec{y}$, which are possibly free in η. Given a structure $\mathscr{A}$, define $m_{\vec{a}\vec{b}} := \eta^{\mathscr{A}}[\vec{a},\vec{b}]$ for tuples $\vec{a},\vec{b}$ from $U(\mathscr{A})$ interpreting $\vec{x}$ and $\vec{y}$ respectively. We consider $M := \left(m_{\vec{a}\vec{b}}\right)$ as an integer matrix whose rows are indexed by $|\vec{x}|$-tuples and whose columns are indexed by $|\vec{y}|$-tuples.

For prime p let M_p denote the matrix of the residue classes of M's matrix entries mod p. We view M_p as a matrix over GF_p. Notice that the rank of M_p is well-defined since it does not depend on the ordering of the rows and columns. In general, we also allow matrices to be indexed by number variables, but they need to be bounded again so that we obtain finite matrices.

Definition 2.1. Given a prime p, a τ-structure $\mathscr{A}$, and a numeric $\mathsf{FO}^+[\tau]$-term or formula η, let $\vec{x}_1$, $\vec{y}_1$ be universe variables, $\vec{x}_2$, $\vec{y}_2$ be number variables and $\vec{t}_1$, $\vec{t}_2$ be tuples of numeric terms bounding the number variables in $\vec{x}_2$ and $\vec{y}_2$, respectively. Then $\mathrm{rk}_p(\vec{x}_1\,\vec{x}_2 \le \vec{t}_1, \vec{y}_1\,\vec{y}_2 \le \vec{t}_2)\ \eta$ is a numeric term denoting the rank of $M_p = \left(\eta^{\mathscr{A}}[\vec{a}_1\vec{a}_2, \vec{b}_1\vec{b}_2] \mod p\right)$ over GF_p.

For prime p, we write $\mathsf{FO{+}rk}_p$ for the extension of FO^+ with the rank operator rk_p, and we let $\mathsf{FO{+}rk}$ be the extension of FO^+ with all the rank operators.

Remark. For any numeric term t, there is a polynomial p so that $t^{\mathscr{A}} \le p(|\mathscr{A}|)$ in any structure $\mathscr{A}$. To see this, use induction over terms and notice that the rank of a matrix is bounded by both its row and column dimension.

Other logics. In this paper, we relate $\mathsf{FO{+}rk}$ to various other logics that play a prominent role in descriptive complexity theory. In particular, we consider *inflationary fixed-point logic* FP, *deterministic transitive closure logic* DTC, *symmetric transitive closure logic* STC, *transitive closure logic* TC, and *infinitary first-order logic with finitely many variables* $\mathsf{L}^{\omega}_{\infty\omega}$. For a detailed discussion we refer to the standard literature [10, 21].

Let φ be an $\mathsf{FO}^+[\tau]$-formula. We define the *counting terms* $\sharp\vec{x}\,\varphi$ to denote the number of assignments to $\vec{x}$ so that φ holds in $\mathscr{A}$. Once again, we require all occurrences of number variables in $\vec{x}$ to be bounded by numeric terms and we write $\sharp\vec{x}\vec{y} \le \vec{t}\,\varphi$ to indicate that the number variables y_i are bounded by numeric terms t_i for each i respectively. Let $\mathsf{FO{+}C}$ denote the extension of FO^+ with counting terms.

We also consider FP over the two-sorted extension of structures with bounded quantification. We write $\mathrm{ifp}_{X \leftarrow \vec{x}\vec{y} \le \vec{t}}\,\varphi$ for the inflationary fixed-point of φ over the relation variable X of mixed type $|\vec{x}|.|\vec{y}|$, where the number variables in $\vec{y}$ are bounded by the numeric terms in $\vec{t}$. Extending this now with counting terms, we obtain the logic $\mathsf{FP{+}C}$ which plays an important role in descriptive complexity theory. Similarly, for prime p, the rank logic $\mathsf{FP{+}rk}_p$ is obtained by extending FP in the two-sorted setting with the rank operator rk_p, and we write $\mathsf{FP{+}rk}$ for the extension of FP with all the rank operators.

Remark. For any formula $\varphi(\vec{x})$ and any prime p, the term $\mathrm{rk}_p(\vec{x},\vec{y})\ \vec{x}{=}\vec{y} \wedge \varphi(\vec{x})$ with $|\vec{y}| = |\vec{x}|$ denotes the number of $|\vec{x}|$-tuples from $U(\mathscr{A})$ so that φ holds. It follows immediately that $\mathsf{FO{+}C} \leqq \mathsf{FO{+}rk}$ and $\mathsf{FP{+}C} \leqq \mathsf{FP{+}rk}$.

Remark. Observe that all the rank logics $\mathsf{FP{+}rk}$, $\mathsf{FO{+}rk}$, $\mathsf{FP{+}rk}_p$ and $\mathsf{FO{+}rk}_p$ (for prime p) are closed under first-order reductions.

Linear systems. Let a τ-structure $\mathscr{A}$ be given along with numeric τ-terms $\psi(\vec{x},\vec{y})$ and $\beta(\vec{x})$. Considering $\psi^{\mathscr{A}}[\vec{a},\vec{b}]$ as a matrix A_ψ and $\beta^{\mathscr{A}}[\vec{c}]$ as a vector $\mathbf{b}_\beta$ over GF_p for some prime p, ψ and β describe the system of linear equations $A_\psi \mathbf{x} = \mathbf{b}_\beta$ on $\mathscr{A}$.

Such a system is solvable if and only if $\mathbf{b}_\beta$ is contained in the span of the column vectors of A_ψ, or in other words, if and only if adding $\mathbf{b}_\beta$ as a new column to A_ψ does not increase the rank of the matrix. Based on this, it is easy to see that the following formula of $\mathsf{FO{+}rk}_p$ defines solvability of the system.

$$\forall z\ \ \mathrm{rk}_p(\vec{x},\vec{y}y')\ \big((y' \neq z)\cdot\beta + (y' = z)\cdot\psi\big)\ \le\ \mathrm{rk}_p(\vec{x},\vec{y})\,\psi$$

In the above formula, by our convention, the formulas $y' = z$ and $y' \neq z$ take on truth values from $\{0,1\} \subset \mathbb{N}_0$. Note that the matrix defined on the left-hand side of the inequality will contain multiple copies of the column vector $\mathbf{b}_\beta$, which of course does not alter the solvability of the system. Atserias et al. [1] have shown that the solvability of linear systems cannot be defined in $\mathsf{FP{+}C}$. Thus, $\mathsf{FP{+}rk}$ is strictly more expressive than $\mathsf{FP{+}C}$.

3. Expressive Power of FO+rk

Graph connectivity. The *symmetric* (s,t)*-reachability* problem is the problem of determining, given a graph G with distinguished vertices s and t, whether there is a path from s to t in the undirected graph underlying G. We show that symmetric (s,t)-reachability can be defined in $\mathsf{FO{+}rk}_p$ for all primes p.

Let $G=(V,E)$ be an undirected graph and let s and t be two vertices in V. For a prime p, let $\mathfrak{S}_{G,s,t}$ be the system of linear equations over GF_p with variables x_v for all $v \in V$ and equations:

- for every edge $e=(u,v) \in E$: $x_u - x_v = 0$,
- $x_s = 1$; $x_t = 0$.

Lemma 3.1. *The linear system $\mathfrak{S}_{G,s,t}$ is solvable over GF_p iff there is no path between s and t in the graph G.*

Proof. This follows from the observation that the edge equations of $\mathfrak{S}_{G,s,t}$ force variables x_u and x_v to take the same value if u and v are in the same connected component of G. □

The matrix of the system $\mathfrak{S}_{G,s,t}$ is easily definable in first-order logic by a numeric term $\eta(x_1x_2,y)$ holding the value 1 at (ss,s), (tt,t), and (uv,u) and -1 at (uv,v) for edges $(u,v) \in E$. Note that every edge equation is stated twice in equivalent ways. For any directed graph G, the underlying undirected graph is also first-order definable. Hence there is a first-order reduction from symmetric (s,t)-reachability to the problem of deciding solvability of linear systems.

Corollary 3.2. *Symmetric (s,t)-reachability is definable in $\mathsf{FO{+}rk}_p$ for all primes p.*

The above method for defining reachability fails in general when applied to directed graphs. We can, however, consider an important special case: the class of all directed graphs whose vertices have out-degree at most 1.

Let $G=(V,E)$ be a directed graph. Define the *deterministic* part E_d of E as all those edges (u,v) from E for which u has out-degree 1. Thus, $E_d \subseteq E$ and all vertices of $G_d := (V,E_d)$ have out-degree at most 1.

Given G and vertices $s,t \in V$, the *deterministic* (s,t)*-reachability* problem asks whether there is a path from s to t in G_d. It is easy to see that after removing any outgoing edge from t, deterministic (s,t)-reachability becomes equivalent to symmetric (s,t)-reachability, thus establishing:

Lemma 3.3. *Deterministic (s,t)-reachability is definable in $\mathsf{FO{+}rk}_p$ for all primes p.*

Corollary 3.4. *On the class of all finite structures, $\mathsf{STC} \lneqq \mathsf{FO{+}rk}_p$ and $\mathsf{DTC} \lneqq \mathsf{FO{+}rk}_p$ for all primes p.*

Cai-Fürer-Immerman graphs. Cai, Fürer and Immerman [5] proved that FP+C does not capture PTIME on the class of all finite structures, thereby settling what had been an important open problem in descriptive complexity theory. For the proof, they constructed a query on a class of graphs that can be defined by a polynomial time computation but not by any sentence of FP+C. In this section we show that this query can be expressed in the logic FO+rk, by considering a first-order definable system of linear equations over GF_2.

We first introduce the class of Cai-Fürer-Immerman (CFI) graphs on which we define the separating query. The following presentation of the graphs is adapted from [5, 9]. Note that unlike the presentation of Dawar et al. [9], who show that the CFI query can be expressed in the logic of choiceless polynomial time, we do not require an ordering on the underlying graphs G.

Definition 3.5 (Cai-Fürer-Immerman graphs)**.** Let $G=(V,E)$ be a connected undirected graph with at least two vertices. We denote the set of edges incident to $v \in V$ by $E(v)$. Let $T \subseteq V$. For each $v \in T$, let

$$I_v := \{v_Z : Z \subseteq E(v), |Z| \equiv 1 \pmod 2\},$$

and for each $v \in V \setminus T$, let

$$I_v := \{v_Z : Z \subseteq E(v), |Z| \equiv 0 \pmod 2\}.$$

Let $\hat{V} := \cup_{v \in V} I_v$, $\hat{E} := \{e_0, e_1 \big| e \in E\}$, $\hat{C} := \{e_c \big| e \in E\}$, and $U^* := \hat{V} \cup \hat{E} \cup \hat{C}$. Define an edge relation on U^*

$$\begin{aligned} E^* := &\big\{\{v_Z, e_1\} : e \in Z\big\} \cup \\ &\big\{\{v_Z, e_0\} : e \in E(v) \setminus Z\big\} \cup \\ &\big\{\{e_i, e_c\} : e \in E(v), i \in \{0,1\}\big\}, \end{aligned}$$

and a unary relation $C^* := \hat{C} \subset U^*$. Finally define $\mathscr{G}^T := (U^*, E^*, C^*)$.

We refer to the sets of vertices $\hat{C}$, $\hat{E}$ and $\hat{V}$ as the *colour nodes*, *outer nodes* and *inner nodes* of $\mathscr{G}^T$, respectively. The *parity* of a CFI graph $\mathscr{G}^T$ is the parity of $|T|$. We say $\mathscr{G}^T$ is *even* if it has even parity and *odd* if it has odd parity. In [5], Cai et al. show the following:

- For a connected graph G, where every vertex has degree at least two, and all $T,S \subseteq V(G)$, the graphs $\mathscr{G}^T$ and $\mathscr{G}^S$ are isomorphic iff they have the same parity.
- While there is a PTIME algorithm that can distinguish between the odd and even CFI graphs of any graph G, there is no fixed formula of FP+C that can do the same.

Now let $G=(V,E)$ be a connected graph where every vertex has degree at least two and let $\mathscr{G}^T$ be a CFI graph constructed from G.

Let $\mathfrak{S}_{\mathscr{G}^T}$ be the system of linear equations over GF_2 with variables x_{e_i} for all $e_i \in \hat{E}$ and x_{v_Z} for all $v_Z \in \hat{V}$, and equations:

- for all $e_i \in \hat{E}$: $x_{e_i} + x_{e_{1-i}} = 1$,
- $\sum_{v_Z \in \hat{V}} x_{v_Z} = 0$,
- for all $v_Z \in \hat{V}$:

  $\sum_{e \in Z} x_{e_1} + \sum_{e \in E(v) \setminus Z} x_{e_0} = \sum_{v_Y \in I_v} x_{v_Y}$.

The following is not hard to establish.

Lemma 3.6. *$\mathfrak{S}_{\mathscr{G}^T}$ is first-order definable over $\mathscr{G}^T$.*

Proof. It can be seen that there are first-order formulas $\varphi_c(x)$, $\varphi_o(x)$ and $\varphi_i(x)$ that define the sets $\hat{C}$, $\hat{E}$ and $\hat{V}$, respectively. Similarly, there is a first-order formula $\theta_p(x,y)$ that says that x and y are distinct outer nodes derived from the same edge $e \in E$, and a first-order formula $\theta_i(x,y)$ that says that x and y are inner nodes derived from the same vertex $v \in V$.

The system $\mathfrak{S}_{\mathscr{G}^T}$ can now be defined by formulas $\varphi(x,y)$ and $\beta(x)$ over $\mathscr{G}^T$ in the following way. The equations for the outer nodes e_i are defined at row indices a for which $\varphi_o[a] = 1$. Similarly, the equations for inner nodes v_Z are defined at row indices a for which $\varphi_i[a] = 1$, using $\theta_i(x,y)$ and the fact that the set of e_1 with $e \in Z$ is exactly the neighbourhood of v_Z in $\mathscr{G}^T$, and the set of e_0 with $e \in E(v) \setminus Z$ can be defined similarly. Finally, the equation that sums all the x_{v_Z} can be defined at row indices a for which $\varphi_c[a] = 1$; there will be multiple copies of this equation, which of course does not affect the solvability of the system. The definition of $\beta(x)$ follows similarly. □

Lemma 3.7. *The system $\mathfrak{S}_{\mathscr{G}^T}$ is solvable iff $\mathscr{G}^T$ is even.*

The preceding lemmas now establish that there is a first-order reduction from the problem of distinguishing odd and even CFI graphs to the problem of deciding solvability of linear systems over GF_2.

Theorem 3.8. *There is a sentence $\varphi_{CFI} \in \mathsf{FO{+}rk}_2$ that holds in structures $\mathscr{G}^T$ when $|T|$ is even but not in structures $\mathscr{G}^T$ when $|T|$ is odd.*

Since the work of Cai et al., other constructions that expose the limitations of FP+C have been given. Gurevich and Shelah [16] define a class of rigid structures known as *multipedes*, and consider the problem of uniformly defining a linear order over this class. They show that this problem, while computable in polynomial time, is not definable by any fixed formula of FP+C. We are able to show that such an order can be defined in FO+rk. The details are left out of the present paper due to lack of space.

4. Descriptive Complexity

It is a classical result of descriptive complexity theory that extensions of first-order logic with various fixed-point operators capture different complexity classes on the class of ordered structures. TC, for instance, captures nondeterministic logspace, while FP captures PTIME in the presence of an order. In this section, we show such a natural correspondence for first-order logic with rank operators, namely that for each prime p, $\mathsf{FO{+}rk}_p$ captures the complexity class $\mathrm{MOD}_p\,\mathrm{L}$ on ordered structures.

We start by reviewing some classical notions from the field of descriptive complexity. For a vocabulary τ with $\leq \in \tau$, we call a τ-structure $\mathscr{A}$ an *ordered structure* if $\mathscr{A}$ interprets $\leq$ as a total order of its universe. We will identify $\mathscr{A}$'s linearly ordered universe with $[0, |\mathscr{A}| - 1] \subset \mathbb{N}_0$.

An ordered τ-structure $\mathscr{A}$ can be encoded as a word over $\{0,1\}$ in a canonical way (see e.g. [10] for details). We write $\mathsf{enc}(\mathscr{A})$ for the canonical encoding of an ordered structure $\mathscr{A}$. The Turing machines we consider use $\{0,1\}$ as their input and work tape alphabet. If K is a class of ordered τ-structures, we say that a Turing machine M decides K if for any ordered τ-structure $\mathscr{A}$,

$$M(\mathsf{enc}(\mathscr{A})) \begin{cases} \text{accepts} & \text{if } \mathscr{A} \in K, \\ \text{rejects} & \text{if } \mathscr{A} \notin K. \end{cases}$$

Since it can be decided in logarithmic space whether a given string in $\{0,1\}^*$ is a valid encoding of a τ-structure, M can be turned into a machine that decides $\{\mathsf{enc}(\mathscr{A}) \,|\, \mathscr{A} \in K\} \subseteq \{0,1\}^*$. For a complexity class C we write $K \in \mathrm{C}$ to mean $\{\mathsf{enc}(\mathscr{A}) \,|\, \mathscr{A} \in K\} \in \mathrm{C}$.

Given a complexity class C, a logic L *captures* C *on ordered structures* if for any vocabulary τ with $\leq \in \tau$ and any class K of ordered τ-structures, $K \in \mathrm{C}$ if and only if there is a sentence φ_K of $\mathsf{L}[\tau]$ that defines K.

Definition 4.1. For a non-deterministic Turing machine M, let $|M(x)|$ denote the number of accepting computation paths on input string x. Let $n \in \mathbb{N}$. A $\mathrm{MOD}_n\,\mathrm{L}$-Turing machine M is a non-deterministic Turing machine with logarithmic workspace which is said to accept an input x whenever $|M(x)| \neq 0 \bmod n$.

The complexity class $\mathrm{MOD}_n\,\mathrm{L}$ consists of all problems $P \subseteq \{0,1\}^*$ for which there is a $\mathrm{MOD}_n\,\mathrm{L}$-Turing machine M deciding P.

$\mathrm{MOD}_2\,\mathrm{L}$ is better known under the name "parity logspace", usually denoted by $\oplus\mathrm{L}$.

Theorem 4.2. *Let p be prime. $\mathsf{FO{+}rk}_p$ captures $\mathrm{MOD}_p\,\mathrm{L}$ on ordered structures.*

Proof. The proof consists of two parts. Firstly, we have to show that for any sentence $\varphi \in \mathsf{FO{+}rk}_p$ there is a $\mathrm{MOD}_p\,\mathrm{L}$-Turing machine M_φ that, given the encoding of a structure

$\mathscr{A}$, decides $\mathscr{A} \models \varphi$. Secondly, given a $\mathrm{MOD}_p\,\mathrm{L}$-Turing machine M, we construct a sentence φ_M that holds in a structure $\mathscr{A}$ if and only if M accepts $\mathsf{enc}(\mathscr{A})$.

For the first part, assume that τ is a vocabulary with $\leq \in \tau$, and that φ is a $\mathsf{FO{+}rk}_p[\tau]$-sentence. In order to deal with rank operators occurring in φ, we need two results on $\mathrm{MOD}_p\,\mathrm{L}$-machines. The first one says that the rank of a matrix over GF_p can be decided by a $\mathrm{MOD}_p\,\mathrm{L}$-machine.

Lemma 4.3 (Buntrock et al. [4])**.** *Let p be prime. There is a $\mathrm{MOD}_p\,\mathrm{L}$ machine M_{rk} which takes as input an integer $r \in \mathbb{N}_0$ and a matrix $A \in \mathsf{GF}_p^{m\times n}$ and decides if $\mathrm{rk}A = r$.*

The second result states that $\mathrm{MOD}_p\,\mathrm{L}$-machines that make oracle queries to a $\mathrm{MOD}_p\,\mathrm{L}$ problem can be simulated by a $\mathrm{MOD}_p\,\mathrm{L}$-Turing machine without oracle queries.

Lemma 4.4 (Hertrampf et al. [19])**.** *Let p be prime. Then* $\mathrm{MOD}_p\,\mathrm{L}^{\mathrm{MOD}_p\,\mathrm{L}} = \mathrm{MOD}_p\,\mathrm{L}$.

It is left to show that there is a $\mathrm{MOD}_p\,\mathrm{L}$-machine M_φ that decides φ by induction over the construction of φ. The argument is fairly tedious but straightforward and will be omitted.

For the other direction, consider a $\mathrm{MOD}_p\,\mathrm{L}$-Turing machine M with space bound $d\cdot\log n$ that decides a class of τ-structures $K \in \mathrm{MOD}_p\,\mathrm{L}$. Without loss of generality we assume that M has only one accepting configuration. We construct a formula φ_M that defines K. By standard methods we can restrict ourselves to structures $\mathscr{A}$ so that $|\mathscr{A}| > \max\{d\log|\mathscr{A}|, q\}$ where q is the number of states of M and choose d' large enough so that all configurations of M can be encoded by d'-tuples of elements from $\mathscr{A}$.

Consider the configuration graph of M, that is, the directed graph G_M with vertices $\vec{a} \in U(\mathscr{A})^{d'}$ and edges $(\vec{a},\vec{b})$ whenever $\vec{b}$ is a successor configuration of $\vec{a}$ under M's transition relation. If s and t denote the unique start and accept configurations, respectively, M accepting $\mathscr{A}$ is equivalent to the condition that the number of paths from s to t in G_M is $\neq 0 \bmod p$. Note that G_M can be assumed to be free of cycles; if it is not, use d variables to add a time mark to each configuration that is increased by 1 at every transition. Since $\mathsf{DTC} \leqq \mathsf{FO{+}rk}_p$ (Corollary 3.4), the following result shows that G_M can be defined in $\mathsf{FO{+}rk}_p$.

Lemma 4.5 (Ebbinghaus and Flum [10])**.** *There are DTC-formulas $\chi_{start}(\vec{x})$, $\chi_{accept}(\vec{x})$, and $\chi_{succ}(\vec{x},\vec{y})$ such that for all ordered τ-structures $\mathscr{A}$ with $|\mathscr{A}| > \max\{d\log|\mathscr{A}|, q\}$ and $\vec{a} \in U(\mathscr{A})^{d'}$,*

- *$\mathscr{A} \models \chi_{start}(\vec{a})$ ($A \models \chi_{accept}(\vec{a})$) if and only if $\vec{a}$ is the encoding of the start (accept) configuration of M,*
- *$\mathscr{A} \models \chi_{succ}(\vec{a},\vec{b})$ if and only if $\vec{b}$ is a valid successor configuration of $\vec{a}$.*

$\chi_{\mathrm{succ}}(\vec{x},\vec{y})$ defines the adjacency matrix A of G_M. Let I denote the identity matrix of the same dimension as A. Then $I-A$ is definable in $\mathsf{FO{+}rk}_p$ by a term $\eta(\vec{x},\vec{y})$, and the term

$$\eta^*(\vec{x},\vec{y}) \;:=\; \big(\neg\chi_{\mathrm{accept}}(\vec{x}) \wedge \neg\chi_{\mathrm{start}}(\vec{y})\big)\cdot\eta(\vec{x},\vec{y})$$

defines $I-A$ with row t and column s set to 0. The formula $\varepsilon(\vec{x},\vec{y}) := \vec{x}=\vec{y} \wedge \neg\chi_{\mathrm{start}}(\vec{x})$ defines the identity matrix of the same dimension as A with one row set to 0. Let

$$\varphi'_M \;:=\; \mathrm{rk}_p(\vec{x},\vec{y})\;\varepsilon(\vec{x},\vec{y}) \;=\; \mathrm{rk}_p(\vec{x},\vec{y})\;\eta^*(\vec{x},\vec{y}).$$

The following completes the proof of Theorem 4.2.

Lemma 4.6. *For any ordered τ-structure $\mathscr{A}$ with $|\mathscr{A}| > \max\{d\log|\mathscr{A}|, q\}$, $\mathscr{A} \models \varphi'_M$ if and only if M accepts $\mathscr{A}$.*

Proof. As G_M is cycle-free, there is no path of length $\geq n^{d'} =: m$, hence $A^m = 0$. Thus, $I-A$ is non-singular over GF_p, with the inverse explicitly given by $(I-A)^{-1} = I + A + A^2 + \ldots + A^{m-1}$, where all arithmetic is over GF_p. Notice that for $k \in \mathbb{N}_0$, the $(i,j)^{\text{th}}$ entry of A^k equals the number of paths modulo p of length k from i to j in G_M. Thus, $(I-A)^{-1}$ is the matrix of the total numbers of paths modulo p. Recall that s and t denote the start and accept configuration, respectively. Then M accepts $\mathscr{A}$ if and only if $(I-A)^{-1}(s,t) \neq 0$.

The adjugate rule says that for any invertible matrix B, the entries b^{-1}_{ij} of its inverse B^{-1} are given by

$$b^{-1}_{ij} = (-1)^{i+j}\det B_{ji} \big/ \det B,$$

where B_{ji} is B with the j^{th} row and the i^{th} column deleted. To check if $(I-A)^{-1}(s,t) \neq 0$, it is therefore enough to test if $(I-A)_{ts}$ has full rank, which is exactly what φ'_M does. □

5. Arity Hierarchy of Rank Operators

Generalized quantifiers were introduced by Lindström in [24] and have been studied as a way to increase the expressiveness of FO by a prescribed query. For any vocabulary $\sigma = (R_1,\ldots,R_k)$ with relations R_i of arity n_i, and any class of σ-structures K, let Q_K denote the *Lindström quantifier* associated with K. For any vocabulary τ, a τ-structure $\mathscr{A}$ satisfies the formula $Q_K\vec{x}_1\ldots\vec{x}_k(\psi_1(\vec{x}_1),\ldots,\psi_k(\vec{x}_k))$ if $(U(\mathscr{A}),\psi_1^{\mathscr{A}},\ldots,\psi_k^{\mathscr{A}}) \in K$ as a σ-structure. If Q is a set of Lindström quantifiers, then $\mathsf{L}(Q)$ denotes the extension of a logic L by all the quantifiers in Q. More information on Lindström quantifiers can be found in [10].

In 1996, Hella [17] proved that for any $n \in \mathbb{N}$, augmenting infinitary first-order logic with all Lindström quantifiers of arity at most n, denoted $\mathsf{L}^{\omega}_{\infty\omega}(Q_n)$, is not expressive enough to define all PTIME queries over the class of all structures (not necessarily ordered). Since $\mathsf{FP{+}C} \leqq \mathsf{L}^{\omega}_{\infty\omega}(Q_1)$, his result extends Cai, Fürer, and Immerman's result discussed in Section 3.

Our aim is to show that the arities of rank operators yield a strict hierarchy. Rank operators are not themselves Lindström quantifiers, but they can be translated into such quantifiers. To be precise, we define quantifiers $\mathrm{rk}_p^{\leq r}$ so that $\mathrm{rk}_p^{\leq r}\,\vec{x}\vec{y}\,(\eta)$ is interpreted as $\mathrm{rk}_p\,\eta[\vec{x},\vec{y}] \leq r$. The *arity* of the quantifier is $|\vec{x}+\vec{y}|$. Writing $\mathscr{R}_n$ for the set of all rank quantifiers of arity at most n and $\mathsf{FP{+}rk}^{[n]}$ for the set of all those $\mathsf{FP{+}rk}$-formulas in which all occurrences of rk-operators are of arity at most n, it can be shown that formulas of $\mathsf{FP{+}rk}^{[n]}$ can be translated into $\mathsf{L}^{\omega}_{\infty\omega}(\mathscr{R}_n)$, so in particular $\mathsf{FP{+}rk}^{[n]} \leqq \mathsf{L}^{\omega}_{\infty\omega}(Q_n)$ for all $n \in \mathbb{N}$.

Theorem 5.1. *For any $n \in \mathbb{N}$ there is an $\mathsf{FO}(\mathscr{R}_{n+1})$ query that is not definable in $\mathsf{L}^{\omega}_{\infty\omega}(Q_n)$. Thus, $\mathsf{FP{+}rk}^{[n]} \lneqq \mathsf{FP{+}rk}^{[n+1]}$ and $\mathsf{FO{+}rk}^{[n]} \lneqq \mathsf{FO{+}rk}^{[n+1]}$ for any n.*

We prove this theorem by showing that Hella's queries, which separate $\mathsf{L}^{\omega}_{\infty\omega}(Q_n)$ from PTIME, can be expressed using a linear system over GF_2 of arity $n+1$. In fact, our proof already shows the strictness of the rk_2 arity hierarchy. By generalizing Hella's construction, strictness can be shown for the arity hierarchy of rk_p for every prime p, though we omit details for lack of space. The following construction is due to Hella [17].

Let $C = \{c_1,\dots,c_{n+1},d_1,\dots,d_{n+1}\}$ be equipped with the quasi order $\prec$ defined by

$$x \prec y \quad :\Leftrightarrow \quad x \in \{c_i,d_i\} \text{ and } y \in \{c_j,d_j\} \text{ for some } 1 \leq i < j \leq n+1,$$

and let $P = \{d_1,\dots,d_{n+1}\} \subset C$. Define relations R^+ and R^- by

$$\begin{aligned}(a_1,\dots,a_{n+1}) \in R^+ :\Leftrightarrow \quad & a_1 \prec \dots \prec a_{n+1} \\ & \text{and } |\{i : a_i \in P\}| \text{ is even}\\ (a_1,\dots,a_{n+1}) \in R^- :\Leftrightarrow \quad & a_1 \prec \dots \prec a_{n+1} \\ & \text{and } |\{i : a_i \in P\}| \text{ is odd}\end{aligned}$$

Now assume that $n \geq 2$ and $G = (V, E^G, <^G)$ is a finite undirected graph which is connected, regular of degree $n+1$, and $<^G$ is a strict linear order on V. For every vertex $u \in V$, fix an enumeration $h_u : \{v \mid (u,v) \in E\} \to [n+1]$ of its $n+1$ neighbors.

Definition 5.2. Let $\tau = (R,E,<)$ be a vocabulary, where R is $(n+1)$-ary and E, $<$ are binary. For any subset $S \subseteq V$, the τ-structure $\mathscr{D}(G,S) = (D_G, R^{\mathscr{D}(G,S)}, E^{\mathscr{D}(G,S)}, <^{\mathscr{D}(G,S)})$ is defined by

- $D_G = V \times C$,
- $R^{\mathscr{D}(G,S)}$ is the set of all tuples $((u,a_1),\dots,(u,a_{n+1}))$ in D_G so that either $u \notin S$ and $(a_1,\dots,a_{n+1}) \in R^+$, or $u \in S$ and $(a_1,\dots,a_{n+1}) \in R^-$,
- $E^{\mathscr{D}(G,S)}$ is the set of all pairs $((u,c_i),(v,c_j))$ and $((u,d_i),(v,d_j))$ in D_G^2 such that $(u,v) \in E$, $i = h_u(v)$, and $j = h_v(u)$,
- $(u,a) <^{\mathscr{D}(G,S)} (v,b)$ if and only if $u <^G v \quad$ or $\quad u = v \wedge a \prec b$

Unlike the CFI graphs, Hella's graph construction does not "twist" the actual edges. Instead, the twists are encoded in $R^{\mathscr{D}(G,S)}$. Notice that $<^{\mathscr{D}(G,S)}$ has width 2, as for every $(u,a) \in D_G$, there is exactly one $(u,b) \in D_G$ with neither $(u,a) <^{\mathscr{D}(G,S)} (u,b)$ nor $(u,b) <^{\mathscr{D}(G,S)} (u,a)$. We call (u,a), (u,b) an *incomparable pair*. Hella proves the following:

Lemma 5.3. *Let $S,T \subseteq V$. The structures $\mathscr{D}(G,S)$ and $\mathscr{D}(G,T)$ are isomorphic if and only if $|S|$ and $|T|$ are of the same parity.*

Thus, there are exactly two non-isomorphic structures $\mathscr{D}(G,S)$ for any graph G. Let $\mathscr{A}(G) = \mathscr{D}(G,\emptyset)$ and $\mathscr{B}(G) = \mathscr{D}(G,\{u\})$ for some $u \in V$.

Theorem 5.4 (Hella 1996)**.** *For any $n \in \mathbb{N}$, there is a family of connected $(n+1)$-regular graphs G_k with $|G_k| = O(k^2)$ so that for any $\mathsf{L}^{\omega}_{\infty\omega}(Q_n)$-sentence φ there is $k_\varphi \in \mathbb{N}$ such that $\mathscr{A}(G_k) \models \varphi \Leftrightarrow \mathscr{B}(G_k) \models \varphi$ for all $k \geq k_\varphi$.*

$\mathscr{A}(G)$ and $\mathscr{B}(G)$ can actually be distinguished by a PTIME computation, so this theorem implies that $\mathsf{FP}(Q_n)$ does not capture PTIME for any $n \in \mathbb{N}$. We show here that $\mathscr{A}(G)$ and $\mathscr{B}(G)$ can be distinguished by a linear system of arity $n+1$.

Let $\mathscr{T} = (V \times C, R^{\mathscr{T}}, E^{\mathscr{T}}, <^{\mathscr{T}})$ be a τ-structure. Then let $\mathfrak{S}$ be the linear system over GF_2 with variables $x_{(u,a)}$ for every $(u,a) \in V \times C$ and the following equations:

- for every incomparable (u,a), (u,b): $x_{(u,a)} + x_{(u,b)} = 1$,
- for each $((u,a),(v,b)) \in E^{\mathscr{T}}$: $x_{(u,a)} + x_{(v,b)} = 0$,
- for every $(n+1)$-tuple $((u,a_1),\dots,(u,a_{n+1})) \in R^{\mathscr{T}}$: $x_{(u,a_1)} + \dots + x_{(u,a_{n+1})} = 0$.

$\mathfrak{S}$ can be defined by formulas $\varphi(\vec{x},y)$ and $\beta(\vec{x})$ over $\mathscr{T}$ with $|\vec{x}| = n$ as follows. For every n-tuple $\vec{a} := ((u,a_1),\dots,(u,a_n))$ with $a_1 \prec \dots \prec a_n$, there is at most one element (u,a_{n+1}) such that $((u,a_1),\dots,(u,a_{n+1})) \in R^{\mathscr{T}}$. Let $\varphi(\vec{a},y)$ express the equation $x_{(u,a_1)} + \dots + x_{(u,a_{n+1})} = 0$, i.e., $\varphi[\vec{a},b] = 1 \Leftrightarrow b = (u,a_i)$ for some $1 \leq i \leq n+1$. The equations for edges and incomparable pairs can be defined at row indexes $v_1 \dots v_n$ for which $v_1 = \dots = v_{n-1}$. Defining the vector $\beta(\vec{x})$ is trivial.

In order to prove Theorem 5.1, we have to decide whether $\mathfrak{S}$ has a solution without increasing the arity of the matrix defined by φ. For this, we need the following lemma from linear algebra, which is easy to prove.

Lemma 5.5. *Let A be a matrix that does not have full column rank. Then the linear system $A\mathbf{x} = \mathbf{b}$ is solvable if and only if for all columns $\mathbf{c}$ of A, adding $\mathbf{b}$ to $\mathbf{c}$ does not increase the rank of A.*

As G is an $(n+1)$-regular graph with $n \geq 1$, G contains a cycle. Let H be such a cycle and let $J = \{(u,a) \in D_G \mid u \in H \text{ has a neighbor } v \text{ in } H \text{ s.t. for some } b : ((u,a),(v,b)) \in E^{\mathscr{T}}\}$. Then it is readily verified that on every row, the sum of the entries in columns indexed by J is 0.

Thus, the matrix defined by φ does not have full column rank and (using Lemma 5.5) there is a sentence $\Psi_{\mathfrak{S}} \in \mathsf{FO}(\mathscr{R}_{n+1})$ defining the solvability of $\mathfrak{S}$. It is easy to verify that $\mathfrak{S}$ is solvable over $\mathscr{A}(G)$ by setting

$$x_{(u,a)} = \begin{cases} 0 & \text{if } a = c_i \text{ for some } i, \\ 1 & \text{if } a = d_i \text{ for some } i. \end{cases}$$

Since $\Psi_{\mathfrak{S}}$ is invariant under isomorphism, $\mathfrak{S}$ is solvable over $\mathscr{D}(G,S)$ whenever $|S|$ is even. The following lemma allows us to complete the proof of Theorem 5.1.

Lemma 5.6. *Let $\mathscr{D}(G,S)$ be as above. Any solution $\vec{x}$ of $\mathfrak{S}$ over $\mathscr{D}(G,S)$ induces an isomorphism $\iota : \mathscr{D}(G,S) \to \mathscr{A}(G)$ by letting*

$$\iota(u,a_i) = \begin{cases} (u,c_i) & \text{if } x_{(u,a_i)} = 0, \\ (u,d_i) & \text{if } x_{(u,a_i)} = 1. \end{cases}$$

6. Linear Algebra in FP+C

We have seen that fixed-point logic with counting is not expressive enough to describe the rank of matrices over finite fields. It is still, however, an open problem whether or not matrix *determinant* is definable in this logic. It has been observed by Rossman that Csanky's algorithm [7] for computing the characteristic polynomial (and thereby, the determinant) of a matrix over any commutative ring of characteristic zero is expressible in the logic of choiceless polynomial time with counting. Blass and Gurevich [2] used this observation to show that the same logic can also express the determinant of any definable matrix over a finite field.

In this section we strengthen the above result by showing that Le Verrier's method for finding the coefficients of the characteristic polynomial of a matrix (cf. [11]), which is the main building block of Csanky's algorithm, can already be expressed in FP+C for both integer and rational matrices, as well as matrices over finite fields. Le Verrier's method calculates M's characteristic polynomial

$$\chi_M(x) := \det(xI - M) = x^n - p_1 x^{n-1} + p_2 x^{n-2} - \cdots \pm p_n$$

by solving the linear system $A\mathbf{x} = \mathbf{b}$ for $\mathbf{x} = (p_n, \ldots, p_1)^t$ with $\mathbf{b} = (\pm\frac{s_n}{n}, \mp\frac{s_{n-1}}{n-1}, \ldots, s_1)^t$, $s_k := \operatorname{tr}(M^k)$ and

$$A = \begin{pmatrix} 1 & -\frac{s_1}{n} & \ldots & \pm\frac{s_{n-2}}{n} & \mp\frac{s_{n-1}}{n} \\ 0 & 1 & \ldots & \mp\frac{s_{n-3}}{n-1} & \pm\frac{s_{n-2}}{n-1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \ldots & 1 & -\frac{s_1}{2} \\ 0 & 0 & \ldots & 0 & 1 \end{pmatrix}. \qquad (1)$$

Characteristic polynomial over $\mathbb{Q}$ and $\mathbb{Z}$. To represent matrices with values from the ring of integers and the field of rationals as finite structures, we follow the convention of Blass et al. [3] and write matrix entries in binary notation. In this way, an integer matrix $A = (a_{ij})$ can be thought of as a ternary relation $A(i,j,k)$ defined by "the coefficient of 2^k in the binary expansion of a_{ij} is 1". Matrices with rational entries can be treated similarly by handling the numerators and denominators separately.

Matrices as above can be represented using a framework similar to the one we introduced in section 2. Here we consider rational matrices; integer matrices can be defined by setting all denominators to 1. Let $\mathscr{A}$ be a τ-structure and let $\eta_n(\vec{x},\vec{y},z)$ and $\eta_d(\vec{x},\vec{y},z)$ be FP+C$[\tau]$-formulas, where $\vec{x}$ and $\vec{y}$ are universe variables and z is a number variable. The determinant and the characteristic polynomial are defined only for square matrices, so we assume $|\vec{x}| = |\vec{y}|$. The role of z is to index the binary expansion of the matrix elements over the number sort, bounded by some numeric term t. Let M denote the square rational matrix in binary representation of bit length $t^{\mathscr{A}}$ defined by η_n and η_d, which separately define the numerators and denominators of M, respectively. Below we will sketch a proof of the following theorem.

Theorem 6.1. *There are numeric* FP+C$[\tau]$*-terms $\theta^n_{char}(x,y)$ and $\theta^d_{char}(x,y)$, where x and y are number variables, which satisfy:*

- *$\theta^n_{char}[k,i] = d$ iff the i-th bit of the* numerator *of the coefficient of x^k in the characteristic polynomial $p_M(x)$ of M over $\mathbb{Q}$ is d;*
- *$\theta^d_{char}[k,i] = d$ iff the i-th bit of the* denominator *of the coefficient of x^k in the characteristic polynomial $p_M(x)$ of M over $\mathbb{Q}$ is d.*

Recall that for any $n \times n$ matrix M, the constant term of $\chi_M(x)$ takes value $(-1)^n \det(M)$.

Corollary 6.2. *There are numeric* FP+C$[\tau]$*-terms $\theta^n_{det}(x)$ and $\theta^d_{det}(x)$, where x is a number variable, which satisfy:*

- *$\theta^n_{det}[i] = d$ iff the i-th bit of the* numerator *of the determinant of M over $\mathbb{Q}$ is d;*
- *$\theta^d_{det}[i] = d$ iff the i-th bit of the* denominator *of the determinant of M over $\mathbb{Q}$ is d.*

In [3], Blass et al. describe various matrix properties and operations that are definable in FP+C (see also [8] for a good summary). These constructions can be adapted to work for matrices over $\mathbb{Z}$ and $\mathbb{Q}$ in the representation we chose above. This gives us the following.

- *Matrix product:* There are terms and formulas of FP+C that define the matrix product MN for any definable matrices M, N.
- *Matrix powers:* There are terms and formulas of FP+C that define the matrix M^k for any definable $n \times n$ matrix M and $k \in [n]$.
- *Trace:* There are terms and formulas of FP+C that define $\mathrm{tr}(M)$ for any definable matrix M.

We also need to justify that we can define the terms $s_k = \mathrm{tr}(M^k)$ over M for each $k \in [n]$. Let m be the maximum absolute value of any integer appearing either as a numerator or denominator of an element of M; clearly $\log_2(m) \leq t$ where t is the number term bounding the binary indices in M. The following bound is not hard to establish.

Lemma 6.3. *$s_k \leq (nm)^k$ for all $k \in [n]$.*

This shows that we have to consider binary indices up to

$$\lceil \log_2((nm)^n) \rceil \leq n \lceil \log_2(n) + \log_2(m) \rceil \leq (n+t)^2.$$

Hence we can define all the required index positions by using quantification over the number sort bounded by the definable numeric term $\kappa := (n+t)^2$.

By now it is clear that we can define the linear system $A\mathbf{x} = \mathbf{b}$ from equation (1) in FP+C. We can express any polynomial-time property of this linear system, because the matrices A and $\mathbf{b}$ are defined on an *ordered* definable subset of the number sort. In particular, we can express Gaussian elimination as a fixed-point formula, use that to solve the system for $\mathbf{x}$ and hence obtain the coefficients of $\chi_M(x)$.

Characteristic polynomial over finite fields. Now consider a τ-structure $\mathscr{A}$ of size n and numeric FP+C$[\tau]$-terms p and $\eta(\vec{x},\vec{y})$, where $\vec{x}$, $\vec{y}$ are universe variables with $|\vec{x}| = |\vec{y}|$. Let M_p denote the matrix modulo p defined by $\eta(\vec{x},\vec{y})$, as in section 2. Le Verrier's method involves division by integers up to n, so it cannot be applied directly over the prime field GF_p. Instead, we map the input to a matrix $\hat{M}$ over the ring of integers, apply Le Verrier's method to $\hat{M}$ over $\mathbb{Z}$, and then reduce the result modulo p to get the specification of the characteristic polynomial over GF_p. This approach was shown to work by Blass and Gurevich in [2].

We assume that the elements of M_p are initially given by numeric terms in the range $[0, p-1]$. The binary expansion of each element of M_p can be defined in FP+C using $p \geq \log p$ bits, thereby defining the matrix $\hat{M}$ (cf. [21]). Likewise, the binary representations $\theta^n_{\mathrm{char}}(x,y)$ and $\theta^n_{\mathrm{det}}(x)$ can be reduced modulo p to an integer in $[0, p-1]$, from which we recover a term using a counting quantifier. Thus we get the following corollary.

Corollary 6.4. *There are numeric FP+C$[\tau]$-terms θ_{det} and $\theta_{char}(x)$, where x is a number variable, which satisfy:*

- *$\theta_{det} = d$ iff the determinant of M_p over GF_p is d;*
- *$\theta_{char}[k] = d$ iff the coefficient of x^k in the characteristic polynomial $p_{M_p}(x)$ of M_p over GF_p is d.*

As a consequence of this result, we can use the adjugate rule to construct a numeric term denoting the inverse to a given invertible matrix $\eta(\vec{x},\vec{y})$ over GF_p.

Note that our framework for representing finite matrices by definable terms and formulas allows us to handle matrices in any of the relational vocabularies for prime fields defined by Blass et al. in [3]. More generally, we can show that the characteristic polynomial of a matrix over any finite field, not necessarily prime, can be defined in FP+C. We defer the details of this construction to the journal version of this paper.

Defining matrix rank over $\mathbb{Q}$ in FP+C. Let M be a matrix over $\mathbb{Q}$, not necessarily square, and let $M^* = M^t M$, where M^t denotes the transpose of M. It can be shown that $\mathrm{rank}\, M = \mathrm{rank}\, M^* = \mathrm{rank}\,(M^*)^2$. With this in mind, the following lemma (cf. [23]) tells us that the rank of M can be inferred directly from its characteristic polynomial.

Lemma 6.5. *Let M be an $n \times n$ matrix over any field. If $\mathrm{rank}\, M = \mathrm{rank}\, M^2$, then $\mathrm{rank}\, M = n-k$ where x^k is the highest power of x that divides the characteristic polynomial $\chi_M(x)$.*

Now let M be presented by definable terms and formulas as before. As all the computation steps described above can be carried out in FP+C, we get the following result.

Corollary 6.6. *There is a numeric FP+C$[\tau]$-term θ_{rank} which satisfies: $\theta_{rank} = r$ iff the rank of M over $\mathbb{Q}$ is r.*

7. Discussion

We have introduced logics with rank operators and demonstrated their surprising expressive power. This work raises many interesting questions. A first natural question that arises is what are the limits of the expressive power of FP+rk? Can we demonstrate a polynomial-time property that is not definable in this logic? Establishing this requires not only finding suitable candidate properties but also developing methods for proving inexpressiblity in FP+rk. Our results in Section 5 give one approach, by translating the operators into Lindström quantifiers, but perhaps, other, game-based methods need to be developed.

An easier testing ground for such methods would be to establish inexpressiblity results for FO+rk. For instance, can we show that the *alternating transitive closure* (ATC) query is not definable in this logic? If it were definable in $\mathsf{FO+rk}_2$, as deterministic transitive closure is, then it would imply that PTIME $= \oplus$L. It might also be possible to show that ATC is not in FO+rk without complexity-theoretic assumptions. This would establish (the expected result) that FO+rk is strictly weaker than FP+rk.

It would also be interesting to investigate the relationship of FP+rk with other logics that have been proposed which extend FP+C while remaining inside poynomial time. In particular, how does FP+rk compare with Choiceless Polynomial Time with Counting ($\tilde{\mathsf{C}}$PT(Card)) as introduced by Blass et al. [3]? It has been shown [9] that $\tilde{\mathsf{C}}$PT(Card) can express the CFI property where the CFI graph $\mathscr{G}^T$ is constructed from an *ordered* graph G, while our construction in Section 3 works even with unordered G. It remains an open question whether the rank of a matrix can be computed or the solvability of systems of linear equations determined in $\tilde{\mathsf{C}}$PT(Card). Indeed, an inclusion either way between FP+rk and $\tilde{\mathsf{C}}$PT(Card) is unknown.

A more positive direction to investigate would be to use FP+rk to express some natural properties that are not known to be in FP+C. One example is the problem of determining whether a given graph has a perfect matching. It is known [3] that there is a sentence of FP+C that defines this property on bipartite graphs, but it is an open question whether or not there is one that defines it on all graphs. Another interesting problem to consider is isomorphism on bounded-degree graphs, which is known to be decidable in polynomial-time by a result of Luks [25] but is also known not to be in FP+C as a consequence of the CFI construction. Indeed, classes of graphs of bounded-degree are an interesting case where we know, in principle, that there is a logic that captures PTIME, since there is a polynomial-time canonization algorithm, but we do not have a natural logic for that purpose. Could it be that FP+rk captures PTIME on such classes?

References

[1] A. Atserias, A. Bulatov, and A. Dawar. Affine systems of equations and counting infinitary logic. In *ICALP*, pages 558–570, 2007.

[2] A. Blass and Y. Gurevich. A quick update on open problems in Blass-Gurevich-Shelah's article 'On polynomial time computations over unordered structures', 2005.

[3] A. Blass, Y. Gurevich, and S. Shelah. On polynomial time computations over unordered structures. *Journal of Symbolic Logic*, 67:1093–1125, 2002.

[4] G. Buntrock, C. Damm, U. Hertrampf, and C. Meinel. Structure and importance of logspace-MOD class. *Math. Syst. Theory*, 25:223–237, 1992.

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

[6] A. Chandra and D. Harel. Structure and complexity of relational queries. *Journal of Computer and System Sciences*, 25:99–128, 1982.

[7] L. Csanky. Fast parallel matrix inversion algorithms. *SIAM J. Comput.*, 5:618–623, 1976.

[8] A. Dawar. On the descriptive complexity of linear algebra. In *WoLLIC '08: Proceedings of the 15th international workshop on Logic, Language, Information and Computation*, pages 17–25. Springer-Verlag, 2008.

[9] A. Dawar, D. Richerby, and B. Rossman. Choiceless polynomial time, counting and the Cai-Fürer-Immerman graphs. *Annals of Pure and Applied Logic*, 152:31–50, 2008.

[10] H.-D. Ebbinghaus and J. Flum. *Finite model theory*. Springer Verlag, Berlin, Germany, 2nd edition, 1999.

[11] D. K. Faddeev and V. N. Faddeeva. *Computational Methods of Linear Algebra,(Translated by RC Williams)*. Freeman, San Francisco, 1963.

[12] M. Grohe. Fixed-point logics on planar graphs. In *Proceedings of the 13th IEEE Symposium on Logic in Computer Science*, pages 6–15, 1998.

[13] M. Grohe. Definable tree decompositions. In *Proceedings of the 23rd IEEE Symposium on Logic in Computer Science*, pages 406–417, 2008.

[14] M. Grohe and J. Mariño. Definability and descriptive complexity on databases of bounded tree-width. In C. Beeri and P. Buneman, editors, *Proceedings of the 7th International Conference on Database Theory*, volume 1540 of *Lecture Notes in Computer Science*, pages 70–82. Springer-Verlag, 1999.

[15] Y. Gurevich. Logic and the challenge of computer science. In E. Börger, editor, *Current trends in theoretical computer science*, pages 1–57. Computer Science Press, 1988.

[16] Y. Gurevich and S. Shelah. On finite rigid structures. *Journal of Symbolic Logic*, 61:61–549, 1996.

[17] L. Hella. Logical hierarchies in PTIME. *Inf. Comput.*, 129(1):1–19, 1996.

[18] L. Hella, Ph.G. Kolaitis, and K. Luosto. Almost everywhere equivalence of logics in finite model theory. *Bulletin of Symbolic Logic*, 2:422–443, 1996.

[19] U. Hertrampf, S. Reith, and H. Vollmer. A note on closure properties of logspace mod classes. *Information Processing Letters*, 75:91–93, 2000.

[20] N. Immerman. Expressibility as a complexity measure: results and directions. In *Proceedings of the 2nd IEEE Symposium on Structure in Complexity Theory*, pages 194–202, 1987.

[21] N. Immerman. *Descriptive Complexity*. Springer, 1999.

[22] N. Immerman and E. Lander. Describing graphs: A first-order approach to graph canonization. In A. Selman, editor, *Complexity theory retrospective*, pages 59–81. Springer-Verlag, 1990.

[23] D. C. Kozen. *The Design and Analysis of Algorithms*. Springer, 1992.

[24] P. Lindström. First order predicate logic with generalized quantifiers. *Theoria*, 32:186–195, 1966.

[25] E. M. Luks. Isomorphism of graphs of bounded valence can be tested in polynomial time. *Journal of Computer and System Sciences*, 25:42–65, 1982.