Rank Logic is Dead, Long Live Rank Logic!

Erich Grädel and Wied Pakusa

Mathematical Foundations of Computer Science, RWTH Aachen University, Germany
{graedel,pakusa}@logic.rwth-aachen.de

Abstract

Motivated by the search for a logic for polynomial time, we study rank logic (FPR) which extends fixed-point logic with counting (FPC) by operators that determine the rank of matrices over finite fields. While FPR can express most of the known queries that separate FPC from PTIME, nearly nothing was known about the limitations of its expressive power.

In our first main result we show that the extensions of FPC by rank operators over different prime fields are incomparable. This solves an open question posed by Dawar and Holm and also implies that rank logic, in its original definition with a distinct rank operator for every field, fails to capture polynomial time. In particular we show that the variant of rank logic FPR* with an operator that uniformly expresses the matrix rank over finite fields is more expressive than FPR.

One important step in our proof is to consider solvability logic FPS which is the analogous extension of FPC by quantifiers which express the solvability problem for linear equation systems over finite fields. Solvability logic can easily be embedded into rank logic, but it is open whether it is a strict fragment. In our second main result we give a partial answer to this question: in the absence of counting, rank operators are strictly more expressive than solvability quantifiers.

1998 ACM Subject Classification F.1.3 Complexity Measures and Classes, F.4.1 Mathematical Logic

Keywords and phrases logic, descriptive complexity, polynomial time, rank logic

Digital Object Identifier 10.4230/LIPIcs.CSL.2015.390

1 Introduction

"Le roi est mort, vive le roi!" has been the traditional proclamation, in France and other countries, to announce not only the death of the monarch, but also the immediate installment of his successor on the throne. The purpose of this paper is to kill the rank logic FPR, in the form in which it has been proposed in [7], as a candidate for a logic for PTIME. The logic FPR extends fixed-point logic by operators rk_p (for every prime p) which compute the rank of definable matrices over the prime field $\mathbb{F}_p$ with p elements. Although rank logic is well-motivated, as a logic that strictly extends fixed-point logic with counting by the ability to express important properties of linear algebra, most notably the solvability of linear equation systems over finite fields, our results show that the choice of having a separate rank operator for every prime p leads to a significant deficiency of the logic. Indeed, it follows from our main theorem that even the uniform rank problem, of computing the rank of a given matrix over an arbitrary prime, cannot be expressed in FPR and thus separates FPR from PTIME. This also reveals that a more general variant of rank logic, which has already been proposed in [14, 15, 17] and which is based on a rank operator that takes not only the matrix but also the prime p as part of the input, is indeed strictly more powerful than FPR. Our result thus installs this new rank logic, denoted FPR*, as the rightful and distinctly more powerful successor of FPR as a potential candidate for a logic for PTIME. A full version of this paper can be found at [11].

24th EACSL Annual Conference on Computer Science Logic (CSL 2015).
Editor: Stephan Kreutzer; pp. 390–404
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

A logic for polynomial time. The question whether there is a logic that expresses precisely the polynomial-time properties of finite structures is an important challenge in the field of finite model theory [10, 12]. The logic of reference for this quest is fixed-point logic with counting (FPC) which captures polynomial time on many interesting classes of structures and which is strong enough to express many of the fundamental techniques which are used in polynomial-time algorithms [5]. Although it has been known for more than twenty years that FPC fails to capture PTIME in general, by the fundamental CFI-construction due to Cai, Fürer, and Immerman [4], we still do not know many properties of finite structures that provably separate FPC from PTIME. The two main sources of such problems are tractable cases of the graph isomorphism problem and queries from the field of linear algebra. First of all, the CFI-construction shows that FPC cannot define the isomorphism problem on graphs with bounded degree *and* bounded colour class size whereas the isomorphism problem is known to be tractable on all classes of graphs with bounded degree *or* bounded colour class size. Secondly, Atserias, Bulatov and Dawar [2] proved that FPC cannot express the solvability of linear equation systems over any finite Abelian group. It follows, that also other problems from the field of linear algebra are not definable in FPC. Interestingly, also the CFI-query can be formulated as linear equation system over $\mathbb{F}_2$ [7].

Rank logic. This latter observation motivated Dawar, Grohe, Holm and Laubner [7] to introduce *rank logic* (FPR) which is the extension of FPC by operators for the rank of definable matrices over prime fields $\mathbb{F}_p$. To illustrate the idea of rank logic, let $\varphi(x, y)$ be a formula (of FPC, say) which defines a binary relation $\varphi^{\mathfrak{A}} \subseteq A \times A$ in an input structure $\mathfrak{A}$. We identify the relation $\varphi^{\mathfrak{A}}$ with the associated adjacency matrix

$$M_\varphi^{\mathfrak{A}} : A \times A \to \{0,1\}, (a,b) \mapsto \begin{cases} 1, & \text{if } (a,b) \in \varphi^{\mathfrak{A}} \\ 0, & \text{if } (a,b) \notin \varphi^{\mathfrak{A}}. \end{cases}$$

In this sense, the formula φ defines in every structure $\mathfrak{A}$ a matrix $M_\varphi^{\mathfrak{A}}$ with entries in $\{0,1\} \subseteq \mathbb{F}_p$. Now, rank logic FPR provides for every prime $p \in \mathbb{P}$ a *rank operator* rk_p which can be used to form a *rank term* $[\mathsf{rk}_p\, \varphi(x,y)]$ whose value in an input structure $\mathfrak{A}$ is the matrix rank of M_φ over $\mathbb{F}_p$ (we remark that rank logic also allows to express the rank of matrices which are indexed by tuples of elements; the precise definition is given in Section 2).

It turns out that rank operators have quite surprising expressive power. For example, they can define the transitive closure of symmetric relations, they can count the number of paths in DAGs modulo p and they can express the solvability of linear equation systems over finite fields (recall that a linear equation system $M \cdot \vec{x} = \vec{b}$ is solvable if, and only if, $\mathsf{rk}(M) = \mathsf{rk}(M \,|\, \vec{b})$) [7]. Furthermore, rank operators can be used to define the isomorphism problem on various classes of structures on which the Weisfeiler-Lehman method (and thus fixed-point logic with counting) fails, e.g. classes of C(ai)-F(ürer)-I(mmerman) graphs [4, 7] and multipedes [13, 14]. The common idea of these isomorphism procedures is to reduce the isomorphism problem of structures to a suitable linear equation system over a finite field. More generally, by a recent result (which is mainly concerned with another candidate of a logic for polynomial time [1]), it follows that FPR captures polynomial time on certain classes of structures of bounded colour class size. In particular, this holds for the class of all structures of colour class size two (to which CFI-graphs and multipedes belong).

While these results clearly show the high potential of rank logic, almost nothing has been known about its limitations. For instance, it has remained open whether rank logic suffices to capture polynomial time, whether rank operators can simulate fixed-point inductions [7] and also whether rank logic can define closely related problems from linear algebra (such as the solvability of linear equations over finite *rings* rather than fields [6]). One particular intriguing question is whether rank operators over different prime fields can simulate each other. In other words: is it possible to reduce the problem of determining the rank of a matrix over $\mathbb{F}_p$ (within fixed-point logic with counting) to the problem of determining the rank of a matrix over $\mathbb{F}_q$ (where p, q are distinct primes)? To attack this problem, Dawar and Holm [8, 14] developed a powerful toolkit of so called *partition games* of which one variant (so called *matrix-equivalence games*) precisely characterises the expressive power of infinitary logic extended by rank quantifiers. By using these games, Holm [14] was able to give a negative answer to the above question for the restricted case of rank operators of dimension one.

In this paper we propose a different method, based on exploiting symmetries rather than game theoretic arguments, to prove new lower bounds for logics with rank operators. In our main result (Theorem 3) we prove that for every prime q there exists a class of structures $\mathcal{K}_q$ on which FPC fails to capture polynomial time and on which rank operators over *every* prime field $\mathbb{F}_p$, $p \neq q$ can be simulated in FPC. On the other hand, rank operators over $\mathbb{F}_q$ can be used to canonise structures in $\mathcal{K}_q$ which means that the extension of fixed-point logic by rk_q-operators captures polynomial time on $\mathcal{K}_q$. From this result we can easily extract the following answers to the open questions outlined above.

(a) Rank logic (as defined in [7]) fails to capture polynomial time (Theorem 2).

(b) The extensions of fixed-point logic by rank operators over different prime fields are incomparable (Theorem 1), cf. [14, 8, 15].

We obtain these classes of structures $\mathcal{K}_q$ by generalising the well-known construction of Cai, Fürer and Immerman [4]. It has been observed that their construction actually is a clever way of encoding a linear equation system over $\mathbb{F}_2$ into an appropriate graph structure (see e.g. [2, 7, 14, 15]). Intuitively, each gadget in the CFI-construction can be seen as an equation (or, equivalently, as a circuit gate) which counts the number of transpositions of adjacent edges modulo two, and the CFI-query is to decide whether the total number of such transpositions is even or odd. Knowing this, it is very natural to ask whether this idea can be generalised to encode linear equation systems over arbitrary finite fields or, more generally, equation systems over arbitrary (Abelian) groups.

In [18], in order to obtain hardness results for the graph isomorphism problem, Torán followed this idea and established a graph construction which simulates mod k-counting gates for all $k \geq 2$. Moreover, in order to separate the fragments of rank logic by operators over different prime fields, Holm presented in [14] an even more general kind of construction which allows the representation of equations over *every* Abelian group G. In fact, we obtain the classes $\mathcal{K}_q$ essentially by using his construction for the special case where $G = \mathbb{F}_q$.

Solvability logic. One important step in our proof is to consider *solvability logic* FPS which is the extension of FPC by quantifiers which can express the solvability of linear equation systems over finite fields (so called *solvability quantifiers*, see [6, 17]). Obviously the logic FPS can easily be embedded into rank logic (as rank operators can be used to solve linear equation systems), but it remains open whether the inclusion FPS $\leq$ FPR is strict. To prove our main result outlined above we show that over certain classes of structures the logics FPS and FPR have precisely the same expressive power. In a more general context this might give some evidence that in the framework of fixed-point logic with counting rank operators can be simulated by solvability quantifiers. On the other hand we show in Section 4 that the extension of first-order logic (without counting) by solvability quantifiers is strictly weaker than the respective extension by rank operators. This last result thus separates solvability quantifiers and rank operators in the absence of counting.

2 Logics with linear-algebraic operators

By $\mathcal{S}(\tau)$ we denote the class of all *finite, relational* structures of signature τ. We assume that the reader is familiar with *first-order logic* (FO) and *inflationary fixed-point logic* (FP). For details see [9, 10]. We write $\mathbb{P}$ for the set of primes and denote the prime field with p elements by $\mathbb{F}_p$. We consider matrices and vectors over *unordered* index sets. Formally, if I and J are non-empty sets, then an $I \times J$-matrix M over $\mathbb{F}_p$ is a mapping $M : I \times J \to \mathbb{F}_p$ and an I-vector $\vec{v}$ over $\mathbb{F}_p$ is a mapping $\vec{v} : I \mapsto \mathbb{F}_p$.

A *preorder* $\preceq$ on A is a reflexive, transitive and total binary relation. It induces a linear order on the classes of the associated equivalence relation $x \sim y := (x \preceq y \wedge y \preceq x)$. We write $A = C_0 \preceq \cdots \preceq C_{n-1}$ to denote the decomposition of A into $\sim$-classes C_i which are linearly ordered by $\preceq$ as indicated. We denote by $\mathrm{Aut}(\mathfrak{A}) \leq \mathrm{Sym}(A)$ the automorphism group of a structure $\mathfrak{A}$ as a subgroup of the symmetric group acting on the set A. We assume that the reader is familiar with the basic notions from (linear) algebra.

We recall the definitions of *first-order logic with counting* FOC and *(inflationary) fixed-point logic with counting* FPC. Formulas of FOC and FPC are evaluated over the *two-sorted extension* of an input structure by a copy of the arithmetic. Following [7] we let $\mathfrak{A}^{\#}$ denote the two-sorted extension of a τ-structure $\mathfrak{A} = (A, R_1, \dots, R_k)$ by the arithmetic $\mathfrak{N} = (\mathbb{N}, +, \cdot, 0, 1)$, i.e. the two-sorted structure $\mathfrak{A}^{\#} = (A, R_1, \dots, R_k, \mathbb{N}, +, \cdot, 0, 1)$ where the universe of the first sort (also referred to as *vertex sort*) is A and the universe of the second sort (also referred to as *number sort* or *counting sort*) is $\mathbb{N}$.

As usual for the two-sorted setting we have typed first-order variables, where Latin letters $x, y, z, \dots$ stand for variables that range over vertices, and Greek letters $\nu, \mu, \dots$ for variables ranging over numbers. For second-order variables we allow mixed types, i.e. a relation symbol R of type $(k, \ell) \in \mathbb{N} \times \mathbb{N}$ stands for a relation $R \subseteq A^k \times \mathbb{N}^\ell$. Of course, already first-order logic over such two-sorted extensions is undecidable. To obtain logics whose data complexity is in polynomial time we restrict the quantification over the number sort by a numeric term t, i.e. $Q\nu \leq t.\varphi$ where $Q \in \{\exists, \forall\}$ and where t is a closed *numeric* term. Similarly, for fixed-point logic FP we bound the numeric components of fixed-point variables R of type (k, ℓ) in all fixed-point definitions $\left[\mathsf{ifp}\, R\bar{x}\bar{\nu} \leq \bar{t} \,.\, (\varphi(\bar{x}, \bar{\nu}))\right](\bar{x}, \bar{\nu})$ by a tuple of closed numeric terms $\bar{t} = (t_1, \dots, t_\ell)$ where each t_i bounds the range of the variable ν_i in the tuple $\bar{\nu}$. For the logics which we consider here the value of such numeric terms (and thus the range of all quantifiers over the number sort) is polynomially bounded in the size of the input structure. Together with the standard argument that inflationary fixed-points can be evaluated in polynomial time and the fact that the matrix rank over any field can be determined in polynomial time (for example by the method of Gaussian elimination), this ensures that all the logics which we introduce in the following have polynomial-time data complexity.

Let $\bar{x}\bar{\nu}$ be a mixed tuple of variables and let $\bar{t}$ be a tuple of closed numeric terms which bounds the range of the numeric variables in $\bar{\nu}$. For a formula φ we define a *counting term* $s = [\#\bar{x}\bar{\nu} \leq \bar{t} \,.\, \varphi]$ whose value $s^{\mathfrak{A}} \in \mathbb{N}$ in a structure $\mathfrak{A}$ corresponds to the number of tuples $(\bar{a}, \bar{n}) \in A^k \times \mathbb{N}^\ell$ such that $\mathfrak{A} \models \varphi(\bar{a}, \bar{n})$ and $n_i \leq t_i^{\mathfrak{A}}$ where $k = |\bar{x}|$ and $\ell = |\bar{\nu}|$ (to be precise, we should write $\mathfrak{A}^{\#}$ instead of $\mathfrak{A}$, but we usually omit the superscript for the sake of better readability). We then define *first-order logic with counting* FOC as the extension of (the above described two-sorted variant of) FO by counting terms. Similarly, by adding counting terms to the logic FP we obtain *(inflationary) fixed-point logic with counting* FPC.

Rank operators. Let $\Theta(\bar{x}\bar{\nu}, \bar{y}\bar{\mu})$ be a numeric term and let $\bar{t}$ and $\bar{s}$ be tuples of closed numeric terms which bound the range of the numeric variables in $\bar{\nu}$ and $\bar{\mu}$, respectively. Given a structure $\mathfrak{A}$ we define $\mathbb{N}^{\leq \bar{t}} := \{\bar{n} \in \mathbb{N}^{|\bar{\nu}|} : n_i \leq t_i^{\mathfrak{A}}\}$. The set $\mathbb{N}^{\leq \bar{s}} \subset \mathbb{N}^{|\bar{\mu}|}$ is defined analogously. The term Θ together with $\bar{t}$ and $\bar{s}$ defines in the structure $\mathfrak{A}$ for $I := A^{|\bar{x}|} \times \mathbb{N}^{\leq \bar{t}}$ and $J := A^{|\bar{y}|} \times \mathbb{N}^{\leq \bar{s}}$ the $I \times J$-matrix M_Θ with values in $\mathbb{N}$ given as $M_\Theta(\bar{a}\bar{n}, \bar{b}\bar{m}) := \Theta^{\mathfrak{A}}(\bar{a}\bar{n}, \bar{b}\bar{m})$.

The *matrix rank operators* compute the rank of the matrix M_Θ over a prime field $\mathbb{F}_p$ for $p \in \mathbb{P}$. First, as in [7], we define for every prime p a matrix rank operator rk_p which allows us to construct a new numeric *rank term* $[\mathsf{rk}_p\,(\bar{x}\bar{\nu} \leq \bar{t}, \bar{y}\bar{\mu} \leq \bar{s})\,.\,\Theta]$ whose value in the structure $\mathfrak{A}$ is the rank of the matrix $(M_\Theta \bmod p)$ over $\mathbb{F}_p$. Secondly, we propose a uniform rank operator rk^* which takes the prime p as an additional input. Formally, with this rank operator rk^* we can construct a rank term $[\mathsf{rk}^*\,(\bar{x}\bar{\nu} \leq \bar{t}, \bar{y}\bar{\mu} \leq \bar{s}, \pi \leq r)\,.\,\Theta]$ where π is an additional free numeric variable whose range is bounded by some closed numeric term r. Given a structure $\mathfrak{A}$ and an assignment $\pi \mapsto p$ for some prime $p \leq r^{\mathfrak{A}}$, the value of the rank term is the matrix rank of $(M_\Theta \bmod p)$ considered as a matrix over $\mathbb{F}_p$ (if $\pi \mapsto n$ for $n \notin \mathbb{P}$, then the value is zero). The rank operator rk^* is a unification for the the family of separate rank operators $(\mathsf{rk}_p)_{p \in \mathbb{P}}$ and has been introduced in [14, 15, 17].

We define, for every set of primes $\Omega \subseteq \mathbb{P}$, the extension FOR_Ω of FOC and the extension FPR_Ω of FPC by matrix rank operators rk_p with $p \in \Omega$. For convenience, we let $\text{FOR} = \text{FOR}_{\mathbb{P}}$ and $\text{FPR} = \text{FPR}_{\mathbb{P}}$. Similarly, we denote by FPR^* the extension of FPC by the uniform rank operator rk^*. We remark, that rank operators can directly simulate counting terms. For example we have that $[\#x\,.\,\varphi(x)] = [\mathsf{rk}_p\,(x, y)\,.\,(x = y \wedge \varphi(x))]$. Hence, we could equivalently define the rank logics $\text{FOR}_\Omega, \text{FPR}_\Omega$ and FPR^* as the extensions of (the two-sorted variants of) FO and FP, respectively.

Solvability quantifiers. We next introduce extensions by quantifiers which directly express the solvability problem for linear equation systems over finite fields. Besides the applications in this paper, an additional advantage of such quantifiers is that they can be generalised for linear equation systems over more general classes of domains, like rings, for which no appropriate notion of matrix rank exists, cf. [6].

Let $\Omega \subseteq \mathbb{P}$ be a set of primes. Then the *solvability logic* FPS_Ω extends the syntax of FPC for every $p \in \Omega$ by the following formula creation rule for *solvability quantifiers* slv_p.

- Let $\varphi(\bar{x}\bar{\nu}, \bar{y}\bar{\mu}, \bar{z}) \in \text{FPS}_\Omega$ and let $\bar{t}$ and $\bar{s}$ be tuples of closed numeric terms with $|\bar{t}| = |\bar{\nu}|$ and $|\bar{s}| = |\bar{\mu}|$. Then also $\psi(\bar{z}) = (\mathsf{slv}_p\,\bar{x}\bar{\nu} \leq \bar{s}, \bar{y}\bar{\mu} \leq \bar{t})\varphi(\bar{x}\bar{\nu}, \bar{y}\bar{\mu}, \bar{z})$ is a formula of FPS_Ω.

The semantics of the formula $\psi(\bar{z})$ is defined similarly as for rank logic. More precisely, let $k = |\bar{x}|$ and $\ell = |\bar{y}|$. To a pair $(\mathfrak{A}, \bar{z} \mapsto \bar{c}) \in \mathcal{S}(\sigma, \bar{z})$ we associate the $I \times J$-matrix M_φ over $\{0, 1\} \subseteq \mathbb{F}_p$ where $I = A^k \times \mathbb{N}^{\leq \bar{s}}$ and $J = A^\ell \times \mathbb{N}^{\leq \bar{t}}$ and where for $\bar{a} \in I$ and $\bar{b} \in J$ we have $M_\varphi(\bar{a}, \bar{b}) = 1$ if, and only if, $\mathfrak{A} \models \varphi(\bar{a}, \bar{b}, \bar{c})$.

Let $\mathbb{1}$ be the I-identity vector over $\mathbb{F}_p$, i.e. $\mathbb{1}(\bar{a}) = 1$ for all $\bar{a} \in I$. Then M_φ and $\mathbb{1}$ determine the linear equation system $M_\varphi \cdot \vec{x} = \mathbb{1}$ over $\mathbb{F}_p$ where $\vec{x} = (x_j)_{j \in J}$ is a J-vector of variables x_j which range over $\mathbb{F}_p$. Finally, $\mathfrak{A} \models \psi(\bar{c})$ if, and only if, $M_\varphi \cdot \vec{x} = \mathbb{1}$ is solvable.

At first glance, the solvability quantifier seem to pose serious restrictions on the syntactic form of definable linear equation systems. Specifically, the coefficient matrix has to be a matrix over $\{0, 1\}$ and the vector of constants is fixed from outside. However, it is not hard to show that general linear equation systems can be brought into this kind of normal form by using quantifier-free first-order transformations (see Lemma 4.1 in [6]).

We write FPS to denote the logic $\text{FPS}_{\mathbb{P}}$ and FPS_p to denote the logic $\text{FPS}_{\{p\}}$ for $p \in \mathbb{P}$. Analogously to the definition of FPR^* we also consider a solvability quantifier slv which gets the prime p as an additional input and which can uniformly simulate all solvability quantifiers slv_p for $p \in \mathbb{P}$. Let FPS^* denote the extension of FPC by this uniform version of a solvability quantifier. The following inclusions follow from the definitions and the fact that rank operators can be used to define the solvability problem for linear equation systems.

$$\begin{array}{ccccccccc} \mathrm{FOR}_p & \leq & \mathrm{FPR}_p & \leq & \mathrm{FPR} & \leq & \mathrm{FPR}^* & \leq & \textsc{Ptime} \\ \underline{\vee} & & \underline{\vee} & & \underline{\vee} & & \underline{\vee} & & \\ \mathrm{FOS}_p & \leq & \mathrm{FPS}_p & \leq & \mathrm{FPS} & \leq & \mathrm{FPS}^* & & \\ \underline{\vee} & & \underline{\vee} & & & & & & \\ \mathrm{FO} & \leq & \mathrm{FPC} & & & & & & \end{array}$$

Finally we remark that, analogously to [7], we defined rank operators and solvability quantifiers for prime fields only. Of course, the definition can easily be generalised to cover all finite fields, i.e. also finite fields of prime power order. However, for the case of solvability quantifiers, Holm was able to prove in [14] that this does not alter the expressive power of the resulting logics since solvability quantifiers over a finite field $\mathbb{F}_q$ of prime power order $q = p^k$ can be simulated by solvability quantifiers over $\mathbb{F}_p$. In fact, a similar reduction can be achieved for rank operators which justifies to focus on rank operators and solvability quantifiers over prime fields.

3 Separation results over different classes of fields

In this section we separate the extensions FPS_Ω of fixed-point logic with counting by solvability quantifier for different sets of primes. Moreover, we transfer these results to the extensions FPR_Ω by rank operators.

▶ **Theorem 1.** *Let $\Omega \neq \Omega'$ be two sets of primes. Then $\mathrm{FPS}_\Omega \neq \mathrm{FPS}_{\Omega'}$ and $\mathrm{FPR}_\Omega \neq \mathrm{FPR}_{\Omega'}$.*

▶ **Theorem 2.** *Rank logic fails to capture polynomial time. We have $\mathrm{FPR} < \mathrm{FPR}^* \leq \textsc{Ptime}$.*

In fact, both theorems are simple consequences of our following main result.

▶ **Theorem 3.** *For every prime q there is a class of structures $\mathcal{K}_q$ such that*
(a) *$\mathrm{FPS}_\Omega = \mathrm{FPC}$ on $\mathcal{K}_q$ for every set of primes Ω with $q \notin \Omega$,*
(b) *$\mathrm{FPR}_\Omega = \mathrm{FPS}_\Omega$ on $\mathcal{K}_q$ for all sets of primes Ω,*
(c) *$\mathrm{FPC} < \textsc{Ptime}$ on $\mathcal{K}_q$, and*
(d) *$\mathrm{FPS}_q = \textsc{Ptime}$ on $\mathcal{K}_q$.*

Proof of Theorem 1. Without loss of generality let $q \in \Omega \setminus \Omega'$. Then by Theorem 3 there exists a class $\mathcal{K}_q$ on which $\mathrm{FPS}_\Omega = \mathrm{FPR}_\Omega = \textsc{Ptime}$ and $\mathrm{FPS}_{\Omega'} = \mathrm{FPR}_{\Omega'} = \mathrm{FPC} < \textsc{Ptime}$. ◀

Proof of Theorem 2. Assume that $\mathrm{FPR} = \textsc{Ptime}$. Then, in particular, $\mathrm{FPR} = \mathrm{FPR}^*$ and there exists a formula $\varphi \in \mathrm{FPR}$ which can uniformly determine the rank of matrices over prime fields, i.e. which can express the uniform rank operator rk^*. As a matter of fact we have $\varphi \in \mathrm{FPR}_\Omega$ for some *finite* set of primes Ω. By using φ we can uniformly express the matrix rank over each prime field $\mathbb{F}_p$ in FPR_Ω. In other words, we have $\mathrm{FPS} \leq \mathrm{FPR} \leq \mathrm{FPR}^* \leq \mathrm{FPR}_\Omega$.

Now let $q \in \mathbb{P} \setminus \Omega$. By Theorem 3 there exists a class of structures $\mathcal{K}_q$ on which $\mathrm{FPR}_\Omega = \mathrm{FPC} < \textsc{Ptime}$. However, the class $\mathcal{K}_q$ can be chosen such that $\textsc{Ptime} = \mathrm{FPS}_q \leq \mathrm{FPR}_\Omega$ on $\mathcal{K}_q$ by Theorem 3 (d) and we obtain the desired contradiction. ◀

The proof of Theorem 2 reveals a deficiency of the logic FPR: each formula can only access rk_p-operators for a finite set Ω of distinct primes p. In fact, the query which we constructed to separate FPR from PTIME can be defined in FPR*. Altogether this suggests to generalise the notion of rank operators and to specify the prime p as a part of the input, as we did for FPR*, and as it was proposed in [14, 15, 17].

The proof of Theorem 3 is structured as follows. We fix a prime q and identify, in a first step, sufficient criteria (i)–(iv) of classes of structures $\mathcal{K} = \mathcal{K}_q$ which guarantee that the relations claimed in (a), (b), (c) and (d) hold. In a second step, we construct a class of structures $\mathcal{K}$ and verify, in a third step, that $\mathcal{K}$ satisfies these sufficient criteria.

Establishing sufficient criteria. We start to find sufficient criteria for part (a) of Theorem 3.

(i) The automorphism groups $\Delta_{\mathfrak{A}} := \mathrm{Aut}(\mathfrak{A})$ of structures $\mathfrak{A} \in \mathcal{K}$ are Abelian q-groups.

(ii) The orbits of ℓ-tuples in structures $\mathfrak{A} \in \mathcal{K}$ can be ordered in FPC: For all $\ell \geq 1$ there exists $\varphi_{\leq}(x_1, \dots, x_\ell, y_1, \dots, y_\ell) \in$ FPC such that for all $\mathfrak{A} \in \mathcal{K}$, the formula $\varphi_{\leq}(\bar{x}, \bar{y})$ defines in $\mathfrak{A}$ a linear preorder $\leq$ on A^ℓ with the property that two ℓ-tuples $\bar{a}, \bar{b} \in A^\ell$ are $\leq$-equivalent if, and only if, they are in the same $\Delta_{\mathfrak{A}}$-orbit.

▶ **Lemma 4.** *If $\mathcal{K}$ satisfies* (i) *and* (ii)*, then* $\mathrm{FPS}_\Omega = \mathrm{FPC}$ *on $\mathcal{K}$ for all $\Omega \subseteq \mathbb{P} \setminus \{q\}$.*

The only interesting step of an inductive translation is the case of a solvability formula

$$\psi(\bar{z}) = (\mathsf{slv}_p\, \bar{x}\bar{\nu} \leq \bar{s}, \bar{y}\bar{\mu} \leq \bar{t})\varphi(\bar{x}\bar{\nu}, \bar{y}\bar{\mu}, \bar{z}).$$

Let $|\bar{x}| = |\bar{y}| = \ell$, $|\bar{\nu}| = |\bar{\mu}| = \lambda$ and $|\bar{z}| = k$. To explain our main argument, we fix a structure $\mathfrak{A} \in \mathcal{K}$ and a k-tuple of parameters $\bar{c} \in (A \uplus \mathbb{N})^k$ which is compatible with the type of $\bar{z}$. According to the semantics of the slv_p-quantifier, the formula φ defines in $(\mathfrak{A}, \bar{z} \mapsto \bar{c})$ an $I \times J$-matrix $M = M^{\mathfrak{A}}_{\bar{c}}$ over $\{0,1\} \subseteq \mathbb{F}_p$ where $I = I^{\mathfrak{A}} := A^\ell \times \mathbb{N}^{\leq \bar{s}} \subseteq A^\ell \times \mathbb{N}^\lambda$ and $J = J^{\mathfrak{A}} := A^\ell \times \mathbb{N}^{\leq \bar{t}} \subseteq A^\ell \times \mathbb{N}^\lambda$ that is defined for $\bar{a} \in I$ and $\bar{b} \in J$ as $M(\bar{a}, \bar{b}) = 1$ if, and only if, $\mathfrak{A} \models \varphi(\bar{a}, \bar{b}, \bar{c})$. Moreover, we have $\mathfrak{A} \models \psi(\bar{c})$ if, and only if, the linear system $M \cdot \vec{x} = \mathbb{1}$ over $\mathbb{F}_p$ is solvable. The key idea of our proof is to use the symmetries of the structure $\mathfrak{A}$ to translate the linear equation system $M \cdot \vec{x} = \mathbb{1}$ into an equivalent linear system for which the solvability problem is FPC-definable.

We set $\Gamma = \Gamma^{\mathfrak{A}}_{\bar{c}} := \mathrm{Aut}(\mathfrak{A}, \bar{c}) \leq \Delta = \Delta_{\mathfrak{A}} = \mathrm{Aut}(\mathfrak{A})$. The group Γ acts on I and J in the natural way. We identify each automorphism $\pi \in \Gamma$ with the corresponding $I \times I$-permutation matrix Π_I and the corresponding $J \times J$-permutation matrix Π_J in the usual way. More precisely, to $\pi \in \Gamma$ we associate the $I \times I$-permutation matrix Π_I with entries $\{0,1\}$ which is defined as $\Pi_I(\bar{a}, \bar{b}) = 1$ if, and only if, $\pi(\bar{a}) = \bar{b}$. Then Γ acts on the set of $I \times J$-matrices by left multiplication with $I \times I$-permutation matrices. Analogously, we let Π_J denote the $J \times J$-permutation matrix with entries $\{0,1\}$ that is defined in the same way as Π_I. Then Γ also acts on the set of $I \times J$-matrices by right multiplication with $J \times J$-permutation matrices. Specifically, for $\pi \in \Gamma$ we have $(\Pi_I \cdot M)(\bar{a}, \bar{b}) = M(\pi(\bar{a}), \bar{b})$ and $(M \cdot \Pi_J^{-1})(\bar{a}, \bar{b}) = M(\bar{a}, \pi(\bar{b}))$. Since M is defined by a formula in the structure $(\mathfrak{A}, \bar{c})$ and since $\Gamma = \mathrm{Aut}(\mathfrak{A}, \bar{c})$ we conclude that $(\Pi_I \cdot M \cdot \Pi_J^{-1})(\bar{a}, \bar{b}) = M(\pi(\bar{a}), \pi(\bar{b})) = M(\bar{a}, \bar{b})$ and thus

$$\Pi_I \cdot M \cdot \Pi_J^{-1} = M \quad \Leftrightarrow \quad \Pi_I \cdot M = M \cdot \Pi_J.$$

This identity leads to the following important observation.

▶ **Lemma 5.** *If $M \cdot \vec{x} = \mathbb{1}$ is solvable, then the system has a* Γ-symmetric *solution, i.e. a solution $\vec{b} \in \mathbb{F}_p^J$ such that $\Pi_J \cdot \vec{b} = \vec{b}$ for all $\pi \in \Gamma$.*

Proof. If $M \cdot \vec{b} = \mathbb{1}$, then also $\Pi_I \cdot (M \cdot \vec{b}) = \mathbb{1}$ and thus $M \cdot (\Pi_J \cdot \vec{b}) = \mathbb{1}$ for all $\pi \in \Gamma$. This shows that Γ acts on the solution space of the linear equation system. Since $\mathcal{K}$ satisfies property (i) we know that Γ is a q-group for a prime $q \neq p$. Thus each Γ-orbit has size q^r for some $r \geq 0$. On the other hand, the number of solutions is a power of p. We conclude that there is at least one Γ-orbit which contains a single solution only. ◀

Let $\vec{b} \in \mathbb{F}_p^J$ be a Γ-symmetric solution. Then the entries of the solution $\vec{b}$ on Γ-orbits are constant: for $j \in J$ and $\pi \in \Gamma$ we have $\vec{b}(\pi(j)) = (\Pi_J \cdot \vec{b})(j) = \vec{b}(j)$. We use property (ii) to show that there is an FPC-formula $\varphi_\leq(\bar{x}, \bar{y})$ which defines for all $\mathfrak{A} \in \mathcal{K}$ and $\bar{c} \in (A \uplus \mathbb{N})^k$ as above a linear preorder $\preceq$ on A^ℓ which identifies Γ-orbits. Note that, in general, $\Gamma = \mathrm{Aut}(\mathfrak{A}, \bar{c})$ is a strict subgroup of $\Delta = \mathrm{Aut}(\mathfrak{A})$. Thus we can not directly apply (ii). However, the Γ-orbits on A^ℓ correspond to the Δ-orbits on $A^{k'+\ell}$ where the first k' entries are fixed to the elements in $\{c_1, \dots, c_k\} \cap A$.

The linear preorder $\preceq$ naturally extends to a preorder on the sets I and J with the same properties. Let us write $J = J_0 \preceq J_1 \preceq \dots \preceq J_{v-1}$ to denote the decomposition of J into the Γ-orbits J_j which are ordered by $\preceq$ as indicated. Moreover, for $j \in [v]$ we let e_j denote the identity vector on the j-th orbit J_j, i.e. the J-vector which defined for $i \in J$ as $e_j(i) = 1$ if $i \in J_j$ and as $e_j(i) = 0$ otherwise. Let E denote the $J \times [v]$-matrix whose j-th column is the vector e_j. It follows that a Γ-symmetric solution $\vec{b}$ can be written as $E \cdot \vec{b}_* = \vec{b}$ for a unique $[v]$-vector $\vec{b}_*$. Together with Lemma 5 this shows the following.

▶ **Lemma 6.** *The system $M \cdot \vec{x} = \mathbb{1}$ is solvable if, and only if, $(M \cdot E) \cdot \vec{x}_* = \mathbb{1}$ is solvable.*

Finally, we observe that the coefficient matrix $M_* := (M \cdot E)$ of the equivalent linear equation system $M_* \cdot \vec{x}_* = \mathbb{1}$ can easily be obtained in FPC and that it is a matrix over the *ordered* set of column indices $[v]$. It is a simple observation that such linear equation systems can be solved in FPC: the linear order on the column set induces (together with some fixed order on $\mathbb{F}_p$) a lexicographical ordering on the set of rows which is, up to duplicates of rows, a linear order on this set. Thus, in general, if we have a linear order on *one* of the index sets of the coefficient matrix this suffices to obtain an equivalent matrix where *both* index sets are ordered, see also [17]. This finishes our proof of Lemma 4.

We proceed to show that the conditions (i) and (ii) also guarantee that rank operators can be reduced to solvability operators over the class

▶ **Lemma 7.** *If $\mathcal{K}$ satisfies* (i) *and* (ii)*, then* $\mathrm{FPR}_\Omega = \mathrm{FPS}_\Omega$ *on $\mathcal{K}$ for all sets of primes Ω.*

Proof. The only interesting case of an inductive translation is the case of rank terms

$$\Upsilon(\bar{z}) = [\mathsf{rk}_p\,(\bar{x}\bar{\nu} \leq \bar{t}, \bar{y}\bar{\mu} \leq \bar{s})\,.\,\Theta(\bar{x}\bar{\nu}, \bar{y}\bar{\mu}, \bar{z})].$$

Let $|\bar{x}| = |\bar{y}| = \ell$, $|\bar{\nu}| = |\bar{\mu}| = \lambda$ and $|\bar{z}| = k$. Let $\mathfrak{A} \in \mathcal{K}$ and let $\bar{c}$ be a k-tuple of parameters $\bar{c} \in (A \uplus \mathbb{N})^k$ which is compatible with $\bar{z}$. The term Θ defines in $(\mathfrak{A}, \bar{z} \mapsto \bar{c})$ for $I^{\mathfrak{A}} = I := A^{|\bar{x}|} \times \mathbb{N}^{\leq \bar{t}}$ and $J^{\mathfrak{A}} = J := A^{|\bar{y}|} \times \mathbb{N}^{\leq \bar{s}}$ the $I \times J$-matrix M over $\mathbb{F}_p$ which is defined as $M(\bar{a}\bar{n}, \bar{b}\bar{m}) := \Theta^{\mathfrak{A}}(\bar{a}\bar{n}, \bar{b}\bar{m}, \bar{c}) \bmod p$.

We proceed to show that we can obtain the matrix rank of M, that is the value $\Upsilon^{\mathfrak{A}}(\bar{c}) \in \mathbb{N}$, by a recursive application of solvability queries. We first make the following key observation.

Claim: There are FPC-formulas $\varphi_\leq(\bar{y}_1\bar{\mu}_1, \bar{y}_2\bar{\mu}_2)$, $\psi_\leq(\bar{v}, \bar{y}_1\bar{\mu}_1, \bar{y}_2\bar{\mu}_2)$ such that for every $\mathfrak{A} \in \mathcal{K}$

(a) $\varphi_\leq^{\mathfrak{A}}$ is a linear preorder $\preceq$ on $J^{\mathfrak{A}}$, and such that
(b) for every $\preceq$-class $[j] \subseteq J^{\mathfrak{A}}$ there exists $\bar{d} \in A^{|\bar{v}|}$ such that $\psi_\leq^{\mathfrak{A}}(\bar{d})$ is a linear order on $[j]$.

Proof of claim: We use property (ii) to choose an FPC-formula $\varphi_\leq$ which defines in all $\mathfrak{A} \in \mathcal{K}$ a linear preorder $\preceq$ on $J^{\mathfrak{A}}$ such that $\preceq$-classes correspond to $\Delta_{\mathfrak{A}}$-orbits Analogously, we choose an FPC-formula $\vartheta_\leq$ which defines in every structure $\mathfrak{A} \in \mathcal{K}$ a linear preorder $\preceq^*$ on $J^{\mathfrak{A}} \times J^{\mathfrak{A}}$ and that induces a linear order on the $\Delta_{\mathfrak{A}}$-orbits.

To obtain $\psi_\leq$, we let $[j] \subseteq J^{\mathfrak{A}}$ be a $\preceq$-class for some $\mathfrak{A} \in \mathcal{K}$. By property (i) we know that $\Delta_{\mathfrak{A}}$ is an Abelian group. Thus, each automorphism $\pi \in \Delta_{\mathfrak{A}}$ which fixes *one* element in the $\Delta_{\mathfrak{A}}$-orbit $[j]$ point-wise fixes *every* element in the class $[j]$. We conclude that the restriction of $\preceq^*$ to elements in $\{j'\} \times [j]$ corresponds to a linear order on $[j]$ for each $j' \in [j]$. ⊣

We are prepared to describe the recursive procedure which allows us to determine the rank of the matrix M in FPS_Ω. We fix formulas $\varphi_\leq$ and $\psi_\leq$ with the properties stated in the claim above. Moreover, let $\preceq$ denote the linear preorder defined by $\varphi_\leq$ on $J = J_0 \preceq J_1 \preceq \cdots \preceq J_{r-1}$. We use the formula $\psi_\leq$ to obtain on each class J_i a family of definable linear orderings (which depend on the choice of different parameters). For $j \in J$ we denote by $\vec{m}_j \in \mathbb{F}_p^I$ the j-th column of the matrix M. Then the rank of M coincides with the dimension of the $\mathbb{F}_p$-vector space which is generated by the set of columns $\{\vec{m}_j : j \in J\}$ of the matrix M.

Now, for $i \in [r]$ we recursively obtain the dimension $d_i \in \mathbb{N}$ of the $\mathbb{F}_p$-vector space generated by $V_i := \{\vec{m}_j : j \in J_0 \cup J_1 \cup \cdots \cup J_i\}$ as follows. First, we use $\psi_\leq$ to fix a linear order on J_i (the following steps are independent of the specific linear order and can thus be performed in parallel for each such order). Using this linear order on J_i we can identify in FPS_Ω a maximal set $W \subseteq \{\vec{m}_j : j \in J_i\}$ of linearly independent columns such that $\langle V_{i-1} \rangle \cap \langle W \rangle = \{\vec{0}\}$. Indeed, if $\langle V_{i-1} \rangle \cap \langle W \rangle = \{\vec{0}\}$, then for $\vec{m} \in \{\vec{m}_j : j \in J_i\}$, $\vec{m} \notin \langle W \rangle$ we have that $\langle V_{i-1} \rangle \cap \langle W \uplus \{\vec{m}\} \rangle = \{\vec{0}\}$ if, and only if, $\vec{m} \notin \langle V_{i-1} \cup W \rangle$. Observe that the conditions $\vec{m} \notin \langle W \rangle$ and $\vec{m} \notin \langle V_{i-1} \cup W \rangle$ correspond to the solvability of a linear equation system over $\mathbb{F}_p$. We claim that $d_i = d_{i-1} + |W|$. Indeed, by the maximality of W and since $\langle V_{i-1} \rangle \cap \langle W \rangle = \{\vec{0}\}$ it follows that $\langle V_i \rangle = \langle V_{i-1} \rangle \oplus \langle W \rangle$. Moreover, W consists of linearly independent columns and is a basis for $\langle W \rangle$.

Since the above described recursion can easily be implemented in FPS_Ω, we conclude that the rank d_{r-1} of the matrix M can be determined in FPS_Ω which completes our proof. ◀

We now focus on the parts (c) and (d) of Theorem 3.

(iii) There exists an FPS_q-definable canonisation procedure on $\mathcal{K}$.

(iv) For all $k \geq 1$ there is a pair $\mathfrak{A}, \mathfrak{B} \in \mathcal{K}$ such that $\mathfrak{A} \ncong \mathfrak{B}$ and $\mathfrak{A} \equiv_k^C \mathfrak{B}$ (that is, $\mathfrak{A}$ and $\mathfrak{B}$ cannot be distinguished in the k-variable fragment of infinitary counting logic $\mathrm{C}^k_{\infty\omega}$).

▶ **Lemma 8.** *If* $\mathcal{K}$ *satisfies* (iii) *and* (iv), *then* $\mathrm{FPC} < \mathrm{FPS}_q = \textsc{Ptime}$ *on* $\mathcal{K}$.

Constructing an appropriate class of structures. We proceed to construct a class of structures $\mathcal{K}$ which satisfies properties (i)–(iv). Our approach is a generalisation of the well-known construction of Cai, Fürer and Immerman [4] for fields $\mathbb{F}_q$, $q \in \mathbb{P}$. The difference to the original construction (which arises as a special case for $q = 2$) is that we replace every edge e from the original graph $\mathcal{G}$ by q copies $e_0, e_1, \ldots, e_{q-1}$ which we arrange on a directed cycle of length q. For $q = 2$ this is equivalent to just taking two non-connected atoms e_0, e_1. While the symmetries of the original CFI-graphs arise by twisting pairs of corresponding edges e_0, e_1, the symmetries of generalised CFI-structures arise by shifting the cycles on $e_0, e_1, \ldots, e_{q-1}$ by some value $x \in \mathbb{F}_q$. In both cases, the resulting twists and cyclic shifts can be propagated along paths in $\mathcal{G}$. We remark that the same kind of generalisations have been studied, for example, in [14, 18]. Due to space limitations, we have to leave out the following proofs which are mostly straightforward adaptations of the arguments for the original construction.

We start from an *(undirected), connected* and *ordered* graph $\mathcal{G} = (V, \leq, E)$. Let C, I and R be binary relation symbols. We set $\tau := \{\preceq, C, I, R\}$. We define for every prime q and every sequence of *gadget values* $\vec{d} = (d_v)_{v \in V} \in [q]^V$ a τ-structure $\mathrm{CFI}_q(\mathcal{G}, \vec{d})$ which we call a *CFI-structure over* $\mathcal{G}$. For the following construction we agree that arithmetic is modulo q so that we can drop the operator "mod q" in statements of the form $x = y \bmod q$ and $x + y \bmod q$ for the sake of better readability. For what follows, let $E(v) \subseteq E$ denote the set of *directed* edges starting in v. Since $\mathcal{G}$ is an undirected graph, this means that for each undirected edge $\{v, w\}$ of $\mathcal{G}$ we have $(v, w) \in E(v)$ and $(w, v) \in E(w)$.

- The *universe* of $\mathrm{CFI}_q(\mathcal{G}, \vec{d})$ consists of *edge nodes* and *equation nodes*.
  - The set of *edge nodes* $\hat{E}$ is defined as $\hat{E} := \bigcup_{e \in E} \hat{e}$ where for every *directed* edge $e \in E$ we let the *edge class* $\hat{e} = \{e_0, e_1, \dots, e_{q-1}\}$ consist of q distinct copys of e. In particular, for every edge $e = (v, w) \in E$ and its reversed edge $e^{-1} := f = (w, v) \in E$ the sets $\hat{e}$ and $\hat{f}$ are disjoint. We say that two such edges (or edge classes) are *related*.
  - The set of *equation nodes* $\hat{V}$ is defined as $\hat{V} := \bigcup_{v \in V} \hat{v}^{\vec{d}(v)}$ where for every vertex $v \in V$ and $d \in [q]$ the *equation class* $\hat{v}^d$ consist of all functions $\rho : E(v) \to [q]$ which satisfy $\sum \rho := \sum_{e \in E(v)} \rho(e) = d$.
- The *linear preorder* $\preceq$ orders the edge classes according to the linear order induced by $\leq$ on E. More precisely, we let $\hat{e} \preceq \hat{f}$ whenever $e \leq f$. Similarly, $\preceq$ orders the equation classes according to the order of $\leq$ on V, i.e. $\hat{v} \preceq \hat{w}$ if $v \leq w$. Moreover, we let $\hat{e} \preceq \hat{v}$ for edge classes $\hat{e}$ and equation classes $\hat{v}$.
- The *cycle relation* C contains a directed cycle of length q on each of the edge classes $\hat{e}$ for $e \in E$, i.e. $C = \{(e_i, e_{i+1}) : i \in [q], e \in E\}$.
- The *inverse relation* I connects two related edge classes by pairing additive inverses. More precisely, let $e = (v, w) \in E$ and $f = (w, v) \in E$. Then I contains all edges (e_x, f_y) with $x + y = 0$ for $x, y \in [q]$.
- The *gadget relation* R is defined as $R := \bigcup_{v \in V} R_v^{\vec{d}(v)}$ where for $v \in V$ and $d \in [q]$ the relation R_v^d is given as

$$R_v^d := \{(\rho, e_{\rho(e)}) : \rho \in \hat{v}^d, e \in E(v)\}.$$

At first glance our construction associates to every graph $\mathcal{G}$ (with the above properties) and to each sequence of gadget values $\vec{d} \in [q]^V$ a different structure $\mathrm{CFI}_q(\mathcal{G}, \vec{d})$. However, for each graph $\mathcal{G}$ with the above properties there really are, up to isomorphism, only q different CFI-structures $\mathrm{CFI}_q(\mathcal{G}, \vec{d})$.

▶ **Lemma 9.** *Let $\vec{d}, \vec{d}_* \in ([q])^V$. Then $\mathrm{CFI}_q(\mathcal{G}, \vec{d}) \cong \mathrm{CFI}_q(\mathcal{G}, \vec{d}_*)$ if, and only if, $\sum \vec{d} = \sum \vec{d}_*$.*

A connected graph $\mathcal{G}$ is *k-connected*, for $k \geq 0$, if $\mathcal{G}$ contains more than k vertices and if $\mathcal{G}$ stays connected when we remove any set of at most k vertices. The *connectivity* $\mathrm{con}(\mathcal{G})$ of $\mathcal{G}$ is the maximal $k \geq 0$ such that $\mathcal{G}$ is k-connected. Moreover, the *connectivity* $\mathrm{con}(\mathfrak{G})$ of a class $\mathfrak{G}$ of connected graphs is the function $\mathrm{con}(\mathfrak{G}) : \mathbb{N} \to \mathbb{N}$ defined as

$$n \mapsto \min_{\mathcal{G} \in \mathfrak{G}, |\mathcal{G}| = n} \mathrm{con}(\mathcal{G}).$$

We are prepared to define the class $\mathcal{K}$: let $\mathfrak{G}$ be a class of *undirected, ordered, connected* graphs such that $\mathrm{con}(\mathfrak{G}) \in \omega(1)$ (for example complete, ordered graphs). Then we set

$$\mathcal{K} = \mathcal{K}_q := \{\mathrm{CFI}_q(\mathcal{G}, \vec{d}) \;:\; \mathcal{G} = (V, \leq, E) \in \mathfrak{G}, \vec{d} \in [q]^V\}.$$

Verifying the required properties. First of all, one can see that the cycle relation C and the preorder $\preceq$ enforce that the automorphism group of a CFI-structure $\mathrm{CFI}_q(\mathcal{G},\vec{d})$ over $\mathcal{G} = (V,\preceq,E) \in \mathfrak{G}$ is a subgroup of $\mathbb{F}_q^E$. Thus property (i) holds for $\mathcal{K}$.

To show that $\mathcal{K}$ satisfies property (ii), we fix the length $\ell \geq 1$ of tuples on which we want to define a linear preorder which identifies $\Delta_{\mathfrak{A}}$-orbits. By the definition of $\mathcal{K}$ it suffices to consider CFI-structures $\mathfrak{A} = \mathrm{CFI}_q(\mathcal{G},\vec{d})$ over graphs $\mathcal{G} = (V,\preceq,E) \in \mathfrak{G}$ with $\mathrm{con}(\mathcal{G}) > (\ell+2)$ since almost all structures in $\mathcal{K}$ satisfy this condition. Then we can show that the equivalence classes of ℓ-tuples in the infinitary logic with counting and $(\ell+2)$ variables $\mathrm{C}^{\ell+2}_{\infty\omega}$ coincides with the $\Delta_{\mathfrak{A}}$-orbits of ℓ-tuples for structures $\mathfrak{A} \in \mathcal{K}$.

▶ **Lemma 10.** *Let $\lambda \leq \ell$ and let $\bar{a},\bar{b} \in A^\lambda$. Then $(\mathfrak{A},\bar{a}) \equiv^C_{\ell+2} (\mathfrak{A},\bar{b})$ if, and only if, there exists $\pi \in Aut(\mathfrak{A})$ such that $\pi(\bar{a}) = \bar{b}$.*

It is well-known that classes of $\mathrm{C}^{\ell+2}_{\infty\omega}$-equivalent tuples can be ordered in FPC, see e.g. [16]. Hence, it follows from our previous lemma that the class $\mathcal{K}$ satisfies property (ii).

▶ **Lemma 11.** *The class $\mathcal{K}$ satisfies the properties* (i) *and* (ii).

We turn our attention to property (iv). In the next lemma we state that for each $k \geq 1$ and each sufficiently connected graph $\mathcal{G} \in \mathfrak{G}$, the logic $\mathrm{C}^k_{\infty\omega}$ cannot distinguish between any pair of CFI-structures over $\mathcal{G}$ (although there exist non-isomorphic CFI-structures over $\mathcal{G}$).

▶ **Lemma 12.** *Let $k \geq 1$ and let $\mathcal{G} = (V,\preceq,E) \in \mathfrak{G}$ such that con($\mathcal{G}$) $> k$. Then for all $\vec{d},\vec{d}_* \in [q]^V$ it holds that $\mathrm{CFI}_q(\mathcal{G},\vec{d}) \equiv^C_k \mathrm{CFI}_q(\mathcal{G},\vec{d}_*)$. Thus, $\mathcal{K}$ satisfies property (iv).*

To complete our proof we establish an FPS_q-definable canonisation procedure on $\mathcal{K}$. The idea is as follows: given a CFI-structure $\mathfrak{A} = \mathrm{CFI}_q(\mathcal{G},\vec{d})$ and a value $z \in [q]$ we construct a linear equation system over $\mathbb{F}_q$ which is solvable if, and only if, $\sum \vec{d} = z$. This linear equation system is FO-definable in $\mathfrak{A}$ which shows that FPS_q can determine the isomorphism class of a CFI-structure over $\mathcal{G}$. Since the graph $\mathcal{G}$ is ordered it is easy to construct an ordered representative from each isomorphism class of CFI-structures over $\mathcal{G}$.

More specifically, let $\mathcal{G} = (V,\preceq,E) \in \mathfrak{G}$, let $\mathfrak{A} = \mathrm{CFI}_q(\mathcal{G},\vec{d}) \in \mathcal{K}$ and let $z \in \mathbb{F}_q$. For our linear equation system we identify each element $e_i \in \hat{E}$ and each vertex $v \in V$ with a variable over $\mathbb{F}_q$, i.e. we let $\mathcal{V} := \hat{E} \uplus V$ be the set of variables. The equations are given as follows:

$$e_{i+1} = e_i + 1 \qquad \text{for all } e_i \in \hat{E} \qquad \text{(E 1)}$$

$$e_i = -f_{-i} \qquad \text{for related edges } e,f \in E \qquad \text{(E 2)}$$

$$v = \sum_{e \in E(v)} e_{\rho(e)} \qquad \text{for all } v \in V, \rho \in \hat{v} \qquad \text{(E 3)}$$

$$z = \sum_{v \in V} v. \qquad \text{(E 4)}$$

It is easy to see that this system is FO-definable in $\mathfrak{A}$. First of all, the equation (E 4) can be defined as a sum over the ordered set V. Moreover, we can express the equations of type (E 1) and (E 2) by using the cycle and inverse relation, respectively. Finally, the equations of type (E 3) can be expressed by using the gadget relation R.

▶ **Lemma 13.** *The above defined system is solvable if, and only if, $\sum \vec{d} = z$.*

▶ **Lemma 14.** *The class $\mathcal{K}$ satisfies the property* (iii).

4 Solvability quantifiers vs. rank operators

In the previous section we obtained separation results for the extensions of FPC by solvability quantifiers (and rank operators) over different sets of primes. One important step of our proof was to construct a class of structures on which the expressive power of the logics FPR_Ω and FPS_Ω coincides. Moreover, as we already mentioned in Section 2, most of the queries which are known to separate fixed-point logic with counting and rank logic can also be expressed in FPS. This naturally leads to the question whether, in general, rank operators can be simulated by solvability quantifiers in fixed-point logic with counting.

In this section we solve a simplified version of this question and show that in the absence of counting, rank operators are strictly more expressive than solvability quantifiers. The reader should recall that rank operators can easily simulate counting terms but this does not hold for solvability quantifiers.

To state our main result formally, we define for every prime p the extension FOS_p of first-order logic (without counting) by solvability quantifiers over $\mathbb{F}_p$. The crucial difference to the extension FOR_p of first-order logic by rank operators rk_p is that FOS_p is a *one-sorted* logic which does not have access to a counting sort.

▶ **Definition 15.** For every prime p, the logic FOS_p results by extending the syntax of FO by the following formula creation rule:

- If $\varphi(\bar{x}, \bar{y}, \bar{z}) \in \mathrm{FOS}_p$, then $\psi(\bar{z}) = (\mathsf{slv}_p\, \bar{x}, \bar{y})\varphi(\bar{x}, \bar{y}, \bar{z})$ is an FOS_p-formula.

The semantics of $\psi(\bar{z})$ are defined as for FPS_p.

We briefly summarise what is known about FOS_p (see also [6, 17]). It follows from [7, 14] that for every prime p, the logic FOS_p subsumes the logic STC and that $\mathrm{FOS}_p \nleq \mathrm{FPC}$. Moreover, on ordered structures, the expressive power of FOS_p can be characterised in terms of a natural complexity class: in [3], Buntrock et. al. introduced the *logarithmic space modulo counting classes* $\mathrm{MOD}_k\mathrm{L}$ for integers $k \geq 2$. Informally, a problem is in $\mathrm{MOD}_k\mathrm{L}$ if there exists a NL-Turing machine which verifies its inputs by producing a number of accepting paths which is not congruent $0 \bmod k$. It turns out that, at least for primes p, the class $\mathrm{MOD}_p\mathrm{L}$ is closed under many natural operations, including all Boolean operations and even logspace Turing reductions [3]. Furthermore, many problems from linear algebra over $\mathbb{F}_p$ are complete for $\mathrm{MOD}_p\mathrm{L}$. In particular this is true for the solvability problem of linear equation systems over $\mathbb{F}_p$ and for computing the matrix rank over $\mathbb{F}_p$ [3].

Building on these insights, Dawar et. al. were able to show in [7] that for all $p \in \mathbb{P}$, the logic FOR_p captures $\mathrm{MOD}_p\mathrm{L}$ on the class of ordered structures. It has been noted in [17] that their proof shows the same correspondence for FOS_p.

▶ **Proposition 16** ([7],[17]). *On ordered structures we have* $\mathrm{FOS}_p = \mathrm{FOR}_p = \mathrm{MOD}_p\mathrm{L}$.

Despite this nice characterisation over ordered structures, the situation over general structures remained unclear. It easily follows that $\mathrm{FOS}_p \leq \mathrm{FOR}_p \leq \mathrm{FPR}_p$, but, so far, it has been open whether one, or both, of these inclusions are strict. In this section we show:

▶ **Theorem 17.** *For all primes p we have $\mathrm{FOS}_p < \mathrm{FOR}_p$ over the class of sets $\mathcal{S}(\varnothing)$.*

In some sense, this result is not very surprising. While FOS_p has to express $\mathcal{S}(\varnothing)$-properties over *unordered* sets, which have the maximal amount of symmetries, FOR_p can use the size of a set as a complete invariant to express properties of $\mathcal{S}(\varnothing)$-structures over the *ordered* numerical sort. However, it is not obvious how one can turn this intuition into a formal argument. In fact, FOS_p has non-trivial expressive power over sets. For instance,

FOS_p can determine the size of sets modulo p^k for every fixed k, while fixed-point logic FP, for example, collapses to first-order logic over sets.

To prove Theorem 17 we recall the following normal form for FOS_p which has been established in Corollary 4.8 of [6].

▶ **Theorem 18.** *Every formula $\vartheta(\bar{z}) \in \mathrm{FOS}_p$ is equivalent to an FOS_p-formula of the form $(\mathsf{slv}_p\, \bar{x}_1, \bar{x}_2)\alpha(\bar{x}_1, \bar{x}_2, \bar{z})$ where $\alpha(\bar{x}_1, \bar{x}_2, \bar{z})$ is quantifier-free.*

Similar to our approach in Section 3, the main idea for separating FOS_p and FOR_p is to exploit the symmetries of definable linear equation systems. More precisely, our plan is to considerably reduce the size of a given linear equation system along an FOR_p-definable transformation. For the remainder of this section, let us fix a quantifier-free formula $\alpha(x_1, \dots, x_k, y_1, \dots, y_\ell) \in \mathrm{FO}(\varnothing)$ and a prime p. According to the semantics of FOS_p, the formula α defines in an input structure $\mathfrak{A} = ([n])$ of size n the $[n]^k \times [n]^\ell$-coefficient matrix M_n which is given for $\bar{a} \in [n]^k, \bar{b} \in [n]^\ell$ as

$$M_n(\bar{a}, \bar{b}) = \begin{cases} 1, & \text{if } \mathfrak{A} \models \alpha(\bar{a}, \bar{b}) \\ 0, & \text{otherwise.} \end{cases}$$

Then $\mathfrak{A} \models (\mathsf{slv}_p\, \bar{x}_1, \bar{x}_2)\alpha(\bar{x}_1, \bar{x}_2)$ if the linear equation system $M_n \cdot \vec{x} = \mathbb{1}$ over $\mathbb{F}_p$ is solvable. For convenience we set $I_n = [n]^k$ and $J_n = [n]^\ell$.

Let $\Gamma = \Gamma_n = \mathrm{Sym}([n])$. Then the group Γ acts on I_n and J_n and we identify the action of $\pi \in \Gamma$ with the multiplication by the associated $I_n \times I_n$-permutation matrix Π_I and the $J_n \times J_n$-permutation matrix Π_J, respectively, as in Section 3. Hence, for $\pi \in \Gamma$ we have

$$\Pi_I \cdot M_n \cdot \Pi_J^{-1} = M_n \quad \Leftrightarrow \quad \Pi_I \cdot M_n = M_n \cdot \Pi_J.$$

For what follows, we fix a prime $q \neq p$ and a subgroup $\Delta \leq \Gamma$ such that $|\Delta| = q^m$ for some $m \geq 0$. The overall strategy is to use the Δ-symmetries of the matrix M_n to strongly reduce the size of the linear equation system $M_n \cdot \vec{x} = \mathbb{1}$. More precisely we claim that for $M_n^* := \sum_{\pi \in \Delta} \Pi_I \cdot M_n$ the linear equation system $M_n \cdot \vec{x} = \mathbb{1}$ is solvable if, and only if, $M_n^* \cdot \vec{x} = \mathbb{1}$ is solvable. First of all we note that for all $\pi \in \Delta$ we have:

- $\Pi_I \cdot M_n^* = \sum_{\lambda \in \Delta} \Pi_I \cdot \Lambda_I \cdot M_n = \sum_{\pi \in \Delta} \Pi_I \cdot M_n = M_n^*$
- $M_n^* \cdot \Pi_J = \sum_{\lambda \in \Delta} \Lambda_I \cdot M_n \cdot \Pi_J = \sum_{\lambda \in \Delta} \Lambda_I \cdot \Pi_I \cdot M_n = M_n^*$.

To verify our original claim assume that $M_n^* \cdot \vec{b} = \mathbb{1}$. Then we have

$$\mathbb{1} = M_n^* \cdot \vec{b} = \big(\sum_{\pi \in \Delta} \Pi_I \cdot M_n\big) \cdot \vec{b} = \big(\sum_{\pi \in \Delta} M_n \cdot \Pi_J\big) \cdot \vec{b} = M_n \cdot \sum_{\pi \in \Delta} (\Pi_J \cdot \vec{b}).$$

For the other direction let $M_n \cdot \vec{b} = \mathbb{1}$. Then $\sum_{\pi \in \Delta} \Pi_I \cdot M_n \cdot \vec{b} = |\Delta| \cdot \mathbb{1}$, hence $(1/|\Delta|) \cdot \vec{b}$ is a solution of the linear equation system $M_n^* \cdot \vec{x} = \mathbb{1}$. Note that for this direction we require that q and p are co-prime as we have to divide by $|\Delta|$.

Since M_n^* satisfies $\Pi_I \cdot M_n^* = M_n^* \cdot \Pi_J = M_n^*$ for all $\pi \in \Delta$ we have

$$M_n^*(\bar{a}, \bar{b}) = M_n^*(\pi(\bar{a}), \bar{b}) = M_n^*(\bar{a}, \pi(\bar{b}))$$

for all $\bar{a} \in I_n, \bar{b} \in J_n$ and $\pi \in \Delta$. In other words, the entries of the $I_n \times J_n$-matrix M_n^* are constant on the Δ-orbits of the index sets I_n and J_n. More specifically, if we let I_n^Δ and J_n^Δ denote the sets of Δ-orbits on I_n and J_n, respectively, then M_n^* can be identified with the matrix (M_n^*/Δ) which is defined as

$$(M_n^*/\Delta) : I_n^\Delta \times J_n^\Delta \to \mathbb{F}_p, ([\bar{a}], [\bar{b}]) \mapsto M_n^*(\bar{a}, \bar{b}).$$

Note that, depending on the size of the group Δ, the sets I_n^Δ and J_n^Δ can be noticeably smaller than the index sets I_n and J_n. Hence our obvious strategy is to choose Δ as large as possible to obtain a compact linear equation system $M_n^* \cdot \vec{x} = \mathbb{1}$ which is equivalent to the given one. It can be shown that for the case $n = q^r$, the size of the maximal q-subgroups Δ_n of Γ_n (the *q-Sylow subgroups*) is exponential in n and that the Δ_n-orbits on I_n and J_n can be described by a tuple of constant length with entries in $[r]$. Moreover, given a set $\mathfrak{A} = ([r])$ it is possible to construct in FOR_p the matrix $M_n^* = (M_n^*/\Delta_n)$ for $n = q^r$. In other words, FOR_p can equivalently express the solvability problem $M_n \cdot \vec{x} = \mathbb{1}$ defined by α in a structure of size $n = q^r$ in an exponentially more succinct structure of size r. The following lemma summarises this fact.

▶ **Lemma 19.** *There exists an* FOC*-term* $\Theta(\bar{\mu}, \bar{\nu})$ *which defines for all* $r \geq q$ *in the structure* $\mathfrak{A} = ([r])$ *the matrix* M_n^* *for* $n = q^r$.

▶ **Definition 20.** Let $\mathcal{K} \subseteq \mathcal{S}(\varnothing)$ be a class of sets. The *q-power* $\mathcal{K}^q \subseteq \mathcal{S}(\varnothing)$ *of* $\mathcal{K}$ consists of all sets $\mathfrak{A} = ([q^r])$ such that $\mathfrak{B} = ([r]) \in \mathcal{K}$.

▶ **Theorem 21.** *Let* $\mathcal{K} \subseteq \mathcal{S}(\varnothing)$. *If* $\mathcal{K}^q$ *is definable in* FOS_p, *then* $\mathcal{K}$ *is definable in* FOR_p.

Proof. If $\mathcal{K}^q$ is FOS_p-definable, then by Theorem 18 by a formula $\varphi = (\mathsf{slv}_p\, \bar{x}_1, \bar{x}_2)\alpha(\bar{x}_1, \bar{x}_2) \in \mathrm{FOS}_p$ where α is quantifier-free.

By using the above construction and Lemma 19, we conclude that the linear equation system $M_n \cdot \vec{x} = \mathbb{1}$ defined by α in an input structure $\mathfrak{A} = ([n])$ of size $n = q^r$ can be transformed into the equivalent system $M_n^* \cdot \vec{x} = \mathbb{1}$ which is FOC-definable in $\mathfrak{B} = ([r])$. Let $\varphi^* \in \mathrm{FOR}_p$ be a formula which expresses the solvability of the linear system $M_n^* \cdot \vec{x} = \mathbb{1}$ in a structure $\mathfrak{B} = ([r])$. Then $\mathfrak{B} \models \varphi^*$ if, and only if, $\mathfrak{A} \models \varphi$ since the linear equation systems $M_n \cdot \vec{x} = \mathbb{1}$ and $M_n^* \cdot \vec{x} = \mathbb{1}$ are equivalent. Hence φ^* defines $\mathcal{K}$. ◀

Proof of Theorem 17. Otherwise we would have $\mathrm{FOS}_p = \mathrm{FOR}_p$. Let $\mathcal{K} \subseteq \mathcal{S}(\varnothing)$ be a class of sets such that $\mathcal{K} \notin \mathrm{FOR}_p$, but such that $(\mathcal{K}^q)^q \in \mathrm{FOR}_p$. Such a class $\mathcal{K}$ is well-known to exist (just combine the fact that, over sets, we have LOGSPACE $\leq \mathrm{FOR}_p \leq$ PTIME and the space-hierarchy theorem). Since $\mathrm{FOS}_p = \mathrm{FOR}_p$ we had $(\mathcal{K}^q)^q \in \mathrm{FOS}_p$ and by Theorem 21 this means that $\mathcal{K}^q \in \mathrm{FOR}_p$. Again, since $\mathrm{FOR}_p = \mathrm{FOS}_p$, we had $\mathcal{K}^q \in \mathrm{FOS}_p$. A second application of Theorem 21 yields $\mathcal{K} \in \mathrm{FOR}_p$ which contradicts our assumptions. ◀

Finally we remark that, in the absence of counting, the same proof works for the extension of fixed-point logic by solvability quantifiers. The simple reason is that fixed-point operators do not increase the expressive power of first-order logic over the empty signature since all definable relations are composed from a constant-sized set of basic building blocks.

5 Discussion

We showed that the expressive power of rank operators over different prime fields is incomparable and we inferred that the version of rank logic FPR with a distinct rank operator rk_p for every prime $p \in \mathbb{P}$ fails to capture polynomial time. In particular our proof shows that FPR cannot express the uniform version of the matrix rank problem where the prime p is part of the input. Moreover, we separated rank operators and solvability quantifiers in the absence of counting.

Of course, an immediate question is whether the extension FPR^* of FPC by the uniform rank operator rk^* suffices to capture polynomial time. We do not believe that this is the case. A natural candidate to separate FPR^* from PTIME is the solvability problem for linear equation systems over finite rings rather than fields [6]. While linear equations systems can be efficiently solved also over rings, there is no notion of matrix rank that seems to be helpful for this purpose. In particular, it is open whether FPR* can define the isomorphism problem for CFI-structures generalised to $\mathbb{Z}_4$. A negative answer to this last question would provide a class of structures on which FPR* is strictly weaker than Choiceless Polynomial Time (which captures PTIME on this class [1]).

Another question concerns the relationship between solvability logic FPS and rank logic FPR*. Our proof of Lemma 7 shows that on every class of structures of bounded colour class size the two logics have the same expressive power. However, over general structures this reduction fails. We only know, by our results from Section 4, that a simulation of rank operators by solvability quantifiers would require counting.

Finally, we think it is worth to explore the connections between our approach and the game-theoretic approach proposed by Dawar and Holm in [8] to see to what extent our methods can be combined. For example, what kind of properties does a variant of their partition games have for infinitary logics with solvability quantifiers?

References

1 F. Abu Zaid, E. Grädel, M. Grohe, and W. Pakusa. Choiceless Polynomial Time on structures with small Abelian colour classes. In *MFCS 2014*, volume 8634 of *Lecture Notes in Computer Science*, pages 50–62. Springer, 2014.

2 A. Atserias, A. Bulatov, and A. Dawar. Affine systems of equations and counting infinitary logic. *Theoretical Computer Science*, 410:1666–1683, 2009.

3 G. Buntrock, U. Hertrampf, C. Damm, and C. Meinel. Structure and importance of logspace-mod-classes. *STACS'91*, pages 360–371, 1991.

4 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.

5 A. Dawar. The nature and power of fixed-point logic with counting. *ACM SIGLOG News*, pages 8–21, 2015.

6 A. Dawar, E. Grädel, B. Holm, E. Kopczynski, and W. Pakusa. Definability of linear equation systems over groups and rings. *Logical Methods in Computer Science*, 9(4), 2013.

7 A. Dawar, M. Grohe, B. Holm, and B. Laubner. Logics with Rank Operators. In *LICS'09*, pages 113–122. IEEE Computer Society, 2009.

8 A. Dawar and B. Holm. Pebble games with algebraic rules. In *Automata, Languages, and Programming*, pages 251–262. Springer, 2012.

9 H.-D. Ebbinghaus and J. Flum. *Finite model theory*. Springer-Verlag, 2nd edition, 1999.

10 E. Grädel et al. *Finite Model Theory and Its Applications*. Springer, 2007.

11 E. Grädel and W. Pakusa. Rank logic is dead, long live rank logic! *CoRR*, abs/1503.05423, 2015.

12 M. Grohe. The quest for a logic capturing PTIME. In *LICS 2008*, pages 267–271, 2008.

13 Y. Gurevich and S. Shelah. On finite rigid structures. *The Journal of Symbolic Logic*, 61(02):549–562, 1996.

14 B. Holm. *Descriptive complexity of linear algebra*. PhD thesis, Univ. of Cambridge, 2010.

15 B. Laubner. *The structure of graphs and new logics for the characterization of Polynomial Time*. PhD thesis, Humboldt-Universität Berlin, 2011.

16 M. Otto. *Bounded Variable Logics and Counting*. Springer, 1997.

17 W. Pakusa. Finite model theory with operators from linear algebra. Staatsexamensarbeit, RWTH Aachen University, 2010.

18 J. Torán. On the hardness of graph isomorphism. *SIAM Journal on Computing*, 33(5):1093–1108, 2004.