

Rank logic is dead, long live rank logic!

Erich Grädel and Wied Pakusa

Mathematical Foundations of Computer Science, RWTH Aachen University
 {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.

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 [15, 16, 18] 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 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, 11]. 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 is strong enough to express most of the algorithmic techniques leading to polynomial-time procedures [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

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

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 contains for every prime $p \in \mathbb{P}$ a *rank operator* rk_p which can be used to form a *rank term* $[\text{rk}_p \varphi(x, y)]$ 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, $\text{rk}(M) = \text{rk}(M | \vec{b})$) [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 [12, 15]. 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, 15] 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 [15] 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. [15, 8, 16].

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, 15, 16]). 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 [20], 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 [15] 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, 18]). 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 $\text{FPS} \leq \text{FPR}$ 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.

Let us briefly sketch the main idea of our proofs which is to exploit the symmetries of definable linear equation systems. To this end, let $M \cdot \vec{x} = \mathbb{1}$ be a linear equation system over some prime field \mathbb{F}_p where M is an $I \times I$ -matrix over \mathbb{F}_p and where $\mathbb{1}$ is the I -identity vector

over \mathbb{F}_p , i.e. $\mathbb{1}(i) = 1$ for all $i \in I$. Moreover, let Γ be a group which acts on I and which stabilises M , i.e. for all $i, j \in I$ and $\pi \in \Gamma$ we have $M(i, j) = M(\pi(i), \pi(j))$. In other words, if we identify the elements $\pi \in \Gamma$ with $I \times I$ -permutation matrices Π then we have $\Pi \cdot M = M \cdot \Pi$. Now let $\vec{b} \in \mathbb{F}_p^I$ be a solution of the linear equation system $M \cdot \vec{x} = \mathbb{1}$. Then we observe that also $\Pi \cdot \vec{b}$ is a solution for $\pi \in \Gamma$ since

$$M \cdot (\Pi \cdot \vec{b}) = (M \cdot \Pi) \cdot \vec{b} = \Pi \cdot (M \cdot \vec{b}) = \Pi \cdot \mathbb{1} = \mathbb{1}.$$

In other words: the solution space of the linear equation system $M \cdot \vec{x} = \mathbb{1}$ is closed under the action of Γ . Such and similar observations will enable us to transform a given linear equation system into a considerably simpler linear system which still is equivalent to the original one.

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 $M : I \times J \rightarrow \mathbb{F}_p$ and an I -vector \vec{v} over \mathbb{F}_p is a mapping $\vec{v} : I \rightarrow \mathbb{F}_p$.

A *(linear) preorder* $\leq \subseteq A \times A$ on A is a reflexive, transitive and total binary relation. A preorder \leq induces a linear order on the classes of the associated equivalence relation $x \sim y := (x \leq y \wedge y \leq x)$. We write $A = C_0 \leq \dots \leq C_{n-1}$ to denote the decomposition of A into \sim -classes C_i which are linearly ordered by \leq as indicated.

We briefly recall the definitions of *first-order logic with counting* FOC and *(inflationary) fixed-point logic with counting* FPC which are the extensions of FO and FP by counting terms. 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, for both, the vertex and the number sort, a collection of typed first-order variables. We agree to use Latin letters x, y, z, \dots for variables which range over the vertices and Greek letters ν, μ, \dots for variables ranging over the numbers. Similarly, for second-order variables R 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

$$\left[\text{ifp } R\bar{x}\bar{\nu} \leq \bar{t}. (\varphi(\bar{x}, \bar{\nu})) \right] (\bar{x}, \bar{\nu})$$

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 $\bar{x}\bar{v}$ 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{v} . For a formula φ we define a *counting term* $s = [\#\bar{x}\bar{v} \leq \bar{t}. \varphi]$ 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 $\ell = |\bar{v}|$.

We 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.

Extensions by rank operators

Next, we recall the notion of rank operators as introduced in [7]. Let $\Theta(\bar{x}\bar{v} \leq \bar{t}, \bar{y}\bar{\mu} \leq \bar{s})$ be a numeric term where \bar{t} and \bar{s} are tuples of closed numeric terms which bound the range of the numeric variables in the tuples \bar{v} and $\bar{\mu}$, respectively. Given a structure \mathfrak{A} we define $\mathbb{N}^{\leq \bar{t}} := \{\bar{n} \in \mathbb{N}^{|\bar{v}|} : n_i \leq t_i^{\mathfrak{A}}\}$. The set $\mathbb{N}^{\leq \bar{s}} \subset \mathbb{N}^{|\bar{\mu}|}$ is defined analogously. The term Θ 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} that is 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 $p \in \mathbb{P}$. Firstly, 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* $[\text{rk}_p(\bar{x}\bar{v} \leq \bar{t}, \bar{y}\bar{\mu} \leq \bar{s}). \Theta]$ 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 more flexible rank operator rk^* which gets the prime p as an additional input. Formally, with this rank operator rk^* we can construct a rank term $[\text{rk}^*(\bar{x}\bar{v} \leq \bar{t}, \bar{y}\bar{\mu} \leq \bar{s}, \pi \leq r). \Theta]$ 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 this rank term is the matrix rank of $(M_\Theta \bmod p)$ considered as a matrix over \mathbb{F}_p . The rank operator rk^* can be seen as a unification for the family of rank operators $(\text{rk}_p)_{p \in \mathbb{P}}$ and has been introduced in [15, 16, 18].

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

$$[\#x. \varphi(x)] = [\text{rk}_p(x, y). (x = y \wedge \varphi(x))].$$

Hence, we could equivalently define the rank logics FOR_Ω , FPR_Ω and FPR^* as the extensions of (the two-sorted variants of) FO and FP, respectively.

Extensions by solvability quantifiers

It is well-known that the extensions of FOC and FPC by matrix rank operators have surprising expressive power which, in particular, goes beyond that of fixed-point logic with counting. For example, it is known that rank operators can easily define the symmetric transitive closure of binary relations and that they can be used to express the structure isomorphism problem on various classes on which the Weisfeiler-Lehman test fails like, for example, classes of Cai, Fürer and Immerman graphs [4, 7]. Interestingly, such results for rank logic were obtained by reducing the respective queries to a *solvability problem for linear*

equation system over finite fields. Although the solvability problem (for linear equation systems) can be defined in rank logic, we propose to study extensions by quantifiers which directly express this solvability problem. One advantage of this approach is that one can naturally define such quantifiers for linear systems over more general classes of algebraic 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 $\varphi(\bar{x}\bar{v}, \bar{y}\bar{\mu}, \bar{z}) \in \text{FPS}_\Omega$ and let \bar{t} and \bar{s} be tuples of closed numeric terms with $|\bar{t}| = |\bar{v}|$ and $|\bar{s}| = |\bar{\mu}|$. Then also $\psi(\bar{z}) = (\text{slv}_p \bar{x}\bar{v} \leq \bar{s}, \bar{y}\bar{\mu} \leq \bar{t})\varphi(\bar{x}\bar{v}, \bar{y}\bar{\mu}, \bar{z})$ 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 $M_\varphi \cdot \bar{x} = \mathbb{1}$ over \mathbb{F}_p where $\bar{x} = (x_j)_{j \in J}$ 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, $M_\varphi \cdot \bar{x} = \mathbb{1}$ 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. Then the following inclusions easily follow from the definitions and the fact that rank operators can be used to define the solvability problem for linear equation systems.

$$\begin{array}{ccccccc}
 \text{FOR}_p & \leq & \text{FPR}_p & \leq & \text{FPR} & \leq & \text{FPR}^* & \leq & \text{PTIME} \\
 \vee & & \vee & & \vee & & \vee & & \\
 \text{FOS}_p & \leq & \text{FPS}_p & \leq & \text{FPS} & \leq & \text{FPS}^* & & \\
 & & \vee & & & & & & \\
 & & \text{FPC} & & & & & &
 \end{array}$$

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 [15] that this does not alter the expressive power of the resulting logics since solvability quantifiers ψ over a finite field \mathbb{F}_q of prime power order $q = p^k$ can be simulated by solvability quantifiers over \mathbb{F}_p . Moreover, a similar reduction can be achieved for rank operators which altogether shows that it suffices 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 of fixed-point logic with counting by solvability quantifiers and rank operators over different prime fields. Specifically, we show that the

expressive power of the logics FPS_Ω is different for all sets of primes Ω . Moreover, we transfer these results to the extensions FPR_Ω by rank operators. In this way we can answer the following open question about rank logic:

It holds that $\text{FPR}_p \neq \text{FPR}_q$ for pairs of different primes p, q . [8, 15, 16]

Another important consequence of our result is that rank logic (in the way it was defined in [7]) does not suffice to capture polynomial time. Let us state these results formally.

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

► **Theorem 2.** *Rank logic fails to capture polynomial time. We have $\text{FPR} < \text{FPR}^* \leq \text{PTIME}$.*

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) $\text{FPS}_\Omega = \text{FPC}$ on \mathcal{K}_q for every set of primes Ω with $q \notin \Omega$,
- (b) $\text{FPR}_\Omega = \text{FPS}_\Omega$ on \mathcal{K}_q for all sets of primes Ω ,
- (c) $\text{FPC} < \text{PTIME}$ on \mathcal{K}_q , and
- (d) $\text{FPS}_q = \text{PTIME}$ on \mathcal{K}_q .

Proof of Theorem 1. Let Ω and Ω' be two sets of primes as above. Without loss of generality let us assume that there exists a prime $q \in \Omega \setminus \Omega'$. Then by Theorem 3 there exists a class \mathcal{K}_q on which $\text{FPS}_\Omega = \text{FPR}_\Omega = \text{PTIME}$ and on which $\text{FPS}_{\Omega'} = \text{FPR}_{\Omega'} = \text{FPC} < \text{PTIME}$. ◀

Proof of Theorem 2. Otherwise assume that $\text{FPR} = \text{PTIME}$. Then, in particular, $\text{FPR} = \text{FPR}^*$ and there exists a formula $\varphi \in \text{FPR}$ 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 $\varphi \in \text{FPR}_\Omega$ 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 $\text{FPS} \leq \text{FPR} \leq \text{FPR}^* \leq \text{FPR}_\Omega$.

Now let $q \in \mathbb{P} \setminus \Omega$. By Theorem 3 there exists a class of structures \mathcal{K}_q on which $\text{FPR}_\Omega = \text{FPC} < \text{PTIME}$. However, the class \mathcal{K}_q can be chosen such that $\text{PTIME} = \text{FPS}_q \leq \text{FPR}_\Omega$ 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 [15, 16, 18].

The remainder of this section is devoted to the proof of Theorem 3. We fix a prime q and proceed as follows. In a first step, we identify properties of classes of structures \mathcal{K}_q which guarantee that the relations claimed in (a), (b), (c) and (d) hold. In a second step, we proceed to show that we can obtain a class of structures \mathcal{K}_q that satisfies all of these sufficient criteria. This together then proves our theorem.

Establishing sufficient criteria

We start by establishing sufficient criteria for the most relevant part of Theorem 3 which is the relation claimed in (a). Assume that we have a class of structures $\mathcal{K}_q = \mathcal{K}$ with the following properties.

- (I) The automorphism groups $\Delta_{\mathfrak{A}} := \text{Aut}(\mathfrak{A})$ 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.

Formally, for each $\ell \geq 1$ there is a formula $\varphi_{\leq}(x_1, \dots, x_\ell, y_1, \dots, y_\ell) \in \text{FPC}$ such that for every structure $\mathfrak{A} \in \mathcal{K}$, 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 $\text{FPS}_\Omega = \text{FPC}$ on \mathcal{K} for all $\Omega \subseteq \mathbb{P} \setminus \{q\}$.*

The proof of this lemma is by induction on the structure of FPS_Ω -formulas. Obviously, the only interesting step is the translation of a solvability formula

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

into an FPC-formula $\vartheta(\bar{z})$ which is equivalent to $\psi(\bar{z})$ on the class \mathcal{K} . Let $|\bar{x}| = |\bar{y}| = \ell$, $|\bar{v}| = |\bar{\mu}| = \lambda$ 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 the variable tuple \bar{z} . According to the semantics of the solvability quantifier, the formula φ defines in $(\mathfrak{A}, \bar{z} \mapsto \bar{c})$ an $I \times J$ -matrix $M = M_{\bar{c}}^{\mathfrak{A}}$ with entries in $\{0, 1\} \subseteq \mathbb{F}_p$ where $I = I_{\bar{c}}^{\mathfrak{A}} := A^\ell \times \mathbb{N}^{\leq s} \subseteq A^\ell \times \mathbb{N}^\lambda$ and $J = J_{\bar{c}}^{\mathfrak{A}} := A^\ell \times \mathbb{N}^{\leq t} \subseteq A^\ell \times \mathbb{N}^\lambda$ that is defined for $\bar{a} \in I$ and $\bar{b} \in J$ as

$$M(\bar{a}, \bar{b}) = \begin{cases} 1, & \text{if } \mathfrak{A} \models \varphi(\bar{a}, \bar{b}, \bar{c}) \\ 0, & \text{else.} \end{cases}$$

By definition we have $\mathfrak{A} \models \psi(\bar{c})$ if, and only if, the linear equation system $M \cdot \bar{x} = \mathbb{1}$ over \mathbb{F}_p is solvable. The key idea is to use the symmetries of the structure \mathfrak{A} to translate the linear equation system $M \cdot \bar{x} = \mathbb{1}$ into an equivalent linear system which is *simpler* in the sense that its solvability can be defined in the logic FPC. The reader should observe that each automorphism $\pi \in \Delta_{\mathfrak{A}} = \text{Aut}(\mathfrak{A})$ naturally induces an automorphism of the two-sorted extension $\mathfrak{A}^\#$ which point-wise fixes every number $n \in \mathbb{N}$. In particular we have $\text{Aut}(\mathfrak{A}) = \text{Aut}(\mathfrak{A}^\#)$.

We set $\Gamma = \Gamma_{\bar{c}}^{\mathfrak{A}} := \text{Aut}(\mathfrak{A}, \bar{c}) \leq \Delta = \Delta_{\mathfrak{A}} = \text{Aut}(\mathfrak{A})$. 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 which is defined as

$$\Pi_I(\bar{a}, \bar{b}) = \begin{cases} 1, & \pi(\bar{a}) = \bar{b} \\ 0, & \text{otherwise.} \end{cases}$$

Then Γ acts on the set of $I \times J$ -matrices by left multiplication with $I \times I$ -permutation matrices. Similarly, we let Π_J denote the $J \times J$ -permutation matrix defined as

$$\Pi_J(\bar{a}, \bar{b}) = \begin{cases} 1, & \pi(\bar{a}) = \bar{b} \\ 0, & \text{otherwise.} \end{cases}$$

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 $\Gamma = \text{Aut}(\mathfrak{A}, \bar{c})$ 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 of size one which proves our claim. ◀

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 proceed to use the property (II) and show that there exists an FPC-formula $\varphi_{\leq}(\vec{x}, \vec{y})$ which defines for all $\mathfrak{A} \in \mathcal{K}$ and $\vec{c} \in (A \uplus \mathbb{N})^k$ as above a linear preorder \leq on A^ℓ which identifies Γ -orbits. Note that, in general, $\Gamma = \text{Aut}(\mathfrak{A}, \vec{c})$ is a strict subgroup of $\Delta = \text{Aut}(\mathfrak{A})$. 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 $\{c_1, \dots, c_k\} \cap A$.

The linear preorder \leq naturally extends to a preorder on the sets I and J with the same properties. Let us write $J = J_0 \leq J_1 \leq \dots \leq J_{v-1}$ to denote the decomposition of J into the Γ -orbits J_j which are ordered by \leq 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

$$e_j(i) := \begin{cases} 1, & \text{if } i \in J_j \\ 0, & \text{else.} \end{cases}$$

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 linear equation system $M \cdot \vec{x} = \mathbb{1}$ is solvable if, and only if, the linear equation system $(M \cdot E) \cdot \vec{x}_* = \mathbb{1}$ 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 [18]. 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 \mathcal{K} . In fact, for this translation we only require the somewhat weaker assumption that we can define in FPC on ℓ -tuples in structures $\mathfrak{A} \in \mathcal{K}$ a linear preorder in which every class can be linearised in FPC by fixing a constant number of parameters. The precise technical requirements will become clear from the proof of the following lemma.

► **Lemma 7.** *If \mathcal{K} satisfies (I) and (II), then $\text{FPR}_\Omega = \text{FPS}_\Omega$ on \mathcal{K} for all sets of primes Ω .*

Proof. We inductively translate FPR_Ω -formulas into formulas of FPS_Ω which are equivalent on the class \mathcal{K} . The only interesting case is the transformation of rank terms

$$\Upsilon(\vec{z}) = [\text{rk}_p(\vec{x}\vec{v} \leq \vec{t}, \vec{y}\vec{\mu} \leq \vec{s}) \cdot \Theta(\vec{x}\vec{v}, \vec{y}\vec{\mu}, \vec{z})].$$

Let $|\bar{x}| = |\bar{y}| = \ell$, $|\bar{v}| = |\bar{\mu}| = \lambda$ 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 the type of the variable tuple \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

$$M(\bar{a}\bar{n}, \bar{b}\bar{m}) := \Theta^{\mathfrak{A}}(\bar{a}\bar{n}, \bar{b}\bar{m}, \bar{c}) \bmod p.$$

According to the semantics of matrix rank operators, the value $\Upsilon^{\mathfrak{A}}(\bar{c}) \in \mathbb{N}$ is the rank of the matrix M . We proceed to show that we can determine the matrix rank of M by a recursive application of solvability queries. To this end we 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 \leq on $J^{\mathfrak{A}}$, and such that
- (b) for every \leq -class $[j] \subseteq J^{\mathfrak{A}}$ there exists a parameter tuple $\bar{d} \in A^{|\bar{v}|}$ such that $\psi_{\leq}^{\mathfrak{A}}(\bar{d})$ is a linear order \leq on $[j]$.

Proof of claim: First of all, we let φ_{\leq} be an FPC-formula which defines in every structure $\mathfrak{A} \in \mathcal{K}$ a linear preorder \leq on $J^{\mathfrak{A}}$ such that \leq -classes correspond to $\Delta_{\mathfrak{A}}$ -orbits. Such a formula exists by our assumption that \mathcal{K} satisfies property (II). Analogously, we choose an FPC-formula ϑ_{\leq} which defines in every structure $\mathfrak{A} \in \mathcal{K}$ a linear preorder \leq^* on $J^{\mathfrak{A}} \times J^{\mathfrak{A}}$ that induces a linear order on the $\Delta_{\mathfrak{A}}$ -orbits.

Now let $[j] \subseteq J^{\mathfrak{A}}$ be a \leq -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 \leq^* to elements in $\{j\} \times [j]$ corresponds to a linear order on $[j]$ for each $j \in [j]$. In this way we obtain an FPC-formula ψ_{\leq} with the desired properties. \dashv

We are now prepared to describe the recursive procedure which allows us to determine the rank of the matrix M in FPS_{Ω} . To this end we fix formulas φ_{\leq} and ψ_{\leq} with the above properties. Moreover, let \leq denote the linear preorder defined by φ_{\leq} on J and let

$$J = J_0 \leq J_1 \leq \dots \leq J_{r-1}.$$

We use the formula ψ_{\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 $\vec{m}_j \in \mathbb{F}_p^J$ 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 $V_i := \{\vec{m}_j : j \in J_0 \cup J_1 \cup \dots \cup J_i\}$ as follows. First, we use ψ_{\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 $\langle V_i \rangle = \langle V_{i-1} \rangle \uplus \langle W \rangle$. 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. \blacktriangleleft

We now focus on the parts (c) and (d) of Theorem 3 and establish sufficient criteria which guarantee that FPC fails to capture PTIME on \mathcal{K} while FPS_q can express every polynomial-time decidable property of \mathcal{K} -structures.

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

(IV) For every $k \geq 1$ there exists a pair of structures $\mathfrak{A} \in \mathcal{K}$ and $\mathfrak{B} \in \mathcal{K}$ such that $\mathfrak{A} \not\equiv \mathfrak{B}$ and $\mathfrak{A} \equiv_k^C \mathfrak{B}$.

► **Lemma 8.** *If \mathcal{K} satisfies (III) and (IV), then $\text{FPC} < \text{FPS}_q = \text{PTIME}$ on \mathcal{K} .*

Proof. It is clear that by property (III) we have $\text{FPS}_q = \text{PTIME}$ on \mathcal{K} . Moreover, if we had $\text{FPC} = \text{PTIME}$ on \mathcal{K} then, by the embedding of FPC into $C_{\infty\omega}^\omega$ and the fact that \mathcal{K} -structures can be canonised in polynomial time, there exists a fixed $k \geq 1$ such that $C_{\infty\omega}^k$ can identify each structure in \mathcal{K} which, in turn, contradicts property (IV). ◀

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 cyclic groups other than \mathbb{F}_2 . To illustrate the main differences, let us briefly recall the idea of the original construction. Starting from an undirected and connected graph $\mathcal{G} = (V, E)$, we take two copies e_0, e_1 of every edge $e \in E$. Moreover, for every vertex $v \in V$ we consider the set $vE \subseteq E$ of edges which are adjacent to v and we add one of the following two constraints to restrict the symmetries of the resulting CFI-graph: either the set of all sets $\{e_{\rho(e)} : e \in vE\}$ with $\rho : vE \rightarrow \mathbb{F}_2$ and $\sum_{e \in vE} \rho(e) = 0$ is stabilised (an *even* node) or the dual set of all sets $\{e_{\rho(e)} : e \in vE\}$ with $\rho : vE \rightarrow \mathbb{F}_2$ and $\sum_{e \in vE} \rho(e) = 1$ is stabilised (an *odd* node). These constraints are encoded by a simple graph gadget. Although it seems that for each of these exponentially many choices we obtain a different CFI-graph, there really are, up to isomorphism, only two such graphs which in turn are determined by the parity of the number of odd nodes. Very roughly, the reason is that we can transpose, or *twist*, two copies e_0, e_1 of each an edge e and move this twist along a path (in the *connected* graph \mathcal{G}) to iteratively resolve pairs of odd nodes.

In order to generalise this construction to \mathbb{F}_q we take for every edge $e \in E$ a *directed cycle* of length q over q copies e_0, e_1, \dots, e_{q-1} of the edge e . We then add similar constraints for sets of incident edges as above, but naturally, instead of having only two different kinds of such constraints, we have one for each value $0, 1, \dots, q-1 \in \mathbb{F}_q$. Now, instead of twisting pairs of edges, we consider *cyclic shifts* of length $\leq q$ on the edge classes e_0, e_1, \dots, e_{q-1} which respect the cycle relation. Again, these shifts can be propagated along paths in the original graph \mathcal{G} and, with a reasoning analogous to the original approach, it turns out that there are, up to isomorphism, only q different types of generalised CFI-graphs over \mathbb{F}_q . We remark that the same kind of generalisations has been studied, for example, in [15, 20].

Let us formalise the above described intuitions. We start with an (*undirected*), *connected* and *ordered* graph $\mathcal{G} = (V, \leq, E)$. Let C, I and R be binary relation symbols. We set $\tau := \{\leq, C, I, R\}$. 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 $\text{CFI}_q(\mathcal{G}, \vec{d})$ 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 construction is illustrated in Figure 1.

- The *universe* of $\text{CFI}_q(\mathcal{G}, \vec{d})$ 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 copies 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 $\rho : E(v) \rightarrow [q]$ which satisfy $\sum \rho := \sum_{e \in E(v)} \rho(e) = d$.
- The *linear preorder* \leq orders the edge classes according to the linear order induced by \leq on E . More precisely, we let $\hat{e} \leq \hat{f}$ whenever $e \leq f$. Similarly, \leq orders the equation classes according to the order of \leq on V , i.e. $\hat{v} \leq \hat{w}$ if $v \leq w$. Moreover, we let $\hat{e} \leq \hat{v}$ 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

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



■ **Figure 1** Generalised CFI-construction for the v -gadget where $q = 3$ and $\vec{d}(v) = 0$

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 $\text{CFI}_q(\mathcal{G}, \vec{d})$. However, for each graph \mathcal{G} with the above properties there really are, up to isomorphism, only q different

CFI-structures $\text{CFI}_q(\mathcal{G}, \vec{d})$. In fact, the value $\sum \vec{d} := \sum_{v \in V} \vec{d}(v)$ completely determines the isomorphism class of a CFI-structure over \mathcal{G} .

To obtain this characterisation, we analyse the automorphism group of CFI-structures and, more generally, the set of isomorphisms between two structures $\mathfrak{A} = \text{CFI}_q(\mathcal{G}, \vec{d}_1)$ and $\mathfrak{B} = \text{CFI}_q(\mathcal{G}, \vec{d}_2)$. For such structures we know that the set \hat{E} of edge nodes, the linear preorder \leq on \hat{E} , the cycle relation C and the inverse relation I do not depend on the sequence of gadget values. This means that each possible isomorphism π which maps \mathfrak{A} to \mathfrak{B} induces an automorphism of the common substructure $\mathfrak{C} := (\hat{E}, (\leq \upharpoonright \hat{E}), C, I)$ which only depends on \mathcal{G} but not on $\vec{d} \in [q]^V$. Thus

$$(\text{Iso}(\mathfrak{A}, \mathfrak{B}) \upharpoonright \hat{E}) \subseteq \Gamma := \text{Aut}(\mathfrak{C}) \leq \text{Sym}(\hat{E}).$$

Let $\pi \in \Gamma$. The linear preorder \leq on \hat{E} and the cycle relation C enforce that π is the composition of cyclic shifts on the individual edge classes \hat{e} , i.e. $\pi \in \prod_{e \in E} (\langle e_0 e_1 \dots e_{q-1} \rangle) \leq \text{Sym}(\hat{E})$. It is convenient to identify the group $\prod_{e \in E} (\langle e_0 e_1 \dots e_{q-1} \rangle)$ with the vector space \mathbb{F}_q^E in the obvious way.

In addition, the inverse relation I enforces that cyclic shifts for pairs of related edge classes are inverse to each other in the following sense: let $e = (v, w) \in E$ and $f = (w, v) \in E$ be a pair of related edges. Assume that we have a permutation $\pi \in \mathbb{F}_q^E$ such that $\pi(e) = x$ and $\pi(f) = y$. We have $(e_0, f_0) \in I$. Hence, if π is supposed to be an automorphism of \mathfrak{C} then we have $\pi(I) = I$ and thus $(e_x, e_y) \in I$ which means that $x + y = 0$.

In conclusion, it follows that $\Gamma \leq \mathbb{F}_q^E$ is the subgroup of \mathbb{F}_q^E which contains all E -vectors $\pi \in \mathbb{F}_q^E$ with the property that $\pi(e) + \pi(f) = 0$ for pairs of related edges $e, f \in E$. Again we remind the reader that Γ only depends on \mathcal{G} but not on $\vec{d} \in [q]^V$. If we want to stress this dependence, then we sometimes write $\Gamma(\mathcal{G})$ but usually we omit \mathcal{G} if the graph is clear from the context.

Now, given a CFI-structure $\mathfrak{A} = \text{CFI}_q(\mathcal{G}, \vec{d})$, we define for each vertex $v \in V$ the v -gadget as the set $\text{gadget}(v) := \hat{v}^{d(v)} \uplus \bigcup_{e \in E(v)} \hat{e}$.

► **Lemma 9.** *Let $\mathfrak{A} = \text{CFI}_q(\mathcal{G}, \vec{d})$ and let $\pi \in \Gamma$. Then there is precisely one extension $\hat{\pi}$ of π to $\hat{E} \uplus \hat{V}$ such that $\hat{\pi}(\mathfrak{A})$ is a CFI-structure over \mathcal{G} .*

Proof. Let $\rho \in \hat{v} = \hat{v}^{d(v)}$ for some $v \in V$. We show that under the assumption that $\hat{\pi}(\mathfrak{A})$ is a CFI-structure over \mathcal{G} the action of π on \hat{E} determines $\hat{\pi}(\rho)$.

We have that $(\rho, e_{\rho(e)}) \in R$ for all $e \in E(v)$. Hence for a potential isomorphism $\hat{\pi}$ we must have that $(\hat{\pi}(\rho), \pi(e_{\rho(e)})) \in R'$ (for a gadget relation R' of a CFI-structure over \mathcal{G}). Since we have $\pi(e_{\rho(e)}) = e_{\rho(e) + \pi(e)}$, it follows by the definition of CFI-structures that the function $\hat{\pi}(\rho) : E(v) \rightarrow [q]$ is determined as $(\hat{\pi}(\rho))(e) = \rho(e) + \pi(e)$ which in turn only depends on the action of π on the edge classes \hat{e} for $e \in E(v)$. ◀

The preceding lemma shows that $\text{Iso}(\mathfrak{A}, \mathfrak{B})$ can be identified with a subset of Γ . In fact, the set $\text{Aut}(\mathfrak{A})$ turns out to be a subgroup of Γ of which $\text{Iso}(\mathfrak{A}, \mathfrak{B})$ is a coset in Γ . Specifically, we saw that every $\pi \in \Gamma$ can uniquely be identified with an isomorphism of CFI-structures over \mathcal{G} by setting $\pi(\rho) = \rho + \pi$ for $\rho \in \hat{v}^d$. As a consequence, this means that $\pi(\hat{v}^d) = \hat{v}^{d_*}$ where $d_* = d + \sum_{e \in E(v)} \pi(e)$ and that

$$\pi(R_v^d) = \{(\rho + \pi, e_{\rho(e) + \pi(e)}) : (\rho, e_{\rho(e)}) \in R_v^d\} = R_v^{d_*}.$$

In particular, π stabilises the relation R_v^d if, and only if, $\sum_{e \in E(v)} \pi(e) = 0$.

► **Lemma 10.** Γ acts on $\{\text{CFI}_q(\mathcal{G}, \vec{d}) : \vec{d} \in [q]^V\}$. For $\pi \in \Gamma$ we have

$$\pi(\text{CFI}_q(\mathcal{G}, \vec{d})) = \text{CFI}_q(\mathcal{G}, \vec{d}_*) \text{ where } \vec{d}_*(v) = (\vec{d}(v) + \sum_{e \in E(v)} \pi(e)).$$

► **Lemma 11.** Let $\vec{d}, \vec{d}_* \in ([q])^V$ be two sequences of gadget values. Then $\text{CFI}_q(\mathcal{G}, \vec{d}) \cong \text{CFI}_q(\mathcal{G}, \vec{d}_*)$ if, and only if, $\sum \vec{d} = \sum \vec{d}_*$.

Proof. Let $\pi \in \Gamma$ such that $\pi(\text{CFI}_q(\mathcal{G}, \vec{d})) = \text{CFI}_q(\mathcal{G}, \vec{d}_*)$. By Lemma 10 this means that $\vec{d}_*(v) = (\vec{d}(v) + \sum_{e \in E(v)} \pi(e))$ for $v \in V$. Thus $\sum_{v \in V} \vec{d}_*(v) = \sum_{v \in V} \vec{d}(v) + \sum_{v \in V} \sum_{e \in E(v)} \pi(e) = \sum_{v \in V} \vec{d}(v) + \sum_{e \in E} \pi(e)$. Since for all pairs of related edges $e, f \in E$ we have $\pi(e) + \pi(f) = 0$ the claim follows.

For the other direction we proceed by induction on the number i of vertices $v \in V$ such that $\vec{d}(v) \neq \vec{d}_*(v)$. If no such vertex exists, then the claim is trivial. Otherwise, because of our assumption, there exist at least two such vertices $v, w \in V$, $v \neq w$. Since \mathcal{G} is connected we find a simple path

$$\bar{p} : v = v_0 \xrightarrow{E} v_1 \xrightarrow{E} v_2 \xrightarrow{E} \dots \xrightarrow{E} v_m = w$$

from v to w of length $m \geq 1$. Consider the following E -vector $\pi \in \mathbb{F}_q^E$ which is defined for $z := \vec{d}_*(v) - \vec{d}(v)$ as

$$\pi(e) := \begin{cases} z, & \text{if } e = (v_i, v_{i+1}), 0 \leq i < m \\ -z, & \text{if } e = (v_{i+1}, v_i), 0 \leq i < m \\ 0, & \text{else.} \end{cases}$$

By the definition of π it follows that $\pi \in \Gamma$. Let $\pi(\text{CFI}_q(\mathcal{G}, \vec{d})) = \text{CFI}_q(\mathcal{G}, \vec{d}_+)$. We claim that the number of $v \in V$ such that $\vec{d}_+(v) \neq \vec{d}_*(v)$ is at most $i - 1$. From Lemma 10 we know that $\vec{d}_+(v) = \vec{d}(v) + \sum_{e \in E(v)} \pi(e)$. For $v \in V$ it follows that

- if $v \notin \{v_0, \dots, v_m\}$, then $\vec{d}_+(v) = \vec{d}(v)$, and
- if $v = v_0$, then $\vec{d}_+(v) = \vec{d}(v) + z = \vec{d}_*(v)$, and
- if $v = v_j$ for $1 \leq j < m$, then

$$\vec{d}_+(v) = \vec{d}(v) + \pi(v_j, v_{j-1}) + \pi(v_j, v_{j+1}) = \vec{d}(v) - z + z = \vec{d}(v), \text{ and}$$

- if $v = v_m$, then $\vec{d}_+(v) = \vec{d}(v) - z$.

Thus the claim follows from the induction hypothesis. ◀

The kind of isomorphism that we constructed in the proof of Lemma 11 plays an important role later on. Thus, for a simple path \bar{p} from v_0 to v_m ($m \geq 1$)

$$\bar{p} : v = v_0 \xrightarrow{E} v_1 \xrightarrow{E} v_2 \xrightarrow{E} \dots \xrightarrow{E} v_m = w$$

as above and a constant $z \in \mathbb{F}_q$ we denote this isomorphism by $\pi[\bar{p}, z] \in \Gamma$. In other words, if we let $\sigma^z[e] \in \Gamma$ for $e \in E$ and $z \in \mathbb{F}_q$ denote the E -vector which is defined as

$$\sigma^z[e](f) = \begin{cases} z, & \text{if } f = e, \\ -z, & \text{if } f = e^{-1}, \\ 0, & \text{else,} \end{cases}$$

then $\pi[\bar{p}, z] = \sigma^z[(v_0, v_1)] + \sigma^z[(v_1, v_2)] + \dots + \sigma^z[(v_{m-1}, v_m)]$. Intuitively, the isomorphism $\pi[\bar{p}, z]$ allows us to simultaneously increase the gadget value at v_0 by z and to decrease the

gadget value at v_m by z while the induced twists are moved along the path \bar{p} through the gadget relations of the vertices v_j , $1 \leq j < m$, whose gadget value does not change. A very important special case arises when \bar{p} is a simple cycle of length $m \geq 3$

$$\bar{p}: v = v_0 \xrightarrow{E} v_1 \xrightarrow{E} v_2 \xrightarrow{E} \dots \xrightarrow{E} v_m = v.$$

Then for all values $z \in \mathbb{F}_q$ the isomorphism $\pi[\bar{p}, z] \in \Gamma$ is an *automorphism* of CFI-structures over \mathcal{G} . We are going to use these automorphisms to show that it is possible to define in FPC an ordering on the orbits of ℓ -tuples as required by property (II). It turns out that it therefore suffices to ensure that the graph \mathcal{G} is sufficiently connected.

Recall that a graph \mathcal{G} is *k-connected*, for $k \geq 1$, 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* $\text{con}(\mathcal{G})$ of a graph \mathcal{G} is the maximal $k \geq 1$ such that \mathcal{G} is k -connected. Moreover, the *connectivity* $\text{con}(\mathfrak{G})$ of a class \mathfrak{G} of graphs is the function $\text{con}(\mathfrak{G}) : \mathbb{N} \rightarrow \mathbb{N}$ defined by

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

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

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

Verifying the required properties

We proceed to show that \mathcal{K} satisfies the required properties (I) - (IV).

First of all, we saw that the automorphism group of each CFI-structure $\text{CFI}_q(\mathcal{G}, \vec{d})$ is a \mathbb{F}_q -vector space, so property (I) clearly holds for the class \mathcal{K} .

The proof that \mathcal{K} satisfies property (II) is more involved. Let us 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 choice of \mathcal{K} it suffices to consider CFI-structures $\mathfrak{A} = \text{CFI}_q(\mathcal{G}, \vec{d})$ over graphs $\mathcal{G} = (V, \leq, E)$ with $\text{con}(\mathcal{G}) > (\ell + 2)$ since almost all structures in \mathcal{K} satisfy this condition. As above let $\Gamma \leq \mathbb{F}_q^E$ denote the group that acts on the set of CFI-structures over \mathcal{G} and let $A := (\hat{V} \uplus \hat{E})$ denote the universe of the CFI-structure \mathfrak{A} .

► **Definition 12.** Let $\lambda \leq \ell$ and let $\bar{a} \in A^\lambda$.

- (i) Let $v \in V$. We say that the vertex v is *marked* (given the parameters \bar{a}) if for some $x \in \{a_1, \dots, a_\lambda\}$ we have $x \in \hat{v}$ ($= \hat{v}^{\bar{a}(v)}$).
- (ii) Let $e = (v, w) \in E$. We say that the edge e is *marked* (given the parameters \bar{a}) if one of the vertices v or w is marked or if for some $x \in \{a_1, \dots, a_\lambda\}$ we have that $x \in \hat{e} \cup \hat{f}$ where $f = (w, v) \in E$ is the edge related with e .

► **Lemma 13.** Let $\lambda \leq \ell$ and let $\bar{a} \in A^\lambda$.

- (a) If $v \in V$ is marked, then the v -gadget can be identified in $\mathbb{C}_{\infty\omega}^{\ell+2}$ (using the parameters \bar{a}), i.e. for every $c \in \text{gadget}(v)$ there exists a formula $\vartheta(\bar{x}, y) \in \mathbb{C}_{\infty\omega}^{\ell+2}$ such that $\vartheta^{\mathfrak{A}}(\bar{a}) = \{c\}$.
- (b) If an edge $e \in E$ is marked, then the edge classes \hat{e} and \hat{f} for $f = e^{-1}$ are identified in $\mathbb{C}_{\infty\omega}^{\ell+2}$ (given the parameters \bar{a}), i.e. for every $c \in \hat{e} \uplus \hat{f}$ there exists a formula $\vartheta(\bar{x}, y) \in \mathbb{C}_{\infty\omega}^{\ell+2}$ such that $\vartheta^{\mathfrak{A}}(\bar{a}) = \{c\}$.

Proof. First of all, it is straightforward (even without using the parameters) to fix the \leq -class of any element $c \in A$ in $C_{\infty\omega}^{\ell+2}$. Secondly, observe that if an element $\rho \in \hat{v}$ is fixed, then we can fix an element in each of the edge classes \hat{e} for $e \in E(v)$ since ρ is R -connected to precisely one vertex in each of these classes. Moreover, if we have fixed an element $x \in \hat{e}$ in some edge class \hat{e} , then we can simply use the cycle relation C to identify each element $c \in \hat{e}$ via its C -distance to a in $C_{\infty\omega}^{\ell+2}$. Finally, the inverse relation I yields a definable bijection between related edge classes. \blacktriangleleft

► **Lemma 14.** *Let $\lambda \leq \ell$, $\bar{a} \in A^\lambda$ and let $v \in V$ be a vertex that is not marked. Then for all edges $e, e' \in E(v)$, $e \neq e'$ which are not marked there exists $\pi \in \text{Aut}(\mathfrak{A}, \bar{a})$ such that $\pi(e) = -\pi(e') \neq 0$ and such that $\pi(f) = 0$ for all $f \in E(v) \setminus \{e, e'\}$.*

Proof. Let $e = (v, w)$ and $e' = (v, w')$ as above. Then the vertices w and w' are not marked.

Consider the graph \mathcal{G}' that results from \mathcal{G} by removing the vertex v and each marked vertex $y \in V$. Let $V' \subseteq V$ denote the vertex set and $E' \subseteq E$ the edge relation of the graph \mathcal{G}' . Moreover, let $M := \{a_1, \dots, a_\lambda\} \cap (\bigcup_{e \in E} \hat{e})$. We observe that $|V| - |V'| \leq \lambda - |M| + 1$.

For every $x \in M$ there is an edge $f \in E$ such that $x \in \hat{f}$. For each such edge f that is also contained in the subgraph \mathcal{G}' we delete one of its endpoints but *neither the vertex w nor the vertex w'* and denote the resulting subgraph by \mathcal{G}'' with vertex set $V'' \subseteq V'$ and edge relation $E'' \subseteq E'$. It still might happen that there is a parameter $x \in M$ such that $x \in \hat{f}$ for $f \in E''$. However, this can only occur if f connects w' and w . Since we removed at most $(|V| - |V'|) + |M| \leq \lambda + 1 \leq (\ell + 1)$ vertices from the graph \mathcal{G} to obtain \mathcal{G}'' and since $\text{con}(\mathcal{G}) > (\ell + 2)$ we know that there is a simple path of length $m \geq 2$ (i.e. the path does not consist of a single edge between w and w') which connects w and w' in \mathcal{G}'' :

$$\bar{p} : w \xrightarrow{E''} v_1 \xrightarrow{E''} v_2 \xrightarrow{E''} \dots \xrightarrow{E''} v_{m-1} \xrightarrow{E''} w'.$$

We extend \bar{p} to a simple cycle \bar{p}_c in \mathcal{G} from v to v by using the edges $(v, w), (v, w') \in E$:

$$\bar{p}_c : v \xrightarrow{E} w \xrightarrow{E} v_1 \xrightarrow{E} v_2 \xrightarrow{E} \dots \xrightarrow{E} v_{m-1} \xrightarrow{E} w' \xrightarrow{E} v.$$

Let $0 \neq z \in [q]$. We claim that $\pi := \pi[\bar{p}_c, z]$ satisfies the desired properties.

By the definition of π it holds that $\pi(e) = z = -\pi(e')$. Let $x \in \{a_1, \dots, a_\lambda\}$. Then we have $x \notin \bigcup_{i=1}^{m-1} \hat{v}_i \cup \hat{w} \cup \hat{w}' \cup \hat{v}$, since none of the vertices v, w and w' is marked and since we removed any other marked vertex $y \in V$ from \mathcal{G} .

Moreover, for $f \in \{(v, w), (w, v), (v, w'), (w', v)\}$ we have that $x \notin \hat{f}$ by our assumption that e, e' are not marked. Also for $f \in \{(w, v_1), (w', v_{m-1})\}$ we have $x \notin \hat{f}$ since otherwise we had removed the vertices v_1 and v_{m-1} from \mathcal{G}' . Finally, for $f \in \bigcup_{i=1}^{m-2} \{(v_i, v_{i+1}), (v_{i+1}, v_i)\}$ we have $x \notin \hat{f}$ since otherwise we had removed one of the endpoints of each such edge f from \mathcal{G}' . Hence $\pi(x) = x$. Finally, since $v \notin V''$ we also have that $\pi(f) = f$ for all $f \in E(v) \setminus \{e, e'\}$. \blacktriangleleft

► **Lemma 15.** *Let $\lambda \leq \ell$ and let $\bar{a}, \bar{b} \in A^\lambda$. Then $(\mathfrak{A}, \bar{a}) \equiv_{\ell+2}^C (\mathfrak{A}, \bar{b})$ if, and only if, there exists $\pi \in \text{Aut}(\mathfrak{A})$ such that $\pi(\bar{a}) = \bar{b}$.*

Proof. We proceed by induction on the maximal position $1 \leq i \leq \lambda$ up to which the tuples \bar{a} and \bar{b} agree, i.e. such that for $1 \leq j < i$ we have $a_j = b_j$ and such that $a_i \neq b_i$. Let $a := a_i$ and $b := b_i$. Then we have to show that there exists an automorphism $\pi \in \text{Aut}(\mathfrak{A}, a_1 \dots a_{i-1})$ such that $\pi(a) = b$. Since \bar{a} and \bar{b} have the same $C_{\infty\omega}^{\ell+2}$ -type we know that a and b belong to the same \leq -class. We choose $v \in V$ such that $a, b \in \text{gadget}(v)$.

In what follows, whenever we speak of *marked vertices* or *marked edges* then we implicitly refer to a marking with respect to the already fixed part of parameters $\{a_1, \dots, a_{i-1}\}$.

Without loss of generality we may assume that the vertex v is not marked (by an element $x \in \{a_1, \dots, a_{i-1}\}$), because otherwise, by Lemma 13, every element in $\text{gadget}(v)$ can uniquely be identified in $C_{\infty\omega}^{\ell+2}$. We distinguish between the two cases where a and b are equation nodes and where a and b are edge nodes.

For the first case let $a, b \in \hat{v}$. There exists a unique $\pi \in \mathbb{F}_q^{E(v)}$ such that $\pi(a) = b$ and such that $\sum_{e \in E(v)} \pi(e) = 0$. Moreover, this vector π can easily be defined in $C_{\infty\omega}^{\ell+2}$ given the elements a and b . Now assume that one of the edges $e = (v, w) \in E(v)$ is marked but that $\pi(e) \neq 0$. Since the edge e is marked, every element in \hat{e} can uniquely be identified in $C_{\infty\omega}^{\ell+2}$ by Lemma 13. However, since a and b are R -connected to *different* elements in \hat{e} (as $\pi(e) \neq 0$) this contradicts the fact that \bar{a} and \bar{b} have the same $C_{\infty\omega}^{\ell+2}$ -type. Thus, for every edge $e \in E(v)$ we either have that $\pi(e) = 0$ or that e is not marked. By induction on the number of edges $e \in E(v)$ with $\pi(e) \neq 0$ we show that π can be extended to an automorphism in $\text{Aut}(\mathfrak{A}, a_1, \dots, a_{i-1})$. Thus let us fix $e \in E(v)$ such that $\pi(e) \neq 0$. Since we have that $\sum_{f \in E(v)} \pi(f) = 0$ there has to be another edge $e' \in E(v)$ with $\pi(e') \neq 0$. We apply Lemma 14 to obtain an automorphism $\sigma \in \text{Aut}(\mathfrak{A}, a_1, \dots, a_{i-1})$ such that $\sigma(e) = \pi(e)$, $\sigma(e') = -\pi(e)$ and $\sigma(f) = 0$ for all $f \in E(v) \setminus \{e, e'\}$. Now consider $(\pi - \sigma) \in \mathbb{F}_q^{E(v)}$. By the induction hypothesis we can extend this vector to an automorphism $\pi_* \in \text{Aut}(\mathfrak{A}, a_1, \dots, a_{i-1})$. But then $(\pi_* + \sigma) \in \text{Aut}(\mathfrak{A}, a_1, \dots, a_{i-1})$ is an extension of π .

For the second case assume that $a, b \in \hat{e}$ for some edge $e \in E(v)$. As above we conclude that the edge e is not marked. Since $\text{con}(\mathcal{G}) > (\ell + 2)$ the minimal degree of each vertex in \mathcal{G} is at least $(\ell + 4)$. Since the vertex v is not marked there has to be another edge $e' \in E(v)$, $e \neq e'$ which is not marked. Thus we can apply Lemma 14 to obtain an automorphism $\pi \in \text{Aut}(\mathfrak{A}, a_1, \dots, a_{i-1})$ such that $\pi(a) = b$ and $\pi(f) = 0$ for all $f \in E(v) \setminus \{e, e'\}$. ◀

It is well-known that classes of $C_{\infty\omega}^{\ell+2}$ -equivalent tuples can be ordered in FPC, see e.g. [17]. Hence, it follows from our previous lemma that the class \mathcal{K} satisfies property (II).

► **Lemma 16.** *The class \mathcal{K} satisfies the properties (I) and (II).*

Let us now turn our attention to property (IV). In the next lemma we are going to show that for each $k \geq 1$ and each sufficiently connected graph $\mathcal{G} \in \mathfrak{G}$, the logic $C_{\infty\omega}^k$ cannot distinguish between any pair of CFI-structures over \mathcal{G} (although there exist non-isomorphic CFI-structures over \mathcal{G}).

► **Lemma 17.** *Let $k \geq 1$ and let $\mathcal{G} = (V, \leq, E) \in \mathfrak{G}$ such that $\text{con}(\mathcal{G}) > k$. Then for all $\vec{d}, \vec{d}_* \in [q]^V$ it holds that*

$$\text{CFI}_q(\mathcal{G}, \vec{d}) \stackrel{C}{\equiv}_k \text{CFI}_q(\mathcal{G}, \vec{d}_*).$$

Thus, the class \mathcal{K} satisfies property (IV).

Proof. Let $\mathfrak{A} = \text{CFI}_q(\mathcal{G}, \vec{d})$ and let $\mathfrak{B} = \text{CFI}_q(\mathcal{G}, \vec{d}_*)$. Without loss of generality we assume that $\mathfrak{A} \not\equiv \mathfrak{B}$. We show that Duplicator wins the k -pebble bijection game on \mathfrak{A} and \mathfrak{B} . Let $z_a := \sum_{v \in V} \vec{d}(v)$, let $z_b := \sum_{v \in V} \vec{d}_*(v)$ and let $z := z_b - z_a$. As above, for $e = (v, w) \in E$ and $y \in [q]$ we let $\sigma^y[e] \in \Gamma = \Gamma(\mathcal{G})$ denote the isomorphism which shifts the edge class \hat{e} by y , the edge class \hat{f} for $f = (w, v)$ by $-y$ and which stabilises all remaining classes, i.e.

$$\sigma^y[e](f) = \begin{cases} z, & \text{if } f = (v, w), \\ -z, & \text{if } f = (w, v), \\ 0, & \text{else.} \end{cases}$$

Given a position $(\mathfrak{A}, a_1, \dots, a_\ell, \mathfrak{B}, b_1, \dots, b_\ell)$ in the k -pebble bijection game, we say that a pair (v, π) with $v \in V$ and $\pi \in \Gamma(\mathcal{G})$ is *good* if:

- the v -gadget is not marked (by the pebbled elements a_1, \dots, a_ℓ in \mathfrak{A} or, equivalently, by the pebbled elements b_1, \dots, b_ℓ in \mathfrak{B}),
- $\pi(a_i) = b_i$ for $1 \leq i \leq \ell$,
- $\pi(\mathfrak{A} \setminus \hat{v}) = \mathfrak{B} \setminus \hat{v}$, and
- $(\sigma^z[e] + \pi)(\mathfrak{A} \upharpoonright \text{gadget}(v)) = \mathfrak{B} \upharpoonright \text{gadget}(v)$ for all $e \in E(v)$.

Intuitively this means that π is nearly an isomorphism between \mathfrak{A} and \mathfrak{B} except for the gadget associated to vertex v . Of course π itself does not induce a bijection between the universes of the two CFI-structures (as otherwise $\mathfrak{A} \cong \mathfrak{B}$). However, for each $e \in E(v)$ we can associate a bijection $\hat{\pi}_e : A \rightarrow B$ to π which is defined as

$$\hat{\pi}_e(x) = \begin{cases} \pi(x), & \text{if } x \notin \hat{v}, \\ (\sigma^z[e] + \pi)(x), & \text{if } x \in \hat{v}. \end{cases}$$

In what follows we show that Duplicator can play in such a way that after each round such a good pair (v, π) exists. Obviously, if Duplicator can maintain this invariant this suffices for her to win the game.

Indeed we can find such a good pair (v, π) by Lemma 11 for the initial position $(\mathfrak{A}, \mathfrak{B})$ of the game. Let us now consider one round of the game which starts from a position $(\mathfrak{A}, a_1, \dots, a_\ell, \mathfrak{B}, b_1, \dots, b_\ell)$ for which a good pair (v, π) exists. First, Spoiler chooses a pair $i \leq k$ of pebbles which he removes from the game board (if the corresponding pebbles are placed at all). Duplicator then answers Spoiler's challenge by providing a bijection $\hat{\pi}_e$ for some edge $e \in E(v)$ which is not marked. Note that such an edge e exists since $\text{con}(\mathcal{G}) > k$ and thus each vertex has degree at least $k + 2$. Spoiler picks a new pair $(a, \hat{\pi}_e(a)) \in A \times B$ of $\hat{\pi}_e$ -related elements on which he places the i -th pair of pebbles. By the properties of π it immediately follows that the resulting mapping $\bar{a}[i \mapsto a] \mapsto \bar{b}[i \mapsto b]$ is a partial isomorphism. However, it might happen that Spoiler placed the i -th pair of pebbles on equation nodes \hat{v} in the gadget associated to vertex v . In this case the pair (v, π) is not good any longer. So assume that Spoiler pebbled a new pair of elements $(a, \pi_e(a)) \in \hat{v} \times \hat{v}$. Since the edge $e = (v, w)$ was not marked we know that w is not marked. Thus it is easy to see that the pair $(w, \sigma^z[e] + \pi)$ is good. \blacktriangleleft

To complete our proof we establish an FPS_q -definable canonisation procedure on the class \mathcal{K} . The idea is as follows: given a CFI-structure $\mathfrak{A} = \text{CFI}_q(\mathcal{G}, \vec{d})$ over a graph \mathcal{G} 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 the structure \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 classes of CFI-structures over \mathcal{G} which concludes our argument.

More specifically, let $\mathcal{G} = (V, \leq, E) \in \mathfrak{G}$, let $\mathfrak{A} = \text{CFI}_q(\mathcal{G}, \vec{d}) \in \mathcal{K}$ 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 of the linear system are given as follows:

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

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

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

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

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

Proof. If $\sum \vec{d} = z$ then it is easy to verify that we obtain a solution $\vec{\sigma} \in \mathbb{F}_q^V$ of the linear system by setting $\vec{\sigma}(e_i) = i$ and $\vec{\sigma}(v) = \vec{d}(v)$. For the other direction, we show that a solution $\vec{\sigma} \in \mathbb{F}_q^V$ of this system defines an isomorphism π between \mathfrak{A} and $\mathfrak{B} = \text{CFI}_q(\mathcal{G}, \vec{d}_+)$ where $\vec{d}_+(v) := \vec{\sigma}(v)$. As a preparation, we let $\delta(e) := \vec{\sigma}(e_i) - i$ for $e \in E$ and some $e_i \in \hat{e}$. Since $\vec{\sigma}$ is a solution, $\delta \in \mathbb{F}_q^E$ is well-defined. Now we obtain the isomorphism π for $e_i \in \hat{E}$ and $\rho \in \hat{V}$ by setting

$$\begin{aligned} \pi(e_i) &\mapsto e_{\vec{\sigma}(e_i)} \\ \pi(\rho) &\mapsto \rho + \delta. \end{aligned}$$

Using the equations (E 1) and (E 2) one easily verifies that π respects the cycle relation C and the inverse relation I . Moreover, let $(\rho, e_{\rho(e)}) \in R$. Then

$$\pi(e_{\rho(e)}) = e_{\vec{\sigma}(e_{\rho(e)})} \text{ and } \vec{\sigma}(e_{\rho(e)}) = \rho(e) + \delta(e).$$

Thus, π also respects R . Finally, by the equations of type (E 3), for all $v \in V$ and $\rho \in \hat{v}$ we have that

$$\sum \rho + \delta = \sum_{e \in E(v)} \vec{\sigma}(e_{\rho(e)}) = \vec{\sigma}(v).$$

This shows that $\vec{\sigma}(v) = d_+(v)$ and that $\sum \vec{d}_+ = \sum_{v \in V} \vec{\sigma}(v) = z$ because of equation (E 4). ◀

► **Lemma 19.** *The class \mathcal{K} satisfies the property (III).*

This finishes our proof of Theorem 3.

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 leads to the interesting question whether, in general, rank operators can be simulated by solvability quantifiers within fixed-point logic with counting. In this context, it

is worthwhile to remark that many other problems from linear algebra are known to sit in between of “solving linear equation systems” and “computing the matrix rank”, for example, deciding whether two matrices are similar or equivalent, see [18, 15, 16].

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

In order to state our main result formally, we first 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 the logic FOS_p is a *one-sorted* logic which does not have access to a counting sort.

► **Definition 20.** For every prime p , the logic FOS_p results by extending the syntax of FO by the following formula creation rule:

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

The semantics of $\psi(\bar{z})$ are defined as above. For completeness, let $k = |\bar{x}|$ and $\ell = |\bar{y}|$. A pair $(\mathfrak{A}, \bar{z} \mapsto \bar{c})$ with $\bar{c} \in A^{|\bar{z}|}$ defines an $I \times J$ -matrix M_φ over $\{0, 1\} \subseteq \mathbb{F}_p$ where $I = A^k$ and $J = A^\ell$ and where $M_\varphi(\bar{a}, \bar{b}) = 1$ 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 . Now we let $\mathfrak{A} \models \psi(\bar{c})$ if, and only if, $M_\varphi \cdot \vec{x} = \mathbb{1}$ is solvable.

Analogously to the definition of FPS in Section 2, the syntactic normal form of definable linear equation systems in the definition of slv_p -quantifier does not lead to a severe restriction (again, see Lemma 4.1 in [6]).

Let us briefly summarise what is known about the logic FOS_p (see also [6, 18]). First of all, it follows from [7] that for every prime p , the logic FOS_p can express the symmetric transitive closure of definable relations. Hence, FOS_p subsumes the logic STC and can express every LOGSPACE-computable property of ordered structures. Secondly, it also follows from [7] that FOS_2 can distinguish between the odd and even version of a CFI-graph, which means that FOS_2 cannot be a fragment of FPC. More generally, by adapting the CFI-construction for other fields one can show that $\text{FOS}_p \not\leq \text{FPC}$ for all $p \in \mathbb{P}$ (see e.g. [15]).

On the domain of 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* MOD_kL for integers $k \geq 2$. Analogously to the case of modulo counting classes for polynomial time, the idea is to say that a problem is in MOD_kL if there exists a non-deterministic logspace Turing machine which verifies its inputs by producing a number of accepting paths which is not congruent 0 mod k . For the formal definition we refer the reader to [3]. It turns out that, at least for primes p , the class MOD_pL is closed under many natural operations, including all Boolean operations and even logspace Turing reductions [3, 14]. Furthermore, many problems from linear algebra over \mathbb{F}_p are complete for MOD_pL . 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 that for all primes p , the logic FOR_p captures MOD_pL on the class of ordered structures. It has been noted in [18] that their proof shows that the same correspondence holds for the logic FOS_p .

► **Proposition 21** ([7],[18]). *On ordered structures we have $\text{FOS}_p = \text{FOR}_p = \text{MOD}_p\text{L}$.*

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

► **Theorem 22.** *For all primes p we have $\text{FOS}_p < \text{FOR}_p$ (over the class of sets $\mathcal{S}(\emptyset)$).*

In some sense, this result is not very surprising. Over the class of sets, the logic FOR_p captures the complexity class MOD_pL since the size of a set is a complete invariant. In contrast, the logic FOS_p cannot access the counting sort and thus had to express properties over pure unordered sets which have the maximal amount of symmetries. However, it is not obvious how one can turn this intuition into a formal argument. Strikingly, FOS_p has non-trivial expressive power over sets. For instance, FOS_p can determine the size of sets modulo p [18], and consequently, modulo p^k for every fixed k (since $n \equiv 0 \pmod{p^k}$ if, and only if, $n \equiv 0 \pmod{p}$ and $\binom{n}{p} \equiv 0 \pmod{p^{k-1}}$). Note that fixed-point logic FP, for example, collapses to first-order logic over sets.

In order to prove Theorem 22 we make use of the following strong normal form for FOS_p which has been established in Corollary 4.8 of [6].

► **Theorem 23.** *Every formula $\vartheta(\bar{z}) \in \text{FOS}_p$ is equivalent to an FOS_p -formula of the form $(\text{slv}_p \bar{x}_1, \bar{x}_2)\alpha(\bar{x}_1, \bar{x}_2, \bar{z})$ 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, we are aiming at considerably reducing the size of an input linear equation system via an FOR_p -definable transformation. For the remainder of this proof, let us fix a quantifier-free formula $\alpha(x_1, \dots, x_k, y_1, \dots, y_\ell) \in \text{FO}(\emptyset)$ 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 $\mathfrak{A} \models (\text{slv}_p \bar{x}_1, \bar{x}_2)\alpha(\bar{x}_1, \bar{x}_2)$ 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 $\Gamma = \Gamma_n = \text{Sym}([n])$. Then the group Γ acts on I_n and J_n in the natural way. As in Section 3 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. 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 which is distinct from p and a subgroup $\Delta \leq \Gamma$ which is a q -group, i.e. $|\Delta| = q^m$ 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

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

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

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

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 much more compact linear equation system $M_n^* \cdot \vec{x} = \mathbb{1}$ which is equivalent to the given one.

Recall that the maximal q -subgroups $\Delta \leq \Gamma$ are the q -Sylow groups of Γ . It is well-known that for the case where $\Gamma = \text{Sym}([n])$ these groups can be obtained via an inductive construction which we want to explain here for the special case of n being a power of q (the general case can be handled similarly, see e.g. [13]). Hence from now on, let us assume that $n = q^r$ for some $r \geq 1$.

First of all, we determine the size of q -Sylow groups of Γ . A simple induction shows that the maximal $t \geq 1$ such that q^t divides $n! = (q^r)!$ is given as

$$t = q^{r-1} + q^{r-2} + \dots + q + 1 = \frac{q^r - 1}{q - 1}.$$

In fact, we can write $(q^r)!$ as $(q^r)! = 1 \cdot \dots \cdot (1 \cdot q) \cdot \dots \cdot (2 \cdot q) \cdot \dots \cdot (q^{r-1} \cdot q)$. Hence $t = t_* + q^{r-1}$ where t_* is the maximal such that q^{t_*} divides $(q^{r-1})!$.

In particular, if we denote for $n = q^r$ a q -Sylow of $\text{Sym}([n])$ by Δ_r , then our argument from above shows that $|\Delta_1| = q$ and that

$$|\Delta_{r+1}| = |\Delta_r|^q \cdot q.$$

As it turns out, this equation already gives a hint about the algebraic structure of Δ_r . Indeed, Δ_{r+1} can be obtained as the *wreath product* of Δ_r and the cyclic group \mathbb{F}_q . Since $\Delta_1 = \mathbb{F}_q$ it follows that Δ_r is the r -fold wreath product of the cyclic group \mathbb{F}_q . We decided to skip the formal definition of the notion of wreath products and rather to directly illustrate this concept for the particular case of the q -Sylow groups of $\Gamma = \text{Sym}([n]) = \text{Sym}([q^r])$.

To obtain an algebraic description of these groups, we inductively construct for $r \geq 1$ a q -Sylow subgroup $\Delta_r \leq \text{Sym}([q^r])$ together with a family of trees \mathcal{T}_i^x for $i = 0, \dots, r$ and $x \in [q^{r-i}]$ such that the following properties hold.

- (I) \mathcal{T}_i^x is a complete q -ary tree of height i whose leaves are labelled with elements from $[n]$. More precisely, the labels of the leaves of \mathcal{T}_i^x form the set $\mathcal{P}_i^x = \{x \cdot q^i, \dots, (x+1) \cdot q^i - 1\}$ (note that \mathcal{P}_i^x is the x -th block of the natural partition of $[n]$ into parts of size q^i).

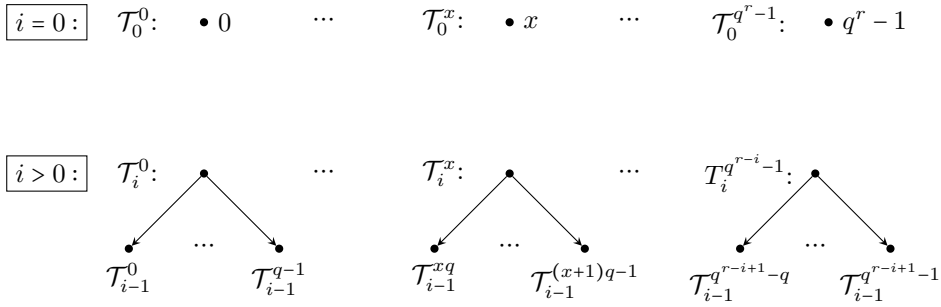
- (II) For all $i \leq r$ the group Δ_r transitively acts on the set $\{\mathcal{T}_i^x : x \in [q^{r-i}]\}$ by applying permutations $\delta \in \Delta_r$ to the labels of the leaves of the tree \mathcal{T}_i^x . Moreover, for each $i \leq r$, the subgroup of Δ_r which point-wise stabilises the trees \mathcal{T}_i^x is a normal subgroup of Δ_r .
- (III) We have $\Delta_1 \leq \Delta_2 \leq \dots \leq \Delta_r$ where Δ_i acts on the set of labels \mathcal{P}_i^0 of the tree \mathcal{T}_i^0 . More generally, for every block \mathcal{P}_i^x , the group Δ_r contains a subgroup $\Delta_r^{i,x} \leq \Delta_r$ which point-wise fixes the elements of all blocks \mathcal{P}_i^y for $y \neq x$ and whose action on \mathcal{P}_i^x corresponds to the action of Δ_i on \mathcal{P}_i^0 .

The inductive construction of the trees \mathcal{T}_i^x is depicted in Figure 2. To understand this construction better, it is quite useful to think of elements $y \in [n]$ as being represented in their q -adic encoding, i.e. $y = y_0 + y_1 \cdot q + \dots + y_{r-1} \cdot q^{r-1}$. Then we have that $y \in \mathcal{P}_r^0 = [n]$ and

- $y \in \mathcal{P}_{r-1}^{y_{r-1}}$
- $y \in \mathcal{P}_{r-2}^{y_{r-2} + y_{r-1} \cdot q}$
- ...
- $y \in \mathcal{P}_0^{y_0 + \dots + y_{r-1} \cdot q^{r-1}} = \mathcal{P}_0^y$.

Hence, the q -adic encoding of y describes the unique path in the tree \mathcal{T}_r^0 from the root to the leaf \mathcal{T}_0^y . The trees \mathcal{T}_i^x clearly satisfy the properties stated in (I).

For the inductive construction of the q -Sylow groups Δ_r we first fix Δ_1 as the cyclic group generated by the natural cyclic shift $\gamma = (0\ 1 \dots q-1)$ on the set $\mathcal{P}_1^0 = \{0, \dots, q-1\}$.



■ **Figure 2** Inductive definition of the trees \mathcal{T}_i^x

We proceed with the inductive step $r \mapsto r+1$. The set $[q]^{r+1}$ splits into q blocks $\mathcal{P}_r^0, \dots, \mathcal{P}_r^{q-1}$ each of size q^r . The group Δ_r acts on \mathcal{P}_r^0 and point-wise fixes the elements from the blocks \mathcal{P}_r^x with $x \neq 0$. Let $\gamma \in \text{Sym}([n])$ for $n = q^{r+1}$ be the following permutation which shifts the segments $\mathcal{P}_r^0, \dots, \mathcal{P}_r^{q-1}$ in a cycle of length q by composing the natural shifts on the sets of residues modulo q^r :

$$\gamma = (0 \dots (q-1)q^r)(1 \dots 1 + (q-1)q^r) \dots (q^r - 1 \dots q^r - 1 + (q-1)q^r).$$

Hence for all $a \in [n]$ we have $\gamma(a) = (a + q^r) \bmod q^{r+1}$. We set $\Delta_r^0 = \Delta_r$ and, more generally, $\Delta_r^x = (\gamma^x)\Delta_r(\gamma^x)^{-1}$ for $x = 0, \dots, q-1$ to obtain q copys of Δ_r which independently act on the segments \mathcal{P}_r^x for $0 \leq x \leq q-1$. Finally, we define Δ_{r+1} as the *semi-direct product* of $(\Delta_r^0 \times \dots \times \Delta_r^{q-1})$ and the cyclic group $\langle \gamma \rangle$ of size q . This means that the group elements of Δ_{r+1} are elements in the set $(\Delta_r^0 \times \dots \times \Delta_r^{q-1} \times \langle \gamma \rangle)$ and that the group operation is given by

$$(\delta_1, \dots, \delta_{q-1}, \alpha) \cdot (\epsilon_1, \dots, \epsilon_{q-1}, \beta) = (\delta_1 \cdot \alpha \epsilon_1 \alpha^{-1}, \dots, \delta_{q-1} \cdot \alpha \epsilon_{q-1} \alpha^{-1}, \alpha \cdot \beta).$$

Since $|\Delta_{r+1}| = |\Delta_r|^q \cdot q$ we conclude that Δ_{r+1} indeed is a q -Sylow subgroup.

From our construction it immediately follows that Δ_{r+1} satisfies the properties stated in (III). To see that Δ_{r+1} also satisfies the properties stated in (II) we start by showing that, for $i \leq r$, Δ_{r+1} transitively acts on $\{\mathcal{T}_i^x : x \in [q^{r+1-i}]\}$. If we split the set $[q^{r+1-i}]$ into q blocks $\mathcal{P}_{r-i}^0, \dots, \mathcal{P}_{r-i}^{q-1}$ of size q^{r-i} then we know from the induction hypothesis that Δ_r^0 transitively acts on the set of trees $\{\mathcal{T}_i^x : x \in \mathcal{P}_{r-i}^0\} = \{\mathcal{T}_i^x : x \in [q^{r-i}]\}$. Moreover, it is easy to verify that for all $x \in [q^{r+1-i}]$ we have $\gamma(\mathcal{T}_i^x) = \mathcal{T}_i^z$ where $z = x + q^{r-i} \pmod{q^{r+1-i}}$. Hence $(\gamma^y)\{\mathcal{T}_i^x : x \in \mathcal{P}_{r-i}^0\} = \{\mathcal{T}_i^x : x \in \mathcal{P}_{r-i}^y\}$ for all $0 \leq y \leq q-1$ which means that Δ_r^y transitively acts on $\{\mathcal{T}_i^x : x \in \mathcal{P}_{r-i}^y\}$ and thus (II) holds.

The crucial step is to understand the action of Δ_r on the sets $I_n = [n]^k$ and $J_n = [n]^\ell$ (for the case where $n = q^r$). In fact, our next aim is to develop a complete invariant for the Δ_r -orbits on these index sets. Recall that the sets of Δ_r -orbits on I_n and J_n provide index sets for the succinct linear equation system $M_n^* \cdot \vec{x} = \mathbb{1}$. To define this invariant, the main idea is to describe the position of a tuple $\bar{a} \in I_n$ (or $\bar{a} \in J_n$, respectively) in the tree $\mathcal{T} := \mathcal{T}_r^0$.

Let us first define the *signature* $\text{sgn}(a, b)$ of a pair $(a, b) \in [n] \times [n]$ as the tuple $(i, z) \in [r+1] \times [q]$ such that the lowest common ancestor of a, b in \mathcal{T} is the root of a tree \mathcal{T}_i^x and such that a is located in a subtree $\mathcal{T}_{i-1}^{xq+y_a}$ for $y_a \in [q]$ and b is located in the subtree $\mathcal{T}_{i-1}^{xq+y_b}$ where $y_b = y_a + z \pmod{q}$. For the special case where $i = 0$ we have $a = b$ and agree to set $z = 0$. With this preparation we define the signature $\text{sgn}(\bar{a})$ of a tuple $\bar{a} = (a_1, \dots, a_\ell) \in J_n$ as the list $\sigma \in ([r+1] \times [q])^{\ell(\ell-1)/2}$ consisting of the individual signatures $\text{sgn}(a_i, a_j)$ for all pairs a_i, a_j with $1 \leq i < j \leq \ell$. The signature of tuples in I_n is defined analogously.

► **Lemma 24.** *Let $\bar{a} \in J_n$. Then $\text{sgn}(\bar{a}) = \text{sgn}(\pi\bar{a})$ for all $\pi \in \Delta_r$.*

Proof. Immediately follows from the construction of Δ_r and the trees \mathcal{T}_i^x . ◀

► **Lemma 25.** *Let $\bar{a}, \bar{b} \in J_n$. If $\text{sgn}(\bar{a}) = \text{sgn}(\bar{b})$, then $\bar{b} \in \Delta_r(\bar{a})$.*

Proof. We proceed by induction on the maximal position $0 \leq i \leq \ell$ such that $a_j = b_j$ for all $j = 1, \dots, i$. The case $i = \ell$ is clear, so assume that $i < \ell$. Let $\bar{a} = (a_1, \dots, a_i, a_{i+1}, \dots, a_\ell)$ and $\bar{b} = (a_1, \dots, a_i, b_{i+1}, \dots, b_\ell)$. We show that there exists a permutation $\delta \in \Delta_r$ which pointwise fixes a_1, \dots, a_i and such that $\delta(a_{i+1}) = b_{i+1}$. Then the claim follows from Lemma 24 together with the induction hypothesis. For $i = 0$ this is easy, because Δ_r acts transitively on $[n]$. If $i > 0$ we choose $a_w \in \{a_1, \dots, a_i\}$ such that $\text{sgn}(a_w, a_{i+1}) = (c, d)$ and such that c is minimal with this property. Obviously we have $c > 0$. By the choice of a_w the lowest common ancestor of a_w and a_{i+1} is the root of a tree \mathcal{T}_c^x . Moreover, a_w is located in a subtree \mathcal{T}_{c-1}^{xq+y} for some $0 \leq y \leq q-1$ and a_{i+1} is located in the subtree \mathcal{T}_{c-1}^{xq+z} where $z = y + d \pmod{q}$. Since $\text{sgn}(\bar{a}) = \text{sgn}(\bar{b})$ also b_{i+1} occurs as the label of a leaf in the subtree \mathcal{T}_{c-1}^{xq+z} . By the minimality assumption on c we know that none of the elements $\{a_1, \dots, a_i\}$ occurs in the tree \mathcal{T}_{c-1}^{xq+z} . Hence, by the properties of the group Δ_r stated in (III) we can find an element $\delta \in \Delta_r$ which point-wise fixes all elements outside the block \mathcal{P}_{c-1}^{xq+z} (in particular, the elements a_1, \dots, a_i) and which moves a_{i+1} to b_{i+1} . ◀

Following our definition from above, the signature $\text{sgn}(\bar{a})$ of an element $\bar{a} \in J_n$ is a tuple of length $\ell(\ell-1)/2$ whose entries are pairs $(i, z) \in [r+1] \times [q]$. We denote the set of all possible sequences of this form by $S_n^\ell = ([r+1] \times [q])^{\ell(\ell-1)/2}$. Of course, not every tuple in $\sigma \in S_n^\ell$ can be realised as the signature $\text{sgn}(\bar{a}) = \sigma$ of an element $\bar{a} \in J_n$. Similarly, we define the set $S_n^k = ([r+1] \times [q])^{k(k-1)/2}$ to capture all possible signatures of elements in I_n .

Since the coefficient matrix M_n^* of the equivalent linear equation system $M_n^* \cdot \vec{x} = \mathbb{1}$ can be defined as a matrix whose index sets are the collections of Δ_r -orbits on I_n and J_n , we

can use the notion of signatures to describe M_n^* as an $(S_n^k \times S_n^\ell)$ -matrix. This fits with our proof plan as the index sets S_n^k and S_n^ℓ of the matrix M_n^* are much smaller than the index sets I_n and J_n of the coefficient matrix M_n of the original linear equation system. However, it still might be the case that the succinctness of the matrix M_n^* does not help, because it is not possible to obtain its entries within FOR_p .

We show that this is not the case. More precisely we show that we can define the matrix M_n^* in FOC in a structure of size r (where we assume that $r \geq q$). Therefore, the main technical step is to show that FOC can count (modulo p) the number of realisations of a potential signature $\sigma \in S_n^k$.

First of all, we need some further notation. A *complete equality type in $k + \ell$ variables* is a consistent set $\tau(x_1, \dots, x_k, x_{k+1}, \dots, x_{k+\ell})$ of literals $x_i = x_j, x_i \neq x_j$ which contains for every pair $i < j$ either the atom $x_i = x_j$ or the literal $x_i \neq x_j$. Note that each quantifier-free formula $\alpha \in \text{FO}(\emptyset)$ can be expressed as a Boolean combination of complete equality types.

In the following main technical lemma we show that in the structure $\mathfrak{A} = ([r])$ we can count (modulo p) the number of realisations of a (potential) signature $\sigma \in S_n^\ell$ in a subtree \mathcal{T}_i^x in FOC. More generally, this is possible if we additionally fix some entries of the tuples which should realise σ in \mathcal{T}_i^x . Here we need another prerequisite: as we want to work with elements from the set $[n] = [q^r]$ in a structure of size r we have to agree on some sort of succinct representation. Of course the natural choice here is to represent numbers $x \in [n]$ in the structure \mathfrak{A} via their q -adic encoding: a binary relation $R \subseteq [r]^2$ which corresponds to a function $R : [r] \rightarrow [q]$ represents the number $x(R) \in [n] = \sum_{i=0}^{r-1} R(i) \cdot q^i$. Note that this encoding requires a linear order on the set $[r]$ (which is *not* the case for the structure \mathfrak{A}). However, as we are working with FOC we can just use the number sort on which a linear order is available. Hence in the following, whenever we specify FOC-formulas or FOC-terms with free variables or with free relation symbols which should represent numbers, then we implicitly assume that these variables are *numeric* variables and that the relation symbols are evaluated over the number sort. The same holds for signatures $\sigma \in S_n^\ell$ which we specify in FOC-formulas by a list of pairs (h_i, d_i) of *numeric* variables of length $\binom{\ell}{2}$.

Before we state our main technical lemma it is helpful to recall that our inductive construction of the trees \mathcal{T}_i^x fits very well with the q -adic encoding of numbers $x \in [n]$. Again, let $x \in [n]$ be given by its q -adic encoding as $x = (x_0, \dots, x_{r-1}) \in [q]^r$, i.e. $x = \sum_{i=0}^{r-1} x_i \cdot q^i$. Then the i -th node on the unique path from the root in the tree $\mathcal{T} = \mathcal{T}_r^0$ to the leaf \mathcal{T}_0^x is the root of the tree \mathcal{T}_{r-i}^y where $y = x_{r-i} + x_{r-i+1}q + \dots + x_{r-1}q^{i-1}$. In other words, the q -adic encoding of x precisely describes the path in the tree \mathcal{T} from the root to the leaf labelled with x where at level $(r - i)$ the i last entries x_{r-i}, \dots, x_{r-1} in the q -adic encoding of x are determined (i.e. x is a member of the block \mathcal{P}_{r-i}^y).

► **Lemma 26.** *For all $\ell \geq 1$ and $0 \leq s \leq \ell$ there exist*

(a) a term $\Theta(i, h_1, d_1, \dots, h_t, d_t) \in \text{FOC}(\{R_x, R_1, \dots, R_s\})$, and
 (b) formulas $\varphi_e(y, z, i, h_1, d_1, \dots, h_t, d_t) \in \text{FOC}(\{R_x, R_1, \dots, R_s\})$ for $e = s + 1, \dots, \ell$,
 where $t = \binom{\ell}{2}$, such that for all $r \geq q$, all $i \leq r$, all $\sigma = ((h_1, d_1), \dots, (h_t, d_t)) \in S_n^\ell$ where $n = q^r$, all $x \in [q^{r-i}]$ and all $a_1, \dots, a_s \in \mathcal{P}_i^x$ the following holds: let $\mathfrak{A} = ([r])$ and let R_x, R_1, \dots, R_s be numerical relations such that R_x represents the (q -adic encoding of the) element $x \in [q^{r-i}]$ and such that each R_i represents the (q -adic encoding of) the element a_i . Then we have that

(i) the value $\Theta^{\mathfrak{A}}(q, i, h_1, d_1, \dots, h_t, d_t)$ of the term Θ in \mathfrak{A} is $|Z| \bmod p$ where

$$Z = \{(a_{s+1}, \dots, a_\ell) \in (\mathcal{P}_i^x)^{\ell-s} : \text{sgn}(a_1, \dots, a_s, a_{s+1}, \dots, a_\ell) = \sigma\}.$$

(ii) if $Z \neq \emptyset$, then the formulas $(\varphi_e)_{s < e \leq \ell}$ define the q -adic representation of witnessing elements $a_{s+1}, \dots, a_\ell \in \mathcal{P}_i^x$, i.e. such that $(a_{s+1}, \dots, a_\ell) \in Z$.

Proof. First of all, by our previous observations it is easy to see that the condition $a_j \in \mathcal{P}_i^x$ for $j = 1, \dots, s$ can be defined in FOC. More generally, we can use the q -adic encoding of the elements a_j to determine $\text{sgn}(a_1, \dots, a_s)$ in FOC. Hence, for the remainder of the proof we assume that $\text{sgn}(a_1, \dots, a_s)$ is consistent with σ and that $a_j \in \mathcal{P}_i^x$ for $j = 1, \dots, s$.

We proceed by induction on ℓ . For $\ell = 1$ it suffices to show that FOC can compute $(n \bmod p)$ where $n = q^r$ in the structure \mathfrak{A} . To see this, recall that p and q are co-prime and thus we can use Lagrange's theorem to conclude that $q^r \equiv q^{r'} \pmod p$ if $r' \equiv r \pmod{p-1}$. Since p is a constant, the claim follows.

Let $\ell \geq 2$. We distinguish between the following two cases. If $s = 0$, then we can partition the set of realisations \bar{a} of σ according to first entry a_1 into $|\mathcal{P}_i^x|$ parts of equal size. It suffices to determine the size of each of these blocks, since we can determine $|\mathcal{P}_i^x| \bmod p$ in FOC similarly as above.

Without loss of generality let us assume that $a_1 = x \cdot q^i$. Since we have given the q -adic encoding of x it is easy to see that we can define the q -adic encoding of xq^i in FOC. This gives us the formula φ_1 . Next, we partition the set of indices $\{2, \dots, \ell\}$ into classes according to the equivalence relation $j_1 \approx j_2$ if $\sigma(1, j_1) = \sigma(1, j_2)$. Let the resulting classes be Y_1, \dots, Y_v and let $\sigma(1, y) = (h_w, d_w)$ for all $y \in Y_w$ and $w = 1, \dots, v$.

We observe that there exists a tuple \bar{a} with $a_1 = x \cdot q^i$ which realises σ in the tree \mathcal{T}_i^x (that is $Z \neq \emptyset$) if, and only if, the following conditions are satisfied:

- for all $w = 1, \dots, v$ we have $h_w \leq i$, and
- for every $Y_w = \{y_1^w, \dots, y_{\ell_w}^w\}$ there is a tuple \bar{a}^w of length ℓ_w which realises σ (restricted to the indices from Y_w) in the subtree $\mathcal{T}_{h_w-1}^{xq^{i-h_w+1}+d_w}$, and
- for all pairs $y_1 \in Y_{w_1}$ and $y_2 \in Y_{w_2}$ with $w_1 \neq w_2$ we have that

$$\sigma(y_1, y_2) = \begin{cases} (h_{w_1}, d_{w_2} - d_{w_1} \bmod q), & \text{if } h_{w_1} = h_{w_2} \\ (h_{w_2}, d_{w_2}) & \text{if } h_{w_1} < h_{w_2} \\ (h_{w_1}, d_{w_1}) & \text{if } h_{w_2} < h_{w_1}. \end{cases}$$

Since ℓ is a constant, the number of possible partitions of $\{2, \dots, \ell\}$ is bounded by a constant as well. It is easy to see that for every possible such partition we can check the first and third condition in FOC. To verify the second condition in FOC we use the induction hypothesis. There are two aspects which have to be discussed with more precision. First of all, we have to handle one particular case separately: indeed, if $h_w = 1$ for all $w = 1, \dots, v$, then we cannot use the induction hypothesis since all elements (including a_1) have to be chosen in the same subtree of height one. However, in this case there is only one realisation (if the third condition is satisfied) so this does not cause any problems. The other difficulty is that we have to define the q -adic encoding of the value $z_w = xq^{i-h_w+1} + d_w$ in FOC. We already noted before that the q -adic representation of xq^{i-h_w+1} can be defined in FOC and since $0 \leq d_w < q$ we can also define the q -adic encoding of z in FOC.

In fact, the induction hypothesis also provides us with a term which counts modulo p the number of possible realisations of σ in the subtrees $\mathcal{T}_{h_w-1}^{z_w}$ restricted to the indices in Y_w together with formulas φ_e which define witnessing elements. Finally, since the overall number of possible realisations of σ in \mathcal{T}_i^x is the product of the realisations restricted to the components Y_w , the claim follows for the case where $s = 0$.

For the general case let $\ell \geq s > 0$ and let $a_1, \dots, a_s \in \mathcal{P}_i^x$ be the components of the tuple \bar{a} that are already fixed. Recall that we can assume without loss of generality that

$\text{sgn}(a_1, \dots, a_s)$ is consistent with σ and that all elements a_1, \dots, a_s are located in the subtree \mathcal{T}_i^x . Since we have fixed the element a_1 , we can proceed as above except for two small changes. First of all, when applying the induction hypothesis we have to respect the remaining fixed elements a_2, \dots, a_s . Moreover, when we form the partitions of $\{2, \dots, \ell\}$ into parts Y_1, \dots, Y_v as above then we have to adapt the position of elements corresponding to the index set Y_w since the element a_1 is not necessarily contained in the tree $\mathcal{T}_{h_w+1}^{xq^{i-h_w+1}}$. However, since we have given the q -adic representation of a_1 we can define in FOC the element $0 \leq d_a < q$ such that a_1 is located in the subtree $\mathcal{T}_{h_w+1}^{xq^{i-h_w+1+d_a}}$. The remaining steps can be performed as above. This finishes our proof. \blacktriangleleft

► **Lemma 27.** *Let $\tau(x_1, \dots, x_k, y_1, \dots, y_\ell) \in \text{FO}(\emptyset)$ be a complete equality type (in $k + \ell$ variables). Then there is an FOC-term $\Theta_\tau(\bar{z}_x, \bar{z}_y)$ such that for all $r \geq q$, all $\sigma_{\bar{a}} \in S_n^k$ and $\sigma_{\bar{b}} \in S_n^\ell$, where $n = q^r$, the value $\Theta^\mathfrak{A}(\sigma_{\bar{a}}, \sigma_{\bar{b}})$ of Θ in $\mathfrak{A} = ([r])$ is*

$$\Theta^\mathfrak{A}(q, \sigma_{\bar{a}}, \sigma_{\bar{b}}) = |\{\bar{b} \in J_n : \text{sgn}(\bar{b}) = \sigma_{\bar{b}}, ([n]) \models \tau(\bar{a}, \bar{b})\}| \text{ mod } p$$

for some (or, equivalently, all) $\bar{a} \in I_n$ with $\text{sgn}(\bar{a}) = \sigma_{\bar{a}}$.

Proof. By Lemma 26 we can first check in FOC that $\sigma_{\bar{a}}$ and $\sigma_{\bar{b}}$ can be realised (otherwise the answer is trivial). Moreover, if τ (restricted to x_1, \dots, x_k) is not consistent with $\sigma_{\bar{a}}$ or if τ (restricted to y_1, \dots, y_ℓ) contradicts $\sigma_{\bar{b}}$, then the answer is trivial as well.

In all other cases, Lemma 26 provides FOC-formulas which define in the structure \mathfrak{A} the q -adic encoding of elements $a_1, \dots, a_k \in [n]$ such that $\text{sgn}(\bar{a}) = \sigma_{\bar{a}}$. Moreover, if τ contains a literal $x_i = y_j$, then we can fix the entry b_j as well. Hence, let us assume without loss of generality that τ contains the literals $x_i \neq y_j$ for all $1 \leq i \leq k$ and $1 \leq j \leq \ell$.

For $Y \subseteq \{1, \dots, \ell\}$ and a partial assignment $\epsilon : \{1, \dots, \ell\} \rightarrow \{a_1, \dots, a_k\}$ with $\text{dom}(\epsilon) \cap Y = \emptyset$ we define the set

$$B_Y^\epsilon = \{\bar{b} \in J_n : \text{sgn}(\bar{b}) = \sigma_{\bar{b}}, \text{ for } i \in \text{dom}(\epsilon) : b_i = \epsilon(i), \text{ for } i \in Y : b_i \neq a_1, \dots, a_k\}.$$

With this notation our overall aim is to determine $(|B_Y^\emptyset| \text{ mod } p)$ for $Y = [\ell]$ in FOC. The first observation is that by Lemma 26 we can determine $(|B_\emptyset^\epsilon| \text{ mod } p)$ for all partial assignments ϵ in FOC. The second observation is that we can construct the values $(|B_Y^\epsilon| \text{ mod } p)$ by induction on $|Y|$ as follows. For $Y \subseteq \{1, \dots, \ell\}$ and a partial assignment ϵ (with $\text{dom}(\epsilon) \cap Y = \emptyset$) we have for all $j \in Y$ that

$$|B_Y^\epsilon| = |B_{Y \setminus \{j\}}^\epsilon| - \sum_{a \in \{a_1, \dots, a_k\}} |B_{Y \setminus \{j\}}^{\epsilon \cup \{j \mapsto a\}}|.$$

In this way we recursively obtain the value $(|B_Y^\emptyset| \text{ mod } p)$ for $Y = [\ell]$. Since ℓ is a constant the recursion depth is bounded by a constant as well and the procedure can be formalised in FOC. \blacktriangleleft

► **Lemma 28.** *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^* where $n = q^r$.*

Proof. Recall that we can view M_n^* as an $(S_n^k \times S_n^\ell)$ -matrix over \mathbb{F}_p . To represent the index sets S_n^k and S_n^ℓ we let $\bar{\mu}$ and $\bar{\nu}$ be tuples of numeric variables of lengths $|\bar{\mu}| = \binom{k}{2}$ and $|\bar{\nu}| = \binom{\ell}{2}$, respectively.

Note that the entry $M_n^*(\sigma_{\bar{a}}, \sigma_{\bar{b}})$ of M_n^* for $\sigma_{\bar{a}} \in S_n^k$ and $\sigma_{\bar{b}} \in S_n^\ell$ is given as

$$M_n^*(\sigma_{\bar{a}}, \sigma_{\bar{b}}) = |\{\bar{b} \in J_n : \text{sgn}(\bar{b}) = \sigma_{\bar{b}}, M_n(\bar{a}, \bar{b}) = 1\}| \text{ mod } p,$$

for some (or, equivalently, all) $\bar{a} \in I_n$ with $\text{sgn}(\bar{a}) = \sigma_{\bar{a}}$. The entry $M_n(\bar{a}, \bar{b})$, in turn, is determined by the quantifier-free formula $\alpha(\bar{x}_1, \bar{x}_2) \in \text{FO}(\emptyset)$. With this preparation, Lemma 27 already shows that we can determine the value $M_n^*(\sigma_{\bar{a}}, \sigma_{\bar{b}})$ for the case where α is a complete equality type. For the general case we just write α as the union of complete equality types and combine the constant number of intermediate results. \blacktriangleleft

► **Definition 29.** Let $\mathcal{K} \subseteq \mathcal{S}(\emptyset)$ be a class of sets. The q -power $\mathcal{K}^q \subseteq \mathcal{S}(\emptyset)$ of \mathcal{K} consists of all sets $\mathfrak{A} = ([q^r])$ such that $\mathfrak{B} = ([r]) \in \mathcal{K}$.

► **Theorem 30.** Let $\mathcal{K} \subseteq \mathcal{S}(\emptyset)$ be a class of sets. If \mathcal{K}^q is definable in FOS_p , then \mathcal{K} is definable in FOR_p .

Proof. If \mathcal{K}^q is definable in FOS_p , then by Theorem 23 we can also find a formula $\varphi = (\text{slv}_p \bar{x}_1, \bar{x}_2)\alpha(\bar{x}_1, \bar{x}_2) \in \text{FOS}_p$ that defines \mathcal{K}^q such that α is quantifier-free.

By using the above construction and Lemma 28, 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 $\varphi^* \in \text{FOR}_p$ 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. \blacktriangleleft

► **Theorem 31.** For all primes p we have $\text{FOS}_p < \text{FOR}_p$ (even over $\mathcal{S}(\emptyset)$).

Proof. Otherwise we had $\text{FOS}_p = \text{FOR}_p$. As above we fix some prime $q \neq p$. Let $\mathcal{K} \subseteq \mathcal{S}(\emptyset)$ be a class of sets such that $\mathcal{K} \notin \text{FOR}_p$, but such that $(\mathcal{K}^q)^q \in \text{FOR}_p$. Such a class \mathcal{K} is well-known to exist (just combine the fact that, over sets, we have $\text{LOGSPACE} \leq \text{FOR}_p \leq \text{PTIME}$ and the space-hierarchy theorem, see e.g. [19]). Since $\text{FOS}_p = \text{FOR}_p$ we had $(\mathcal{K}^q)^q \in \text{FOS}_p$ and by Theorem 30 this means that $\mathcal{K}^q \in \text{FOR}_p$. Again, since $\text{FOR}_p = \text{FOS}_p$, we had $\mathcal{K}^q \in \text{FOS}_p$. A second application of Theorem 30 yields $\mathcal{K} \in \text{FOR}_p$ which contradicts our assumptions. \blacktriangleleft

Finally we remark that the proof also works for the extension of fixed-point logic by solvability quantifiers but in the absence of counting. 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 M. Grohe. The quest for a logic capturing PTIME . In *LICS 2008*, pages 267–271, 2008.
- 12 Y. Gurevich and S. Shelah. On finite rigid structures. *The Journal of Symbolic Logic*, 61(02):549–562, 1996.
- 13 M. Hall. *The theory of groups*. American Mathematical Soc., 1976.
- 14 U. Hertrampf, S. Reith, and H. Vollmer. A note on closure properties of logspace mod classes. *Information Processing Letters*, 75(3):91–93, 2000.
- 15 B. Holm. *Descriptive complexity of linear algebra*. PhD thesis, University of Cambridge, 2010.
- 16 B. Laubner. *The structure of graphs and new logics for the characterization of Polynomial Time*. PhD thesis, Humboldt-Universität Berlin, 2011.
- 17 M. Otto. Bounded variable logics and counting: A study in finite models. 1997.
- 18 W. Pakusa. Finite model theory with operators from linear algebra. Staatsexamensarbeit, RWTH Aachen University, 2010.
- 19 C. Papadimitriou. *Computational complexity*. Addison-Wesley, 1995.
- 20 J. Torán. On the hardness of graph isomorphism. *SIAM Journal on Computing*, 33(5):1093–1108, 2004.