The Journal of Symbolic Logic
Volume 84, Number 1, March 2019

RANK LOGIC IS DEAD, LONG LIVE RANK LOGIC!

ERICH GRÄDEL AND WIED PAKUSA

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, almost 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 article 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.

Received May 13, 2016.
2010 *Mathematics Subject Classification.* 68Q19, 03C13.
Key words and phrases. finite model theory, descriptive complexity, logic, linear algebra.

0022-4812/19/8401-0003
DOI:10.1017/jsl.2018.33

1.1. 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 which is strong enough to express most of the algorithmic techniques leading to polynomial-time procedures [5]. On the other hand, we know that FPC fails to capture PTIME, which follows from a fundamental construction due to Cai, Fürer, and Immerman [4] (today, this construction is known as the CFI-construction). It turns out that two main sources for PTIME-problems that are hard for FPC 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. Second, 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 important problems from the field of linear algebra are not definable in FPC. In particular, this holds for the CFI-query which can be formulated by means of linear equation systems over $\mathbb{F}_2$ [7].

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

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

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

It turns out that rank operators have quite surprising expressive power. For example, they can define the transitive closure of symmetric relations, they can count the number of paths in DAGs modulo p and they can express the solvability of linear equation systems over finite fields (recall that a linear equation system $M \cdot \vec{x} = \vec{b}$ is solvable if, and only if, $\mathsf{rk}(M) = \mathsf{rk}(M \mid \vec{b})$) [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, such as CFI-graphs [4, 7] and multipedes [12, 15]. The common idea of these isomorphism procedures is to reduce the isomorphism problem for a pair of structures to a suitable linear equation system over a finite field. More generally, by a recent result from [1] (which is mainly concerned with another candidate of a logic for polynomial time), it follows that FPR captures polynomial time on certain

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

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]. A particularly intriguing question asks 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 article 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 3.3) 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 3.2).
(b) The extensions of fixed-point logic by rank operators over different prime fields are incomparable (Theorem 3.1), cf. [8, 15, 16].

We obtain these classes of structures $\mathcal{K}_q$ by generalising the CFI-construction. It has been observed by many researchers that the CFI-construction can be viewed as 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 [21], 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 linear

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

equations over *every* (finite) Abelian group G. Essentially, we obtain our classes $\mathcal{K}_q$ by using Holm's construction for the special case where $G = \mathbb{F}_q$.

1.3. 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 FPS $\leq$ FPR is strict. To prove our main result we show that over certain classes of structures the logics FPS and FPR have precisely the same expressive power. On the other hand we show in Section 4 that the extensions of first-order logic and fixed-point logic (without counting) by solvability quantifiers are strictly weaker than the respective extension by rank operators. Hence we separate solvability quantifiers and rank operators in the absence of counting.

A central idea of our proofs is to exploit symmetries of definable linear equation systems. To illustrate this 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-characteristic 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 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 equivalent system.

§2. Logics with linear-algebraic operators. By $\mathrm{FStr}(\tau)$ 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 nonempty sets, then an $I \times J$-matrix M over $\mathbb{F}_p$ is a mapping $M : I \times J \to \mathbb{F}_p$ and an I-vector $\vec{v}$ over $\mathbb{F}_p$ is a mapping $\vec{v} : I \mapsto \mathbb{F}_p$.

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

We 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 $\mathfrak{A} = (A, R_1, \ldots, R_k)$ by the arithmetic $\mathfrak{N} = (\mathbb{N}, +, \cdot, 0, 1)$, i.e., the two-sorted structure $\mathfrak{A}^{\#} = (A, R_1, \ldots, R_k, \mathbb{N}, +, \cdot, 0, 1)$ 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 $\mathbb{N}$.

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

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 $x, y, z, \ldots$ for variables which range over the vertices and Greek letters $\nu, \mu, \ldots$ 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

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

by a tuple of closed numeric terms $\bar{t} = (t_1, \ldots, t_\ell)$ where each t_i bounds the range of the variable ν_i in the tuple $\bar{\nu}$. For the logics which we consider here the value of such numeric terms (and thus the range of all quantifiers over the number sort) is polynomially bounded in the size of the input structure. Together with the standard argument that inflationary fixed-points can be evaluated in polynomial time and the fact that the matrix rank over any field can be determined in polynomial time (for example by the method of Gaussian elimination), this ensures that all the logics which we introduce in the following have polynomial-time data complexity.

Let $\bar{x}\bar{\nu}$ be a mixed tuple of variables and let $\bar{t}$ be a tuple of closed numeric terms which bounds the range of the numeric variables in $\bar{\nu}$. For a formula φ we define a *counting term* $s = [\#\bar{x}\bar{\nu} \leq \bar{t}\,.\,\varphi]$ whose value $s^{\mathfrak{A}} \in \mathbb{N}$ in a structure $\mathfrak{A}$ corresponds to the number of tuples $(\bar{a}, \bar{n}) \in A^k \times \mathbb{N}^\ell$ such that $\mathfrak{A} \models \varphi(\bar{a}, \bar{n})$ and $n_i \leq t_i^{\mathfrak{A}}$ where $k = |\bar{x}|$ and $\ell = |\bar{\nu}|$.

First-order logic with counting FOC is the extension of (the above described two-sorted variant of) FO by counting terms. Similarly, by adding counting terms to FP we obtain (*inflationary*) *fixed-point logic with counting* FPC.

2.1. Extensions by rank operators. Next, we recall the notion of rank operators as introduced in [7]. Let $\Theta(\bar{x}\bar{\nu} \leq \bar{t}, \bar{\nu}\bar{\mu} \leq \bar{s})$ 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 $\bar{\nu}$ and $\bar{\mu}$, respectively. Given a structure $\mathfrak{A}$ we define $\mathbb{N}^{\leq \bar{t}} := \{\bar{n} \in \mathbb{N}^{|\bar{\nu}|} : n_i \leq t_i^{\mathfrak{A}}\}$. The set $\mathbb{N}^{\leq \bar{s}} \subset \mathbb{N}^{|\bar{\mu}|}$ is defined analogously. The term Θ defines in the structure $\mathfrak{A}$ for $I := A^{|\bar{x}|} \times \mathbb{N}^{\leq \bar{t}}$ and $J := A^{|\bar{\nu}|} \times \mathbb{N}^{\leq \bar{s}}$ 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 $p \in \mathbb{P}$. First, as in [7], we define for every prime p a matrix rank operator rk_p which allows us to construct a new numeric *rank term* $[\mathsf{rk}_p\,(\bar{x}\bar{\nu} \leq \bar{t}, \bar{\nu}\bar{\mu} \leq \bar{s})\,.\,\Theta]$ whose value in the structure $\mathfrak{A}$ is the rank of the matrix $(M_\Theta \bmod p)$ over $\mathbb{F}_p$. Second, we consider a *uniform* rank operator rk^* which gets the prime p as an additional input. Formally, with this rank operator rk^* we can construct a rank term $[\mathsf{rk}^*\,(\bar{x}\bar{\nu} \leq \bar{t}, \bar{\nu}\bar{\mu} \leq \bar{s}, \pi \leq r)\,.\,\Theta]$ 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^*

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

can be seen as a unification for the family of rank operators $(\mathsf{rk}_p)_{p\in\mathbb{P}}$ 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 $p \in \Omega$. We set $\mathrm{FOR} = \mathrm{FOR}_\mathbb{P}$ and $\mathrm{FPR} = \mathrm{FPR}_\mathbb{P}$. Similarly, we denote by FPR^* the extension of FPC by the uniform rank operator rk^*. We remark that rank operators can simulate counting terms. For example we have that

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

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

2.2. 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 define the symmetric transitive closures of binary relations and that they can be used to express the structure isomorphism problem on 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 $\varphi(\bar{x}\bar{v}, \bar{y}\bar{\mu}, \bar{z}) \in \mathrm{FPS}_\Omega$ and let $\bar{t}$ and $\bar{s}$ be tuples of closed numeric terms with $|\bar{t}\,| = |\bar{v}|$ and $|\bar{s}\,| = |\bar{\mu}|$. Then $\psi(\bar{z}) = (\mathsf{slv}_p\,\bar{x}\bar{v} \leq \bar{s}, \bar{y}\bar{\mu} \leq \bar{t})\varphi(\bar{x}\bar{v}, \bar{y}\bar{\mu}, \bar{z})$ 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 $(\mathfrak{A}, \bar{z} \mapsto \bar{c}) \in \mathrm{FStr}(\sigma, \bar{z})$ 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-characteristic vector over $\mathbb{F}_p$, i.e., $\mathbb{1}(\bar{a}) = 1$ for all $\bar{a} \in I$. Then M_φ and $\mathbb{1}$ determine the linear equation system $M_\varphi \cdot \vec{x} = \mathbb{1}$ over $\mathbb{F}_p$ where $\vec{x} = (x_j)_{j\in J}$ is a J-vector of variables x_j which range over $\mathbb{F}_p$. Finally, $\mathfrak{A} \models \psi(\bar{c})$ if, and only if, $M_\varphi \cdot \vec{x} = \mathbb{1}$ is solvable over $\mathbb{F}_p$.

At first glance, the solvability quantifiers slv_p seem to impose severe restrictions on the syntactic form of definable linear equation systems. Indeed, they require that *every* linear equation is of the form $\sum_{j\in J} a_j \cdot x_j = 1$ with coefficients a_j from the set $\{0,1\} \subseteq \mathbb{F}_p$. However, this is no restriction at all, since every definable linear equation system can be transformed into this kind of syntactic normal form via a quantifier-free first-order reduction (see Lemma 4.1 in [6]).

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

We write FPS to denote the logic $\mathrm{FPS}_{\mathbb{P}}$ and FPS_p to denote the logic $\mathrm{FPS}_{\{p\}}$ 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.

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

Here, FOS_p denotes the extension of first-order logic by solvability quantifiers slv_p over the field $\mathbb{F}_p$ but *without counting*. For the precise definition of FOS_p see Definition 4.1.

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 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, see [19]. Hence, it suffices to consider rank operators and solvability quantifiers over prime fields.

§3. Separation results over different prime 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 incomparable for different sets of primes Ω. Moreover, we generalise these results to the extensions FPR_Ω by rank operators. In this way we answer the following open question about rank logic:

It holds that $\mathrm{FPR}_p \neq \mathrm{FPR}_q$ 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 3.1. *Let $\Omega \neq \Omega'$ be two distinct sets of primes. Then* $\mathrm{FPS}_\Omega \neq \mathrm{FPS}_{\Omega'}$ *and* $\mathrm{FPR}_\Omega \neq \mathrm{FPR}_{\Omega'}$.

THEOREM 3.2. *Rank logic fails to capture polynomial time. We have*

$$\mathrm{FPR} < \mathrm{FPR}^* \leq \textsc{Ptime}.$$

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

THEOREM 3.3. *For every prime q there is a class of structures $\mathcal{K}_q$ such that*

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

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

PROOF OF THEOREM 3.1. Let Ω and Ω' be two distinct sets of primes. Without loss of generality let us assume that there exists a prime $q \in \Omega \setminus \Omega'$. Then by Theorem 3.3 there exists a class $\mathcal{K}_q$ on which $\mathrm{FPS}_\Omega = \mathrm{FPR}_\Omega = \textsc{Ptime}$ and on which $\mathrm{FPS}_{\Omega'} = \mathrm{FPR}_{\Omega'} = \mathrm{FPC} < \textsc{Ptime}$. ⊣

PROOF OF THEOREM 3.2. Otherwise assume that $\mathrm{FPR} = \textsc{Ptime}$. Then, in particular, $\mathrm{FPR} = \mathrm{FPR}^*$ and there exists a formula $\varphi \in \mathrm{FPR}$ which can uniformly determine the rank of matrices over prime fields, i.e., which can express the uniform rank operator rk^*. As a matter of fact we have $\varphi \in \mathrm{FPR}_\Omega$ for a *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. Hence we have $\mathrm{FPS} \leq \mathrm{FPR} \leq \mathrm{FPR}^* \leq \mathrm{FPR}_\Omega$.

Now let $q \in \mathbb{P} \setminus \Omega$. By Theorem 3.3 there exists a class of structures $\mathcal{K}_q$ on which $\mathrm{FPR}_\Omega = \mathrm{FPC} < \textsc{Ptime}$. However, the class $\mathcal{K}_q$ can be chosen such that $\textsc{Ptime} = \mathrm{FPS}_q \leq \mathrm{FPR}_\Omega$ on $\mathcal{K}_q$ by Theorem 3.3(d) and we obtain the desired contradiction. ⊣

The proof of Theorem 3.2 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 $\textsc{Ptime}$ 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 3.3. 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 proves our theorem.

3.1. Group actions. In the course of this article, we heavily make use of *group actions*. However, we only need basic notions and results which we summarise now. For more details on group actions the reader may consult any standard textbook on groups, such as [13]. In this article, all groups are *finite* permutation groups. For a (finite) set Ω we denote by $\mathrm{Sym}(\Omega)$ the symmetric group acting on the set Ω, and a permutation group is a subgroup of $\mathrm{Sym}(\Omega)$ for some (finite) set Ω. In general, a group $\Gamma \leq \mathrm{Sym}(\Omega)$ *acts on a set* V if there exists a group homomorphism $h\colon \Gamma \to \mathrm{Sym}(V)$. In other words, Γ acts on a set V if we can identify the elements of Γ with permutations on V (in a structure-preserving way). In this article, we will only encounter the following two types of group actions. Let $\Gamma \leq \mathrm{Sym}(\Omega)$. Then, first of all, Γ naturally acts on the set $V = \Omega^k$ consisting of all k-tuples over Ω, for $k \geq 1$, simply by applying permutations $\gamma \in \Gamma$ to tuples $\bar{v} = (v_1, \ldots, v_k) \in V$ component-wise, that is $\gamma(\bar{v}) = (\gamma(v_1), \ldots, \gamma(v_k))$. Second, we can extend this group action to higher-order objects. For instance, assume that instead we consider the set $V = 2^{\Omega^k}$ consisting of all k-ary relations over Ω. Then we can naturally define a (canonical) action of Γ on V by setting $\gamma(R) = \{\gamma(\bar{r}) : \bar{r} \in R\}$ for $R \in V$ (where $\gamma(\bar{r})$ is defined for tuples $r \in R$ as above). In a slightly more general way, we can let Γ act on matrices which are indexed by tuples of elements from Γ, as we will see soon.

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

Now, let Γ be a group which acts on a set V. Given an element $v \in V$, the *orbit* of v is the set $\Gamma(v) = \{\gamma(v) : v \in V\} \subseteq V$. It turns out that the orbits of elements $v \in V$ form a partition of V. One can say much more. Let $\mathrm{Stab}(v) \leq \Gamma$ be the *stabiliser* (*subgroup*) of an element $v \in V$, that is the subgroup of Γ that consists of all $\gamma \in \Gamma$ satisfying $\gamma(v) = v$. Then the *orbit-stabiliser theorem* relates the index of the subgroup $\mathrm{Stab}(v)$ in Γ with the size of the orbit of $v \in V$; we have

$$|\Gamma| = |\mathrm{Stab}(v)| \cdot |\Gamma(v)|.$$

In particular, it follows that the size of an orbit divides the size of the group. We will heavily make use of this fact in the following way. Recall that a group Γ is called a *p-group*, for a prime $p \in \mathbb{P}$, if $|\Gamma| = p^k$, for some $k \geq 0$. Then, if we assume that a p-group Γ acts on a set V, then, by the orbit-stabiliser theorem, we know that all orbits under this action are of size p^ℓ for some $\ell \geq 0$. In particular, if $|V|$ and p are coprime, then the action must have at least one fixed point, i.e., an orbit of size one. This fact will play a central role in our proof later on.

3.2. Reducing rank operators to solvability quantifiers. We start to establish sufficient criteria for the most relevant part of Theorem 3.3 which is the statement in (a). Assume we have a class of structures $\mathcal{K}_q = \mathcal{K}$ with the following properties.

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

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

Lemma 3.4. *If $\mathcal{K}$ satisfies* (I) *and* (II), *then* $\mathrm{FPS}_\Omega = \mathrm{FPC}$ *over $\mathcal{K}$ 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

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

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 $M = M^\mathfrak{A}_{\bar{c}}$ with entries in $\{0, 1\} \subseteq \mathbb{F}_p$ where $I = I^\mathfrak{A}_{\bar{c}} := A^\ell \times \mathbb{N}^{\leq \bar{s}} \subseteq A^\ell \times \mathbb{N}^\lambda$ and $J = J^\mathfrak{A}_{\bar{c}} := A^\ell \times \mathbb{N}^{\leq \bar{t}} \subseteq A^\ell \times \mathbb{N}^\lambda$ that is defined for $\bar{a} \in I$ and $\bar{b} \in J$ as

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

By definition we have $\mathfrak{A} \models \psi(\bar{c})$ if, and only if, the linear equation system $M \cdot \vec{x} = \mathbb{1}$ 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 $M \cdot \vec{x} = \mathbb{1}$ into an equivalent linear system which is *simpler* in the sense that its solvability can be defined in the logic FPC.

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

The reader should observe that each automorphism $\pi \in \Delta_{\mathfrak{A}} = \mathrm{Aut}(\mathfrak{A})$ 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 $\mathrm{Aut}(\mathfrak{A}) = \mathrm{Aut}(\mathfrak{A}^{\#})$.

We set $\Gamma = \Gamma^{\mathfrak{A}}_{\bar{c}} := \mathrm{Aut}(\mathfrak{A}, \bar{c}) \leq \Delta = \Delta_{\mathfrak{A}} = \mathrm{Aut}(\mathfrak{A})$. The group Γ acts on I and J in the natural way. We identify each automorphism $\pi \in \Gamma$ with the corresponding $I \times I$-permutation matrix Π_I and the corresponding $J \times J$-permutation matrix Π_J in the usual way. More precisely, to $\pi \in \Gamma$ we associate the $I \times I$-permutation matrix Π_I 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 $\Gamma = \mathrm{Aut}(\mathfrak{A}, \bar{c})$ we conclude that $(\Pi_I \cdot M \cdot \Pi_J^{-1})(\bar{a}, \bar{b}) = M(\pi(\bar{a}), \pi(\bar{b})) = M(\bar{a}, \bar{b})$ and thus

$$\Pi_I \cdot M \cdot \Pi_J^{-1} = M,$$

or, stated equivalently,

$$\Pi_I \cdot M = M \cdot \Pi_J.$$

This last identity leads to the following central observation.

LEMMA 3.5. *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 $\varphi_{\preceq}(\bar{x}, \bar{y})$ which defines for all $\mathfrak{A} \in \mathcal{K}$ and $\bar{c} \in (A \uplus \mathbb{N})^k$ as above a linear preorder $\preceq$ on A^ℓ which identifies Γ-orbits. Note that, in general, $\Gamma = \mathrm{Aut}(\mathfrak{A}, \bar{c})$ is a strict subgroup of $\Delta = \mathrm{Aut}(\mathfrak{A})$. Thus we cannot 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 $\{c_1, \ldots, c_k\} \cap A$.

The linear preorder $\preceq$ naturally extends to a preorder on the sets I and J with the same properties. Let us write $J = J_0 \preceq J_1 \preceq \cdots \preceq J_{v-1}$ to denote the decomposition

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

of J into the Γ-orbits J_j which are ordered by $\preceq$ as indicated. Moreover, for $j \in [v]$ we let e_j denote the characteristic 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 3.5 this shows the following.

LEMMA 3.6. *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 3.4.

We proceed to show that the conditions (I) and (II) 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 requirements will become clear from the proof of the following lemma.

LEMMA 3.7. *If $\mathcal{K}$ satisfies* (I) and (II), then $\mathrm{FPR}_\Omega = \mathrm{FPS}_\Omega$ over $\mathcal{K}$ for all sets of primes $\Omega \subseteq \mathbb{P}$.

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

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

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 $I^{\mathfrak{A}} = I := A^{|\bar{x}|} \times \mathbb{N}^{\le \bar{t}}$ and $J^{\mathfrak{A}} = J := A^{|\bar{y}|} \times \mathbb{N}^{\le \bar{s}}$ 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 $\varphi_\preceq(\bar{y}_1\bar{\mu}_1, \bar{y}_2\bar{\mu}_2)$, $\psi_\le(\bar{v}, \bar{y}_1\bar{\mu}_1, \bar{y}_2\bar{\mu}_2)$ such that for every $\mathfrak{A} \in \mathcal{K}$*

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

(a) $\varphi_{\preceq}^{\mathfrak{A}}$ *is a linear preorder* $\preceq$ *on* $J^{\mathfrak{A}}$, *and such that*
(b) *For every* $\preceq$*-class* $[j] \subseteq J^{\mathfrak{A}}$ *there exists 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 $\varphi_{\preceq}$ be an FPC-formula which defines in every structure $\mathfrak{A} \in \mathcal{K}$ a linear preorder $\preceq$ on $J^{\mathfrak{A}}$ such that $\preceq$-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 $\vartheta_{\preceq}$ which defines in every structure $\mathfrak{A} \in \mathcal{K}$ a linear preorder $\preceq^*$ 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 $\preceq$-class for some $\mathfrak{A} \in \mathcal{K}$. By property (I) we know that $\Delta_{\mathfrak{A}}$ is an Abelian group. Thus, each automorphism $\pi \in \Delta_{\mathfrak{A}}$ which fixes *one* element in the $\Delta_{\mathfrak{A}}$-orbit $[j]$ point-wise fixes *every* element in the class $[j]$. We conclude that the restriction of $\preceq^*$ to elements in $\{j\} \times [j]$ corresponds to a linear order on $[j]$ for each $j \in [j]$. In this way we obtain an FPC-formula $\psi_{\leq}$ with the desired properties. ⊣

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

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

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

We now focus on the parts (c) and (d) of Theorem 3.3 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 $\mathfrak{A} \ncong \mathfrak{B}$ and $\mathfrak{A} \equiv_k^C \mathfrak{B}$.

Lemma 3.8. *If* $\mathcal{K}$ *satisfies* (III) *and* (IV), *then* $\mathrm{FPC} < \mathrm{FPS}_q = \text{Ptime}$ *on* $\mathcal{K}$.

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

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

3.3. A generalised Cai–Fürer–Immerman construction. It remains to construct a class of structures $\mathcal{K}$ which satisfies (I)–(IV). Our approach is a generalisation of the 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 with an undirected and connected graph $\mathcal{G} = (V, E)$, we first take two copies e_0, e_1 of every edge $e \in E$ for the universe of the associated CFI-graph. For every vertex $v \in V$ we let $vE \subseteq E$ denote the set of edges which are incident with v. The crucial idea of the CFI-construction is to consider, for every vertex $v \in V$, 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 $\rho : vE \to \mathbb{F}_2$ 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 $\rho : vE \to \mathbb{F}_2$ and $\sum_{e \in vE} \rho(e) = 1$ is stabilised (an *odd* node). This restricts the symmetries of the resulting CFI-graphs (which are obtained by twisting the atoms e_0, e_1 for edges $e \in E$) in a clever way.

The constraints for even and odd nodes are encoded by simple graph gadgets. Although it seems that for the same undirected graph $\mathcal{G}$ we obtain exponentially many different CFI-graphs (for each $v \in V$ we can choose one out of two possible constraints), there really are, up to isomorphism, only two such graphs which are determined by the parity of the number of odd nodes. The reason is that if we twist two copies e_0, e_1 of an edge e, then we can move the resulting twist along a path (in the *connected* graph $\mathcal{G}$) to iteratively balance out 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 $e_0, e_1, \ldots, e_{q-1}$ of the edge e. We then add similar constraints for sets of incident edges as above, but instead of having only two different kinds of such constraints, we have one for each of the possible field elements $0, 1, \ldots, q-1 \in \mathbb{F}_q$. Now, instead of twisting pairs of edges, we consider *cyclic shifts* on the edge classes $e_0, e_1, \ldots, e_{q-1}$. 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 construction has been used, for example, in [15, 21].

Formally, we start with a *connected* and *ordered* graph $\mathcal{G} = (V, \leq, E)$ (with symmetric edge relation). We set $\tau := \{\preceq, C, I, R\}$ for binary relation symbols C, I and R. We define, for every $q \in \mathbb{P}$, and for every sequence of *gadget values* $\vec{d} = (d_v)_{v \in V} \in [q]^V$, a τ-structure $\mathrm{CFI}_q(\mathcal{G}, \vec{d})$ which we call a *CFI-structure over* $\mathcal{G}$. For the following construction we implicitly assume that arithmetic is modulo q so that we drop the operator "mod q" in statements of the form $x = y \bmod q$ and $x + y \bmod q$ for better readability. In 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 an 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.

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

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

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

At first glance our construction associates with every graph $\mathcal{G}$ (with the above properties) and to each sequence of gadget values $\vec{d} \in [q]^V$ a different structure $\mathrm{CFI}_q(\mathcal{G}, \vec{d})$. However, for each graph $\mathcal{G}$ with the above properties there really are, up to isomorphism, only q different CFI-structures $\mathrm{CFI}_q(\mathcal{G}, \vec{d})$. 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 $\mathfrak{A} = \mathrm{CFI}_q(\mathcal{G}, \vec{d_1})$ and $\mathfrak{B} = \mathrm{CFI}_q(\mathcal{G}, \vec{d_2})$. For such structures we know that the set $\hat{E}$ of edge nodes, the linear preorder $\preceq$ 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 $\mathfrak{C} := (\hat{E}, (\preceq \restriction \hat{E}), C, I)$ which only depends on $\mathcal{G}$ but not on $\vec{d} \in [q]^V$. Thus

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

Let $\pi \in \Gamma$. The linear preorder $\preceq$ 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., $\pi \in \prod_{e \in E} \langle (\, e_0 e_1 \ldots e_{q-1} \,) \rangle \leq \mathrm{Sym}(\hat{E})$. It is convenient to identify the group $\prod_{e \in E} \langle (\, e_0 e_1 \ldots e_{q-1} \,) \rangle$ 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

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

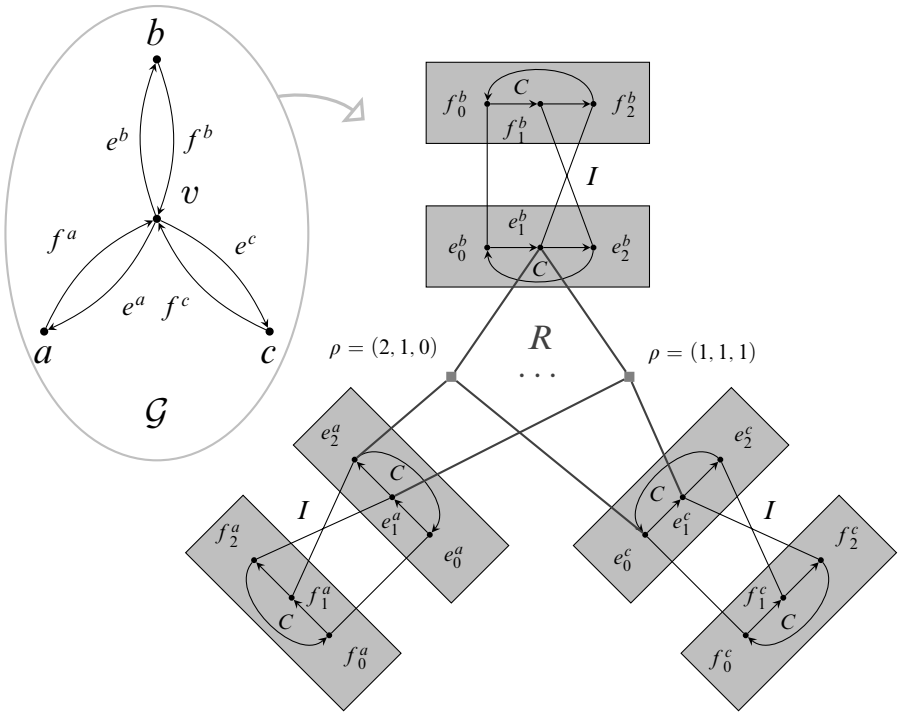

FIGURE 1. Generalised CFI-construction for the v-gadget where $q = 3$ and $\vec{d}(v) = 0$.

$\pi \in \mathbb{F}_q^E$ 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 $\mathfrak{A} = \mathrm{CFI}_q(\mathcal{G}, \vec{d})$, we define for each vertex $v \in V$ the *v-gadget* as the set $\mathrm{gadget}(v) := \hat{v}^{d(v)} \uplus \bigcup_{e \in E(v)} \hat{e}$.

LEMMA 3.9. *Let $\mathfrak{A} = \mathrm{CFI}_q(\mathcal{G}, \vec{d})$ 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 $\rho \in \hat{v} = \hat{v}^{\vec{d}(v)}$ 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 $\hat{\pi}(\rho) : E(v) \to [q]$ is determined as $(\hat{\pi}(\rho))(e) =$

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

$\rho(e) + \pi(e)$ 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 $\mathrm{Iso}(\mathfrak{A}, \mathfrak{B})$ can be identified with a subset of Γ. In fact, the set $\mathrm{Aut}(\mathfrak{A})$ is a subgroup of Γ and $\mathrm{Iso}(\mathfrak{A}, \mathfrak{B})$ is a coset in Γ. Specifically, we saw that $\pi \in \Gamma$ can uniquely be associated 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 3.10. *Γ acts on $\{\mathrm{CFI}_q(\mathcal{G}, \vec{d}) : \vec{d} \in [q]^V\}$. For $\pi \in \Gamma$ we have*

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

LEMMA 3.11. *Let $\vec{d}, \vec{d_*} \in [q]^V$ be sequences of gadget values. Then $\mathrm{CFI}_q(\mathcal{G}, \vec{d}) \cong \mathrm{CFI}_q(\mathcal{G}, \vec{d_*})$ if, and only if, $\sum \vec{d} = \sum \vec{d_*}$.*

PROOF. Let $\pi \in \Gamma$ such that $\pi(\mathrm{CFI}_q(\mathcal{G}, \vec{d})) = \mathrm{CFI}_q(\mathcal{G}, \vec{d_*})$. By Lemma 3.10 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

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

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 $\pi(\mathrm{CFI}_q(\mathcal{G}, \vec{d})) = \mathrm{CFI}_q(\mathcal{G}, \vec{d_+})$. 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 3.10 we know that $\vec{d_+}(v) = \vec{d}(v) + \sum_{e \in E(v)} \pi(e)$. For $v \in V$ it follows that

- if $v \notin \{v_0, \ldots, v_m\}$, 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. ⊣

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

The kind of isomorphism that we constructed in the proof of Lemma 3.11 plays an important role later on. Thus, for a simple path $\bar{p}$ from v_0 to v_m ($m \geq 1$)

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

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 $\pi[\bar{p}, z] = \sigma^z[(v_0, v_1)] + \sigma^z[(v_1, v_2)] + \cdots + \sigma^z[(v_{m-1}, v_m)]$. 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 $m \geq 3$

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

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 $\mathcal{G}$ is and stays connected if we remove any set of less than k vertices. The *connectivity* $\mathrm{con}(\mathcal{G})$ of a graph $\mathcal{G}$ is the maximal $k \geq 1$ such that $\mathcal{G}$ is k-connected. Moreover, the *connectivity* $\mathrm{con}(\mathfrak{G})$ of a class $\mathfrak{G}$ of graphs is the function $\mathrm{con}(\mathfrak{G}) : \mathbb{N} \to \mathbb{N}$ defined by

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

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

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

3.4. Orbits in generalised Cai, Fürer, Immerman structures. Next we show that $\mathcal{K}$ satisfies the required properties (I)–(IV).

First of all, we saw that the automorphism group of each CFI-structure $\mathrm{CFI}_q(\mathcal{G}, \vec{d})$ 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 $\mathfrak{A} = \mathrm{CFI}_q(\mathcal{G}, \vec{d})$ over graphs $\mathcal{G} = (V, \leq, E)$ with $\mathrm{con}(\mathcal{G}) > (\ell + 2)$ since almost all structures in $\mathcal{K}$ satisfy this condition. As above let $\Gamma \leq \mathbb{F}_q^E$ denote the group that actson the set of

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

CFI-structures over $\mathcal{G}$ and let $A := (\hat{V} \uplus \hat{E})$ denote the universe of the CFI-structure $\mathfrak{A}$.

DEFINITION 3.12. 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 $x \in \{a_1, \ldots, a_\lambda\}$ we have $x \in \hat{v}$ $(= \hat{v}^{\vec{d}(v)})$.

(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 $x \in \{a_1, \ldots, a_\lambda\}$ we have that $x \in \hat{e} \cup \hat{f}$ where $f = (w, v) \in E$ is the edge related with e.

LEMMA 3.13. *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 $\mathrm{C}^{\ell+2}_{\infty\omega}$ (using the parameters $\bar{a}$), i.e., for every $c \in$gadget(v) there exists a formula $\vartheta(\bar{x}, y) \in \mathrm{C}^{\ell+2}_{\infty\omega}$ 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 $\mathrm{C}^{\ell+2}_{\infty\omega}$ (given the parameters $\bar{a}$), i.e., for every $c \in \hat{e} \uplus \hat{f}$ there exists a formula $\vartheta(\bar{x}, y) \in \mathrm{C}^{\ell+2}_{\infty\omega}$ such that $\vartheta^{\mathfrak{A}}(\bar{a}) = \{c\}$.*

PROOF. First of all, it is straightforward (even without using the parameters) to fix the $\preceq$-class of any element $c \in A$ in $\mathrm{C}^{\ell+2}_{\infty\omega}$. Second, 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 $\mathrm{C}^{\ell+2}_{\infty\omega}$. Finally, the inverse relation I yields a definable bijection between related edge classes. ⊣

LEMMA 3.14. *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 $\pi \in Aut(\mathfrak{A}, \bar{a})$ 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 $M := \{a_1, \ldots, a_\lambda\} \cap (\bigcup_{e \in E} \hat{e})$. 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 $\mathrm{con}(\mathcal{G}) > (\ell + 2)$ 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 $\mathcal{G}''$:

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

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

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

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

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 $x \in \{a_1, \ldots, a_\lambda\}$. 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 $f \notin E(v) \setminus \{e, e'\}$. ⊣

LEMMA 3.15. *Let $\lambda \leq \ell$ and let $\bar{a}, \bar{b} \in A^\lambda$. Then $(\mathfrak{A}, \bar{a}) \equiv^C_{\ell+2} (\mathfrak{A}, \bar{b})$ if, and only if, there exists $\pi \in Aut(\mathfrak{A})$ such that $\pi(\bar{a}) = \bar{b}$.*

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 $\pi \in \mathrm{Aut}(\mathfrak{A}, a_1 \ldots a_{i-1})$ such that $\pi(a) = b$. Since $\bar{a}$ and $\bar{b}$ have the same $\mathrm{C}^{\ell+2}_{\infty\omega}$-type we know that a and b belong to the same $\preceq$-class. We choose $v \in V$ such that $a, b \in \mathrm{gadget}(v)$.

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 $\{a_1, \ldots, a_{i-1}\}$.

Without loss of generality we may assume that the vertex v is not marked (by an element $x \in \{a_1, \ldots, a_{i-1}\}$), because otherwise, by Lemma 3.13, every element in $\mathrm{gadget}(v)$ can uniquely be identified in $\mathrm{C}^{\ell+2}_{\infty\omega}$. 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 $\mathrm{C}^{\ell+2}_{\infty\omega}$ 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 $\mathrm{C}^{\ell+2}_{\infty\omega}$ by Lemma 3.13. 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 $\mathrm{C}^{\ell+2}_{\infty\omega}$-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 $\mathrm{Aut}(\mathfrak{A}, a_1, \ldots, a_{i-1})$. 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 3.14 to obtain an automorphism $\sigma \in \mathrm{Aut}(\mathfrak{A}, a_1, \ldots, a_{i-1})$ 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 $\pi_* \in \mathrm{Aut}(\mathfrak{A}, a_1, \ldots, a_{i-1})$. But then $(\pi_* + \sigma) \in \mathrm{Aut}(\mathfrak{A}, a_1, \ldots, a_{i-1})$ 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 $\mathrm{con}(\mathcal{G}) > (\ell + 2)$ the minimal

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

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 3.14 to obtain an automorphism $\pi \in \mathrm{Aut}(\mathfrak{A}, a_1, \ldots, a_{i-1})$ 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 $\mathrm{C}^{\ell+2}_{\infty\omega}$-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 3.16. *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 $\mathrm{C}^k_{\infty\omega}$ cannot distinguish between any pair of CFI-structures over $\mathcal{G}$ (although there exist nonisomorphic CFI-structures over $\mathcal{G}$).

LEMMA 3.17. *Let $k \geq 1$ and let $\mathcal{G} = (V, \leq, E) \in \mathfrak{G}$ such that $con(\mathcal{G}) > k$. Then for all $\vec{d}, \vec{d_*} \in [q]^V$ it holds that*

$$\mathrm{CFI}_q(\mathcal{G}, \vec{d}) \equiv^C_k \mathrm{CFI}_q(\mathcal{G}, \vec{d_*}).$$

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

PROOF. Let $\mathfrak{A} = \mathrm{CFI}_q(\mathcal{G}, \vec{d})$ and let $\mathfrak{B} = \mathrm{CFI}_q(\mathcal{G}, \vec{d_*})$. Without loss of generality we assume that $\mathfrak{A} \ncong \mathfrak{B}$. 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 $(\mathfrak{A}, a_1, \ldots, a_\ell, \mathfrak{B}, b_1, \ldots, b_\ell)$ 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 $a_1, \ldots, a_\ell$ in $\mathfrak{A}$ or, equivalently, by the pebbled elements $b_1, \ldots, b_\ell$ 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
- $(\sigma^z[e] + \pi)(\mathfrak{A} \restriction \mathrm{gadget}(v)) = \mathfrak{B} \restriction \mathrm{gadget}(v)$ 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 with 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 $\hat{\pi}_e : A \to B$ 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.

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

Indeed we can find such a good pair (v, π) by Lemma 3.11 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 $(\mathfrak{A}, a_1, \ldots, a_\ell, \mathfrak{B}, b_1, \ldots, b_\ell)$ 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 $\mathrm{con}(\mathcal{G}) > k$ 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 with 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. It follows that the pair $(w, \sigma^z[e] + \pi)$ is good. ⊣

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

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

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

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

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

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

LEMMA 3.18. *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 $\vec{\sigma} \in \mathbb{F}_q^{\mathcal{V}}$ 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 $\vec{\sigma} \in \mathbb{F}_q^{\mathcal{V}}$ of this system defines an isomorphism π between $\mathfrak{A}$ and $\mathfrak{B} = \mathrm{CFI}_q(\mathcal{G}, \vec{d}_+)$ where $\vec{d}_+(v) := \vec{\sigma}(v)$. As a preparation, we let $\delta(e) := \vec{\sigma}(e_i) - i$

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

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

$$\pi(e_i) \mapsto e_{\sigma(e_i)},$$
$$\pi(\rho) \mapsto \rho + \delta.$$

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 3.19. *The class $\mathcal{K}$ satisfies the property* (III).

This finishes our proof of Theorem 3.3.

§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 over which we can show that the expressive power of FPR_Ω and FPS_Ω coincides. This naturally leads to the question whether, in general, rank operators can be simulated by solvability quantifiers in the framework of fixed-point logic with counting. In fact, as we already mentioned in Section 2, most of the queries which are known to separate FPC from rank logic can also be expressed in FPS. In particular, this is interesting because 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 [15, 16, 18]. Clearly, if solvability logic happens to be equivalent to rank logic, then these intermediate problems would be definable in solvability logic as well.

In this section we solve a simplified version of this question and show that, at least in the absence of counting terms, rank operators are strictly more expressive than 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 4.1. For every prime p, the logic FOS_p results by extending the syntax of FO by the following formula creation rule:

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

The semantics of $\psi(\bar{z})$ are defined as above. 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, $\mathfrak{A} \models \varphi(\bar{a}, \bar{b}, \bar{c})$.

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

Let $\mathbb{1}$ be the I-characteristic 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 for linear equation systems in the definition of slv_p-quantifier is no severe restriction (again, see Lemma 4.1 in [6]).

In our main result in this section we show that for every prime $p \in \mathbb{P}$ there is a query $\mathcal{K} \subseteq \mathrm{FStr}(\emptyset)$ over the class of sets which can be expressed in FOR_p but not in FOS_p. In particular, this shows that over $\mathrm{FStr}(\emptyset)$ it holds that

$$\mathrm{FOS}_p < \mathrm{FOR}_p.$$

In some sense this result is not surprising. By contrast to FOR_p, the logic FOS_p does not have access to a counting sort and thus has to express properties of $\mathrm{FStr}(\emptyset)$-structures over pure unordered sets. However, it is not clear how one can turn this intuition into a formal argument. In fact, the logic FOS_p already has nontrivial 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 (recall that $n \equiv 0 \mod p^k$ if, and only if, $n \equiv 0 \mod p$ and $\binom{n}{p} \equiv 0 \mod p^{k-1}$). By contrast, fixed-point logic FP, for example, collapses to first-order logic over sets.

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. Second, 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 $\mathrm{FOS}_p \not\leq \mathrm{FPC}$ 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* $\mathrm{MOD}_k\mathrm{L}$ 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 $\mathrm{MOD}_k\mathrm{L}$ if there exists a nondeterministic 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 $\mathrm{MOD}_p\mathrm{L}$ 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 $\mathrm{MOD}_p\mathrm{L}$. In particular this is true for the solvability problem of linear equation systems over $\mathbb{F}_p$ and for computing the matrix rank over $\mathbb{F}_p$ [3].

Building on these insights, Dawar et. al. showed that for all primes $p \in \mathbb{P}$, the logic FOR_p captures $\mathrm{MOD}_p\mathrm{L}$ on the class of ordered structures. It has been noted in [18] that their proof implies the same correspondence for the logic FOS_p.

PROPOSITION 4.2 ([7], [18]). *Over ordered structures it holds that*

$$\mathrm{FOS}_p = \mathrm{FOR}_p = \mathrm{MOD}_p\mathrm{L}.$$

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

Despite this precise characterisation over the class of ordered structures, the situation over general structures remained unclear. It easily follows that $\mathrm{FOS}_p \leq \mathrm{FOR}_p \leq \mathrm{FPR}$, but, so far, it has been open whether one, or even both, of these inclusions are strict. We give the following partial answer:

THEOREM 4.3. *For all primes p we have $\mathrm{FOS}_p < \mathrm{FOR}_p$ (over the class of sets $FStr(\emptyset)$).*

Before we proceed, let us briefly sketch our proof strategy for Theorem 4.3. We fix a prime $q \in \mathbb{P}, q \neq p$. We then select a class of sets $\mathcal{K}$ which is undefinable in FOR_p, but such that a highly padded variant $\mathcal{K}' = \{([(q^{q^r}])) : ([r]) \in \mathcal{K}\}$ becomes FOR_p-definable. The existence of such a class $\mathcal{K}$ follows from well-known complexity-theoretic arguments (hierarchy theorems). As our main step, we then show that this particular kind of padding (double-exponentiation with basis q) does not extend the expressive power of FOS_p, in contrast to the case of FOR_p. We conclude that FOR_p and FOS_p have different expressiveness.

In order to prove Theorem 4.3, it will be convenient to make use of the following strong normal form for FOS_p which has been established in Corollary 4.8 of [6]:

THEOREM 4.4. *Every formula $\vartheta(\bar{z}) \in \mathrm{FOS}_p$ is equivalent to an FOS_p-formula of the form $(\mathsf{slv}_p\, \bar{x}_1, \bar{x}_2)\alpha(\bar{x}_1, \bar{x}_2, \bar{z})$ where $\alpha(\bar{x}_1, \bar{x}_2, \bar{z})$ is quantifier-free.*

Let us remark that in order to prove our separation result (Theorem 4.3), we only require the normal form stated in Theorem 4.4 for the case of FOS_p-sentences. 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 $\alpha(x_1, \ldots, x_k, y_1, \ldots, y_\ell) \in \mathrm{FO}(\emptyset)$ 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 $\mathfrak{A} \models (\mathsf{slv}_p\, \bar{x}_1, \bar{x}_2)\alpha(\bar{x}_1, \bar{x}_2)$ if the linear equation system $M_n \cdot \vec{x} = \mathbb{1}$ over $\mathbb{F}_p$ is solvable. For convenience we set $I_n = [n]^k$ and $J_n = [n]^\ell$.

Let $\Gamma = \Gamma_n = \mathrm{Sym}([n])$. Then the group Γ acts on I_n and J_n 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. Since M_n is defined by a first-order formula over the empty signature, we conclude that $(\Pi_I \cdot M_n \cdot \Pi_J^{-1})(\bar{a}, \bar{b}) = M_n(\pi(\bar{a}), \pi(\bar{b})) = M_n(\bar{a}, \bar{b})$ and thus $\Pi_I \cdot M_n \cdot \Pi_J^{-1} = M_n$, which can equivalently be written as

$$\Pi_I \cdot M_n = M_n \cdot \Pi_J.$$

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

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

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

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

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 which means that we can *independently* change the indices of rows and columns (within Δ-orbits) without affecting the entry of M_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 (well-)defined as

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

The matrix (M_n^*/Δ) is nothing but a more compact representation of the matrix M_n^* where we shrink each subblock that is indexed by Δ-orbits on I and J to a single entry (a block of size 1×1). Accordingly, the matrix (M_n^*/Δ) gives rise to a new, and more compact, linear equation system

$$(M_n^*/\Delta) \cdot \bar{y} = \mathbb{1}$$

which is solvable, if and only if, the original system $M_n \cdot \bar{x} = \mathbb{1}$ is solvable. To see this, note that we can obtain the matrix (M_n^*/Δ) from the matrix M_n^* by iteratively deleting repeated rows and columns. These operations do not alter the solvability of the linear equation system. As we said, depending on the size of the group Δ, the sets I_n^Δ and J_n^Δ are (much) smaller than the index sets I_n and J_n. More precisely, by the orbit-stabiliser theorem (see Section 3.1), choosing a large group Δ guarantees that we obtain a relatively small linear equation system $(M_n^*/\Delta) \cdot \vec{y} = \mathbb{1}$ which is equivalent to the original one. In the following Section 4.1 we set out to construct such "large" groups Δ. As it will become evident in our proof later, there are certain additional requirements that Δ should meet. Besides of inducing a small number of different orbits on I_n and J_n, we want that the group Δ satisfies the following:

- First of all, in order to be able to switch from the original system $M_n \cdot \bar{x} = \mathbb{1}$ to the reduced system $(M_n^*/\Delta) \cdot \bar{y} = \mathbb{1}$, we require that the size $|\Delta|$ of Δ is coprime to p. We already took care of this by choosing Δ to be a q-group (for a prime q different from p), but a more liberal choice is possible in this respect.
- The reduction $M_n \mapsto (M_n^*/\Delta)$ itself should be of low complexity, that is it should be definable in first-order logic with counting (FOC). To this end, we

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

aim to select Δ in such a way that the structure of Δ-orbits on I_n and J_n is "simple".

As it turns out, one way to meet both requirements is to let Δ be a q-Sylow group that acts on a pure set of size q^r.

4.1. Constructing large groups. Recall that the maximal q-subgroups $\Delta \leq \Gamma$ are called the *q-Sylow groups* of Γ. It is well-known that for the case where $\Gamma = \mathrm{Sym}([n])$ 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

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

In fact, we can write $(q^r)!$ as $(q^r)! = 1 \cdots (1 \cdot q) \cdots (2 \cdot q) \cdots (q^{r-1} \cdot q)$. 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 $\mathrm{Sym}([n])$ 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 $\Gamma = \mathrm{Sym}([n]) = \mathrm{Sym}([q^r])$.

To obtain an algebraic description of these groups, we inductively construct for $r \geq 1$ a q-Sylow subgroup $\Delta_r \leq \mathrm{Sym}([q^r])$ together with a family of trees $\mathcal{T}_i^x$ for $i = 0, \ldots, r$ 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 $\mathcal{P}_i^x = \{x \cdot q^i, \ldots, (x+1) \cdot q^i - 1\}$ (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 tress $\mathcal{T}_i^x$. 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 $\Delta_1 \leq \Delta_2 \leq \cdots \leq \Delta_r$ 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. It is useful to think of elements $y \in [n]$ as being given in their q-adic representation, i.e., $y = y_0 + y_1 \cdot q + \cdots + y_{r-1} \cdot q^{r-1}$. Then we have that $y \in \mathcal{P}_r^0 = [n]$ and

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

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

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 $\gamma = (0\,1\,\cdots\,q-1)$ on the set $\mathcal{P}_1^0 = \{0,\dots,q-1\}$.

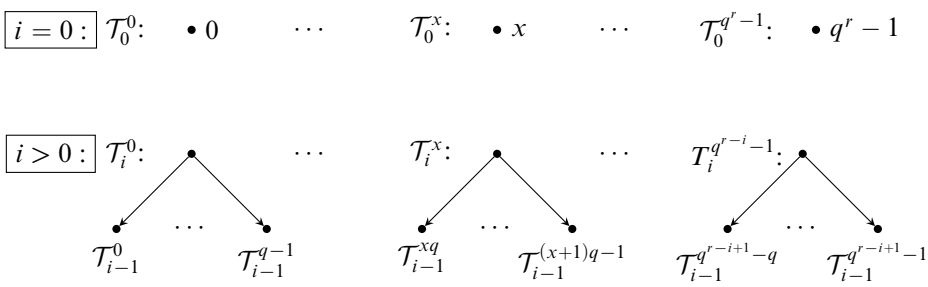

FIGURE 2. Inductive definition of the trees $\mathcal{T}_i^x$.

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 $\gamma \in \mathrm{Sym}([n])$ 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:

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

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 copies of Δ_r which independently act on the segments $\mathcal{P}_r^x$ for $0 \le x \le q-1$. Finally, we define Δ_{r+1} as the *semi-direct product* of $(\Delta_r^0 \times \cdots \times \Delta_r^{q-1})$ 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 $(\Delta_r^0 \times \cdots \times \Delta_r^{q-1} \times \langle\gamma\rangle)$ and that the group operation is given by

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

Since $|\Delta_{r+1}| = |\Delta_r|^q \cdot q$ we conclude that Δ_{r+1} 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 $i \le r$, 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 $z = x + q^{r-i} \bmod q^{r+1-i}$. Hence $(\gamma^y)\{\mathcal{T}_i^x : x \in \mathcal{P}_{r-i}^0\} =$

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

$\{\mathcal{T}_i^x : x \in \mathcal{P}_{r-i}^y\}$ 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* $\mathrm{sgn}(a,b)$ 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 and 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 $y_b = y_a + z \mod q$. For the special case where $i = 0$ we have $a = b$ and agree to set $z = 0$. With this preparation we define the signature $\mathrm{sgn}(\bar{a})$ of a tuple $\bar{a} = (a_1, \ldots, a_\ell) \in J_n$ as the list $\sigma \in ([r+1] \times [q])^{\ell(\ell-1)/2}$ consisting of the individual signatures $\mathrm{sgn}(a_i, a_j)$ 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 4.5. *Let $\bar{a} \in J_n$. Then $sgn(\bar{a}) = sgn(\pi\bar{a})$ for all $\pi \in \Delta_r$.*

PROOF. Immediately follows from the construction of Δ_r and the trees $\mathcal{T}_i^x$. ⊣

LEMMA 4.6. *Let $\bar{a}, \bar{b} \in J_n$. If $sgn(\bar{a}) = sgn(\bar{b})$, 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 $j = 1, \ldots, i$. The case $i = \ell$ is clear, so assume that $i < \ell$. Let $\bar{a} = (a_1, \ldots, a_i, a_{i+1}, \ldots, a_\ell)$ and $\bar{b} = (a_1, \ldots, a_i, b_{i+1}, \ldots, b_\ell)$. We show that there exists a permutation $\delta \in \Delta_r$ which pointwise fixes $a_1, \ldots, a_i$ and such that $\delta(a_{i+1}) = b_{i+1}$. Then the claim follows from Lemma 4.5 together with the induction hypothesis. For $i = 0$ this is easy, because Δ_r acts transitively on $[n]$. If $i > 0$ we choose $a_w \in \{a_1, \ldots, a_i\}$ such that $\mathrm{sgn}(a_w, a_{i+1}) = (c,d)$ 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 $z = y + d \mod q$. Since $\mathrm{sgn}(\bar{a}) = \mathrm{sgn}(\bar{b})$ 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 non of the elements $\{a_1, \ldots, a_i\}$ 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 $a_1, \ldots, a_i$) and which moves a_{i+1} to b_{i+1}. ⊣

4.2. Defining sizes of orbits in FOC. Following our definition from above, the signature $\mathrm{sgn}(\bar{a})$ 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 $\mathrm{sgn}(\bar{a}) = \sigma$ 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

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

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 $\tau(x_1, \ldots, x_k, x_{k+1}, \ldots, x_{k+\ell})$ 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 $\alpha \in \mathrm{FO}(\emptyset)$ 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 $R : [r] \to [q]$ 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 $x = (x_0, \ldots, x_{r-1}) \in [q]^r$, 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 $y = x_{r-i} + x_{r-i+1}q + \cdots + x_{r-1}q^{i-1}$. 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 $x_{r-i}, \ldots, x_{r-1}$ in the q-adic encoding of x are determined (i.e., x is a member of the block $\mathcal{P}_{r-i}^y$).

LEMMA 4.7. *For all $\ell \geq 1$ and $0 \leq s \leq \ell$ there exist*

(a) *A term $\Theta(i, h_1, d_1, \ldots, h_t, d_t) \in \mathrm{FOC}(\{R_x, R_1, \ldots, R_s\})$ and*
(b) *Formulas $\varphi_e(y, z, i, h_1, d_1, \ldots, h_t, d_t) \in \mathrm{FOC}(\{R_x, R_1, \ldots, R_s\})$ for $e = s+1, \ldots, \ell$,*

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

where $t = \binom{\ell}{2}$, *such that for all* $r \geq q$, *all* $i \leq r$, *all* $\sigma = ((h_1, d_1), \ldots, (h_t, d_t)) \in S_n^\ell$ *where* $n = q^r$, *all* $x \in [q^{r-i}]$ *and all* $a_1, \ldots, a_s \in \mathcal{P}_i^x$ *the following holds: let* $\mathfrak{A} = ([r])$ *and let* $R_x, R_1, \ldots, R_s$ *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* $\Theta^{\mathfrak{A}}(q, i, h_1, d_1, \ldots, h_t, d_t)$ *of the term* Θ *in* $\mathfrak{A}$ *is* $|Z|$ *mod* p *where*

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

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

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

We proceed by induction on ℓ. For $\ell = 1$ it suffices to show that FOC can compute (n mod 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 $q^r \equiv q^{r'} \bmod p$ if $r' \equiv r \bmod (p-1)$. 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 $\{2, \ldots, \ell\}$ 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 $Y_1, \ldots, Y_v$ and let $\sigma(1, y) = (h_w, d_w)$ for all $y \in Y_w$ and $w = 1, \ldots, v$.

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 $w = 1, \ldots, v$ we have $h_w \leq i$, and
- For every $Y_w = \{y_1^w, \ldots, y_{\ell_w}^w\}$ 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 $\{2, \ldots, \ell\}$ 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 $w = 1, \ldots, v$, then we cannot use the induction hypothesis

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

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 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 $a_1, \ldots, a_s \in \mathcal{P}_i^x$ be the components of the tuple $\bar{a}$ that are already fixed. Recall that we can assume without loss of generality that $\operatorname{sgn}(a_1, \ldots, a_s)$ is consistent with σ and that all elements $a_1, \ldots, a_s$ 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 $a_2, \ldots, a_s$. Moreover, when we form the partitions of $\{2, \ldots, \ell\}$ into parts $Y_1, \ldots, Y_v$ 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. ⊣

LEMMA 4.8. *Let $\tau(x_1, \ldots, x_k, y_1, \ldots, y_\ell) \in \mathrm{FO}(\emptyset)$ 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 $\Theta_\tau^{\mathfrak{A}}(\sigma_{\bar{a}}, \sigma_{\bar{b}})$ of Θ_τ in $\mathfrak{A} = ([r])$ is*

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

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

PROOF. By Lemma 4.7 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 $x_1, \ldots, x_k$) is not consistent with $\sigma_{\bar{a}}$ or if τ (restricted to $y_1, \ldots, y_\ell$) contradicts $\sigma_{\bar{b}}$, then the answer is trivial as well.

In all other cases, Lemma 4.7 provides FOC-formulas which define in the structure $\mathfrak{A}$ the q-adic encoding of elements $a_1, \ldots, a_k \in [n]$ such that $\operatorname{sgn}(\bar{a}) = \sigma_{\bar{a}}$. 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 $Y \subseteq \{1, \ldots, \ell\}$ and a partial assignment $\epsilon : \{1, \ldots, \ell\} \to \{a_1, \ldots, a_k\}$ with $\operatorname{dom}(\epsilon) \cap Y = \emptyset$ we define the set

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

In this notation our aim is to determine $(|B_Y^\emptyset| \bmod p)$ for $Y = [\ell]$ in FOC. The first observation is that by Lemma 4.7 we can determine $(|B_\emptyset^\epsilon| \bmod p)$ for all partial assignments ϵ in FOC. The second observation is that we can construct the values

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

$(|B^{\epsilon}_{Y}| \bmod p)$ by induction on $|Y|$ as follows. For $Y \subseteq \{1,\dots,\ell\}$ and a partial assignment ϵ (with $\mathrm{dom}(\epsilon) \cap Y = \emptyset$) we have for all $j \in Y$ that

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

In this way we recursively obtain the value $(|B^{\emptyset}_{Y}| \bmod p)$ 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. ⊣

LEMMA 4.9. *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 $(S^k_n \times S^\ell_n)$-matrix over $\mathbb{F}_p$. To represent the index sets S^k_n and S^ℓ_n 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.

The entry $M^*_n(\sigma_{\bar{a}},\sigma_{\bar{b}})$ of M^*_n for $\sigma_{\bar{a}} \in S^k_n$ and $\sigma_{\bar{b}} \in S^\ell_n$ is given as

$$M^*_n(\sigma_{\bar{a}},\sigma_{\bar{b}}) = |\{\bar{b} \in J_n : \mathrm{sgn}(\bar{b}) = \sigma_{\bar{b}}, M_n(\bar{a},\bar{b}) = 1\}| \cdot |\,\mathrm{Stab}(\bar{b})| \mod p,$$

for some (or, equivalently, all) $\bar{a} \in I_n, \bar{b} \in J_n$ with $\mathrm{sgn}(\bar{a}) = \sigma_{\bar{a}}$ and $\mathrm{sgn}(\bar{b}) = \sigma_{\bar{b}}$. The entry $M_n(\bar{a},\bar{b})$, in turn, is determined by the quantifier-free formula $\alpha(\bar{x}_1,\bar{x}_2) \in \mathrm{FO}(\emptyset)$. Lemma 4.8 shows that we can determine the value of the left-hand side of the above equation for the case where α is a complete equality type. For the general case, we write α as the union of complete equality types and combine the constant number of intermediate results. Moreover, we can determine $|\,\mathrm{Stab}(\bar{b})|$ by Lemma 4.7 (which shows that the size of the orbit of $\bar{b}$ is definable) and by the orbit-stabiliser theorem. ⊣

DEFINITION 4.10. Let $\mathcal{K} \subseteq \mathrm{FStr}(\emptyset)$ be a class of sets. The *q-power* $\mathcal{K}^q \subseteq \mathit{FStr}(\emptyset)$ *of* $\mathcal{K}$ consists of all sets $\mathfrak{A} = ([q^r])$ such that $\mathfrak{B} = ([r]) \in \mathcal{K}$.

THEOREM 4.11. *Let* $\mathcal{K} \subseteq \mathit{FStr}(\emptyset)$ *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 4.4 we can also find a formula $\varphi = (\mathsf{slv}_p\,\bar{x}_1,\bar{x}_2)\alpha(\bar{x}_1,\bar{x}_2) \in \mathrm{FOS}_p$ that defines $\mathcal{K}^q$ such that α is quantifier-free.

By using the above construction and Lemma 4.9, we conclude that the linear equation system $M_n \cdot \vec{x} = \mathbb{1}$ defined by α in an input structure $\mathfrak{A} = ([n])$ of size $n = q^r$ can be transformed into the equivalent system $M^*_n \cdot \vec{x} = \mathbb{1}$ which is FOC-definable in $\mathfrak{B} = ([r])$. Let $\varphi^* \in \mathrm{FOR}_p$ be a formula which expresses the solvability of the linear system $M^*_n \cdot \vec{x} = \mathbb{1}$ in a structure $\mathfrak{B} = ([r])$.

Then $\mathfrak{B} \models \varphi^*$ if, and only if, $\mathfrak{A} \models \varphi$ since the linear equation systems $M_n \cdot \vec{x} = \mathbb{1}$ and $M^*_n \cdot \vec{x} = \mathbb{1}$ are equivalent. ⊣

THEOREM 4.12. *For all* $p \in \mathbb{P}$ *we have* $\mathrm{FOS}_p < \mathrm{FOR}_p$ (*already over* $\mathit{FStr}(\emptyset)$).

PROOF. Suppose for the sake of a contradiction that $\mathrm{FOS}_p = \mathrm{FOR}_p$. As above we fix some prime $q \neq p$. Let $\mathcal{K} \subseteq \mathrm{FStr}(\emptyset)$ be a class of sets such that $\mathcal{K} \notin \mathrm{FOR}_p$, but such that $(\mathcal{K}^q)^q \in \mathrm{FOR}_p$. Such a class $\mathcal{K}$ is well-known to exist. In fact, it follows from the space-hierarchy theorem, see, e.g., [20], that there exists a language $L \subseteq \{1^n : n \in \mathbb{N}\}$ such that $L \in \textsc{Space}(2^{cn})$ and $L \notin \textsc{Pspace}$. But then for an appropriate prime q we have that $L' = \{q^{q^n} : 1^n \in L\} \in \textsc{Logspace}$. Since, over

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

sets, we have $\text{LOGSPACE} \leq \text{FOR}_p \leq \text{PTIME} \leq \text{PSPACE}$, this shows that we can choose $\mathcal{K} = \{([n]) : 1^n \in L\}$.

Now, since we assumed that $\text{FOS}_p = \text{FOR}_p$ we have $(\mathcal{K}^q)^q \in \text{FOS}_p$ and by Theorem 4.11 this means that $\mathcal{K}^q \in \text{FOR}_p$. Again, since $\text{FOR}_p = \text{FOS}_p$, we have $\mathcal{K}^q \in \text{FOS}_p$. A second application of Theorem 4.11 yields $\mathcal{K} \in \text{FOR}_p$, which contradicts our assumptions. ⊣

Let us remark that the same proof also works for the extension of *fixed-point logic* by solvability quantifiers (but still in the absence of counting). The simple reason is that, in the absence of counting, fixed-point operators do not increase the expressive power of first-order logic over the empty signature, since all definable relations consist of constantly many basic building blocks (and thus we can evaluate fixed points already in first-order logic). In other words, if we denote by FPS_p^- the extension of fixed-point logic by solvability quantifiers slv_p over $\mathbb{F}_p$ (without counting), then we have $\text{FOS}_p = \text{FPS}_p^-$ over $\text{FStr}(\emptyset)$.

THEOREM 4.13. *For all primes p, we have $\text{FPS}_p^- < \text{FOR}_p$ over $\text{FStr}(\emptyset)$.*

Finally, another interesting consequence is that there exists an FPC-definable query over $\text{FStr}(\emptyset)$ which cannot be defined in FPS_p^-. This immediately follows from our proof of Theorem 4.11, since the solvability of the linear equation system $M_n^* \cdot \vec{x} = \mathbb{1}$ matrix can also be expressed in FPC (we interpret the coefficient matrix M_n^* over the second *ordered* sort). Note that, in contrast, we have no proof which shows that FPC cannot be embedded into FOR_p.

§5. **Discussion.** We have shown 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 efficiently be 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 3.7 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

https://doi.org/10.1017/jsl.2018.33 Published online by Cambridge University Press

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*, ***Mathematical Foundations of Computer Science 2014*** (E. Csuhaj-Varjù, M. Dietzfelbinger, and Z. Ésik, editors), Lecture Notes in Computer Science, vol. 8634, Springer, Berlin, 2014, pp. 50–62.

[2] A. Atserias, A. Bulatov, and A. Dawar, *Affine systems of equations and counting infinitary logic*. ***Theoretical Computer Science***, vol. 410 (2009), pp. 1666–1683.

[3] G. Buntrock, U. Hertrampf, C. Damm, and C. Meinel, *Structure and importance of logspace-mod-classes*, ***Symposium on Theoretical Aspects of Computer Science '91*** (C. Choffrut and M. Jantzen, editors), Springer, Berlin, 1991, pp. 360–371.

[4] J. Cai, M. Fürer, and N. Immerman, *An optimal lower bound on the number of variables for graph identification*. ***Combinatorica***, vol. 12 (1992), no. 4, pp. 389–410.

[5] A. Dawar, *The nature and power of fixed-point logic with counting*. ***SIGLOG News***, vol. 2 (2015), pp. 8–21.

[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***, vol. 9 (2013), no. 4.

[7] A. Dawar, M. Grohe, B. Holm, and B. Laubner, *Logics with rank operators*, ***Proceedings of the 24th Annual Symposium on Logic in Computer Science, LICS 2009***, IEEE Computer Society, Washington, DC, 2009, pp. 113–122.

[8] A. Dawar and B. Holm, *Pebble games with algebraic rules*, ***Automata, Languages, and Programming-ICALP 2012*** (A. Czumaj, K. Mehlhorn, A. Pitts, and R. Wattenhofer, editors), Springer, Berlin, 2012, pp. 251–262.

[9] H.-D. Ebbinghaus and J. Flum, ***Finite Model Theory***, second ed., Springer-Verlag, Berlin, 1999.

[10] E. Grädel, P. Kolaitis, L. Libkin, M. Marx, J. Spencer, M Vardi, Y. Venema, and S. Weinstein, ***Finite Model Theory and Its Applications***, Springer, 2007.

[11] M. Grohe, *The quest for a logic capturing PTIME*, ***Proceedings of the 23rd Annual Symposium on Logic in Computer Science, LICS 2008***, IEEE Computer Society, Washington, DC, 2008, pp. 267–271.

[12] Y. Gurevich and S. Shelah, *On finite rigid structures*, this Journal, vol. 61 (1996), no. 02, pp. 549–562.

[13] M. Hall, ***The Theory of Groups***, American Mathematical Society, Providence, RI, 1976.

[14] U. Hertrampf, S. Reith, and H. Vollmer, *A note on closure properties of logspace mod classes*. ***Information Processing Letters***, vol. 75 (2000), no. 3, pp. 91–93.

[15] B. Holm, ***Descriptive complexity of linear algebra***, Ph.D. thesis, University of Cambridge, 2010.

[16] B. Laubner, ***The structure of graphs and new logics for the characterization of polynomial time***, Ph.D. thesis, Humboldt-Universität Berlin, 2011.

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

[18] W. Pakusa, ***Finite model theory with operators from linear algebra***, Staatsexamensarbeit, RWTH Aachen University, 2010.

[19] ———, ***Linear equation systems and the search for a logical characterisation of polynomial time***, Ph.D. thesis, RWTH Aachen University, 2016.

[20] C. Papadimitriou, ***Computational Complexity***, Addison-Wesley, Boston, 1995.

[21] J. Torán, *On the hardness of graph isomorphism*. ***SIAM Journal on Computing***, vol. 33 (2004), no. 5, pp. 1093–1108.

MATHEMATICAL FOUNDATIONS OF COMPUTER SCIENCE
RWTH AACHEN UNIVERSITY
D-52056 AACHEN, GERMANY
E-mail: graedel@logic.rwth-aachen.de
E-mail: pakusa@logic.rwth-aachen.de