Pebble Games for Rank Logics

Anuj Dawar and Bjarki Holm

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

Abstract. We show that equivalence in finite-variable infinitary logic with rank operators can be characterised in terms of pebble games based on set partitions. This gives us a game-based method for proving lower bounds for FOR and IFPR, the extensions of first-order and fixed-point logic with rank operators, respectively. As an illustration of the game method, we establish that over finite structures, $\mathrm{IFPR}_p^{[2]} \neq \mathrm{IFPR}_q^{[2]}$ for distinct primes p and q, where $\mathrm{IFPR}_p^{[m]}$ is the restriction of IFPR that only has operators for defining rank of matrices of arity at most m over GF_p.

1 Introduction

The question of whether there is a logical characterisation of the complexity class PTIME remains the fundamental open problem of descriptive complexity. Most attempts to answer this question have focused on finding suitable extensions of first-order logic that can describe exactly all properties decidable in PTIME. In this way, Immerman and Vardi independently showed that on inputs equipped with a linear order, inflationary fixed-point logic (IFP) expresses exactly the properties in PTIME [6, 9]. In the absence of an order, IFP is too weak to express all properties in PTIME. In particular, it fails to define very simple cardinality properties. This immediate deficiency is easily solved by extending the logic with counting terms, which gives us fixed-point logic with counting (IFPC), which was at one time conjectured to be a logic for PTIME. However, Cai, Fürer and Immerman later showed that this logic still falls short of capturing PTIME [2].

Since the result of Cai *et al.*, a number of examples have been constructed of polynomial-time decidable properties that are not expressible in IFPC. Recently it was observed that all these examples can be reduced to the problem of determining the solvability of systems of linear equations (see [1, 3, 4]). Over finite fields, this can be further reduced to the problem of computing the *rank* of a definable (unordered) matrix. Computing rank can be understood as a generalised form of counting where, rather than counting the cardinality of a definable set, one is allowed to count the dimension of a definable vector space. This suggests that the key weakness of IFPC is that the form of counting it incorporates is too weak. In [4], Dawar *et al.* proposed fixed-point logic with rank (IFPR), an extension of IFP with operators for computing the rank of a matrix over a fixed prime field. It is shown in [4] that IFPR can express various polynomial-time properties known to separate IFPC from PTIME. It is an open question whether IFPR captures PTIME.

Despite some positive results on the expressive power of logical rank operators, not much is known about their limitations. For instance, it is not even known whether first-order logic with rank (FOR) is strictly less expressive than IFPR over finite structures, although that would seem likely. To establish such separations we seem to lack general methods for proving inexpressiblity in FOR and IFPR. This leads us to consider variations of Ehrenfeucht-Fraïssé-style pebble games, which form an essential tool for analysing expressiveness of other extensions of first-order logic, such as IFP and IFPC.

In this abstract we introduce a new pebble game based on set partitions that characterises expressivity in an infinitary logic with rank quantifiers, which subsumes both FOR and IFPR. This type of game turns out to be quite generic, with standard games for both IFP and IFPC occurring as special cases. As an illustration of the game method, we establish that over finite structures, $\mathrm{IFPR}_p^{[2]} \neq \mathrm{IFPR}_q^{[2]}$ for distinct primes p and q, where $\mathrm{IFPR}_p^{[m]}$ is the restriction of IFPR that only has operators for defining rank of matrices of arity at most m over GF_p. This partially resolves one of the open questions posed in [4]. Due to space constraints, details of all proofs are omitted.

2 Rank Logics

We assume that all structures are finite and all vocabularies are finite and relational. For a logic $\mathcal{L}$, we write $\mathbf{A} \equiv^{\mathcal{L}} \mathbf{B}$ to denote that the structures $\mathbf{A}$ and $\mathbf{B}$ are not distinguished by any sentence of $\mathcal{L}$. We write $|\mathbf{A}|$ for the universe of a structure $\mathbf{A}$ and write $\|\mathbf{A}\|$ for the cardinality of $|\mathbf{A}|$. We often denote tuples $(v_1, \ldots, v_k)$ by $\boldsymbol{v}$ and denote their length by $|\boldsymbol{v}|$.

Inflationary fixed-point logic (IFP) is obtained by adding to first-order logic the ability to define inflationary fixed-points of inductive definitions. It is easily shown that on finite structures, IFP fails to express very simple cardinality queries. We define counting terms $\#x\varphi$ to denote the number of elements that satisfy the formula φ. By adding to IFP rules for building counting terms, we obtain inflationary fixed-point logic with counting (IFPC). For a detailed discussion of these logics we refer to the standard literature [5, 7].

Definable matrices over finite fields. We write $[m]$ to denote the set $\{0, \ldots, m-1\}$, for $m \geq 1$. For sets I and J, an $I \times J$ matrix over the prime field GF_p can be seen as a function $M : I \times J \to [p]$. Here the rows of M are indexed by I and the columns of M are indexed by J. Observe that the sets I and J are not necessarily ordered. Natural matrix properties, such as singularity and rank, are invariant under permutations of rows and columns, and are therefore well-defined in the context of unordered row and column sets.

Using our notation for describing matrices, a formula $\varphi(\boldsymbol{x}, \boldsymbol{y})$ interpreted in a structure $\mathbf{A}$ defines a GF_2 matrix $M_\varphi^{\mathbf{A}} : A^{|\boldsymbol{x}|} \times A^{|\boldsymbol{y}|} \to \{0, 1\}$ given by $M_\varphi^{\mathbf{A}}(\boldsymbol{a}, \boldsymbol{b}) = 1$ if, and only if, $(\mathbf{A}, \boldsymbol{a}, \boldsymbol{b}) \models \varphi$. More generally, let $\Phi = (\varphi_1(\boldsymbol{x}, \boldsymbol{y}), \ldots, \varphi_l(\boldsymbol{x}, \boldsymbol{y}))$ be an l-tuple of formulas, with $1 \leq l < p$ and p prime. Interpreted in a structure $\mathbf{A}$, these formulas define a matrix $M_\Phi^{\mathbf{A}} : A^{|\boldsymbol{x}|} \times A^{|\boldsymbol{y}|} \to [p]$ given by

$$M_\Phi^{\mathbf{A}}(\boldsymbol{a}, \boldsymbol{b}) = \sum_{i=1}^{l} i M_{\varphi_i}^{\mathbf{A}}(\boldsymbol{a}, \boldsymbol{b}) \pmod{p}.$$

For example, for any formula $\varphi(x)$, the formula $(x = y \wedge \varphi(x))$ interpreted in a structure $\mathbf{A}$ defines a square diagonal matrix, with 1 in position $(a, a) \in A \times A$ on the diagonal if, and only if, $(\mathbf{A}, a) \models \varphi$.

Fixed-point logic with rank. We recall the basic definition of rank logics. To simplify the transition to infinitary rank logics later, our presention of rank operators differs slightly from that of Dawar *et al.* [4], although the two definitions can be seen to be equivalent. Specifically, in [4] we consider matrices over GF_p defined by a single number term modulo p, instead of looking at tuples of formulas as we do below.

Inflationary fixed-point logic with rank (IFPR) has two sorts of variables: $x_1, x_2, \ldots$ ranging over the domain elements of the structure, and $\nu_1, \nu_2, \ldots$ ranging over the non-negative integers. All quantification of number variables has to be bounded. Thus, is if ν is a number variable, its binding quantifier must appear in the form $(\forall \nu \leq t\ \varphi)$ or $(\exists \nu \leq t\ \varphi)$ for a numeric term t and a formula φ. In addition, we also have second-order variables $X_1, X_2, \ldots$, each of which has a type which is a finite string in $\{\mathrm{element}, \mathrm{number}\}^*$. Thus, if X is a variable of type (element,number), it is to be interpreted by a binary relation relating elements to numbers. We write $\mathbf{ifp}_{X \leftarrow \boldsymbol{x}\boldsymbol{\nu} \leq \boldsymbol{t}} \varphi$ for the inflationary fixed-point of φ over the relation variable X of type $(\mathrm{element}^{|\boldsymbol{x}|}, \mathrm{number}^{|\boldsymbol{\nu}|})$, where the number variables in $\boldsymbol{\nu}$ are bounded by the numeric terms in $\boldsymbol{t}$. By closing first-order logic under the formation of inflationary fixed-points, we get IFP in a two-sorted setting. The logic IFPR is obtained by extending the formula-formation rules of IFP with a rule for building *rank terms* in the following way:

> for prime p and $l \in \{1, \ldots, p-1\}$, if $\Phi = (\varphi_1(\boldsymbol{x}, \boldsymbol{y}), \ldots, \varphi_l(\boldsymbol{x}, \boldsymbol{y}))$ is an l-tuple of formulas and $\boldsymbol{x}$ and $\boldsymbol{y}$ are tuples of variables of the first sort, then $\mathrm{rk}_p(\boldsymbol{x}, \boldsymbol{y})\Phi$ is a term.

The intended semantics is that $\mathrm{rk}_p(\boldsymbol{x}, \boldsymbol{y})\Phi$ denotes the rank (i.e. the member of the number sort) over GF_p of the matrix defined by the formulas Φ. More generally, we can define rank terms for formulas Φ with number variables. In this case, all free number variables have to be bounded by number terms, as is described in more detail in [4]. The arity of a rank operator $\mathrm{rk}_p(\boldsymbol{x}, \boldsymbol{y})$ is $|\boldsymbol{x}| + |\boldsymbol{y}|$, where $\boldsymbol{x}$ and $\boldsymbol{y}$ are assumed to be tuples of distinct variables.

We write $\mathrm{IFPR}^{[m]}$ for the fragment of IFPR in which all rank operators have arity at most m and write IFPR_p to denote the fragment where only rank operators rk_p are allowed. Putting the two together, we obtain logics $\mathrm{IFPR}_p^{[m]}$ where only rank operators rk_p of arity at most m are allowed.

It is easy to see that rank logics can express the cardinality of any definable set. Indeed, for a formula $\varphi(x)$ and prime p, the rank term $\mathrm{rk}_p(x,y)(x = y \wedge \varphi(x))$ is equivalent to the counting term $\#x\varphi$, as the rank of a diagonal matrix is exactly the number of non-zero entries along the diagonal. This immediately implies that each of the rank logics IFPR_p is at least as expressive as IFPC.

Infinitary rank logics. For each natural number i and prime p, we consider a quantifier rk_p^i where $\mathbf{A} \models \mathrm{rk}_p^i \boldsymbol{xy}\,(\varphi_1, \ldots, \varphi_{p-1})$ if, and only if, the rank of the $|\mathbf{A}|^{|\boldsymbol{x}|} \times |\mathbf{A}|^{|\boldsymbol{y}|}$ matrix defined by $(\varphi_1(\boldsymbol{x},\boldsymbol{y}), \ldots, \varphi_{p-1}(\boldsymbol{x},\boldsymbol{y}))$ over $\mathbf{A}$ is i. Here the rank is taken over GF_p. Let R^k denote k-variable infinitary logic with rank quantifiers. The logic R^ω is given by $R^\omega = \bigcup_{k\in\omega} R^k$. That is, R^ω consists of infinitary rank formulas in which each formula has only finitely many variables. We let R_p^k denote the sublogic of R^k where only rank quantifiers of the form rk_p^i are allowed. We also write $R^{k;[m]}$ and $R_p^{k;[m]}$ to denote the fragments of R^k and R_p^k, respectively, with rank quantifiers of arity at most m, where the arity of a quantifier $\mathrm{rk}_p^i \boldsymbol{xy}$ is $|\boldsymbol{x}|+|\boldsymbol{y}|$. Clearly, $m \leq k$. It can be shown that every formula of $\mathrm{IFPR}_p^{[m]}$ is equivalent to one of $R_p^{\omega;[m]} = \bigcup_{k\in\omega} R_p^{k;[m]}$. Hence, $\mathrm{IFPR} \subseteq \mathrm{R}^\omega$. It is shown in [4] that for any $m \geq 2$, $R^{k;[m]}$ is strictly less expressive than $R^{k;[m+1]}$. Hence also $\mathrm{IFPR}^{[m]} \subsetneq \mathrm{IFPR}^{[m+1]}$.

3 Games for Logics with Rank

We give a game characterisation of equivalence in the logics $R_p^{k;[m]}$. To describe the game we will use the following notation. Let I and J be finite sets, $\mathbf{P}$ a set partition of $I \times J$, and $\gamma : \mathbf{P} \to [p]$ a labeling of the parts in $\mathbf{P}$, with p prime. Then $M_\gamma^{\mathbf{P}}$ denotes the $I \times J$ matrix over GF_p defined by

$$M_\gamma^{\mathbf{P}}(i,j) = \alpha \in [p] \Leftrightarrow \exists P \in \mathbf{P}\ \big((i,j) \in P \wedge \gamma(P) = \alpha\big).$$

We first consider the game for $R_p^{k;[m]}$ when $m = 2$. The game board of the k-pebble 2-ary rank partition game over GF_p consists of two structures $\mathbf{A}$ and $\mathbf{B}$ and k pairs of pebbles $(a_i, b_i), 1 \leq i \leq k$. The pebbles $a_1, \ldots, a_l$ are initially placed on the elements of an l-tuple $\boldsymbol{s}$ of elements in $\mathbf{A}$, and the pebbles $b_1, \ldots, b_l$ on an l-tuple $\boldsymbol{t}$ in $\mathbf{B}$, $l \leq k$. There are two players, Spoiler and Duplicator. At each round, Spoiler picks up two pairs of corresponding pebbles (a_i, b_i) and (a_j, b_j) for some i and j. Duplicator has to respond by choosing (a) partitions $\mathbf{P}$ of $A \times A$ and $\mathbf{Q}$ of $B \times B$, with $|\mathbf{P}| = |\mathbf{Q}|$; and (b) a bijection $f : \mathbf{P} \to \mathbf{Q}$, such that for all labelings $\gamma : \mathbf{P} \to [p]$,

$$\mathrm{rk}_p(M_\gamma^{\mathbf{P}}) = \mathrm{rk}_p(M_{f(\gamma)}^{\mathbf{Q}}).$$

Here $f(\gamma) : \mathbf{Q} \to [p]$ is the labeling of $\mathbf{Q}$ defined by $f(\gamma)(Q) = \gamma(f^{-1}(Q)$ for all $Q \in \mathbf{Q}$. Spoiler next picks a part $P \in \mathbf{P}$, and places the pebbles (a_i, a_j) on an element in $P \subseteq A \times A$ and places the pebbles (b_i, b_j) on an element in $f(P) \subseteq B \times B$. This completes one round in the game. If, after this exchange, the partial map $f : \mathbf{A} \to \mathbf{B}$ given by $a_i \mapsto b_i$ is not a partial isomorphism, or Duplicator is unable to produce the required partitions, then Spoiler has won the game; otherwise it can continue for another round.

For the more general case of m-ary rank quantifiers over GF_p, we modify the above game so that at each round, Spoiler starts by choosing two integers r and s with $r + s = m$. He then picks up m pebbles in some order from $\mathbf{A}$ and the m corresponding pebbles in the same order from $\mathbf{B}$. Duplicator has to respond by choosing partitions $\mathbf{P}$ and $\mathbf{Q}$ of $A^r \times A^s$ and $B^r \times B^s$, respectively, and a bijection $f : \mathbf{P} \to \mathbf{Q}$ between the two partitions. The rest of the round proceeds in exactly the same way as above, with Spoiler finally choosing a part $P \in \mathbf{P}$ and placing the m pebbles in $\mathbf{A}$ on an element in P (in the order they were chosen earlier) and the corresponding m pebbles in $\mathbf{B}$ on an element in $f(P)$ (in the same order). We denote the *k-pebble m-ary rank partition game over* GF_p played on structures $\mathbf{A}$ and $\mathbf{B}$ by $\mathcal{G}_p^{k;[m]}(\mathbf{A}, \mathbf{B})$.

Theorem 1. *Duplicator has a strategy for playing $\mathcal{G}_p^{k;[m]}(\mathbf{A}, \mathbf{B})$ forever if, and only if, $\mathbf{A} \equiv^{R_p^{k;[m]}} \mathbf{B}$.*

Write L^k to denote k-variable infinitary logic and C^k to denote the extension of L^k with counting quantifiers (see [7] for more details). The idea behind the rank partition game can also be used to give alternative characterisations of the relations $\equiv^{L^k}$ and $\equiv^{C^k}$. At each round in the *k-pebble cardinality partition game* on $\mathbf{A}$ and $\mathbf{B}$, the Spoiler picks up a pair of pebbles (a_i, b_i) for some i. Duplicator has to respond by choosing (a) partitions $\mathbf{P}$ of A and $\mathbf{Q}$ of B, with $|\mathbf{P}| = |\mathbf{Q}|$; and (b) a bijection $f : \mathbf{P} \to \mathbf{Q}$, such that for all parts $P \in \mathbf{P}$: $|P| = |f(P)|$. Spoiler then picks a part $P \in \mathbf{P}$, and places a_i on an element in $P \subseteq A$ and places b_i on an element in $f(P) \subseteq B$. This completes one round in the game. If Duplicator fails to produce the required partitions or the partial map defined by the pebbled elements is not a partial isomorphism, then Spoiler wins the game. Otherwise it can continue for another round. It can be shown that Duplicator has a strategy to play this game forever if, and only if, $\mathbf{A} \equiv^{C^k} \mathbf{B}$. Similarly, we can define the *k-pebble partition game* in exactly the same way as above, except we drop the requirement that the corresponding parts have to have the same size, i.e. Duplicator does not have to show that $|P| = |f(P)|$ for all $P \in \mathbf{P}$. It can be shown that Duplicator has a strategy to play this game forever if, and only if, $\mathbf{A} \equiv^{L^k} \mathbf{B}$. These two games can be seen as special cases of the generic rank partition game, which of course reflects the fact that the corresponding infinitary logics are both certain restrictions of infinitary rank logic.

4 Separation Results

The rank partition game can be used to delimit the expressive power of the rank logics restricted to a fixed arity and prime p. Specifically, using the game, we can show the following.

Theorem 2. *For all primes p and q where $q \equiv 1 \pmod{p}$, there is a property of finite graphs which is definable in $\mathrm{FOR}_q^{[2]}$ but not in $R_p^{\omega;[2]}$.*

The basic idea of the proof is as follows. For all primes p and q where $q \equiv 1 \pmod{p}$, and each $k \geq 2$, we construct a pair of non-isomorphic graphs $(\mathbf{A}_k^q, \mathbf{B}_k^q)$ which can be separated by a sentence of $\mathrm{FOR}_q^{[2]}$. We then show that Duplicator has a winning strategy in the game $\mathcal{G}_p^{k;[2]}(\mathbf{A}_k^q, \mathbf{B}_k^q)$, which shows that the classes of graphs $(\mathbf{A}_k^q)_{k\geq 2}$ and $(\mathbf{B}_k^q)_{k\geq 2}$ are not definable in $R_p^{\omega;[2]}$. The graphs $(\mathbf{A}_k^q, \mathbf{B}_k^q)$ are based on a construction of Torán [8]. This is essentially a way of encoding an arithmetic circuit modulo q into a given graph G. For instance, for $q = 2$ we get the graphs defined by Cai *et al.* [2] used to separate IFPC from PTIME. By starting with graphs G of large enough treewidth, we can ensure that for each k, Duplicator can hide the difference between $\mathbf{A}_k^q$ and $\mathbf{B}_k^q$ when playing the k-pebble rank partition game. Note that $q \equiv 1 \pmod{p}$ is required only for technical reasons in the proof; we believe the same method can be generalised for all distinct primes p and q. This gives us the following corollary, which partially resolves one of the open questions posed in [4].

Corollary 1. *For all primes p and q where $q \equiv 1 \pmod{p}$, $\mathrm{IFPR}_p^{[2]} \neq \mathrm{IFPR}_q^{[2]}$.*

References

1. A. Atserias, A. Bulatov, and A. Dawar. Affine systems of equations and counting infinitary logic. *Theor. Comput. Sci.*, 410:1666–1683, 2009.
2. J-Y. Cai, M. Fürer, and N. Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12(4):389–410, 1992.
3. A. Dawar. On the descriptive complexity of linear algebra. In *WoLLIC '08*, volume 5110 of *LNCS*, pages 17–25. Springer, 2008.
4. A. Dawar, M. Grohe, B. Holm, and B. Laubner. Logics with rank operators. In *Proc. 24th IEEE Symp. on Logic in Computer Science*, pages 113–122, 2009.
5. H. D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 1999.
6. N. Immerman. Relational queries computable in polynomial time. *Information and Control*, 68:86–104, 1986.
7. M. Otto. *Bounded Variable Logics and Counting — A Study in Finite Models*, volume 9 of *LNL*. Springer, 1997.
8. J. Torán. On the hardness of graph isomorphism. *SIAM Journal on Computing*, 33(5):1093–1108, 2004.
9. M. Y. Vardi. The complexity of relational query languages. In *Proc. of the 14th ACM Symp. on the Theory of Computing*, pages 137–146, 1982.