Approximations of Isomorphism and Logics with Linear-Algebraic Operators

Anuj Dawar
University of Cambridge, UK
anuj.dawar@cl.cam.ac.uk

Erich Grädel
RWTH Aachen University, Germany
graedel@logic.rwth-aachen.de

Wied Pakusa
RWTH Aachen University, Germany
pakusa@logic.rwth-aachen.de

Abstract

Invertible map equivalences are approximations of graph isomorphism that refine the well-known Weisfeiler-Leman method. They are parameterized by a number k and a set Q of primes. The intuition is that two equivalent graphs $G \equiv^{\mathrm{IM}}_{k,Q} H$ cannot be distinguished by means of partitioning the set of k-tuples in both graphs with respect to *any* linear-algebraic operator acting on vector spaces over fields of characteristic p, for any $p \in Q$. These equivalences have first appeared in the study of rank logic, but in fact they can be used to delimit the expressive power of any extension of fixed-point logic with linear-algebraic operators. We define $\mathrm{LA}^k(Q)$, an infinitary logic with k variables and all linear-algebraic operators over finite vector spaces of characteristic $p \in Q$ and show that $\equiv^{\mathrm{IM}}_{k,Q}$ is the natural notion of elementary equivalence for this logic. The logic $\mathrm{LA}^\omega(Q) = \bigcup_{k\in\omega} \mathrm{LA}^k(Q)$ is then a natural upper bound on the expressive power of any extension of fixed-point logics by means of Q-linear-algebraic operators.

By means of a new and much deeper algebraic analysis of a generalized variant, for any prime p, of the CFI-structures due to Cai, Fürer, and Immerman, we prove that, as long as Q is not the set of *all* primes, there is no k such that $\equiv^{\mathrm{IM}}_{k,Q}$ is the same as isomorphism. It follows that there are polynomial-time properties of graphs which are not definable in $\mathrm{LA}^\omega(Q)$, which implies that no extension of fixed-point logic with linear-algebraic operators can capture PTIME, unless it includes such operators for all prime characteristics. Our analysis requires substantial algebraic machinery, including a homogeneity property of CFI-structures and Maschke's Theorem, an important result from the representation theory of finite groups.

2012 ACM Subject Classification Theory of computation → Finite Model Theory

Keywords and phrases Finite Model Theory, Graph Isomorphism, Descriptive Complexity, Algebra

Digital Object Identifier 10.4230/LIPIcs.ICALP.2019.112

Category Track B: Automata, Logic, Semantics, and Theory of Programming

Related Version A full version of this paper is available at `https://arxiv.org/abs/1902.06648`.

1 Introduction

The graph isomorphism problem (or more generally, the structure isomorphism problem) is an important computational problem which is also very interesting from the point of view of complexity theory. It is not known to be in P nor known to be NP-complete. It is known to be solvable in quasi-polynomial time by Babai's algorithm [3].

An important theoretical approach to understanding the nature of the graph isomorphism problem is the Weisfeiler-Leman method. For each positive integer k, the k-dimensional Weisfeiler-Leman method (k-WL method for short) defines an equivalence relation $\equiv^k$ which

46th International Colloquium on Automata, Languages, and Programming (ICALP 2019).
Editors: Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi;
Article No. 112; pp. 112:1–112:14

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

over-approximates isomorphism in the sense that if $G \cong H$ for a pair of graphs G and H, then $G \equiv^k H$ for any k. The relations form a refining family in the sense that if $G \not\equiv^k H$ then $G \not\equiv^{k'} H$ for all $k' > k$. Thus, the equivalence relation gets finer with increasing k and approaches isomorphism in the limit. Moreover, if G and H are n-vertex graphs then $G \equiv^n H$ if, and only if, $G \cong H$. For each fixed k, the equivalence relation $\equiv^k$ is decidable in polynomial time, indeed in time $n^{O(k)}$. Thus, if there were a fixed k such that $\equiv^k$ were the same as isomorphism, we would have a polynomial-time algorithm for graph isomorphism. However, we know this is not the case. Cai, Fürer and Immerman [6] showed that there are pairs of non-isomorphic graphs G and H with $O(k)$ vertices such that $G \equiv^k H$. We call the construction of such graphs the CFI construction.

The Weisfeiler-Leman equivalences arise naturally in the study of graphs in many different guises. We have definitions based on combinatorics (such as Babai's original definition, see [6]); in logic as the equivalences induced by bounded variable fragments of first-order logic with counting; linear programming (see [2, 21]); and algebra (as in the original definition of Weisfeiler and Leman, extended to dimension k in [13]). The equivalences have proved to be of central importance in the area of descriptive complexity theory. In particular, they delimit the power of fixed-point logic with counting (FPC), an important logic in the study of symmetric polynomial-time computation (see [9]). On many important classes of structures, it turns out that there is a fixed k for which k-WL suffices to distinguish all non-isomorphic graphs. Most significantly, Grohe [20] has shown that for any proper minor-closed class $\mathcal{C}$ of graphs, there is a k such that $\equiv^k$ coincides with isomorphism on graphs in $\mathcal{C}$.

Despite its importance in the interplay of graph structure theory and logic, and its theoretical significance in understanding the graph isomorphism problem, the Weisfeiler-Leman method does not give the most efficient algorithms for solving the isomorphism problem. The CFI construction demonstrates that using the WL method to decide isomorphism would yield an algorithm of complexity $n^{\Omega(n)}$ which is asymptotically no better than trying all permutations and far removed from the quasi-polynomial time algorithms known. This has inspired the search for other structured families of equivalences (see for example [4, 16]). One particularly interesting such family are the *invertible-map equivalences* defined in [14]. This gives, for each k and each set Q of prime numbers, an equivalence relation $\equiv^{\mathrm{IM}}_{k,Q}$. The precise definition is given in Section 2 but the intuition is that if $G \equiv^{\mathrm{IM}}_{k,Q} H$, then G and H are not distinguishable by a refinement of k-tuples given by linear operators acting on vector spaces over fields of characteristic p, for any $p \in Q$. The reason for considering such equivalences stems from the realisation that the CFI-construction codes in graph form the problem of solving equations over $\mathbb{F}_2$ – the 2-element field (see [1]). It can then be shown that the family of equivalences $\equiv^{\mathrm{IM}}_{k,\{2\}}$ properly refine the Weisfeiler-Leman equivalences in that $G \equiv^{\mathrm{IM}}_{k',\{2\}} H$ for sufficiently large k' implies $G \equiv^k H$ for all k, yet $G \not\equiv^{\mathrm{IM}}_{3,\{2\}} H$ for the pairs G, H obtained in the CFI construction.

Furthermore, for any finite Q, the relation $\equiv^{\mathrm{IM}}_{k,Q}$ is decidable in time $n^{O(k)}$. We can also vary Q with n. For instance, we could let Q_s be the collection of all primes up to $s(n)$ for some growing function s. In this case $\equiv^{\mathrm{IM}}_{k,Q_s}$ is decidable in time $s(n)n^{O(k)}$. It is therefore an interesting question whether the family of equivalence relations is (like the Weisfeiler-Leman equivalences) infinitely refining. Do increasing values of k yield ever finer equivalence relations? The rôle of the parameter Q is also worth investigating. If there were a fixed polynomial s and constant k for which $\equiv^{\mathrm{IM}}_{k,Q_s}$ was the same as isomorphism, we would have a polynomial-time test for isomorphism. Even if we could prove this for k growing poly-logarithmically, and s quasi-polynomial, this would yield a new (and more systematic) quasi-polynomial algorithm for isomorphism. We have no reason to conjecture that either of these upper bounds holds, but they have not been ruled out.

One reason for the interest in the invertible-map equivalences is the connection with logic. In the long-running quest for a logic for PTIME (see [19]), an important direction is the study of extensions of fixed-point logic with rank operators (FPR) [12] or other algebraic operators (see [10]). The relations $\equiv^{\mathrm{IM}}_{k,Q}$ were introduced first as a tool to study the expressive power of FPR. It was shown in [14] that for every formula φ of FPR (as originally defined in [12]) there is a k and a finite Q such that the class of models of φ is closed under $\equiv^{\mathrm{IM}}_{k,Q}$. For the more powerful rank logic FPR^* defined in [18], we can show that for any formula φ, there is a k and a polynomial s such that φ is invariant under $\equiv^{\mathrm{IM}}_{k,Q_s}$. This implies, in particular that, if we could show that there is no fixed k such that $\equiv^{\mathrm{IM}}_{k,Q}$ is the same as isomorphism when Q is the set of all primes, we could, by means of padding and results from [14], separate FPR^* from PTIME. In short, any advance in understanding the structure of these equivalence relations is a significant step for resolving important questions.

The equivalence relations tell us about more than just rank logic. They can be used to delimit the expressive power of any extension of fixed-point logic with linear-algebraic operators. In this paper we introduce $\mathrm{LA}^k(Q)$, an infinitary logic with k variables and all linear-algebraic operators (which we define formally below) over finite vector spaces of characteristic $p \in Q$. This is the logic for which $\equiv^{\mathrm{IM}}_{k,Q}$ is the natural notion of elementary equivalence. Then, $\mathrm{LA}^\omega(Q) = \bigcup_{k\in\omega} \mathrm{LA}^k(Q)$ is a natural upper bound on the expressive power of any extension of fixed-point logics by means of Q-linear-algebraic operators.

Our main results can now be stated as follows. As long as Q is not the set of all primes, there is no k such that $\equiv^{\mathrm{IM}}_{k,Q}$ is the same as isomorphism. From this, it follows that there are classes of graphs which are not definable in $\mathrm{LA}^\omega(Q)$. Moreover, we can construct polynomial-time decidable such classes. This implies that any logic with linear-algebraic operators, unless it includes such operators for all prime characteristics, does not capture PTIME. Note, this does not separate FPR^* from PTIME, due to the restriction on Q, but it shows that if FPR^* is to capture PTIME, we need to use the set of all primes.

Establishing the result requires significant technical innovation. In particular, we develop novel algebraic machinery that has not previously been deployed in the field of finite model theory. As noted above, the CFI construction codes, in graph form, the problem of solving systems of linear equations over $\mathbb{F}_2$. We can give a similar construction that codes linear equations over $\mathbb{F}_p$ for any prime p. Such a construction was given in [22], where it was used to establish that the resulting non-isomorphic graphs were not distinguished by a variant of $\equiv^{\mathrm{IM}}_{k,\{q\}}$ for any $q \neq p$, where the matrix operations are restricted to a particularly simple form. A more refined analysis of the construction was used in [18] to separate the expressive power of FPR from that of FPR^*. To be precise, they showed that the formulas of FPR that do not use an operator with the prime p are no more expressive than formulas of FPC over these graphs. Our result uses the same graph construction but brings significant new algebraic machinery to its analysis.

We are able to show, in this paper, that, on graphs obtained by the CFI construction for $\mathbb{F}_p$, the distinguishing power of $\equiv^{\mathrm{IM}}_{k,Q}$, where $p \notin Q$, is no greater than $\equiv^{k'}$ for some fixed k'. Note that the graphs are definitely distinguished in $\equiv^{\mathrm{IM}}_{k,Q}$ when $p \in Q$. We establish the result by showing that on these graphs, the equivalence relation $\equiv^{\mathrm{IM}}_{k,\{q\}}$ is itself definable in FPC when $q \neq p$. This is done by implementing a matrix similarity test in FPC, based on the module isomorphism algorithm of Chistov et al. [8]. There are two key ingredients by which this yields an FPC definition. The first is that, on the graphs obtained in the construction, the equivalence relation $\equiv^k$ (now understood as an equivalence relation on k-tuples of vertices rather than on graphs) coincides with the partition into automorphism orbits, for sufficiently large but constantly bounded k. We say that the graphs are C^k-homogeneous for large

enough k. The second ingredient is that, because the automorphism groups of the graphs are Abelian p-groups, this partition induces a matrix algebra over $\mathbb{F}_q$, when $q \neq p$, which is semisimple and so admits a nice decomposition by Maschke's theorem. Maschke's theorem is a central result in the representation theory of finite groups, which states conditions under which a linear-algebraic representation of a finite group admits a decomposition into irreducible representations. It is a powerful tool and we hope that its use opens the door to further applications of representation theory in the context of finite model theory. Indeed, we see a major contribution of the present work as being the introduction of Maschke's theorem and related tools into the subject.

Much technical detail is omitted due to space reasons. Full proofs, and more detailed background on the algebra we use, can be found in the full version of the paper that is available [11].

2 The Invertible Map Equivalence and Linear-Algebraic Logics

The invertible map equivalence relation was introduced in [14, 22] as a family of approximations of isomorphism. It was shown that it is at least as fine an approximation as that induced by the infinitary logic with rank quantifiers, introduced in [12]. Dawar and Holm posed the question whether there is a logic which corresponds to the invertible map equivalences. Here we answer the question by showing that these equivalence relations are the right notions of elementary equivalence for an infinitary logic extended with *all* linear algebraic operations.

We begin by defining the equivalence relations $\equiv^{\mathrm{IM}}_{k,Q}$ for $k \in \mathbb{N}$ and Q a set of prime numbers. It is worth reviewing the definition of the counting-logic equivalence $\equiv^k$ first. This is not only an equivalence relation among finite structures, it also induces an equivalence on the set of A^k (the set k-tuples over A) inside (any) structure $\mathfrak{A}$ that yields an approximation to the partition of A^k into orbits.

On a structure $\mathfrak{A}$, the relation $\equiv^k$ can be obtained by an iterative refinement process. Suppose we are given a partition $\mathcal{P} = \{P_i\}_{i\in I}$ of A^k indexed by a set I. Now, we say that a pair of tuples $\bar{a}_1$ and $\bar{a}_2$ are *$\mathcal{P}$-similar* if they are in the same part of $\mathcal{P}$ *and* for each $i \in I$ and each $j \in [k]$ the sets $\{b \in A \mid \bar{a}_1[b/j] \in P_i\}$ and $\{b \in A \mid \bar{a}_2[b/j] \in P_i\}$ have the same number of elements. The equivalence relation $\equiv^k$ can then be characterised as the coarsest partition $\mathcal{P}$ of A^k that refines the partition into atomic types, such that any two tuples in the same part of $\mathcal{P}$ are $\mathcal{P}$-similar. This means that we can arrive at this partition by starting with the partition of A^k into atomic types and repeatedly refine it until we get a partition $\mathcal{P}$ for which the notions of $\mathcal{P}$-equivalence and $\mathcal{P}$-similarity are the same.

We now modify this in two ways to obtain the definition of $\equiv^{\mathrm{IM}}_{k,Q}$. First we define similarity not in terms of the substitution of a single element b into a tuple $\bar{a} \in A^k$ but of an ℓ-tuple $\bar{b} \in A^\ell$ for some $\ell < k$. So, for each injective function $\gamma : [\ell] \to [k]$, let $\bar{a}[\bar{b}/\gamma]$ denote the tuple in A^k obtained from $\bar{a}$ by simultaneously substituting b_i in position $\gamma(i)$ for all $i \in [\ell]$. If Γ denotes the set of all injective functions from $[\ell]$ to $[k]$, we say tuples $\bar{a}_1$ and $\bar{a}_2$ are *$\mathcal{P}$-similar* if they are in the same part of $\mathcal{P}$ *and* for each $\gamma \in \Gamma$ and each $i \in I$, the sets $\{\bar{b} \in A^\ell \mid \bar{a}_1[\bar{b}/\gamma] \in P_i\}$ and $\{\bar{b} \in A^\ell \mid \bar{a}_2[\bar{b}/\gamma] \in P_i\}$ have the same size. Taking the coarsest relation that is stable in this sense still gives us $\equiv^k$ (though see [15] for some nuances when comparing with the Weisfeiler-Leman equivalences).

For our purposes, we want a different notion of similarity. Assume that $\ell = 2m$ for some m. We can view any set $C \subseteq A^\ell$ as giving us an $A^m \times A^m$ 0-1 matrix, which we denote M. So the entry in row $\bar{b}_1 \in A^m$ and column $\bar{b}_2 \in A^m$ of M is 1 if, and only if, the ℓ-tuple $\bar{b}_1\bar{b}_2$ is in C. Hence, given, as before, a partition $\mathcal{P} = \{P_i\}_{i\in I}$ of A^k, and an

injective function $\gamma : [\ell] \to [k]$, each tuple $\bar{a}$ induces a partition of tuples $\bar{b}$ in A^ℓ according to which part P_i contains $\bar{a}[\bar{b}/\gamma]$. We think of this partition as a collection $(M_i^{\bar{a}})_{i\in I}$ of 0-1 matrices. For a prime number p, we say that two tuples $\bar{a}_1$ and $\bar{a}_2$ are *$\mathcal{P}$-p-m-similar* if they are in the same part of $\mathcal{P}$ *and* for every γ there is an invertible matrix $S \in \mathbb{F}_p^{A^m \times A^m}$ such that for each $i \in I$ we have $SM_i^{\bar{a}_1}S^{-1} = M_i^{\bar{a}_2}$. In other words, the sequences $(M_i^{\bar{a}_1})_{i\in I}$ and $(M_i^{\bar{a}_2})_{i\in I}$ are simultaneously similar, witnessed by S. We say the tuples are *$\mathcal{P}$-p-similar* if they are $\mathcal{P}$-p-m-similar for all $m \leq k/2$. The equivalence relation $\equiv^{\mathrm{IM}}_{k,p}$ is then the coarsest partition $\mathcal{P}$ that refines the partition into atomic types and such that any two tuples in the same part of $\mathcal{P}$ are $\mathcal{P}$-p-similar. Finally, for a set Q of prime numbers, $\bar{a}_1 \equiv^{\mathrm{IM}}_{k,Q} \bar{a}_2$ if, and only if, $\bar{a}_1 \equiv^{\mathrm{IM}}_{k,p} \bar{a}_2$ for each $p \in Q$. So, $\equiv^{\mathrm{IM}}_{k,Q}$ is the coarsest common refinement of the relations $(\equiv^{\mathrm{IM}}_{k,p})_{p\in Q}$.

Given a fixed set Q of primes with $|Q| = s$, it is possible to compute, for a structure $\mathfrak{A}$ with n elements, the partition of A^k into $\equiv^{\mathrm{IM}}_{k,Q}$ equivalence classes in time $sn^{O(k)}$. To see this, we note that the equivalence relation can be obtained by an iterated refinement process. First, let $\mathcal{P}_0$ be the partition of A^k into atomic types. Then, for each i, let $\mathcal{P}_{i+1}$ be the partition which places two tuples in the same class if, and only if, they are $\mathcal{P}_i$-p-similar for all $p \in Q$. This refinement process converges in at most n^k steps to the partition into $\equiv^{\mathrm{IM}}_{k,Q}$-equivalence classes. At each stage we compute, for each tuple $\bar{a} \in A^k$ and each injective function $\gamma\colon [2m] \to [k]$, the partition of A^{2m} into types, where $m = \lfloor k/2 \rfloor$. This suffices because $\mathcal{P}$-p-m-similarity implies $\mathcal{P}$-p-m'-similarity for all $m' < m$. Having computed the partition, we need to check for each pair of tuples and for each p in Q, whether the induced partitions are simultaneously similar. For this, we use the simultaneous matrix similarity test of Chistov et al. [8]. Since this runs in polynomial time, it follows that the whole procedure can be completed in time $sn^{O(k)}$.

Linear-Algebraic Logic. The study of logics with linear-algebraic operators over finite fields was initiated in [12], where FPR, the fixed-point logic with rank operators, was first introduced. As with fixed-point logics generally, the expressive power of FPR is naturally analysed by seeing it as a fragment of an infinitary logic, in this case with rank *quantifiers*. The notion of elementary equivalence that corresponds to this logic was given in terms of a game characterisation in [14], where the invertible map equivalences were also introduced. Here, we define, for any set Q of primes, an infinitary logic $\mathrm{LA}^\omega(Q)$ with quantifiers for *all* linear-algebraic operators over finite fields of characteristics in Q. This logic is not really intended for practical use. Instead it is designed to be strong enough so that inexpressibility results for $\mathrm{LA}^\omega(Q)$ carry over to any well-defined logic that extends first-order or fixed-point logic by any kind of linear-algebraic operators over Q.

We begin with a definition of linear-algebraic operators. Let $\mathbb{F}$ be a field and let $\mathcal{B}$ be a (non-empty, finite) set that serves as a supply of abstract basis elements. We consider the $\mathbb{F}$-vector space $\mathbb{F}^{\mathcal{B}}$. For each subset $K \subseteq \mathcal{B}$ we identify the vector space $\mathbb{F}^K$ with a subspace of $\mathbb{F}^{\mathcal{B}}$ in the natural way: since $\mathbb{F}^{\mathcal{B}} = \mathbb{F}^K \oplus \mathbb{F}^{\mathcal{B}\setminus K}$ we can (implicitly) set $\mathbb{F}^K = \mathbb{F}^K \oplus \{0\}$.

Let $m \geq 1$. Generally speaking, an *m-ary linear-algebraic operator* is just a function f that defines a linear-algebraic property $f(M_1, \ldots, M_m)$ of m-tuples of $\mathbb{F}$-linear transformations M_i on (subspaces of) $\mathbb{F}^{\mathcal{B}}$. To make things more precise, let $K_i, L_i \subseteq \mathcal{B}$, for $i \in [m]$, denote pairs of (non-empty) subsets of basis elements. We set $V_i = \mathbb{F}^{K_i}$ and $W_i = \mathbb{F}^{L_i}$. We consider m-tuples $(M_1, \ldots, M_m)$ consisting of $\mathbb{F}$-linear mappings $M_i\colon V_i \to W_i$ which are represented succinctly in terms of m-tuples $(M_1, \ldots, M_m)$ of $L_i \times K_i$-matrices with entries in $\mathbb{F}$. Then an m-ary linear-algebraic operator over $\mathbb{F}$ is a function f that takes such sequences $(M_1, \ldots, M_m)$ to some kind of linear-algebraic information $f(M_1, \ldots, M_m)$ about the sequence.

Now, to say that f outputs a "linear-algebraic information" means that the output of f is invariant under $\mathbb{F}$-vector space isomorphisms. Formally, let $\mathcal{C}$ be another (abstract) set of basis elements, where $|\mathcal{B}| = |\mathcal{C}|$, let $K'_i, L'_i \subseteq \mathcal{C}$ where $|K_i| = |K'_i|$ and $|L_i| = |L'_i|$ for $i \in [m]$, and let $(N_1, \dots, N_m)$ be a sequence of matrices $N_i \colon L'_i \times K'_i \to \mathbb{F}$, $i \in [m]$, analogously to the above. Moreover, let $V'_i = \mathbb{F}^{K'_i}$ and $W'_i = \mathbb{F}^{L'_i}$ for $i \in [m]$. Then we say that $(N_1, \dots, N_m)$ results from $(M_1, \dots, M_m)$ by means of an $\mathbb{F}$-vector space isomorphism if we can find an invertible $\mathbb{F}$-linear mapping $S \colon \mathbb{F}^{\mathcal{B}} \to \mathbb{F}^{\mathcal{C}}$ such that the following holds:

- For all $i \in [m]$, S maps each of the subspaces V_i and W_i in $\mathbb{F}^{\mathcal{B}}$ to the respective subspaces V'_i and W'_i in $\mathbb{F}^{\mathcal{C}}$. That is, if we represent S in terms of a $\mathcal{C} \times \mathcal{B}$-matrix with entries in $\mathbb{F}$, then we have that for each of the subblocks $K'_i \times K_i$, $i \in [m]$, the restriction $S \restriction_{(K'_i \times K_i)} \colon K'_i \times K_i \to \mathbb{F}$ of the matrix S to this block is invertible and we have that $S(a, b) = 0$ for all $a \in \mathcal{C} \setminus K'_i$ and $b \in K_i$ (and the analogous holds for all subblocks $L'_i \times L_i$ and the corresponding restrictions $S \restriction_{(L'_i \times L_i)} \colon L'_i \times L_i \to \mathbb{F}$ of S to the subblocks $L'_i \times L_i$).
- For each $i \in [m]$, the $\mathbb{F}$-vector space isomorphism S simultaneously transforms all linear operators $M_i \colon V_i \to W_i$ to the corresponding operators $N_i \colon V'_i \to W'_i$, that is for all $i \in [m]$ we have: $N_i \cdot S = S \cdot M_i$. Note that if we want to read this as a matrix equation, then we formally have to replace the matrix S by its restrictions to the subblocks $K'_i \times K_i$ and $L'_i \times L_i$ as we described above, that is $S \restriction_{(L'_i \times L_i)} \cdot M_i = N_i \cdot S \restriction_{(K'_i \times K_i)}$.

We require that a linear algebraic operator f outputs the same result for all pairs of matrix sequences $(M_1, \dots, M_m)$ and $(N_1, \dots, N_m)$ that are related via an $\mathbb{F}$-vector space isomorphism S (as above), that is $f(M_1, \dots, M_m) = f(N_1, \dots, N_m)$. This condition guarantees that f is not able to distinguish between isomorphic objects and here, in the realm of linear algebra, isomorphisms are $\mathbb{F}$-vector space isomorphisms. Besides this we do not put any kind of additional restrictions on f. For instance, f may not even be a computable function. Note that, though in introducing the function f, we considered a fixed set $\mathcal{B}$, really f defines, for any $\mathcal{B}$, a function on m-tuples of linear operators over subspaces of $\mathbb{F}^{\mathcal{B}}$. Without this, the notion of invariance would not make sense.

Now, we can associate with f a family of *Lindström quantifiers*. For simplicity, we restrict our attention to operators of a specific form without loss of generality. For an explanation of why no generality is lost, we refer the reader to the full paper [11]. Specifically, we assume that $K_i = L_i = \mathcal{B}$ for all i in the above definition, and we assume that the matrices are all 0-1 matrices. In other words, f is defined for a tuple of *square* 0-1 matrices all with the same index set.

Let τ_m denote a vocabulary with m distinct binary relations. Given an operator that defines such an f for each finite $\mathcal{B}$, for each $t \in \mathbb{N}$ we define a class of structures $\mathcal{K}^t_f$ in the vocabulary τ_m. We can think of an index set $\mathcal{B}$ with a collection $M_1, \dots, M_m$ of 0-1 $\mathcal{B} \times \mathcal{B}$ matrices as a τ_m-structure $(\mathcal{B}, M_1, \dots, M_m)$. The class $\mathcal{K}^t_f$ is then the collection of those τ_m-structures where $f(M_1, \dots, M_m) \geq t$. For each $\ell \geq 1$ we then have a quantifier $\mathcal{Q}^{t,\ell}_f$ such that if $I(\bar{x})$ is an $L[\sigma, \tau_m]$-interpretation of dimension ℓ, then $\mathcal{Q}^{t,\ell}_f I(\bar{x})$ is a formula true in a σ structure $\mathcal{A}$ if $I(\mathcal{A}) \in \mathcal{K}^t_f$.

The infinitary logic LA is defined as the closure of first-order logic under *infinitary* disjunction and conjunction, along with quantification $\mathcal{Q}^{t,\ell}_f$ for any linear algebraic operator f over any finite field. That is, if Φ is any set of formulas of LA, then $\bigvee \Phi$ and $\bigwedge \Phi$ are both formulas of LA. And, if f is an m-ary linear algebraic operator over a finite field, and $\Theta(\bar{x})$ is an ℓ-ary LA-interpretation of σ_m in τ, then $\mathcal{Q}^{t,\ell}_f \bar{x} \Theta$ is an LA τ-formula. We are interested in various fragments of the logic LA for which we introduce notation in the following definition.

▶ **Definition 1.** LA^k *is the collection of formulas of* LA *that contain at most* k *distinct variables, and* $\mathrm{LA}^\omega = \bigcup_{k\in\omega} \mathrm{LA}^k$. *For any set* Q *of primes, we write* $\mathrm{LA}(Q)$, $\mathrm{LA}^k(Q)$ *and* $\mathrm{LA}^\omega(Q)$ *to denote the restrictions of these logics to using only linear-algebraic operators over fields of characteristic* $p \in Q$.

If $\mathcal{L}$ *is any of the logics* LA, LA^ω, LA^k, $\mathrm{LA}(Q)$, $\mathrm{LA}^\omega(Q)$ *or* $\mathrm{LA}^k(Q)$, *and* $\ell \in \mathbb{N}$ *we write* ℓ-$\mathcal{L}$ *to denote the fragment of* $\mathcal{L}$ *where all algebraic quantifiers are* $\mathcal{Q}_f^{t,\ell}$ *for some* t *and* f. *In other words, interpretations are restricted to be of dimension* ℓ.

As pointed out above, the LA-logics are merely interesting from a theoretical point of view: in a precise sense they form a maximal extension of infinitary logics by linear-algebraic operators. On the other hand, we do not know of any (non-trivial) property that has a natural definition in the LA-logic, but not already in (infinitary) rank logic for instance (and it is open whether such example exists at all). Having introduced the linear-algebraic logic LA^ω and the invertible-map equivalences $\equiv^{\mathrm{IM}}_{k,Q}$, we are in a position to formulate their tight relationship.

▶ **Theorem 2.** *Let* $k \geq 2$ *be a positive integer and* Q *a set of prime numbers. For any finite structure* $\mathfrak{A}$ *and* $\bar{a}, \bar{b} \in A^k$, *the following are equivalent:*

1. $\bar{a} \equiv^{IM}_{k,Q} \bar{b}$; *and*
2. *for every formula* φ *of* $\mathrm{LA}^k(Q)$, $\mathfrak{A} \models \varphi[\bar{a}]$ *if, and only if,* $\mathfrak{A} \models \varphi[\bar{b}]$.

3 Cai-Fürer-Immerman Structures and Logic

In this section we describe a generalised variant of the *CFI-construction* due to Cai, Fürer, and Immerman [6]. It associates with each

- connected, 3-regular, and ordered (undirected) graph $G = (V, E, \leq)$,
- every prime field $\mathbb{F}_p$, $p \in \mathbb{P}$ (in this article, $\mathbb{P}$ denotes the set of all primes), and
- every vector $\lambda \in \mathbb{F}_p^V$

a structure we call the *CFI-structure* $\mathsf{CFI}[G; p; \lambda]$. Its signature is $\tau_{\mathsf{CFI}} = \{\preceq, R, C, I\}$ where R is a ternary relation symbol and where $\preceq, I, C$ are binary relation symbols. The universe A of $\mathsf{CFI}[G; p; \lambda]$ is $A = E \times \mathbb{F}_p$. The linear order $\leq$ on the vertex set V extends to a linear order on the edge set E (as the lexicographic order, for example). We use this linear order on E to define the following *total preorder* $\preceq$ on A: $(e, x) \preceq (f, y)$ if $e \leq f$. Note that $\preceq$ induces a linear order on the corresponding equivalence classes $e^p = e \times \mathbb{F}_p$. Clearly, each of these classes e^p is of size p. Since G is undirected every edge $e = (v, w) \in E$ comes with its corresponding *dual edge* $f = (w, v) \in E$. In what follows, we use the notation $e^{-1} = f$ to denote the dual of the edge $e \in E$. The relations I and C are defined as follows.

- The *cycle relation* C defines the cyclic structure of the additive group of $\mathbb{F}_p$ on each of the edge classes e^p. More precisely,

$$C = \bigcup_{e\in E} \{((e, x), (e, x + 1 \bmod p)) : x \in \mathbb{F}_p\}.$$

- The *inverse relation* I relates additive inverses for dual edges. Formally,

$$I = \bigcup_{e\in E} \{((e, x), (e^{-1}, -x) : x \in \mathbb{F}_p\}.$$

Note that while the cycle relation C defines a directed graph, the inverse relation I is symmetric. Furthermore, observe that the relations $\preceq, C$ and I are defined independently of the load vector λ and so only depend on the underlying graph G and the prime field $\mathbb{F}_p$. In

contrast, the *CFI-relation* $R = R^\lambda$ is defined using the load vector λ as follows. For each $v \in V$, we let $vE \subseteq V$ denote the set of neighbours of v in G, that is $E(v) = \{v\} \times vE \subseteq E$ is the set of edges outgoing from v. Since G is 3-regular we have that $|vE| = 3$ for each $v \in V$. For $v \in V$ let $vE = \{w_1, w_2, w_3\}$ where $w_1 < w_2 < w_3$. The *CFI-relation* $R^\lambda(v)$ at vertex v is defined as follows:

$$R^\lambda(v) = \{((w_1, x_1), (w_2, x_2), (w_3, x_3)) : x_1 + x_2 + x_3 = \lambda(v) \bmod p\}.$$

The full CFI-relation R^λ of $\mathsf{CFI}\,[G; p; \lambda]$ is given as $R^\lambda = \bigcup_{v \in V} R^\lambda(v)$.

▶ **Theorem 3.** *Two CFI-structures* $\mathsf{CFI}\,[G; p; \lambda], \mathsf{CFI}\,[G; p; \sigma]$ *over the same graph* G *are isomorphic, if, and only if,*

$$\sum \lambda := \sum_{v \in V} \lambda(v) = \sum_{v \in V} \sigma(v) =: \sum \sigma.$$

The CFI-construction unfolds its full power when it is based on a family of underlying graphs that is highly connected. A good choice is to take 3-regular *expander graphs* with $\mathcal{O}(n)$ vertices, as such graphs have a linear lower bound on the size of their separators (which means that we cannot disconnect the graphs into components of size $\leq n/2$ by removing fewer than $\Omega(n)$ vertices).

▶ **Theorem 4** (see e.g. Example 2.2 in [23])**.** *There exists a family of 3-regular, connected expander graphs* $\mathcal{F} = \{G_n : n \in \mathbb{N}\}$ *such that each graph* G_n*,* $n \in \mathbb{N}$*, has* $\mathcal{O}(n)$ *vertices.*

Of course, we can also assume that the graphs in $\mathcal{F}$ are *ordered* just by adding to each graph $G_n = (V_n, E_n) \in \mathcal{F}$ an arbitrary linear order on V_n. From this family $\mathcal{F}$ of 3-regular, connected, ordered expander graphs G_n with $\mathcal{O}(n)$ many vertices we construct, for every $p \in \mathbb{P}$, the CFI-class $\mathsf{CFI}\,[\mathcal{F}; p]$ consisting of all CFI-structures over graphs from $\mathcal{F}$ that is

$$\mathsf{CFI}\,[\mathcal{F}; p] = \bigcup_{n,\lambda} \mathsf{CFI}\,[G_n; p; \lambda].$$

The *CFI-problem* (over $\mathcal{F}$ and $p \in \mathbb{P}$) is to decide, given a structure $\mathsf{CFI}\,[G; p; \lambda] \in \mathsf{CFI}\,[\mathcal{F}; p]$ whether $\sum \lambda = 0$. For the original form of the CFI-construction, it was shown in [6] that this problem is undefinable in counting logic with sublinearly many variables. Also the generalization to more powerful variants, and in particular to our class $\mathsf{CFI}\,[\mathcal{F}; p]$ is well-known.

▶ **Theorem 5.** *For any two structures* $\mathsf{CFI}\,[G_n; p; \lambda], \mathsf{CFI}\,[G_n; p; \sigma] \in \mathsf{CFI}\,[\mathcal{F}; p]$ *we have*

$$\mathsf{CFI}\,[G_n; p; \lambda] \equiv^{\Omega(n)} \mathsf{CFI}\,[G_n; p; \sigma].$$

Thus, from the perspective of counting logic (with $\Omega(n)$ many variables) CFI-structures over the same underlying graph G_n look the same although, for load vectors λ and σ with $\sum \lambda \neq \sum \sigma$, we know that $\mathsf{CFI}\,[G_n; p; \lambda]$ and $\mathsf{CFI}\,[G_n; p; \sigma]$ are not isomorphic.

▶ **Definition 6.** *Let* $\ell \geq 1$*. We say that a structure* $\mathfrak{A}$ *with automorphism group* Γ *is* ℓ-homogeneous *if for all* $k \geq 1$ *and all* k*-tuples* $\bar{a}, \bar{b} \in A^k$ *we have that*

$$(\mathfrak{A}, \bar{a}) \equiv^{\ell \cdot k} (\mathfrak{A}, \bar{b}) \text{ if, and only if, } \Gamma(\bar{a}) = \Gamma(\bar{b}).$$

In other words, the equivalence relation $\equiv^{\ell \cdot k}$ *refines* k*-tuples in* $\mathfrak{A}$ *up to orbits. Moreover, we say that a class* $\mathcal{K}$ *of structures is* homogeneous *if for some constant* $\ell \geq 1$ *each structure* $\mathfrak{A} \in \mathcal{K}$ *is* ℓ*-homogeneous.*

▶ **Theorem 7.** *For every prime p, the class $\mathsf{CFI}[\mathcal{F};p]$ is homogeneous.*

This theorem has been established in [17]. Homogeneity of CFI-structures is very useful because it implies that counting logic (indeed, FPC) can order k-tuples up to orbits. There are counting-type formulae $\mathrm{CT}_{\ell\cdot k}(\bar{x},\bar{y}) \in \mathrm{FPC}$ (see [24]) that define a linear preorder on k-tuples which distinguishes between all pairs of k-tuples in different orbits, and these formulae use only $\mathcal{O}(\ell\cdot k)$ many variables. One key consequence of homogeneity is that on the class of CFI structures, the relations $\equiv^k$ and $\equiv^{\mathrm{IM}}_{k,Q}$ coincide for k above some constant threshold. Indeed, $\equiv^{\mathrm{IM}}_{k,Q}$ is always at least as fine as $\equiv^k$ and no finer than the equivalence given by the partition into automorphism orbits. When the former and the latter are the same, $\equiv^{\mathrm{IM}}_{k,Q}$ must be the same. In particular, this means that the counting-type formulas $\mathrm{CT}_{\ell\cdot k}(\bar{x},\bar{y})$ define a pre-order on the $\equiv^{\mathrm{IM}}_{k,Q}$ equivalence classes.

4 Indistinguishability of CFI-structures using Linear-Algebraic Operators

In this section, we state our main technical result, that CFI-structures over a prime field $\mathbb{F}_p$ cannot be distinguished by means of *any* linear-algebraic operator over a field $\mathbb{F}$ with $\mathrm{char}(\mathbb{F}) \neq p$ if we apply such linear-algebraic operators to $\mathrm{C}^{\Omega(n)}$-definable matrices. For what follows, recall that we consider CFI-structures over a fixed class of expander graphs $\mathcal{F} = \{G_n : n \in \mathbb{N}\}$ where each graph G_n has $\mathcal{O}(n)$ vertices and is ordered, connected, and three-regular.

The exact formulation of the results requires some background from associative algebra. This can be found in the monograph [26], for example. The definitions and results we use are also summarized in the full version of this paper [11, Sec. 5].

The partition of 2ℓ-tuples in a structure $\mathfrak{A}$ into $\equiv^k$-classes (when $k \geq 3\ell$) induces a *coherent configuration* in the sense of [7, Chap. 3]. This implies, in particular, that if we think of this partition as a collection $M_1,\dots,M_s$ of 0-1 matrices with rows and columns indexed by A^ℓ then, for any field $\mathbb{F}$, they form the basis of an $\mathbb{F}$-algebra. That is to say, they are the basis of an $\mathbb{F}$-vector space that is also closed under matrix multiplication. We denote this algebra $\mathsf{Alg}[\mathfrak{A};\ell;\mathrm{C}^k;\mathbb{F}]$ and the *ordered* collection of matrices that forms its basis $\mathrm{Basis}[\mathfrak{A};\ell;\mathrm{C}^k]$. Note that the latter does not depend on the choice of $\mathbb{F}$.

We say that two structures $\mathfrak{A}$ and $\mathfrak{B}$ are $(\mathbb{F};\ell;\mathrm{C}^k)$-isomorphic if $\mathfrak{A} \equiv^k \mathfrak{B}$ and, if $\mathcal{M} = \mathrm{Basis}[\mathfrak{A};\ell;\mathrm{C}^k] = (M_1,\dots,M_s)$ and $\mathcal{N} = \mathrm{Basis}[\mathfrak{B};\ell;\mathrm{C}^k] = (N_1,\dots,N_s)$, then there is an invertible map $S : \mathbb{F}^{A^\ell} \to \mathbb{F}^{B^\ell}$ such that for each $i \in [s]$, $SM_iS^{-1} = N_i$. In short, the two sequences of matrices $\mathrm{Basis}[\mathfrak{A};\ell;\mathrm{C}^k]$ and $\mathrm{Basis}[\mathfrak{B};\ell;\mathrm{C}^k]$ are *simultaneously similar* as witnessed by S. In particular, the $\mathbb{F}$-algebras $\mathsf{Alg}[\mathfrak{A};\ell;\mathrm{C}^k;\mathbb{F}]$ and $\mathsf{Alg}[\mathfrak{B};\ell;\mathrm{C}^k;\mathbb{F}]$ are isomorphic.

For CFI-structures, $\mathfrak{A} = \mathsf{CFI}[G_n;p;\lambda]$ and $\mathfrak{B} = \mathsf{CFI}[G_n;p;\sigma]$, by the homogeneity property, we know that the partition into $\equiv^k$-classes is the same as the partition into automorphism orbits. This allows us to show that when such structures are $(\mathbb{F};\ell;\mathrm{C}^k)$-isomorphic, they cannot be distinguished by any $\mathbb{F}$-linear-algebraic operators. Hence, the key technical theorem we prove is the following.

▶ **Theorem 8.** *There is $\epsilon > 0$ s.t. for large enough $n > 0$ the following holds. Let $\mathfrak{A} = \mathsf{CFI}[G_n;p;\lambda]$ and $\mathfrak{B} = \mathsf{CFI}[G_n;p;\sigma]$ denote CFI-structures over G_n and let $\mathbb{F}$ be a field such that $\mathrm{char}(\mathbb{F}) \neq p$. Then $\mathfrak{A}$ and $\mathfrak{B}$ are $(\mathbb{F};\ell;\mathrm{C}^k)$-isomorphic where $\ell = \lfloor \epsilon n \rfloor$ and $k = 3\ell$.*

Theorem 8 is a consequence of Theorem 9 where we show that, for the above scenario, the simultaneous similarity of the counting-logic bases is definable in counting logic. To state Theorem 9, we introduce some terminology. Consider two sequences of matrices

$\mathcal{M} = (M_i)_{i\in[s]}$ and $\mathcal{N} = (N_i)_{i\in[s]}$, where each M_i is an $I \times I$ matrix over $\mathbb{F}$ and each N_i is a $J \times J$ matrix over $\mathbb{F}$. Here I and J are two arbitrary index sets of the same size. We define the set $\mathrm{H}_{\mathcal{M},\mathcal{N}}$ of $I \times J$-matrices X over $\mathbb{F}$ which satisfy $M_i X = X N_i$ for all $i \in [s]$. Note that the two sequences $\mathcal{M} = (M_i)_{i\in[s]}$ and $\mathcal{N} = (N_i)_{i\in[s]}$ are simultaneously similar if, and only if, $\mathrm{H}_{\mathcal{M},\mathcal{N}}$ contains an *invertible* matrix.

Note that $\mathrm{H}_{\mathcal{M},\mathcal{N}}$ is an $\mathbb{F}$-vector space. Next, consider the set $\mathrm{C}_{\mathcal{M}}$ of $I \times I$-square matrices Z over $\mathbb{F}$ such that $M_k Z = Z M_k$ for all $k \in K$. The set $\mathrm{C}_{\mathcal{M}}$ is called the *centraliser* of the matrix family $\mathcal{M}$. It is easy to verify that $\mathrm{C}_{\mathcal{M}}$ forms an $\mathbb{F}$-algebra. Moreover, by considering matrix multiplication (from the left) by elements from $\mathrm{C}_{\mathcal{M}}$, the $\mathbb{F}$-vector space $\mathrm{H}_{\mathcal{M},\mathcal{N}}$ turns into a $\mathrm{C}_{\mathcal{M}}$-module. With this, we can state the technical result.

▶ **Theorem 9.** *Let $t \geq 3$ be a constant such that all CFI-structures in $\mathsf{CFI}[\mathcal{F}; p]$ are t-homogeneous for all $p \in \mathbb{P}$. Then there exists a constant $c \geq 1$ such that the following holds. Let $\ell \geq 1$ and let $k \geq t\ell$. Then for each $p \in \mathbb{P}$ there exists a C^{ck}-sentence φ such that for all pairs of CFI-structures $\mathfrak{A} = \mathsf{CFI}[G_n; p; \lambda]$ and $\mathfrak{B} = \mathsf{CFI}[G_n; p; \sigma]$ over the same underlying graph $G_n \in \mathcal{F}$ we have that $(\mathfrak{A}, \mathfrak{B}) \models \varphi$ if, and only if, over every field $\mathbb{F}$ with $\mathrm{char}(\mathbb{F}) \neq p$, the $C_{\mathcal{M}}$-module $H_{\mathcal{M},\mathcal{N}}$ contains an invertible matrix $S \in H_{\mathcal{M},\mathcal{N}}$ where $\mathcal{M} = \mathit{Basis}[\mathfrak{A}, \ell, k]$ and $\mathcal{N} = \mathit{Basis}[\mathfrak{B}, \ell, k]$.*

We can derive Theorem 8 from Theorem 9 as follows. First of all, let $c \geq 1$ and $t \geq 3$ be the constants according to Theorem 9. Let $p \in \mathbb{P}$. Then, by Theorem 5, we can find $\delta > 0$ such that for all large enough $n > 1$ we have $\mathfrak{A} \equiv^{\lfloor \delta n \rfloor} \mathfrak{B}$ where $\mathfrak{A} = \mathsf{CFI}[G_n; p; \lambda]$ and $\mathfrak{B} = \mathsf{CFI}[G_n; p; \sigma]$ are two CFI-structures over $\mathbb{F}_p$ and the same underlying expander graph $G_n \in \mathcal{F}$ with $\mathcal{O}(n)$ many vertices. Let $\epsilon = \frac{1}{tc}\delta$. Then $(\mathfrak{A}, \mathfrak{A}) \equiv^{\lfloor t c \epsilon n \rfloor} (\mathfrak{A}, \mathfrak{B})$. Let $\mathbb{F}$ be a field such that $\mathrm{char}(\mathbb{F}) \neq p$. Let $\ell = \lfloor \epsilon n \rfloor$ and $k = \lfloor t \epsilon n \rfloor$. We consider $\mathcal{M} = \mathrm{Basis}[\mathfrak{A}, \ell; \mathrm{C}^k]$ and $\mathcal{N} = \mathrm{Basis}[\mathfrak{B}, \ell; \mathrm{C}^k]$. Since the formula φ according to Theorem 9 contains at most $ck = c \cdot \lfloor t \epsilon n \rfloor \leq \lfloor \delta n \rfloor$ many variables, this formula cannot distinguish between the ordered pairs of CFI-structures $(\mathfrak{A}, \mathfrak{A})$ and $(\mathfrak{A}, \mathfrak{B})$. On the other hand, by its properties stated in Theorem 9, φ would need to distinguish between $(\mathfrak{A}, \mathfrak{A})$ and $(\mathfrak{A}, \mathfrak{B})$ if no invertible matrix $S \in \mathrm{H}_{\mathcal{M},\mathcal{N}}$ would exist. Indeed, note that the $\mathrm{C}_{\mathcal{M}}$-module $\mathrm{H}_{\mathcal{M},\mathcal{M}}$ contains an invertible matrix $S \in \mathrm{H}_{\mathcal{M},\mathcal{M}}$ over every field $\mathbb{F}$ for trivial reasons; for instance it contains the permutation matrix that corresponds to the identity automorphism of $\mathfrak{A}$. Hence, we can conclude that $\mathrm{H}_{\mathcal{M},\mathcal{N}}$ contains an invertible matrix which shows that $\mathfrak{A}$ and $\mathfrak{B}$ are $(\mathbb{F}; \ell; \mathrm{C}^k)$-isomorphic, and thus Theorem 8 follows, because $(\mathbb{F}; \ell; \mathrm{C}^k)$-isomorphic structures are also $(\mathbb{F}; \ell; \mathrm{C}^{3\ell})$-isomorphic since $k \geq 3\ell$.

We now outline a proof strategy for Theorem 9. The full proof is rather long, and is presented in detail in the full paper [11]. First, we fix a *prime field* $\mathbb{F}$ with $\mathrm{char}(\mathbb{F}) \neq p$. We construct a sentence $\varphi_{\mathbb{F}} \in \mathrm{C}^\omega$, with at most $c \cdot k$ many variables, which holds in the ordered pair $(\mathfrak{A}, \mathfrak{B})$ of CFI-structures $\mathfrak{A}$ and $\mathfrak{B}$ if, and only if, $\mathrm{H}_{\mathcal{M},\mathcal{N}}$ (considered as a $\mathrm{C}_{\mathcal{M}}$-module over the $\mathbb{F}$-algebra $\mathrm{C}_{\mathcal{M}}$) contains an invertible matrix S. The desired sentence φ according to Theorem 9 is then the conjunction over all sentences $\varphi_{\mathbb{F}}$ for *prime fields* $\mathbb{F}$ with $\mathrm{char}(\mathbb{F}) \neq p$.

Step (I). As a first step we show that it suffices to restrict our considerations to *prime fields*. This is because the matrix families $\mathcal{M}$ and $\mathcal{N}$ we are interested in only contain 0-1 matrices. Such families are simultaneously similar over a field $\mathbb{F}$ if, and only if, they are simultaneously similar over the prime subfield of $\mathbb{F}$. The restriction is important because we need to use the result (originally proved in [18]) about defining solutions to system of linear equations. The result is that if a system of linear equations over a prime field $\mathbb{F}$ is represented by a homogeneous structure $\mathfrak{A}$ with Abelian automorphism group, and such that

the order of the automorphism group of $\mathfrak{A}$ is co-prime with char($\mathbb{F}$), then a solution to the system can be defined by a formula of counting logic. This is because we can show that in this case, there must exist a solution that is symmetric, i.e. invariant under the action of the automorphism group. In fact, we need here a stronger result saying that not only single solutions, but whole solution spaces are definable in counting logic. This was recently shown in [17] and our full paper [11] also contains a proof sketch.

Now, the CFI structures in $\mathsf{CFI}[\mathcal{F}; p]$ have Abelian automorphism groups of order a power of p, so systems of linear equations suitably defined from them will have the required property. Our aim is to reduce the problem of deciding whether $\mathrm{H}_{\mathcal{M},\mathcal{N}}$ contains an invertible matrix to solving a system of linear equations over $\mathbb{F}$. The condition $M_i X = X N_i$ for all i easily yields a system of such equations with unknowns for the enrtries of X. The question is how to enforce that X is invertible.

Step (II). To carry out the reduction, we use the result from [8] that if $\mathrm{H}_{\mathcal{M},\mathcal{N}}$ contains an invertible matrix, then it is *cyclic* as a $\mathrm{C}_{\mathcal{M}}$-module. This means that $\mathrm{H}_{\mathcal{M},\mathcal{N}}$ is generated by a single element: there is a matrix $X \in \mathrm{H}_{\mathcal{M},\mathcal{N}}$ such that $\mathrm{C}_{\mathcal{M}} X = \{ZX : Z \in \mathrm{C}_{\mathcal{M}}\} = \mathrm{H}_{\mathcal{M},\mathcal{N}}$. This gives a necessary but not sufficient condition for the existence of an invertible matrix in $\mathrm{H}_{\mathcal{M},\mathcal{N}}$. To obtain a necessary and sufficient condition, we use the particular structure of the matrix families $\mathcal{M}$ and $\mathcal{N}$ that follow from the fact that they are generated by the $\equiv^k$-equivalence classes. Roughly speaking, the equivalence relation $\equiv^k$ partitions the row-column index sets of these matrices into classes, and we can think of the matrices as linear combinations of "small" matrices, which are over these individual blocks. This means that the invertible S we are looking for can also be decomposed into the sum of smaller block matrices. Our families $\mathcal{M}$ and $\mathcal{N}$ have the property that they are *locally simultaneously similar*, i.e. we can find similarity transformations for each small block. With this, it becomes possible to prove that the cyclicity of $\mathrm{H}_{\mathcal{M},\mathcal{N}}$ is both necessary and sufficient for the existence of an invertible matrix. It also means that we can restrict ourselves to a certain substructure of this module. To be precise, we let $\mathrm{C}^{\mathrm{D}}_{\mathcal{M}}$ be the subalgebra of $\mathrm{C}_{\mathcal{M}}$ consisting of those matrices that are zero outside the relevant blocks. Then, $\mathrm{H}^{\mathrm{D}}_{\mathcal{M},\mathcal{N}}$ is similarly the collection of matrices in $\mathrm{H}_{\mathcal{M},\mathcal{N}}$ that are zero outside the relevant blocks and this can be seen as a $\mathrm{C}^{\mathrm{D}}_{\mathcal{M}}$-module. For the particular matrix families $\mathcal{M} = \mathrm{Basis}[\mathfrak{A}, \ell, k]$ and $\mathcal{N} = \mathrm{Basis}[\mathfrak{B}, \ell, k]$, we are able to show that they are simultaneously similar if, and only if, $\mathrm{H}^{\mathrm{D}}_{\mathcal{M},\mathcal{N}}$ is cyclic as a $\mathrm{C}^{\mathrm{D}}_{\mathcal{M}}$-module.

As a consequence, to check whether $\mathrm{H}_{\mathcal{M},\mathcal{N}}$ contains an invertible matrix, it suffices to check whether the $\mathrm{C}^{\mathrm{D}}_{\mathcal{M}}$-module $\mathrm{H}^{\mathrm{D}}_{\mathcal{M},\mathcal{N}}$ is cyclic.

Step (III). The third step is the core of our whole argument. We combine results on the FPC-definability of the automorphism groups and orbits of CFI-structures with Maschke's Theorem, an important result from the representation theory of finite groups, to show that the $\mathbb{F}$-algebra $\mathrm{C}^{\mathrm{D}}_{\mathcal{M}}$ is *semisimple*.

Recall that for a finite group G and a field $\mathbb{F}$, the *group algebra* $\mathbb{F}[G]$ is the $\mathbb{F}$-algebra whose elements are formal sums of the form $\sum_{g \in G} r_g g$ with coefficients $r_g \in \mathbb{F}$. Addition and scalar multiplication are defined component-wise and multiplication is defined by convolution on the group elements. Maschke's theorem tells us that $\mathbb{F}[G]$ is semisimple if, and only if, char($\mathbb{F}$) does not divide the order of G.

For an algebra A, an A-module M is called *simple* if every submodule of M is trivial (either 0 or M itself) and *semisimple* if it is the direct sum of simple modules. From the semi-simplicity of $\mathrm{C}^{\mathrm{D}}_{\mathcal{M}}$ we are able to show that the $\mathrm{C}^{\mathrm{D}}_{\mathcal{M}}$-module $\mathrm{H}^{\mathrm{D}}_{\mathcal{M},\mathcal{N}}$ is semisimple.

Step (IV). The key property of semisimple modules is that they have an essentially canonical decomposition as the sum of simple modules. So, $\mathrm{H}^{\mathrm{D}}_{\mathcal{M},\mathcal{N}}$ can be decomposed as the sum of simple modules, and the isomorphism types of modules that occur in that decomposition and their respective multiplicities completely determine the isomorphism type of $\mathrm{H}^{\mathrm{D}}_{\mathcal{M},\mathcal{N}}$. Moreover, we show that we can define generating sets for the respective submodules in counting logic by using at most $c \cdot k$ many variables. This is because, essentially, these generating sets can be obtained as the solution sets of a system of linear equations, and we are able to use the results mentioned in Step (I) above.

Step (V). Finally, we construct the formula $\varphi_{\mathbb{F}}$. By (II), the formula $\varphi_{\mathbb{F}}$ needs to verify that the *semisimple* $\mathrm{C}^{\mathrm{D}}_{\mathcal{M}}$-module $\mathrm{H}^{\mathrm{D}}_{\mathcal{M},\mathcal{N}}$ is cyclic. We approach this problem by expressing a more general query, namely we determine the full isomorphism type of the module $\mathrm{H}^{\mathrm{D}}_{\mathcal{M},\mathcal{N}}$ by means of a formula of counting logic. First of all, we start by determining the isomorphism types of all simple subalgebras of $\mathrm{C}^{\mathrm{D}}_{\mathcal{M}}$. This we can easily do in counting logic because $\mathrm{C}^{\mathrm{D}}_{\mathcal{M}}$ has an (FPC-definable) ordered basis. This implies that the isomorphism type of $\mathrm{H}^{\mathrm{D}}_{\mathcal{M},\mathcal{N}}$ is (uniquely) determined by the multiplicities of the simple subalgebras of $\mathrm{C}^{\mathrm{D}}_{\mathcal{M}}$ as they occur in a decomposition of $\mathrm{H}^{\mathrm{D}}_{\mathcal{M},\mathcal{N}}$ into a direct sum of simple submodules. By using our decomposition from Step (IV), we can easily determine those multiplicities componentwise, since we can linearly order (again in an FPC-definable way) each of the "small" submodules that occur in the decomposition of $\mathrm{H}^{\mathrm{D}}_{\mathcal{M},\mathcal{N}}$. In this way we can determine the multiplicities for each individual component which add up to the total multiplicities for the whole module $\mathrm{H}^{\mathrm{D}}_{\mathcal{M},\mathcal{N}}$. Since the isomorphism type determines the cyclicity of the module, we can obtain our desired formula $\varphi_{\mathbb{F}}$ by selecting modules with appropriate isomorphism types.

5 Main results

In this section we spell out the consequences of the main technical result, Theorem 8, for approximations of isomorphism and for logics with linear-algebraic operators.

With regard to the relations $\equiv^{\mathrm{IM}}_{k,Q}$ as approximations of isomorphism, it follows immediately that as long as $Q \neq \mathbb{P}$, i.e. Q is not the set of all primes, there is no k for which $\equiv^{\mathrm{IM}}_{k,Q}$ coincides with isomorphism on all structures.

▶ **Corollary 10.** *If $Q \neq \mathbb{P}$, there is no fixed k such that $\equiv^{IM}_{k,Q}$ coincides with isomorphism on all structures.*

Proof. Fix a prime $p \notin Q$. Then, for each k, we have, by Theorem 8 a pair of structures $\mathfrak{A} = \mathsf{CFI}\,[G_n; p; \lambda]$ and $\mathfrak{B} = \mathsf{CFI}\,[G_n; p; \sigma]$ that are $(\mathbb{F}_q; \ell; \mathrm{C}^k)$-isomorphic, for all $q \neq p$, though $\sum \lambda \neq \sum \sigma$. It follows that $\mathfrak{A} \equiv^{\mathrm{IM}}_{k,Q} \mathfrak{B}$, but $\mathfrak{A} \ncong \mathfrak{B}$, by Theorem 3. ◀

The consequences for the expressive power of the logic LA^{ω} are also immediate.

▶ **Corollary 11.** *If $Q \neq \mathbb{P}$, there is a class of structures that is not definable in $\mathrm{LA}^{\omega}(Q)$.*

Proof. Fix a prime $p \notin Q$ and consider the class $\mathcal{C}$ of structures of the form $\mathsf{CFI}\,[G_n; p; \lambda]$ where $\sum \lambda = 0$. This is an isomorphism-closed class of structures by Theorem 3. Suppose it were defined by a sentence φ of $\mathrm{LA}^{\omega}(Q)$. Let ℓ be the maximum dimension of an interpretation used with any quantifier in φ and choose k such that $k \geq 3\ell$ and k is greater than the number of variables in φ. Then, by Theorem 8, we have a structure $\mathfrak{A} = \mathsf{CFI}\,[G_n; p; \lambda] \in \mathcal{C}$ which is $(\mathbb{F}_q; \ell; \mathrm{C}^k)$-isomorphic to every structure $\mathsf{CFI}\,[G_n; p; \sigma]$. Letting $\mathfrak{B}$ be such a structure where $\sigma \neq 0$, we have that $\mathfrak{B} \models \varphi$, contradicting the assumption that φ defines $\mathcal{C}$. ◀

It should be noted that the class of structures $\mathcal{C}$ defined in the proof of Corollary 11 is decidable in polynomial time. This is because the class can be decided by solving systems of linear equations, for example by Gaussian elimination. Thus, we know that there exists a PTIME property that $\mathrm{LA}^{\omega}(Q)$ cannot express as long as $Q \neq \mathbb{P}$. Since this logic subsumes any extension of fixed-point logic with Q-linear algebraic operators, we also have the following conclusion.

▶ **Corollary 12.** *If $Q \neq \mathbb{P}$, no extension of fixed-point logic with Q-linear algebraic operators captures PTIME.*

We can say more. The class $\mathcal{C}$ is not just decidable in PTIME, but also definable in *choiceless polynomial time* (CPT) (see [25]). We do not define the class CPT here but details may be found in [5]. Thus, the following corollary is immediate.

▶ **Corollary 13.** *If $Q \neq \mathbb{P}$, no extension of fixed-point logic with Q-linear algebraic operators captures CPT.*

On the other hand it remains an intriguing open question whether CPT captures all of rank logic, for example.

References

1 A. Atserias, A. Bulatov, and A. Dawar. Affine Systems of Equations and Counting Infinitary Logic. *Theoretical Computer Science*, 410:1666–1683, 2009.

2 A. Atserias and E. N. Maneva. Sherali-Adams relaxations and indistinguishability in counting logics. *SIAM J. Comput.*, 42:112–137, 2013.

3 L. Babai. Graph Isomorphism in Quasipolynomial Time [extended abstract]. In *Proc. 48th Annual ACM SIGACT Symp. Theory of Computing, STOC*, pages 684–697, 2016.

4 A. Barghi and I Ponomarenko. Non-Isomorphic Graphs with Cospectral Symmetric Powers. *Electr. J. Comb.*, 16(1), 2009.

5 A. Blass, Y. Gurevich, and S. Shelah. On Polynomial Time Computation Over Unordered Structures. *Journal of Symbolic Logic*, 67(3):1093–1125, 2002.

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

7 P.J. Cameron. *Permutation Groups*. London Mathematical Society Student Texts. Cambridge University Press, 1999.

8 A. Chistov, G. Ivanyos, and M. Karpinski. Polynomial Time Algorithms for Modules over Finite Dimensional Algebras. In *Proceedings of ISSAC '97*, pages 68–74. ACM, 1997.

9 A. Dawar. The nature and power of fixed-point logic with counting. *ACM SIGLOG News*, 2(1):8–21, 2015.

10 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, Special Issue dedicated to CSL 2012*, 2013. URL: `http://www.lmcs-online.org/ojs/viewarticle.php?id=1325`.

11 A. Dawar, E. Grädel, and W. Pakusa. Approximations of Isomorphism and Logics with Linear-Algebraic Operators. arXiv abs/1902.06648. `arXiv:1902.06648`.

12 A. Dawar, M. Grohe, B. Holm, and B. Laubner. Logics with Rank Operators. In *Proceedings of LICS 2009*, pages 113–122, 2009.

13 A. Dawar and B. Holm. Tractable Approximations of Graph Isomorphism. forthcoming.

14 A. Dawar and B. Holm. Pebble Games with Algebraic Rules. *Fundam. Inform.*, 150(3-4):281–316, 2017.

15 A. Dawar and D. Vagnozzi. Generalizations of k-Weisfeiler-Leman Partitions and Related Graph Invariants. forthcoming.

16 H. Derksen. The Graph Isomorphism Problem and approximate categories. *J. Symb. Comput.*, 59:81–112, 2013. doi:10.1016/j.jsc.2013.06.002.

17 E. Grädel, M. Grohe, B. Pago, and W. Pakusa. A Finite-Model-Theoretic View on Propositional Proof Complexity. *Logical Methods in Computer Science*, 15:1:4:1–4:53, 2019.

18 E. Grädel and W. Pakusa. Rank logic is dead, long live rank logic! *Journal of Symbolic Logic*, 2019.

19 M. Grohe. The quest for a logic capturing PTIME. In *Proceedings of the 23rd IEEE Symposium on Logic in Computer Science (LICS'08)*, pages 267–271, 2008.

20 M. Grohe. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*. Cambridge University Press, 2017.

21 M. Grohe and M.Otto. Pebble Games and linear equations. *J. Symb. Log.*, 80:797–844, 2015.

22 B. Holm. *Descriptive Complexity of Linear Algebra*. PhD thesis, University of Cambridge, 2010.

23 S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.

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

25 W. Pakusa. *Linear Equation Systems and the Search for a Logical Characterisation of Polynomial Time*. PhD thesis, RWTH Aachen University, 2016.

26 R.S. Pierce. *Associative Algebras*. Graduate Texts in Mathematics. Springer, 1982.