On the Relative Power of Linear Algebraic Approximations of Graph Isomorphism

Anuj Dawar
Department of Computer Science and Technology, University of Cambridge, UK

Danny Vagnozzi
Department of Computer Science and Technology, University of Cambridge, UK

Abstract

We compare the capabilities of two approaches to approximating graph isomorphism using linear algebraic methods: the *invertible map tests* (introduced by Dawar and Holm) and proof systems with algebraic rules, namely *polynomial calculus*, *monomial calculus* and *Nullstellensatz calculus*. In the case of fields of characteristic zero, these variants are all essentially equivalent to the Weisfeiler-Leman algorithms. In positive characteristic we show that the distinguishing power of the monomial calculus is no greater than the invertible map method by simulating the former in a fixed-point logic with solvability operators. In turn, we show that the distinctions made by this logic can be implemented in the Nullstellensatz calculus.

2012 ACM Subject Classification Theory of computation → Finite Model Theory; Theory of computation → Proof complexity; Theory of computation → Complexity theory and logic

Keywords and phrases Graph isomorphism, proof complexity, invertible map tests

Digital Object Identifier 10.4230/LIPIcs.MFCS.2021.37

Related Version *Full Version*: https://arxiv.org/abs/2103.16294

Funding Research funded in part by EPSRC grant EP/S03238X/1.

Acknowledgements We want to thank Martin Grohe, Benedikt Pago and Gregory Wilsenach for useful discussions.

1 Introduction

The *graph isomorphism problem* consists in deciding whether there is an edge-preserving bijection between the vertex sets of two given graphs. Computationally, this problem is polynomial-time equivalent to finding the partition into orbits of the action of the automorphism group of a given graph on its vertex set. More generally, it is polynomial-time equivalent to computing the partition into orbits of the *induced* action of the automorphism group on the k^{th} power of the vertex set for any fixed k [20] (we shall refer to this partition as the *k-orbit partition* for a graph). The complexity of these problems is notoriously unresolved: while there are reasons to believe that they are not NP-complete, it is still an open problem as to whether they are in P. The best known upper bound to their computational time is quasi-polynomial, which follows from a breakthrough by Babai [2].

There has been a recent surge of interest in linear-algebraic approaches to the graph isomorphism problem (see for example [4, 12, 17, 19, 22]). In this paper, we consider two distinct methods for incorporating algorithms for solving linear systems into graph isomorphism solvers and compare them. The first is based on the use of algebraic proofs systems, such as the polynomial calculus and the second are generalizations of the Weisfeiler-Leman method, based on stability conditions and coherent algebras.

For every $k \in \mathbb{N}$, the *k-Weisfeiler-Leman* algorithm is a generalization of naïve colour refinement, giving an approximation of the k-orbit partition. For a graph with vertex set V, each of these algorithms outputs in time $|V|^{O(k)}$ a canonical labelled partition of V^k satisfying a stability condition and respecting local isomorphism. Informally, they can be seen as forming a family of algorithms, each defining a notion of equivalence on both graphs and tuples of vertices of graphs. A result by Cai, Fürer and Immerman [7] shows how to construct graphs Γ_k of size $O(k)$ for which the k-Weisfeiler-Leman algorithm fails to produce the k-orbit partition. In the same paper, it is shown that the equivalence classes of the output partition of the k-Weisfeiler-Leman algorithm coincide with the equivalence classes of k-tuples of vertices distinguished by *counting logic* formulae with at most $k+1$ variables. Thus, one deduces from the tight connection made by Immerman and Lander (see Theorem 2.3 in [9]), that the equivalence notions defined by the Weisfeiler-Leman family of algorithms delimit the expressive power of *fixed point logic with counting* (FPC). Intuitively, one such limitation is the expressibility of solvability of systems of linear equations over finite fields, since the above mentioned constructions by Cai, Fürer and Immerman are essentially graph encodings of systems of linear equations over $\mathbb{Z}_2$ [1]. This has therefore prompted research into families of algorithms graded by the natural numbers, whose notion of equivalence on tuples of vertices of graphs is conceptually a linear algebraic invariance over some field $\mathbb{F}$. One such family is that of the *invertible map tests* over a field $\mathbb{F}$, first defined in [12]. For any graph, the k^{th} algorithm of this family also produces a canonical labelled partition of k-tuples of its vertices, satisfying a stability condition and respecting local isomorphism, thus giving another notion of equivalence on both graphs and k-tuples of vertices thereof. For a fixed characteristic, the output of the k-invertible map tests is independent of the field; as such, $\mathbb{F}$ will hereafter be taken to be a prime field without loss of generality.One can claim that the family of equivalences on tuples defined by the Weisfeiler-Leman algorithms and that defined by the invertible map tests over $\mathbb{Q}$ *simulate* each other in the following sense. For every $k \in \mathbb{N}$, there is some $k' \in \mathbb{N}$ such that for any graph, any pair of k-tuples of its vertices distinguished by the k-Weisfeiler-Leman algorithm are distinguished by the k'-invertible map test over $\mathbb{Q}$. Conversely, for every $k \in \mathbb{N}$ there is a $k' \in \mathbb{N}$ such that for any graph, any pair of k-tuples of vertices distinguished by the k-invertible map test over $\mathbb{Q}$ is distinguished by the k'-Weisfeiler-Leman algorithm. If the characteristic of $\mathbb{F}$ is positive the former statement holds, but not the latter; indeed, it is shown in [18] and [10] how one can construct graphs $\Gamma_{k,p}$ for each $k \in \mathbb{N}$ and prime number p, for which the 3-invertible map test over $\mathbb{Z}_p$ outputs the 3-orbit partition, but the output of the k-invertible map test over $\mathbb{Z}_q$ with $q \neq p$ is strictly coarser than the k-orbit partition. A recent construction due to Lichter [19] suggests a way of getting graphs, for any value of k, on which the k-orbit partition cannot be obtained by the k-invertible map test over $\mathbb{Z}_q$ for any q whatsoever.

Another approach to approximating the orbit partition is that of algebraic proof systems [3, 8]. These systems are the subject of very active study in the area of proof complexity. They have been studied specifically in the context of graph isomorphism in [4] and [17]. In particular, the proof systems studied are the *polynomial calculus* (PC), and the weaker *Nullstellensatz calculus* (NC) and *monomial calculus* (MC). Each of these gives, for a fixed field $\mathbb{F}$ a set of rules $\mathcal{R}$ dictating how new polynomials, with coefficients in $\mathbb{F}$, may be derived from an initial set of polynomials, which we shall refer to as *axioms*. In the context of graph isomorphism, we encode any graph Γ on a vertex set V as a set of axioms $\mathrm{Ax}(\Gamma) \subset \mathbb{F}[\{x_{uv} | u, v \in V\}]$, i.e. a collection of polynomials over variables x_{uv} corresponding to potential edges in the graph. An $\mathcal{R}$-derivation of the polynomial $x_{u_1v_1}x_{u_2v_2}\dots x_{u_kv_k}$ can then be seen as a proof that the tuples $\vec{u}, \vec{v} \in V^k$ are distinguishable. Say that such a derivation has *degree* d if all

polynomials occurring in the derivation have degree at most d. For each of the calculi, fixing the degree d gives us a polynomial-time algorithm for checking the existence of a derivation and hence a polynomial-time approximation of the orbit partition for graphs. Again, we may restrict $\mathbb{F}$ to a prime field without loss of generality.

The question we address in this paper is how the approximations of the orbit partition obtained by these algebraic proof systems compare with those we get from the invertible map test. In the case of fields of characteristic zero, the answer is quite clear, as both approaches yield algorithms that are (up to constant factors) equivalent to the Weisfeiler-Leman algorithms. This is shown for the invertible map tests in [13] and for the polynomial calculus in [17]. In the case of positive characteristic, we show (in Section 4) the definability of derivations of MC in FPS(p), an extension of fixed-point logic with quantifiers for the solvability of systems of linear equations over fields of characteristic p. This implies, in particular, that the approximation of the orbit partition obtained by MC in characteristic p is no finer than that obtained by the invertible map test in characteristic p.

▶ **Theorem 1.** *For any prime number p, $k \in \mathbb{N}$ and $\vec{u}, \vec{v} \in V^k$, there is a $k' \in \mathbb{N}$ such that if $x_{u_1v_1} \dots x_{u_kv_k}$ has a degree k MC derivation over $\mathbb{Z}_p$ from $\mathrm{Ax}(\Gamma)$, then $\vec{u}$ and $\vec{v}$ are distinguished by the k'-invertible map test over $\mathbb{Z}_p$.*

In the other direction, we show that NC is able to simulate (as far as the graph isomorphism problem is concerned) PC in characteristic zero and at least MC in positive characteristic.

▶ **Theorem 2.**

1. *For any $k \in \mathbb{N}$ and $\vec{u}, \vec{v} \in V^k$, there is a $k' \in \mathbb{N}$ such that if $x_{u_1v_1} \dots x_{u_kv_k}$ has a degree k PC derivation of over $\mathbb{Q}$ from $\mathrm{Ax}(\Gamma)$, then there is also a degree k' NC derivation over $\mathbb{Q}$ from $\mathrm{Ax}(\Gamma)$.*
2. *For any prime number p, $k \in \mathbb{N}$ and $\vec{u}, \vec{v} \in V^k$, there is a $k' \in \mathbb{N}$ such that if $x_{u_1v_1} \dots x_{u_kv_k}$ has a degree k MC derivation over $\mathbb{Z}_p$ from $\mathrm{Ax}(\Gamma)$, then it also has a degree k' NC derivation over $\mathbb{Z}_p$ from $\mathrm{Ax}(\Gamma)$.*

From this, a strengthening of Theorem 6.3 in [4] also follows. Let V be the vertex set of $\Gamma_{k,p}$ as above.

▶ **Theorem 3.** *If $\vec{u}, \vec{v} \in V^3$ are not in the same equivalence class of the 3-orbit partition of $\Gamma_{k,p}$, then $x_{u_1v_1}x_{u_2v_2}x_{u_3v_3}$ has a degree 3 NC derivation over $\mathbb{Z}_p$ from $\mathrm{Ax}(\Gamma_{k,p})$.*

Due to lack of space we omit the proofs of a number of results. These can be found in the extended version [14].

Notational conventions

All sets are finite unless stated otherwise. Given two sets V and I, a tuple in V^I is denoted by $\vec{v}$, and its i^{th} entry by v_i, for each $i \in I$. We use the notation $(v_i)_{i\in I}$ to denote the element of V^I with i^{th} element equal to v_i. We set $[k] = \{1, 2, \dots, k\} \subset \mathbb{N}$ and define $[k]^{(r)} = \{\vec{x} \in [k]^r \mid x_i \neq x_j \ \forall i, j \in [r],\ i \neq j\}$ for $r \leq k$. For any $\vec{v} \in V^k$, $\vec{u} \in V^r$, and $\vec{i} \in [k]^{(r)}$ we define $\vec{v}\langle\vec{i}, \vec{u}\rangle \in V^k$ to be the tuple with entries

$$(\vec{v}\langle\vec{i}, \vec{u}\rangle)_j = \begin{cases} u_{i_s} & \text{if } j = i_s \text{ for some } s \in [r] \\ v_j & \text{otherwise.} \end{cases}$$

In other words, $\vec{v}\langle\vec{i}, \vec{u}\rangle$ is the tuple obtained from $\vec{v}$ by substituting the elements of $\vec{u}$ in the positions specified by $\vec{i}$. Given two tuples $\vec{v} \in V^r$ and $\vec{w} \in V^s$, their *concatenation* is denoted by $\vec{v} \cdot \vec{w} \in V^{r+s}$. For a relation $R \subseteq V^2$ we define the *adjacency matrix* of R to be the $V \times V$ matrix whose (u, v) entry is 1 if $(u, v) \in R$ and 0 otherwise.

2 Preliminaries

2.1 Labelled partitions and refinement operators

A *labelled partition* of a set A is a map $\gamma : A \to X$, where X is a set of elements sometimes referred to as *colours*, and $\gamma(a)$ as the colour of a in γ. Denote the class of labelled partitions of A by $\mathcal{P}(A)$. For partitions $\mathfrak{R}$ and $\mathfrak{S}$ on A we write $\mathfrak{R} \preceq_A \mathfrak{S}$ and say that $\mathfrak{S}$ *refines* $\mathfrak{R}$ if, whenever $a, b \in A$ are in the same equivalence class of $\mathfrak{S}$, they are also in the same equivalence class of $\mathfrak{R}$. We extend the partial order $\preceq_A$ to labelled paritions by writing $\gamma \preceq_A \rho$ to mean that the equivalence relation $\{(a,b) \mid \rho(a) = \rho(b)\}$ refines the relation $\{(a,b) \mid \gamma(a) = \gamma(b)\}$. Note that this does not require that the co-domains of γ and ρ are the same or indeed related in any way.

Define an action of $\mathrm{Sym}(k)$ on V^k by setting $\vec{v}^\pi$ to be the element of V^k with i^{th} entry $v_{\pi^{-1}(i)}$. $\gamma \in \mathcal{P}(V^k)$ is said to be *invariant* if $\gamma(\vec{u}) = \gamma(\vec{v})$ implies $\gamma(\vec{u}^\pi) = \gamma(\vec{v}^\pi)$ for all $\pi \in \mathrm{Sym}(k)$ and $\vec{u}, \vec{v} \in V^k$.

For $t \in [k]$, the *t-projection* of γ is defined to be the labelled partition $\mathrm{pr}_t\gamma \in \mathcal{P}(V^t)$ such that for all $\vec{u} \in V^t$

$$\mathrm{pr}_t\gamma(\vec{u}) = \gamma(u_1, u_2, \ldots, u_t, \ldots, u_t).$$

For $\vec{v} \in V^k$, with t, k as above, we similarly define $\mathrm{pr}_t\vec{v}$ to be the t-tuple $(v_1, v_2, \ldots, v_t)$.

▶ **Definition 4** (Refinement operator). *A k-*refinement operator *is a map R which, for any set V, assigns to each $\gamma \in \mathcal{P}(V^k)$ a partition $R \circ \gamma \in \mathcal{P}(V^k)$ such that $\gamma \preceq R \circ \gamma$ and it is monotone; that is, $\gamma \preceq \rho \implies R \circ \gamma \preceq R \circ \rho$.*

We say that $\gamma \in \mathcal{P}(V^k)$ is *R-stable* if $R \circ \gamma = \gamma$. Given an $X \in \mathcal{P}(V^k)$, define a sequence of labelled partitions by $X^0 = X$ and $X^{i+1} = R \circ X^i$. Then, there is some s such that for all i, $i \geq s$ implies that X^i is R-stable. For the minimal such s we denote X^s by $[X]^R$.

In order to define the refinement operator leading to the invertible map test (in Section 3.1 below), we need the notion of character vectors. Let $\vec{v} \in V^k$ be a k-tuple of vertices, $\vec{i} \in [k]^{(2r)}$ a $2r$-tuple of indices and $\gamma \in \mathcal{P}(V^k)$ a labelled partition of V^k . For a pair $\vec{x}, \vec{y} \in V^r$ of r-tuples of vertices, $\gamma(\vec{v}\langle\vec{i}, \vec{x} \cdot \vec{y}\rangle)$ is the colour of the tuple obtained by substituting $\vec{x}, \vec{y}$ into $\vec{v}$ in the positions specified by i. For each $\sigma \in \mathrm{Im}(\gamma)$, we define the $V^r \times V^r$ matrix χ_σ with 01 entries as the adjacency matrix of the relation $\{(\vec{x}, \vec{y}) \in (V^r)^2 \mid \gamma(\vec{v}\langle\vec{i}, \vec{x} \cdot \vec{y}\rangle) = \sigma\}$. The $(\vec{i}, \vec{v})$-*character vector* of γ is then defined to be the tuple $\vec{\chi} = (\chi_\sigma)_{\sigma \in \mathrm{Im}(\gamma)}$.

2.2 Extensions of first order and inflationary fixed point logics

We assume the reader has some familiarity with first-order (FO) and fixed-point logics (FP), and logical interpretations. Details can be found in [16]. Throughout the paper, for a logic $\mathcal{L}$, we denote by $\mathcal{L}_k$ the class of all $\mathcal{L}$-formulae (over some pre-specified vocabulary) with at most k variables. We use $\mathcal{C}$ to denote the *counting logics* as in [21]. Let $\mathfrak{A}$ be a structure with universe V and fix $\vec{u}, \vec{v} \in V^k$. We say that some $\mathcal{L}_k$ formula $\phi(\vec{z})$ *distinguishes* $(\mathfrak{A}, \vec{z} \mapsto \vec{u})$ from $(\mathfrak{A}, \vec{z} \mapsto \vec{v})$ if either $\mathfrak{A} \models \phi(\vec{u})$ and $\mathfrak{A} \not\models \phi(\vec{v})$ or $\mathfrak{A} \not\models \phi(\vec{u})$ and $\mathfrak{A} \models \phi(\vec{v})$. For a logic $\mathcal{L}$ we denote its extension via *solvability quantifiers* slv_p over a finite field of characteristic p by $\mathcal{L}{+}\mathrm{S}(p)$. Let $\phi(\vec{x}, \vec{y}, \vec{z})$ be a formula of such a logic, where $\vec{x}, \vec{y}, \vec{z}$ are i, j, k-tuples of variables respectively. Then $\mathrm{slv}_p(\vec{x}\vec{y}.\phi(\vec{x}, \vec{y}, \vec{z}))$ is also a $\mathcal{L}{+}\mathrm{S}(p)$ formula. See [11, 17] for more about these quantifiers. The semantics of this quantifier is as follows. To each structure with universe V and k-tuple $\vec{v} \in V^k$, we associate the $V^i \times V^j$ matrix $S^{\vec{v}}_\phi$ over $\{0,1\} \subseteq \mathbb{Z}_p$ with $(\vec{r}, \vec{s})$-entry equal to 1 if, and only if, $\mathfrak{A} \models \phi(\vec{r}, \vec{s}, \vec{v})$. Then, $(\mathfrak{A}, \vec{z} \mapsto \vec{v}) \models \mathrm{slv}_p(\vec{x}\vec{y}.\phi(\vec{x}, \vec{y}, \vec{z}))$ if, and only if, there is some $\vec{a} \in \mathbb{Z}_p^{V^j}$ such that $S^{\vec{v}}_\phi \vec{a} = \mathbf{1}$.

When $\mathcal{L}$ is FP or FO, we denote $\mathcal{L}{+}\mathrm{S}(p)$ by $\mathrm{FPS}(p)$ or $\mathrm{FOS}(p)$ respectively.

3 Refinement operators and proof systems with algebraic rules

We give an overview of the refinement operators and proof systems of interest.

3.1 The invertible map tests

The equivalence relations on tuples of vertices induced by the invertible map tests have been originally introduced in [12], under the guise of a pebble game with algebraic rules on a pair of graphs. Algorithmically, one can find these equivalence classes by computing a fixed point of the *invertible map operators* $\mathrm{IM}^{\mathbb{F}}_{k,r}$, as defined in Sections 8 of [13]. Each of these is a k-refinement operator such that for $\gamma \in \mathcal{P}(V^k)$, the colour of $\vec{v} \in V^k$ in the partition $\mathrm{IM}^{\mathbb{F}}_{k,r} \circ \gamma$ is given by a tuple whose entries are $\gamma(\vec{v})$ and the equivalence classes under matrix conjugation of the $(\vec{i}, \vec{v})$-character vectors of γ, for all $\vec{i} \in [k]^{(2r)}$. Without going into the details, one can show that γ is $\mathrm{IM}^{\mathbb{F}}_{k,r}$-stable if for all $\vec{u}, \vec{v} \in V^k$ and $\vec{i} \in [k]^{(2r)}$

$$\gamma(\vec{u}) = \gamma(\vec{v}) \implies \exists M \in \mathrm{GL}_{V^r} \text{ s.t. } \forall \sigma \in \mathrm{Im}(\gamma)(\mathbb{F}), M\chi_\sigma M^{-1} = \xi_\sigma$$

where $(\chi_\sigma)_{\sigma \in \mathrm{Im}(\gamma)}$ and $(\xi_\sigma)_{\sigma \in \mathrm{Im}(\gamma)}$ are the $(\vec{i}, \vec{v})$ and $(\vec{i}, \vec{u})$-character vectors of γ respectively.

For a graph Γ on V, let $\alpha_{k,\Gamma}$ be a canonical labelled partition of V^k into atomic types of Γ.[1] Define the k-refinement operator $\mathrm{IM}^{\mathbb{F}}_k$ so that for $\gamma \in \mathcal{P}(V^k)$, the colour of $\vec{v} \in V^k$ in $\mathrm{IM}^{\mathbb{F}}_k \circ \gamma$ is given by a tuple whose entries are the colours of $\vec{v}$ in $\mathrm{IM}^{\mathbb{F}}_{k,r} \circ \gamma$, for all $r \leq k/2$. Formally,

$$\mathrm{IM}^{\mathbb{F}}_k \circ \gamma(\vec{v}) = (\mathrm{IM}^{\mathbb{F}}_{k,1} \circ \gamma(\vec{v}), \mathrm{IM}^{\mathbb{F}}_{k,2} \circ \gamma(\vec{v}), \ldots, \mathrm{IM}^{\mathbb{F}}_{k,\lfloor k/2 \rfloor} \circ \gamma(\vec{v})).$$

Then, the output of the k-invertible map test over $\mathbb{F}$ is the labelled partition $[\alpha_{k,\Gamma}]^{\mathrm{IM}^{\mathbb{F}}_k}$. It is explained in Proposition 4.6 in [13] how one can obtain this partition in time $|V|^{O(k)}$ by iteratively applying $\mathrm{IM}^{\mathbb{F}}_k$ to $\alpha_{k,\Gamma}$. Note that for a fixed characteristic, the choice of $\mathbb{F}$ is irrelevant: indeed, if k-tuples $\vec{A}, \vec{B} \in \mathrm{Mat}_V(\mathbb{F})^k$ are related by matrix conjugation over $\mathbb{F}$ if, and only if, they are related by matrix conjugation over any field extension of $\mathbb{F}$ [15]. Since the entries of the character vectors are 01-matrices, we may assume, without loss of generality, that $\mathbb{F}$ is a prime field. Hereafter, we shall then indicate the operators $\mathrm{IM}^{\mathbb{F}}_{k,r}$ and $\mathrm{IM}^{\mathbb{F}}_k$ by $\mathrm{IM}^{c}_{k,r}$ and IM^{c}_k respectively, where c is the characteristic of $\mathbb{F}$.

3.2 Counting logics operators

It is useful to express the partition of k-tuples into equivalence classes under finite variable counting logics as the fixed point of a refinement operator. For this purpose, the k-refinement operators $\mathrm{C}_{k,r}$, for $r < k$, have been defined so that for any $\gamma \in \mathcal{P}(V^k)$, the colour of $\vec{v} \in V^k$ in $\mathrm{C}_{k,r} \circ \gamma$ is given by a tuple whose entries are $\gamma(\vec{v})$ and the multisets of colours in γ of the tuples which can be obtained by substituting an r-tuple into $\vec{v}$ (see Section 4 of [13]). In particular, γ is $\mathrm{C}_{k,r}$-stable if, and only if, for all $\vec{i} \in [k]^{(r)}$ and $\sigma \in \mathrm{Im}(\gamma)$, the size of the set $\{\vec{x} \in V^r \mid \gamma(\vec{v}\langle\vec{i}, \vec{x}\rangle) = \sigma\}$ is independent of the choice of $\vec{v}$ from within its equivalence class in γ. As such, for a graph Γ on V, $\vec{u}, \vec{v} \in V^k$ are in the same equivalence class of $[\alpha_{k,\Gamma}]^{\mathrm{C}_{k,1}}$ if, and only if, there are no $\mathcal{C}_k$ formulae distinguishing $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{u})$ from $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{v})$. The combinatorial properties of $\mathrm{C}_{k,r}$-stable partitions can be used to show the relation between

[1] By *canonical* we mean *invariant under isomorphism*. That is, if Γ and Γ' are graphs on V and V' respectively and $\vec{u} \in V^k$ and $\vec{v} \in (V')^k$, then $\alpha_{k,\Gamma}(\vec{u}) = \alpha_{k,\Gamma'}(\vec{v})$ if, and only if, the mapping $u_i \to v_i$ is an isomorphism of the subgraphs induced by the vertices in $\vec{u}$ and $\vec{v}$.

the distinguishing powers of finite variable fragments of counting logics and the invertible map tests. In short, the invertible map test over fields of characteristic zero is not more distinguishing than counting logic, but over fields of positive characteristic it is. To be precise, with $\Gamma, \vec{u}, \vec{v}$ as above, the following holds:

> For any field $\mathbb{F}$, if the k-invertible map test over $\mathbb{F}$ does not distinguish $\vec{u}$ and $\vec{v}$ in Γ, then there are no $\mathcal{C}_{k-1}$-formulae distinguishing $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{u})$ from $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{v})$.
> If the k-invertible map test over $\mathbb{Q}$ distinguishes $\vec{u}$ from $\vec{v}$ in Γ, then there is some $\mathcal{C}_{2k-1}$-formula distinguishing $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{u})$ from $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{v})$. [2]

3.3 Solvability operators

In order to construct a refinement operator whose stable points reflect the properties of FOS(p), we consider a weakened version of the invertible map operators, whose action on labelled partitions can be computed solely by solving systems of linear equations. To achieve this, we proceed by defining an equivalence relation $\sim_{\text{sol}}$ on the character vectors, which can be seen as a relaxation of the conjugation relation $\sim$.

Let $\mathfrak{P}_V(\mathbb{F}) = \{A \in \text{Mat}_V(\mathbb{F}) \mid \sum_{w \in V} A_{wv} = \sum_{w \in V} A_{uw} = 1, \forall u, v \in V\}$. Or, equivalently, $\mathfrak{P}_V(\mathbb{F})$ is the set of matrices $A \in \text{Mat}_V(\mathbb{F})$ such that $\mathbf{1}_V$ is an eigenvector of both A and A^t, with corresponding eigenvalue 1. Set $\mathbb{J}_V$ to be the $V \times V$ all-ones matrix and for a set I, let

$$\mathcal{X}_V^I(\mathbb{F}) = \{\vec{A} \in \text{Mat}_V(\mathbb{F})^I \mid \sum_{s \in I} A_s = \mathbb{J}_V \text{ and } \forall i \in I, \exists j, A_i^t = A_j\}.$$

Note that if γ is invariant and $\vec{i} \in [k]^{(2r)}$, then any $(\vec{i}, \vec{v})$- character $\vec{\chi}$ of γ is an element of $\mathcal{X}_V^{\text{Im}(\gamma)}(\mathbb{F})$, since

$$\sum_{\sigma \in \text{Im}(\gamma)} \chi_\sigma = \mathbb{J}_{V^r} \tag{1}$$

and by invariance of γ, for any $\sigma \in \text{Im}(\gamma)$ there is some σ' such that $(\chi_\sigma)^t = \chi_{\sigma'}$. Define the relation $\sim_{\text{sol}}$ on $\mathcal{X}_V^I(\mathbb{F})$ as follows: $\vec{A} \sim_{\text{sol}} \vec{B}$ if there is some $S \in \mathfrak{P}_V(\mathbb{F})$ such that $A_i S = S B_i$ for all $i \in I$.

▶ **Lemma 5.** *$\sim_{\text{sol}}$ is an equivalence relation on $\mathcal{X}_V^I(\mathbb{F})$.*

Proof. Clearly, $\vec{A} \sim_{\text{sol}} \vec{A}$, since $\mathbb{I}_V \in \mathfrak{P}_V(\mathbb{F})$ and $A_i \mathbb{I}_V = \mathbb{I}_V A_i$ for all $i \in I$. Suppose $\vec{A} \sim_{\text{sol}} \vec{B}$. Let $S \in \mathfrak{P}_V(\mathbb{F})$ satisfy $A_i S = S B_i$ for all $i \in I$. Then $B_i^t S^t = S^t A_i^t$ and thus, from the definition of $\mathcal{X}_V^I$, $B_i S^t = S^t A_i$ for all $i \in I$. Since $S^t \in \mathfrak{P}_V(\mathbb{F})$, $\vec{B} \sim_{\text{sol}} \vec{A}$. Finally, suppose $\vec{A} \sim_{\text{sol}} \vec{B}$ and $\vec{B} \sim_{\text{sol}} \vec{C}$. Let $S, T \in \mathfrak{P}_V(\mathbb{F})$ satisfy $A_i S = S B_i$ an $B_i T = T C_i$ for all $i \in I$. Then $A_i ST = S B_i T = STC_i$. Since S, T, S^t and T^t must all have $\mathbf{1}_{V^r}$ as eigenvector, with corresponding eigenvalue 1, so must ST and $(ST)^t$. Hence, $ST \in \mathfrak{P}_V(\mathbb{F})$ and $\vec{A} \sim_{\text{sol}} \vec{C}$. ◀

---

[2] This is a direct consequence of the following generalizations of Lemmata 7.1 and 7.3 in [13] respectively: for all $k, r \in \mathbb{N}$ with $2r < k$,

1. The k-projection of a graph-like $\text{IM}_{k+r,r}^c$-stable partition is $\text{C}_{k,r}$-stable for any characteristic c.
2. The k-projection of a graph-like $\text{C}_{k+r,r}$-stable partition is $\text{IM}_{k,r}^0$-stable.

The authors prove it only for the case $r = 1$, but a similar argument holds for any $r \in \mathbb{N}$.

For $k, r \in \mathbb{N}$ with $2r \leq k$, a field $\mathbb{F}$, and an invariant $\gamma \in \mathcal{P}(V^k)$, we define the *solvability operators* $\mathrm{S}^{\mathbb{F}}_{k,r}$ by setting $\mathrm{S}^{\mathbb{F}}_{k,r} \circ \gamma$ to be the labelled partition for which the colour of $\vec{v} \in V^k$ is a tuple whose entries are $\gamma(\vec{v})$ and the equivalence classes under the relation $\sim_{\mathrm{sol}}$ of the $(\vec{i}, \vec{v})$-character vectors of γ, for all $\vec{i} \in [k]^{(2r)}$. Formally:

$$\begin{array}{rccl} \mathrm{S}^{\mathbb{F}}_{k,r} \circ \gamma : & V^k & \to & \mathrm{Im}(\gamma) \times (\mathcal{X}^{\mathrm{Im}(\gamma)}_{V^r}(\mathbb{F})/\sim_{\mathrm{sol}})^{[k]^{(2r)}} \\ & \vec{v} & \mapsto & (\gamma(\vec{v}), (\vec{\chi}_{\vec{i}})_{\vec{i} \in [k]^{(2r)}}), \end{array}$$

where $\vec{\chi}_{\vec{i}}$ is the $(\vec{i}, \vec{v})$-character vector of γ.

As before, since the entries of the matrices in the character vector are all 0 and 1, we may restrict $\mathbb{F}$ to being a prime field without loss of generality, and denote $\mathrm{S}^{\mathbb{F}}_{k,r}$ by $\mathrm{S}^{c}_{k,r}$, where $c = \mathrm{char}(\mathbb{F})$. It is easy to show that $\mathrm{S}^{c}_{k,r}$ is monotone on the class of invariant partitions of V^k and is thus a k-refinement operator when considered with this domain restriction.

For the remainder of this section, we assume that $\gamma \in \mathcal{P}(V^k)$ is invariant and that $\vec{\chi}$ and $\vec{\xi}$ are the $(\vec{i}, \vec{v})$ and $(\vec{i}, \vec{u})$-character vectors of γ respectively, for some fixed $\vec{i} \in [k]^{(2r)}$. The following is a direct consequence of the definition of $\mathrm{S}^{c}_{k,r}$.

▶ **Proposition 6.** *Let $\mathbb{F}$ be the prime field of characteristic c. For all $k, r \in \mathbb{N}$, with $2r \leq k$, $\mathrm{S}^{c}_{k,r} \circ \gamma(\vec{u}) = \mathrm{S}^{c}_{k,r} \circ \gamma(\vec{v})$ if, and only if, $\gamma(\vec{u}) = \gamma(\vec{v})$ and for each $\vec{i} \in [k]^{(2r)}$ there exist some $M \in \mathfrak{P}_{V^r}(\mathbb{F})$ such that for all $\sigma \in \mathrm{Im}(\gamma)$, $\chi_\sigma M = M\xi_\sigma$. In particular, γ is $\mathrm{S}^{c}_{k,r}$-stable if, and only if, for all $\vec{u}, \vec{v} \in V^k$ and $\vec{i} \in [k]^{(2r)}$*

$$\gamma(\vec{u}) = \gamma(\vec{v}) \implies \exists M \in \mathfrak{P}_{V^r}(\mathbb{F}) \text{ s.t. } \forall \sigma \in \mathrm{Im}(\gamma), \chi_\sigma M = M\xi_\sigma \ \forall \sigma \in \mathrm{Im}(\gamma).$$

We now show some useful properties of the operators $\mathrm{S}^{c}_{k,r}$.

▶ **Lemma 7.** *An $\mathrm{IM}^{c}_{k,r}$-stable partition is $\mathrm{S}^{c}_{k,r}$-stable.*

Proof. Let $\mathbb{F}$ be the prime field of characteristic c. Suppose γ is $\mathrm{IM}^{c}_{k,r}$-stable and let $\gamma(\vec{u}) = \gamma(\vec{v})$. Then, for all $\vec{i} \in [k]^{(2r)}$, there is some $M \in GL_{V^r}(\mathbb{F})$ such that $M^{-1}\chi_\sigma M = \xi_\sigma$ for all $\sigma \in \mathrm{Im}(\gamma)$. By equation 1, $M\mathbb{J}_{V^r} = \mathbb{J}_{V^r}M$ and thus, both M and M^t have $\mathbf{1}_{V^r}$ as eigenvector with corresponding non-zero eigenvalue $\lambda \in \mathbb{F}$. Since, $\frac{1}{\lambda}M \in \mathfrak{P}_{V^r}(\mathbb{F})$, the result follows. ◀

▶ **Lemma 8.** *If γ is $\mathrm{S}^{c}_{k,r}$-stable the following hold:*

1. *If $\gamma(\vec{u}) = \gamma(\vec{v})$, and γ is graph-like, then for all $\vec{i} \in [k]^{(2r)}$ and $\sigma \in \mathrm{Im}(\gamma)$, $\{\vec{w} \in V^r \mid \gamma(\vec{u}\langle\vec{i}, \vec{w} \cdot \vec{w}\rangle) = \sigma\}$ is non-empty if, and only if, $\{\vec{w} \in V^r \mid \gamma(\vec{v}\langle\vec{i}, \vec{w} \cdot \vec{w}\rangle) = \sigma\}$ is non-empty.*
2. *If $c = 0$, then γ is $\mathrm{C}_{k,2r}$-stable.*

Note that if $c = 0$, the second statement implies the first (refer to Section 4 and 5 of [13] for more details on properties of $\mathrm{C}_{k,r}$-stability).

Proof. Suppose $\{\vec{w} \in V^r \mid \gamma(\vec{u}\langle\vec{i}, (\vec{w} \cdot \vec{w})\rangle) = \sigma\}$ is non-empty. Since γ is graph-like, χ_σ must have all the non-zero entries on the diagonal. Hence, for any $M \in \mathfrak{P}_{V^r}(\mathbb{F})$, $\chi_\sigma M$ and, consequently $M\xi_\sigma$ must be non-zero. As such, $\{\vec{w} \in V^r \mid \gamma(\vec{v}\langle\vec{i}, (\vec{w} \cdot \vec{w})\rangle) = \sigma\}$ is non-empty. The converse can be argued by symmetry, thus showing (1).

For (2), it suffices to show that: if $A, B \in \mathrm{Mat}_V(\mathbb{F})$ are 01-matrices and $AM = MB$ for some $M \in \mathfrak{P}_V(\mathbb{F})$ then A and B have the same number of non-zero entries if $\mathrm{char}(\mathbb{F}) = 0$. Indeed, note that if α and β are the number of non-zero entries of A and B respectively, then $\alpha\mathbb{J}_V = \mathbb{J}_V A\mathbb{J}_V$ and $\beta\mathbb{J}_V = \mathbb{J}_V B\mathbb{J}_V$. Since $M\mathbb{J}_V = \mathbb{J}_V M = \mathbb{J}_V$,

$$\alpha\mathbb{J}_V = \mathbb{J}_V A\mathbb{J}_V = \mathbb{J}_V AM\mathbb{J}_V = \mathbb{J}_V MB\mathbb{J}_V = \mathbb{J}_V B\mathbb{J}_V = \beta\mathbb{J}_V,$$

whence $\alpha = \beta$. ◀

In particular, if $r = 1$, statement (1) above implies that there is some $x \in V$ such that $\mathrm{pr}_{k-1}\gamma(\mathrm{pr}_{k-1}\vec{u}\langle i, x\rangle) = \sigma$ if, and only if, there is some $y \in V$, such that $\mathrm{pr}_{k-1}\gamma(\mathrm{pr}_{k-1}\vec{v}\langle i, y\rangle) = \sigma$. Furthermore, from (2) and Lemma 5.7 in [13], $\mathrm{pr}_{k-1}\gamma$ is C_{k-1} stable.

▶ **Corollary 9.** *Let $\gamma = [\alpha_{k,\Gamma}]^{\mathrm{S}^c_{k,1}}$ for some graph Γ on V. If $\mathrm{pr}_{k-1}\gamma(\vec{u}) = \mathrm{pr}_{k-1}\gamma(\vec{v})$, there are no first-order formulae distinguishing $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{u})$ from $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{v})$. In addition, if $c = 0$, there are no $\mathcal{C}_{k-1}$-formulae distinguishing $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{u})$ from $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{v})$.*

Similarly to the invertible map operators, we define the k-refinement operator S^c_k so that for $\gamma \in \mathcal{P}(V^k)$, the colour of $\vec{v} \in V^k$ in $\mathrm{S}^c_k \circ \gamma$ is given by a tuple whose entries are the colours of $\vec{v}$ in $\mathrm{S}^c_{k,r} \circ \gamma$, for all $r < k/2$; that is,

$$\mathrm{S}^c_k \circ \gamma(\vec{v}) = (\mathrm{S}^c_{k,1} \circ \gamma(\vec{v}), \mathrm{S}^c_{k,2} \circ \gamma(\vec{v}), \ldots, \mathrm{S}^c_{k,\lfloor k/2 \rfloor} \circ \gamma(\vec{v})).$$

The next statement will be the crux of our main results.

▶ **Theorem 10.** *For any prime number p, if $\gamma = [\alpha_{k,\Gamma}]^{\mathrm{S}^p_k}$ and $\gamma(\vec{u}) = \gamma(\vec{v})$, there are no $FOS_{k-1}(p)$ formulae distinguishing $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{u})$ from $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{v})$.*

Proof. We proceed by induction on the structure of $\mathrm{FOS}(p)$ formulae $\phi(\vec{z})$, where $\vec{z}$ is a k-tuple of pairwise distinct variables. If $\phi(\vec{z})$ contains only atomic formulae, boolean connectives and first order quantifiers, the statement holds by Corollary 9. Assume that for some $\phi(\vec{z}) \in \mathrm{FOS}_k(p)$, $\mathfrak{A}_\Gamma \models \phi(\vec{u}) \iff \mathfrak{A}_\Gamma \models \phi(\vec{v})$, and suppose $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{v}) \models \mathrm{slv}_p[\vec{x}\vec{y}.\phi(\vec{z}\langle\vec{i}, \vec{x} \cdot \vec{y}\rangle)]$, where $\vec{i} \in [k]^{(2r)}$ and $\vec{x}, \vec{y}$ are r-tuples of distinct variables (distinct from variables in the tuple $\vec{z}$). Let $S^{\vec{v}}$ be the adjacency matrix of the relation $\{(\vec{a}, \vec{b}) \in V^r \times V^r \mid (\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{v}) \models \phi(\vec{v}\langle\vec{i}, \vec{a} \cdot \vec{b}\rangle)\}$, and similarly define $S^{\vec{u}}$. Then, there is some $\vec{a} \in \mathbb{Z}_p^{V^r}$ such that $S^{\vec{v}}\vec{a} = \mathbf{1}_{V^r}$. By the induction hypothesis, $S^{\vec{u}} = \sum_{\sigma \in I} \chi_\sigma$ and $S^{\vec{v}} = \sum_{\sigma \in I} \xi_\sigma$ for some $I \subseteq \mathrm{Im}(\gamma)$. Since there exists $M \in \mathfrak{P}_{V^r}(\mathbb{Z}_p)$ such that $\chi_\sigma M = M\xi_\sigma$ for all $\sigma \in \mathrm{Im}(\gamma)$, we have

$$S^{\vec{v}}\vec{a} = \mathbf{1}_{V^r} \implies MS^{\vec{v}}\vec{a} = \mathbf{1}_{V^r} \implies S^{\vec{u}}(M\vec{a}) = \mathbf{1}_{V^r},$$

from which it follows that $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{u}) \models \mathrm{slv}_p[\vec{x}\vec{y}.\phi(\vec{z}\langle\vec{i}, \vec{x} \cdot \vec{y}\rangle)]$. Using a symmetric argument, we conclude that $\mathrm{slv}_p[\vec{x}\vec{y}.\phi(\vec{z}\langle\vec{i}, \vec{x} \cdot \vec{y}\rangle)]$ does not distinguish $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{u})$ from $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{v})$ and the result follows by induction. ◀

3.4 Polynomial, monomial and Nullstellensatz calculi

The idea behind *polynomial calculus* (PC), *monomial calculus* (MC) and *Nullstellensatz calculus* (NC) is that of encoding Boolean formulae as multivariate polynomials and concluding that they are inconsistent if the polynomials do not have a common root. The PC inference rules for a set of axioms $A \subset \mathbb{F}[x_1, ..., x_n]$ are as follows:

1. $\frac{}{f}$ for all $f \in A$.
2. *Multiplication rule*: $\frac{f}{xf}$ for all derived polynomials f, and variables $x \in \{x_1, ..., x_n\}$.
3. *Linearity rule*: $\frac{f,g}{\lambda f + \mu g}$ for all derived polynomials f, g and $\lambda, \mu \in \mathbb{F}$.

The inference rules for MC are obtained by restricting f in the multiplication rule to be an axiom times a monomial or a monomial. By further restricting f to be an axiom times a monomial, one obtains the NC inference rules. The *degree* of a PC (or MC or NC) derivation is the maximum degree of all polynomials involved in the derivation, and a PC (or MC or NC, respectively) *refutation* of A is a derivation of 1 from A using PC (or MC or NC, respectively) rules. We are really interested in roots where the variables are assigned 01-values, so as to encode the Boolean framework. To enforce this, we assume that the axioms always includes the polynomials $x^2 - x$ for all $x \in \{x_1, \ldots, x_n\}$. We may therefore restrict our focus to multilinear polynomials exclusively.

We denote by PC_k the proof system using the same inference rules as PC with the added constraint that all derivations must have degree at most k, and we use the same convention for MC_k and NC_k. Though these proof systems with bounded degree are not complete, their refutations can be decided in polynomial time in the number of variables.

For a graph Γ on V we define $\mathrm{Ax}(\Gamma) \subseteq \mathbb{F}[\{x_{uv} | u, v \in V\}]$ to be the set of axioms containing the following polynomials:

1. $\sum_{u \in V} x_{uv} - 1$ for all $v \in V$.
2. $\sum_{u \in V} x_{vu} - 1$ for all $v \in V$.
3. $x_{uv}x_{u'v'}$ if the map $u \mapsto u', v \mapsto v'$ is not a local isomorphism in Γ.
4. $x_{uv}^2 - x_{uv}$ for all $u, v \in V$.

For $\vec{u}, \vec{v} \in V^k$, we further define $\mathrm{Ax}(\Gamma_{\vec{u}\to\vec{v}})$ to contain the above plus $x_{v_i u_i} - 1$ for all $i \in [k]$.

When considering these axiom we assume, without loss of generality, that $\mathbb{F}$ is a prime field. Let $\equiv^c_{\mathrm{PC}_k}$ be the relation on V^k, where $\vec{u} \equiv^c_{\mathrm{PC}_k} \vec{v}$ if there is no degree k PC refutation of $\mathrm{Ax}(\Gamma_{\vec{u}\to\vec{v}})$ over the prime field of characteristic c, and similarly define $\equiv^c_{\mathrm{MC}_k}$ and $\equiv^c_{\mathrm{NC}_k}$.

▶ **Lemma 11.** *$\equiv^c_{\mathrm{PC}_k}$, $\equiv^c_{\mathrm{MC}_k}$ and $\equiv^c_{\mathrm{NC}_k}$ are equivalence relations on V^k.*

Proof. Clearly, $\vec{u} \equiv^c_{\mathrm{PC}_k} \vec{u}$. Indeed, the polynomials in the ideal generated by $\mathrm{Ax}(\Gamma_{\vec{u}\to\vec{u}})$ have the common root $x_{rs} = \delta_{rs}$, where δ_{rs} is the Kronecker delta. Hence, the ideal generated by $\mathrm{Ax}(\Gamma_{\vec{u}\to\vec{u}})$ is non-trivial. The set $\mathrm{Ax}(\Gamma)$ is invariant under the transformation $x_{rs} \to x_{sr}$ for all $r, s \in V$. Thus, $\vec{u} \equiv^c_{\mathrm{PC}_k} \vec{v} \implies \vec{v} \equiv^c_{\mathrm{PC}_k} \vec{u}$.

Suppose $\vec{u} \equiv^c_{\mathrm{PC}_k} \vec{v}$ and $\vec{v} \equiv^c_{\mathrm{PC}_k} \vec{w}$. Let π be the map $v_i \to w_i$ for $i \in [k]$. Note that π is well defined, for if $v_i = v_j$ for some $i \neq j$ and $\pi(v_i) \neq \pi(v_j)$, then $x_{v_i w_i} x_{v_j w_j} \in \mathrm{Ax}(\Gamma)$ - a contradiction. For $A \subseteq V^2$, define A^π as follows:

$$A^\pi = \begin{cases} \{(r, \pi(s)) \mid (r, s) \in A\} & \text{if for all } (r,s) \in A \text{ there is some } j \text{ such that } s = v_j; \\ \{(r, \pi^{-1}(s)) \mid (r, s \in A)\} & \text{if for all } (r,s) \in A \text{ there is some } j \text{ such that } s = w_j; \text{ and} \\ A & \text{otherwise.} \end{cases}$$

For a multilinear polynomial $f = \sum a_A X_A$, let $p^\pi = \sum a_A X_{A^\pi}$. We show by induction on the PC inference rules that there is a PC_k derivation of p if, and only if, there is PC_k derivation of f^π. Indeed, if f is in $\mathrm{Ax}(\Gamma)$, then so is f^π. Suppose there is a PC_k derivation of f, g, f^π and g^π. Then there is a PC_k derivation of $(\lambda f + \mu g)^\pi = \lambda f^\pi + \mu g^\pi$. Finally, suppose the degree of f is less than k, and suppose, without loss of generality that $f = X_A$ for some $A \subseteq V^2$. Then, there is a PC_k derivation of $X_{A \cup \{(r,s)\}}$ for any $r, s \in V$. By checking case by case, it follows that there is a PC_k derivation of $(X_{A\cup\{(r,s)\}})^\pi$. Since $(A^\pi)^\pi = A$ and hence, $(f^\pi)^\pi = f$, there is a PC_k derivation of f if, and only if, there is a PC_k derivation of f^π. In particular, since there is no PC_k derivation of $X_{\{(v_i,w_i)|i\in[k]\}}$, there is no derivation of $(X_{\{(u_i,v_i)|i\in[k]\}})^\pi = X_{\{(u_i,w_i)|i\in[k]\}}$. Whence, $\vec{u} \equiv^c_{\mathrm{PC}_k} \vec{w}$. ◀

It is easy to see that the relation $\equiv^c_{\mathrm{PC}_k}$ refines $\equiv^c_{\mathrm{MC}_k}$ which, in turn, refines $\equiv^c_{\mathrm{NC}_k}$, since a NC_k refutation is a MC_k refutation which is also a PC_k refutation. More precisely:

▶ **Lemma 12.** *For any graph on V and $\vec{u}, \vec{v} \in V^k$, $\vec{u} \equiv^c_{\mathrm{PC}_k} \vec{v} \implies \vec{u} \equiv^c_{\mathrm{MC}_k} \vec{v} \implies \vec{u} \equiv^c_{\mathrm{NC}_k} \vec{v}$.*

For $c = 0$, Grohe et al. have characterized these relations in terms of counting logics:

> Let Γ be a graph on V and $\vec{u}, \vec{v} \in V^k$. Then $\vec{u} \equiv^0_{\mathrm{MC}_k} \vec{v}$ if, and only if, no $\mathcal{C}_k$ formula distinguishes $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{u})$ from $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{v})$ (Theorem 4.4 in [4]). Furthermore, if $\vec{u} \not\equiv^0_{\mathrm{PC}_k} \vec{v}$, there a $k' = O(k)$, such that some $\mathcal{C}_{k'}$ formula distinguishes $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{u})$ from $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{v})$ (Theorem 6.6 in [17]).

Our main results attempt to give a similar characterization for $c > 0$ in terms of logics with solvability quantifiers.

4 Definability of monomial calculus refutations over finite fields

At the core of the proof of Theorem 1 is the definability of monomial calculus refutations in FPS(p). More precisely, the main objective of this section is to prove the following statement (we will explain what we mean by structural encoding in Section 4.1).

▶ **Lemma 13.** *Let $\mathfrak{A}$ be a structural encoding of a finite set of polynomials P of degree at most d, over a finite field $\mathbb{F}$ of positive characteristic p. For any $k \in \mathbb{N}$ there is a* FPS(p) *formula $\phi_{d,k}$ such that $\mathfrak{A} \models \phi_{d,k}$ if, and only if, there is an* MC_k *refutation of P over $\mathbb{F}$.* [3]

For the sake of argument, we assume that $\mathbb{F} = \mathbb{Z}_p$. We first recall how to express the solvability of linear equations with coefficients other than 0 and 1.

4.1 Defining solvability of linear equations over finite fields

For each prime number p, let $\mathtt{LIN}_p$ be a relational vocabulary with the following symbols:

1. A binary relational symbol $\mathtt{A}_q$ for each $q \in \mathbb{Z}_p$.
2. A unary relational symbol $\mathtt{b}_q$ for each $q \in \mathbb{Z}_p$.

Let $A \in \mathrm{Mat}_{E\times V}(\mathbb{Z}_p)$ and $\vec{b} \in \mathbb{Z}_p^E$. A $\mathtt{LIN}_p$-structure $\mathfrak{A}$ with universe $V \cup E$ (V for variables and E for equations) is a structural encoding of the system of linear equations $A\vec{x} = \vec{b}$ if, for all $e \in E, v \in V$, $\mathfrak{A} \models \mathtt{A}_q(e, v)$ if $A_{ev} = q$ and $\mathfrak{A} \models \mathtt{b}_q(e)$ if $b_e = q$.

Recall Lemma 4.1 in [11].

▶ **Lemma 14.** *There is a quantifier free interpretation $\mathcal{I}$ of $\mathtt{LIN}_p$ into $\mathtt{LIN}_p$ such that if $\mathfrak{A}$ encodes the system of linear equations $A\vec{x} = \vec{b}$, then:*

1. *$\mathcal{I}(\mathfrak{A})$ encodes a system of linear equations $A'\vec{y} = \mathbf{1}$, where $\mathbf{1}$ is the all 1s vector of appropriate length and A' is a 01-matrix.*
2. *$A'\vec{y} = \mathbf{1}$ has a solution if, and only if, $A\vec{x} = \vec{b}$ has a solution.*

Thus, $\mathcal{I}(\mathfrak{A}) \models \mathrm{slv}_p(xy.\mathtt{A}_1(x, y))$ if, and only if, $A\vec{x} = \vec{b}$ has a solution and hence, there is a FOS(p) formula Φ such that $\mathfrak{A} \models \Phi$ if, and only if, the system encoded by $\mathfrak{A}$ has a solution.

4.2 Idea of proof of Lemma 13

Deciding whether a set of axioms has a monomial calculus refutations of a given degree can be understood as the following procedure. The input is a finite set of multilinear polynomials P from the ring $\mathbb{F}[x_1, \ldots, x_r]$, and the output is `REFUTE` or `NOREFUTE`. We denote the multilinear monomial $x_{a_1}x_{a_2}\ldots x_{a_r}$ by X_A where $A = \{a_1, a_2 \ldots, a_r\}$, so that for a polynomial f, $X_A f = x_{a_1}x_{a_2}\ldots x_{a_r}f$. In this form, $X_\emptyset$ denotes 1.

INPUT: $P \subset \mathbb{F}[x_1, \ldots, x_r]$
OUTPUT: `REFUTE` or `NOREFUTE`.
Initialize $\mathcal{S} = \{X_A f \mid f \in P, \deg(X_A f) \leq k\}$.
while $\mathrm{span}_{\mathbb{F}}\mathcal{S}$ has changed since last round or $1 \notin \mathrm{span}_{\mathbb{F}}\mathcal{S}$ **do**
  $\mathcal{M} \leftarrow \{A \subseteq [r] \mid |A| < k, X_A \in \mathrm{span}_{\mathbb{F}}\mathcal{S}\}$.
  $\mathcal{S} \leftarrow \mathcal{S} \cup \{X_B \mid |B| \leq k, \exists A \in \mathcal{M}, A \subseteq B\}$.
end while
if $1 \in \mathrm{span}_{\mathbb{F}}\mathcal{S}$ **then**
  OUTPUT `REFUTE`
else output `NOREFUTE`.
end if

[3] Note that it is possible to derive a similar result independent of the parameter d. Since $d = 2$ for axioms of the form $\mathrm{Ax}(\Gamma)$, the statement of Lemma 13 suffices for our purpose.

Note that to verify the condition of the `while` loop, one need not store in memory the set $\text{span}_{\mathbb{F}}\mathcal{S}$ (whose size is exponential in the input); this can be done by checking the solvability of linear equations. The number of iterations of the `while` loop is at most the number of multilinear monomials of degree at most k, thus ensuring that the procedure runs in polynomial time. Crucially, at each iteration of the `while` loop, the $\mathbb{F}$-span of $\mathcal{S}$ has a *canonical* generating set.

Recall that we are assuming that $\mathbb{F} = \mathbb{Z}_p$. By viewing polynomials over $\mathbb{Z}_p$ as vectors in the standard basis given by monomials, we encode P as a structure $\mathfrak{A}$ with universe V over the vocabulary $\texttt{POLY}_p = (\texttt{Var}, \texttt{U}, \texttt{C}_0, \texttt{C}_1, \ldots, \texttt{C}_{p-1})$, where:

1. $V = P \cup \{x_i \mid i \in [r]\} \cup \{1\}$.
2. `Var` and `U` are unary relational symbol with $\mathfrak{A} \models \texttt{Var}(v)$ if, and only if, $v \in \{x_i \mid i \in [r]\}$, and $\mathfrak{A} \models \texttt{U}(v)$ if, and only if, $v = 1$.
3. $\texttt{C}_q$ for $q \in \mathbb{Z}_p$ are $(d+1)$-ary relations, where d is the maximal degree of the polynomials in P. These encode P in matrix form; that is, $\mathfrak{A} \models \texttt{C}_q(u, v_1, \ldots, v_d)$ if, and only if, $u \in P$, each v_i is equal to a variable or 1, and the coefficient of the monomial $v_1 v_2 \ldots v_d$ in u is equal to q.

We now use a k-ary interpretation to obtain a structure whose universe is the set of multilinear polynomials of degree at most k. Formally, let $\texttt{POLY}^*_p = (\texttt{A}_0, \texttt{A}_1, \ldots, \texttt{A}_{p-1}, \texttt{U}^*, \texttt{Mon}, \texttt{Sub})$ be a vocabulary where `Mon` and $\texttt{U}^*$ are unary relational symbols, `Sub` is a binary relational symbol, and $\texttt{A}_q$ are as in the vocabulary $\texttt{LIN}_p$. One can then define an interpretation $\mathcal{J}$ from $\texttt{POLY}_p$ into $\texttt{POLY}^*_p$ such that:

1. The universe of $\mathcal{J}(\mathfrak{A})$ are elements $\vec{v} \in V^k$ where either all v_i are equal to 1, all v_i are variables, or v_1 is a polynomial in P and $v_2, \ldots, v_k$ are either all equal to 1 or are variables such that $|\{v_i \mid 2 \leq i \leq k\}| \leq k - \deg(v_1)$. The relation $\equiv_{\mathcal{J}}$ partitions the universe into equivalence classes uniquely determined by the set of entries of each tuple. If the entries of $\vec{v}$ are all equal to 1, we indicate its class by $X_\emptyset$, and similarly, if $v_i = x_{a_i}$ for all $i \in [k]$, we indicate its class by X_A where $A = \{a_i \mid i \in [k]\}$. If $v_1 = f$ for some $f \in P$ and $v_2, \ldots, v_k$ are all equal to 1, we indicate the class of $\vec{v}$ by f and if $v_i = a_i$ for $2 \leq i \leq k$, we indicate its class by the pair (X_A, f) where $A = \{a_i \mid 2 \leq i \leq k\}$. We leave it to the reader to check that such an equivalence relation can be defined with first-order formulae.
2. $\mathfrak{A} \models \texttt{U}(\vec{v})$ and $\mathfrak{A} \models \texttt{Mon}(\vec{v})$ if, and only if, the equivalence class of $\vec{v}$ is $X_\emptyset$ and X_A with $|A| \geq 1$ respectively.
3. $\mathfrak{A} \models \texttt{A}_q(\vec{u}, \vec{v})$ if, and only if, the equivalence class of $\vec{u}$ and $\vec{v}$ are X_A and (X_B, f) for some $f \in P$ respectively, and the coefficient of X_A in $X_B f$ equals q.
4. $\mathfrak{A} \models \texttt{Sub}(\vec{u}, \vec{v})$ if, and only if, the equivalence classes of $\vec{u}$ and $\vec{v}$ are X_A and X_B respectively and $A \subseteq B$.

The last thing required for proving Lemma 13 is showing the definability of monomials in the set $\mathcal{S}$ *after* each iteration of the `while` loop. In what follows, set $\mathcal{T} = \{X_A f \mid f \in P, A \subseteq [r], \deg(X_A f) \leq k\}$.

Proof of Lemma 13. Let $\psi(z)$ be some FPS(p) formula over the vocabulary $\texttt{POLY}^*_p$. There is an interpretation $\mathcal{K}(t)$ of $\texttt{POLY}^*_p$ into $\texttt{LIN}_p$ such that $\mathcal{K}(\mathcal{J}(\mathfrak{A}), t \mapsto X_A)$ encodes the system of linear equations determining whether X_A is in the $\mathbb{Z}_p$-span of the set

$$\mathcal{T}_\psi = \mathcal{T} \cup \{X_B \mid (\mathfrak{A}, z \mapsto X_B) \models \psi(z) \wedge \texttt{Mon}(z)\}.$$

By Lemma 14, there is a FPS(p) formula $\theta_\psi(z)$, depending on ψ, such that $(\mathcal{J}(\mathfrak{A}), z \mapsto X_A) \models \theta_\psi(z)$ if, and only if, the monomial X_A is in the $\mathbb{Z}_p$-span of $\mathcal{T}_\psi$.

In particular, replacing ψ with some unary relational variable Z,

$$(\mathcal{J}(\mathfrak{A}), z \mapsto X_A) \models \exists y.\theta_Z(y) \wedge \mathtt{Sub}(y, z)$$

holds if, and only if, there is some $B \subseteq A$ such that X_B is in the $\mathbb{Z}_p$-span of $\mathcal{T}_Z$. Whence,

$$(\mathcal{J}(\mathfrak{A}), z' \mapsto X_\emptyset) \models \mathrm{ifp}_{Z,z}(\exists y.\theta_Z(y) \wedge \mathtt{Sub}(y, z))(z') \tag{2}$$

if, and only if, there is an MC_k refutation of P over $\mathbb{Z}_p$. By applying the Interpretation Lemma to the above formula and the interpretation $\mathcal{J}$, the desired result follows. ◀

Note that in formula 2, the solvability quantifier is included in the formula $\theta_Z(y)$.

▶ **Corollary 15.** *For any $k \in \mathbb{N}$, there is some k' such that for any graph Γ on V and $\vec{u}, \vec{v} \in V^k$, if there is a MC_k refutation of $\mathrm{Ax}(\Gamma_{\vec{u}\to\vec{v}})$ over $\mathbb{Z}_p$, then there is some $FOS_{k'}(p)$ formula $\phi(\vec{z})$ distinguishing $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{u})$ from $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{v})$.*

Proof. By Lemma 13 the equivalence classes of $\equiv^p_{\mathrm{MC}_k}$ are definable in $\mathrm{FPS}(p)$ and hence, by the embedding of $\mathrm{FPS}(p)$ in infinitary $\mathrm{FOS}(p)$, the statement follows. ◀

Combining the latter with Theorem 10 and Lemma 7, one deduces Theorem 1.

5 Nullstellensatz refutations and S^p_k-stability

The focus of this section is the proof of the following statement.

▶ **Lemma 16.** *Let Γ be a graph on V and let $\gamma \in \mathcal{P}(V^k)$. If for all $\vec{u}, \vec{v} \in V^k$, $\gamma(\vec{u}) = \gamma(\vec{v})$ if, and only if, $\vec{u} \equiv^p_{\mathrm{NC}_k} \vec{v}$, then γ is S^p_k-stable.*

Theorems 2 and 3 are its direct consequences.

5.1 Proof of Lemma 16

In what follows, $\gamma \in \mathcal{P}(V^k)$ satisfies the assumptions of Lemma 16, $\vec{u}, \vec{v} \in V^k$ and $\vec{\chi}$ and $\vec{\xi}$ are the $(\vec{i}, \vec{u})$ and $(\vec{i}, \vec{v})$-character vectors of γ respectively, where we may assume without loss of generality that $\vec{i} = (k, k-1, \ldots, k-2r+1) \in [k]^{(2r)}$, for some r with $2r \leq k$. For $\vec{w}, \vec{z} \in V^l$, denote by $X_{\vec{w}\vec{z}}$ the monomial $x_{w_1 z_1} x_{w_2 z_2} \ldots x_{w_l z_l}$ (which need not be multilinear).

Recall that γ is $\mathrm{S}^p_{k,r}$-stable, if, and only if, for every $\vec{u}, \vec{v}$ there is some matrix $T \in \mathfrak{P}_{V^r}(\mathbb{Z}_p)$ such that $\chi_\sigma T = T\xi_\sigma$, for all $\sigma \in \mathrm{Im}(\gamma)$. That is, the following system of equations is solvable in the variables $T_{\vec{w}\vec{z}}$:

$$\sum_{\vec{a}\in V^r} (\chi_\sigma)_{\vec{w}\vec{a}} T_{\vec{a}\vec{z}} - \sum_{\vec{a}\in V^r} T_{\vec{w}\vec{a}} (\xi_\sigma)_{\vec{a}\vec{z}} = 0 \quad \text{for } \sigma \in \mathrm{Im}(\gamma),\ \vec{w}, \vec{z} \in V^r$$

$$\sum_{\vec{a}\in V^r} T_{\vec{w}\vec{a}} - 1 = 0 \quad \text{and} \quad \sum_{\vec{a}\in V^r} T_{\vec{a}\vec{z}} - 1 = 0 \quad \text{for } \vec{w}, \vec{z} \in V^r,$$

where $(\chi_\sigma)_{\vec{w}\vec{a}}$ is equal to 1 if $\gamma(\vec{u}\langle\vec{i}, \vec{w}\cdot\vec{a}\rangle) = \sigma$ and 0 otherwise (similarly for ξ_σ).

We show that there is a NC_k derivation from $\mathrm{Ax}(\Gamma_{\vec{v}\to\vec{u}})$ over $\mathbb{Z}_p$ of the following multilinear polynomials (Lemma 19):

$$\sum_{\vec{a}|\gamma(\vec{u}\langle\vec{i},\vec{w}\cdot\vec{a}\rangle)=\sigma} X_{\vec{a}\vec{z}} - \sum_{\vec{a}|\gamma(\vec{v}\langle\vec{i},\vec{a}\cdot\vec{z}\rangle)=\sigma} X_{\vec{w}\vec{a}} \quad \text{for } \sigma \in \mathrm{Im}(\gamma),\ \vec{w},\vec{z} \in V^r \tag{3}$$

$$\sum_{\vec{a}\in V^r} X_{\vec{w}\vec{a}} - 1 \ \text{ and } \ \sum_{\vec{a}\in V^r} X_{\vec{a}\vec{z}} - 1 \ \text{ for } \vec{w},\vec{z} \in V^r. \tag{4}$$

We may view each monomial (apart from the constant term) as a distinct linear variable, so that all of the above are linear polynomials. Since $\gamma(\vec{u}) = \gamma(\vec{v})$, there is no NC_k refutation of $\mathrm{Ax}(\Gamma_{\vec{v}\to\vec{u}})$ and hence, no linear combination of 3 and 4 gives the constant polynomial 1. It follows that if viewed as linear polynomials, 3 and 4 have a common root, thus showing Lemma 16.

▶ **Lemma 17.** *If $\gamma(\vec{u}\langle\vec{i},\vec{w}\cdot\vec{z}\rangle) \neq \gamma(\vec{v}\langle\vec{i},\vec{w}'\cdot\vec{z}'\rangle)$, then there is a NC_k derivation over $\mathbb{Z}_p$ from $\mathrm{Ax}(\Gamma_{\vec{v}\to\vec{u}})$ of $X_{\vec{w}\vec{w}'}X_{\vec{z}\vec{z}'}$.*

Proof. Set $Y = X_{\vec{w}\vec{w}'}X_{\vec{z}\vec{z}'}$ and let $X_{\vec{u}'\vec{v}'}$ be the degree $k-2r$ monomial where $u'_j = u_j$ and $v'_j = v_j$ for $j \in [k-2r]$. Then, there is a NC_k derivation of $Y(X_{\vec{u}'\vec{v}'} - 1)$, for indeed

$$Y(X_{\vec{u}'\vec{v}'}-1) = Y(x_{u_1v_1}-1)+Yx_{u_1v_1}(x_{u_2v_2}-1)+\ldots+Y(x_{u_1v_1}\ldots x_{u_{k-2r-1}v_{k-2r-1}})(x_{u_{k-2r}v_{k-2r}}-1).$$

By assumption, $\gamma(\vec{u}\langle\vec{i},\vec{w}\cdot\vec{z}\rangle) \neq \gamma(\vec{v}\langle\vec{i},\vec{w}'\cdot\vec{z}'\rangle)$ and hence, there is an NC_k derivation of $YX_{\vec{u}'\vec{v}'}$. Subtracting the latter from $Y(X_{\vec{u}'\vec{v}'} - 1)$ yields the desired statement. ◀

▶ **Lemma 18.** *For any $\vec{w},\vec{z} \in V^r$ and $\vec{s} \in V^t$ there is a NC_{t+r} derivation over $\mathbb{Z}_p$ from $\mathrm{Ax}(\Gamma_{\vec{v}\to\vec{u}})$ of*

$$X_{\vec{w}\vec{z}}\Big(\sum_{\vec{a}\in V^t} X_{\vec{s}\vec{a}} - 1\Big) \ \text{ and } \ X_{\vec{w}\vec{z}}\Big(\sum_{\vec{a}\in V^t} X_{\vec{a}\vec{s}} - 1\Big).$$

Proof. We proceed by induction on t. For $t = 1$, $X_{\vec{w}\vec{z}}\big(\sum_{a\in V} x_{sa} - 1\big)$ is the product of a monomial and an axiom, so has a NC_{r+1} derivation.

Assume $X_{\vec{w}\vec{z}}\big(\sum_{\vec{a}\in V^t} X_{\vec{s}\vec{a}} - 1\big)$ has a NC_{t+r} derivation. It can be easily verified that if a polynomial f has a NC_r derivation from some set of axioms, then xf has a NC_{r+1} derivation for any variable x. Thus, $x_{s'a'}X_{\vec{w}\vec{z}}\big(\sum_{\vec{a}\in V^t} X_{\vec{s}\vec{a}} - 1\big)$ has a NC_{t+r+1} derivation. Finally

$$\sum_{a'\in V} x_{s'a'}X_{\vec{w}\vec{z}}\Big(\sum_{\vec{a}\in V^t} X_{\vec{s}\vec{a}} - 1\Big) = X_{\vec{w}\vec{z}}\Big(\sum_{\vec{a}\in V^t, a'\in V} X_{(\vec{s}\cdot s')(\vec{a}\cdot a')} - 1\Big) = X_{\vec{w}\vec{z}}\Big(\sum_{\vec{a}\in V^{t+1}} X_{(\vec{s}\cdot s')\vec{a}} - 1\Big)$$

as required. ◀

▶ **Lemma 19.** *There is a NC_k derivation over $\mathbb{Z}_p$ from $\mathrm{Ax}(\Gamma_{\vec{v}\to\vec{u}})$ of the polynomials in formulae (3) and (4).*

Proof. For $\vec{a} \in V^r$, set $\mathcal{N}(\vec{u},\vec{w}) = \{\vec{a} \in V^r \mid \gamma(\vec{u}\langle\vec{i},\vec{w}\cdot\vec{a}\rangle) = \sigma\}$ and $\mathcal{N}(\vec{v},\vec{z}) = \{\vec{a} \in V^r \mid \gamma(\vec{v}\langle\vec{i},\vec{a}\cdot\vec{z}\rangle) = \sigma\}$ (note the slight asymmetry). By Lemma 18, there is a NC_{2r} (and hence, NC_k, since $2r \leq k$) derivation of $X_{\vec{a}\vec{z}}\big(\sum_{\vec{a}'\in V^r} X_{\vec{s}\vec{a}'} - 1\big)$ for every $\vec{a},\vec{s},\vec{z} \in V^r$. By subtracting from the above all monomials $X_{\vec{a}\vec{z}}X_{\vec{s}\vec{a}'}$ for which $\vec{a}' \notin \mathcal{N}(\vec{v},\vec{z})$ (which have a NC_k derivation by Lemma 17) one gets $X_{\vec{a}\vec{z}}\big(\sum_{\vec{a}'\in\mathcal{N}(\vec{v},\vec{z})} X_{\vec{s}\vec{a}'} - 1\big)$. Adding these for all $\vec{a} \in \mathcal{N}(\vec{u},\vec{w})$ yields

$$\sum_{\vec{a}\in\mathcal{N}(\vec{u},\vec{w})} X_{\vec{a}\vec{z}}\Big(\sum_{\vec{a}'\in\mathcal{N}(\vec{v},\vec{z})} X_{\vec{s}\vec{a}'} - 1\Big). \tag{5}$$

A similar argument shows that there is a NC_k derivation of

$$\sum_{\vec{a}\in\mathcal{N}(\vec{v},\vec{z})} X_{\vec{w}\vec{a}}\Big(\sum_{\vec{a}'\in\mathcal{N}(\vec{u},\vec{w})} X_{\vec{a}'\vec{s}} - 1\Big). \tag{6}$$

Subtracting (5) from (6) yields (3).

The polynomials in (4) can be derived by setting $r = 0$ in Lemma 18. ◀

5.2 Generalized Cai-Fürer-Immerman constructions

In 1992, Cai, Fürer and Immerman provided, for $k \geq 1$, a family of pairs of non-isomorphic graphs $(\mathcal{G}_k, \mathcal{H}_k)$ which cannot be distinguished by $\mathcal{C}_k$-formulae. These graphs really encode a system of linear equations over $\mathbb{Z}_2$, the solvability of which can be decided in polynomial time. These structures can be generalized to encode a systems of linear equations over an arbitrary finite field. Loosely speaking, the *generalized Cai-Fürer-Immerman* construction for the field $\mathbb{Z}_p$ provides, for each $k \in \mathbb{N}$, p non-isomorphic 3-regular graphs $\mathcal{G}_k^{(1)}, \ldots, \mathcal{G}_k^{(p)}$. These delimit the power of well known linear algebra based polynomial-time approximations of graph isomorphism. Let $\Gamma_{k,p}$ be the disjoint union of the graphs $\mathcal{G}_k^{(1)}, \ldots, \mathcal{G}_k^{(p)}$ and let V denote its vertex set.

▶ **Theorem 20** (Theorems 8.1 and 8.2 in [10]). *If $\Gamma = \Gamma_{k,p}$ and $q \neq p$, then the equivalence classes of $[\alpha_{k,\Gamma}]^{\mathrm{IM}_k^q}$ do not coincide with those of the k-orbit partition for Γ. If $q = p$, the equivalence classes of $[\alpha_{3,\Gamma}]^{\mathrm{IM}_3^q}$ coincide with those of the 3-orbit partition for Γ.*

Originally, the generalized Cai-Fürer-Immerman constructions were introduced to delimit the expressive power of the extension of fixed point logics with rank operators over a finite field $\mathrm{FPR}(p)$ and, correspondingly, the distinguishing power of the extension of first order logic by said operators, $\mathrm{FOR}(p)$ (see Chapter 7 of [18]). The distinguishing power of the invertible map test in characteristic p is at least that of $\mathrm{FOR}(p)$, so in one direction Theorem 20 provides a lower bound for this logic. In the other direction, we can show that the orbit partition on $\Gamma_{k,p}$ can already be defined in $\mathrm{FPR}(p)$. Indeed, it can be defined in the apparently weaker logic $\mathrm{FPS}(p)$, giving the following result, of which Theorem 3 is a direct consequence.

▶ **Theorem 21.** *Let $\Gamma = \Gamma_{k,p}$. If $p \neq q$, there exist $\vec{u}, \vec{v} \in V^k$ in different equivalence classes of the k-orbit partition of Γ, such that there are no $FOS_k(q)$-formulae distinguishing $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{u})$ from $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{v})$. If $q = p$ and $\vec{u}, \vec{v} \in V^2$, then $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{u})$ is distinguished from $(\mathfrak{A}_\Gamma, \vec{z} \mapsto \vec{v})$ by some $FOS_2(q)$-formula if, and only if, $\vec{u}, \vec{v}$ are in different equivalence classes of the 2-orbit partition of Γ.*

6 Conclusions: where does polynomial calculus lie?

The invertible map tests can be thought of a family of algorithms, each of which distinguishes tuples of vertices of graphs according to the most general linear algebraic invariants expressible with a bounded number of variables in a logic. As bounded degree PC, MC and NC refutations can be decided solely by using basic field operations, one expects that the equivalences defined by the invertible map tests simulate those defined by the above mentioned proof systems. Theorem 1 gives a partial proof of this conjecture, leaving it open as to whether the invertible map tests can simulate bounded degree PC refutations, when taken over some finite field.

Our approach to this question was to attempt to define the above proof systems in the simplest logics which could express the solvability of systems of linear equations. The proof of Lemma 13 hints at the flaws of this choice. Deciding whether or not there is a PC_k refutation of $P \subseteq \mathbb{F}[x_1, \ldots, x_r]$ can be understood as the following procedure, similar to that in Section 4.2.

INPUT: $P \subset \mathbb{F}[x_1, \ldots, x_r]$
OUTPUT: `REFUTE` or `NOREFUTE`.
Initialize $\mathcal{S} = \{X_A f \mid f \in P, \deg(X_A f) \leq k\}$.
while $\text{span}_{\mathbb{F}}\mathcal{S}$ has changed since last round or $1 \notin \text{span}_{\mathbb{F}}\mathcal{S}$ **do**
  Find a set $\mathcal{B}$ generating the $\mathbb{F}$-space $\{f \in \text{span}_{\mathbb{F}}\mathcal{S} \mid \deg(f) < k\}$.
  $\mathcal{S} \leftarrow \mathcal{S} \cup \{X_A f \mid f \in \mathcal{B}, \deg(X_A f) \leq k\}$.
end while
if $1 \in \text{span}_{\mathbb{F}}\mathcal{S}$ output `REFUTE`. **then**
else output `NOREFUTE`
end if

This procedure runs in polynomial time, as one can find $\mathcal{B}$ by using Gaussian elimination (there is no need to store the set $\text{span}_{\mathbb{F}}\mathcal{S}$ as a generating set suffices), and the number of iterations of the `while` loop is bounded by the number of monomials of degree at most k in the variables $\{x_1, \ldots, x_r\}$. If $\mathbb{F}$ is finite, this procedure is a priori *not* definable in FPS(p), as it is not immediate whether there is a canonical choice for the set $\mathcal{B}$ (its counterpart in the procedure for monomial calculus refutations was the set $\mathcal{M}$ of monomials in the span of $\mathcal{S}$). Put otherwise, defining the set $\mathcal{B}$ requires defining the solution space of a system of linear equations over a field of positive characteristic p, rather than just determining the solvability of the system and it is not clear if this can be done in FPS(p). The FPC definability of bounded degree polynomial calculus over the field $\mathbb{Q}$ (Theorem 4.9 in [17]) relies in fact on the FPC definability of solution spaces of linear equations over $\mathbb{Q}$ (Theorem 4.11 in [17]).

Let us view the problem from the viewpoint of proof complexity. It follows from Lemma 16, that if PC_k refutations over $\mathbb{Z}_p$ are definable in FPS(p), then for every k, there is some k' such that if $\text{Ax}(\Gamma_{\vec{u}\to\vec{v}})$ has a PC_k refutation over $\mathbb{Z}_p$, then it has a $\text{NC}_{k'}$ refutation over $\mathbb{Z}_p$. It is known that for all n and for any field, there is a set of axioms on $n(n+1)$ variables which can be refuted by PC_3 but require degree $\Omega(n)$ to be refuted by NC (Theorem 6 in [5]). Furthermore, this lower bound can be shown to be optimal. On the other hand, Buss et al. have shown that NC derivations can be used to simulate *tree-like* PC derivations (see Theorems 5.3 and 5.4 in [6]) with only a small increase in degree. For a set of axioms of the form $\text{Ax}(\Gamma_{\vec{u}\to\vec{v}})$, it is not known if any PC_k refutation of such can be converted into a tree-like refutation without incurring in an unbounded increase in degree.

References

1 A. Atserias, A. Bulatov, and A. Dawar. Affine systems of equations and counting infinitary logic. *Theoretical Computer Science*, 410(18):1666–1683, 2009. Automata, Languages and Programming (ICALP 2007).

2 L. Babai. Graph isomorphism in quasipolynomial time [extended abstract]. In *Proc. 48th Annual ACM SIGACT Symp. Theory of Computing, STOC*, pages 684–697, 2016.

3 P. Beame, R. Impagliazzo, J. Krajicek, T. Pitassi, and P. Pudlak. Lower bounds on Hilbert's Nullstellensatz and propositional proofs. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, volume 73, pages 794–806, 1994.

4 C. Berkholz and M. Grohe. Limitations of algebraic approaches to graph isomorphism testing. *CoRR*, abs/1502.05912, 2015.

5 S. R. Buss. Lower bounds on Nullstellensatz proofs via designs. In *in Proof Complexity and Feasible Arithmetics, P. Beame and S. Buss, eds., American Mathematical Society*, pages 59–71. American Math. Soc, 1998.

6 S. R. Buss, R. Impagliazzo, J. Krajícek, P. Pudlak, A. Razborov, and J. Sgall. Proof complexity in algebraic systems and bounded depth Frege systems with modular counting. *Computational Complexity*, 6:256–298, January 1997. `doi:10.1007/BF01294258`.

7 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.

8 M. Clegg, J. Edmonds, and R. Impagliazzo. Using the Gröbner basis algorithm to find proofs of unsatisfiability. *Proceedings of STOC'96*, March 2000.

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

10 A. Dawar, E. Grädel, and W. Pakusa. Approximations of isomorphism and logics with linear algebraic operators. In *46th International Colloquium on Automata, Languages, and Programming, ICALP*, pages 112:1–112:14, 2019. doi:10.4230/LIPIcs.ICALP.2019.112.

11 A. Dawar, E. Graedel, B. Holm, E. Kopczyński, and W. Pakusa. Definability of linear equation systems over groups and rings. *Logical Methods in Computer Science*, 9, April 2012. doi:10.2168/LMCS-9(4:12)2013.

12 A. Dawar and B. Holm. Pebble games with algebraic rules. In Artur Czumaj, Kurt Mehlhorn, Andrew Pitts, and Roger Wattenhofer, editors, *Automata, Languages, and Programming*, pages 251–262, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

13 A. Dawar and D. Vagnozzi. Generalizations of k-Weisfeiler-Leman stabilization. *Moscow Journal of Number Theory and Combinatorics*, 2020.

14 A. Dawar and D. Vagnozzi. On the relative power of algebraic approximations of graph isomorphism. *arXiv*, 2021. arXiv:2103.16294.

15 C. de Seguins Pazzis. Invariance of simultaneous similarity and equivalence of matrices under extension of the ground field. *Linear Algebra and its Applications*, 433, February 2009. doi:10.1016/j.laa.2010.03.022.

16 H-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 2nd edition, 1999.

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), 2019.

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

19 M. Lichter. Separating rank logic from polynomial time. In *Proc. 36th ACM/IEEE Symp. on Logic in Computer Science (LICS)*, 2021.

20 R. Mathon. A note on the graph isomorphism counting problem. *Information Processing Letters*, 8:131–136, 1979.

21 M. Otto. *Bounded Variable Logics and Counting – A Study in Finite Models*, volume 9 of *Lecture Notes in Logic*. Springer-Verlag, 1997.

22 G. Tinhofer. Graph isomorphism and theorems of Birkhoff type. *Computing*, 36:285–300, 1986.