


Symmetric Proofs in the Ideal Proof System

Anuj Dawar 

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

Erich Grädel 

Mathematical Foundations of Computer Science, RWTH Aachen University, Germany

Leon Kullmann 

RWTH Aachen University, Germany

Benedikt Pago 

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

Abstract

We consider the Ideal Proof System (IPS) introduced by Grochow and Pitassi and pose the question of which tautologies admit symmetric proofs, and of what complexity. The symmetry requirement in proofs is inspired by recent work establishing lower bounds in other symmetric models of computation. We link the existence of symmetric IPS proofs to the expressive power of logics such as fixed-point logic with counting and Choiceless Polynomial Time, specifically regarding the graph isomorphism problem. We identify relationships and tradeoffs between the symmetry of proofs and other parameters of IPS proofs such as size, degree and linearity. We study these on a number of standard families of tautologies from proof complexity and finite model theory such as the pigeonhole principle, the subset sum problem and the Cai-Fürer-Immerman graphs, exhibiting non-trivial upper bounds on the size of symmetric IPS proofs.

2012 ACM Subject Classification Theory of computation → Proof complexity; Theory of computation → Algebraic complexity theory; Theory of computation → Finite Model Theory

Keywords and phrases proof complexity, algebraic complexity, descriptive complexity, symmetric circuits, graph isomorphism

Digital Object Identifier 10.4230/LIPIcs.MFCS.2025.40

Related Version *Full Version:* <https://arxiv.org/abs/2504.16820> [11]

Funding *Anuj Dawar:* Funded by UK Research and Innovation (UKRI) under the UK government's Horizon Europe funding guarantee: grant number EP/X028259/1.

Benedikt Pago: Funded by UK Research and Innovation (UKRI) under the UK government's Horizon Europe funding guarantee: grant number EP/X028259/1.

Acknowledgements We are grateful to Iddo Tzameret for familiarising us with the current trends in IPS lower bounds and for valuable suggestions for future directions of this project.

1 Introduction

A central project within the subject of proof complexity [2, 32, 25] is to define ever more powerful proof systems for propositional logic for which it is possible to establish superpolynomial lower bounds on the size of proofs. The *Ideal Proof System* (IPS), introduced by Grochow and Pitassi in [21], is a powerful *algebraic* proof system that subsumes many previous algebraic and propositional proof systems such as the polynomial calculus and Frege proof systems. No superpolynomial lower bounds are known for the unrestricted IPS and it has been shown, for instance, that it efficiently simulates powerful proof systems such as extended Frege, for which such lower bounds are a long-standing open problem. On the other hand, lower bounds have been shown for interesting restrictions of IPS.



© Anuj Dawar, Erich Grädel, Leon Kullmann, and Benedikt Pago;
licensed under Creative Commons License CC-BY 4.0

50th International Symposium on Mathematical Foundations of Computer Science (MFCS 2025).

Editors: Paweł Gawrychowski, Filip Mazowiecki, and Michał Skrzypczak; Article No. 40; pp. 40:1–40:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In IPS, like other algebraic proof systems, propositional formulas are coded as polynomial equations and a proof (or refutation) is a certificate showing that a given equation system has no solution. Equivalently, the certificate shows that the input polynomials, or *axioms*, have no common zero. In IPS, the certificate is an *algebraic circuit* which witnesses that the unit polynomial 1 is in the ideal generated by the axioms, and we measure the complexity of the certificate in terms of the size of the circuit. Here, an algebraic circuit is a representation of a polynomial as a circuit with constants and variables as inputs and addition and multiplication gates. In order to make progress on showing lower bounds for IPS, it has been suggested [21] to study fragments of IPS given by restrictions on the circuits considered. A successful recent line of work has established super-polynomial lower bounds for fragments such as bounded-depth circuits, multilinear formulas and so-called read-once oblivious algebraic branching programs (roABPs) [17, 18, 23]. The common method behind these results has been named the *functional lower bound* technique.

In the present paper, we consider a natural restriction on IPS proofs based on *symmetric circuits*. Non-trivial lower bounds for symmetric algebraic circuits have recently been established [15, 16, 13] and this raises the question of whether, when the input polynomials have natural symmetries, the corresponding proof can also be made symmetric. Or, perhaps more interestingly, whether the lower bound methods for symmetric algebraic circuits can be possibly combined with functional lower bound techniques to obtain lower bounds for symmetric IPS proofs.

The present paper lays the foundations for this. We formalise the symmetry requirement in IPS proofs and show that any proof can be symmetrised though the interaction with linear IPS and other natural fragments is subtle. Beyond this, our work has two main contributions. First, we show how the study of symmetric IPS fits into the context of a larger project that investigates the role of symmetry in lower bounds in a number of areas of complexity theory (see [9]). This project is centred around the concept of *symmetric computation* in the sense of symmetric circuits and logics from finite model theory [10, 12]. Specifically, it seeks to lift game-based methods from finite model theory which are used to show inexpressiveness results in logic to lower bounds in other models of computing. There is a well-established connection between the expressive power of such logics and propositional proof complexity. In particular, it has been shown that, by encoding propositional axioms as relational structures in a suitable way, the problem of the existence of a resolution refutation of a given width, or of a polynomial calculus refutation of a fixed degree are expressible in existential fixed-point logic and fixed-point logic with counting (FPC) respectively and are indeed *complete* for these logics under weak, symmetry-preserving reductions [29]. Hence, in a precise sense these logics exactly characterise the power of the corresponding proof systems and from lower bounds on the expressive power of the logics we can extract lower bounds in proof complexity. Moreover, the stronger logic Choiceless Polynomial Time (CPT) can, in a symmetry-preserving way, be simulated in the bounded-degree extended polynomial calculus [27]. CPT is of great importance in descriptive complexity as a candidate logic for capturing PTIME [28, 19].

As the results we show in Section 5 establish, symmetric IPS proofs offer a suitable unifying formalism in which the results from [29, 27] can be recast. The expressive power of FPC corresponds to the bounded-degree symmetric IPS proofs, and the stronger logic CPT can only be simulated by symmetric IPS proofs of unbounded degree. Importantly, the symmetries of the IPS proofs we consider are the inherent symmetries of the input structure, which are the symmetries under which the corresponding FPC or CPT computation is invariant.

The second focus of this paper is upper bounds for the size of symmetric IPS refutations of various families of instances that possess natural symmetries. One example, where upper (and lower) bounds actually follow from results in finite model theory, is a formulation of the *graph isomorphism problem* as a system of polynomial equations [3] over \mathbb{Q} . The equations expressing that two graphs G and H are isomorphic are symmetric under the automorphism groups of G and H . We show that the complexity of symmetric IPS refutations with these symmetry groups matches the descriptive complexity of distinguishing G and H : Graphs that are distinguishable in FPC, or equivalently by the k -dimensional *Weisfeiler-Leman algorithm* [24] for fixed k , admit bounded-degree polynomial-size symmetric IPS refutations. Graphs that are distinguishable in the stronger logic CPT, but not in FPC, admit polynomial-size symmetric IPS refutations, but not of bounded degree. This separates in particular the bounded-degree fragment of symmetric IPS from its unbounded-degree version (Theorem 13).

A related family of instances we consider is a system of linear equations over the finite field \mathbb{F}_2 that expresses the isomorphism problem on Cai-Fürer-Immerman (CFI) graphs [6, 1] which is known to be inexpressible in FPC. Even though the instances exhibit a large number of interesting symmetries and are the core of numerous finite model theory lower bounds, it turns out quite surprisingly that they do admit polynomial-size symmetric IPS refutations. However, we are not able to exhibit a *linear* symmetric IPS refutation of the CFI equations of less than exponential size, making them a candidate for separating the linear fragment from the full symmetric IPS. Here the linear (or *Hilbert-like*) fragment is as defined in [21] (see also Section 3.2).

Besides the CFI equations, we consider standard hard examples from proof complexity, namely the *pigeonhole principle* (PHP) and the *subset sum* problem. The former is the classic tautology asserting that there is no injective function from a set of $n + 1$ elements to a set of n elements. The latter is simply the equation $\sum_{i=1}^n x_i - \beta = 0$, together with the equations $x_i^2 - x_i = 0$ for all i , which is unsatisfiable when $\beta > n$. The pigeonhole principle has been used to obtain lower bounds for resolution [22] and bounded-degree polynomial calculus [30]. This is again a promising candidate for showing lower bounds on the size of symmetric IPS proofs as the tautologies are rich in symmetries and we know of no sub-exponential size symmetric IPS refutations for PHP.

In contrast, for subset sum, we do obtain polynomial size symmetric IPS refutations, which is significant because subset sum and its variations are the key instances where the functional lower bound method [17, 18, 23] has been used to obtain lower bounds. Thus, symmetry alone is a limitation of a rather different nature than others, e.g. multilinear-formula, constant depth or roABP-IPS where functional lower bounds have been established.

Our upper bounds on proof sizes are summarised in the table below. In this table, the entry “none” indicates that no proofs with the given restrictions exist. In the first row, k is an arbitrary positive integer and c is a function of k . In the other rows c is a fixed constant. Our upper bounds for the graph isomorphism problem depend on the graphs: We consider graphs that are distinguishable by the k -dimensional Weisfeiler Leman (WL) algorithm versus graphs that are distinguished by CPT; the latter distinguishes strictly more graphs [14]. It should be noted that even the exponential lower bounds are non-trivial as the trivial bound for a symmetric circuit may be $n!$.

Proof System	FPC-definable problems	Graph isomorphism	CFI	Subset sum	Pigeonhole principle
$\text{deg}_k\text{-sym-IPS}$	$\mathcal{O}(n^c)$	$\mathcal{O}(n^c)$ if k -WL-distinguishable	none	none	none
$\text{sym-IPS}_{\text{LIN}}$	$\mathcal{O}(n^c)$	$\mathcal{O}(n^c)$ if CPT-distinguishable	$\mathcal{O}(2^n)$	$\mathcal{O}(n^c)$	$\mathcal{O}(3^n \cdot n)$
sym-IPS	$\mathcal{O}(n^c)$	$\mathcal{O}(n^c)$ if CPT-distinguishable	$\mathcal{O}(n^c)$	$\mathcal{O}(n^c)$	$\mathcal{O}(3^n \cdot n)$

2 Preliminaries

We denote by $[n]$ the set $\{1, \dots, n\}$. We call a set X , with the action of a group Γ on it a Γ -set. The action of Γ on X has a *natural extension* to a number of other sets defined from X . In particular, the following are Γ -sets: X^k for any $k \in \mathbb{N}$ ($\pi \in \Gamma$ takes (x_1, \dots, x_k) to $(\pi(x_1), \dots, \pi(x_k))$); the set A^X of functions from X to any set A (here $\pi \in \Gamma$ takes $f : X \rightarrow A$ to $f \circ \pi^{-1}$); and the collection $\mathbb{F}[X]$ of polynomials over X with coefficients from a field \mathbb{F} .

If X is a Γ -set and $p \in \mathbb{F}[X]$ a polynomial, then we say that p is Γ -symmetric or Γ -invariant if $\pi(p) = p$ for every $\pi \in \Gamma$. A set of polynomials $\mathcal{F} \subseteq \mathbb{F}[X]$ is Γ -invariant if $\{\pi(f) \mid f \in \mathcal{F}\} = \mathcal{F}$ for every $\pi \in \Gamma$. We write $\mathbf{Sym}(X)$ for the *symmetric group* on X and for a finite relational structure \mathfrak{A} we write $\mathbf{Aut}(\mathfrak{A})$ for the *automorphism group* of \mathfrak{A} .

In what follows, X denotes a finite set of variables, \mathbb{F} denotes a field, and $\mathbb{F}[X]$ is the corresponding ring of polynomials in the variables X . Instances for the algebraic proof systems we consider are *systems of polynomial equations*. A system of polynomial equations is a finite set $\mathcal{F} \subseteq \mathbb{F}[X]$, and a *solution* for \mathcal{F} is an assignment $s : X \rightarrow \mathbb{F}$ such that $f(s(\vec{x})) = 0$ for each $f(\vec{x}) \in \mathcal{F}$. We say that \mathcal{F} is satisfiable if it has a solution, and unsatisfiable otherwise. If we wish to encode Boolean satisfiability problems, where the assignments are in $\{0, 1\}$ only, then we include in \mathcal{F} the *Boolean axioms* $x^2 - x$ for each $x \in X$.

For a system of polynomial equations \mathcal{F} , the *degree* $\text{deg}(\mathcal{F})$ of \mathcal{F} is the maximum degree of any polynomial in \mathcal{F} . The *size* $|\mathcal{F}|$ of \mathcal{F} is the total number of variables and equations, i.e. if $\mathcal{F} = \{f_j \in \mathbb{F}[X] \mid 1 \leq j \leq m\}$, then $|\mathcal{F}| := m + |X|$.

2.1 (Symmetric) Algebraic Circuits

We study the complexity of a polynomial not only in terms of its degree or its number of monomials but also via the size of the smallest algebraic circuit representing it.

► **Definition 1** (Algebraic circuit). *An algebraic circuit over variables X and a field \mathbb{F} is a connected directed acyclic graph $C = (D, W)$ with a labelling $\lambda : D \rightarrow \{+, \times\} \cup X \cup \mathbb{F}$ such that $\lambda(g) \in X \cup \mathbb{F}$ if g has in-degree 0, and $\lambda(g) \in \{+, \times\}$, otherwise.*

The nodes in D are called *gates* and the edges in W are called *wires*. Gates with in-degree 0 are called *input gates* and gates with out-degree 0 are called *output gates*. Unless stated otherwise, there is only one output gate. Every gate that is not an input gate is an *internal gate*. The *size* of a circuit C is $|C| := |D|$.

We view an algebraic circuit C over X and \mathbb{F} as computing a polynomial $C(\vec{x}) \in \mathbb{F}[X]$ by evaluating gates to polynomials and treating $+, \times$ as ring operations. We write C for the circuit and $C(\vec{x})$ for the polynomial it computes.

The *semantic degree* of an algebraic circuit is the maximum degree of a polynomial computed by any of its subcircuits: $\text{deg}(C) := \max(\{\text{deg}(C_g(\vec{x})) \mid g \in D\})$.

We are mostly concerned with *symmetric* algebraic circuits, as they are studied for example in [15]. If Γ is a group acting on a variable set X , then a circuit C over X is Γ -*symmetric* if we can extend the action of Γ on the input gates to automorphisms of the whole circuit. An automorphism of a circuit is a permutation of the gates that preserves its structure as a DAG, as well as all labels. Formally, $\pi \in \Gamma$ *extends* to an automorphism of $C = (D, W)$ if there is a $\sigma \in \mathbf{Sym}(D)$ such that σ is an automorphism of the DAG C that preserves labels of internal gates, and for every input gate $g \in D$ with label $\lambda(g) \in X \cup \mathbb{F}$, we have $\lambda(\sigma(g)) = \pi(\lambda(g))$. Here, the action of Γ on \mathbb{F} is trivial, i.e. if $\lambda(g) \in \mathbb{F}$, then $\pi(\lambda(g)) = \lambda(g)$. It is easy to check that the polynomial computed by a Γ -symmetric algebraic circuit is itself invariant under the action of Γ on its variables.

2.2 Logics from finite model theory

Fixed-point logic with counting

Formulas of fixed-point logic with counting (FPC) are evaluated in finite structures expanded with a numeric sort as the domain for counting terms. For every finite structure \mathfrak{A} , let $\mathfrak{A}^* := \mathfrak{A} \uplus (\mathbb{N}, <, +, \cdot, 0, e)$, where e is interpreted as $|A|$. The variables in FPC are typed so that each variable either has as its domain the point sort (i.e. \mathfrak{A}) or the numeric sort of \mathfrak{A}^* . Point variables are denoted with Roman letters and numeric variables with Greek letters.

The syntax of FPC is that of first-order logic with the following extensions:

- **Quantifiers over numeric variables:** If μ is a numeric variable, then quantification over μ is only allowed in the form $Q\mu < t.\varphi$, where $Q \in \{\exists, \forall\}$ and t is a numeric term. This ensures that every numeric variable has a fixed range of polynomial size (with respect to the size of the point sort) – otherwise, it would not be possible to evaluate FPC-formulas in polynomial time.
- **Counting quantifiers:** If $\varphi(\vec{x}, \vec{\mu})$ is an FPC-formula, y a point variable and t a numeric term, then $\exists^{\geq t} y \varphi(\vec{x}, \vec{\mu})$ is a formula which is true in \mathfrak{A}^* if the number of satisfying assignments to y is at least the value of t .
- **Fixed-point operators:** Let Z be a second-order variable, φ an FPC-formula in which Z occurs only positively, and t a numeric term. Then $[\mathbf{lfp} \ Z \vec{x} \vec{\mu} < t. \varphi(Z, \vec{x}, \vec{\mu})](\vec{x}, \vec{\mu})$ is a formula of FPC. It is satisfied in $(\mathfrak{A}^*, \vec{x} \vec{\mu} \mapsto \vec{a} \vec{b})$ if $\vec{a} \vec{b}$ is in the least fixed-point of the sequence defined by $Z_0 = \emptyset, Z_{i+1} = \{\vec{c} \vec{d} \mid \mathfrak{A}^* \models \varphi(Z_i, \vec{c}, \vec{d})\}$.

There also exists the dual operator $[\mathbf{gfp} \ Z \vec{x} \vec{\mu} < t. \varphi(Z, \vec{x}, \vec{\mu})](\vec{x}, \vec{\mu})$, which computes the greatest fixed-point.

This is one way to present FPC – there are other equivalent presentations, for example where the numeric sort in \mathfrak{A}^* is finite, or where counting *terms* instead of counting *quantifiers* are used. For more details and background on FPC, we refer to [26] or [8].

Choiceless Polynomial Time

Choiceless Polynomial Time (with counting) can be viewed as an extension of FPC with higher-order data structures. It properly extends the expressive power of FPC and it remains an open question whether it can express all polynomial-time decidable properties of finite structures. It seeks to overcome the limitation of FPC that the stages of a fixed-point computation are relations of a fixed-arity. In CPT, the computation stages are much more expressive objects, namely *hereditarily finite sets*. These are arbitrarily nested finite sets whose atoms are elements of the respective input structure. CPT has an iteration mechanism with definable state-updates similar to FPC, with the difference that in each step, new

hereditarily finite sets may be constructed. Thus, computations are not guaranteed to reach a fixed-point. To guarantee polynomial time evaluation, every CPT-sentence Π comes with an explicit polynomial bound $p(n): \mathbb{N} \rightarrow \mathbb{N}$. On an input structure \mathfrak{A} , the evaluation of Π is aborted if it takes longer than $p(|\mathfrak{A}|)$ many steps. Likewise, the constructed h.f. sets are required to have size at most $p(|\mathfrak{A}|)$. In this article, we write $\text{CPT}(p(n))$ for the fragment of CPT consisting of sentences whose explicit polynomial bound is $p(n)$. Since CPT is a logic and the state-updates are definable, the h.f. sets constructed during a computation are naturally symmetric under the automorphisms of the input structure.

We refrain from giving a formal definition of CPT, since the details are not needed here and it would be quite lengthy. A concise survey can be found in [19]. The work that originally introduced CPT as an abstract state machine model is [4]; for a more “logic-like” presentation of CPT, see *BGS-logic* [31].

3 Algebraic Proof Systems

Algebraic proof systems are formalisms for refuting the satisfiability of polynomial equation systems. A refutation is a certificate of unsatisfiability checkable in deterministic or randomized polynomial time. Usually, the certificate is based on Hilbert’s Nullstellensatz: It states that for any polynomial system of equations over an algebraically closed field \mathbb{F} , $\mathcal{F} \subseteq \mathbb{F}[X]$, \mathcal{F} is unsatisfiable if, and only if, 1 is in the ideal generated by \mathcal{F} , which means that there are $g_1, g_2, \dots, g_m \in \mathbb{F}[X]$ such that $\sum_{i \leq m} f_i g_i = 1$. Thus, a refutation of \mathcal{F} in an algebraic proof system is typically a systematic proof of the existence of such g_1, g_2, \dots, g_m . The proof systems we are concerned with are (variants of) the *polynomial calculus* (PC) and the *Ideal Proof System* (IPS). In the following, whenever we consider unsatisfiable polynomial systems of equations over a field \mathbb{F} and their refutations, we always implicitly mean the algebraic closure of \mathbb{F} – otherwise the algebraic proof systems are not complete because the Nullstellensatz does not hold.

The efficiency of different proof systems is compared using the standard notion of *p-simulation*: Let P_1, P_2 be two proof systems and μ_1, μ_2 be complexity measures for them, that is, functions that map refutations to natural numbers. Then we write $P_2 \leq_p P_1$ (with respect to the measures μ_1, μ_2 , which will usually be clear from the context) if for every system of polynomial equations \mathcal{F} , and every P_2 -refutation R_2 , there exists a P_1 -refutation R_1 of \mathcal{F} with $\mu_1(R_1) \leq \text{poly}(\mu_2(R_2))$. We write $P_1 \equiv_p P_2$ if $P_2 \leq_p P_1$ and $P_1 \leq_p P_2$.

3.1 Polynomial Calculus

The polynomial calculus [7] predates the IPS as an algebraic proof system. Unlike the IPS, the PC is a more classical rule-based system consisting in a set of sound inference rules for systematically deriving the polynomials in the ideal generated by the input \mathcal{F} .

► **Definition 2** (Polynomial Calculus (PC)). *Let $\mathcal{F} \subseteq \mathbb{F}[X]$ be a system of polynomial equations over a field \mathbb{F} with variables in X .*

The inference rules of the polynomial calculus are:

- *Axiom:* $\frac{f}{f}$ for all $f \in \mathcal{F}$.
- *Multiplication:* $\frac{f}{xf}$ for any $f \in \mathbb{F}[X], x \in X$.
- *Linear Combination:* $\frac{f}{af + bg}$ for any $f, g \in \mathbb{F}[X], a, b \in \mathbb{F}$.

A PC refutation of \mathcal{F} is a sequence $(p_1, p_2, \dots, p_n = 1)$ of polynomials such that p_n is the 1-polynomial, and each p_i is either a Boolean axiom, an axiom from \mathcal{F} , or is the result of the application of a proof rule to one or two polynomials $p_j, p_{j'}$ with $j, j' < i$.

The PC is a sound and complete proof system, that is, a polynomial equation system \mathcal{F} is unsatisfiable if, and only if, there exists a refutation for \mathcal{F} . The complexity of this refutation may in general be super-polynomial.

We consider two complexity measures for the PC. The *number of lines* of a PC refutation $R = (p_1, p_2, \dots, p_n = 1)$ is n , i.e. the length of the derivation sequence. The *degree* $\deg(R)$ of R is the maximum degree of any of the p_j , for $1 \leq j \leq n$. For $k \in \mathbb{N}$, we denote by \deg_k -PC the restriction of the PC where only refutations of degree $\leq k$ are allowed. This is no longer a complete proof system; for instance, it cannot prove the pigeon hole principle [30].

3.2 Ideal Proof System

The Ideal Proof System (IPS) introduced by Grochow and Pitassi [21] is a more general formalism for proving in a verifiable way that 1 is in the ideal generated by a set \mathcal{F} of axioms. In this proof system, there exist no derivation rules, just a certificate. This takes the form of an algebraic circuit. IPS in particular p -simulates the polynomial calculus [21, Proposition 3.4].

► **Definition 3** (Ideal Proof System (IPS), [21]). *Let \mathcal{F} be the polynomial equation system $\{f_1(\vec{x}), f_2(\vec{x}), \dots, f_m(\vec{x})\}$ in $\mathbb{F}[X]$ and let $\overline{\mathbb{F}}$ be the algebraic closure of \mathbb{F} . Let $Y = \{y_1, y_2, \dots, y_m\}$ be a fresh set of variables. An IPS certificate of unsatisfiability of \mathcal{F} over $\overline{\mathbb{F}}$ is a polynomial $C(\vec{x}, \vec{y}) \in \mathbb{F}[X \uplus Y]$ such that (1) $C(\vec{x}, \vec{0}) = 0$, and (2) $C(\vec{x}, \vec{f}) = 1$. An IPS proof of unsatisfiability (i.e. a refutation) of \mathcal{F} is an algebraic circuit C with variables $X \uplus Y$ and constants \mathbb{F} computing an IPS-certificate of unsatisfiability.*

Condition (1) ensures that $C(\vec{x}, \vec{f})$ is in the ideal generated by \mathcal{F} . The unrestricted IPS is sound and complete, i.e. there is an IPS refutation of \mathcal{F} if, and only if, \mathcal{F} is unsatisfiable [21]. Note that this is not necessarily a proof system in the sense of Cook and Reckhow as it is not clear that certificates can be verified in polynomial time. Verifying that a given circuit indeed computes a valid certificate requires polynomial identity testing (PIT), which is in randomized polynomial time (but not known to be in P).

We define the *size* $|C|$ of an IPS proof $C = (D, W)$ as $|C| := \min(|D|, |X| + |Y|)$. That is, we define the size of a refutation to be at least the instance size. This would be inadequate for sublinear size refutations but in the symmetric setting that we study here, the smallest possible proof size is generally $|X| + |Y|$.

For families $(\mathcal{F}_n)_{(n \in \mathbb{N})}$ of instances and corresponding refutations $(C_n)_{(n \in \mathbb{N})}$, we are mainly interested in IPS refutations of *polynomial size* $|C_n| \leq p(|\mathcal{F}_n|)$. When we speak of the complexity of IPS proofs, we usually mean this complexity measure.

For every $k \in \mathbb{N}$, \deg_k -IPS denotes the restriction of the IPS in which the semantic degree $\deg(C)$ (see Section 2.1) of any refutation C is at most k .

Other natural restrictions of the IPS are obtained by allowing only circuits from certain circuit classes, such as bounded depth, or as in our case, symmetric circuits. By the Nullstellensatz, there always exists a \vec{y} -linear (also called *Hilbert-like* in [21]) IPS refutation $C(\vec{x}, \vec{y})$ for every unsatisfiable \mathcal{F} , i.e. $C(\vec{x}, \vec{y}) = \sum_{i=1}^m y_i g_i(\vec{x})$ for some $\vec{g} \in \mathbb{F}[X]$. In other words, the \vec{y} -linear IPS is a sound and complete fragment of the general IPS. We denote it IPS_{LIN} . It is shown in [21] that IPS_{LIN} p -simulates the general IPS on instances that are themselves representable with polynomial-size algebraic circuits. For fragments of the IPS with restricted circuit classes, it is however not clear that every proof can efficiently be simulated by a \vec{y} -linear one. Indeed, for symmetric proofs, as we show, IPS_{LIN} turns out to be a true restriction.

The unrestricted IPS is remarkably powerful: It is shown in [21, Theorem 3.5] that IPS over any field of characteristic q p -simulates any Frege proof system with MOD_q -connectives. This is even true for Extended Frege, where extension axioms may be used as in the Extended Polynomial Calculus. Frege systems are standard textbook propositional proof systems such as the sequent calculus.

4 Symmetric IPS Proofs

We now introduce the symmetry restriction on IPS proofs. That is to say, we consider when a set \mathcal{F} of polynomials that is Γ -invariant for a group Γ acting on its variables, admits a Γ -symmetric circuit as an IPS refutation. We start with a formal definition.

► **Definition 4** (Symmetric IPS). *Let Γ be a group and X a Γ -set of variables. Let $\mathcal{F} = \{f_1(\vec{x}), \dots, f_m(\vec{x})\} \subseteq \mathbb{F}[X]$ be a Γ -invariant set of polynomials and let $Y = \{y_1, \dots, y_m\}$ be a Γ -set of variables with the following action: For every $\pi \in \Gamma$ and $i \in [m]$, $\pi(y_i) = y_j$ for the j with $\pi(f_i) = f_j$.*

A Γ -symmetric IPS proof of unsatisfiability of \mathcal{F} is a Γ -symmetric algebraic circuit with variables $X \uplus Y$ computing an IPS certificate $C(\vec{x}, \vec{y})$ of \mathcal{F} .

Note that if $\Gamma = \{\text{id}\}$, then any IPS refutation of \mathcal{F} is also a Γ -symmetric IPS refutation of \mathcal{F} . The complexity of a symmetric IPS refutation in general heavily depends on the choice of Γ . The bigger Γ is and hence the more symmetric the circuits are required to be, the bigger we can expect the proof size to be.

We generally write Γ -sym-IPS to mean the proof system allowing only Γ -symmetric proofs. Of course, this only makes sense for sets \mathcal{F} of polynomials that are Γ -invariant. For these, as we show below, Γ -sym-IPS is a complete proof system. Where Γ is clear from context, we may write just sym-IPS.

Let Γ and Γ' be groups acting on X with $\Gamma \leq \Gamma'$. We say that Γ' -sym-IPS p -simulates Γ -sym-IPS if for every Γ' -invariant \mathcal{F} and a Γ -sym-IPS refutation C of \mathcal{F} , there is a Γ' -sym-IPS refutation of \mathcal{F} with size polynomial in the size of C . We also use this notion of p -simulation in the context of restricted proof systems, such as linear or bounded degree systems. For instance, in Theorem 7 below, we show that (in suitably defined cases) $\deg_k\text{-IPS} \leq_p \deg_k\text{-sym-IPS}$. This means that for each k , there is a fixed polynomial p such that whenever a Γ -invariant \mathcal{F} has any (i.e. $\{\text{id}\}$ -symmetric) IPS proof of size n and degree k , it also has a Γ -symmetric IPS proof of size $p(n)$ and degree k .

4.1 Completeness of symmetric IPS

A first natural question that arises when we restrict ourselves to symmetric proofs is whether it is the case that every Γ -symmetric collection \mathcal{F} of polynomials has a Γ -symmetric proof. We show that this is indeed the case, in the sense that any IPS proof can be suitably symmetrised, though this may entail an exponential blowup in the size of the proof. On the other hand, there are \mathcal{F} for which there is no symmetric *linear* proof, which contrasts with the fact that linear IPS (without symmetry requirements) is complete.

We first prove the completeness of the symmetric IPS and then provide a simple counter-example for completeness for sym-IPS_{LIN}.

► **Theorem 5.** *Let \mathcal{F} be a Γ -symmetric system of polynomial equations. If \mathcal{F} is unsatisfiable, then there is a Γ -symmetric IPS refutation of \mathcal{F} .*

Proof. Since IPS is complete, there is a certificate $C(\vec{x}, \vec{y})$ of unsatisfiability of \mathcal{F} , computed by some algebraic circuit C . We construct a Γ -symmetric circuit C^{sym} with the same semantics: For every $\pi \in \Gamma$, we introduce a copy of $\pi(C)$ in such a way that all these $\pi(C)$, for all $\pi \in \Gamma$, are identified at their input gates and otherwise disjoint. To finish the construction, we add a multiplication gate g_\times as the output of C^{sym} . It multiplies the outputs of all circuits $\pi(C)$, for all $\pi \in \Gamma$. The resulting circuit C^{sym} is Γ -symmetric by construction because every $\pi \in \Gamma$ extends to a circuit automorphism that maps each subcircuit $\pi'(C)$ to $(\pi \circ \pi')(C)$. We can also verify that C^{sym} is again a refutation:

$$C^{sym}(\vec{x}, \vec{0}) = \prod_{\pi \in \Gamma} \pi(C(\vec{x}, \vec{0})) = \prod_{\pi \in \Gamma} \underbrace{C(\pi\vec{x}, \pi\vec{0})}_{=0} = 0. \quad (1)$$

$$C^{sym}(\vec{x}, \vec{f}) = \prod_{\pi \in \Gamma} \pi(C(\vec{x}, \vec{f})) = \prod_{\pi \in \Gamma} \underbrace{C(\pi\vec{x}, \pi\vec{f})}_{=1} = 1. \quad (2)$$

In the last equality of (1) and (2) we used that 0 and 1 are Γ -symmetric polynomials, so if $C(\vec{x}, \vec{0}) = 0$, then also $C(\pi\vec{x}, \pi\vec{0}) = 0$ for every $\pi \in \Gamma$ (and likewise for 1). In the penultimate equality of (2), we also use the Γ -invariance of \mathcal{F} and the fact that the action of Γ on Y is exactly as given by the lift of its action on X to \mathcal{F} . ◀

This relatively naive construction blows up the size of the circuit by a factor of $|\Gamma|$, which may be as large as $|X|! \cdot |Y|!$. So even though there always exists a symmetric refutation, this may in general be much larger than the smallest asymmetric refutation.

The following example shows that there are cases where symmetric linear refutations do not exist, regardless of the circuit size.

► **Example 6.** Let the variable set be $X = \{x_1, x_1^*, x_2, x_2^*\}$ and let $\Gamma \leq \mathbf{Sym}(X)$ be the group that is generated by $\{\pi, \pi^*\}$ defined as follows: $\pi = (x_1 \ x_2) \circ (x_1^* \ x_2^*)$, and $\pi^* = (x_1 \ x_1^*) \circ (x_2 \ x_2^*)$. Note that this is the Klein 4 group. That is, π exchanges 1 and 2, and π^* exchanges non-star with star. Consider the following equations over \mathbb{F}_2 :

$$\begin{aligned} (1) \ x_1 + x_2 &= 1 & (3) \ x_1^* + x_2 &= 1 & (5) \ x_1 + x_1^* &= 1 \\ (2) \ x_1^* + x_2^* &= 1 & (4) \ x_1 + x_2^* &= 1 & (6) \ x_2 + x_2^* &= 1 \end{aligned}$$

The equations are partitioned into three Γ -orbits. Equation (1) and (2) form an orbit, equation (3) and (4) as well, and equation (5) and (6). The system is unsatisfiable over \mathbb{F}_2 (and its algebraic closure) because (1) + (4) + (6) is an IPS certificate of unsatisfiability. Nonetheless, there is no linear Γ -symmetric refutation $C(\vec{x}, \vec{y})$: Every monomial in such a certificate $C(\vec{x}, \vec{y})$ would contain exactly one \vec{y} -variable. In order to have $C(\vec{x}, \vec{f}) = 1$, there must be a degree-1 monomial in $C(\vec{x}, \vec{y})$, i.e. consisting only of a single \vec{y} -variable y_i . But then, the entire orbit of y_i must appear in $C(\vec{x}, \vec{y})$ due to symmetry. As each orbit of equations has even size and the field has characteristic 2 the constant terms in $C(\vec{x}, \vec{f})$ sum up to 0. Hence, there is no symmetric \vec{y} -linear polynomial with $C(\vec{x}, \vec{f}) = 1$.

On the other hand, we can show that as long we are working over a field \mathbb{F} of characteristic either 0 or coprime with the order of Γ , then any Γ -invariant set \mathcal{F} of polynomials has a Γ -symmetric linear refutation. This follows from Corollary 9 below.

4.2 Symmetry in Bounded-Degree IPS

Recall that for any constant $k \in \mathbb{N}$, $\text{deg}_k\text{-IPS}$ is the restriction of IPS where the polynomials computed at each gate have degree at most k . This bounded-degree fragment is of special interest because of its relation to the bounded-degree polynomial calculus, whose expressive power is related to important logics from finite model theory and also to the well-known

40:10 Symmetric Proofs in the Ideal Proof System

Weisfeiler-Leman graph isomorphism algorithm (see Section 5). We show that constant-degree IPS proofs can be symmetrised efficiently, under certain assumptions on the field and the symmetry group (which are satisfied in most interesting cases). Thus, in the bounded-degree regime, requiring proofs to be symmetric is essentially no restriction.

► **Theorem 7.** *Let \mathbb{F} be a field and let Γ be a group such that either $\text{char}(\mathbb{F}) = 0$, or \mathbb{F} has positive characteristic and $|\Gamma|$ and $\text{char}(\mathbb{F})$ are coprime. Let $k \in \mathbb{N}$ be a constant. Then for every Γ -invariant polynomial equation system $\mathcal{F} \subseteq \mathbb{F}[X]$ that possesses a \deg_k -IPS refutation C , there also exists a Γ -symmetric refutation C^{sym} with $|C^{\text{sym}}| \leq \mathcal{O}(|\mathcal{F}|^k) \leq \mathcal{O}(|C|^k)$. If C is \vec{y} -linear, then so is C^{sym} .*

Proof. Let $C(\vec{x}, \vec{y})$ be a certificate of unsatisfiability of the polynomial equation system \mathcal{F} , computed by a circuit with semantic degree k . Let $M \subseteq \mathbb{F}[X \cup Y]$ be the set of monomials appearing in $C(\vec{x}, \vec{y})$. Then the certificate $C(\vec{x}, \vec{y})$ can be written as $\sum_{m \in M} c_m \cdot m(\vec{x}, \vec{y})$, where $c_m \in \mathbb{F}$ and every $m \in M$ has degree at most k . We define $C^{\text{sym}}(\vec{x}, \vec{y}) := |\Gamma|^{-1} \cdot \sum_{\pi \in \Gamma} \pi C(\vec{x}, \vec{y})$. This polynomial $C^{\text{sym}}(\vec{x}, \vec{y})$ is also an IPS certificate of unsatisfiability of \mathcal{F} since

$$(1) \quad C^{\text{sym}}(\vec{x}, \vec{0}) = |\Gamma|^{-1} \cdot \sum_{\pi \in \Gamma} \pi C(\vec{x}, \vec{0}) = 0, \text{ and}$$

$$(2) \quad C^{\text{sym}}(\vec{x}, \vec{f}) = |\Gamma|^{-1} \cdot \sum_{\pi \in \Gamma} \pi C(\vec{x}, \vec{f}) = |\Gamma|^{-1} \cdot |\Gamma| \cdot 1 = 1.$$

Note that $|\Gamma|^{-1}$ is defined in \mathbb{F} since either the characteristic of \mathbb{F} is zero, or it is coprime with $|\Gamma|$. Let ΓM be the closure of M under the action of Γ , i.e. $\Gamma M := \bigcup_{\pi \in \Gamma} \pi(M)$. Now $C^{\text{sym}}(\vec{x}, \vec{y})$ has the form

$$C^{\text{sym}}(\vec{x}, \vec{y}) = |\Gamma|^{-1} \sum_{\pi \in \Gamma} \sum_{m \in M} c_m \cdot \pi(m(\vec{x}, \vec{y})) = |\Gamma|^{-1} \sum_{m \in \Gamma M} \left(\sum_{\pi \in \Gamma} c_{\pi^{-1}(m)} \right) \cdot m$$

The number of distinct monomials in $C^{\text{sym}}(\vec{x}, \vec{y})$ is at most $\mathcal{O}((|X| + |Y| + 1)^k)$ because all monomials have degree at most k . Thus, the circuit that just expresses $C^{\text{sym}}(\vec{x}, \vec{y})$ as a sum of monomials has polynomial size, and it is also symmetric because $C^{\text{sym}}(\vec{x}, \vec{y})$ is symmetric by construction. By our convention (see Section 3.2), $|C| \geq |X| + |Y|$, so $|C^{\text{sym}}| \leq \text{poly}(|C|)$. If C is \vec{y} -linear, then so is C^{sym} because if M contains only \vec{y} -linear monomials, then this is still true for ΓM . ◀

► **Corollary 8.** *Let \mathbb{F} be a field and let Γ be a group such that either $\text{char}(\mathbb{F}) = 0$, or \mathbb{F} has positive characteristic and $|\Gamma|$ and $\text{char}(\mathbb{F})$ are coprime. Let $k \in \mathbb{N}$ be a constant. Then for this field and symmetry group, $\deg_k\text{-PC} \leq_p \deg_k\text{-sym-IPS}_{\text{LIN}}$.*

Proof. Any degree- k PC refutation can be translated in a straightforward way into a degree- k IPS_{LIN} -refutation (see [21, Proposition 3.4]). This can be efficiently Γ -symmetrised using Theorem 7. ◀

► **Corollary 9.** *Let \mathbb{F} be a field and let Γ be a group such that either $\text{char}(\mathbb{F}) = 0$, or \mathbb{F} has positive characteristic and $|\Gamma|$ and $\text{char}(\mathbb{F})$ are coprime. Then, any Γ -invariant unsatisfiable set of polynomials \mathcal{F} over this field has a Γ -symmetric linear refutation.*

Proof. If \mathcal{F} is an unsatisfiable polynomial equation system, then it has a PC refutation by the completeness of polynomial calculus. Let k be its degree. By Corollary 8, there exists a Γ - \deg_k -sym- IPS_{LIN} -refutation. ◀

5 Applications in Finite Model Theory and Graph Isomorphism

We now turn to other well-studied symmetry-invariant formalisms and show that they are subsumed by the symmetric IPS in a certain sense. These formalisms are logics from finite model theory, specifically *fixed-point logic with counting* (FPC) and *Choiceless Polynomial Time* (CPT). Drawing on previous work [29, 27], we show how their expressive power relates to the power of sym-IPS, specifically with respect to the graph isomorphism problem.

5.1 Simulating fixed-point logic with counting in sym-IPS

The evaluation of any sentence ψ in fixed-point logic with counting in a given finite structure \mathfrak{A} can be simulated by a bounded-degree symmetric IPS proof over \mathbb{Q} in the following sense. For every fixed FPC-sentence ψ and structure \mathfrak{A} , there is an axiom system $\mathcal{F}_\psi(\mathfrak{A})$ that expresses the existence of a winning strategy in the model-checking game for $\mathfrak{A} \models \psi$. This is an instance of a *threshold safety game* [20]. An IPS refutation of $\mathcal{F}_\psi(\mathfrak{A})$ is then a witness for the fact that $\mathfrak{A} \models \psi$. The axiom system $\mathcal{F}_\psi(\mathfrak{A})$ is FOC-interpretable in \mathfrak{A} , meaning that it is FO-definable in \mathfrak{A} extended with a numeric sort (see Section 2.2). In total, the problem of deciding the existence of a bounded-degree sym-IPS refutation for a given axiom system is complete for FPC under efficient symmetry-preserving reductions:

► **Theorem 10.** *For every FPC-sentence ψ with signature τ , there exists an FOC-definable mapping \mathcal{F}_ψ that takes every finite τ -structure \mathfrak{A} to a polynomial equation system $\mathcal{F}_\psi(\mathfrak{A})$ over \mathbb{Q} such that:*

1. $|\mathcal{F}_\psi(\mathfrak{A})| \leq \text{poly}(|\mathfrak{A}|)$.
2. $\text{Aut}(\mathfrak{A})$ has a natural action on the variables of $\mathcal{F}_\psi(\mathfrak{A})$, and $\mathcal{F}_\psi(\mathfrak{A})$ is $\text{Aut}(\mathfrak{A})$ -invariant.
3. $\mathcal{F}_\psi(\mathfrak{A})$ is unsatisfiable if, and only if, $\mathfrak{A} \models \psi$.
4. If $\mathcal{F}_\psi(\mathfrak{A})$ is unsatisfiable, then $\mathcal{F}_\psi(\mathfrak{A})$ has an $\text{Aut}(\mathfrak{A})$ -symmetric deg_2 -IPS_{LIN} refutation of size $\text{poly}(|\mathcal{F}_\psi(\mathfrak{A})|)$.

This is mainly a consequence of Theorem 4.4 in [29]. There, the desired mapping \mathcal{F}_ψ is constructed and it is shown that $\mathcal{F}_\psi(\mathfrak{A})$ has a degree-2 PC refutation R over \mathbb{Q} if and only if $\mathfrak{A} \models \psi$. By Corollary 8, there also exists a $\text{Aut}(\mathfrak{A})$ -symmetric deg_2 -IPS_{LIN}-refutation of size $\text{poly}(|R|)$, and even of size $\text{poly}(|\mathcal{F}_\psi(\mathfrak{A})|)$ (see Theorem 7). This proves Item 4 from the theorem. The first two items follow from the properties of FOC-interpretations and the third item is due to the construction in [29].

5.2 Symmetric IPS proofs of graph non-isomorphism

Now we pass on from FPC to a stronger logic, namely *Choiceless Polynomial Time* (CPT) and show that this, too, can in a certain sense be simulated in sym-IPS_{LIN}. The greater model-theoretic expressiveness of CPT is reflected on the proof system side in the fact that we (provably) need to go beyond the bounded-degree regime. Another difference to the simulation of FPC is that we now consider a fixed problem, namely *graph isomorphism*. We show that sym-IPS_{LIN} efficiently distinguishes all graphs (using their natural symmetries as the symmetry group) that are also distinguishable in CPT.

Distinguishing graphs in an algebraic proof system

Whenever we speak of *graphs* in this section, they may be vertex and edge coloured. Given two graphs G and H , there is a standard system of polynomial equations $\mathcal{F}_{\text{iso}}(G, H)$ [3, 27] whose solutions encode isomorphisms between G and H . Thus, G and H are non-isomorphic

if, and only if, $\mathcal{F}_{\text{iso}}(G, H)$ is unsatisfiable. The variable set of $\mathcal{F}_{\text{iso}}(G, H)$ is $X := \{x_{vw} \mid v \in V(G), w \in V(H), v \sim w\}$, where $\sim \subseteq V(G) \times V(H)$ is the relation that contains all (v, w) such that v and w have the same colour. The polynomials of $\mathcal{F}_{\text{iso}}(G, H)$ are:

$$\begin{aligned} & \sum_{v \in V(G): v \sim w} x_{vw} - 1 && \text{for each } w \in V(H) \\ & \sum_{w \in V(H): v \sim w} x_{vw} - 1 && \text{for each } v \in V(G) \text{ and} \\ & x_{vw}x_{v'w'} && \text{for all } v, v' \in V(G), w, w' \in V(H) \text{ with } v \sim w, v' \sim w' \\ & && \text{such that } vv' \mapsto ww' \text{ is not a local isomorphism.} \end{aligned}$$

The Boolean axioms $x^2 - x$ for all $x \in X$ are also part of $\mathcal{F}_{\text{iso}}(G, H)$. The idea behind this formulation is that a satisfying Boolean assignment to the variables in X encodes an isomorphism from G to H in the sense that $v \in V(G)$ is mapped to $V(H)$ if, and only if, x_{vw} is assigned to 1. Note that $\mathbf{Aut}(G) \times \mathbf{Aut}(H)$ acts naturally on X : Let $(\pi_G, \pi_H) \in \mathbf{Aut}(G) \times \mathbf{Aut}(H)$. Then $(\pi_G, \pi_H)(x_{vw}) = x_{\pi_G(v)\pi_H(w)}$. It is not hard to see that $\mathcal{F}_{\text{iso}}(G, H)$ is invariant under this group action: The polynomials associated with vertices in G and H are clearly symmetric and the polynomials $x_{vw}x_{v'w'}$ that forbid local non-isomorphisms depend on the edges and non-edges, which are preserved by $\mathbf{Aut}(G) \times \mathbf{Aut}(H)$.

When we say that an algebraic proof system distinguishes two graphs G and H , we mean that it admits a refutation of $\mathcal{F}_{\text{iso}}(G, H)$. If \mathcal{K} is a class of graphs, then we say that $\text{sym-IPS}_{\text{LIN}}$ *efficiently distinguishes* all non-isomorphic graphs in \mathcal{K} if there exists a polynomial $p(n)$ such that for any two non-isomorphic $G, H \in \mathcal{K}$, there exists an $\mathbf{Aut}(G) \times \mathbf{Aut}(H)$ -symmetric IPS_{LIN} -refutation C of $\mathcal{F}_{\text{iso}}(G, H)$ of size $|C| \leq p(|\mathcal{F}_{\text{iso}}(G, H)|)$ over the field \mathbb{Q} .

Distinguishing graphs in Choiceless Polynomial Time

For any pair of non-isomorphic graphs G and H , there is some formula (say of first-order logic) that distinguishes them. For a class of graphs, we are interested in obtaining bounds (say on the number of variables or other parameters) of the minimum distinguishing formulas for pairs of non-isomorphic graphs from the class. Here we are particularly interested in the number of variables of an FPC formula, or the resource bounds of a CPT sentence. And, we relate these to bounds on the IPS refutation of $\mathcal{F}_{\text{iso}}(G, H)$.

► **Definition 11** (Distinguishing graphs in CPT, [27]). *Let \mathcal{K} be a class of graphs. We say that CPT distinguishes all graphs in \mathcal{K} if there exists a polynomial $p(n)$ and a constant $k \in \mathbb{N}$ such that for any two non-isomorphic $G, H \in \mathcal{K}$, there exists a sentence $\Pi \in \text{CPT}(p(n))$ with $\leq k$ variables such that $G \models \Pi$ and $H \not\models \Pi$.*

Recall from Section 2.2 that $\text{CPT}(p(n))$ is the fragment of CPT whose sentences have resource bound at most $p(n)$. Similarly, the k -variable fragment of FPC distinguishes all graphs in \mathcal{K} if there exists a distinguishing FPC-sentence with $\leq k$ variables for all $G \not\cong H$ in \mathcal{K} .

► **Theorem 12.** *Let \mathcal{K} be a graph class.*

1. *If there is a $k \in \mathbb{N}$ such that the k -variable fragment of FPC distinguishes all non-isomorphic graphs in \mathcal{K} , then so does symmetric \deg_k -IPS.*
2. *If CPT distinguishes all non-isomorphic graphs in \mathcal{K} , then $\text{sym-IPS}_{\text{LIN}}$ efficiently distinguishes them.*

Note that the situation in the first statement is equivalent to the non-isomorphic graphs in \mathcal{K} being distinguishable by the well-known k -dimensional Weisfeiler-Leman algorithm [24]. The first part then follows from [3, Theorem 4.4], which states that k -Weisfeiler-Leman-distinguishable graphs can also be distinguished in the degree- k PC. This can be p -simulated

by \deg_k -sym-IPS by Corollary 8. To prove the second part, we use Theorem 1 from [27]. It shows that if all non-isomorphic graphs in \mathcal{K} are CPT-distinguishable in the sense of Definition 11, then they are also distinguishable in the degree-3 *extended polynomial calculus* (EPC), and the refutations have polynomial size. As discussed in the conclusion of [27], this EPC refutation is symmetric in the sense that its extension axioms are closed under the action of $\mathbf{Aut}(G) \times \mathbf{Aut}(H)$. To conclude the second item of Theorem 12, we show that any symmetric bounded-degree EPC refutation can be simulated efficiently in $\text{sym-IPS}_{\text{LIN}}$. The proof can be found in the full version [11, Theorem 22].

Another result from [27] immediately gives us the following separation between the bounded-degree (symmetric) IPS and its unbounded version.

► **Theorem 13.** *There exists a sequence $(G_n, H_n)_{n \in \mathbb{N}}$ of pairs of non-isomorphic graphs such that $\mathcal{F}_{\text{iso}}(G_n, H_n)$ has a polynomial-size $\mathbf{Aut}(G) \times \mathbf{Aut}(H)$ -symmetric IPS_{LIN} -refutation but there is no $k \in \mathbb{N}$ such that for all $n \in \mathbb{N}$, $\mathcal{F}_{\text{iso}}(G_n, H_n)$ has a \deg_k -sym-IPS-refutation.*

Proof. Theorem 3 in [27] yields exactly this statement for the symmetric degree-3 EPC and the degree- k PC. By [11, Theorem 22], the symmetric degree-3 EPC can be p -simulated by $\text{sym-IPS}_{\text{LIN}}$. Now suppose for a contradiction that $\mathcal{F}_{\text{iso}}(G_n, H_n)$ had a degree- k sym-IPS-refutation for some constant k , for all $n \in \mathbb{N}$. This is in particular a degree- k IPS refutation. Using existing results, it can be shown that on instances of constant degree (which we have here), bounded-degree IPS refutations can be simulated in bounded-degree PC. A proof of this is given in Theorem of the full version [11]. This yields a contradiction to [27, Theorem 3], which states that the bounded-degree PC cannot refute $\mathcal{F}_{\text{iso}}(G_n, H_n)$ for all $n \in \mathbb{N}$. ◀

The graphs (G_n, H_n) that are used as hard instances for the bounded-degree IPS are the well-known *Cai-Fürer-Immerman* (CFI) graphs [3, 6], equipped with a linear order on their base graphs. These are standard examples of graphs that are indistinguishable in bounded-variable counting logic and hence FPC. The above theorem tells us that these graphs are in fact efficiently distinguishable in $\text{sym-IPS}_{\text{LIN}}$. This is because there is a CPT-sentence which can tell G_n and H_n apart, for all $n \in \mathbb{N}$ – that sentence uses a sophisticated “circuit-like” construction due to [14]. Via Theorem 12, this translates into a polynomial-size $\text{sym-IPS}_{\text{LIN}}$ -refutation. In the next section, we also study a well-known formulation of the CFI graph isomorphism problem as a system of linear equations and show that also that presentation of the problem admits polynomial-size $\text{sym-IPS}_{\text{LIN}}$ -refutations.

6 Upper bounds

6.1 The Cai-Fürer-Immerman equations

The algebraic formulation of the isomorphism problem of Cai-Fürer-Immerman graphs is the following system of equations over \mathbb{F}_2 (see e.g. [1]).

► **Definition 14** (CFI equations). *Let $G = (V, E)$ be a 3-regular undirected connected graph. Let $u \in V(G)$ be some fixed distinguished vertex and let $a \in \{0, 1\}$. The variable set is $X = \{x_i^e \mid e \in E, i \in \mathbb{F}_2\}$. The equations are*

$$\begin{aligned}
 x_i^e + x_j^f + x_k^g &= i + j + k \pmod{2} && \text{for every } v \in V \setminus \{u\}, \{e, f, g\} = E(v), \\
 &&& \text{and every } i, j, k \in \mathbb{F}_2 \\
 x_i^e + x_j^f + x_k^g &= i + j + k + a \pmod{2} && \text{for vertex } u, \{e, f, g\} = E(u), \\
 &&& \text{and every } i, j, k \in \mathbb{F}_2 \\
 x_0^e + x_1^e &= 1 && \text{for every } e \in E
 \end{aligned}$$

This, together with the Boolean axioms for every variable, is the linear equation system $\mathcal{F}_{\text{CFI}}(G, u, a)$ over \mathbb{F}_2 .

This system is satisfiable if, and only if, $a = 0$ [1]. The typical symmetries that are associated with CFI graphs (over linearly ordered base graphs) are called “edge flips”. With regards to the equation system, this means the following. Let Γ be the subgroup of the Boolean vector space (\mathbb{F}_2^E, \oplus) which consists only of those vectors π such that $\sum_{e \in E(v)} \pi(e) = 0 \pmod 2$ for every $v \in V$. The action of this group on X is given by $\pi(x_i^e) = x_{i+\pi_e}^e$, where addition is in \mathbb{F}_2 , and π_e denotes the entry of π at index $e \in E$. It is easy to check that $\mathcal{F}_{\text{CFI}}(G, u, a)$ is Γ -invariant.

► **Theorem 15.** *Let $(G_n)_{n \in \mathbb{N}}$ be an arbitrary family of 3-regular graphs, and let Γ_n be the subgroup of $\mathbb{F}_2^{E(G_n)}$ defined above. For every $n \in \mathbb{N}$, there exists a non- \tilde{y} -linear Γ_n -symmetric \mathbb{F}_2 -sym-IPS-refutation of the unsatisfiable equation system $\mathcal{F}_{\text{CFI}}(G_n, u_n, 1)$ (regardless of the choice of $u_n \in V(G_n)$), which has size at most $\text{poly}(|\mathcal{F}_{\text{CFI}}(G_n, u_n, 1)|)$.*

In particular, the theorem is true if $(G_n)_{n \in \mathbb{N}}$ is a family of unbounded treewidth. Such families of graphs are the ones for which the equation systems $\mathcal{F}_{\text{CFI}}(G_n, u_n, 1)$ are not distinguishable from their satisfiable counterparts $\mathcal{F}_{\text{CFI}}(G_n, u_n, 0)$ in bounded-variable counting logic [1]. This also means that $\mathcal{F}_{\text{CFI}}(G_n, u_n, 1)$ has no bounded-degree PC refutation: The existence of such a refutation is definable in bounded-variable counting logic [29], and hence, bounded-variable counting logic would be able to distinguish $\mathcal{F}_{\text{CFI}}(G_n, u_n, 1)$ from $\mathcal{F}_{\text{CFI}}(G_n, u_n, 0)$ if bounded-degree PC could refute $\mathcal{F}_{\text{CFI}}(G_n, u_n, 1)$. Thus, this theorem provides another example for the separation of the bounded-degree PC/IPS from the unbounded-degree version.

The construction of the refutation is quite involved, so we have to defer the proof of Theorem 15 to the full version. We just remark that the circuit for the IPS certificate is deeply nested and computes polynomials of linear degree, so it is challenging to ensure that each gate only has a small number of automorphic images under the action of Γ . In fact, we have not been able to accomplish this with a \tilde{y} -linear certificate, so the smallest possible sym-IPS_{LIN}-refutation we know is exponential:

► **Theorem 16.** *Let the setting be as in Theorem 15. For every $n \in \mathbb{N}$, there exists a Γ_n -symmetric \mathbb{F}_2 -sym-IPS_{LIN}-refutation of the unsatisfiable equation system $\mathcal{F}_{\text{CFI}}(G_n, u_n, 1)$, which has size at most $\mathcal{O}(2^{|E_n|})$.*

It is an interesting open question whether the CFI equations provide an exponential separation between sym-IPS and sym-IPS_{LIN}, that is, whether the upper bound in the above theorem can be improved or not. As discussed in the previous section, the graph isomorphism formulation of the CFI equations does admit a small sym-IPS_{LIN}-refutation but it may well be that this is not the case for $\mathcal{F}_{\text{CFI}}(G_n, u_n, 1)$.

6.2 Subset sum

► **Definition 17** (Subset sum, [17]). *Let $n \in \mathbb{N}$, let \mathbb{F} be a field with $\text{char}(\mathbb{F}) > n$, and let $\beta \in \mathbb{F} \setminus \{0, \dots, n\}$. The subset sum instance $\mathcal{F}_{\text{sum}(\vec{x})}(n, \mathbb{F}, \beta)$ has variable set $X = \{x_1, \dots, x_n\}$, and the axiom $\sum_{i=1}^n x_i - \beta = 0$, along with the Boolean axioms $x_i^2 - x_i = 0$ for all $i \in [n]$. The lifted subset sum instance $\mathcal{F}_{\text{sum}(\vec{x}\vec{y})}(n, \mathbb{F}, \beta)$ has variable set $X \cup \{y_1, \dots, y_n\}$ and the axiom $\sum_{i=1}^n x_i y_i - \beta = 0$, along with the Boolean axioms for all variables in $X \cup Y$.*

It is clear that $\mathcal{F}_{\text{sum}(\vec{x})}(n, \mathbb{F}, \beta)$ and $\mathcal{F}_{\text{sum}(\vec{x}\vec{y})}(n, \mathbb{F}, \beta)$ are unsatisfiable for any choice of n, \mathbb{F}, β as required in the definition. Moreover, $\mathcal{F}_{\text{sum}(\vec{x})}(n, \mathbb{F}, \beta)$ is **Sym_n**-symmetric with respect to the obvious action on X , and $\mathcal{F}_{\text{sum}(\vec{x}\vec{y})}(n, \mathbb{F}, \beta)$ is **Sym_n**-symmetric with respect to the

simultaneous action on $X \cup Y$. It is proven in [17, Proposition 5.3] that $\mathcal{F}_{\text{sum}(\vec{x})}(n, \mathbb{F}, \beta)$ has no \deg_k -IPS_{LIN} refutation for any $k < n$. Moreover, [17] shows $\mathcal{F}_{\text{sum}(\vec{x})}(n, \mathbb{F}, \beta)$ to be hard for *sparse IPS* (where circuits are just allowed to be sums of monomials) and $\mathcal{F}_{\text{sum}(\vec{x}\vec{y})}(n, \mathbb{F}, \beta)$ to be hard for roABPs with a fixed variable order, and depth-3 powering formulas.

Subset sum and its liftings are thus a natural starting point in the quest for sym-IPS lower bounds, especially because the variable set is “maximally symmetric” and so it might be expected that the size of any symmetric refutation must be large. However, it turns out that at least for the two subset sum variants we study here, polynomial-size symmetric refutations do in fact exist. There are more complex liftings of the subset axiom that have been used for lower bounds against stronger fragments of IPS such as bounded product-depth circuits [18, 23], and these may be promising candidates for sym-IPS lower bounds, too.

► **Theorem 18.** *The polynomial equation system $\mathcal{F}_{\text{sum}(\vec{x})}(n, \mathbb{F}, \beta)$ has asym-IPS_{LIN}-refutation of size at most $\text{poly}(|\mathcal{F}_{\text{sum}}(n, \mathbb{F}, \beta)|)$, for all n, \mathbb{F}, β such that the system is unsatisfiable. The same is true for $\mathcal{F}_{\text{sum}(\vec{x}\vec{y})}(n, \mathbb{F}, \beta)$. Moreover, there is no constant $k \in \mathbb{N}$ such that \deg_k -IPS can refute $\mathcal{F}_{\text{sum}(\vec{x})}(n, \mathbb{F}, \beta)$ and $\mathcal{F}_{\text{sum}(\vec{x}\vec{y})}(n, \mathbb{F}, \beta)$ for all $n \in \mathbb{N}$.*

The proof of the theorem can be found in the full version. In short, we show that the refutation given in Proposition B.1 in the appendix of [17] can be realised with polynomial-size symmetric circuits. The key ingredient are the elementary symmetric polynomials $S_{n,k} = \sum_{\substack{S \subseteq [n] \\ |S|=k}} \prod_{i \in S} x_i$. They admit efficient symmetric circuits by [33].

6.3 Pigeonhole principle

► **Definition 19.** *For $n, m \in \mathbb{N}$, the n -to- m pigeonhole principle is the polynomial equation system $\mathcal{F}_{\text{PHP}}(n, m)$. Its variable set is $X = \{x_{ij} \mid i \in [n], j \in [m]\}$ and its equations are the Boolean axioms together with:*

$$\begin{aligned} \sum_{j \in [m]} x_{ij} - 1 &= 0 \text{ for every } i \in [n] \\ x_{ij}x_{i'j} &= 0 \text{ for every } j \in [m], i \neq i' \in [n] \end{aligned}$$

Any $\{0, 1\}$ -valued solution to this system gives an injective function from $[n]$ to $[m]$ that maps i to j if $x_{ij} = 1$. The equation $\sum_{j \in [m]} x_{ij} - 1 = 0$ guarantees that each value i is mapped to exactly one j and $x_{ij}x_{i'j} = 0$ ensures that distinct i and i' are not mapped to the same value. Whenever $n > m$, $\mathcal{F}_{\text{PHP}}(n, m)$ is thus unsatisfiable. This is the case over \mathbb{Q} , but also over every finite field. It is easy to see that $\mathcal{F}_{\text{PHP}}(n, m)$ is invariant under $\mathbf{Sym}_n \times \mathbf{Sym}_m$, where the group action on X is: $(\pi, \sigma)(x_{ij}) = x_{\pi(i)\sigma(j)}$. In this section, we focus on the pigeonhole principle $\mathcal{F}_{\text{PHP}}(n+1, n)$. It can be checked that over finite fields, $\mathcal{F}_{\text{PHP}}(n+1, n)$ does not admit a symmetric \vec{y} -linear refutation for all n (analogous to Example 6). Therefore, we only consider $\mathcal{F}_{\text{PHP}}(n+1, n)$ over \mathbb{Q} .

With no symmetry restriction in place, it is – not surprisingly – possible to refute $\mathcal{F}_{\text{PHP}}(n+1, n)$ efficiently in the IPS. Indeed, the IPS p -simulates any Frege proof system [21], and the pigeonhole principle, formulated in propositional logic has a polynomial-size Frege proof [5]. The proof constructed in [5], however, proceeds along a linear order on $[n]$. Thus, a naive symmetrisation of it would require size $\mathcal{O}(n!)$. We show that we can do much better than that, although we do not obtain a symmetric refutation of subexponential size. It is plausible that this is impossible, and we leave the precise complexity of symmetrically refuting the pigeonhole principle as an intriguing open problem.

► **Theorem 20.** *There is a $(\mathbf{Sym}_{n+1} \times \mathbf{Sym}_n)$ -sym-IPS_{LIN} refutation of $\mathcal{F}_{\text{PHP}}(n+1, n)$ of size $\mathcal{O}(3^n \cdot n)$ over the field \mathbb{Q} .*

The key part of the refutation is to compute, for every subset $D \subseteq [n+1]$ the sum over all monomials that encode injections from the pigeons in D to the holes. Formally, let $B_D(\vec{x}) := \sum_{\gamma: D \hookrightarrow [n]} \prod_{i \in D} x_{i\gamma(i)}$. The polynomials $B_D(\vec{x})$ can also be viewed as the sums of certain *permanents*. The permanent of an $n \times n$ -matrix is the polynomial $\sum_{\pi \in \mathbf{Sym}_n} \prod_{i \in [n]} x_{i\pi(i)}$. The polynomial $B_D(\vec{x})$ is the sum over the permanents of all $|D| \times |D|$ -submatrices of $D \times [n]$. This hints at the potential hardness of symmetric refutations for the PHP because we know that the permanent admits no symmetric circuit representation of subexponential size [15]. In the proof of Theorem 20 in the appendix, we construct an $\mathcal{O}(3^n)$ -size symmetric circuit for computing the B_D for all $D \subseteq [n+1]$ and show how this yields a refutation for $\mathcal{F}_{\text{PHP}}(n+1, n)$.

7 Conclusion and future work

This article initiates the study of how symmetry in IPS proofs affects their complexity. We identify the following promising directions for future research: Firstly, we would like to obtain matching lower bounds for the exponential upper bounds we have established – that is, for *linear* symmetric IPS refutations of the CFI equations (which would show an exponential gap between linear and non-linear refutations), and for the pigeonhole principle. Secondly, the question in how far the functional lower bound method [17, 18, 23] can be combined with our framework deserves a deeper investigation. We have shown that the subset sum axiom and one of its liftings do admit small symmetric proofs but this does not rule out a lower bound via more complex subset sum variants such as in [18, 23]. Also, it is worth studying if the symmetry restriction we consider here, combined with the functional lower bound method from [18, 23] can extend the scope of that technique. So far, the functional lower bound method has only been applied to instances with a *single axiom*, and it provably fails on encodings of Boolean formulas [23]; can these limitations of the functional lower bound method be overcome by restricting to symmetric refutations? Finally, we ask if the connection between IPS lower bounds for Boolean CNFs and the separation of VP and VNP established by Grochow and Pitassi [21] also holds in a symmetric sense. That is, for suitably defined symmetric analogues of VP and VNP, would super-polynomial lower bounds against symmetric refutations of Boolean CNFs entail a separation of “symmetric VNP” from “symmetric VP”?

References

- 1 Albert Atserias, Andrei Bulatov, and Anuj Dawar. Affine systems of equations and counting infinitary logic. *Theoretical Computer Science*, 410(18):1666–1683, 2009. doi:10.1016/J.TCS.2008.12.049.
- 2 Paul Beame and Toniann Pitassi. Propositional proof complexity: Past, present, and future. *Current Trends in Theoretical Computer Science Entering the 21st Century*, pages 42–70, 2001.
- 3 Christoph Berkholz and Martin Grohe. Limitations of algebraic approaches to graph isomorphism testing. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *Automata, Languages, and Programming*, pages 155–166, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg. doi:10.1007/978-3-662-47672-7_13.
- 4 Andreas Blass, Yuri Gurevich, and Saharon Shelah. Choiceless polynomial time. *Annals of Pure and Applied Logic*, 100(1):141–187, 1999. doi:10.1016/S0168-0072(99)00005-6.
- 5 Samuel R Buss. Polynomial size proofs of the propositional pigeonhole principle. *The Journal of Symbolic Logic*, 52(4):916–927, 1987. doi:10.2307/2273826.

- 6 J. Cai, M. Fürer, and N. Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12:389–410, 1992. doi:10.1007/BF01305232.
- 7 Matthew Clegg, Jeffery Edmonds, and Russell Impagliazzo. Using the Groebner basis algorithm to find proofs of unsatisfiability. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 174–183, 1996.
- 8 Anuj Dawar. The nature and power of fixed-point logic with counting. *ACM SIGLOG News*, 2(1):8–21, 2015. doi:10.1145/2728816.2728820.
- 9 Anuj Dawar. Symmetric computation (invited talk). In *28th EACSL Annual Conference on Computer Science Logic, CSL 2020*, 2020. doi:10.4230/LIPIcs.CSL.2020.2.
- 10 Anuj Dawar. Limits of symmetric computation (invited talk). In *51st International Colloquium on Automata, Languages, and Programming, ICALP 2024*, volume 297 of *LIPIcs*, pages 1:1–1:8, 2024. doi:10.4230/LIPIcs.ICALP.2024.1.
- 11 Anuj Dawar, Erich Grädel, Leon Kullmann, and Benedikt Pago. Symmetric Proofs in the Ideal Proof System, 2025. doi:10.48550/arXiv.2504.16820.
- 12 Anuj Dawar and Benedikt Pago. A logic for P: are we nearly there yet? *ACM SIGLOG News*, 11:35–60, 2024. doi:10.1145/3665453.3665459.
- 13 Anuj Dawar, Benedikt Pago, and Tim Seppelt. Symmetric algebraic circuits and homomorphism polynomials. *arXiv*, abs/2502.06740, 2025. doi:10.48550/arXiv.2502.06740.
- 14 Anuj Dawar, David Richerby, and Benjamin Rossman. Choiceless Polynomial Time, Counting and the Cai–Fürer–Immerman graphs. *Annals of Pure and Applied Logic*, 152(1-3):31–50, 2008. doi:10.1016/J.APAL.2007.11.011.
- 15 Anuj Dawar and Gregory Wilsenach. Symmetric Arithmetic Circuits. In *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*, volume 168 of *LIPIcs*, pages 36:1–36:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ICALP.2020.36.
- 16 Anuj Dawar and Gregory Wilsenach. Lower bounds for symmetric circuits for the determinant. In *13th Innovations in Theoretical Computer Science Conference, ITCS*, volume 215 of *LIPIcs*, pages 52:1–52:22. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.ITCS.2022.52.
- 17 Michael A. Forbes, Amir Shpilka, Iddo Tzameret, and Avi Wigderson. Proof Complexity Lower Bounds from Algebraic Circuit Complexity. In Ran Raz, editor, *31st Conference on Computational Complexity (CCC 2016)*, volume 50 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 32:1–32:17, Dagstuhl, Germany, 2016. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CCC.2016.32.
- 18 Nashlen Govindasamy, Tuomas Hakoniemi, and Iddo Tzameret. Simple hard instances for low-depth algebraic proofs. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 188–199. IEEE, 2022. doi:10.1109/FOCS54457.2022.00025.
- 19 Erich Grädel and Martin Grohe. Is Polynomial Time Choiceless? In *Fields of Logic and Computation II*, pages 193–209. Springer, 2015. doi:10.1007/978-3-319-23534-9_11.
- 20 Erich Grädel and Stefan Hegselmann. Counting in Team Semantics. In Jean-Marc Talbot and Laurent Regnier, editors, *25th EACSL Annual Conference on Computer Science Logic (CSL 2016)*, volume 62 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 35:1–35:18, Dagstuhl, Germany, 2016. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CSL.2016.35.
- 21 Joshua A. Grochow and Toniann Pitassi. Circuit Complexity, Proof Complexity, and Polynomial Identity Testing: The Ideal Proof System. *J. ACM*, 65(6), November 2018. doi:10.1145/3230742.
- 22 Armin Haken. The intractability of resolution. *Theoretical computer science*, 39:297–308, 1985. doi:10.1016/0304-3975(85)90144-6.
- 23 Tuomas Hakoniemi, Nutan Limaye, and Iddo Tzameret. Functional lower bounds in algebraic proofs: Symmetry, lifting, and barriers. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1396–1404, 2024. doi:10.1145/3618260.3649616.

- 24 Sandra Kiefer. The Weisfeiler-Leman algorithm: an exploration of its power. *ACM SIGLOG News*, 7(3):5–27, 2020. doi:10.1145/3436980.3436982.
- 25 Jan Krajíček. *Proof Complexity*, volume 170. Cambridge University Press, 2019.
- 26 Martin Otto. *Bounded Variable Logics and Counting*, volume 9 of *Lecture Notes in Logic*. Springer, 1997.
- 27 Benedikt Pago. Finite Model Theory and Proof Complexity Revisited: Distinguishing Graphs in Choiceless Polynomial Time and the Extended Polynomial Calculus. In Bartek Klin and Elaine Pimentel, editors, *31st EACSL Annual Conference on Computer Science Logic (CSL 2023)*, volume 252 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 31:1–31:19, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CSL.2023.31.
- 28 Wied Pakusa. *Linear Equation Systems and the Search for a Logical Characterisation of Polynomial Time*. PhD thesis, RWTH Aachen, 2015.
- 29 Wied Pakusa, Benedikt Pago, Martin Grohe, and Erich Grädel. A Finite-Model-Theoretic View on Propositional Proof Complexity. *Logical Methods in Computer Science*, 15, 2019. doi:10.23638/LMCS-15(1:4)2019.
- 30 Alexander A Razborov. Lower bounds for the polynomial calculus. *computational complexity*, 7:291–324, 1998. doi:10.1007/S000370050013.
- 31 Benjamin Rossman. Choiceless Computation and Symmetry. In *Fields of Logic and Computation*, pages 565–580. Springer, 2010. doi:10.1007/978-3-642-15025-8_28.
- 32 Nathan Segerlind. The Complexity of Propositional Proofs. *Bulletin of symbolic Logic*, 13(4):417–481, 2007. doi:10.2178/BSL/1203350879.
- 33 Amir Shpilka and Avi Wigderson. Depth-3 arithmetic circuits over fields of characteristic zero. *Computational Complexity*, 10:1–27, 2001. doi:10.1007/PL00001609.