Depth Hierarchies for Symmetric Algebraic Circuits*

Anuj Dawar† Prateek Dwivedi‡ Benedikt Pago† Tim Seppelt§

3rd November 2026

Abstract

Algebraic complexity studies the cost of computing families of polynomials by algebraic circuits, with the central goal of proving super-polynomial circuit lower bounds and separations between different models of algebraic computation. Since such lower bounds remain elusive for general circuits, a common strategy is to focus on restricted models. A recent line of work (ICALP 2020, ITCS & STOC 2026) introduces and studies symmetry as a syntactic circuit restriction, showing exponential symmetric lower bounds for the permanent, as well as separations between symmetric analogues of algebraic formulas (VF), skew circuits (VS), and circuits (VP).

Motivated by the progress of this research programme, we investigate the role of *depth* as a complexity parameter in symmetric circuits. For every fixed $\Delta \geq 1$, we exhibit a family of polynomials computable by symmetric circuits of polynomial size and depth $2\Delta + 3$, but requiring size $n^{\Omega(\log n)}$ at depth Δ. For a parameter called *support depth*, which captures the semantically relevant concept of depth for symmetric circuits, this complexity gap arises already between Δ and $\Delta + 1$.

We also define symmetric versions of the bounded-depth algebraic complexity classes VAC^k, for all $k \in \mathbb{N}$, that we call symVAC^k, and establish two strict hierarchies: Firstly, a hierarchy within symVAC^0, corresponding to constant-depth symmetric computation, and secondly, a hierarchy of the form $\mathsf{symVAC}^k \subsetneq \mathsf{symVAC}^{k+1}$ for all k. The latter shows that symmetric circuits of depth $(\log n)^{k+1}$ are strictly more powerful than depth-$(\log n)^k$ circuits. Our results are obtained via a graph-theoretic characterisation of the symmetric complexity classes in terms of homomorphism polynomials. This also allows us to place the bounded-depth classes within the landscape of symmetric algebraic complexity developed by Dwivedi, Pago and Seppelt (STOC 2026). Finally, we complement this picture by introducing symVNP, a symmetric analogue of VNP, and establishing symVNP-completeness of the permanent under symmetric projections.

*A Lean formalisation of the main results is available at https://laxarchive.org/lax-461425/.

†University of Cambridge, United Kingdom. {anuj.dawar, benedikt.pago}@cl.cam.ac.uk.

‡Universität des Saarlandes, Germany, and Université Savoie Mont Blanc, France. dwivedi@cs.uni-saarland.de.

§IT-Universitetet i København, Denmark. tise@itu.dk.

Contents

1. Introduction **1**
- 1.1. Results . . . 2
- 1.2. Techniques . . . 5
- 1.3. Future work . . . 6

2. Preliminaries **7**
- 2.1. Circuits and polynomial families . . . 7
- 2.2. Matrix symmetry and orbit size . . . 7
- 2.3. Rigid circuits, supports, and support depth . . . 7
- 2.4. Homomorphism polynomials and graph parameters . . . 8
- 2.5. Symmetric algebraic complexity classes . . . 9

3. Symmetric constant-depth hierarchy **9**
- 3.1. Characterisations . . . 10
- 3.2. Separation . . . 11

4. Symmetric depth reduction **11**

5. Symmetric polylogarithmic-depth hierarchy **12**
- 5.1. Constant-depth reduction for formulas . . . 13
- 5.2. Bounded pathwidth in logarithmic support depth . . . 13
- 5.3. Separating symVAC^1 from symVP . . . 14

6. Symmetric algebraic branching programs **14**

7. Symmetric VNP **16**
- 7.1. symVNP contains all matrix-symmetric polynomials . . . 17
- 7.2. Hypercube sums of symmetric formulas . . . 19
- 7.3. Completeness of the permanent under symmetric projections . . . 19

Acknowledgements **21**

References **22**

A. Reduced circuits and supports **27**

B. Material omitted in Section 3 **28**
- B.1. Circuits for patterns of bounded forestwidth . . . 28
- B.2. Homomorphism polynomials from symmetric circuits . . . 30
  - B.2.1. Labelled graphs . . . 30
  - B.2.2. Eliminating many labels at once . . . 32
  - B.2.3. Induction on the circuit . . . 33
- B.3. Combinatorics of bounded-depth tree decompositions . . . 33
  - B.3.1. Richness . . . 34
  - B.3.2. Oddomorphisms reflect richness . . . 35
- B.4. Non-uniform descriptive complexity monotonicity . . . 36
  - B.4.1. CFI graphs for non-induced minors . . . 36
  - B.4.2. Proof of Theorem 3.7 . . . 39

B.5. Separation 40

C. Material omitted in Section 4 **41**
C.1. Omitted material on reduction to logarithmic depth 41
C.2. Material omitted on bounded pathwidth and logarithmic depth 42
C.3. Separations of symVAC^k and symVP 43
C.3.1. Set-up 43
C.3.2. Translating structures into weighted graphs 44
C.3.3. The separation 46

1. Introduction

Algebraic complexity theory studies families of polynomials in terms of the algebraic circuit size that is required to represent them. An algebraic circuit is a directed acyclic graph whose internal vertices correspond to addition and multiplication gates and whose input gates are labelled with variables or constants from a fixed underlying field. It computes a polynomial in the input variables in the natural way, and the circuit size is defined as the total number of gates and wires it contains. In 1979, Valiant postulated the conjecture that the *permanent* polynomial, i.e., the determinant without signs, does not admit polynomial-size algebraic circuits [Val79]. A central part of the algebraic complexity programme is thus to develop techniques for proving circuit size lower bounds. The common approach is to begin with restricted circuit models, where lower bounds are easier to establish. A substantial line of work has focussed on proving lower bounds for example on monotone or multilinear circuits and formulas [KS14]. Arguably, these results tell us more about the weakness of the restricted models than about the true complexity of the permanent: In fact, many of them, like the exponential lower bound for multilinear formulas computing the permanent [Raz04], apply equally to the determinant, even though the determinant is known to admit polynomial-size circuits.

In recent years, a new circuit restriction has been investigated [DW22; DW20; DPS26a; DPS26b], in which determinant and permanent provably do exhibit a polynomial versus exponential complexity gap, in line with Valiant's conjecture. This model is that of *symmetric circuits*. Not only has the separation between determinant and permanent [DW20] been established for this model, but also the separations of symmetric analogues of the major classes in algebraic complexity, that is, VF, VBP/ VS, and VP [DPS26b] (corresponding to algebraic formulas, branching programs or skew circuits, and general circuits, respectively). Briefly, the notion of symmetry means the following: Whenever a polynomial $p \in \mathbb{K}[X]$ is invariant under a permutation group Γ acting on the variables X, then it makes sense to require Γ-symmetry also of a circuit C that computes p. A circuit C is Γ-*symmetric* if, for every permutation $\pi \in \Gamma$ that acts on the input variables, there also exists a permutation σ of the gates of C that is consistent with π on the input gates, and preserves the wire structure of the circuit (that is, is an automorphism of C). For a formal definition, see Section 2.2. In this paper and in most of the works referenced above, the group Γ is always $\mathbf{Sym}_n \times \mathbf{Sym}_n$, and it acts via row and column permutations on variables $X_n := \{x_{ij} \mid i, j \in [n]\}$ that are arranged in an $n \times n$ matrix. This notion is called *matrix symmetry*, and it applies to matrix polynomials such as the permanent and many other natural polynomials[1] from algebraic complexity.

The recent lower bounds and separation results for symmetric circuits establish this model as an important and promising concept that merits further investigation. In this paper, we focus on the interplay of symmetry and *depth* in algebraic circuits. Bounded-depth circuit classes are a natural notion that has been studied extensively both in the algebraic and the Boolean setting. We prove *depth hierarchies* for symmetric circuits and place them in the existing symmetric algebraic circuit complexity landscape. A depth hierarchy is a statement along these lines: For every $\Delta \in \mathbb{N}$, there exist families of polynomials that are computable by depth-Δ circuits of polynomial size but require super-polynomial size to be computed by a circuit of depth $< \Delta$. Such depth hierarchy results were first shown for restricted circuit models, such as constant-depth multilinear formulas [RY09; Chi+18], and recently for general constant-depth algebraic circuits in the breakthrough by Limaye, Srinivasan and Tavenas [LST25]. Other known results in this style are the separation

[1] The determinant is symmetric, too, but only under the *simultaneous* action of $\mathbf{Sym}_n$ on the rows and columns, not under all independent row and column permutations.

between multilinear NC^1 and NC^2 circuits [Raz06], and depth hierarchies of various strengths in the Boolean setting, that is, for AC^0 circuits [Sip83; Hås87; Hås+17].

1.1. Results

Our main results establish depth hierarchies for symmetric algebraic circuits, both within the constant-depth regime between Δ and $\Delta + 1$, as well as between depths $(\log n)^k$ and $(\log n)^{k+1}$ for all constants k. The former is a symmetric analogue of the depth hierarchy inside VAC^0, which is the class of polynomials computable by constant-depth polynomial-size circuits, and the latter shows a hierarchy between symmetric analogues of the classes VAC^k and VAC^{k+1} for all k.

These results are not implied by the hierarchies that are known in the non-symmetric setting, and are in some sense incomparable: While the symmetric model is syntactically restricted, its complexity measure is more lenient than for standard circuits. Namely, the natural size parameter of symmetric circuits is their *orbit size*, which is the number of gates in its largest orbit under the action of the symmetry group. Thus, a circuit can be considered efficient under symmetry aspects even when it contains a super-polynomial number of gates, as long as for each gate, there is at most a polynomial number of gates that are symmetric to it. A circuit with super-polynomial orbit size has in particular super-polynomial size, so lower bound results for symmetric circuits apply to the standard notion of size (i.e., total number of gates), too. Our depth hierarchy holds both with respect to the orbit size and standard circuit size measure.

Lean formalisation on LaxLean-Archive

Theorem 1.1. *Fix any $\Delta \geq 1$. There is a family $(p_n)_{n\in\mathbb{N}}$ of matrix-symmetric polynomials such that*

1. *Every matrix-symmetric circuit of depth $\leq \Delta$ computing p_n has size at least $n^{\Omega(\log_2 n)}$.*
2. *There is a matrix-symmetric circuit of depth $2\Delta + 3$ and size $O(n^{\Delta+2})$ computing p_n.*

Stated in terms of the notion of *support depth*, which is the semantically relevant depth parameter for symmetric circuits (see Section 2.3), the separation is in fact between Δ and $\Delta + 1$, see Theorem 3.1 for the formal statement. This is similar to other depth hierarchy results where *product depth*, rather than depth, is the relevant measure. Moreover, support depth and depth agree up to a factor of 2 at polynomial orbit size, as we will see in Corollary 3.6.

Theorem 1.1 motivates introducing bounded-depth symmetric algebraic complexity classes and locating them in the symmetric algebraic complexity hierarchy

$$\mathsf{symVF} \subsetneq \mathsf{symVS} \subsetneq \mathsf{symVP}$$

that was established in [DPS26b]. Using sophisticated machinery from finite model theory [GLN25] and the theory of homomorphism indistinguishability [Rob26], we place the bounded-depth classes in the existing picture, and unconditionally establish the separations depicted in Figure 1.

Definition 1.2. For $\Delta \geq 0$, let symVAC^0_Δ denote the class of families (p_n) of matrix-symmetric polynomials for which there is a constant c such that, for all sufficiently large n, the polynomial p_n is computed by a matrix-symmetric circuit of orbit size $\leq n^c$ and support depth $\leq \Delta$. Let $\mathsf{symVAC}^0 := \bigcup_{\Delta\geq 0} \mathsf{symVAC}^0_\Delta$.

For $k \geq 1$, let symVAC^k denote the class of families (p_n) of matrix-symmetric polynomials for which there is a constant c such that, for all sufficiently large n, the polynomial p_n is computed by a matrix-symmetric circuit of orbit size $\leq n^c$ and support depth $\leq c \log^k n$. Let $\mathsf{symVAC} := \bigcup_{k\geq 1} \mathsf{symVAC}^k$.

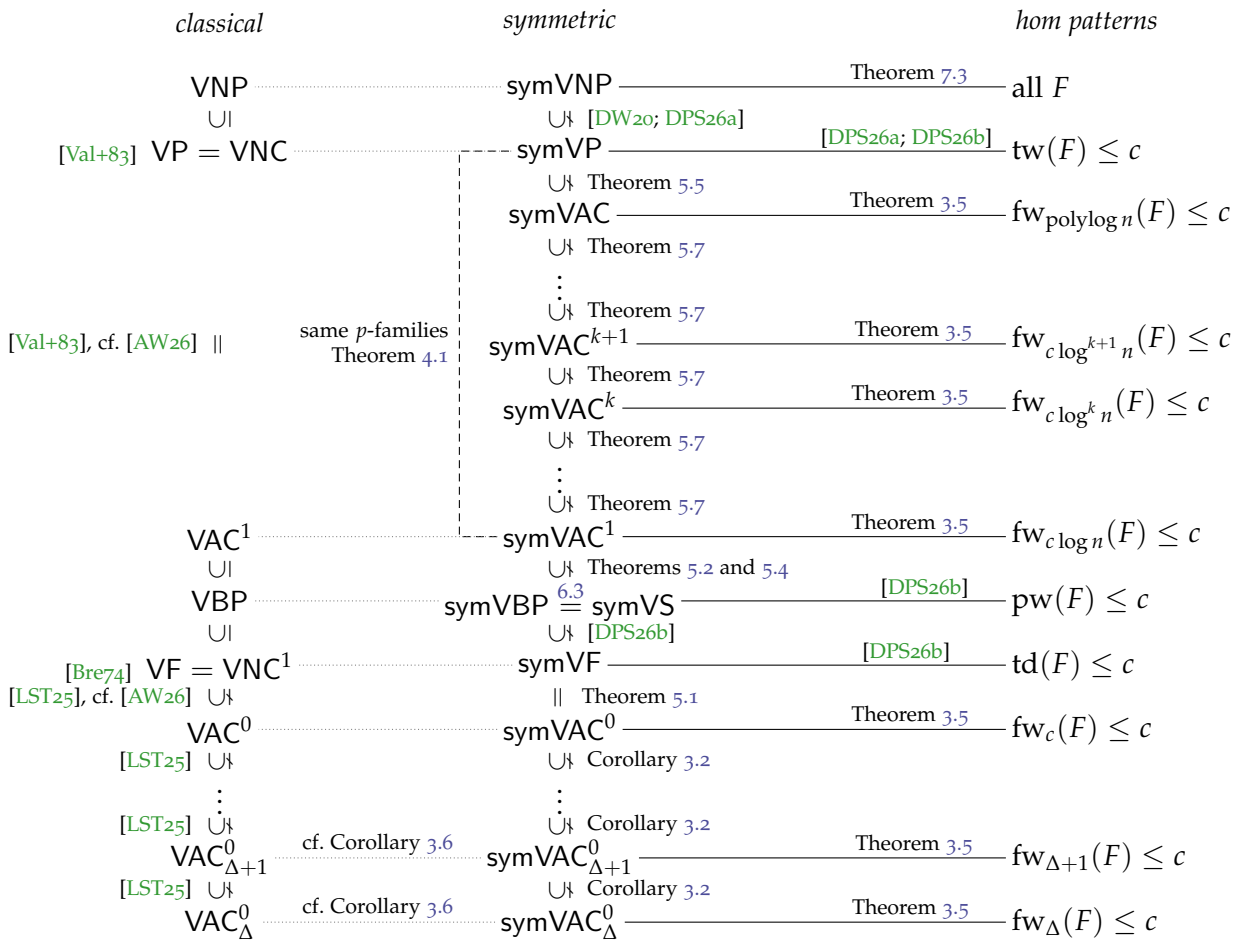


Figure 1: The classical landscape of algebraic complexity classes (left) and the symmetric landscape developed here (middle). Dotted lines connect symmetric classes to their classical counterparts (by identifying size with orbit size, depth with support depth). On the classical side, all classes consist of p-families, that is, families of polynomials with polynomially bounded degree: VP and VNP by definition, VAC^0, $\mathsf{VF} = \mathsf{VNC}^1$, and VBP automatically, and VAC^1 and VNC are restricted to p-families. Each symmetric class consists precisely of the linear combinations of homomorphism polynomials $\hom_{F,n}$ for patterns F satisfying the condition on the right, for a constant c. Restricted to p-families, symVP and symVAC^1 are equal.

Symmetric algebraic complexity classes and homomorphism polynomials. The classes symVF, symVS, and symVP comprise matrix-symmetric polynomials that are computable by matrix-symmetric formulas, skew circuits, and circuits of polynomial orbit size, respectively. The main technical contribution of [DW20; DPS26b] was a characterisation of these syntactic circuit classes in terms of a graph-theoretic semantics: Each of these classes consists precisely of linear combinations of *homomorphism polynomials* of patterns of bounded treedepth, bounded pathwidth and bounded treewidth, respectively, cf. Theorem 2.9. A homomorphism polynomial

$$\hom_{F,n}(X) = \sum_{h\colon V(F)\to[n]} \prod_{ab\in E(F)} x_{h(a),h(b)}$$

is parametrised by a pattern graph F and the size n of the target graph. Its variables can be instantiated with the adjacency matrix of an n-vertex graph G, and $\hom_{F,n}(G)$ evaluates to the number of homomorphisms from F to G. The separation of the symmetric algebraic complexity classes then follows from the separation of the graph distinguishing power of the corresponding classes of homomorphism polynomials.

The right column of Figure 1 describes the structural graph parameter of the homomorphism polynomials whose linear combinations capture the respective symmetric algebraic complexity class. For example, the correspondence between symVP and $\mathrm{tw}(F) \leq c$ of [DPS26a] is to be read as follows: A family $(p_n)_{n\in\mathbb{N}}$ is in symVP if and only if there is a fixed constant $c \in \mathbb{N}$ such that for every $n \in \mathbb{N}$

$$p_n(X) = \sum_{F \text{ with } \mathrm{tw}(F)\leq c} \alpha_{F,n} \cdot \hom_{F,n}(X), \tag{1}$$

for some choice of coefficients $\alpha_{F,n} \in \mathbb{K}$.

In Theorem 3.5, we prove a characterisation of the polynomials computed by matrix-symmetric circuits of a given orbit size and support depth. The graph parameter arising here is Δ-*forestwidth* fw_Δ, a graph parameter inspired by Δ-treewidth of Bhargav, Chen, Curticapean, and Dwivedi [Bha+25]: For $\Delta \geq 0$, a graph F has Δ-*treewidth* $\mathrm{tw}_\Delta(F) \leq k$ if it admits a tree decomposition (T, β) of width $\leq k$ such that the decomposition tree is of depth at most Δ. For example, $\mathrm{tw}_0(F) = |V(F)| - 1$, as $\Delta = 0$ forces the tree decomposition to consist of a single vertex, and $\mathrm{tw}_1(F)$ is closely related to the vertex-cover number of F, see [Bha+25]. Bounded-depth tree decompositions are studied in finite model theory [Adl+26] and also arise as divide-and-conquer schemes for iterated matrix multiplication [LST25]. Δ-forestwidth differs from Δ-treewidth only in the additional technical requirement that the root bag be empty.

In order to describe the bounded-depth setting, our characterisation Theorem 3.5 needs to be more precise than previous results [DPS26b; DPS26a], which it fully subsumes. In particular, Theorem 3.5 does not only hold asymptotically but at every fixed host size n, see Section 3 for a discussion. This allows us to overcome the limitation of [DPS26b; DPS26a] to circuits of polynomial orbit size.

Comparing classical and symmetric hierarchies. The depth hierarchy within symVAC^0 parallels the one within VAC^0, though it follows from completely different arguments. As noted earlier, it is not just a special case of the result for VAC^0, among other things because of the distinction of size versus orbit size. The separations between symVAC^k and symVAC^{k+1} have an essentially uninteresting analogue in the non-symmetric setting: Since the definition of VAC^k bounds total circuit size rather than orbit size, there is an upper bound of $2^{O(\log^{k+1} n)}$ on the degree of the polynomials that VAC^k can compute—therefore, VAC^{k+1} contains polynomials that are not in VAC^k simply because their degree is too large. The definition of symVAC^k, however, puts no

upper bound on the total circuit size, but rather on the orbit size. This does not entail an upper bound on the degree, so the separation between symVAC^k and symVAC^{k+1} is a genuine result, not an immediate consequence of the definition. It builds on quite sophisticated lower bounds on the iteration number of the Weisfeiler–Leman algorithm for graph isomorphism [GLN25]. This iteration number translates into symmetric circuit depth via our homomorphism polynomial characterisation.

A new separation that we establish in the symmetric setting and which is not known to hold for non-symmetric circuits is $\mathsf{symVBP} = \mathsf{symVS} \subsetneq \mathsf{symVAC}^1$. Even the inclusion, not just the separation, is surprising, given that no degree restrictions are imposed in these classes. For example, symVS contains the homomorphism polynomials of paths with 2^n vertices, whose natural representation in symVS has depth $2^{\Omega(n)}$. It is not at all straightforward to show that these polynomials are also in symVAC^1, i.e., require only logarithmic depth. For p-families, this follows from Theorem 4.1, but for families of higher degree, this requires substantially different techniques from invariant theory, as our graph-theoretic criterion alone does not yield this.

Symmetric VBP and VNP. As depicted in the figure, we also provide definitions of the symmetric classes symVNP and symVBP, for which no natural definitions were known so far. While symVS was defined in terms of symmetric skew circuits, symVBP comprises the families of polynomials computable by symmetric algebraic branching programs (ABPs) of polynomial orbit size. In the general non-symmetric setting, skew circuits and ABPs are polynomially equivalent. As we prove in Theorem 6.3, via the homomorphism polynomial characterisation, this carries over to the symmetric setting, so that we also have $\mathsf{symVS} = \mathsf{symVBP}$. In order to define the class symVNP, we place a natural symmetry condition on the definition of VNP in terms of hypercube sums over VP-polynomials, see Definition 7.1. This definition turns out to yield the correct symmetric analogue of VNP in the sense that it satisfies the following neat properties: Firstly, symVNP is equal to the class of *all* matrix-symmetric polynomials, see Theorem 7.3. Secondly, the permanent is complete for symVNP under *symmetric* projections (Theorem 7.7)—this is in analogy to the VNP-completeness of the permanent under not necessarily symmetric projections. Symmetric formulas are strictly weaker than symmetric circuits [DPS26b, Theorem A]. Nevertheless, Corollary 7.4 shows that their Boolean sums define the same class. Thus, as in classical VNP, restricting witnesses to formulas does not reduce expressive power.

1.2. Techniques

On a high level, the proof of Theorem 1.1 follows the strategy established in [DPS26b; DPS26a] for separating symmetric algebraic complexity classes. Its two ingredients are a *characterisation* of a symmetric algebraic complexity class in terms of homomorphism polynomials for patterns from some graph classes of restricted structure and a *separation* result asserting that a certain polynomial cannot be written as a linear combination of homomorphism polynomials for patterns from some graph class. However, the ideas in [DPS26b; DPS26a] are insufficient for Theorem 1.1 and need to be substantially refined.

Characterisations. The characterisations are insufficient as they only apply to circuits of polynomial orbit size. Our lower bounds require considering circuits of orbit size $n^{\Omega(\log n)}$ and thus a tighter analysis. Moreover, the most technical ingredient of [DPS26b; DPS26a], a lemma describing how support-changing gates affect the inductively constructed linear combination of homomorphism polynomials, needs to be replaced entirely. For details, see Section B.2.

Separation. Even though every matrix-symmetric polynomial can be written as a linear combination of homomorphism polynomials, this linear combination is by no means unique; see [DPS26b, Example 3.1] and [Cur22] for examples. Figure 1 also illustrates this: By Theorems 5.2, 3.5, and 2.9, it holds that $\mathsf{symVBP} \subseteq \mathsf{symVAC}^1$, i.e., every polynomial that can be written as a linear combination of homomorphism polynomials for patterns of pathwidth $O(1)$ can also be written as such a combination for patterns of $O(\log n)$-forestwidth $O(1)$. However, there are families of graphs of constant pathwidth and super-constant $O(\log n)$-forestwidth. Nonetheless, their homomorphism polynomials lie in symVAC^1.

Dawar, Pago, and Seppelt [DPS26a] introduced the *non-uniform descriptive complexity monotonicity* for ruling out such failures of homomorphism polynomials of distinct patterns to be linearly independent. For example, in order to show that a matrix-symmetric polynomial p_n is not in symVP, one may (recalling (1)) construct n-vertex graphs G and H that are *homomorphism indistinguishable* over all graphs of treewidth $\leq c$, i.e., admit the same number of homomorphisms from every such graph, but differ in the value of $p_n(G) \neq p_n(H)$. The construction of G and H in [DPS26a] crucially relies on the fact that the bound on the treewidth $\leq c$ is constant while the host size n is unbounded. It therefore cannot yield lower bounds as in Theorem 1.1 or lower bounds against symVAC^k where 'treewidth $\leq c$' is replaced by '$c \log^k n$-forestwidth $\leq c$', a condition depending on the host size n. In Theorem 3.7, we overcome this issue using techniques from structural graph theory and homomorphism indistinguishability [Rob26]. Our upgraded non-uniform descriptive complexity monotonicity supersedes [DPS26b; DPS26a] and has a more direct proof.

1.3. Future work

We highlight the following three directions for future work that may build on the results we have obtained. Firstly, our logarithmic depth reduction for p-families (Theorem 4.1) raises the question of whether symmetric circuits can be reduced further, to depth 3 or 4, and at what cost in orbit size. In the classical setting, depth reduction began with balancing circuits [Hya79; Val+83] and formulas [Bre74; BCE95]. Later work gave sharper bounds for homogeneous formulas [Fou+26] and syntactically multilinear circuits [KOS19]. The depth-four [AV08; Koi12; Tav15] and depth-three reductions over the rationals [Gup+16; Tav15] allow sufficiently strong lower bounds for shallow circuits to be lifted to general circuits. This is the motivation for seeking analogous reductions that preserve symmetry.

Another promising direction is to port the symmetry framework to *algebraic proof complexity*. In [Daw+25], a symmetric version of the *Ideal Proof System* (IPS), which is based on algebraic circuits, was introduced. Can symmetries be exploited in order to prove strong lower bounds against this proof system? An active recent line of work has established lower bounds for very restricted fragments of IPS, for example ro-ABP or constant depth multilinear IPS proofs [For+21; GHT22; HLT24; Elb+26]. A natural next step in this direction would be to attempt lower bounds for bounded depth symmetric IPS proofs. The machinery we have developed in this paper may serve as a tool to approach this problem.

Finally, we believe that our definition of the class symVBP is a first step towards combining the symmetry-based lower bound approach with an ongoing programme on stronger (multilinear read-once oblivious) ABP lower bounds, e.g., [Cha+24; BDS25; Fab+26]. Our definition and homomorphism polynomial characterisation of symmetric ABPs may be used in combination with other techniques to make new progress in this line of research, too.

2. Preliminaries

We work over a field $\mathbb{K}$ of characteristic zero. Write $[n] = \{1, \dots, n\}$ and let $\mathbf{Sym}_n$ be its permutation group. We follow the circuit and graph-polynomial conventions of [DPS26a; DPS26b].

2.1. Circuits and polynomial families

An *algebraic circuit* is a finite directed acyclic graph whose input gates are labelled by variables or field constants and whose internal gates are labelled by $+$ or $\times$. Input gates with the same label are identified. We allow unbounded fan-in and parallel wires, counted with multiplicity in each operation. There is a unique output gate, and every gate lies on a path to it. The *size* is the number of gates and wires, counting multiplicities. The *depth* is the maximum number of wires on an input-to-output path. Likewise, the depth of a rooted tree or forest counts edges, so a single node has depth 0.

A *formula* has no parallel wires and its internal gates induce a tree; input gates may be shared. A circuit is *skew* if each multiplication gate has at most one incoming wire from an internal gate.

A function is *p-bounded* if it is bounded by a polynomial in its arguments. A *p-family* is a polynomial family whose number of variables and degree are p-bounded. The classes VP and VF consist of p-families computed by circuits and formulas, respectively, of polynomial total size. A p-family (f_n) belongs to VNP if there are a p-family $(g_n) \in \mathsf{VP}$ and a p-bounded function q such that

$$f_n(X_n) = \sum_{\boldsymbol{e} \in \{0,1\}^{q(n)}} g_n(X_n, \boldsymbol{e}).$$

2.2. Matrix symmetry and orbit size

For $n \geq 1$, put $\mathcal{X}_n = \{x_{i,j} : i, j \in [n]\}$. The group $\Gamma_n = \mathbf{Sym}_n \times \mathbf{Sym}_n$ acts on $\mathcal{X}_n$ by $(\sigma, \tau)x_{i,j} = x_{\sigma(i),\tau(j)}$. A polynomial is *matrix-symmetric* if it is invariant under this action. For a group Γ acting on a set Ω and $S \subseteq \Omega$, write $\mathbf{Stab}_\Gamma(S)$ for the setwise stabiliser and $\mathbf{Stab}^\bullet{}_\Gamma(S)$ for the pointwise stabiliser. We omit Γ when it is clear.

Let Γ act on the variables of a circuit C, fixing field constants. Write $\lambda(v)$ for the label of a gate v. A circuit automorphism preserves internal operation labels and the directions and multiplicities of wires. Such an automorphism σ *extends* $\gamma \in \Gamma$ if $\lambda(\sigma(v)) = \gamma(\lambda(v))$ for every input gate v.

Definition 2.1 (Symmetric circuits)**.** A circuit is Γ-*symmetric* if every $\gamma \in \Gamma$ extends to a circuit automorphism. For $\Gamma = \Gamma_n$ acting on $\mathcal{X}_n$, we call it matrix-symmetric.

Lean formalisation on LaxLean-Archive

Let $\mathbf{Aut}_\Gamma(C)$ be the group of all automorphisms extending some element of Γ. The *orbit size* is

$$\mathrm{maxOrb}(C) = \max_{v \in V(C)} \big|\{\sigma(v) : \sigma \in \mathbf{Aut}_\Gamma(C)\}\big|.$$

2.3. Rigid circuits, supports, and support depth

The group Γ_n acts on $[n] \uplus [n]$ and on $\mathcal{X}_n$.

Definition 2.2. A matrix-symmetric circuit C is *rigid* if every element of Γ_n extends to a unique automorphism of C. It is *reduced* if no two distinct input gates have the same label and no two distinct internal gates have the same label and the same multiset of children.

If C is rigid, then Γ_n acts on the gates of C, and the orbit size of C is the maximal index of a stabiliser $\mathbf{Stab}(g) \leq \Gamma_n$ of a gate g.

Lemma 2.3. *Every reduced matrix-symmetric circuit is rigid.*

Lean formalisation on LaxLean-Archive

Lemma 2.4 (Rigidification)**.** *For every matrix-symmetric circuit C, there is a reduced matrix-symmetric circuit C' computing the same polynomial whose size, depth, and orbit size are at most those of C.*

Lean formalisation on LaxLean-Archive

The proofs are given in Section A. We call the circuit C' constructed there the *rigidification* of C. The rigidification lemma here has appeared in similar forms in the literature on symmetric circuits before, e.g., in [DPS26a; DPS26b]. In all these instances, the lemma actually yields not only a rigid, but also a *reduced* circuit (which was not always made explicit). Reducedness is in fact a stronger condition than rigidity, and the support theorems for symmetric circuits, as presented below, are actually based on this stronger assumption.

Let C be rigid. A set $S \subseteq [n] \uplus [n]$ *supports* a gate g if the pointwise stabiliser $\mathbf{Stab}^\bullet(S)$ fixes g. We write S_L and S_R for the two sides of S. If g has a support S with $|S_L|, |S_R| < n/2$, then g has a unique minimal support $\sup(g) \subseteq S$, which is contained in every support of g, and $\sup(\pi g) = \pi \sup(g)$ for all $\pi \in \Gamma_n$ [DPS26a, Lemmas 2.1 and 2.2]. The *support size* $\mathrm{maxSup}(C)$ of C is the least k such that every gate of C has a support of size $\leq k$. If $\mathrm{maxSup}(C) < n/2$, then every gate g of C has a minimal support and $\mathrm{maxSup}(C) = \max_g |\sup(g)|$. In this case, a wire from h to g is *support-changing* if $\sup(h) \not\subseteq \sup(g)$, and the *support depth* of C is the maximum number of support-changing wires on a path from the output gate to an input gate. If $\mathrm{maxSup}(C) \geq n/2$, we let the support depth of C be its depth. For an arbitrary matrix-symmetric circuit, we define its support size and support depth as those of its rigidification. In any case, the support depth of a circuit is at most its depth.

Observation 2.5. *Every rigid matrix-symmetric circuit of support size $\leq k$ has orbit size $\leq n^k$.*

Proof. Let S be a support of size $\leq k$ of a gate g. Since $\mathbf{Stab}^\bullet(S)$ fixes g, the orbit of g has size at most the number of images of S under Γ_n, which is $\leq n^{|S|}$. □

Conversely, small orbit size implies small support size. In contrast to [DPS26a; DPS26b], we allow circuits of superpolynomial orbit size. For a function $f\colon \mathbb{N} \to \mathbb{N}$ bounding the orbit size, let

$$\kappa_f(n) := \min\left\{ r \in \mathbb{N} : \binom{n}{r} > f(n) \right\}. \tag{2}$$

If $f(n) \in n^{O(1)}$, then $\kappa_f(n) \in O(1)$, i.e. polynomial orbit size corresponds to constant support size, as in [DPS26a; DPS26b]. The proof of the following theorem is given in Section A.

Theorem 2.6 (Support theorem)**.** *Let $f\colon \mathbb{N} \to \mathbb{N}$. Let $n \geq 20$ such that $3\kappa_f(n) \leq n + 2$. Then every gate of a reduced matrix-symmetric circuit of orbit size $\leq f(n)$ has a minimal support of size $\leq \kappa_f(n) - 1$.*

Lean formalisation on LaxLean-Archive

2.4. Homomorphism polynomials and graph parameters

A *pattern* is a finite bipartite multigraph F with a fixed bipartition $A \uplus B$. For $n \geq 1$, its *n-th homomorphism polynomial* is

$$\hom_{F,n}(X) = \sum_{\substack{h_A\colon A \to [n] \\ h_B\colon B \to [n]}} \prod_{ab \in E(F)} x_{h_A(a), h_B(b)}.$$

The product counts edges with multiplicity. Evaluated on the bi-adjacency matrix of a bipartite graph, this polynomial counts bipartition-respecting homomorphisms.

Observation 2.7. *If $F + K_1$ is obtained from F by adding an isolated vertex to the left part of the bipartition, then $\hom_{F+K_1,n} = n \hom_{F,n}$ for all $n \geq 1$. By adjusting coefficients, we may assume that the patterns have no isolated vertices.*

A *tree decomposition* of a graph F is a tree T with bags $\beta(t) \subseteq V(F)$ covering $V(F)$. Each edge has both endpoints in some bag, and the bags containing any fixed vertex induce a connected subtree. Its width is $\max_t |\beta(t)| - 1$. Minimising this width gives the *treewidth* $\mathrm{tw}(F)$; requiring T to be a path gives the *pathwidth* $\mathrm{pw}(F)$. A forest decomposition uses a forest in place of T.

An *elimination forest* of F is a rooted forest on $V(F)$ in which the endpoints of every edge are comparable in the ancestor order. The *treedepth* $\mathrm{td}(F)$ is the minimum possible maximum number of vertices on a root-leaf path in such a forest. Thus $\mathrm{td}(K_1) = 1$, while the depth of a one-node rooted tree is 0. These graph parameters ignore edge multiplicities.

2.5. Symmetric algebraic complexity classes

We recall the complexity classes introduced by Dwivedi, Pago, and Seppelt [DPS26b].

Definition 2.8 ([DPS26b])**.** The classes symVP, symVS, and symVF contain all families (p_n) of matrix-symmetric polynomials computed by matrix-symmetric circuits, skew circuits, and formulas, respectively, of orbit size $n^{O(1)}$.

These classes impose no degree or total-size bound. The class symVBP of symmetric algebraic branching programs is defined in Section 6. Theorem 6.3 shows that $\mathsf{symVBP} = \mathsf{symVS}$.

Theorem 2.9 ([DPS26b; DPS26a])**.** *Let (p_n) be a family of matrix-symmetric polynomials.*

1. $(p_n) \in \mathsf{symVP}$ *if, and only if, there is a constant k such that, for all sufficiently large n, p_n is a linear combination of homomorphism polynomials for patterns of treewidth $\leq k$.*
2. $(p_n) \in \mathsf{symVS}$ *if, and only if, there is a constant k such that, for all sufficiently large n, p_n is a linear combination of homomorphism polynomials for patterns of pathwidth $\leq k$.*
3. $(p_n) \in \mathsf{symVF}$ *if, and only if, there is a constant k such that, for all sufficiently large n, p_n is a linear combination of homomorphism polynomials for patterns of treedepth $\leq k$.*

3. Symmetric constant-depth hierarchy

In this section, we prove the following depth hierarchy theorem for matrix-symmetric circuits, which makes Theorem 1.1 precise.

Lean formalisation on LaxLean-Archive

Theorem 3.1. *Let $\Delta \geq 1$ and $n \geq 20$. Let $T_{\Delta,n}$ denote the full n-ary tree of depth Δ.*

1. *Every matrix-symmetric circuit of support depth $\leq \Delta$ computing $\hom_{T_{\Delta,n},n}$ has orbit size, and hence size, at least*
$$n^{\log_2 n - \Delta \log_2 \log_2 n - 3}.$$
2. *There is a matrix-symmetric circuit of support depth $\Delta + 1$, depth $2\Delta + 3$, orbit size $\leq n^2$, and size $\leq 13 n^{\Delta+2}$ computing $\hom_{T_{\Delta,n},n}$.*

Thus, for fixed Δ, the size gap between support depth Δ and $\Delta + 1$ for these polynomials is $n^{(1-o(1))\log_2 n}$. Theorem 3.1 has the following immediate corollary.

Corollary 3.2. *For every $\Delta \geq 0$, it holds that $\mathsf{symVAC}^0_\Delta \subsetneq \mathsf{symVAC}^0_{\Delta+1}$.*

The second assertion of Theorem 3.1 is standard and proven in Theorem B.2. For the first assertion, we substantially refine the machinery of [DPS26a; DPS26b]. The proof is conducted in two steps: In Section 3.1, we give a translation between matrix-symmetric algebraic circuits and homomorphism polynomials such that support size and support depth of the former correspond to width and depth of tree decompositions of the patterns of the latter. In Section 3.2, we show that homomorphism polynomials for $T_{\Delta,n}$ cannot be written as linear combinations of homomorphism polynomials of patterns of Δ-forestwidth $o(\log n)$.

3.1. Characterisations

The first step is to characterise the polynomials computed by bounded-depth matrix-symmetric circuits in terms of homomorphism polynomials. The graph-theoretic parameter of the patterns arising here is a modification of the Δ-*treewidth*[2] of Bhargav, Chen, Curticapean, and Dwivedi [Bha+25].

Definition 3.3. Let $\Delta \geq 1$. The Δ-*forestwidth* $\mathrm{fw}_\Delta(F)$ of F is the least $k \geq 0$ such that F admits a tree decomposition (T, β) with root r such that T has depth $\leq \Delta$, $\beta(r) = \emptyset$, and $|\beta(t)| \leq k+1$ for all $t \in V(T)$. Equivalently, deleting r yields a forest decomposition of F in which every root-leaf path has at most Δ nodes.

Lean formalisation on LaxLean-Archive

Our Theorem 3.1 follows from the following observation.

Example 3.4 ([Bha+25, Theorem 17]). *Let $\Delta, k \geq 1$. Let $T_{\Delta,k}$ denote the full k-ary tree of depth Δ. Then $\mathrm{fw}_{\Delta+1}(T_{\Delta,k}) = 1$ whereas $\mathrm{fw}_\Delta(T_{\Delta,k}) \geq k-1$.*

The main result facilitating all separations in Figure 1 is the following characterisation of the polynomials computed by bounded-depth matrix-symmetric circuits. The characterisation is exact when the circuits are measured by their supports rather than their orbit size, see Section 2.3.

Theorem 3.5. *Let $\Delta \geq 1$, $k \geq 2$, $n \geq 2k+1$. For a polynomial $p \in \mathbb{K}[\mathcal{X}_n]$, the following are equivalent:*

1. *p is computed by a reduced matrix-symmetric circuit of support size $\leq k$ and support depth $\leq \Delta$.*
2. *p is a linear combination of homomorphism polynomials $\hom_{F,n}$ for patterns F with $\mathrm{fw}_\Delta(F) \leq k-1$.*

Lean formalisation on LaxLean-Archive

Theorem 3.5 can be translated to a statement on orbit size via Theorem 2.6. Indeed, if $n \geq 20$, $n^2 \leq f(n)$, and $3\kappa_f(n) \leq n+2$, then every polynomial computed by a matrix-symmetric circuit of orbit size $\leq f(n)$ and support depth $\leq \Delta$ is a linear combination of homomorphism polynomials for patterns F with $\mathrm{fw}_\Delta(F) \leq \kappa_f(n) - 2$. Conversely, for $k \geq 1$ and $n \geq 2k+3$, linear combinations of homomorphism polynomials for patterns with $\mathrm{fw}_\Delta(F) \leq k$ admit matrix-symmetric circuits of orbit size $\leq n^{k+1}$ and support depth $\leq \Delta$. The two directions do not match in terms of orbit size: for instance, for a fixed pattern F of treewidth t and all sufficiently large n, every matrix-symmetric circuit computing $\hom_{F,n}$ has orbit size $\geq \binom{n}{t+1}$.

As a corollary, we observe that support depth and depth agree up to a factor of two, at the expense of the circuit size. Clearly, every matrix-symmetric circuit of depth $\leq \Delta$ has support depth $\leq \Delta$. Corollary 3.6 gives the converse:

Corollary 3.6. *Let $n \geq 20$, $\Delta \geq 1$, and $n^2 \leq s$ such that $\kappa := \kappa_s(n)$ satisfies $3\kappa \leq n+2$. Every polynomial computed by a matrix-symmetric circuit of orbit size $\leq s$ and support depth $\leq \Delta$ is computed by a reduced matrix-symmetric circuit of depth $\leq 2\Delta+2$, orbit size $\leq n^{\kappa-1}$, support size $\leq \kappa - 1$, and support depth $\leq \Delta$.*

[2] In the definition of Δ-treewidth of [Bha+25], the forest decomposition is taken to be a tree decomposition.

Proof. Let p be computed by such a circuit. By Lemma 2.4, we may suppose that it is reduced. By Theorem 2.6, it has support size $\leq \kappa - 1$, where $\kappa \geq 3$ as $s \geq n^2 > \binom{n}{2}$, and $n \geq 2(\kappa-1)+1$ as $3\kappa \leq n+2$. By Theorem 3.5, p is a linear combination of $\hom_{F,n}$ for patterns F with $\mathrm{fw}_\Delta(F) \leq \kappa - 2$. The circuit constructed for this linear combination in Section B.1 has the claimed properties by Theorem B.2 and Observation 2.5, where the output sum gate adds one to the depth. □

3.2. Separation

Towards Theorem 3.1, it remains to argue that the homomorphism polynomials for $T_{\Delta,n}$ cannot be written as linear combinations of homomorphism polynomials of patterns of Δ-forestwidth $o(\log n)$. This will be implied by the following general theorem.

Lean formalisation on LaxLean-Archive

Theorem 3.7. *Let $\mathcal{F}$ be a class of simple graphs. Suppose that $\mathcal{F}$ contains the 1-vertex graph K_1, that $\mathcal{F}$ is minor-closed, and that $\mathcal{F}$ is closed under weak oddomorphism. Let $G \notin \mathcal{F}$ be a connected graph. Let F be a bipartite multigraph containing G as a minor. Let $N \geq \gamma(G) := \sum_{v \in V(G)} 2^{\deg_G(v)-1}$. Then $\hom_{F,N} \notin \mathfrak{H}_N(\mathcal{F})$.*

Here, for a graph class $\mathcal{F}$ and $n \geq 1$, we write $\mathfrak{H}_n(\mathcal{F})$ for the $\mathbb{K}$-linear span of the homomorphism polynomials $\hom_{F,n}$ for all bipartite multigraphs F whose underlying simple graph is in $\mathcal{F}$.

The statement playing the same role as Theorem 3.7 in [DPS26a] differs in that it only permits G to be an induced minor of F, and its size bounds are less tight and thus insufficient for our needs. Moreover, the proof of Theorem 3.7 is more direct.

The assumption in Theorem 3.7 that $\mathcal{F}$ is closed under weak oddomorphism requires an explanation: A *weak oddomorphism* [Rob26] is a homomorphism between two graphs satisfying certain parity constraints. These constraints arise from analysing the CFI construction [CFI92], the most versatile method for constructing highly indistinguishable graphs. A graph class $\mathcal{F}$ is *closed under weak oddomorphisms* if, for every $F \in \mathcal{F}$ and every weak oddomorphism $F \to G$, it holds that $G \in \mathcal{F}$.

Roberson [Rob26] showed that graph classes $\mathcal{F}$ that are closed under weak oddomorphisms (and under disjoint unions and taking summands) are *homomorphism distinguishing closed*, i.e., for every $F' \notin \mathcal{F}$, there exist G and H that admit the same number of homomorphisms $\hom(F,G) = \hom(F,H)$ from every $F \in \mathcal{F}$ but $\hom(F',G) \neq \hom(F',H)$. In this case, $\hom_{F',N} \notin \mathfrak{H}_N(\mathcal{F})$ for $N \geq |V(G)|, |V(H)|$. Our Theorem 3.7 combines this result with a compact reduction to graph minors inspired by [Sep24b].

The classes of graphs of treewidth $\leq k$, pathwidth $\leq k$, and treedepth $\leq k$ are closed under weak oddomorphisms [Neu24; Sep24a; Adl+26]. These results underpin the $\mathsf{symVF} \subsetneq \mathsf{symVS} \subsetneq \mathsf{symVP}$ hierarchy of Dwivedi, Pago, and Seppelt [DPS26b]. We prove the analogous statement (in an approximate sense) for the class of graphs of Δ-forestwidth $\leq k$ in Corollary B.18. This requires explicitly describing the forbidden minors of graphs of Δ-forestwidth $\leq k$. Combining Theorem 3.7 and Section 3.1 yields Theorem 3.1 in Section B.5.

4. Symmetric depth reduction

Hyafil [Hya79] showed that polynomials of degree d computed by circuits of size s have circuits of depth $O(\log s \log d)$ and size $s^{O(\log d)}$. Valiant, Skyum, Berkowitz, and Rackoff [Val+83] reduced the size to $\mathrm{poly}(s,d)$, with depth $O(\log d \log(sd))$ and fan-in two; their proof yields depth $O(\log d)$ if sum gates may have unbounded fan-in, see e.g. [AW26]. That is, $\mathsf{VP} \subseteq \mathsf{VAC}^1$, noting that by definition VP contains only p-families. We prove the analogue for matrix-symmetric circuits, where orbit size takes the role of size.

Theorem 4.1. symVP *and* symVAC^1 *contain the same p-families.*

The proof of Theorem 4.1 differs from [Val+83]: it combines Theorem 3.5 with the balanced tree decompositions of Theorem 4.2.

Theorem 4.2 (Bodlaender [Bod89, Theorem 4.2 and proof of Theorem 4.1])**.** *Let $n, k \geq 1$. Every n-vertex graph F of treewidth $\leq k$ admits a tree decomposition with bags of size $\leq 3(k+1)$ whose tree is a rooted binary tree of depth $\leq 2\lceil \log_{\frac{5}{4}}(2n) \rceil$.*

Proof of Theorem 4.1. By definition, $\mathsf{symVAC}^1 \subseteq \mathsf{symVP}$, so we only need to show the other inclusion. Let $(p_n) \in \mathsf{symVP}$ be a p-family. By Theorem 2.9, there exists a constant $k \geq 1$ such that, for all sufficiently large n, p_n can be written as a linear combination of homomorphism polynomials

$$p_n = \sum_F \alpha_{F,n} \hom_{F,n}$$

for bipartite multigraphs F of treewidth $\mathrm{tw}(F) \leq k$. The degree of $\hom_{F,n}$ is equal to the number of edges of F. Since $|V(F)| \leq 2|E(F)|$ as F contains wlog no isolated vertices by Observation 2.7, we have that the number of vertices in F is bounded by a polynomial. By Theorem 4.2, the patterns F admit a tree decomposition of width $\leq 3(k+1)$ and depth $O(\log n)$. By the easy direction of Theorem 3.5, i.e., Theorem B.2, it follows that $(p_n) \in \mathsf{symVAC}^1$. □

Following the same idea as in the above proof but making the parameter dependencies more precise, we show the following symmetric analogue of [Val+83]. In contrast to Theorem 4.1, Theorem 4.3 also applies to circuits of superpolynomial orbit size. The theorem is proven in Section C.1.

Theorem 4.3. *Let $\varepsilon > 0$. Let (p_n) be a family of matrix-symmetric polynomials of degree $\leq d(n)$ computed by matrix-symmetric circuits of orbit size $\leq s(n) \leq 2^{n^{1-\varepsilon}}$. Then (p_n) is computed by matrix-symmetric circuits of depth $O(\log d(n))$ and orbit size $\leq \max\{s(n), n^2\}^{3/\varepsilon}$ for all sufficiently large n.*

We finally note that the depth $O(\log d(n))$ in Theorem 4.3 cannot be improved, even if the orbit size may grow polynomially. Let P_m denote the m-vertex path.

Theorem 4.4. *Let $a \geq 1$ and $d\colon \mathbb{N} \to \mathbb{N}$ with $2 \leq d(n) \leq n^a$.*

1. *The polynomials $\hom_{P_{d(n)},n}$ have degree $< d(n)$ and are computed by matrix-symmetric circuits of orbit size $\leq n^3$ and depth $O(\log d(n))$.*
2. *For every c, every matrix-symmetric circuit of orbit size $\leq n^c$ computing $\hom_{P_{d(n)},n}$ has support depth, and hence depth, $\Omega_{c,a}(\log d(n))$.*

5. Symmetric polylogarithmic-depth hierarchy

We explore the relationship between the bounded-depth complexity classes symVAC^k for $k \geq 0$ and symVF, symVBP, and symVP. Using graph-theoretic arguments, we show that $\mathsf{symVAC}^0 = \mathsf{symVF}$ in Theorem 5.1. Using invariant-theoretic techniques, we show that $\mathsf{symVBP} \subseteq \mathsf{symVAC}^1$ in Theorem 5.2. Using the separation techniques from Section 3.2, we show that this inclusion is proper, i.e., $\mathsf{symVBP} \subsetneq \mathsf{symVAC}^1$, in Theorem 5.4. Finally, we show strictness of the symVAC^k-hierarchy and that $\mathsf{symVAC} = \bigcup_{k \geq 1} \mathsf{symVAC}^k$ is a proper subclass of symVP. Here, we employ iteration number lower bounds for the Weisfeiler–Leman algorithm by Grohe, Lichter, and Neuen [GLN25].

5.1. Constant-depth reduction for formulas

We first prove a symmetric analogue of a result of Brent [Bre74] asserting that every formula of size s can be rewritten to a formula of depth $O(\log s)$ and size $\mathrm{poly}(s)$.

Theorem 5.1. $\mathsf{symVAC}^0 = \mathsf{symVF}$.

Proof. By transforming tree decompositions into elimination forests and vice versa, cf. [Adl+26], every non-empty graph F and every $\Delta \geq 1$ satisfy

$$\mathrm{fw}_{\mathrm{td}(F)}(F) \leq \mathrm{td}(F) - 1 \qquad \text{and} \qquad \mathrm{td}(F) \leq (\mathrm{fw}_{\Delta}(F) + 1)\Delta. \tag{3}$$

Moreover, $\mathrm{fw}_{\Delta}(F) \leq \mathrm{fw}_{\Delta'}(F)$ for $\Delta \geq \Delta'$.

Let $(p_n) \in \mathsf{symVF}$. By Theorem 2.9, for some $\Delta \geq 2$ and all sufficiently large n, p_n is a linear combination of homomorphism polynomials for patterns of treedepth $\leq \Delta$. By (3), these patterns satisfy $\mathrm{fw}_{\Delta}(F) \leq \mathrm{fw}_{\mathrm{td}(F)}(F) \leq \Delta - 1$. By Theorem 3.5 and Observation 2.5, p_n is computed by matrix-symmetric circuits of support depth $\leq \Delta$ and orbit size $\leq n^{\Delta}$ for $n \geq 2\Delta + 1$. Hence, $(p_n) \in \mathsf{symVAC}^0_{\Delta}$.

Conversely, let $(p_n) \in \mathsf{symVAC}^0_{\Delta}$ be computed by matrix-symmetric circuits of orbit size $\leq n^c$ and support depth $\leq \Delta$, where $c \geq 2$ is an integer and $\Delta \geq 1$. For all sufficiently large n, we have $\kappa_{n^c}(n) \leq c + 1$ and $3(c+1) \leq n + 2$, so the rigidification has support size $\leq c$ by Theorem 2.6. By Theorem 3.5, p_n is a linear combination of homomorphism polynomials for patterns with $\mathrm{fw}_{\Delta}(F) \leq c - 1$, and thus $\mathrm{td}(F) \leq c\Delta$ by (3). Hence, $(p_n) \in \mathsf{symVF}$ by Theorem 2.9. □

5.2. Bounded pathwidth in logarithmic support depth

We now turn to the relationship between symVAC^1 and symVBP. By Theorem 2.9 and Theorem 3.5, the complexity classes symVAC^1 and symVBP are characterised as the polynomials that can be written, for some constant c, as linear combinations of homomorphism polynomials for patterns F_n of $\mathrm{fw}_{c \log n}(F_n) \leq c$ and $\mathrm{pw}(F_n) \leq c$, respectively. We show the following.

Theorem 5.2. $\mathsf{symVBP} \subseteq \mathsf{symVAC}^1$.

This is notable since the exponentially long paths P_{2^n} have pathwidth 1, but for every constant c and all sufficiently large n, it holds that $\mathrm{fw}_{c \log n}(P_{2^n}) \geq c + 1$ by Example B.11 and Lemma B.10. The proof of Theorem 5.2 is therefore not combinatorial but invariant-theoretic. The key ingredient is the following consequence of the Nagata–Higman theorem [Hig56] with Razmyslov's bound [Raz74] on the nilpotency index.

Theorem 5.3 ([Pro76, Theorems 3.2–3.4], [Raz74])**.** *Suppose that $\mathbb{K}$ has characteristic zero. Let $N, m \geq 1$. Let Q be a commutative $\mathbb{K}$-algebra and $A_1, \ldots, A_m \in Q^{N \times N}$. For a word $w = w_1 \cdots w_\ell$ over $[m]$, write $A_w := A_{w_1} \cdots A_{w_\ell}$. Then every A_w is a linear combination of matrices A_u with $|u| < N^2$ whose coefficients are polynomials over $\mathbb{K}$ in the traces $\mathrm{tr}(A_v)$ with $|v| \leq N^2$.*

We encode patterns of bounded pathwidth as series compositions of bilabelled graphs, cf. [MR20; GRS25; Sep24a]. To a bilabelled graph, we may associate a homomorphism matrix counting homomorphisms with prescribed label images. We apply Theorem 5.3 to these matrices. The theorem then yields that every homomorphism polynomial of a bounded pathwidth pattern is a polynomial in traces of words in homomorphism matrices of polynomial length. Applying traces amounts to identifying in- and out-labels of bilabelled graphs and thus preserving bounded pathwidth. Finally, we apply Theorem 4.2 to reduce to logarithmic depth. The full proof of Theorem 5.2 is conducted in Section C.2. Finally, we observe that the inclusion of Theorem 5.2 is strict.

Theorem 5.4. $\mathsf{symVBP} \subsetneq \mathsf{symVAC}^1$.

Proof. Let B_n be the complete binary tree with n leaves, as in [DPS26b]. Since $\mathrm{pw}(B_n) \in \Theta(\log n)$ is unbounded, $(\hom_{B_n,n}) \notin \mathsf{symVBP}$ by [DPS26b]. The bags $\{v, \mathrm{parent}(v)\}$ together with an empty root bag show that $\mathrm{fw}_\Delta(B_n) \leq 1$ for $\Delta \in O(\log n)$. By Theorem 3.5 and Observation 2.5, $\hom_{B_n,n}$ is computed by matrix-symmetric circuits of orbit size $\leq n^2$ and support depth $O(\log n)$. □

5.3. Separating symVAC^1 from symVP

Using the robust lower bounds on the iteration number of the Weisfeiler–Leman algorithm of Grohe, Lichter, and Neuen [GLN25], we show that some families in symVP require superpolylogarithmic support depth. By Theorem 4.1, the witnesses of this separation are not p-families.

Theorem 5.5. $\mathsf{symVAC} \subsetneq \mathsf{symVP}$.

Theorem 5.5 is inferred from the following result.

Theorem 5.6 ([GLN25, Theorem 1.2])**.** *There are absolute constants $k_0 \in \mathbb{N}$ and $\alpha, \varepsilon > 0$ such that, for every $d \geq k \geq k_0$ and every $n \geq \alpha d^8 k^6$, there is a pair of k-ary n-element relational structures $\mathfrak{A}$ and $\mathfrak{B}$ that are distinguished by k-WL, but d-WL does not distinguish $\mathfrak{A}$ and $\mathfrak{B}$ after $n^{\varepsilon k}$ refinement rounds.*

We fix $k := k_0$, choose $d \approx n^{1/8}$, and transfer $\mathfrak{A}$ and $\mathfrak{B}$ to a pair of $N \times N$ matrices with $N := (n+1)^k$. By Theorem 3.5, circuits of orbit size 2^{N^η} and support depth Δ compute linear combinations of homomorphism polynomials for patterns F with $\mathrm{fw}_\Delta(F) < N^\eta$. For small $\eta > 0$, we have $kN^\eta \leq d$, so these polynomials cannot distinguish the two matrices unless $kN^\eta \Delta > n^{\varepsilon k}$, since d-WL does not distinguish $\mathfrak{A}$ and $\mathfrak{B}$ within $n^{\varepsilon k}$ refinement rounds. Hence, $\Delta \geq N^\theta$ for some $\theta > 0$. The proof is given in Section C.3. A version of the lower bounds of Grohe, Lichter, and Neuen [GLN25] with a tunable number of rounds separates the levels of symVAC.

Theorem 5.7. *For every $k \geq 1$, $\mathsf{symVAC}^k \subsetneq \mathsf{symVAC}^{k+1}$.*

6. Symmetric algebraic branching programs

We introduce symmetric algebraic branching programs, the symmetric analogue of algebraic branching programs. We then show that the resulting class symVBP coincides with symVS.

We state the results for square variable matrices $\mathcal{X}_n$ and polynomial families (f_n). All symmetric classes in this section refer to this setting, with $\Gamma_n = \mathbf{Sym}_n \times \mathbf{Sym}_n$. The same arguments apply to rectangular variable matrices, with orbit size polynomial in the sum of the two dimensions.

An *algebraic branching program* (ABP) in variables X is a finite directed acyclic graph $G = (V, E, \ell)$ with a unique source s and a unique sink t. Each edge e is labelled by an affine linear form $\ell(e)$ over $\mathbb{K}$ in X. The polynomial computed by G is the sum, over all directed paths from s to t, of the product of their edge labels.

Definition 6.1 (Symmetric ABPs)**.** Let $G = (V, E, \ell)$ be an ABP in variables X and let $\Gamma \leq \mathbf{Sym}(X)$. We call G Γ-*symmetric* if every $\pi \in \Gamma$ extends to an automorphism σ of (V, E) fixing s and t such that

$$\ell(\sigma(e)) = \pi(\ell(e)) \qquad \text{for every } e \in E.$$

Here π permutes the variables and fixes field constants. Let $\mathbf{Aut}_\Gamma(G)$ be the group of all automorphisms extending some $\pi \in \Gamma$. The *orbit size* of G is

$$\mathrm{maxOrb}(G) = \max_{v \in V} \big|\{\sigma(v) : \sigma \in \mathbf{Aut}_\Gamma(G)\}\big|.$$

Definition 6.2 (symVBP)**.** The class symVBP contains all families of matrix-symmetric polynomials $(p_n)_{n\in\mathbb{N}}$ in variables $\mathcal{X}_n$ that can be represented by Γ_n-symmetric ABPs whose orbit size is polynomial in n.

Recall from Section 2.5 that symVS consists of matrix-symmetric polynomial families computed by symmetric skew circuits of polynomial orbit size. In classical algebraic complexity, VBP = VS: polynomial-size ABPs and skew circuits compute the same polynomial families [Mah13]. The following theorem extends this equivalence to the symmetric setting.

Theorem 6.3. symVBP = symVS.

The following two lemmas prove both directions.

Lemma 6.4. symVBP ⊆ symVS.

Proof. We use the standard conversion from ABPs to skew circuits [Mah13, Section 3, p. 11] and verify that it preserves symmetry and polynomial orbit size. Let G be a symmetric ABP of orbit size B, and write $\ell(u,v) = a_{uv,0} + \sum_{x\in X} a_{uv,x}x$. For each vertex v, introduce a sum gate P_v. Give P_s the single input 1, and, for $v \neq s$, set

$$P_v = \sum_{(u,v)\in E} \left(a_{uv,0}P_u + \sum_{x\in X} a_{uv,x}xP_u \right).$$

Each term is a separate product gate feeding directly into P_v. Retain terms with zero coefficients. Every product gate has just one internal child, so the circuit is skew. Induction over G shows that P_v computes the path sum from s to v. Choose P_t as the output gate.

Every extending automorphism of G acts on this circuit by permuting the vertex and variable indices. To bound orbit size, we must consider all extending circuit automorphisms. Any such automorphism permutes the sum gates P_v, fixes P_s as the only sum gate with an input child, and fixes the output P_t. A product gate has parent P_v and unique internal child P_u, which identify the endpoints of its ABP edge. Its input labels identify the coefficient and the variable, if present. Since we retained all terms, including those with zero coefficients, this recovers every edge and its affine label. Thus every extending circuit automorphism induces an extending ABP automorphism, and sum-gate orbits have size at most B. Each product gate is determined by an ABP edge (u,v) and the choice of the constant term or a variable $x \in X$. The induced ABP automorphism gives at most B possible images for each of u and v, and each resulting edge has $|X|+1$ product gates. Thus every product-gate orbit has size at most $B^2(|X|+1)$. With input gates shared by label, variable inputs have orbits of size at most $|X|$, and constant inputs are fixed. Hence the circuit has polynomial orbit size. □

Lemma 6.5. symVS ⊆ symVBP.

Proof. Let $(f_n) \in$ symVS. By Theorem 2.9(2), there is a constant k such that, for all sufficiently large n, f_n is a finite linear combination of homomorphism polynomials $\hom_{F,n}$ with $\mathrm{pw}(F) \leq k$. For each remaining n, a homomorphism expansion uses patterns with at most n vertices on each side, hence pathwidth less than $2n$. Enlarging k covers these finitely many cases.

Fix n and one such pattern F. We recall the standard path-decomposition construction underlying the upper bound in [KPR22, Theorem 17]. The ABP chooses the images of the pattern vertices along a path decomposition. A state remembers only the images of vertices in the current bag.

Take a path decomposition with bags of size at most $k+1$, starting and ending with the empty bag. Between consecutive bags, forget vertices that leave before introducing those that enter.

Process each edge occurrence once, at a bag containing both endpoints, by repeating that bag. This gives a sequence of introduction, forgetting, and edge-processing steps without increasing the bag size.

For each bag, create a layer with one vertex for every assignment of its left and right vertices to the left and right copies of $[n]$, respectively. Connect consecutive layers according to the step:

- To introduce a vertex, connect each assignment to all its extensions, with edge label 1.
- To forget a vertex, connect each assignment to its restriction, with edge label 1.
- To process an edge ab, connect each assignment h to its copy in the next layer, with edge label $x_{h(a),h(b)}$.

The empty endpoint bags give a unique source and sink. Each path chooses one image for every pattern vertex. An image is retained until that vertex is forgotten, after all its incident edges have been processed. Thus paths correspond bijectively to the maps defining $\hom_{F,n}$, with the same products of edge labels. This also accounts for isolated vertices and parallel edge occurrences. Denote this ABP by G_F.

Claim 6.5a. The ABP G_F is matrix-symmetric and has orbit size at most n^{k+1}.

Proof of Claim. Applying $(\sigma, \tau) \in \Gamma_n$ to the left and right images, respectively, gives an extending automorphism. Every extending automorphism preserves layers, since the layer number is the directed distance from the source. Each layer contains at most n^{k+1} vertices, so every vertex orbit has at most this size. ⌟

If $f_n = 0$, use a single edge labelled 0. Otherwise, to combine the summands without enlarging the orbit bound, give their ABPs different path lengths. For each term $c_F \hom_{F,n}$, add an initial edge labelled c_F and a final edge labelled 1. Pad with edges labelled 1 until the branches have distinct source-to-sink path lengths, then identify their sources and sinks. The resulting path sum is f_n. Every extending automorphism preserves path lengths, hence each branch and its layers. The orbit bound remains n^{k+1}, proving $(f_n) \in \mathsf{symVBP}$. □

7. Symmetric VNP

We introduce symVNP, the analogue of VNP for symmetric algebraic circuits.

As in the previous section, we state the definitions and results for square matrices $\mathcal{X}_n$ under Γ_n, with all symmetric classes understood for one-parameter families. The results also hold for rectangular matrices.

Definition 7.1. The class symVNP consists of all families of polynomials $(f_n) \in \mathbb{K}[\mathcal{X}_n]$ for which there exist a p-bounded function $q\colon \mathbb{N} \to \mathbb{N}$, a family of polynomials (g_n) over variables $\mathcal{X}_n \cup \mathcal{Y}_n$ where $|\mathcal{Y}_n| = q(n)$, and a group homomorphism $\rho_n\colon \Gamma_n \to \mathbf{Sym}(\mathcal{Y}_n)$ such that:

$$f_n(\mathcal{X}_n) \;=\; \sum_{\boldsymbol{e} \in \{0,1\}^{q(n)}} g_n(\mathcal{X}_n, \boldsymbol{e}),$$

the family (g_n) is computed by a circuit family (C_n) of polynomially bounded orbit size, and for every $\pi \in \Gamma_n$, the permutation on $\mathcal{X}_n \cup \mathcal{Y}_n$ defined by applying π to $\mathcal{X}_n$ and $\rho_n(\pi)$ to $\mathcal{Y}_n$ extends to a circuit automorphism σ of the circuit C_n computing g_n. Orbit size is measured with respect to this joint Γ_n-action.

Unlike the definitions of symVP and the other symmetric complexity classes, this definition does not require the output polynomials to be matrix-symmetric. Instead, matrix symmetry follows

from the joint symmetry of the witness and summation over the entire Boolean cube, as the following observation shows. Following [DPS26a; DPS26b], write $\mathfrak{G}_n \subseteq \mathbb{K}[\mathcal{X}_n]$ for the space of matrix-symmetric polynomials.

Observation 7.2. *If $(f_n) \in$ symVNP, then $f_n \in \mathfrak{G}_n$ for every n.*

Proof. Fix n and write $X = \mathcal{X}_n$, $Y = \mathcal{Y}_n$, $f = f_n$, $g = g_n$, and $\rho = \rho_n$ for a witness in the definition. Let $\gamma \in \Gamma_n$. The joint permutation of $X \cup Y$ extends to a circuit automorphism fixing the unique output gate. Hence

$$g(\gamma X, \rho(\gamma)Y) = g(X, Y).$$

Since $\rho(\gamma)$ permutes Y, it induces a bijection on the set of Boolean assignments $\{0,1\}^Y$. Reindexing the sum gives

$$f(\gamma X) = \sum_{\boldsymbol{e} \in \{0,1\}^Y} g(\gamma X, \boldsymbol{e}) = \sum_{\boldsymbol{e}' \in \{0,1\}^Y} g(\gamma X, \rho(\gamma)\boldsymbol{e}') = \sum_{\boldsymbol{e}' \in \{0,1\}^Y} g(X, \boldsymbol{e}') = f(X).$$

Thus f is matrix-symmetric, so $f \in \mathfrak{G}_n$. □

Later, we will show that the permanent is complete for symVNP under symmetric projections.

7.1. symVNP contains all matrix-symmetric polynomials

By Observation 7.2, every family $(f_n) \in$ symVNP satisfies $f_n \in \mathfrak{G}_n$ for all n. We now prove the converse, showing that symVNP consists precisely of all families of matrix-symmetric polynomials. We express each $f \in \mathfrak{G}_n$ as a linear combination of homomorphism polynomials and encode the vertex maps using Boolean variables, giving witnesses computed by symmetric formulas of polynomial orbit size.

Theorem 7.3. *Over a field $\mathbb{K}$ of characteristic zero, a polynomial family (f_n) belongs to* symVNP *if and only if $f_n \in \mathfrak{G}_n$ for every n. More precisely, every $f \in \mathfrak{G}_n$ can be written as*

$$f(X) = \sum_{\boldsymbol{e} \in \{0,1\}^{2n^2}} g(X, \boldsymbol{e}),$$

where g admits a symmetric formula whose gate orbits have size at most $4n^2$, under a joint action on X and the auxiliary variables.

Proof. The forward implication is Observation 7.2. For the converse, fix n and let $f \in \mathfrak{G}_n$. By the argument of [DPS26a, Lemma 4], f is a finite linear combination of homomorphism polynomials of patterns with at most n vertices on each side. The argument applies over $\mathbb{K}$ since it only requires division by nonzero integers. If a pattern F has a left and b right vertices, padding it with isolated vertices to obtain $\widehat{F}$ with n vertices on each side gives $\hom_{\widehat{F},n} = n^{2n-a-b} \hom_{F,n}$ by Observation 2.7 and its right-side analogue. Absorbing these nonzero factors into the coefficients and denoting the padded patterns again by F, we can therefore write

$$f(X) = \sum_F c_F \hom_{F,n}(X),$$

with all patterns F on the same bipartition $A \uplus B$, where $|A| = |B| = n$.

Introduce variables $y_{a,i}$ for $(a,i) \in A \times [n]$ and $y_{b,j}$ for $(b,j) \in B \times [n]$. Let Y be this common set of $2n^2$ variables. For each pattern F, put

$$W_F(X,Y) \;=\; \prod_{ab \in E(F)} \left(\sum_{i=1}^{n} \sum_{j=1}^{n} y_{a,i}\, y_{b,j}\, x_{i,j} \right), \tag{4}$$

counting edges with multiplicity, and define

$$V_A(Y) = \prod_{a \in A} \left[\underbrace{\left(\sum_{i=1}^{n} y_{a,i} \right)}_{\substack{\text{rules out} \\ \text{all zeros}}} \underbrace{\prod_{1 \le i < k \le n} (1 - y_{a,i} y_{a,k})}_{\text{at most one 1}} \right], \quad V_B(Y) = \prod_{b \in B} \left[\left(\sum_{j=1}^{n} y_{b,j} \right) \prod_{1 \le j < \ell \le n} (1 - y_{b,j} y_{b,\ell}) \right],$$

$$g_F(X,Y) \;=\; W_F(X,Y)\, V_A(Y)\, V_B(Y).$$

On Boolean assignments, $V_A V_B$ is 1 precisely when each vertex chooses exactly one image, and is 0 otherwise: the sum rules out choosing no image, and the pairwise factors rule out choosing two. Such assignments correspond bijectively to maps $h\colon A \to [n]$, $h\colon B \to [n]$. On the assignment encoding h, each edge factor in W_F equals $x_{h(a),h(b)}$. Hence

$$\sum_{\boldsymbol{e} \in \{0,1\}^Y} g_F(X, \boldsymbol{e}) = \sum_{\substack{h\colon A \to [n] \\ h\colon B \to [n]}} \prod_{ab \in E(F)} x_{h(a),h(b)} = \hom_{F,n}(X).$$

Claim 7.3a. Each $\hom_{F,n}$ above is a Boolean-cube sum of a polynomial g_F admitting a symmetric formula with orbit size at most $4n^2$.

Proof of Claim. The Boolean-sum identity is established above, and $|Y| = 2n^2$. Compute the displayed expression as a formula, using fresh copies for every occurrence of an internal gate and sharing only input gates. In Equation (4), every edge occurrence contributes one factor $\sum_{i,j} y_{a,i} y_{b,j} x_{i,j}$. If two edges are parallel, their factors are identical polynomials, but we compute them with separate copies.

For each copy of the edge factor $P = \sum_{i,j} y_{a,i} y_{b,j} x_{i,j}$ in Equation (4), compute $t^{-1}(tP)$ instead. This adds two multiplication gates without changing P. Choose all constants t, t^{-1} distinct and absent from the original formula. This is possible since $\mathbb{K}$ is infinite. Every extending automorphism fixes these constants. Each has only one parent, so both new gates are fixed. The root of this copy of P is therefore fixed too. Thus identical copies cannot be exchanged.

Extend $(\sigma, \tau) \in \Gamma_n$ to Y by

$$y_{a,i} \mapsto y_{a,\sigma(i)}, \qquad y_{b,j} \mapsto y_{b,\tau(j)},$$

keeping the pattern vertices fixed. Together with its action on X, this permutes the gates computing $y_{a,i} y_{b,j} x_{i,j}$ for each fixed ab, and the computations of $1 - y_{a,i} y_{a,k}$ and $1 - y_{b,j} y_{b,\ell}$ for each fixed a and b. Conversely, every extending automorphism preserves the marked edge-factor copies. Within each copy, the variable inputs determine the term indices; in the vertex factors they determine the vertex and the unordered pair of indices. Thus the bounds apply to all extending automorphisms: the term gates have orbits of size at most n^2, and the pair-indexed gates have orbits of size at most $\binom{n}{2}$. Every addition or multiplication combining all these terms stays fixed, since its inputs are merely permuted. Input gates have orbits of size at most n^2 or n, and constants are fixed, so every gate orbit has size at most $4n^2$. ⌟

Finally, set $g = \sum_F c_F g_F$, using the same variables Y in every summand. By linearity,

$$\sum_{e \in \{0,1\}^Y} g(X, e) = \sum_F c_F \hom_{F,n}(X) = f(X).$$

Take a separate formula for each g_F, sharing only input gates. Compute its contribution as $t_F^{-1}(t_F c_F g_F)$. Choose all constants t_F, t_F^{-1} distinct and absent from these formulas and the coefficients c_F. Every extending automorphism fixes these constants and their unique parent gates. Hence it preserves each summand. Adding one sum gate gives a formula for g. The added gates are fixed, so every gate orbit still has size at most $4n^2$. Since $|Y| = 2n^2$, this proves the theorem. □

7.2. Hypercube sums of symmetric formulas

In classical algebraic complexity, VNP is the class of hypercube sums of VP witnesses. Valiant [Val82] proved that using VF witnesses gives the same class; see also [Büroo, Theorem 2.13] for a proof. Malod and Portier [MPo8] gave an alternative proof based on parse trees. The following corollary shows that the same phenomenon holds for symVNP.

Here $\Sigma_{\{0,1\}}$symVP denotes Boolean sums of circuit witnesses satisfying the joint symmetry and orbit-size conditions in Definition 7.1, with polynomially many auxiliary variables. We use $\Sigma_{\{0,1\}}$symVF for the same definition with formula witnesses.

Corollary 7.4. *Over a field $\mathbb{K}$ of characteristic zero, let $(f_n) \in$ symVNP. There exist a p-bounded function $q\colon \mathbb{N} \to \mathbb{N}$ and polynomials $h_n \in \mathbb{K}[\mathcal{X}_n \cup Y_n]$, with $|Y_n| = q(n)$, such that*

$$f_n(X) = \sum_{e \in \{0,1\}^{Y_n}} h_n(X, e).$$

Each h_n admits a formula symmetric under a joint Γ_n-action on $\mathcal{X}_n \cup Y_n$ extending the usual row and column action on $\mathcal{X}_n$. These formulas have polynomially bounded gate orbits. In short, we write

$$\text{symVNP} = \Sigma_{\{0,1\}}\text{symVP} = \Sigma_{\{0,1\}}\text{symVF}.$$

Proof. The first equality is the definition of symVNP. Since every symmetric formula is a symmetric circuit, $\Sigma_{\{0,1\}}$symVF $\subseteq$ symVNP. Conversely, let $(f_n) \in$ symVNP. By Observation 7.2, each f_n belongs to $\mathfrak{G}_n$. By Theorem 7.3, it has a symmetric formula witness using $2n^2$ Boolean variables and with gate orbits of size at most $4n^2$. Thus $(f_n) \in \Sigma_{\{0,1\}}$symVF, proving the reverse inclusion. □

7.3. Completeness of the permanent under symmetric projections

For an $n \times n$ variable matrix X, the *permanent polynomial* is

$$\text{perm}_n(X) = \sum_{\pi \in \mathbf{Sym}_n} \prod_{i=1}^{n} x_{i,\pi(i)}.$$

Proposition 7.5. *The permanent family belongs to* symVNP *over any field $\mathbb{K}$.*

Proof. Ryser's formula gives

$$\text{perm}_n(X) = \sum_{S \subseteq [n]} (-1)^{n-|S|} \prod_{i=1}^{n} \left(\sum_{j \in S} x_{i,j} \right).$$

Indeed, expanding the products, the monomial $\prod_i x_{i,h(i)}$ for a map $h\colon [n] \to [n]$ has coefficient $\sum_{S \supseteq h([n])} (-1)^{n-|S|}$, which is 1 if h is surjective and 0 otherwise. The surviving maps are exactly the permutations.

Introduce $Y = \{y_1, \dots, y_n\}$ and put

$$g_n(X, Y) = \left(\prod_{j=1}^{n} (2y_j - 1) \right) \prod_{i=1}^{n} \left(\sum_{j=1}^{n} x_{i,j} y_j \right).$$

A Boolean assignment $\boldsymbol{e}$ encodes the set $S = \{j : e_j = 1\}$. On this assignment, $\prod_j (2e_j - 1) = (-1)^{n-|S|}$ and $\sum_j x_{i,j} e_j = \sum_{j \in S} x_{i,j}$. Hence

$$\operatorname{perm}_n(X) = \sum_{\boldsymbol{e} \in \{0,1\}^n} g_n(X, \boldsymbol{e}).$$

The circuit given by this expression is symmetric under the action $\rho_n(\sigma, \tau)(y_j) = y_{\tau(j)}$ on Y. Together with $x_{i,j} \mapsto x_{\sigma(i),\tau(j)}$, this sends the gate computing $x_{i,j} y_j$ to the one computing $x_{\sigma(i),\tau(j)} y_{\tau(j)}$, the sum for row i to the sum for row $\sigma(i)$, and the computation of $2y_j - 1$ to that of $2y_{\tau(j)} - 1$. The two products and the output gate are fixed. These maps preserve operations and wires and extend every (σ, τ) to a circuit automorphism.

The gates indexed by (i, j) have orbits of size at most n^2, those indexed by i or j have orbits of size at most n, and all other gates are fixed. Thus g_n has polynomial orbit size (indeed, total size $O(n^2)$) and uses only n auxiliary variables. □

We now show that the permanent is also complete for symVNP under symmetric projections. The construction uses one directed cycle for each monomial and all its relabellings.

Definition 7.6 (Symmetric projection to the permanent)**.** Let $\Gamma = \Gamma_n$ act on $X = \mathcal{X}_n$, and fix an action $\rho\colon \Gamma \to \mathbf{Sym}(Y)$ on a finite auxiliary variable set Y. A polynomial $h \in \mathbb{K}[X \cup Y]$ is a *symmetric projection of the permanent* if $h = \operatorname{perm}_A(M)$ for some $A \geq 1$ and an $A \times A$ matrix M of affine linear forms over $\mathbb{K}$ in $X \cup Y$, and every $\gamma \in \Gamma$ admits permutations $\alpha_\gamma, \beta_\gamma \in \mathbf{Sym}_A$ such that

$$M_{\alpha_\gamma(i),\beta_\gamma(j)} = \gamma \cdot M_{i,j} \qquad (i, j \in [A]).$$

Here γ acts on Y by $\rho(\gamma)$ and fixes field constants. For a polynomial in $\mathbb{K}[X]$, take $Y = \emptyset$.

For polynomial families, we apply this definition to each f_n, without any bound on the dimension $A = A(n)$ of the matrix M. The permanent family is *complete* for symVNP under symmetric projections if it belongs to symVNP and every family $(f_n) \in$ symVNP admits such projections. The construction below uses only variables and field constants as matrix entries.

Theorem 7.7 (Symmetric completeness of the permanent)**.** *Over a field $\mathbb{K}$ of characteristic zero, the permanent family is complete for* symVNP *under symmetric projections.*

Recall that a *directed cycle cover* is a collection of vertex-disjoint directed cycles covering all vertices, with loops allowed. The permanent of a weighted adjacency matrix sums the weights of these covers, where each weight is the product of the selected edge labels.

Proof. By Observation 7.2, it suffices to construct a symmetric projection for each matrix-symmetric polynomial $f \in \mathbb{K}[X]$, where $X = \mathcal{X}_n$. Put $\Gamma = \Gamma_n$ and write

$$f = \sum_{r \in R} c_r \prod_{j=1}^{d_r} z_{r,j}, \qquad z_{r,j} \in X.$$

We encode the average of the polynomials $\gamma(f)$ using the standard correspondence between paths and cycle covers [MV97, Section 6.2].

Start with a vertex v. For each $r \in R$ and $\gamma \in \Gamma$, add a directed cycle $C_{r,\gamma}$ through v, with edge labels

$$\frac{c_r}{|\Gamma|}, \quad \gamma(z_{r,1}), \ldots, \gamma(z_{r,d_r}), \quad 1.$$

These cycles share only v. Keep separate copies for all $\gamma \in \Gamma$, even when their monomials agree, so that group multiplication can permute the copies. Add a loop of weight 1 at every vertex other than v, and let M be the weighted adjacency matrix, with absent edges represented by 0. A constant monomial gives a two-cycle; for $f = 0$, take the 1×1 zero matrix.

Every cycle cover selects exactly one cycle $C_{r,\gamma}$ through v and uses loops at all remaining vertices. Summing their weights gives

$$\operatorname{perm}(M) = \sum_{\gamma \in \Gamma} \sum_{r \in R} \frac{c_r}{|\Gamma|} \prod_{j=1}^{d_r} \gamma(z_{r,j}) = \frac{1}{|\Gamma|} \sum_{\gamma \in \Gamma} \gamma(f) = f,$$

where the last equality holds because f is matrix-symmetric.

For $\delta \in \Gamma$, let η_δ fix v and send each other vertex of $C_{r,\gamma}$ to the vertex in the same position on $C_{r,\delta\gamma}$, counting from v along the cycle. An edge labelled $\gamma(z_{r,j})$ is then sent to one labelled $\delta\gamma(z_{r,j})$, its image under δ. Constant labels and the added loops are preserved. Thus

$$M_{\eta_\delta(u),\eta_\delta(w)} = \delta \cdot M_{u,w}.$$

Taking $\alpha_\delta = \beta_\delta = \eta_\delta$ in Definition 7.6 proves that f is a symmetric projection of the permanent.

Applying the construction for every n and using the permanent's membership in symVNP from Proposition 7.5 proves completeness. □

Acknowledgements

Prateek Dwivedi is supported by the PRIME programme of the German Academic Exchange Service (DAAD) with funds from the German Federal Ministry of Education and Research (BMBF). He also thanks Université Savoie Mont Blanc for its hospitality and the University of Cambridge for hosting his visit to Anuj Dawar and Benedikt Pago, during which the ideas for this project first took shape.

Tim Seppelt is funded by the European Union (CountHom, No. 101077083). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

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

AI statement

We used Claude Code with Sonnet 5.5, Opus 5.5 and Fable 5.1 to assist with carrying out technical arguments following human instructions, literature review, and formalising the results in Lean [Daw+26]. The tool materially affected Sections 3, 5.2, and 5.3 and the corresponding appendices. The authors verified the correctness and originality of all content, including references.

References

[Adl+26] Isolde Adler, Eva Fluck, Tim Seppelt, and Gian Luca Spitzer. 'Going deep and going wide: Counting logic and homomorphism indistinguishability over graphs of bounded treedepth and treewidth'. In: *Logical Methods in Computer Science* Volume 22, Issue 2, 33 (June 2026) (cited on pages 4, 11, 13, 43, 44).

[AV08] Manindra Agrawal and V. Vinay. 'Arithmetic Circuits: A Chasm at Depth Four'. In: *49th Annual IEEE Symposium on Foundations of Computer Science*. 2008, pp. 67–75 (cited on page 6).

[AW26] Robert Andrews and Avi Wigderson. 'Constant-Depth Arithmetic Circuits for Linear Algebra Problems'. In: *SIAM Journal on Computing* (Jan. 2026), FOCS24–236–FOCS24–295 (cited on pages 3, 11).

[BDS25] C. S. Bhargav, Prateek Dwivedi, and Nitin Saxena. 'Lower bounds for the sum of small-size algebraic branching programs'. In: *Theor. Comput. Sci.* 1041 (2025), p. 115214 (cited on page 6).

[Bha+25] C.S. Bhargav, Shiteng Chen, Radu Curticapean, and Prateek Dwivedi. 'Monotone Bounded-Depth Complexity of Homomorphism Polynomials'. In: *50th International Symposium on Mathematical Foundations of Computer Science (MFCS 2025)*. Ed. by Paweł Gawrychowski, Filip Mazowiecki, and Michał Skrzypczak. Vol. 345. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025, 19:1–19:18 (cited on pages 4, 10, 28, 34).

[Bod89] Hans L. Bodlaender. 'NC-algorithms for graphs with small treewidth'. In: *Graph-Theoretic Concepts in Computer Science*. Ed. by G. Goos et al. Vol. 344. Berlin, Heidelberg: Springer Berlin Heidelberg, 1989, pp. 1–10 (cited on page 12).

[Bre74] Richard P. Brent. 'The Parallel Evaluation of General Arithmetic Expressions'. In: *Journal of the ACM* 21.2 (Apr. 1974), pp. 201–206 (cited on pages 3, 6, 13).

[BCE95] Nader H. Bshouty, Richard Cleve, and Wayne Eberly. 'Size-Depth Tradeoffs for Algebraic Formulas'. In: *SIAM Journal on Computing* 24.4 (1995), pp. 682–705 (cited on page 6).

[Büro0] Peter Bürgisser. *Completeness and Reduction in Algebraic Complexity Theory*. Ed. by E. Becker, M. Bronstein, H. Cohen, D. Eisenbud, and R. Gilman. Vol. 7. Algorithms and Computation in Mathematics. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000 (cited on page 19).

[CFI92] Jin-Yi Cai, Martin Fürer, and Neil Immerman. 'An optimal lower bound on the number of variables for graph identification'. In: *Combinatorica* 12.4 (1992), pp. 389–410 (cited on pages 11, 36, 46).

[Cha+24] Prerona Chatterjee, Deepanshu Kush, Shubhangi Saraf, and Amir Shpilka. 'Lower Bounds for Set-Multilinear Branching Programs'. In: *39th Computational Complexity Conference, CCC 2024, Ann Arbor, MI, USA, July 22-25, 2024*. Ed. by Rahul Santhanam. Vol. 300. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024, 20:1–20:20 (cited on page 6).

[Chi+18] Suryajith Chillara, Christian Engels, Nutan Limaye, and Srikanth Srinivasan. 'A Near-Optimal Depth-Hierarchy Theorem for Small-Depth Multilinear Circuits'. In: *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. Paris, France: IEEE, Oct. 2018, pp. 934–945 (cited on page 1).

[Cur22] Radu Curticapean. 'Determinants from Homomorphisms'. In: *30th Annual European Symposium on Algorithms (ESA 2022)*. Vol. 244. Leibniz International Proceedings in Informatics (LIPIcs). 2022, 38:1–38:7 (cited on page 6).

[Daw+26] Anuj Dawar, Prateek Dwivedi, Benedikt Pago, and Tim Seppelt. *Depth Hierarchies for Symmetric Algebraic Circuits*. Lax Archive, lax-461425. 2026 (cited on page 21).

[Daw+25] Anuj Dawar, Erich Grädel, Leon Kullmann, and Benedikt Pago. 'Symmetric Proofs in the Ideal Proof System'. In: *50th International Symposium on Mathematical Foundations of Computer Science, MFCS 2025, Warsaw, Poland, August 25-29, 2025*. Ed. by Pawel Gawrychowski, Filip Mazowiecki, and Michal Skrzypczak. Vol. 345. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025, 40:1–40:18 (cited on page 6).

[DJR21] Anuj Dawar, Tomáš Jakl, and Luca Reggio. 'Lovász-Type Theorems and Game Comonads'. In: *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS*. IEEE, 2021, pp. 1–13 (cited on page 44).

[DPS26a] Anuj Dawar, Benedikt Pago, and Tim Seppelt. 'Symmetric Algebraic Circuits and Homomorphism Polynomials'. In: *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*. Ed. by Shubhangi Saraf. Vol. 362. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2026, 46:1–46:15 (cited on pages 1, 3–11, 17, 28, 30, 31, 36, 38, 39).

[DW20] Anuj Dawar and Gregory Wilsenach. 'Symmetric Arithmetic Circuits'. In: *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*. Ed. by Artur Czumaj, Anuj Dawar, and Emanuela Merelli. Vol. 168. Leibniz International Proceedings in Informatics (LIPIcs). ISSN: 1868-8969. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020, 36:1–36:18 (cited on pages 1, 3, 4).

[DW22] Anuj Dawar and Gregory Wilsenach. 'Lower Bounds for Symmetric Circuits for the Determinant'. In: *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*. Vol. 215. Leibniz International Proceedings in Informatics (LIPIcs). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022, 52:1–52:22 (cited on page 1).

[DM96] John D. Dixon and Brian Mortimer. *Permutation Groups*. Vol. 163. Graduate Texts in Mathematics. New York: Springer, 1996 (cited on page 27).

[DPS26b] Prateek Dwivedi, Benedikt Pago, and Tim Seppelt. 'Lower Bounds in Algebraic Complexity via Symmetry and Homomorphism Polynomials'. In: *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*. Salt Lake City UT USA: ACM, June 2026, pp. 631–640 (cited on pages 1–11, 14, 17).

[Elb+26] Tal Elbaz, Nashlen Govindasamy, Jiaqi Lu, and Iddo Tzameret. 'Lower Bounds against the Ideal Proof System in Finite Fields'. In: *Proceedings of the 58th Annual ACM Symposium on Theory of Computing, STOC 2026, Salt Lake City, UT, USA, June 22-26, 2026*. Ed. by Aditya Bhaskara and Artur Czumaj. ACM, 2026, pp. 1–12 (cited on page 6).

[Fab+26] Théo Borém Fabris, Nutan Limaye, Srikanth Srinivasan, and Amir Yehudayoff. 'Multilinear Algebraic Branching Programs and the Min-Partition Rank Method'. In: *41st Computational Complexity Conference, CCC 2026, Lisbon, Portugal, August 3-6, 2026*. Ed. by Dana Moshkovitz. Vol. 383. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2026, 22:1–22:20 (cited on page 6).

[For+21] Michael A. Forbes, Amir Shpilka, Iddo Tzameret, and Avi Wigderson. 'Proof Complexity Lower Bounds from Algebraic Circuit Complexity'. In: *Theory Comput.* 17 (2021), pp. 1–88 (cited on page 6).

[Fou+26] Hervé Fournier, Nutan Limaye, Guillaume Malod, Srikanth Srinivasan, and Sébastien Tavenas. 'Towards Optimal Depth-Reductions for Algebraic Formulas'. In: *Computational Complexity* 35, 3 (2026) (cited on page 6).

[GHT22] 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)*. 2022, pp. 188–199 (cited on page 6).

[GLN25] Martin Grohe, Moritz Lichter, and Daniel Neuen. 'The Iteration Number of the Weisfeiler-Leman Algorithm'. In: *ACM Transactions on Computational Logic* 26.1 (Jan. 2025), pp. 1–31 (cited on pages 2, 5, 12, 14, 43, 44, 47, 48).

[GRS25] Martin Grohe, Gaurav Rattan, and Tim Seppelt. 'Homomorphism Tensors and Linear Equations'. In: *Advances in Combinatorics* (Apr. 2025) (cited on page 13).

[Gup+16] Ankit Gupta, Pritish Kamath, Neeraj Kayal, and Ramprasad Saptharishi. 'Arithmetic Circuits: A Chasm at Depth 3'. In: *SIAM J. Comput.* 45.3 (2016), pp. 1064–1079 (cited on page 6).

[HLT24] 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, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*. ACM, 2024, pp. 1396–1404 (cited on page 6).

[Hås87] Johan Håstad. *Computational limitations of small-depth circuits*. Cambridge, MA, USA: MIT Press, 1987 (cited on page 2).

[Hås+17] Johan Håstad, Benjamin Rossman, Rocco A. Servedio, and Li-Yang Tan. 'An Average-Case Depth Hierarchy Theorem for Boolean Circuits'. In: *J. ACM* 64.5 (2017), 35:1–35:27 (cited on page 2).

[Hig56] Graham Higman. 'On a conjecture of Nagata'. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 52.1 (Jan. 1956). Ed. by P. Hall, pp. 1–4 (cited on page 13).

[Hya79] Laurent Hyafil. 'On the Parallel Evaluation of Multivariate Polynomials'. In: *SIAM Journal on Computing* 8.2 (May 1979), pp. 120–123 (cited on pages 6, 11).

[KS14] Neeraj Kayal and Ramprasad Saptharishi. 'A Selection of Lower Bounds for Arithmetic Circuits'. In: *Perspectives in Computational Complexity*. Ed. by Manindra Agrawal and Vikraman Arvind. Cham: Springer International Publishing, 2014, pp. 77–115 (cited on page 1).

[Koi12] Pascal Koiran. 'Arithmetic circuits: The chasm at depth four gets wider'. In: *Theoretical Computer Science* 448 (2012), pp. 56–65 (cited on page 6).

[KPR22] Balagopal Komarath, Anurag Pandey, and Chengot Sankaramenon Rahul. 'Monotone Arithmetic Complexity of Graph Homomorphism Polynomials'. In: *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*. Vol. 229. Leibniz International Proceedings in Informatics (LIPIcs). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022, 83:1–83:20 (cited on page 15).

[KOS19] Mrinal Kumar, Rafael Oliveira, and Ramprasad Saptharishi. 'Towards Optimal Depth Reductions for Syntactically Multilinear Circuits'. In: *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*. Vol. 132. Leibniz International Proceedings in Informatics (LIPIcs). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019, 78:1–78:15 (cited on page 6).

[LST25] Nutan Limaye, Srikanth Srinivasan, and Sébastien Tavenas. 'Superpolynomial Lower Bounds Against Low-Depth Algebraic Circuits'. In: *Journal of the ACM* 72.4 (2025), pp. 1–35. ECCC: TR21-081 (cited on pages 1, 3, 4).

[Mah13] Meena Mahajan. *Algebraic Complexity Classes*. arXiv:1307.3863 [cs]. Dec. 2013 (cited on page 15).

[MV97] Meena Mahajan and V. Vinay. 'Determinant: Combinatorics, Algorithms, and Complexity'. In: *Chicago Journal of Theoretical Computer Science* 1997.5 (1997) (cited on page 21).

[MP08] Guillaume Malod and Natacha Portier. 'Characterizing Valiant's Algebraic Complexity Classes'. In: *Journal of Complexity* 24.1 (2008), pp. 16–38 (cited on page 19).

[MR20] Laura Mančinska and David E. Roberson. 'Quantum isomorphism is equivalent to equality of homomorphism counts from planar graphs'. In: *IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. 2020, pp. 661–672 (cited on page 13).

[Neu24] Daniel Neuen. 'Homomorphism-Distinguishing Closedness for Graphs of Bounded Tree-Width'. In: *41st International Symposium on Theoretical Aspects of Computer Science (STACS 2024)*. Ed. by Olaf Beyersdorff, Mamadou Moustapha Kanté, Orna Kupferman, and Daniel Lokshtanov. Vol. 289. Leibniz International Proceedings in Informatics (LIPIcs). ISSN: 1868-8969. Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024, 53:1–53:12 (cited on page 11).

[NS26] Daniel Neuen and Tim Seppelt. 'Distinguishing Graphs by Counting Homomorphisms from Sparse Graphs'. In: *41st Annual Symposium on Logic in Computer Science (LICS 2026)*. Ed. by Claudia Faggian and Joost-Pieter Katoen. Vol. 380. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2026, 75:1–75:29 (cited on page 35).

[Pro76] C Procesi. 'The invariant theory of $n \times n$ matrices'. In: *Advances in Mathematics* 19.3 (Mar. 1976), pp. 306–381 (cited on page 13).

[Raz04] Ran Raz. 'Multi-linear formulas for permanent and determinant are of super-polynomial size'. In: *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*. Ed. by László Babai. ACM, 2004, pp. 633–641 (cited on page 1).

[Raz06] Ran Raz. 'Separation of Multilinear Circuit and Formula Size'. In: *Theory Comput.* 2.6 (2006), pp. 121–135 (cited on page 2).

[RY09] Ran Raz and Amir Yehudayoff. 'Lower Bounds and Separations for Constant Depth Multilinear Circuits'. In: *Comput. Complex.* 18.2 (2009), pp. 171–207 (cited on page 1).

[Raz74] Ju P Razmyslov. 'Trace Identities of Full Matrix Algebras Over a Field of Characteristic Zero'. In: *Mathematics of the USSR-Izvestiya* 8.4 (Aug. 1974), pp. 727–760 (cited on page 13).

[Rob26] David E. Roberson. 'Oddomorphisms and homomorphism indistinguishability over graphs of bounded degree'. In: *Journal of Combinatorial Theory, Series B* 181 (Nov. 2026), pp. 1–61 (cited on pages 2, 6, 11, 33, 35, 36).

[Sep24a] Tim Seppelt. 'Homomorphism Indistinguishability'. Dissertation. Aachen: RWTH Aachen University, 2024 (cited on pages 11, 13).

[Sep24b] Tim Seppelt. 'Logical equivalences, homomorphism indistinguishability, and forbidden minors'. In: *Information and Computation* 301 (Dec. 2024), p. 105224 (cited on pages 11, 38).

[Sip83] Michael Sipser. 'Borel Sets and Circuit Complexity'. In: *Proceedings of the 15th Annual ACM Symposium on Theory of Computing, 25-27 April, 1983, Boston, Massachusetts, USA*. Ed. by David S. Johnson, Ronald Fagin, Michael L. Fredman, David Harel, Richard M. Karp, Nancy A. Lynch, Christos H. Papadimitriou, Ronald L. Rivest, Walter L. Ruzzo, and Joel I. Seiferas. ACM, 1983, pp. 61–69 (cited on page 2).

[Tav15] Sébastien Tavenas. 'Improved bounds for reduction to depth 4 and depth 3'. In: *Inf. Comput.* 240 (2015), pp. 2–11 (cited on page 6).

[Val79] L. G. Valiant. 'Completeness classes in algebra'. In: *Proceedings of the eleventh annual ACM symposium on Theory of computing - STOC '79*. Atlanta, Georgia, United States: ACM Press, 1979, pp. 249–261 (cited on page 1).

[Val+83] L. G. Valiant, S. Skyum, S. Berkowitz, and C. Rackoff. 'Fast Parallel Computation of Polynomials Using Few Processors'. In: *SIAM Journal on Computing* 12.4 (Nov. 1983), pp. 641–644 (cited on pages 3, 6, 11, 12).

[Val82] Leslie G. Valiant. 'Reducibility by Algebraic Projections'. In: *Logic and Algorithmic: An International Symposium Held in Honor of Ernst Specker*. Vol. 30. Monographies de L'Enseignement Mathématique. 1982, pp. 365–380 (cited on page 19).

A. Reduced circuits and supports

We prove Lemmas 2.3 and 2.4 and Theorem 2.6. For a gate g, write $M(g)$ for the multiset of its children and $m(g,h)$ for the multiplicity of the wire from h to g. The *height* of a gate is the maximal number of wires on a path from it to an input gate.

Lemma 2.3. *Every reduced matrix-symmetric circuit is rigid.*

Lean formalisation on LaxLean-Archive

Proof. Let C be a Γ_n-symmetric reduced circuit. Two automorphisms extending the same element of Γ_n differ by an automorphism ρ of C that fixes all input gates, since input labels are distinct. By induction on the height, ρ fixes every gate g: the gate $\rho(g)$ has the same label as g and the multiset of children $\rho(M(g)) = M(g)$, so $\rho(g) = g$ since C is reduced. Thus, ρ is the identity automorphism, and hence, there do not exist two distinct automorphisms of C that extend the same element of Γ_n. □

Lemma 2.4 (Rigidification)**.** *For every matrix-symmetric circuit C, there is a reduced matrix-symmetric circuit C' computing the same polynomial whose size, depth, and orbit size are at most those of C.*

Lean formalisation on LaxLean-Archive

Proof. We define an equivalence relation $\equiv$ on the gates of C by induction on the height. Input gates are equivalent if they have the same label. Internal gates g and g' are equivalent if they have the same height and label and $m(g,c) := \sum_{h \in c} m(g,h) = m(g',c)$ for every class c of gates of smaller height. Let C' have as gates the classes from which the class $[\text{out}]$ of the output gate can be reached, where $[g]$ carries the label of g and has a wire of multiplicity $m(g,c)$ from every class c. By induction on the height, all gates in a class compute the same polynomial that is computed by the class in C'. The circuit C' is reduced since classes with the same label and the same children have the same height in C and are collapsed into a single gate in C'. Moreover, C' has at most as many gates and wires as C, and the same depth.

Every automorphism σ of C preserves heights, labels, and multiplicities, and thus maps classes to classes. It induces an automorphism σ' of C', which acts on the input labels as σ does. Hence, C' is matrix-symmetric. Since C' is rigid by Lemma 2.3, σ' is the unique extension of the element of Γ_n extended by σ. Hence, the map $g \mapsto [g]$ maps every $\mathbf{Aut}_\Gamma(C)$-orbit onto an orbit of C', and every orbit of C' arises this way. So the orbit size does not increase. □

The only group-theoretic machinery that we are going to need is a classical result on subgroups of small index. Write $\mathbf{Alt}_{(T)} \leq \mathbf{Sym}_{(T)} \leq \mathbf{Sym}_n$ for the pointwise stabilisers of $T \subseteq [n]$ in the alternating and the symmetric group.

Theorem A.1 (cf. [DM96, Theorem 5.2B])**.** *Let $n \geq 9$ and $1 \leq r \leq n/2$. If $H \leq \mathbf{Sym}_n$ has index less than $\min\{\frac{1}{2}\binom{n}{\lfloor n/2 \rfloor}, \binom{n}{r}\}$, then $\mathbf{Alt}_{(T)} \leq H$ for some $T \subseteq [n]$ with $|T| < r$.*

Lean formalisation on LaxLean-Archive

Proof. By [DM96, Theorem 5.2B], since $n \geq 5$ and the index of H is less than $\binom{n}{r}$, one of the following holds: (i) $\mathbf{Alt}_{(T)} \leq H$ for some $T \subseteq [n]$ with $|T| < r$; (ii) $n = 2m$ is even and H has index $\frac{1}{2}\binom{n}{m}$; or (iii) one of six exceptional cases with $n \leq 8$ holds. Case (ii) is excluded by the bound on the index, and case (iii) by $n \geq 9$. □

Recall the definition of $\kappa_f(n)$ from Section 2.3:

$$\kappa_f(n) := \min\left\{ r \in \mathbb{N} : \binom{n}{r} > f(n) \right\}. \tag{5}$$

Theorem 2.6 (Support theorem)**.** *Let $f : \mathbb{N} \to \mathbb{N}$. Let $n \geq 20$ such that $3\kappa_f(n) \leq n + 2$. Then every gate of a reduced matrix-symmetric circuit of orbit size $\leq f(n)$ has a minimal support of size $\leq \kappa_f(n) - 1$.*

Lean formalisation on LaxLean-Archive

Proof. Write $\kappa := \kappa_f(n)$, and let C be a reduced matrix-symmetric circuit of orbit size at most $f(n)$. Note that $f(n) < \frac{1}{2}\binom{n}{\lfloor n/2 \rfloor}$ and $\kappa < n/2$ as $n \geq 20$ and $3\kappa \leq n+2$. If C contains no variable, then all gates are fixed by Γ_n and have support $\emptyset$. Otherwise, an input gate $x_{i,j}$ has orbit size $n^2 \leq f(n)$, so $\kappa \geq 2$.

By induction on the height, we construct supports S_g with $|(S_g)_L|, |(S_g)_R| \leq \kappa - 1$. For input gates $x_{i,j}$, let $S_g := \{i\} \uplus \{j\}$, and for constants, let $S_g := \emptyset$. Let g be an internal gate and $H := \{\pi \in \mathbf{Sym}_n \mid (\pi, \mathrm{id}) \in \mathbf{Stab}(g)\}$. Then $[\mathbf{Sym}_n : H] \leq [\Gamma_n : \mathbf{Stab}(g)] \leq f(n) < \min\{\frac{1}{2}\binom{n}{\lfloor n/2 \rfloor}, \binom{n}{\kappa}\}$. By Theorem A.1 for $r := \kappa$, there is $T_L \subseteq [n]$ with $|T_L| \leq \kappa - 1$ and $A := \mathbf{Alt}_{(T_L)} \times \{\mathrm{id}\} \leq \mathbf{Stab}(g)$.

We show that $\mathbf{Sym}_{(T_L)} \times \{\mathrm{id}\} \leq \mathbf{Stab}(g)$ using that C is reduced. Let h be a child of g. Since $|[n] \setminus (T_L \cup (S_h)_L)| \geq n - 2\kappa + 2 \geq 2$, there is a transposition τ of two points outside $T_L \cup (S_h)_L$. Then $(\tau, \mathrm{id}) \in \mathbf{Stab}^\bullet(S_h) \leq \mathbf{Stab}(h)$. As $\mathbf{Sym}_{(T_L)} \times \{\mathrm{id}\} = A \cup A(\tau, \mathrm{id})$, for every $\sigma \in \mathbf{Sym}_{(T_L)} \times \{\mathrm{id}\}$ there is $\alpha \in A$ such that $\sigma h = \alpha h$. Since α fixes g, it follows that $m(g, \sigma h) = m(g, \alpha h) = m(g, h)$. Hence, $\sigma(M(g)) = M(g)$, so σg and g have the same label and the same children. Since C is reduced, $\sigma g = g$.

Analogously, there is $T_R \subseteq [n]$ with $|T_R| \leq \kappa - 1$ and $\{\mathrm{id}\} \times \mathbf{Sym}_{(T_R)} \leq \mathbf{Stab}(g)$. Hence, $S_g := T_L \uplus T_R$ supports g. Since $\kappa - 1 < n/2$, the minimal support $\mathrm{sup}(g) \subseteq S_g$ exists.

Write $k_L := |\mathrm{sup}_L(g)|$ and $k_R := |\mathrm{sup}_R(g)|$, so that $k_L, k_R \leq \kappa - 1$. By equivariance of minimal supports, the orbit of g maps onto the Γ_n-orbit of $(\mathrm{sup}_L(g), \mathrm{sup}_R(g))$. Hence,

$$\binom{n}{k_L + k_R} \leq \binom{n}{k_L}\binom{n - k_L}{k_R} \leq \binom{n}{k_L}\binom{n}{k_R} \leq f(n) < \binom{n}{\kappa}.$$

Suppose that $k_L + k_R \geq \kappa$. Since $k_L + k_R \leq 2\kappa - 2 \leq n - \kappa$ and $\binom{n}{r} \geq \binom{n}{\kappa}$ for all $\kappa \leq r \leq n - \kappa$, this is a contradiction. Hence, $k_L + k_R \leq \kappa - 1$. □

B. Material omitted in Section 3

B.1. Circuits for patterns of bounded forestwidth

We construct circuits for single homomorphism polynomials, following [Bha+25, Lemma 19] and [DPS26a]. For a linear combination $\sum_i \alpha_i \hom_{F_i,n}$, one applies the construction below to every F_i, multiplies the output gates by the constants α_i, adds an output sum gate, and reduces via Lemma 2.4. The proof of Theorem B.2 applies verbatim, except that the depth grows by one. This yields the implication from 2 to 1 in Theorem 3.5.

Lemma B.1 (Normalised decompositions)**.** *Let $\Delta, k \geq 1$ and let F be a multigraph without isolated vertices with $\mathrm{fw}_\Delta(F) \leq k - 1$. Then there are a tree decomposition (T, β) of F with root r_0 and a map $\mathrm{rep}\colon E(F) \to V(T)$ with $e \subseteq \beta(\mathrm{rep}(e))$ for every edge e such that*

1. *$\beta(r_0) = \emptyset$, T has depth at most Δ, and $|\beta(t)| \leq k$ for every node t,*
2. *for every node t and every $v \in \beta(t)$, there is an edge e incident with v such that $\mathrm{rep}(e)$ is t or a descendant of t,*
3. *$\beta(c) \not\subseteq \beta(t)$ for every node $c \neq r_0$ with parent t, and*
4. *$|V(T)| \leq |V(F)| + 1$.*

We say that an edge e is *represented* at the node $\mathrm{rep}(e)$. So the second item asserts that every vertex in a bag is incident with an edge that is represented at this bag or below it.

Proof. Take a tree decomposition of F of depth $\leq \Delta$ with bags of size $\leq k$ whose root r_0 has $\beta(r_0) = \emptyset$. Choose $\mathrm{rep}(e)$ as any node whose bag contains e, for every edge e, counted with multiplicity. The first item holds. For the second item, delete v from $\beta(t)$ whenever it fails. For every vertex v, the nodes that keep v form an upwards closed subset of the subtree of nodes whose bag contains v, so they still induce a subtree. This subtree contains $\mathrm{rep}(e)$ for every edge e incident with v, and it is non-empty as v is not isolated. For the third item, as long as $\beta(c) \subseteq \beta(t)$ for some node $c \neq r_0$ with parent t, attach the children of c to t, represent the edges represented at c at t instead, and delete c. This preserves the first two items. For the fourth item, map every node $c \neq r_0$ to some vertex in $\beta(c) \setminus \beta(t)$, where t is the parent of c. This map is injective since c is the topmost node whose bag contains this vertex. □

Lean formalisation on LaxLean-Archive

Theorem B.2. *Let $\Delta \geq 1$, $k \geq 2$, and $n \geq 2k+1$. Let F be a bipartite multigraph with $\mathrm{fw}_\Delta(F) \leq k-1$. Then $\hom_{F,n}$ is computed by a reduced matrix-symmetric circuit of support size $\leq k$, support depth $\leq \Delta$, orbit size $\leq n^k$, depth $\leq 2\Delta+1$, and size $\leq 5n^k(|V(F)|+|E(F)|)+1$.*

Proof. Decompositions. If F has j isolated vertices and F' is obtained by deleting them, then $\hom_{F,n} = n^j \hom_{F',n}$ by Observation 2.7 and $\mathrm{fw}_\Delta(F') \leq \mathrm{fw}_\Delta(F)$. Hence, we may suppose that F has no isolated vertices: at the end, we add the constant n^j as a further child of the output gate, which is a product gate, so that the depth does not grow. Let $A \uplus B$ be the bipartition of F. Let (T, β) with root r_0 and rep be as in Lemma B.1. Write $\mathrm{level}(t) \leq \Delta$ for the distance of a node t from r_0, $t' \succeq t$ if t' is a descendant of t or $t' = t$, and $V_t := \bigcup_{t' \succeq t} \beta(t')$.

The circuit. A map $\gamma \colon X \to [n] \uplus [n]$ on $X \subseteq V(F)$ is *admissible* if it maps $A \cap X$ to the left and $B \cap X$ to the right copy of $[n]$. For every node t and admissible γ on $\beta(t)$, and for every node $c \neq r_0$ with parent t and admissible δ on $\beta(t) \cap \beta(c)$, define gates

$$R_{t,\gamma} := \prod_{e=ab \in \mathrm{rep}^{-1}(t)} x_{\gamma(a),\gamma(b)} \cdot \prod_{c \text{ child of } t} S_{c,\gamma|_{\beta(t)\cap\beta(c)}}, \qquad S_{c,\delta} := \sum_{\gamma' \supseteq \delta \text{ admissible on } \beta(c)} R_{c,\gamma'},$$

where empty products are the constant 1. The output gate is $R_{r_0,\emptyset}$. Write $P_{t,\gamma}$ and $Q_{c,\delta}$ for the polynomials computed by $R_{t,\gamma}$ and $S_{c,\delta}$. By induction on T,

$$P_{t,\gamma} = \sum_{h \supseteq \gamma \text{ admissible on } V_t} \prod_{e \in E(F),\, \mathrm{rep}(e) \succeq t} x_{h(e)} \quad \text{and} \quad Q_{c,\delta} = \sum_{\gamma' \supseteq \delta} P_{c,\gamma'},$$

where $x_{h(e)} := x_{h(a),h(b)}$ for $e = ab$. For the inductive step, note that $V_c \cap V_{c'} \subseteq \beta(t)$ and $V_c \cap \beta(t) = \beta(c) \cap \beta(t)$ for distinct children c, c' of t, so the extensions of γ to the sets V_c can be chosen independently. Since $V_{r_0} = V(F)$, the output gate computes $P_{r_0,\emptyset} = \hom_{F,n}$.

Depth and size. Write $d(t)$ for the depth of the subtree of T rooted at t. By induction, $R_{t,\gamma}$ has height $\leq 2d(t)+1$ and $S_{c,\delta}$ has height $\leq 2d(c)+2$. Since $d(r_0) \leq \Delta$, the circuit has depth $\leq 2\Delta+1$. For every node t, there are at most $n^{|\beta(t)|} \leq n^k$ gates $R_{t,\gamma}$, each with $|\mathrm{rep}^{-1}(t)|$ wires from input gates and one wire from each child of t. For every node $c \neq r_0$, there are at most n^k gates $S_{c,\delta}$, which receive at most $n^{|\beta(c)|} \leq n^k$ wires in total. Together with at most $n^2 \leq n^k$ input gates and the constants 1 and n^j with one wire, the size is at most

$$\begin{aligned} &n^k\big(|V(T)| + (|E(F)| + |V(T)| - 1) + 2(|V(T)| - 1) + 1\big) + 3 \\ &\leq n^k\big(4|V(F)| + |E(F)| + 3\big) \leq 5n^k\big(|V(F)| + |E(F)|\big), \end{aligned}$$

using $|V(T)| \leq |V(F)| + 1$ by Lemma B.1, $n^k \geq 3$, and $|V(F)| \geq 2$ and $|E(F)| \geq 1$.

Symmetry. The group Γ_n acts on the gates via $R_{t,\gamma} \mapsto R_{t,\pi\circ\gamma}$, $S_{c,\delta} \mapsto S_{c,\pi\circ\delta}$, and naturally on input gates, and it fixes the remaining gates. This preserves labels and wires, so the circuit is

matrix-symmetric. The orbit of $R_{t,\gamma}$ has at most $n^{|\beta(t)|} \leq n^k$ elements, as does the orbit of $S_{c,\delta}$, and input gates have orbits of size $\leq n^2 \leq n^k$. Moreover, $\pi P_{t,\gamma} = P_{t,\pi\circ\gamma}$ and $\pi Q_{c,\delta} = Q_{c,\pi\circ\delta}$ for all $\pi \in \Gamma_n$. Let C be the reduced circuit constructed in Lemma 2.4, and write $[g]$ for the gate of C corresponding to a gate g. The automorphism of C extending $\pi \in \Gamma_n$ maps $[g]$ to $[\pi g]$ by the proof of Lemma 2.4 and since C is rigid. Hence, C has orbit size $\leq n^k$. By Lemma 2.4, the depth and size bounds carry over to C.

Supports. We show that $\sup([R_{t,\gamma}]) = \gamma(\beta(t))$ and $\sup([S_{c,\delta}]) = \delta(\beta(t) \cap \beta(c))$. Let $g := [R_{t,\gamma}]$ and $U := \gamma(\beta(t))$. Every $\pi \in \mathbf{Stab}^\bullet(U)$ satisfies $\pi \circ \gamma = \gamma$ and thus fixes g. Since $|U| \leq k < n/2$, the minimal support $\sup(g) \subseteq U$ exists. Suppose that there is $p_0 \in U \setminus \sup(g)$, say on the left side, and let $v \in \beta(t)$ with $\gamma(v) = p_0$. Since $|U| < n$, there is a left point $p_1 \notin U$. The transposition τ of p_0 and p_1 fixes $\sup(g)$ pointwise and thus fixes g, so $[R_{t,\tau\circ\gamma}] = g$ and $\tau P_{t,\gamma} = P_{t,\tau\circ\gamma} = P_{t,\gamma}$. By Lemma B.1, every monomial of $P_{t,\gamma}$ contains a variable in row p_0, so every monomial of $\tau P_{t,\gamma}$ contains a variable in row p_1. However, mapping all left vertices in $V_t \setminus \beta(t)$ to p_0 yields a monomial of $P_{t,\gamma}$ with positive coefficient that contains no variable in row p_1, a contradiction. The same argument applies to $[S_{c,\delta}]$, using the second item of Lemma B.1 at c. The constants and the input gates $x_{i,j}$ have minimal supports $\emptyset$ and $\{i\} \uplus \{j\}$, respectively. In particular, C has support size $\leq k$.

Support depth. The children of $[R_{t,\gamma}]$ are input gates $x_{\gamma(a),\gamma(b)}$ for $a, b \in \beta(t)$, constants, and gates $[S_{c,\gamma|_{\beta(t)\cap\beta(c)}}]$. Their supports are contained in $\gamma(\beta(t))$, so no wire into $[R_{t,\gamma}]$ is support-changing. Hence, only wires into gates $[S_{c,\delta}]$ are support-changing. By induction on the height of gates in C, the support depth of $[R_{t,\gamma}]$ is at most $\Delta - \mathrm{level}(t)$ and that of $[S_{c,\delta}]$ is at most $\Delta - \mathrm{level}(c) + 1$ for every representative. Here, we use that all representatives of a gate of C have the same children in C. Since $\mathrm{level}(r_0) = 0$, the circuit C has support depth $\leq \Delta$. □

B.2. Homomorphism polynomials from symmetric circuits

We prove the implication from 1 to 2 in Theorem 3.5. We follow the approach of [DPS26a, Section 4] via homomorphism polynomials of labelled graphs. Since the support size k may grow with n and we keep track of the depth of decompositions, we deviate from [DPS26a] as follows.

- We use rooted tree decompositions whose root bag consists *exactly* of the labelled non-isolated vertices, and isolated vertices occur in no bag. In [DPS26a], the labelled vertices are only required to lie in a common bag. Our convention makes gluing free in depth and adding fresh labels free in width. As a consequence, bags have size at most k rather than $2k$ as in [DPS26a].
- In [DPS26a], labels are summed or multiplied out one at a time. This would cost one level of depth per label, i.e. up to k levels per gate. We instead eliminate all labels of an orbit of children at once by a single restricted sum over injective assignments, at the cost of one level (Lemma B.6).
- To that end, products over orbits are expressed via power sums (Lemma B.5), replacing the orbit-counting argument for products and the induction for restricted products of [DPS26a].
- We work with reduced circuits and Theorem 2.6 instead of rigid circuits of polynomial orbit size, and the depth of the decomposition grows only at support-changing wires (Lemma B.7).

B.2.1. Labelled graphs

Let $\ell, r \in \mathbb{N}$. An (ℓ, r)-*labelled graph* $\boldsymbol{F} = (F, \boldsymbol{a}, \boldsymbol{b})$ consists of a bipartite multigraph F with bipartition $A \uplus B$ and tuples $\boldsymbol{a} \in A^\ell$ and $\boldsymbol{b} \in B^r$ [DPS26a, Definition 4.1]. For $\boldsymbol{v} \in [n]^\ell$ and $\boldsymbol{w} \in [n]^r$,

let

$$\boldsymbol{F}_n(\boldsymbol{v},\boldsymbol{w}) := \sum_{\substack{h\colon A\uplus B\to[n]\uplus[n]\\ h(\boldsymbol{a})=\boldsymbol{v},\ h(\boldsymbol{b})=\boldsymbol{w}}} \prod_{ab\in E(F)} x_{h(a),h(b)},$$

where h maps A to the left and B to the right copy of $[n]$. For $\ell = r = 0$, this is $\hom_{F,n}$. Write F^+ for the graph obtained from F by deleting all isolated vertices and $L^+(\boldsymbol{F})$ for the set of labelled vertices of F^+.

For $k \geq 1$ and $\Delta \geq 0$, let $\mathcal{T}^{\Delta,k}(\ell,r)$ denote the class of all (ℓ,r)-labelled graphs $\boldsymbol{F}$ such that F^+ admits a rooted tree decomposition of depth $\leq \Delta$ with bags of size $\leq k$ whose root bag is $L^+(\boldsymbol{F})$. We call such a decomposition *root-exact*. Let $\mathfrak{T}_n^{\Delta,k}(\ell,r)$ denote the $\mathbb{K}$-linear span of the maps $\boldsymbol{F}_n\colon [n]^\ell \times [n]^r \to \mathbb{K}[\mathcal{X}_n]$ for $\boldsymbol{F} \in \mathcal{T}^{\Delta,k}(\ell,r)$. Clearly, $\mathfrak{T}_n^{\Delta,k}(\ell,r) \subseteq \mathfrak{T}_n^{\Delta+1,k}(\ell,r)$. For a set I of label positions, let $\mathcal{T}_I^{\Delta,k}(\ell,r)$ denote the class of all $\boldsymbol{F} \in \mathcal{T}^{\Delta,k}(\ell,r)$ such that every vertex in $L^+(\boldsymbol{F})$ carries a label at a position in I, and let $\mathfrak{T}_{n,I}^{\Delta,k}(\ell,r)$ denote the span of the corresponding maps. If $\ell + r \leq k$, then taking I to be the set of all positions shows that these restrictions are vacuous.

The *gluing* $\boldsymbol{F} \odot \boldsymbol{F}'$ of $\boldsymbol{F}, \boldsymbol{F}' \in \mathcal{T}^{\Delta,k}(\ell,r)$ is obtained from the disjoint union of F and F' by identifying a_i with a_i' and b_j with b_j' for all i and j. Then $(\boldsymbol{F} \odot \boldsymbol{F}')_n$ is the pointwise product of $\boldsymbol{F}_n$ and $\boldsymbol{F}'_n$ [DPS26a]. Write $\boldsymbol{J}$ for the edgeless graph whose labels lie on distinct vertices, and $\boldsymbol{D}^{i,j}$ for the edgeless graph whose labels lie on distinct vertices except that the labels at the positions $i \neq j$ on the same side share a vertex. Then $\boldsymbol{J}_n \equiv 1$, and $\boldsymbol{D}_n^{i,j}(\boldsymbol{v},\boldsymbol{w})$ is 1 if the entries at positions i and j agree and 0 otherwise.

Lemma B.3. *Let $k \geq 1$, $\Delta \geq 0$, $\ell, r \in \mathbb{N}$, and $\boldsymbol{F}, \boldsymbol{F}' \in \mathcal{T}^{\Delta,k}(\ell,r)$.*

1. *Every edgeless (ℓ,r)-labelled graph is in $\mathcal{T}_I^{0,k}(\ell,r)$ for every set I of positions.*
2. *If $\boldsymbol{F}, \boldsymbol{F}' \in \mathcal{T}_I^{\Delta,k}(\ell,r)$ for a set I of at most k positions, then $\boldsymbol{F} \odot \boldsymbol{F}' \in \mathcal{T}_I^{\Delta,k}(\ell,r)$.*
3. *Permuting the labels of $\boldsymbol{F}$ on each side, and adding labels on new isolated vertices, yields graphs in $\mathcal{T}^{\Delta,k}$. If $\boldsymbol{F} \in \mathcal{T}_I^{\Delta,k}$, then the resulting graphs are in $\mathcal{T}_{I'}^{\Delta,k}$, where I' denotes the new positions of the labels at the positions in I.*
4. *Deleting labels from $\boldsymbol{F}$ yields a graph in $\mathcal{T}^{\Delta+1,k}$.*

In particular, for every set I of at most k positions, $\mathfrak{T}_{n,I}^{\Delta,k}(\ell,r)$ is a $\mathbb{K}$-algebra under pointwise multiplication with unit $\boldsymbol{J}_n$. This applies to $\mathfrak{T}_n^{\Delta,k}(\ell,r)$ if $\ell + r \leq k$.

Proof. 1. Use a single node with empty bag. 2. Identify the roots of root-exact decompositions of $\boldsymbol{F}$ and $\boldsymbol{F}'$ and apply the quotient map to all bags. Since gluing adds no edges, a vertex of $F \odot F'$ is non-isolated if and only if one of its preimages is, so the new root bag is $L^+(\boldsymbol{F} \odot \boldsymbol{F}')$. Every vertex in this bag carries a label at a position in I, and distinct vertices carry distinct labels, so the root bag has size at most $|I| \leq k$. The other bags do not grow. Unlabelled vertices occur in one of the two subtrees only, and the occurrences of a labelled non-isolated vertex are connected via the common root. 3. The decomposition remains valid. 4. Add a new root above the old one whose bag consists of the remaining labelled non-isolated vertices, a subset of the old root bag. □

Lemma B.4. *Let $k, \Delta \geq 1$. An unlabelled graph F is in $\mathcal{T}^{\Delta,k}(0,0)$ if, and only if, $\mathrm{fw}_\Delta(F) \leq k-1$.*

Proof. A root-exact decomposition of F^+ has an empty root bag. Attaching a child with bag $\{v\}$ to the root for every isolated vertex v yields a decomposition witnessing $\mathrm{fw}_\Delta(F) \leq k-1$, since $\Delta \geq 1$. Conversely, deleting the isolated vertices from all bags of such a decomposition yields a root-exact decomposition of F^+. □

B.2.2. Eliminating many labels at once

Let $\ell, \ell', r, r' \in \mathbb{N}$. For $\boldsymbol{v} \in [n]^\ell$ and $\boldsymbol{w} \in [n]^r$, let $Y(\boldsymbol{v}, \boldsymbol{w})$ denote the set of pairs $(\boldsymbol{y}, \boldsymbol{z}) \in [n]^{\ell'} \times [n]^{r'}$ of injective tuples such that $\boldsymbol{y}$ avoids the entries of $\boldsymbol{v}$ and $\boldsymbol{z}$ avoids the entries of $\boldsymbol{w}$. For a map φ on $[n]^{\ell+\ell'} \times [n]^{r+r'}$, define the *restricted sum*

$$(\Sigma^{\neq}\varphi)(\boldsymbol{v}, \boldsymbol{w}) \coloneqq \sum_{(\boldsymbol{y},\boldsymbol{z}) \in Y(\boldsymbol{v},\boldsymbol{w})} \varphi(\boldsymbol{v}\boldsymbol{y}, \boldsymbol{w}\boldsymbol{z}).$$

The following expression of a product in terms of power sums is a consequence of Newton's identities.

Lemma B.5 (Products from power sums)**.** *Let R be a commutative $\mathbb{Q}$-algebra, $N \geq 1$, $y_1, \ldots, y_N \in R$, and $p_m \coloneqq \sum_{i=1}^N y_i^m$. Then*

$$\prod_{i=1}^N y_i = \frac{1}{N!} \sum_\pi \mu(\pi) \prod_{P \in \pi} p_{|P|} \quad \text{where} \quad \mu(\pi) \coloneqq \prod_{P \in \pi} (-1)^{|P|-1}(|P|-1)!$$

and π ranges over all partitions of $[N]$.

Proof. For a partition π, let $g(\pi) \coloneqq \sum_{h\colon \pi \to [N]} \prod_i y_{h(\pi(i))} = \prod_{P \in \pi} p_{|P|}$, where $\pi(i)$ is the part containing i, and let $f(\pi)$ be the same sum restricted to injective h. Then $g(\pi) = \sum_{\sigma \geq \pi} f(\sigma)$, where σ ranges over coarsenings of π. By Möbius inversion on the partition lattice, $N! \prod_i y_i = f(\hat{0}) = \sum_\sigma \mu(\sigma) g(\sigma)$. □

The following lemma eliminates the last ℓ' left and r' right labels at once.

Lemma B.6. *Let $k, c \geq 1$ and $\Delta \geq 0$ with $\ell + r \leq k$, let I be a set of at most k positions, and let $\varphi \in \mathfrak{T}_{n,I}^{\Delta,k}(\ell + \ell', r + r')$. There are $\psi_\Sigma, \psi_\Pi \in \mathfrak{T}_n^{\Delta+1,k}(\ell, r)$ with the following property. Let $\boldsymbol{v} \in [n]^\ell$ and $\boldsymbol{w} \in [n]^r$ be injective, and let $\mathcal{P}$ be a partition of $Y(\boldsymbol{v}, \boldsymbol{w})$ into classes each of size exactly c such that $(\boldsymbol{y}, \boldsymbol{z}) \mapsto \varphi(\boldsymbol{v}\boldsymbol{y}, \boldsymbol{w}\boldsymbol{z})$ is constant on every class $P \in \mathcal{P}$. For $P \in \mathcal{P}$, write $\varphi_P \coloneqq \varphi(\boldsymbol{v}\boldsymbol{y}, \boldsymbol{w}\boldsymbol{z})$ for any $(\boldsymbol{y}, \boldsymbol{z}) \in P$. Then*

$$\psi_\Sigma(\boldsymbol{v}, \boldsymbol{w}) = \sum_{P \in \mathcal{P}} \varphi_P \quad \text{and} \quad \psi_\Pi(\boldsymbol{v}, \boldsymbol{w}) = \prod_{P \in \mathcal{P}} \varphi_P.$$

Proof. Let δ be the pointwise product of the maps $\boldsymbol{J}_n - \boldsymbol{D}_n^{i,j}$ over all pairs of distinct positions i, j on the same side of which at least one is among the last ℓ' left or r' right positions. By Lemma B.3, $\delta \in \mathfrak{T}_{n,I}^{0,k}(\ell + \ell', r + r')$, and $\delta(\boldsymbol{v}\boldsymbol{y}, \boldsymbol{w}\boldsymbol{z})$ is the indicator of $(\boldsymbol{y}, \boldsymbol{z}) \in Y(\boldsymbol{v}, \boldsymbol{w})$. For every $\varphi' \in \mathfrak{T}_{n,I}^{\Delta,k}(\ell + \ell', r + r')$, the map $\Sigma^{\neq}\varphi'$ is obtained from $\delta \odot \varphi'$ by deleting the last ℓ' left and r' right labels, so $\Sigma^{\neq}\varphi' \in \mathfrak{T}_n^{\Delta+1,k}(\ell, r)$ by Lemma B.3.

Let $\psi_\Sigma \coloneqq \frac{1}{c}\Sigma^{\neq}\varphi$. Under the assumptions of the lemma, $\psi_\Sigma(\boldsymbol{v}, \boldsymbol{w}) = \frac{1}{c}\sum_{(\boldsymbol{y},\boldsymbol{z}) \in Y(\boldsymbol{v},\boldsymbol{w})} \varphi(\boldsymbol{v}\boldsymbol{y}, \boldsymbol{w}\boldsymbol{z}) = \sum_{P \in \mathcal{P}} \varphi_P$. For injective $\boldsymbol{v}$ and $\boldsymbol{w}$, the size of $Y(\boldsymbol{v}, \boldsymbol{w})$ depends only on n, ℓ, ℓ', r, and r'. If c does not divide it, then no partition $\mathcal{P}$ as above exists and we let $\psi_\Pi \coloneqq 0$. Otherwise, let $N \coloneqq |Y(\boldsymbol{v}, \boldsymbol{w})|/c$, let $Q_m \coloneqq \frac{1}{c}\Sigma^{\neq}(\varphi^{\odot m})$ for $m \geq 1$, and let $\psi_\Pi \coloneqq \frac{1}{N!}\sum_\pi \mu(\pi) \bigodot_{B \in \pi} Q_{|B|}$, where π ranges over all partitions of $[N]$. Since $\varphi^{\odot m} \in \mathfrak{T}_{n,I}^{\Delta,k}(\ell + \ell', r + r')$ and $\ell + r \leq k$, we have $\psi_\Pi \in \mathfrak{T}_n^{\Delta+1,k}(\ell, r)$ by Lemma B.3. Under the assumptions of the lemma, $|\mathcal{P}| = N$ and $Q_m(\boldsymbol{v}, \boldsymbol{w}) = \frac{1}{c}\sum_{(\boldsymbol{y},\boldsymbol{z}) \in Y(\boldsymbol{v},\boldsymbol{w})} \varphi(\boldsymbol{v}\boldsymbol{y}, \boldsymbol{w}\boldsymbol{z})^m = \sum_{P \in \mathcal{P}} \varphi_P^m$ for all $m \geq 1$. That is, $Q_m(\boldsymbol{v}, \boldsymbol{w})$ is the m-th power sum p_m of the N elements φ_P for $P \in \mathcal{P}$ of $\mathbb{K}[\mathcal{X}_n]$. Since pointwise products are evaluated pointwise, Lemma B.5 applied to these elements yields

$$\psi_\Pi(\boldsymbol{v}, \boldsymbol{w}) = \frac{1}{N!}\sum_\pi \mu(\pi) \prod_{B \in \pi} Q_{|B|}(\boldsymbol{v}, \boldsymbol{w}) = \frac{1}{N!}\sum_\pi \mu(\pi) \prod_{B \in \pi} p_{|B|} = \prod_{P \in \mathcal{P}} \varphi_P.$$ □

We stress that the costs in Lemma B.6 depend neither on ℓ' and r' nor on k.

B.2.3. Induction on the circuit

Let C be a reduced matrix-symmetric circuit with support size $k < n/2$, so that every gate g has a minimal support. Fix an ordering $\vec{\sup}(g) = (\vec{\sup}_L(g), \vec{\sup}_R(g))$ of $\sup(g)$ for every gate g, and write $\mathrm{sc}(g)$ for the maximal number of support-changing wires on a path from g to an input gate. For $\gamma \in \Gamma_n$, we have $p_{\gamma g} = \gamma p_g$ and $\sup(\gamma g) = \gamma \sup(g)$.

Lemma B.7. *For every gate g of C, there is $\varphi_g \in \mathfrak{T}_n^{\mathrm{sc}(g),k}(|\sup_L(g)|, |\sup_R(g)|)$ such that $p_g = \varphi_g(\vec{\sup}(g))$.*

Proof. Input gates. The gate $x_{i,j}$ has minimal support $\{i\} \uplus \{j\}$, and $k \geq 2$. Let φ be the map of the $(1,1)$-labelled single edge, whose root bag is the edge. A constant c has empty support, and we let $\varphi := c\boldsymbol{J}_n$.

Internal gates. Let g be internal, $\boldsymbol{u} := \vec{\sup}(g)$, and $G := \mathbf{Stab}^\bullet(\sup(g)) \leq \mathbf{Stab}(g)$. Then G permutes the children of g and preserves wire multiplicities. Let O be a G-orbit of children of g, let m_O be the common multiplicity of the wires from O to g, and let $h \in O$. Let $N := \sup(h) \setminus \sup(g)$, ordered as in $\vec{\sup}(h)$. By Lemma B.3, there is $\psi_h \in \mathfrak{T}_{n,I}^{\mathrm{sc}(h),k}$, where I is the set of the $|\sup(h)| \leq k$ positions of the elements of $\sup(h)$, with label positions for $\boldsymbol{u}$ followed by positions for N such that $\psi_h(\boldsymbol{u}, \boldsymbol{y}) = \varphi_h(\vec{\sup}(h)[N \mapsto \boldsymbol{y}])$ for all tuples $\boldsymbol{y}$, obtained by permuting the labels of φ_h and adding labels for $\sup(g) \setminus \sup(h)$ on isolated vertices. Here and below, we suppress the split of the label positions into left and right ones.

Let $Y_{\boldsymbol{u}} := Y(\boldsymbol{u})$ be the set of injective tuples $\boldsymbol{y}$ for the positions of N that avoid $\sup(g)$ on each side. Since G acts as the full symmetric group on the complement of $\sup(g)$ on each side, it acts transitively on $Y_{\boldsymbol{u}}$. For $\boldsymbol{y} \in Y_{\boldsymbol{u}}$, choose $\gamma \in G$ mapping N to $\boldsymbol{y}$ and let $h_{\boldsymbol{y}} := \gamma h$. This does not depend on the choice of γ, since two choices differ by an element of $\mathbf{Stab}^\bullet(\sup(g) \cup N) \leq \mathbf{Stab}^\bullet(\sup(h)) \leq \mathbf{Stab}(h)$. The map $\boldsymbol{y} \mapsto h_{\boldsymbol{y}}$ is G-equivariant and onto O, so all its fibres have the same size c. Moreover, $p_{h_{\boldsymbol{y}}} = \gamma p_h = \varphi_h(\gamma \vec{\sup}(h)) = \psi_h(\boldsymbol{u}, \boldsymbol{y})$.

If $N = \emptyset$, then $O = \{h\}$ since $G \leq \mathbf{Stab}^\bullet(\sup(h))$, and we let $\psi_\Sigma^O := \psi_\Pi^O := \psi_h \in \mathfrak{T}_n^{\mathrm{sc}(g),k}$. Otherwise, the wires from O to g are support-changing, so $\mathrm{sc}(g) \geq \mathrm{sc}(h) + 1$. By Lemma B.6, applied to ψ_h with the fibres of $\boldsymbol{y} \mapsto h_{\boldsymbol{y}}$ as classes, there are ψ_Σ^O and ψ_Π^O in $\mathfrak{T}_n^{1+\mathrm{sc}(h),k} \subseteq \mathfrak{T}_n^{\mathrm{sc}(g),k}$ such that

$$\psi_\Sigma^O(\boldsymbol{u}) = \sum_{h' \in O} p_{h'} \quad \text{and} \quad \psi_\Pi^O(\boldsymbol{u}) = \prod_{h' \in O} p_{h'}.$$

If g is a sum gate, let $\varphi_g := \sum_O m_O \psi_\Sigma^O$. If g is a product gate, let φ_g be the pointwise product of the maps $(\psi_\Pi^O)^{\odot m_O}$. Then $\varphi_g \in \mathfrak{T}_n^{\mathrm{sc}(g),k}$ by Lemma B.3, since $|\sup(g)| \leq k$, and $\varphi_g(\boldsymbol{u}) = p_g$. □

Theorem B.8. *Let $\Delta \geq 1$, $k \geq 2$, and $n \geq 2k+1$. If p is computed by a reduced matrix-symmetric circuit of support size $\leq k$ and support depth $\leq \Delta$, then p is a linear combination of homomorphism polynomials $\hom_{F,n}$ for patterns F with $\mathrm{fw}_\Delta(F) \leq k-1$.*

Lean formalisation on LaxLean-Archive

Proof. Since $k < n/2$, Lemma B.7 applies. The output gate is fixed by Γ_n, so its minimal support is empty. Hence, $p \in \mathfrak{T}_n^{\Delta,k}(0,0)$, and the claim follows from Lemma B.4. □

B.3. Combinatorics of bounded-depth tree decompositions

In this section, we introduce a dual notion to Δ-forestwidth certifying that a graph has high Δ-forestwidth. We then show that this notion is reflected by weak oddomorphisms. These are graph homomorphisms with certain parity constraints introduced by Roberson [Rob26] arising

in the analysis of CFI graphs. We use such graphs in Theorem 3.7, the main tool for proving symmetric lower bounds on homomorphism polynomials.

B.3.1. Richness

We introduce a dual definition to Δ-forestwidth following [Bha+25, Lemma 18].

Definition B.9. Let $k \geq 1$. A graph F is $(1,k)$-*rich* if some connected component of F has at least $k+1$ vertices. For $\Delta \geq 2$, a connected graph F is (Δ,k)-*rich* if it contains $k+1$ pairwise vertex-disjoint connected subgraphs each of which is $(\Delta-1,k)$-rich. A disconnected graph is (Δ,k)-*rich* if it has a (Δ,k)-rich connected component.

Lemma B.10. *Let $\Delta, k \geq 1$. If a graph F is (Δ,k)-rich, then $\mathrm{fw}_\Delta(F) \geq k$.*

Proof. The proof essentially follows [Bha+25, Lemma 18]. It suffices to treat a rich connected component, so let F be connected. Suppose that F has a tree decomposition of depth at most Δ with empty root bag and all bags of size at most k. As the root bag is empty and F is connected, F lies in the bags of the subtree (T, r, β) rooted at a single child r of the root, which has depth at most $\Delta - 1$. Induction on Δ. For $\Delta = 1$, T is a single node, so $|V(F)| \geq k+1$ forces a bag of size at least $k+1$.

Let $\Delta \geq 2$ and $R := \beta(r)$, so $|R| \leq k$. Among the $k+1$ disjoint $(\Delta-1,k)$-rich connected subgraphs $F_1, \ldots, F_{k+1}$ of F, some F_i is disjoint from R. A vertex outside R occurs in the bags of exactly one child subtree of r, and two adjacent vertices outside R occur in the same one, because a bag containing both cannot be $\beta(r)$. Hence $V(F_i)$ lies in the bags of one child subtree T_c, and adding a root with an empty bag above $(T_c, c, \beta|_{T_c} \cap V(F_i))$ yields a tree decomposition of F_i of depth at most $\Delta - 1$ with empty root bag and bags of size at most k. So $\mathrm{fw}_{\Delta-1}(F_i) \leq k-1$, contradicting the induction hypothesis. □

Let P_m denote the m-vertex path.

Example B.11. *Let $\Delta, k \geq 1$. The path $P_{(k+1)^\Delta}$ is (Δ,k)-rich.*

Proof. We show by induction on Δ that every path with at least $(k+1)^\Delta$ vertices is (Δ,k)-rich. For $\Delta = 1$, the path is connected and has at least $k+1$ vertices. For $\Delta \geq 2$, the path is connected and splits into $k+1$ pairwise vertex-disjoint subpaths with at least $(k+1)^{\Delta-1}$ vertices each, which are $(\Delta-1,k)$-rich by induction. □

Example B.12 ([Bha+25])**.** *Let $\Delta, k \geq 1$. The full $(k+1)$-ary tree $T_{\Delta,k+1}$ of depth Δ is (Δ,k)-rich.*

Proof. This follows by induction on Δ: The tree $T_{1,k+1}$ is connected and has $k+2$ vertices. The tree $T_{\Delta+1,k+1}$ is connected and the subtrees rooted at the $k+1$ children of its root are disjoint copies of $T_{\Delta,k+1}$. In particular, $\mathrm{fw}_\Delta(T_{\Delta,k+1}) \geq k$ by Lemma B.10, cf. Example 3.4. □

Lemma B.13. *Let $\Delta, k \geq 1$. If a connected graph F has a (Δ,k)-rich minor, then F is (Δ,k)-rich.*

Proof. Let P be a rich minor of F with a model $(X_v)_{v \in V(P)}$ of disjoint connected branch sets in F. We may assume P connected. For $\Delta = 1$, $|V(F)| \geq |V(P)| \geq k+1$. For $\Delta \geq 2$, let $P_1, \ldots, P_{k+1}$ be disjoint $(\Delta-1,k)$-rich connected subgraphs of P; then $F_i := F[\bigcup_{v \in V(P_i)} X_v]$ are disjoint connected subgraphs of F having P_i as a minor, hence $(\Delta-1,k)$-rich by induction, and F is connected. □

Lemma B.14. *Let $\Delta, k \geq 1$. Every (Δ,k)-rich graph F has a (Δ,k)-rich tree P as a minor with $|V(P)| = (k+1)^\Delta$ and $\mathrm{maxdeg}(P) \leq \Delta k$.*

Proof. Induction on Δ, within a rich component F. For $\Delta = 1$, any subtree of F with $k+1$ vertices is $(1,k)$-rich and has maximum degree at most k.

Let $\Delta \geq 2$ and let $H_1, \dots, H_{k+1}$ be disjoint $(\Delta-1,k)$-rich connected subgraphs of F. By induction, each H_i has a $(\Delta-1,k)$-rich tree minor P_i with $(k+1)^{\Delta-1}$ vertices and maximum degree at most $(\Delta-1)k$. The branch sets of models of the P_i in the H_i are pairwise disjoint connected subsets of $V(F)$. Since F is connected, we may enlarge them to a partition of $V(F)$ into connected sets by repeatedly adding an uncovered vertex to a branch set containing one of its neighbours. The enlarged sets still form models of the P_i in F; let Y_i be the union of the branch sets of the model of P_i, a connected set. The graph obtained from F by contracting each of $Y_1, \dots, Y_{k+1}$ to a single vertex is connected, so it has a spanning tree T with k edges. For each edge ij of T, pick an edge of F between Y_i and Y_j; it joins the branch set of some vertex of P_i to the branch set of some vertex of P_j. Let P be the disjoint union of the P_i together with the k edges between the corresponding vertices. Then P is a minor of F and a tree, since it is obtained from the tree T by replacing every vertex by a tree. It contains the disjoint $(\Delta-1,k)$-rich subtrees P_i, so it is (Δ,k)-rich, and it has $(k+1)\cdot(k+1)^{\Delta-1} = (k+1)^{\Delta}$ vertices. As T has $k+1$ vertices, at most k of the added edges are incident to vertices of any fixed P_i, so the maximum degree of P is at most $(\Delta-1)k + k = \Delta k$. □

B.3.2. Oddomorphisms reflect richness

Lean formalisation on LaxLean-Archive

Definition B.15 ([Rob26]). Let F and G be graphs. Let $\varphi\colon F \to G$ be a homomorphism. A vertex $a \in V(F)$ is *φ-even / φ-odd* if $|N_F(a) \cap \varphi^{-1}(x)|$ is even/odd for every $x \in N_G(\varphi(a))$. The homomorphism φ is an *oddomorphism* if

1. every vertex $a \in V(F)$ is φ-even or φ-odd,
2. for every $x \in V(G)$, the set $\varphi^{-1}(x)$ contains an odd number of φ-odd vertices.

The homomorphism φ is a *weak oddomorphism* if there is a subgraph $F' \subseteq F$ such that $\varphi|_{F'}\colon F' \to G$ is an oddomorphism.

Lemma B.16 ([Rob26, Theorem 3.13, Lemma 5.2]). *Let $\varphi\colon F \to G$ be an oddomorphism.*

1. *If G is connected, then there exists a connected subgraph F' of F such that $\varphi|_{F'}\colon F' \to G$ is an oddomorphism.*
2. *For every $U \subseteq V(G)$, the restriction of φ to $F[\varphi^{-1}(U)]$ is an oddomorphism $F[\varphi^{-1}(U)] \to G[U]$.*

Proof. For the second item, let $F_U := F[\varphi^{-1}(U)]$. For $a \in \varphi^{-1}(U)$ and $x \in U$, it holds that $N_{F_U}(a) \cap \varphi^{-1}(x) = N_F(a) \cap \varphi^{-1}(x)$. Hence, every φ-even or φ-odd vertex of F in $\varphi^{-1}(U)$ is even or odd, respectively, for the restriction of φ to F_U. Moreover, the fibres of φ over the vertices in U are unchanged. □

Theorem B.17. *Let $\Delta, k \geq 1$. Let $F \to G$ be a weak oddomorphism. If G is (Δ,k)-rich, then F is (Δ,k)-rich.*

The proof is similar to that of [NS26, Theorem 3.6].

Proof. A weak oddomorphism restricts to an oddomorphism from a subgraph of F, and a graph with a rich subgraph is rich, so it suffices to show: if $\varphi\colon F \to G$ is an oddomorphism and G is (Δ,k)-rich, then F is (Δ,k)-rich. We induct on Δ. Let G_0 be a rich component of G. By Lemma B.16(2) with $U = V(G_0)$ and then Lemma B.16(1), some connected subgraph $F_0 \subseteq F$ admits an oddomorphism onto G_0; it suffices to show that F_0 is rich. So we may assume that F and G are connected.

For $\Delta = 1$, every fibre of an oddomorphism is non-empty, so $|V(F)| \geq |V(G)| \geq k+1$. For $\Delta \geq 2$, let $G_1, \dots, G_{k+1}$ be pairwise disjoint $(\Delta-1,k)$-rich connected subgraphs of G. For each i, by Lemma B.16(2) with $U = V(G_i)$, the restriction of φ to $F_i := F[\varphi^{-1}(V(G_i))]$ is an oddomorphism onto $G[V(G_i)]$, which is connected and $(\Delta-1,k)$-rich. By induction F_i is $(\Delta-1,k)$-rich, so it contains a $(\Delta-1,k)$-rich connected subgraph D_i. The D_i are pairwise disjoint, since they lie over disjoint vertex sets, and F is connected. Hence F is (Δ,k)-rich. □

We collect the facts about (Δ,k)-rich graphs proven here and thereby verify the hypotheses of Theorem 3.7.

Corollary B.18. *Let $\Delta, k \geq 1$. The class $\mathcal{F}$ of all graphs that are not (Δ,k)-rich contains K_1 and is closed under taking minors and weak oddomorphisms.*

Proof. Every (Δ,k)-rich graph has at least $k+1 \geq 2$ vertices. Hence, $K_1 \in \mathcal{F}$. If a graph F has a (Δ,k)-rich minor, then a rich connected component of this minor is a minor of a connected component of F, which is (Δ,k)-rich by Lemma B.13. Hence, $\mathcal{F}$ is minor-closed. By Theorem B.17, $\mathcal{F}$ is closed under weak oddomorphisms. □

B.4. Non-uniform descriptive complexity monotonicity

We prove Theorem 3.7.

Lean formalisation on LaxLean-Archive

Theorem 3.7. *Let $\mathcal{F}$ be a class of simple graphs. Suppose that $\mathcal{F}$ contains the 1-vertex graph K_1, that $\mathcal{F}$ is minor-closed, and that $\mathcal{F}$ is closed under weak oddomorphism. Let $G \notin \mathcal{F}$ be a connected graph. Let F be a bipartite multigraph containing G as a minor. Let $N \geq \gamma(G) := \sum_{v \in V(G)} 2^{\deg_G(v)-1}$. Then $\hom_{F,N} \notin \mathfrak{H}_N(\mathcal{F})$.*

The proof of Theorem 3.7 follows [DPS26a, Theorem 6.1], where a graph $F \notin \mathcal{F}$ is distinguished from $\mathcal{F}$ by means of the CFI graphs of an *induced* minor of F that does not belong to $\mathcal{F}$. The size of these CFI graphs is exponential in the maximum degree of the induced minor, which may be much larger than the maximum degree of the forbidden minor G it contains. We avoid this blow-up by modifying the CFI graphs such that they can be used for minors which are not necessarily induced. Furthermore, we work with bipartite graphs with fixed bipartitions throughout, which renders the dependence of the bound in [DPS26a] on the chromatic number of F obsolete.

B.4.1. CFI graphs for non-induced minors

We recall the CFI graphs of [Rob26], cf. [CFI92]. For a connected graph G and $v \in V(G)$, write $E(v) \subseteq E(G)$ for the set of edges incident to v. For $W \in \mathbb{Z}_2^{V(G)}$, the graph G_W has vertices (v,T) for $v \in V(G)$ and $T \in \mathbb{Z}_2^{E(v)}$ such that $\sum_{e \in E(v)} T(e) = W(v)$. Two vertices (u,S) and (v,T) are adjacent if $uv \in E(G)$ and $S(uv) = T(uv)$. Up to isomorphism, G_W depends only on $\sum_{v \in V(G)} W(v) \in \mathbb{Z}_2$ [Rob26, Corollary 3.7]. We write G_0 and G_1 for the two resulting graphs. Each of them has $\gamma(G)$ vertices. The map $\rho\colon G_i \to G$, $(v,T) \mapsto v$ is a homomorphism. For a homomorphism $\psi\colon F \to G$, write $\hom_\psi(F, G_i)$ for the number of homomorphisms $h\colon F \to G_i$ such that $\rho \circ h = \psi$.

Lemma B.19 ([Rob26, Theorem 3.6 and proof of Theorem 3.13])**.** *Let G be a connected graph and $\psi\colon F \to G$ be a homomorphism. Then $\hom_\psi(F, G_0) \geq \hom_\psi(F, G_1)$ with strict inequality if, and only if, ψ is a weak oddomorphism.*

For $i \in \{0,1\}$, let $\widehat{G}_i$ denote the graph obtained from G_i by adding all edges $(u,S)(v,T)$ for distinct non-adjacent vertices $u, v \in V(G)$. That is, the non-edges of G are replaced by complete bipartite graphs between the respective gadgets. Clearly, $\widehat{G}_i$ has $\gamma(G)$ vertices.

Lemma B.20. *Let G be a connected graph and F be a graph. Then $\hom(F, \widehat{G}_0) \geq \hom(F, \widehat{G}_1)$ with strict inequality if, and only if, there exist a subgraph F' of F with $V(F') = V(F)$ and a weak oddomorphism $\varphi\colon F' \to G$ such that $\varphi(a) \neq \varphi(b)$ for all $ab \in E(F)$.*

Proof. Write K for the complete graph on $V(G)$. Since $\widehat{G}_i$ has no edges between vertices (v,S) and (v,T), the map $\rho\colon \widehat{G}_i \to K$, $(v,T) \mapsto v$ is a homomorphism. For a homomorphism $\psi\colon F \to K$, let F_ψ denote the subgraph of F with vertex set $V(F)$ and edge set $\{ab \in E(F) \mid \psi(a)\psi(b) \in E(G)\}$. Then ψ is a homomorphism $F_\psi \to G$. Let $h\colon V(F) \to V(\widehat{G}_i)$ be a map such that $\rho \circ h = \psi$. If $ab \in E(F) \setminus E(F_\psi)$, then $\psi(a)$ and $\psi(b)$ are distinct and non-adjacent in G, so $h(a)h(b) \in E(\widehat{G}_i)$. If $ab \in E(F_\psi)$, then $h(a)h(b) \in E(\widehat{G}_i)$ if, and only if, $h(a)h(b) \in E(G_i)$. Hence, h is a homomorphism $F \to \widehat{G}_i$ if, and only if, it is a homomorphism $F_\psi \to G_i$. It follows that

$$\hom(F, \widehat{G}_i) = \sum_{\psi\colon F \to K} \hom_\psi(F_\psi, G_i). \tag{6}$$

By Lemma B.19, this implies that $\hom(F, \widehat{G}_0) \geq \hom(F, \widehat{G}_1)$ and that the inequality is strict if, and only if, there is a homomorphism $\psi\colon F \to K$ such that $\psi\colon F_\psi \to G$ is a weak oddomorphism. If such a ψ exists, then $F' := F_\psi$ and $\varphi := \psi$ are as desired. Conversely, let F' and φ be as in the lemma. Then φ is a homomorphism $F \to K$ and F' is a subgraph of F_φ. Since φ restricts to an oddomorphism on a subgraph of F', the homomorphism $\varphi\colon F_\varphi \to G$ is a weak oddomorphism. □

Next, we lift Lemma B.20 to bipartite graphs containing G as a minor. For bipartite graphs F and B with fixed bipartitions, write $\hom_{\text{bip}}(F, B)$ for the number of homomorphisms $F \to B$ that map the left side of F to the left side of B and the right side of F to the right side of B. For a simple graph H, let $B(H)$ denote the simple bipartite graph with left side $V(H) \times \{0\}$ and right side $V(H) \times \{1\}$ in which $(u,0)$ and $(v,1)$ are adjacent if $u = v$ or $uv \in E(H)$. That is, $B(H)$ is obtained from the bipartite double cover of H by adding the perfect matching $(v,0)(v,1)$ for $v \in V(H)$. For a graph F, write $\Gamma(F)$ for the set of all partitions $\mathcal{R}$ of $V(F)$ such that $F[R]$ is connected for all $R \in \mathcal{R}$. For $\mathcal{R} \in \Gamma(F)$, the simple graph $F/\mathcal{R}$ has vertex set $\mathcal{R}$ and two distinct parts are adjacent if there is an edge of F between them.

Lemma B.21. *For every simple bipartite graph F with fixed bipartition and every simple graph H,*

$$\hom_{\text{bip}}(F, B(H)) = \sum_{\mathcal{R} \in \Gamma(F)} \hom(F/\mathcal{R}, H). \tag{7}$$

Note that the graphs $F/\mathcal{R}$ and H on the right-hand side of Equation (7) are not necessarily bipartite and that the right-hand side does not depend on the bipartition of F.

Proof. Write $A \uplus A'$ for the fixed bipartition of F. Call a map $h\colon V(F) \to V(H)$ *admissible* if $h(a) = h(a')$ or $h(a)h(a') \in E(H)$ for all $aa' \in E(F)$. By definition of $B(H)$, the bipartition-respecting homomorphisms $F \to B(H)$ correspond bijectively to the admissible maps h via $a \mapsto (h(a), 0)$ for $a \in A$ and $a' \mapsto (h(a'), 1)$ for $a' \in A'$. It remains to give a bijection between the admissible maps and the pairs $(\mathcal{R}, g)$ of a partition $\mathcal{R} \in \Gamma(F)$ and a homomorphism $g\colon F/\mathcal{R} \to H$.

Given an admissible map h, let $\mathcal{R}$ be the set of vertex sets of the connected components of the graphs $F[h^{-1}(v)]$ for $v \in V(H)$. Then $\mathcal{R} \in \Gamma(F)$ and h is constant on every part. Let

$g\colon \mathcal{R} \to V(H)$ send a part to the image of its vertices under h. If two distinct parts $R, S \in \mathcal{R}$ are adjacent in $F/\mathcal{R}$, then there is an edge $rs \in E(F)$ with $r \in R$ and $s \in S$. Since r and s are adjacent and lie in distinct connected components of the fibres of h, we have $h(r) \neq h(s)$ and thus $g(R)g(S) = h(r)h(s) \in E(H)$ as h is admissible. Hence, g is a homomorphism $F/\mathcal{R} \to H$.

Conversely, given $\mathcal{R} \in \Gamma(F)$ and a homomorphism $g\colon F/\mathcal{R} \to H$, let h send every vertex in a part $R \in \mathcal{R}$ to $g(R)$. This map is admissible: the endpoints of an edge of F lie in the same part or in two parts that are adjacent in $F/\mathcal{R}$. It remains to observe that the partition constructed from h as above is $\mathcal{R}$. Every part $R \in \mathcal{R}$ induces a connected subgraph on which h is constant. If there is an edge of F between two distinct parts R and S, then $g(R) \neq g(S)$ since H does not have loops. Hence, the connected components of the fibres of h are precisely the parts of $\mathcal{R}$. The two constructions are mutually inverse. □

Lemma B.21 is a variant of the identity

$$\hom(F, H \cdot K) = \sum_{\mathcal{R} \in \Gamma(F)} \hom(F/\mathcal{R}, H) \prod_{R \in \mathcal{R}} \hom(F[R], K) \tag{8}$$

of [Sep24b, Theorem 16] for the lexicographic product $H \cdot K$ of simple graphs H and K, which is used in [DPS26a] for K a clique whose size is the chromatic number of F. Here, $H \cdot K$ has vertex set $V(H) \times V(K)$ and (u, i) and (v, j) are adjacent if $uv \in E(H)$, or $u = v$ and $ij \in E(K)$. The proofs of both identities are analogous. For bipartite F and $K = K_2$, the factors $\hom(F[R], K_2)$ in Equation (8) are equal to 2, so the two right-hand sides differ only in the positive weights $2^{|\mathcal{R}|}$. The left-hand sides differ as follows: The graph $B(H)$ is the spanning subgraph of $H \cdot K_2$ comprising the edges between $V(H) \times \{0\}$ and $V(H) \times \{1\}$. While a homomorphism $F \to H \cdot K_2$ may choose for every part $R \in \mathcal{R}$ which side of the bipartition of $F[R]$ is mapped to $V(H) \times \{0\}$, this choice is dictated by the fixed bipartition of F for the homomorphisms counted by $\hom_{\text{bip}}(F, B(H))$. The advantage of Lemma B.21 lies in the size of the host graphs: in order to evaluate a homomorphism polynomial at $H \cdot K_2$, one has to pass to its bipartite double cover, which has $2|V(H)|$ vertices on either side, whereas $B(H)$ has only $|V(H)|$ vertices on either side. This saves a factor of 2 in Theorem 3.7.

Lemma B.22. *Let $\mathcal{C}$ be a minor-closed graph class. Let G be a connected graph such that no graph in $\mathcal{C}$ admits a weak oddomorphism to G. Let $B_i := B(\widehat{G}_i)$ for $i \in \{0, 1\}$. Then*

1. *B_0 and B_1 are simple bipartite graphs with $\gamma(G)$ vertices on either side,*
2. *$\hom_{\text{bip}}(F', B_0) = \hom_{\text{bip}}(F', B_1)$ for all bipartite $F' \in \mathcal{C}$ with any fixed bipartition, and*
3. *$\hom_{\text{bip}}(F, B_0) > \hom_{\text{bip}}(F, B_1)$ for every bipartite graph F with fixed bipartition that contains G as a minor.*

Proof. The first assertion is immediate. By Lemma B.21, for every bipartite graph F' with fixed bipartition and $i \in \{0, 1\}$,

$$\hom_{\text{bip}}(F', B_i) = \sum_{\mathcal{R} \in \Gamma(F')} \hom(F'/\mathcal{R}, \widehat{G}_i). \tag{9}$$

If $F' \in \mathcal{C}$, then all subgraphs of all graphs $F'/\mathcal{R}$ are minors of F' and hence in $\mathcal{C}$. By assumption, none of them admits a weak oddomorphism to G. By Lemma B.20, $\hom(F'/\mathcal{R}, \widehat{G}_0) = \hom(F'/\mathcal{R}, \widehat{G}_1)$ for all $\mathcal{R}$. This yields the second assertion.

For the third assertion, first observe that, by Equation (9) and Lemma B.20, $\hom_{\text{bip}}(F', B_0) \geq \hom_{\text{bip}}(F', B_1) \geq 1$ for every bipartite graph F' with fixed bipartition. The second inequality

holds since, for every vertex x of $\widehat{G}_1$, mapping the left side of F' to $(x,0)$ and the right side to $(x,1)$ is a homomorphism. Let F be a bipartite graph with fixed bipartition containing G as a minor. Since G is connected, it is a minor of a connected component C of F. Since $\hom_{\text{bip}}(F,B_i)$ is the product of the numbers $\hom_{\text{bip}}(C',B_i)$ over all connected components C' of F, it suffices to show that $\hom_{\text{bip}}(C,B_0) > \hom_{\text{bip}}(C,B_1)$. To that end, it suffices to argue that $\hom(C/\mathcal{R},\widehat{G}_0) > \hom(C/\mathcal{R},\widehat{G}_1)$ for some $\mathcal{R} \in \Gamma(C)$.

Since G is a minor of C, there exist pairwise disjoint sets $X_v \subseteq V(C)$ for $v \in V(G)$ such that each $C[X_v]$ is connected and there is an edge of C between X_u and X_v for all $uv \in E(G)$. Since C is connected, we may suppose that the X_v partition $V(C)$ by repeatedly adding an uncovered vertex to a set X_v containing one of its neighbours. Let $\mathcal{R} := \{X_v \mid v \in V(G)\} \in \Gamma(C)$. Identifying X_v with v, the graph $C/\mathcal{R}$ has vertex set $V(G)$ and its edge set contains $E(G)$. By Lemma B.20, $\hom(C/\mathcal{R},\widehat{G}_0) > \hom(C/\mathcal{R},\widehat{G}_1)$. □

In contrast to [DPS26a, Theorem 6.1], the graph $C/\mathcal{R}$ in the last step of the proof is an induced minor of F which contains G as a spanning subgraph but may have many more edges. These edges are irrelevant for $\widehat{G}_i$ by Lemma B.20. Furthermore, the matching edges $(v,0)(v,1)$ of $B(\widehat{G}_i)$ replace the lexicographic products with cliques and the additional universal vertices of [DPS26a], whose purpose is to account for contractions and deletions. This is possible since the patterns and the hosts are bipartite graphs with fixed bipartitions: a connected subgraph of the pattern may be mapped onto a single matching edge.

B.4.2. Proof of Theorem 3.7

Proof of Theorem 3.7. Since $K_1 \in \mathcal{F}$, the graph G has at least two vertices. Since $\mathcal{F}$ is closed under weak oddomorphisms and $G \notin \mathcal{F}$, no graph in $\mathcal{F}$ admits a weak oddomorphism to G. Let F° denote the underlying simple graph of F with the bipartition of F. Write $F^\circ = F' + \ell K_1$ where F' does not have any isolated vertices and $\ell \geq 0$. Since G is connected and has at least two vertices, G is a minor of F'.

Let B_0 and B_1 be the bipartite graphs from Lemma B.22 for $\mathcal{C} := \mathcal{F}$ and G. Let $m := N - \gamma(G)$ and let B_i' be obtained from B_i by adding m isolated vertices on either side, for $i \in \{0,1\}$. We may suppose that B_0' and B_1' have left side $[N] \times \{0\}$ and right side $[N] \times \{1\}$.

Claim B.22a. $\hom_{\text{bip}}(F'',B_0') = \hom_{\text{bip}}(F'',B_1')$ for all bipartite $F'' \in \mathcal{F}$ with any fixed bipartition.

Proof of Claim. Write $C_1,\dots,C_r$ for the connected components of F''. Since $\mathcal{F}$ is minor-closed, $C_1,\dots,C_r \in \mathcal{F}$. Let $e_j := m$ if C_j is a single vertex and $e_j := 0$ otherwise. This is the number of bipartition-respecting homomorphisms from C_j to the $2m$ additional vertices of B_i'. By Lemma B.22,

$$\hom_{\text{bip}}(F'',B_0') = \prod_{j=1}^{r} \left(\hom_{\text{bip}}(C_j,B_0) + e_j\right) = \prod_{j=1}^{r} \left(\hom_{\text{bip}}(C_j,B_1) + e_j\right) = \hom_{\text{bip}}(F'',B_1'). \qquad \lrcorner$$

Claim B.22b. $\hom_{\text{bip}}(F^\circ,B_0') \neq \hom_{\text{bip}}(F^\circ,B_1')$.

Proof of Claim. Since F' does not have isolated vertices, no homomorphism $F' \to B_i'$ has an isolated vertex of B_i' in its image. Hence, $\hom_{\text{bip}}(F^\circ,B_i') = N^\ell \hom_{\text{bip}}(F',B_i') = N^\ell \hom_{\text{bip}}(F',B_i)$. Since G is a minor of F', the claim follows from Lemma B.22. $\lrcorner$

Let $X_i \in \{0,1\}^{N\times N}$ denote the bi-adjacency matrix of B_i' for $i \in \{0,1\}$. Since the entries of X_i are 0 or 1, evaluating the homomorphism polynomial of a bipartite multigraph at X_i yields the number of bipartition-respecting homomorphisms from its underlying simple graph to B_i'. By

Claim B.22a, the linear map $p \mapsto p(X_0) - p(X_1)$ vanishes on all generators of $\mathfrak{H}_N(\mathcal{F})$ and hence on $\mathfrak{H}_N(\mathcal{F})$. By Claim B.22b and since $\mathbb{K}$ has characteristic zero, it does not vanish on $\hom_{F,N}$. Hence, $\hom_{F,N} \notin \mathfrak{H}_N(\mathcal{F})$. □

B.5. Separation

We now put the pieces together and prove Theorem 3.1. To that end, we first deduce the following corollary of Corollary B.18 and Theorem 3.7.

Corollary B.23. *Let $\Delta \geq 1$ and $k \geq 2$. Let G be a connected (Δ, k)-rich graph and F be a bipartite multigraph containing G as a minor. Let $N \geq \max\{\gamma(G), 4(k+1), 20\}$. Then every matrix-symmetric circuit of support depth $\leq \Delta$ computing $\hom_{F,N}$ has orbit size at least $\binom{N}{k+1}$.*

Proof. Towards a contradiction, let C be such a circuit of orbit size less than $\binom{N}{k+1}$. Let C' be the reduced circuit constructed from C in Lemma 2.4. Since $N \geq 4(k+1) \geq 12$, the orbit size of C' is at most $f := \binom{N}{k+1} - 1 < \binom{N}{k+2}/2 \leq \binom{N}{\lfloor N/2 \rfloor}/2$ and $\kappa_f(N) \leq k+1 \leq N/4$. By Theorem 2.6, $\operatorname{maxSup}(C') \leq k < N/2$. Hence, C' has support depth $\leq \Delta$ by definition. By Theorem 3.5 for C', the polynomial $\hom_{F,N}$ is a linear combination of homomorphism polynomials $\hom_{F',N}$ for patterns F' with $\operatorname{fw}_\Delta(F') \leq k-1$. Let $\mathcal{F}$ denote the class of all simple graphs that are not (Δ, k)-rich. By Lemma B.10, $\hom_{F,N} \in \mathfrak{H}_N(\mathcal{F})$. Since $G \notin \mathcal{F}$ is a connected minor of F and $N \geq \gamma(G)$, this contradicts Theorem 3.7, which is applicable by Corollary B.18. □

Lean formalisation on LaxLean-Archive

Theorem 3.1. *Let $\Delta \geq 1$ and $n \geq 20$. Let $T_{\Delta,n}$ denote the full n-ary tree of depth Δ.*

1. *Every matrix-symmetric circuit of support depth $\leq \Delta$ computing $\hom_{T_{\Delta,n},n}$ has orbit size, and hence size, at least*
$$n^{\log_2 n - \Delta \log_2 \log_2 n - 3}.$$
2. *There is a matrix-symmetric circuit of support depth $\Delta+1$, depth $2\Delta+3$, orbit size $\leq n^2$, and size $\leq 13n^{\Delta+2}$ computing $\hom_{T_{\Delta,n},n}$.*

Proof. For the second assertion, recall that $\operatorname{fw}_{\Delta+1}(T_{\Delta,n}) = 1$ by Example 3.4. By Theorem B.2 for $\Delta+1$ and $k := 2$, the polynomial $\hom_{T_{\Delta,n},n}$ admits a matrix-symmetric circuit of support depth $\leq \Delta+1$, orbit size $\leq n^2$, depth $\leq 2(\Delta+1)+1 = 2\Delta+3$, and size $\leq 5n^2(|V(T_{\Delta,n})| + |E(T_{\Delta,n})|) + 1$. Since $n \geq 20$, we have $|V(T_{\Delta,n})| = \sum_{j=0}^{\Delta} n^j \leq \frac{n}{n-1} n^\Delta \leq \frac{20}{19} n^\Delta$ and $|E(T_{\Delta,n})| = |V(T_{\Delta,n})| - 1$. Hence, the size is at most $5n^2 \cdot \frac{40}{19} n^\Delta + 1 \leq 13n^{\Delta+2}$.

For the first assertion, let $\Delta \geq 1$ and $n \geq 20$. Write $L := \log_2 n$ and $\ell := \log_2 L$, and let $k := \lfloor L - (\Delta-1)\ell - 3 \rfloor$. If $k < 2$, then $L - \Delta\ell - 3 < 2 - \ell \leq 0$ since $\ell \geq 2$, and the asserted bound is trivial. Hence, suppose that $k \geq 2$. Then $n \geq 32$ and $k+1 \leq L \leq n/4$.

Let $G := T_{\Delta,k+1}$, which is a subgraph of $T_{\Delta,n}$. By Example B.12, G is (Δ, k)-rich. The root of G has degree $k+1$, its $\sum_{j=1}^{\Delta-1} (k+1)^j \leq 2(k+1)^{\Delta-1}$ other inner vertices have degree $k+2$, and its $(k+1)^\Delta$ leaves have degree 1. Hence, using $k+1 \leq 2^k$,

$$\gamma(G) \leq 2^k + 2^{k+2}(k+1)^{\Delta-1} + (k+1)^\Delta \leq 2^{k+3}(k+1)^{\Delta-1} \leq 2^{L-(\Delta-1)\ell} L^{\Delta-1} = n.$$

Let C be a matrix-symmetric circuit of support depth $\leq \Delta$ computing $\hom_{T_{\Delta,n},n}$. Since G is a connected (Δ, k)-rich subgraph of $T_{\Delta,n}$ and $n \geq \max\{\gamma(G), 4(k+1)\}$, Corollary B.23 for Δ and k applies. Hence, the orbit size of C is at least $\binom{n}{k+1} \geq \left(\frac{n}{k+1}\right)^{k+1}$. Since $L - (\Delta-1)\ell - 3 < k+1 \leq L$,

$$\log_2 \binom{n}{k+1} \geq (k+1)(L-\ell) \geq (L - (\Delta-1)\ell - 3)(L-\ell) \geq L(L - \Delta\ell - 3).$$

Thus, the orbit size of C is at least $n^{L-\Delta\ell-3}$. □

C. Material omitted in Section 4

C.1. Omitted material on reduction to logarithmic depth

Theorem 4.3. *Let $\varepsilon > 0$. Let (p_n) be a family of matrix-symmetric polynomials of degree $\le d(n)$ computed by matrix-symmetric circuits of orbit size $\le s(n) \le 2^{n^{1-\varepsilon}}$. Then (p_n) is computed by matrix-symmetric circuits of depth $O(\log d(n))$ and orbit size $\le \max\{s(n), n^2\}^{3/\varepsilon}$ for all sufficiently large n.*

Proof. Let n be sufficiently large, and write $p := p_n$, $d := d(n)$, $s := \max\{s(n), n^2\}$, $\kappa := \kappa_s(n)$, and $r := \lceil n^{1-\varepsilon} \rceil$. Since $s \ge n^2 > \binom{n}{2}$, we have $\kappa \ge 3$. Since $\binom{n}{r} \ge (n/r)^r \ge (n^\varepsilon/2)^{n^{1-\varepsilon}} > 2^{n^{1-\varepsilon}} \ge s$, we have $\kappa \le r$, i.e. $\kappa - 1 \le n^{1-\varepsilon}$. Hence $n \ge 20$ and $n \ge 6\kappa - 5$. By minimality of κ, $(n/(\kappa-1))^{\kappa-1} \le \binom{n}{\kappa-1} \le s$, and $n/(\kappa-1) \ge n^\varepsilon$ yields $n^{\varepsilon(\kappa-1)} \le s$, i.e. $n^{3\kappa-3} \le s^{3/\varepsilon}$.

Let C be a matrix-symmetric circuit of orbit size $\le s$ computing p. By Lemma 2.4, we may suppose that C is reduced. Let $\Delta \ge 1$ bound its depth, so that C has support depth $\le \Delta$. As $3\kappa \le n + 2$, Theorem 2.6 shows that C has support size $\le \kappa - 1$. By Theorem 3.5 with $k := \kappa - 1 \ge 2$, we have $p = \sum_F \alpha_F \hom_{F,n}$ for patterns F with $\mathrm{fw}_\Delta(F) \le \kappa - 2$. In particular, $\mathrm{tw}(F) \le \kappa - 2$. Since $\hom_{F,n}$ is homogeneous of degree $|E(F)|$, comparing homogeneous components shows that we may keep only the patterns with $|E(F)| \le d$. Patterns without edges contribute constants. For the others, we may delete isolated vertices by Observation 2.7, so that $|V(F)| \le 2d$.

By Theorem 4.2, every such F has a rooted tree decomposition with bags of size $\le 3(\kappa-1) = 3\kappa - 3$ and depth $\le q := 2\lceil \log_{5/4}(4d) \rceil$. Adding a root with an empty bag shows $\mathrm{fw}_{q+1}(F) \le 3\kappa - 4$. By Theorem B.2 with $k := 3\kappa - 3$, using $n \ge 6\kappa - 5 = 2k + 1$, $\hom_{F,n}$ is computed by a reduced matrix-symmetric circuit of depth $\le 2q + 3$ and orbit size $\le n^{3\kappa-3} \le s^{3/\varepsilon}$. Taking the linear combination as in Section B.1 increases the depth by one, so the depth is $\le 2q + 4 \in O(\log d)$. □

Let P_m denote the m-vertex path.

Lemma C.1. *Let $\Delta, k \ge 2$, $d \ge 1$, and $N \ge \max\{4(k+1), 20\}$ be such that $(k+1)^\Delta \le \min\{d, N/2\}$. Then every matrix-symmetric circuit of support depth $\le \Delta$ computing $\hom_{P_d,N}$ has orbit size at least $\binom{N}{k+1}$.*

Proof. Let $G := P_{(k+1)^\Delta}$. By Example B.11, G is (Δ, k)-rich. Since G has maximum degree 2, $\gamma(G) \le 2(k+1)^\Delta \le N$. The claim follows from Corollary B.23. □

Theorem 4.4. *Let $a \ge 1$ and $d\colon \mathbb{N} \to \mathbb{N}$ with $2 \le d(n) \le n^a$.*

1. *The polynomials $\hom_{P_{d(n)},n}$ have degree $< d(n)$ and are computed by matrix-symmetric circuits of orbit size $\le n^3$ and depth $O(\log d(n))$.*
2. *For every c, every matrix-symmetric circuit of orbit size $\le n^c$ computing $\hom_{P_{d(n)},n}$ has support depth, and hence depth, $\Omega_{c,a}(\log d(n))$.*

Proof. Write $d := d(n)$. For the upper bound, decompose P_d recursively: a node for a subpath with endpoints a and b and midpoint m has the bag $\{a, m, b\}$ and children for the subpaths from a to m and from m to b, and subpaths with one edge are leaves. Together with a root with an empty bag, this shows $\mathrm{fw}_{\lceil \log_2 d \rceil + 1}(P_d) \le 2$. By Theorem B.2 with $k := 3$, the polynomial $\hom_{P_d,n}$ is computed by a matrix-symmetric circuit of orbit size $\le n^3$ and depth $\le 2\lceil \log_2 d \rceil + 3$ for $n \ge 7$.

For the lower bound, we may suppose that $c \ge 2$ is an integer. Let n be sufficiently large, so that $n^c < \binom{n}{c+1}$ and $n \ge \max\{4(c+1), 20\}$, and let $\Delta := \lfloor \log_{c+1} \min\{d, n/2\} \rfloor$. If $\Delta \ge 2$, then

Lemma C.1 for $k := c$ shows that the support depth is $> \Delta \geq \log_{c+1} \min\{d, n/2\} - 1$, since $(c+1)^\Delta \leq \min\{d, n/2\}$. If $\Delta \leq 1$, then $\log_{c+1} \min\{d, n/2\} - 1 < 1$, and the support depth is ≥ 1 since $\hom_{P_d,n}$ is not constant. Finally, if $d > n/2$, then $\log_{c+1}(n/2) \geq \log_{c+1} n - 1 \geq \frac{1}{a} \log_{c+1} d - 1$ as $d \leq n^a$. □

C.2. Material omitted on bounded pathwidth and logarithmic depth

Let $c \geq 0$ and let $\Sigma := \{\ell_1, \dots, \ell_{c+1}, r_1, \dots, r_{c+1}\}$ be a set of *slots*, where the ℓ_i are *left* and the r_j are *right*. A *bilabelled graph* $\boldsymbol{F} = (F, \alpha, \beta)$ consists of a bipartite multigraph F and maps $\alpha, \beta \colon \Sigma \to V(F)$, the *in-* and *out-labels*, sending left slots to left and right slots to right vertices. Let $L_n := [n]^\Sigma$ and $N := |L_n| = n^{2(c+1)}$. The *homomorphism matrix* $\boldsymbol{F}_n \in \mathbb{K}[\mathcal{X}_n]^{L_n \times L_n}$ is

$$(\boldsymbol{F}_n)_{a,b} := \sum_{\substack{h \colon V(F) \to [n] \uplus [n] \\ h \circ \alpha = a,\ h \circ \beta = b}} \prod_{xy \in E(F)} x_{h(x),h(y)},$$

where h maps left vertices to the left and right vertices to the right copy of $[n]$. That is, $\boldsymbol{F}_n$ is the map of the $(2(c+1), 2(c+1))$-labelled graph with labels α and β, cf. Section B.2.1, arranged as a matrix. Clearly, $\mathbf{1}^T \boldsymbol{F}_n \mathbf{1} = \hom_{F,n}$ and $\operatorname{tr}(\boldsymbol{F}_n) = \hom_{F^\circ,n}$, where F° is obtained from F by identifying $\alpha(s)$ with $\beta(s)$ for all $s \in \Sigma$.

The *series composition* $\boldsymbol{F} \cdot \boldsymbol{F}'$ of $\boldsymbol{F} = (F, \alpha, \beta)$ and $\boldsymbol{F}' = (F', \alpha', \beta')$ is obtained from the disjoint union of F and F' by identifying $\beta(s)$ with $\alpha'(s)$ for all $s \in \Sigma$, with in-labels α and out-labels β'. Then $(\boldsymbol{F} \cdot \boldsymbol{F}')_n = \boldsymbol{F}_n \boldsymbol{F}'_n$, since homomorphisms from $F \cdot F'$ correspond to pairs of homomorphisms from F and F' that agree on the identified vertices, and summing over their common value $b = h \circ \beta = h' \circ \alpha'$ is matrix multiplication.

Let $\boldsymbol{I}$ be the edgeless bilabelled graph with vertex set Σ and $\alpha = \beta = \mathrm{id}$. The *atomic* graphs are the following $m := (c+1)^2 + 2(c+1)$ bilabelled graphs. For a left slot s and a right slot t, let $\boldsymbol{E}_{s,t}$ be $\boldsymbol{I}$ with an edge st added. For a slot s, let $\boldsymbol{R}_s$ be $\boldsymbol{I}$ with a new vertex s' on the side of s added and $\beta(s) := s'$. Then $(\boldsymbol{E}_{s,t})_n$ is the diagonal matrix with entries $x_{a(s),a(t)}$, and $(\boldsymbol{R}_s)_n$ is the 0/1-matrix whose (a,b)-entry is 1 if and only if a and b agree outside s. For a word $u = u_1 \cdots u_\ell$ over the atomic graphs, let $\boldsymbol{F}_u = (F_u, \alpha_u, \beta_u) := \boldsymbol{I} \cdot u_1 \cdots u_\ell$ and $F_u^\circ := (F_u)^\circ$.

Lemma C.2. *Let u be a word of length ℓ over the atomic graphs.*

1. $(\boldsymbol{F}_u)_n = (u_1)_n \cdots (u_\ell)_n$, $\mathbf{1}^T (\boldsymbol{F}_u)_n \mathbf{1} = \hom_{F_u,n}$, *and* $\operatorname{tr}((\boldsymbol{F}_u)_n) = \hom_{F_u^\circ,n}$.
2. *F_u and F_u° have at most $2(c+1) + \ell$ vertices,* $\operatorname{pw}(F_u) \leq 2c+1$, *and* $\operatorname{pw}(F_u^\circ) \leq 4c+3$.
3. *For every bipartite multigraph F with $\operatorname{pw}(F) \leq c$, there is a word u such that F_u is the disjoint union of F and $2(c+1)$ isolated vertices.*

Proof. 1. This follows from the above, as $\boldsymbol{I}_n$ is the identity matrix.

2. Every atomic graph $\boldsymbol{R}_s$ adds one vertex, and $\boldsymbol{E}_{s,t}$ adds none. For $i = 0, \dots, \ell$, let X_i be the image of the out-labels of $\boldsymbol{F}_{u_1 \cdots u_i}$, viewed in F_u. A vertex that loses its out-label never regains it, and the edge added by $u_i = \boldsymbol{E}_{s,t}$ has both endpoints in X_i. Hence, the bags $X_0, \dots, X_\ell$ of size $\leq 2(c+1)$ form a path decomposition of F_u. Since X_0 is the image of α_u and every vertex of F_u° obtained by identification contains a vertex of X_0, adding X_0 to every bag yields a path decomposition of F_u°.

3. Let $B_1, \dots, B_p$ be the bags of a path decomposition of F of width $\leq c$. The vertices on either side of F correspond to intervals of $[p]$, at most $c+1$ of which overlap in any point. Colour them greedily with the $c+1$ slots of their side, so that vertices sharing a bag receive distinct slots.

Assign every edge of F to a bag containing it. Let u be the concatenation, for $i = 1, \dots, p$, of the atomic graphs $\boldsymbol{E}_{s,t}$ for the edges assigned to B_i, where s and t are the slots of their endpoints, followed by the atomic graphs $\boldsymbol{R}_s$ for the slots s of the vertices in $B_i \setminus B_{i+1}$, where $B_{p+1} := \emptyset$. Identify the vertices of F with slot s, in the order of their intervals, with the successive vertices carrying the out-label s: the first with the vertex s of $\boldsymbol{I}$, and each further one with the vertex added by the $\boldsymbol{R}_s$ after the last bag of its predecessor. As intervals with the same slot are disjoint, the out-label s sits on the vertex of B_i with slot s while the atomic graphs for B_i are composed, so every edge of F is added between its endpoints. The vertex carrying the out-label s at the end is isolated, which gives the $2(c+1)$ isolated vertices. □

Lemma C.3. *Let $c \geq 0$, $n \geq 2$, $k := 12(c+1)$, and $\Delta := 2\lceil (4c+6)\log_{5/4} n \rceil + 1$. For every bipartite multigraph F with $\mathrm{pw}(F) \leq c$, the polynomial $\hom_{F,n}$ is a linear combination of homomorphism polynomials $\hom_{F',n}$ for patterns F' with $\mathrm{fw}_\Delta(F') \leq k-1$.*

Proof. By Lemma C.2 and Observation 2.7, $\hom_{F,n} = n^{-2(c+1)}\mathbf{1}^T(\boldsymbol{F}_w)_n\mathbf{1}$ for some word w. By Lemma C.2 and Theorem 5.3 applied to the m matrices of the atomic graphs with $N = n^{2(c+1)}$, the matrix $(\boldsymbol{F}_w)_n$ is a linear combination of matrices $(\boldsymbol{F}_u)_n$ with $|u| < N^2$ whose coefficients are polynomials in traces $\mathrm{tr}((\boldsymbol{F}_v)_n)$ with $|v| \leq N^2$. By Lemma C.2, $\mathbf{1}^T(\boldsymbol{F}_w)_n\mathbf{1}$ is thus a linear combination of products $\hom_{F_u,n}\prod_i \hom_{F_{v_i}^\circ,n} = \hom_{F',n}$, where F' is the disjoint union of F_u and the $F_{v_i}^\circ$. Each of these components has pathwidth, and hence treewidth, $\leq 4c+3$ and at most $2(c+1)+N^2$ vertices, where $2(2(c+1)+N^2) \leq 3N^2 \leq n^{4c+6}$. By Theorem 4.2, each component has a tree decomposition with bags of size $\leq 12(c+1) = k$ and depth $\leq \Delta - 1$. Adding a root with an empty bag above all of them shows that $\mathrm{fw}_\Delta(F') \leq k-1$. □

Theorem 5.2. $\mathsf{symVBP} \subseteq \mathsf{symVAC}^1$.

Proof. By Theorems 6.3 and 2.9, every family in symVBP is, for all sufficiently large n, a linear combination of homomorphism polynomials for patterns of pathwidth $\leq c$ for some constant c. By Lemma C.3, each of them is a linear combination of homomorphism polynomials for patterns F with $\mathrm{fw}_\Delta(F) \leq k-1$, where $k = 12(c+1)$ and $\Delta = O(c\log n)$. By Theorem 3.5 and Observation 2.5, these linear combinations are computed by matrix-symmetric circuits of support depth $\leq \Delta$ and orbit size $\leq n^k$ for all $n \geq 2k+1$. □

C.3. Separations of symVAC^k and symVP

We prove Theorems 5.5 and 5.7. The idea is to evaluate matrix-symmetric polynomials on weighted encodings of relational structures that witness the Weisfeiler–Leman iteration number lower bounds of Grohe, Lichter, and Neuen [GLN25].

C.3.1. Set-up

Let Σ be a finite set of labels. A Σ-*labelled bipartite graph* G consists of disjoint sets $A(G)$ and $B(G)$ and relations $E_\sigma(G) \subseteq A(G) \times B(G)$ for $\sigma \in \Sigma$. We view it as a relational structure with universe $A(G) \cup B(G)$, unary relations A and B, and binary relations E_σ. A homomorphism $F \to G$ maps $A(F)$ to $A(G)$ and $B(F)$ to $B(G)$ and preserves every E_σ; we write $\hom(F,G)$ for their number. Write C_d^r for the sentences of first-order logic with counting quantifiers with at most d variables and quantifier rank at most r, and $\mathsf{C}_d := \bigcup_r \mathsf{C}_d^r$.

For $K, q \geq 0$, let $\mathcal{T}_q^K$ denote the class of graphs F that admit a rooted tree decomposition (T,β) such that $|\beta(t)| \leq K$ for all $t \in V(T)$ and the bags along every root-leaf path of T contain at most q vertices in total [Adl+26]. We first translate between this notion and Definition 3.3.

Lemma C.4. *Let $K, \Delta, q \geq 1$, and let F be a graph.*

1. *If $\mathrm{fw}_\Delta(F) \leq K-1$, then $F \in \mathcal{T}^K_{\Delta K}$.*
2. *If $F \in \mathcal{T}^K_q$, then $\mathrm{fw}_q(F) \leq K-1$.*

Proof. 1. A decomposition witnessing $\mathrm{fw}_\Delta(F) \leq K-1$ has an empty root bag and depth $\leq \Delta$, so every root-leaf path meets at most Δ non-empty bags of size $\leq K$.

2. Let (T, β) be as in the definition of $\mathcal{T}^K_q$. As long as some node t has a child c with $\beta(c) \subseteq \beta(t)$, contract the edge tc into t, keeping the bag $\beta(t)$; this preserves the properties of (T, β). If the root bag is non-empty, add a new root with an empty bag. Now every node t other than the root has a vertex $v_t \in \beta(t)$ that is not in the bag of its parent. Since the nodes whose bags contain v_t form a subtree of T, the vertices v_t for the non-root nodes t on a root-leaf path are pairwise distinct. Hence, every root-leaf path has at most q non-root nodes, i.e. the decomposition has depth $\leq q$. □

We use the following known correspondence.

Theorem C.5 ([DJR21; Adl+26])**.** *Let G and H be Σ-labelled bipartite graphs and $K, q \geq 1$. Then G and H are C^q_K-equivalent if and only if $\hom(F, G) = \hom(F, H)$ for all Σ-labelled bipartite graphs $F \in \mathcal{T}^K_q$.*

C.3.2. Translating structures into weighted graphs

We translate the relational structures of [GLN25] to weighted graphs on which we evaluate the homomorphism matrices. First, we handle edge colours via interpolation. Then, we reduce from general relational structures to edge-coloured graphs.

Edge colours via interpolation. Let G be a Σ-labelled bipartite graph such that every pair $(a, b) \in A(G) \times B(G)$ lies in at most one relation $E_\sigma(G)$, and let $N \geq |A(G)|, |B(G)|$. For $\omega \in \mathbb{K}^\Sigma$, let $M^\omega_G \in \mathbb{K}^{N \times N}$ be the matrix whose rows and columns are indexed by $A(G)$ and $B(G)$, padded by zero rows and columns, with entry ω_σ at $(a, b) \in E_\sigma(G)$ and 0 at all other positions. Matrix-symmetric polynomials do not depend on the choice of the indexing. Let F be a bipartite multigraph with bipartition $A \uplus B$ and without isolated vertices. For a map $\lambda \colon E(F) \to \Sigma$, let F^λ be the Σ-labelled bipartite graph with $A(F^\lambda) = A$, $B(F^\lambda) = B$, and $(a, b) \in E_\sigma(F^\lambda)$ if some edge $e = ab$ has $\lambda(e) = \sigma$. Expanding the definition yields

$$\hom_{F,N}(M^\omega_G) = \sum_{\lambda \colon E(F) \to \Sigma} \Big(\prod_{e \in E(F)} \omega_{\lambda(e)} \Big) \hom(F^\lambda, G). \tag{10}$$

Here, we use that a homomorphism $F^\lambda \to G$ maps all parallel edges of F to the same pair, which lies in at most one relation.

Lemma C.6. *Let G and H be Σ-labelled bipartite graphs as above with $|A(G)| = |A(H)|$ and $|B(G)| = |B(H)|$, and let $N \geq |A(G)|, |B(G)|$. Let L be a Σ-labelled bipartite graph with $\hom(L, G) \neq \hom(L, H)$. Then there are a bipartite multigraph F whose underlying simple graph is a subgraph of the Gaifman graph of L, and $\omega \in \mathbb{K}^\Sigma$ such that $\hom_{F,N}(M^\omega_G) \neq \hom_{F,N}(M^\omega_H)$.*

Proof. Deleting an isolated vertex of L divides both $\hom(L, G)$ and $\hom(L, H)$ by the same number $|A(G)|$ or $|B(G)|$, which is non-zero since $\hom(L, G) \neq \hom(L, H)$, so we may suppose that L has no isolated vertices. Let F_0 be the bipartite multigraph with an edge $e = ab$ labelled $\lambda(e) := \sigma$ for every $\sigma \in \Sigma$ and $(a, b) \in E_\sigma(L)$, so that $L = F_0^\lambda$. Let $e_1, \ldots, e_s$ be the edges of F_0, and let

F be obtained by replacing e_i by 2^{i-1} parallel edges. By (10), $\hom_{F,N}(M^\omega_D)$ is a polynomial in ω for $D \in \{G, H\}$. Labellings of F that are not constant on the parallel copies of some e_i contribute 0, since every pair lies in at most one relation of D. The remaining labellings correspond to labellings μ of F_0, contribute $\prod_i \omega_{\mu(e_i)}^{2^{i-1}} \hom(F_0^\mu, D)$, and are determined by their monomial via the binary expansions of the exponents. So the coefficient of $\prod_i \omega_{\lambda(e_i)}^{2^{i-1}}$ is $\hom(L, D)$. Hence, $\hom_{F,N}(M^\omega_G) - \hom_{F,N}(M^\omega_H)$ is a non-zero polynomial in ω, and some $\omega \in \mathbb{K}^\Sigma$ does not annihilate it. □

Lemma C.7. *Let G, H, and N be as in Lemma C.6, and let $K, q \geq 1$. If G and H are C^q_K-equivalent, then $p(M^\omega_G) = p(M^\omega_H)$ for all $\omega \in \mathbb{K}^\Sigma$ and all linear combinations p of $\hom_{F,N}$ for bipartite multigraphs $F \in \mathcal{T}^K_q$.*

Proof. Isolated vertices of F contribute the same factor N on both sides by Observation 2.7, and deleting them preserves membership in $\mathcal{T}^K_q$. For F without isolated vertices and $\mu\colon E(F) \to \Sigma$, the Gaifman graph of F^μ is the underlying simple graph of F, so $F^\mu \in \mathcal{T}^K_q$ and $\hom(F^\mu, G) = \hom(F^\mu, H)$ by Theorem C.5. The claim follows by (10). □

Incidence graphs. Let $\mathfrak{D}$ be a relational structure with universe U over a finite vocabulary τ of arity $\leq m$, and let $\tau_s \subseteq \tau$ consist of the relation symbols of arity s. The incidence graph $\mathrm{Inc}(\mathfrak{D})$ is the Σ-labelled bipartite graph with $\Sigma := \{(s, J, S) \mid s \in [m],\ \emptyset \neq J \subseteq [s],\ S \subseteq \tau_s\}$, $A(\mathrm{Inc}(\mathfrak{D})) := U$, and $B(\mathrm{Inc}(\mathfrak{D})) := \bigcup_{s\in[m]} U^s$. For $u \in U$ and $\bar{t} \in U^s$ with $J := \{j \in [s] \mid t_j = u\} \neq \emptyset$, the pair $(u, \bar{t})$ lies in $E_{(s,J,S)}$ for $S := \{R \in \tau_s \mid \bar{t} \in R^\mathfrak{D}\}$, and in no other relation. Every pair lies in at most one relation. If $\mathfrak{A}$ and $\mathfrak{B}$ both have n elements, then the sides of $\mathrm{Inc}(\mathfrak{A})$ and $\mathrm{Inc}(\mathfrak{B})$ have the same sizes n and $\sum_{s\in[m]} n^s < (n+1)^m$.

Lemma C.8. *Let $\mathfrak{A}$ and $\mathfrak{B}$ be structures over a finite vocabulary of arity $\leq m$, and let $K, q \geq 1$.*

1. *If $\mathrm{Inc}(\mathfrak{A})$ and $\mathrm{Inc}(\mathfrak{B})$ are C^{q+1}_{K+1}-equivalent, then $\mathfrak{A}$ and $\mathfrak{B}$ are C^q_K-equivalent.*
2. *If $\mathfrak{A}$ and $\mathfrak{B}$ are C^{mq}_{mK}-equivalent, then $\mathrm{Inc}(\mathfrak{A})$ and $\mathrm{Inc}(\mathfrak{B})$ are C^q_K-equivalent.*

Proof. 1. We translate every C^q_K-sentence into a C^{q+1}_{K+1}-sentence by relativising every quantifier $\exists^{\geq\nu} x$ to $A(x)$ and replacing every atom $R(z_1, \ldots, z_s)$ by $\exists y \bigwedge_{j\in[s]} \bigvee_{(s,J,S)\in\Sigma,\, j\in J,\, R\in S} E_{(s,J,S)}(z_j, y)$, for a fresh variable y. The result has at most $K+1$ variables and quantifier rank at most $q+1$, since only atoms are replaced, and it holds in $\mathrm{Inc}(\mathfrak{D})$ if and only if the original sentence holds in $\mathfrak{D}$.

2. We translate every C^q_K-sentence over $\mathrm{Inc}(\mathfrak{D})$ into a C^{mq}_{mK}-sentence over $\mathfrak{D}$ via a m-dimensional interpretation. Every vertex of $\mathrm{Inc}(\mathfrak{D})$ is represented by a *type* and a m-tuple over U: a vertex $u \in U$ by the type el and $(u, \ldots, u)$, and a vertex $\bar{t} \in U^s$ by the type (s, S), where $S := \{R \in \tau_s \mid \bar{t} \in R^\mathfrak{D}\}$, and the m-tuple obtained from $\bar{t}$ by repeating its last entry. Distinct vertices of the same type have distinct representing tuples. For every type π, there is a quantifier-free formula $\chi_\pi(\bar{y})$ over τ that holds exactly for the representing tuples of type π, in both structures. Translate a formula with at most K variables by assigning to each variable y a block $\bar{y}$ of m variables and by a case distinction over the finitely many types of its free variables. Atoms translate into quantifier-free formulas: $A(y)$ and $B(y)$ become true or false depending on the type of y; $y = y'$ becomes $\bar{y} = \bar{y}'$ if y and y' have the same type and false otherwise; and $E_{(s,J,S)}(x, y)$ becomes $\bigwedge_{j\in J} x_1 = y_j \wedge \bigwedge_{j\in[s]\setminus J} x_1 \neq y_j$ if x has type el and y has type (s, S), and false otherwise. A quantifier $\exists^{\geq\nu} y$ becomes the disjunction over all $(\nu_\pi)_\pi$ with $\sum_\pi \nu_\pi = \nu$ of the conjunction over all types π of $\exists^{\geq\nu_\pi} \bar{y}\, (\chi_\pi(\bar{y}) \wedge \ldots)$, where $\ldots$ is the translation for type π of y. Counting quantifiers over m-tuples are expressible in counting logic with the m quantified variables and quantifier rank m. The translation therefore has at most

mK variables and quantifier rank at most mq, and it holds in $\mathfrak{D}$ if and only if the original sentence holds in $\mathrm{Inc}(\mathfrak{D})$. □

We transfer pairs of structures to pairs of matrices.

Lemma C.9. *Let $m, K_0, q_0 \geq 1$ and $n, N \geq 1$ with $N \geq (n+1)^m$. Let $\mathfrak{A}$ and $\mathfrak{B}$ be relational structures with n elements over a finite vocabulary of arity $\leq m$ that are not $\mathsf{C}_{K_0}^{q_0}$-equivalent. Then there are matrices $M, M' \in \mathbb{K}^{N \times N}$ and a bipartite multigraph F with $\mathrm{fw}_{q_0+1}(F) \leq K_0$ such that*

1. *$\hom_{F,N}(M) \neq \hom_{F,N}(M')$, and*
2. *$\hom_{F',N}(M) = \hom_{F',N}(M')$ for every bipartite multigraph F' with $\mathrm{fw}_\Delta(F') \leq K-1$ for some $K, \Delta \geq 1$ such that $\mathfrak{A}$ and $\mathfrak{B}$ are $\mathsf{C}_{mK}^{mK\Delta}$-equivalent.*

Proof. By Lemma C.8 (contrapositively), $\mathrm{Inc}(\mathfrak{A})$ and $\mathrm{Inc}(\mathfrak{B})$ are not $\mathsf{C}_{K_0+1}^{q_0+1}$-equivalent. By Theorem C.5, some Σ-labelled bipartite graph $L \in \mathcal{T}_{q_0+1}^{K_0+1}$ satisfies $\hom(L, \mathrm{Inc}(\mathfrak{A})) \neq \hom(L, \mathrm{Inc}(\mathfrak{B}))$. Let F and ω be as in Lemma C.6 for L and N, and let $M := M^\omega_{\mathrm{Inc}(\mathfrak{A})}$ and $M' := M^\omega_{\mathrm{Inc}(\mathfrak{B})}$. Then item 1 holds. Since $\mathcal{T}_{q_0+1}^{K_0+1}$ is closed under subgraphs, $F \in \mathcal{T}_{q_0+1}^{K_0+1}$, so $\mathrm{fw}_{q_0+1}(F) \leq K_0$ by Lemma C.4.

For item 2, let F' be a bipartite multigraph and $K, \Delta \geq 1$ such that $\mathrm{fw}_\Delta(F') \leq K-1$ and $\mathfrak{A}$ and $\mathfrak{B}$ are $\mathsf{C}_{mK}^{mK\Delta}$-equivalent. Then $F' \in \mathcal{T}_{\Delta K}^{K}$ by Lemma C.4. By Lemma C.8, $\mathrm{Inc}(\mathfrak{A})$ and $\mathrm{Inc}(\mathfrak{B})$ are $\mathsf{C}_{K}^{K\Delta}$-equivalent. By Lemma C.7, $\hom_{F',N}(M) = \hom_{F',N}(M')$. □

C.3.3. The separation

By [CFI92, Theorem 5.2], two d-tuples receive the same colour after r rounds of d-WL if, and only if, they satisfy the same C_{d+1}^r-formulas. Hence, if d-WL does not distinguish two structures after r rounds, then they are C_{d+1}^r-equivalent. If k-WL distinguishes two structures, then they are not C_{k+1}-equivalent, since counting the tuples of a given colour requires only k further quantifiers. Thus, the structures of Theorem 5.6 are not C_{k+1}-equivalent but $\mathsf{C}_{d+1}^{\lfloor n^{\varepsilon k} \rfloor}$-equivalent. We show the following stronger version of Theorem 5.5.

Theorem C.10. *There are $\eta, \theta > 0$, an infinite set $\mathcal{N} \subseteq \mathbb{N}$, and a family $(p_N) \in \mathsf{symVP}$ such that the following holds. For every sufficiently large $N \in \mathcal{N}$, every matrix-symmetric circuit of orbit size $\leq 2^{N^\eta}$ computing p_N has support depth $\geq N^\theta$. In particular, $(p_N) \notin \mathsf{symVAC}$.*

Proof. Let $k := k_0$, α, and ε be as in Theorem 5.6, and let $\eta := \min\{\varepsilon, 1/(8k)\}/3$ and $\theta := \varepsilon/3$. For $n \geq \alpha k^{14}$, let $d(n) := \lfloor (n/(\alpha k^6))^{1/8} \rfloor \geq k$, let $\mathfrak{A}_n$ and $\mathfrak{B}_n$ be as in Theorem 5.6 for $d(n)$, k, and n, and let $N(n) := (n+1)^k$. Let $\mathcal{N} := \{N(n) \mid n \geq \alpha k^{14}\}$. For $N = N(n) \in \mathcal{N}$, let M_N, M'_N, and F_N be as in Lemma C.9 for $m := k$, $K_0 := k+1$, and some q_0, and let $p_N := \hom_{F_N,N}$. For $N \notin \mathcal{N}$, let $p_N := 0$.

Upper bound. By Lemma C.9, $\mathrm{fw}_{q_0+1}(F_N) \leq k+1$. By Theorem B.2, p_N is computed by matrix-symmetric circuits of orbit size $\leq N^{k+2}$ for $N \geq 2k+5$. Hence, $(p_N) \in \mathsf{symVP}$.

Lower bound. Let $N = N(n) \in \mathcal{N}$ be sufficiently large, and let C be a matrix-symmetric circuit of orbit size $\leq f := 2^{N^\eta}$ and support depth $\leq \Delta$ computing p_N, where $\Delta \geq 1$. By Lemma 2.4, we may suppose that C is reduced. Let $\kappa := \kappa_f(N)$. As $f < \binom{N}{\lfloor N/2 \rfloor}$, we have $\kappa \leq N/2$, so $2^{\kappa-1} \leq \binom{N}{\kappa-1} \leq f$ by minimality of κ. Hence $\kappa - 1 \leq N^\eta$, and in particular $3\kappa \leq N+2$. Let $K := \max\{\kappa-1, 2\} \leq N^\eta$. By Theorems 3.5 and 2.6, p_N is a linear combination of $\hom_{F,N}$ for bipartite multigraphs F with $\mathrm{fw}_\Delta(F) \leq K-1$. Since $n+1 = N^{1/k}$ and $\eta < 1/(8k)$, we have $kK \leq kN^\eta \leq d(n)$. If $kK\Delta \leq \lfloor n^{\varepsilon k} \rfloor$, then $\mathfrak{A}_n$ and $\mathfrak{B}_n$ are $\mathsf{C}_{kK}^{kK\Delta}$-equivalent, and Lemma C.9 yields $p_N(M_N) = p_N(M'_N)$, contradicting

$\hom_{F_N,N}(M_N) \neq \hom_{F_N,N}(M'_N)$. Hence, $\Delta > \lfloor n^{\varepsilon k} \rfloor/(kK) \geq (N^{1/k}-1)^{\varepsilon k}/(2kN^{\eta}) \geq N^{\theta}$, as $\eta + \theta < \varepsilon$. Applying this to $\Delta := \max\{\Delta', 1\}$, where Δ' is the support depth of C, yields $\Delta' \geq N^{\theta}$. Since every family in symVAC is computed by circuits of polynomial orbit size and polylogarithmic support depth, $(p_N) \notin \mathsf{symVAC}$. □

Theorem 5.5 follows from Theorem C.10.

For Theorem 5.7, we use a version of Theorem 5.6 with a tunable number of rounds, at the price of constantly many variables.

Lemma C.11 (cf. [GLN25, Lemmas 5.1, 5.2, and 5.15]). *There are constants $\ell_0 \geq 3$ and $\xi \geq 1$ such that for all $h \geq \ell_0$ and $r \geq 1$ there are relational structures $\mathfrak{A}$ and $\mathfrak{B}$ with the same number $n \leq 2\xi h^2 r$ of elements over a finite vocabulary of arity $\leq \ell_0$ that are C^r_h-equivalent but not $\mathsf{C}^{\xi hr}_{\ell_0}$-equivalent.*

Proof. The round bound is not stated in [GLN25], so we inspect the construction in the proof of [GLN25, Lemma 5.2]. Let $R_0 \geq 9$ be the constant of [GLN25, Corollary 5.12], let $\ell_0 := R_0 + 1$, and let δ be the constant of [GLN25, Lemma 5.2]. Let $h \geq \ell_0$ and $r \geq 1$, and let $\Lambda := 2(h+1)(r+1)$. For $\ell_{\mathrm{hi}} := h$, the proof of [GLN25, Lemma 5.2] constructs a layered graph $Z = (V, W, E)$ with layers $V_0, \dots, V_\Lambda$ and $W_1, \dots, W_\Lambda$ such that, for all $i \in [\Lambda]$, the graph $Z[V_i \cup W_i]$ is a perfect matching, every $w \in W_i$ has ℓ_0 neighbours, and these lie in $V_{i-1} \cup V_i$. It considers the set of XOR-constraints

$$\mathcal{C} := \{(N_Z(w), 0) \mid w \in W\} \cup \{(\{v\}, 0) \mid v \in V_0\} \cup \{(\{v_\Lambda\}, 1)\}$$

of arity $\leq \ell_0$ for some $v_\Lambda \in V_\Lambda$, and shows that $|V| \leq \delta h^2 r$ and that Falsifier does not win the r-round h-pebble game on $\mathcal{C}$. Let $\mathfrak{A} := \mathfrak{A}(\mathcal{C})$ and $\mathfrak{B} := \mathfrak{B}(\mathcal{C})$ be the structures of [GLN25, Section 5.1]. They have $2|V| \leq 2\delta h^2 r$ elements and arity $\leq \ell_0$, and they are C^r_h-equivalent by [GLN25, Lemma 5.1].

We count the rounds of the strategy of Falsifier in the ℓ_0-pebble game from the proof of [GLN25, Lemma 5.15]. In the first round, Falsifier pebbles v_Λ. If Verifier answers 0, then $(\{v_\Lambda\}, 1)$ is violated. Otherwise, suppose that the current position assigns 1 to some $v_i \in V_i$ with $i \geq 1$, and let $w_i \in W_i$ be the matching partner of v_i. Then $N_Z(w_i) \cap V_i = \{v_i\}$. Falsifier keeps the pebbles on the current pebbled vertices in $N_Z(w_i)$, including v_i, and pebbles the remaining vertices of $N_Z(w_i)$ one per round. This takes at most $\ell_0 - 1$ rounds and at most ℓ_0 pebbles. Unless $(N_Z(w_i), 0)$ is violated, some $v_{i-1} \in N_Z(w_i) \cap V_{i-1}$ is assigned 1. Once some $v_0 \in V_0$ is assigned 1, the constraint $(\{v_0\}, 0)$ is violated. Hence, Falsifier wins within $1 + \Lambda(\ell_0 - 1) \leq 1 + 8(\ell_0 - 1)hr \leq 8\ell_0 hr$ rounds, using $(h+1)(r+1) \leq 4hr$. By [GLN25, Lemma 5.1], $\mathfrak{A}$ and $\mathfrak{B}$ are not $\mathsf{C}^{8\ell_0 hr}_{\ell_0}$-equivalent. The claim follows with $\xi := \max\{\delta, 8\ell_0\}$. □

Proof of Theorem 5.7. Let ℓ_0 and ξ be as in Lemma C.11, let $k \geq 1$, and write $\log = \log_2$. For $N \geq 2$, let $h(N) := \lfloor \sqrt{\log N} \rfloor$ and $r(N) := h(N)\lceil \log^k N \rceil$. Let $N_0 \geq 2\ell_0 + 3$ be such that $h(N) \geq \ell_0$ and $(2\xi h(N)^2 r(N) + 1)^{\ell_0} \leq N$ for all $N \geq N_0$; it exists since the left-hand side is polylogarithmic in N. For $N \geq N_0$, write $h := h(N)$ and $r := r(N)$, let $\mathfrak{A}_N$ and $\mathfrak{B}_N$ be as in Lemma C.11 for h and r, and let M_N, M'_N, and F_N be as in Lemma C.9 for $m := \ell_0$, $K_0 := \ell_0$, and $q_0 := \xi hr$.

The family. Let $p_N := \hom_{F_N,N}$ for $N \geq N_0$ and $p_N := 0$ otherwise. By Lemma C.9, $\mathrm{fw}_{\xi hr+1}(F_N) \leq \ell_0$. By Theorem B.2, p_N is computed by a matrix-symmetric circuit of orbit size $\leq N^{\ell_0+1}$ and support depth $\leq \xi hr + 1 = \xi h^2 \lceil \log^k N \rceil + 1 \leq 3\xi \log^{k+1} N$. Hence, $(p_N) \in \mathsf{symVAC}^{k+1}$.

The lower bound. Suppose that $(p_N) \in \mathsf{symVAC}^k$, witnessed by a constant $c \geq 1$. Let N be sufficiently large, and let C be a matrix-symmetric circuit of orbit size $\leq f := N^c$ and support depth $\leq \Delta := \lceil c \log^k N \rceil$ computing p_N. By Lemma 2.4, we may suppose that C is reduced. Let $\kappa := \kappa_f(N)$. Since $\binom{N}{\lfloor c \rfloor + 1} > N^c$, we have $\kappa \leq c + 1$ and hence $3\kappa \leq N + 2$.

Let $K := \max\{\kappa - 1, 2\} \leq c + 1$. By Theorems 3.5 and 2.6, p_N is a linear combination of $\hom_{F,N}$ for bipartite multigraphs F with $\mathrm{fw}_\Delta(F) \leq K - 1$. Since $h \to \infty$, we have $\ell_0 K \leq h$ and $\ell_0 K\Delta \leq \ell_0(c+1)^2\lceil \log^k N\rceil \leq r$. Hence, $\mathfrak{A}_N$ and $\mathfrak{B}_N$ are $\mathsf{C}^{\ell_0 K\Delta}_{\ell_0 K}$-equivalent, and Lemma C.9 yields $p_N(M_N) = p_N(M'_N)$, contradicting $\hom_{F_N,N}(M_N) \neq \hom_{F_N,N}(M'_N)$. Hence, $(p_N) \notin \mathsf{symVAC}^k$. □

Remark C.12. Without inspecting the proofs of [GLN25], a weaker hierarchy follows from Theorem 5.6 and [GLN25, Theorem 1.1]: with k_0 and ε as in Theorem 5.6, $\mathsf{symVAC}^j \subsetneq \mathsf{symVAC}^{\lfloor j/\varepsilon\rfloor + k_0}$ for every $j \geq 1$. In particular, the hierarchy $(\mathsf{symVAC}^j)_j$ does not collapse.

Indeed, let $k := k_0$ and $g := \lfloor j/(\varepsilon k)\rfloor + 1$, so that $g\varepsilon k > j$ and $gk \leq j/\varepsilon + k$. For large N, let $n := \lfloor \log^g N\rfloor$ and $d := \lfloor (n/(\alpha k^6))^{1/8}\rfloor$, and let $\mathfrak{A}_N$ and $\mathfrak{B}_N$ be as in Theorem 5.6 for d, k, and n. By [GLN25, Theorem 1.1], k-WL stabilises on $\mathfrak{A}_N$ and $\mathfrak{B}_N$ after $O(kn^{k-1}\log n)$ rounds. Hence, k-WL distinguishes them after this many rounds, and they are not $\mathsf{C}^{q_0}_{k+1}$-equivalent for $q_0 := n^k$ by [GLN25, Corollary 2.2], for large N. Let M_N, M'_N, and F_N be as in Lemma C.9 for $m := k$, $K_0 := k + 1$, and q_0, and let $p_N := \hom_{F_N,N}$. As in the proof of Theorem 5.7, Theorem B.2 yields circuits of orbit size $\leq N^{k+2}$ and support depth $\leq n^k + 1 = O(\log^{gk} N)$, so $(p_N) \in \mathsf{symVAC}^{gk}$. Conversely, for circuits of orbit size $\leq N^c$ and support depth $\leq \lceil c\log^j N\rceil$, the argument of the lower bound in the proof of Theorem 5.7 applies, since $\mathfrak{A}_N$ and $\mathfrak{B}_N$ are $\mathsf{C}^{\lfloor n^{\varepsilon k}\rfloor}_{d+1}$-equivalent, $d \to \infty$, and $n^{\varepsilon k} = \Omega(\log^{g\varepsilon k} N)$ with $g\varepsilon k > j$. Hence, $(p_N) \notin \mathsf{symVAC}^j$.

The gap between j and $\lfloor j/\varepsilon\rfloor + k_0$ arises since the upper bound of [GLN25, Theorem 1.1] on the number of rounds is polynomially larger than the lower bound $n^{\varepsilon k}$. Lemma C.11 closes this gap.