Kronecker products and iterated matrix multiplication

Ahmed Ittihad Hasib, Christian Ikenmeyer, Dimitrios Tsintsilidas

September 15, 2026

Abstract

Geometric complexity theory is an approach initiated by Mulmuley and Sohoni (SIAM J. Comp. 2001, 2008) towards unconditional algebraic complexity lower bounds by comparing two orbit closures, the determinant orbit closure and the padded permanent orbit closure. Some of its central conjectures have been disproved, starting in 2009 (Briand–Orellana–Rosas, comput. compl.). We attempt to set geometric complexity theory on a more robust footing, so in this paper we propose a change of the two orbit closures. Our proposed change is the result of a careful search for an algebraic model of computation that incorporates desired algebro-geometric and invariant-theoretic properties, namely a complexity-theoretically complete point that is characterized by its reductive stabilizer and whose endomorphism orbit over the tensor algebra is Zariski-closed (this combination of properties had not been achieved so far), while the second orbit closure should also have as many of these properties as possible. We find that iterated matrix multiplication, interpreted as a matrix of homogeneous polynomials, satisfies all such properties, and it is the only such model that we know of. It comes with the additional benefit that it is definable over arbitrary commutative semirings, and not just for polynomials, but also for the tensor algebra and the exterior algebra. This gives us the first new orbit closure for geometric complexity theory, replacing the orbit closure of the determinant.

For the second orbit closure, we search for a hard counterpart for iterated matrix multiplication, which should preserve as many of its desirable properties as possible. We start by observing for the first time that the Kronecker product of tensors is the operation that converts the determinant polynomial into Cayley's first hyperdeterminant, and that the Kronecker product appears in the construction of several VNP-complete polynomials that originate in algebra. The key property of the Kronecker product is its equivariance. Using the equivariance we obtain a one-line proof for two results of Gurvits (STOC 2003, MFCS 2005) that the mixed discriminant and Cayley's first hyperdeterminant are as hard as the permanent, which gives VNP-hardness in characteristic different from 2. We raise iterated matrix multiplication (which we call the computant) to its Kronecker square, which results in the hypercomputant, whose entries are VNP-complete and VW[1]-complete polynomials, where we prove the hardness in a gadget-free way via the equivariance of the Kronecker product. In contrast, in the literature VNP-hardness is usually proved via gadget constructions. Using gadgets, we sharpen our construction and then use the equivariance to give short proofs of the VNP-completeness of the mixed discriminant and Cayley's first hyperdeterminant over all commutative rings, strengthening the two aforementioned results of Gurvits, and in the same way we obtain the VNP-completeness over all commutative semirings for the hyperpermanent and hyperhafnian of order at least 3, and the VNP-completeness of the hyperpfaffian of even order at least 6 over all commutative rings. For the tensor algebra this gives a version of "noncommutative VNP", and for polynomials over the nonnegative real numbers this gives a version of "monotone VNP", each with the hypercomputant as the complete object. Inspired by Strassen's study of the matrix multiplication exponent, we take a parameterized complexity viewpoint, parameterized by the degree as in Bläser–Engels (IPEC 2019), and using standard techniques we determine exactly the complexity exponents in the noncommutative setting and in the monotone setting, and compare both settings. We prove that the exponents in the monotone setting cannot exceed the exponents in the noncommutative setting, and notably, the exponents are not always the same.

We see these results as evidence that the hypercomputant is a good counterpart to the computant, but the hypercomputant is not characterized by its stabilizer. Instead, it is a sum of $d!$ many isotypic parts, and we prove that each of them is characterized by its stabilizer. They govern the parameterized complexity growth exponents, i.e., the complexity-theoretic decomposition matches the representation-theoretic decomposition. These parts serve as the second orbit closures for our proposed adjustment of geometric complexity theory.

Towards implementing a full geometric complexity theory approach, we determine the stabilizers of the computant and the hypercomputant, we determine the representation-theoretic multiplicities in the ambient space, and determine the representation-theoretic multiplicities of the coordinate rings of the orbit of the computant. We also prove the polystability of the computant, the hypercomputant, and its isotypic parts. The stabilizers of the isotypic parts seem to have an intriguing structure related to the combinatorial structure of the monotone situation, and will be the focus of future work.

1 Introduction

Determinant versus permanent. The search for unconditional computational complexity lower bounds is among the hardest challenges in computational complexity theory. Reasoning about Turing machines is difficult and in order to make the search for lower bounds more feasible, researchers proposed to focus on proving circuit lower bounds instead, see [AB09, Ch. 14]. For problems from algebra it is reasonable to first focus on the required size of an arithmetic circuit instead of a Boolean circuit, see for example Strassen's breakthrough subcubic matrix multiplication algorithm [Str69].

An arithmetic circuit C is a directed acyclic graph whose indegree 0 vertices (called *inputs*) are labeled with affine linear polynomials, and whose other vertices (called *computation gates*) have indegree 2 and are each labeled with an addition or multiplication operation $+, \times$. An arithmetic circuit computes a polynomial at every vertex by induction over the circuit structure. The size of C is defined as the number of computation gates of C. For a matrix F of polynomials we define $\mathsf{cc}(F)$ to be the size of a smallest arithmetic circuit computing each matrix entry of F. Strassen showed that $\mathsf{cc}(X_n^{(1)} \cdot X_n^{(2)}) = O(n^{2.81})$, where $X_n^{(k)} = \big(x_{i,j}^{(k)}\big)_{\substack{1\leq i\leq n\\ 1\leq j\leq n}}$ is an $n \times n$ matrix of variables, and there has been a long list of papers improving the best known exponent.

Valiant [Val79] defined an algebraic analogue of NP (or more precisely, #P), which is now called VNP, and he gave several equivalent definitions [Val82]. In this paper, a sequence (f) of polynomials is defined to be in VNP if there exist Boolean single-output circuits B_n on $\mathsf{poly}(n)$ many inputs such that $f_n = \sum_{B_n(\mathbf{b})=1} x^{\mathbf{b}}$, or if (f) is a so-called p-projection of such a sequence (h), i.e., $(f) \leq_{\mathrm{p}} (h)$, where we write $(f) \leq_{\mathrm{p}} (h)$ if there exists $q(n) \in \mathsf{poly}(n)$ and affine linear polynomials ℓ_i such that $\forall n : f_n = h_{q(n)}(\ell_1, \ell_2, \ldots)$. Clearly, the permanent polynomial $\mathrm{per}_n = \sum_{\pi\in\mathfrak{S}_n} \prod_{i=1}^n x_{i,\pi(i)}$ and the Hamiltonian cycle polynomial $\mathrm{HC}_n = \sum_{\substack{\pi\in\mathfrak{S}_n\\ \pi \text{ cycle of length } n}} \prod_{i=1}^n x_{i,\pi(i)}$ are contained in VNP. The definition generalizes over any commutative semiring, but we note that some equivalent definitions generalize differently, for example there is a generalization for which the permanent over the nonnegative real numbers is not contained in VNP, see [Yeh19]. A sequence of polynomials (h) is VNP-complete if $(f) \leq_{\mathrm{p}} (h)$ for all $(f) \in$ VNP.

Based on Hyafil's parallelization result [Hya79], Valiant [Val79] realized that the computation of the determinant polynomial $\det_d = \sum_{\pi\in\mathfrak{S}_d} \mathrm{sgn}(\pi) \prod_{i=1}^d x_{i,\pi(i)}$ is complete (under so-called quasipolynomial projections) for all problems of reasonable degree with polynomially bounded arithmetic circuit complexity. More precisely, we have the following three major algebraic complexity classes of efficient computation.

- A sequence of polynomials (f) is defined to be in VP if $\deg(f_n)$ and the number of variables of f_n and $\mathsf{cc}(f_n)$ are all in $\mathsf{poly}(n) = \bigcup_{d\in\mathbb{N}} O(n^d)$.
- A circuit whose graph is a tree is called an arithmetic formula, and the size of the smallest arithmetic formula computing f is denoted by $\mathsf{fc}(f)$. A sequence of polynomials (f) is defined to be in VF if the number of variables of f_n and $\mathsf{fc}(f_n)$ are all in $\mathsf{poly}(n)$. Note that $\mathsf{fc}(f_n) \in \mathsf{poly}(n)$ implies $\deg(f_n) \in \mathsf{poly}(n)$.
- The determinantal complexity $\mathsf{dc}(f)$ is the smallest r such that f can be written as the determinant of an $r \times r$ matrix whose entries are affine linear polynomials. This quantity is always finite, see [Val79]. The class of sequences (f) of polynomials for which the number of variables and $\mathsf{dc}(f_n)$ are in $\mathsf{poly}(n)$ goes under many names, for example VBP. Note that $\mathsf{dc}(f_n) \in \mathsf{poly}(n)$ implies $\deg(f_n) \in \mathsf{poly}(n)$.

We have $\mathsf{VF} \subseteq$ VBP $\subseteq \mathsf{VP}$, and both inclusions are nontrivial, see [Val79] and see e.g. [MV97] and references therein. If one allows quasipolynomial (i.e., $n^{\mathsf{poly}(\log n)}$) formula size, determinantal complexity, and circuit size, then all three definitions coincide, which follows from Hyafil's result. Valiant used this to pointedly contrast the determinant with the permanent, for which he showed VNP-completeness, provided the characteristic of the field is not 2, because in characteristic 2 we have $\det_n = \mathrm{per}_n$. This issue with the characteristic of the field may seem irrelevant, but it crucially prevents all field-independent lower-bound techniques to work. Modern mathematics is striving for field-independent proofs whenever possible, see e.g. [Eis13, AFP$^+$19, DV22], to name a few. Valiant also showed that HC_n is VNP-complete over arbitrary commutative semirings, which however spawned significantly less follow-up work, maybe partly because the definition of HC_n looks more ad-hoc combinatorial and less algebraic in nature. In fact, the question whether VBP $\neq$ VNP is widely known as the determinant versus permanent question. Over fields of characteristic different from 2 it is equivalent to $\mathsf{dc}(\mathrm{per}_n) \notin \mathsf{poly}(n)$. Notably, the notion of a circuit does not appear in the determinant versus permanent question, and both polynomials can be studied with tools from classical algebra.

Geometric complexity theory. Motivated by the geometric proof of lower bounds in [Mul99], in [MS01, MS08] Mulmuley and Sohoni rephrase the determinant versus permanent question as an orbit closure question. The translation is almost literal. The determinantal orbit closure complexity $\mathsf{docc}(f)$ of a polynomial f is the smallest n such that $z^n f(x_1/z, x_2/z, \ldots)$ can be written as the determinant of an $n \times n$ matrix of homogeneous linear polynomials, or a limit of such polynomials. Setting $z = 1$ we see that $\mathsf{docc}(f) \le \mathsf{dc}(f)$. Over $\mathbb{C}$ the orbit closure notation is often used, i.e., $z^n f(x_1/z, x_2/z, \ldots) \in \overline{\mathrm{GL}_{n^2}\mathrm{det}_n}$ (Zariski closure and Euclidean closure coincide, see e.g. [Kra84, AI.7.2 Folgerung]), which means the same. Mulmuley and Sohoni conjectured $\mathsf{docc}(\mathrm{per}_m) \notin \mathsf{poly}(m)$, and they realized that the lower bound $\mathsf{docc}(\mathrm{per}_m) > n$ is equivalent to the non-containment of the two orbit closures $\overline{\mathrm{GL}_{n^2} z^{n-m}\mathrm{per}_m} \not\subseteq \overline{\mathrm{GL}_{n^2}\mathrm{det}_n}$.

Mulmuley–Sohoni's idea is to prove such a non-containment by using the symmetries (i.e., the stabilizers) $\mathrm{stab}(\mathrm{det}_n)$ of det_n and $\mathrm{stab}(\mathrm{per}_m)$ of per_m. The stabilizer $\mathrm{stab}(\mathrm{det}_n)$ is the set of linear transformations (i.e., sending variables to homogeneous linear combinations of variables) that fix det_n, and analogously for per_m. The stabilizer $\mathrm{stab}(\mathrm{det}_n)$ has been determined by Frobenius in 1897 [Fro97, §7, Satz 1], $\mathrm{stab}(\mathrm{per}_m)$ in [MM62] by Marcus and May. In fact, Mulmuley–Sohoni point out that det_n and per_m are *characterized by their stabilizers*, i.e., every polynomial that is fixed by all linear transformations in $\mathrm{stab}(\mathrm{det}_n)$ must be a scalar multiple of det_n, and analogously for per_m. Numerous objects in algebraic complexity theory are characterized by their stabilizers, for example the power sum polynomial $x_1^d + \cdots + x_n^d$, the squarefree monomial $\mathrm{mon}_d := x_1 \cdots x_d$, the matrix multiplication tensor and the unit tensor (see [BI11]).

We explain now how the stabilizers are used in the geometric complexity theory approach, and why it is desired to work with points that are characterized by their stabilizer, and that "padding" polynomials by multiplying them by a power of a homogeneous linear polynomial should be avoided. This paragraph is more technically involved, and we refer the reader to [BLMW11] and the references therein. This paragraph is motivational, and we will not need the technical details in this paper. Let $\mathbb{C}[\mathbb{C}[x_{1,1}, \ldots, x_{n,n}]_n]$ denote the vector space of multivariate polynomials in the coefficients of homogeneous degree n polynomials in n^2 variables (as in [vdBDG$^+$25]). Let $I(\overline{\mathrm{GL}_{n^2}\mathrm{det}_n}) \subseteq \mathbb{C}[\mathbb{C}[x_{1,1}, \ldots, x_{n,n}]_n]$ denote the ideal of those polynomials that vanish identically on $\overline{\mathrm{GL}_{n^2}\mathrm{det}_n}$. Let $\mathbb{C}[\overline{\mathrm{GL}_{n^2}\mathrm{det}_n}] := \mathbb{C}[\mathbb{C}[x_{1,1}, \ldots, x_{n,n}]_n]/I(\overline{\mathrm{GL}_{n^2}\mathrm{det}_n})$ denote the coordinate ring. Analogously for $\overline{\mathrm{GL}_{n^2}(z^{n-m}\mathrm{per}_m)}$. If for the sake of contradiction we have $\overline{\mathrm{GL}_{n^2}(z^{n-m}\mathrm{per}_m)} \subseteq \overline{\mathrm{GL}_{n^2}\mathrm{det}_n}$, then this induces a surjection $\mathbb{C}[\overline{\mathrm{GL}_{n^2}\mathrm{det}_n}] \twoheadrightarrow \mathbb{C}[\overline{\mathrm{GL}_{n^2}(z^{n-m}\mathrm{per}_m)}]$ that is just given by the restriction of the domain of definition. Both sides of the surjection carry a linear action of the general linear group GL_{n^2}, and the surjection is equivariant. This implies by Schur's lemma that for every degree δ the multiplicity of every irreducible GL_{n^2}-representation in $\mathbb{C}[\overline{\mathrm{GL}_{n^2}\mathrm{det}_n}]_\delta$ cannot be smaller than the multiplicity of the same irreducible GL_{n^2}-representation in $\mathbb{C}[\overline{\mathrm{GL}_{n^2}(z^{n-m}\mathrm{per}_m)}]_\delta$. The ring $\mathbb{C}[\mathrm{GL}_{n^2}\mathrm{det}_n]$ is a localization of $\mathbb{C}[\overline{\mathrm{GL}_{n^2}\mathrm{det}_n}]$ at a so-called fundamental invariant [BI17]. In particular, we have an equivariant injection $\mathbb{C}[\overline{\mathrm{GL}_{n^2}\mathrm{det}_n}] \hookrightarrow \mathbb{C}[\mathrm{GL}_{n^2}\mathrm{det}_n]$, hence the representation-theoretic multiplicities of $\mathbb{C}[\mathrm{GL}_{n^2}\mathrm{det}_n]$ are an upper bound for those of $\mathbb{C}[\overline{\mathrm{GL}_{n^2}\mathrm{det}_n}]$. A deep understanding of the specific fundamental invariant also lets one go in the other direction, as demonstrated in [IK20], see also [DIP20]. The representation-theoretic structure of the ring $\mathbb{C}[\mathrm{GL}_{n^2}\mathrm{det}_n] \simeq \mathbb{C}[\mathrm{GL}_{n^2}\mathrm{det}_n]^{\mathrm{stab}(\mathrm{det}_n)}$ can be determined via classical branching rules and character theory. Whether this representation-theoretic approach is feasible is an open question, but in many situations when the points are characterized by their stabilizer no information is lost when going from orbits to the representation-theoretic multiplicities of their coordinate rings, see [LP90, Yu16]. Mulmuley and Sohoni conjectured that already the occurrence pattern of the representations in $\mathbb{C}[\overline{\mathrm{GL}_{n^2}\det_n}]$ and $\mathbb{C}[\overline{\mathrm{GL}_{n^2}(z^{n-m})\mathrm{per}_m}]$ would separate VBP and VNP. A first analysis of the representation theory of $\mathbb{C}[\overline{\mathrm{GL}_{n^2}(z^{n-m}\mathrm{per}_m)}]$ was done by Kadish and Landsberg [KL14] and the representation theory shows strange behavior due to the fact that per_m is multiplied ("padded") by a power of z. In [IP17, BIP19] it was finally proved that comparing the occurrence patterns in $\mathbb{C}[\overline{\mathrm{GL}_{n^2}(z^{n-m}\mathrm{per}_m)}]$ and $\mathbb{C}[\overline{\mathrm{GL}_{n^2}(\det_n)}]$ is not sufficient to separate VBP and VNP. The padding is a crucial element of this no-go proof.

1.I Fixing the padding issue: Width

One can try to overcome this issue by replacing the determinant with a polynomial that is double-indexed by a complexity parameter n and a degree d, the main objects of study in [BE19, DGI$^+$24]. Such families are well-known in algebraic geometry, for example the symmetric border rank (also called border Waring rank) of a homogeneous degree d polynomial f is the smallest n such that $f \in \overline{\mathrm{GL}_n(x_1^d + \cdots + x_n^d)}$. Or the border Chow rank of f is the smallest n such that $f \in \overline{\mathrm{GL}_{nd}(x_1^{(1)} \cdots x_1^{(d)} + \cdots + x_n^{(1)} \cdots x_n^{(d)})}$. These notions capture the computational complexity of f well for $d = 3$ (see [CHI$^+$18]), but there are two main ways of converting these notions into a notion that captures the computational complexity for general d. To see these two variants of measuring complexity, we rewrite Waring rank and Chow rank as follows. The Waring rank of a homogeneous degree d polynomial f is the smallest n such that f can be written as

$$\begin{aligned} f &= \begin{pmatrix}\ell_1 & \ell_2 & \dots & \ell_n\end{pmatrix} \cdot \mathrm{diag}(\ell_1, \ell_2, \dots, \ell_n) \cdots \mathrm{diag}(\ell_1, \ell_2, \dots, \ell_n) \cdot \begin{pmatrix}\ell_1 & \ell_2 & \dots & \ell_n\end{pmatrix}^\mathsf{T} \\ &= \mathrm{tr}\left(\mathrm{diag}(\ell_1, \ell_2, \dots, \ell_n) \cdots \mathrm{diag}(\ell_1, \ell_2, \dots, \ell_n)\right) \end{aligned}$$

for homogeneous linear polynomials ℓ_i. The Chow rank of f is the smallest n such that f can be written as

$$\begin{aligned} f &= \begin{pmatrix}\ell_1^{(1)} & \ell_2^{(1)} & \dots & \ell_n^{(1)}\end{pmatrix} \cdot \mathrm{diag}(\ell_1^{(2)}, \ell_2^{(2)}, \dots, \ell_n^{(2)}) \cdots \mathrm{diag}(\ell_1^{(d-1)}, \dots, \ell_n^{(d-1)}) \cdot \begin{pmatrix}\ell_1^{(d)} & \ell_2^{(d)} & \dots & \ell_n^{(d)}\end{pmatrix}^\mathsf{T} \\ &= \mathrm{tr}\left(\mathrm{diag}(\ell_1^{(1)}, \ell_2^{(1)}, \dots, \ell_n^{(1)}) \cdots \mathrm{diag}(\ell_1^{(d)}, \ell_2^{(d)}, \dots, \ell_n^{(d)})\right) \end{aligned}$$

for homogeneous linear polynomials $\ell_i^{(k)}$. If one allows arbitrary matrices instead of diagonal matrices, one obtains two different notions, each with advantages and disadvantages, as follows.

• For a homogeneous degree d polynomial f, the *homogeneous algebraic branching program width* or just *width* $w(f)$ for short is defined as the smallest n such that there exist homogeneous linear polynomials $\ell_{i,j}^{(k)}$ with

$$f = \begin{pmatrix}\ell_{n,1}^{(1)} & \cdots & \ell_{n,n}^{(1)}\end{pmatrix} \begin{pmatrix} \ell_{1,1}^{(2)} & \cdots & \ell_{1,n}^{(2)} \\ \vdots & \ddots & \vdots \\ \ell_{n,1}^{(2)} & \cdots & \ell_{n,n}^{(2)} \end{pmatrix} \begin{pmatrix} \ell_{1,1}^{(3)} & \cdots & \ell_{1,n}^{(3)} \\ \vdots & \ddots & \vdots \\ \ell_{n,1}^{(3)} & \cdots & \ell_{n,n}^{(3)} \end{pmatrix} \cdots \begin{pmatrix} \ell_{1,1}^{(d-1)} & \cdots & \ell_{1,n}^{(d-1)} \\ \vdots & \ddots & \vdots \\ \ell_{n,1}^{(d-1)} & \cdots & \ell_{n,n}^{(d-1)} \end{pmatrix} \begin{pmatrix} \ell_{1,n}^{(d)} \\ \ell_{2,n}^{(d)} \\ \vdots \\ \ell_{n,n}^{(d)} \end{pmatrix}.$$

The width is a lower bound for the border Waring rank, see [BDI21, Thm 4.2], and the width was promoted in [Ike22, DGI$^+$24, Ike25] and was used successfully in [BCL23] for answering Smale's 17th problem. A homogeneous degree d polynomial f has width at most n if and only if f is a restriction of the so-called *iterated matrix multiplication polynomial*

$$\Gamma_{n,d} := (x_{n,1}^{(1)} \ \cdots \ x_{n,n}^{(1)}) \cdot X_n^{(2)} X_n^{(3)} \cdots X_n^{(d-1)} \cdot (x_{1,n}^{(d)} \ \cdots \ x_{n,n}^{(d)})^\mathsf{T},$$

in compact notation $f \le \Gamma_{n,d}$, which means that f arises from $\Gamma_{n,d}$ by replacing variables by homogeneous linear polynomials. At a first glance, this might indicate that the orbit closure $\overline{\mathrm{GL}_{2n+(d-2)n^2}\Gamma_{n,d}}$ should be the focus of a new geometric complexity theory approach, but $\Gamma_{n,d}$ comes with the drawback that its stabilizer is not reductive. In fact, the Lie algebra corresponding to the connected component of the identity in the stabilizer of $\Gamma_{n,d}$ is not reductive, as a short calculation with the computer algebra system GAP shows, see §11.I, confirming a statement by Landsberg. This adds unwanted difficulties to any attempt to realize a geometric complexity theory approach using $\Gamma_{n,d}$, and in fact to any approach using representation theory.

• This failure can be resolved via the second notion as follows. Define $\Xi_{n,d} := X_n^{(1)} \cdots X_n^{(d)}$, observe that $\Gamma_{n,d} = [\Xi_{n,d}]_{n,n}$ is the bottom-right entry, and define the trace tr-$\Xi_{n,d} := \mathrm{tr}(\Xi_{n,d})$. The polynomial tr-$\Xi_{n,d}$ has been studied in [Ges16] and its stabilizer is indeed reductive, and tr-$\Xi_{n,d}$ is characterized by its stabilizer. This would indicate that the orbit closure $\overline{\mathrm{GL}_{n^2d}\text{tr-}\Xi_{n,d}}$ should be the focus of interest in a new geometric complexity theory. However, also in this case a useful property gets lost: If we consider non-commutative polynomials (also called tensors), i.e., polynomials in which the variables do not commute, then Nisan [Nis91] proved that the width is determined as the maximum of d many ranks of so-called flattening matrices. In particular, for every fixed n, the set of width at most n non-commutative polynomials is Zariski-closed (and hence also Euclidean-closed), see [For16]. Such closed models of computation are very rare, but exist, see [TS16, BIL+21] for more examples. In this sense, it is fair to say that we *fully understand* the width "up to the action of the symmetric group $\mathfrak{S}_d$". However, this is not the case when the trace is taken, i.e., restrictions of tr-$\Xi_{n,d}$ are taken instead of restrictions of $\Gamma_{n,d}$, see [BIM+20], which confirms a conjecture of Forbes. Hence, tr-$\Xi_{n,d}$ is not an optimal choice for geometric complexity theory either. Prior to our work, it was unclear how to resolve this dilemma.

2 Our contributions

2.I The best of both worlds

In this paper we propose a third way that incorporates both properties, i.e., reductive stabilizer (and characterization by the stabilizer) and closedness for tensors. To achieve this, instead of focusing on single polynomials as in [Val79], we focus on matrices of polynomials as in [Str69] or the recent [Ike25]. Indeed, $\Xi_{n,d}$ is a matrix of polynomials. For a matrix of polynomials F we define the width $w(F)$ as the smallest n such that $F \leq \Xi_{n,d}$. Here $\leq$ means replacing variables by homogeneous linear combinations of variables, and also allowing multiplications by matrices of constants from the left and from the right, see §5.IV and Definition 6.5. The definition of $\leq$ for matrices of homogeneous polynomials is chosen carefully, and the multiplications from the left and right ensure the additional desired property that $\Xi_{n,d}$ is characterized by its stabilizer, which gets lost if we do not allow these multiplications, or if we only allow conjugation, see §11.I. We determine the stabilizer of $\Xi_{n,d}$ over $\mathbb{C}$ as $\mathrm{GL}_n \times \mathrm{GL}_n^{d-1} \times \mathrm{GL}_n$ (up to scalars), in particular the stabilizer is reductive, and we prove that $\Xi_{n,d}$ is characterized by its stabilizer, see §11.I. Moreover, we prove that Nisan's result on noncommutative polynomials (i.e., tensors) can be lifted to matrices of noncommutative polynomials, see Theorem 8.8. This shows that $\Xi_{n,d}$ combines the desirable properties of both worlds. A first orbit closure for a new geometric complexity theory is hence $\overline{(\mathrm{GL}_n \times \mathrm{GL}_{n^2d} \times \mathrm{GL}_n)\Xi_{n,d}}$. As a further sanity check, we also determine that $\Xi_{n,d}$ is polystable, see Proposition 11.8. Moreover, we determine the multiplicities in the coordinate ring $\mathbb{C}[(\mathrm{GL}_n \times \mathrm{GL}_{n^2d} \times \mathrm{GL}_n)\Xi_{n,d}]$ of the orbit in (11.7), which are sums of products of multi-Littlewood–Richardson coefficients and multi-Kronecker coefficients, both well-known (but not both well-understood) quantities in classical representation theory, see e.g. [KT99] and [Sta00, Problem 10]. Since we are working with matrices instead of just polynomials, we determine the representation-theoretic multiplicities in the ambient space in §11.II.

We remark that $\Gamma_{n,d}$ is an entry of $\Xi_{n,d}$, and while $\Xi_{n,d}$ has the better properties, we use $\Gamma_{n,d}$ to compare our results to the existing literature on polynomials. One could alternatively use tr-$\Xi_{n,d}$, because $\Gamma_{n,d} \leq \text{tr-}\Xi_{n,d} \leq \Gamma_{n,d}^{\oplus n}$.

We analyze some of these properties only over $\mathbb{C}$, as it is in fact not always clear what the correct generalization should be, but $\Xi_{n,d}$ itself is defined over arbitrary commutative semirings, i.e., "rings without subtraction", see §5.II. Dropping the commutativity of the semiring can be done, but it involves several caveats concerning the multilinear algebra, see Remark 5.1. Moreover, $\Xi_{n,d}$ can be defined over the tensor algebra, i.e., non-commutative polynomials, but even over the exterior algebra. It is unusual to see the exterior algebra defined over commutative semirings, but this is exactly the level of generality in mathlib [mC20], which we take as a confirmation of our choice. We aim to keep this level of generality, which restricts the number of options that we have for the second geometric complexity theory orbit closure, i.e., the orbit closure for the hard object. We base this second orbit closure on Kronecker products, as we observe that several VNP-complete polynomials can be written as Kronecker products, and Kronecker products help us to lift several VNP-completeness results over arbitrary commutative (semi-)rings.

2.II Kronecker products

To set up a geometric complexity theory approach, we need a hard matrix of polynomials. We observe for the first time that many VNP-complete polynomials that come from algebra are in fact Kronecker products of polynomials in the following sense. For two multi-index variables $x_{\mathbf{i}}$ and $x_{\mathbf{j}}$ the Kronecker product $x_{\mathbf{i}} \boxtimes x_{\mathbf{j}}$ is defined to be the variable $x_{\mathbf{i},\mathbf{j}}$, i.e., the variable with concatenated multi-index. For two degree d monomials, the Kronecker product is defined as

$$(x_{i_1} \cdots x_{i_d}) \boxtimes (x_{j_1} \cdots x_{j_d}) := \textstyle\sum_{\pi \in \mathfrak{S}_d} (x_{i_{\pi(1)}} \boxtimes x_{j_1}) \cdots (x_{i_{\pi(d)}} \boxtimes x_{j_d}).$$

Via bilinear extension we obtain the definition of the Kronecker product of two homogeneous degree d polynomials. This is well-defined and symmetric (up to renaming the variables), see §7. In characteristic zero, homogeneous degree d polynomials correspond bijectively to symmetric tensors of order d, and the Kronecker product of polynomials corresponds to the well-known Kronecker product of symmetric tensors, see §7.I. One should think of our definition as the canonical way of removing an unnecessary factor of $d!$, which is important when working characteristic-freely. For example, written in an "overly symmetric way" the permanent equals $\frac{1}{d!} \sum_{\pi,\sigma \in \mathfrak{S}_d} \prod_{i=1}^d x_{\pi(i),\sigma(i)}$, which either requires the division by $d!$, or would result in the zero polynomial if the characteristic divides $d!$. We note that this division by $d!$ is what causes the issues when defining the symmetric asymptotic spectrum of symmetric tensors in [CFTZ21], but our definition resolves this issue when working with homogeneous polynomials instead.

In the literature one can find several examples of VNP-complete polynomials that can be written as a Kronecker product, see Example 7.10:

1. The permanent polynomial: $\mathrm{per}_d := \mathrm{mon}_d \boxtimes \mathrm{mon}_d$, see e.g., [Val79], [Min84, eq. (3.1)].
2. The hyperpermanent (also called the multidimensional permanent): $\mathrm{mon}_d^{\boxtimes m}$, see e.g., [DG87, Tic15, CP16].
3. The rectangular permanent: $\mathrm{mon}_d \boxtimes e_{n,d}$, which is just called the permanent in [Min84, eq. (1.1)], see also [BE19, Exa 4.6].
4. The mixed discriminant: $\det_d \boxtimes \mathrm{mon}_d$, see e.g., [LC83, Pan85, Bap89, Gur05, CP16].
5. Cayley's first hyperdeterminant: $\det_d^{\boxtimes m}$, [Cay46, Bar95, HL13, CP16, BI17, AY23], also called the combinatorial hyperdeterminant, also called the Pascal determinant in [Lan11, §8.3] and [Lan15], or the quantum permanent in [Gur04, §3].

Equivariance gives improved hardness results. The Kronecker product is *equivariant*, see Propositions 7.8 and 7.19 (cp. [IOT25]), which means that a restriction of a factor gives a restriction of the product, for example $\mathrm{mon}_d \le \det_d$ implies by equivariance that $\mathrm{mon}_d \boxtimes \mathrm{mon}_d \le \mathrm{mon}_d \boxtimes \det_d \le \det_d \boxtimes \det_d$, which are two results with separate proofs in the literature, see [Gur04, Gur05], which show that the mixed discriminant and Cayley's first hyperdeterminant are VNP-complete over characteristic different from 2. We define the polynomial $Q_{2d} := \prod_{i=1}^d \big(y_i^{(0)} z_i^{(0)} + y_i^{(1)} z_i^{(1)}\big)$ and prove that $\mathrm{mon}_{2d} \boxtimes Q_{2d}$ is VNP-complete over all commutative semirings in §9.I. Note that we get this stronger hardness result in very close "proximity" to the permanent: While $\mathrm{per}_{2d} = \Gamma_{1,2d} \boxtimes \Gamma_{1,2d}$, we have $Q_{2d} \le \Gamma_{1,2d} \boxtimes \Gamma_{2,2d}$. Since $Q_{2d} \le \det_{2d}$ via block diagonal matrices of 2×2 blocks, it follows by equivariance that $\mathrm{mon}_{2d} \boxtimes Q_{2d} \le \mathrm{mon}_{2d} \boxtimes \det_{2d}$, hence the mixed discriminant is VNP-complete over all commutative rings. Thus, also Cayley's first hyperdeterminant is VNP-complete over all commutative rings. This strengthens Gurvits' results. Analogously, $Q_{2d} \le \mathrm{per}_{2d}$, thus the hyperpermanent $\mathrm{mon}_d^{\boxtimes k}$ is VNP-complete over all commutative semirings for any fixed $k \ge 3$. We also remark that for the hyperhafnian the proof $\mathrm{mon}_d^{\boxtimes 2} \le \mathrm{HypHaf}_d^{[2]}$ from [Bar16, §4.1] generalizes to $\mathrm{mon}_d^{\boxtimes r} \le \mathrm{HypHaf}_d^{[r]}$ for $r \ge 3$, and for the hyperpfaffian the proof $\mathrm{mon}_d^{\boxtimes 2} \le \mathrm{HypPfaf}_d^{[4]}$ from

[GIM+20] generalizes to $\mathrm{mon}_d^{\boxtimes r} \leq \mathrm{HypPfaf}_d^{[2r]}$, which gives VNP-completeness over all commutative (semi-)rings for $r \geq 3$.

In fact, we also give a basis-independent definition of the Kronecker product, which reveals that it is actually a product of a symmetric tensor h with a polynomial f, and the polynomial f can be replaced by a tensor (for non-commutative computation) or by an element of the exterior algebra, see §7.III, which is based on [IOT25], but [IOT25] only consider polynomials. We elaborate on this in §4.

The hypercomputant matrix. The Kronecker product is not only defined for polynomials, but also for matrices of polynomials, e.g., $\mathsf{h}\Xi_{n,d} := \Xi_{n,d}^{\boxtimes 2}$ is an $n^2 \times n^2$ matrix of homogeneous degree d elements that we call the hypercomputant. For polynomials, in terms of monomials, we have (with 0-indexed rows and columns)

$$[\mathsf{h}\Xi_{n,d}]_{an+a',bn+b'} = \sum_{\pi\in\mathfrak{S}_d} \sum_{\substack{1\leq i_1,\ldots,i_{d-1}\leq n\\ 1\leq j_1,\ldots,j_{d-1}\leq n\\ i_0=a,\ i_d=b}} (x^{(\pi(1))}_{i_{\pi(1)-1},i_{\pi(1)}} \boxtimes x^{(1)}_{a',j_1}) \cdots (x^{(\pi(d))}_{i_{\pi(d)-1},i_{\pi(d)}} \boxtimes x^{(d)}_{j_{d-1},b'}).$$

The Boolean decision problem corresponding to any one of the entries is the following: Given a set S of ordered pairs (e_1, e_2) ($e_1 = e_2$ is allowed) of edges in the fully multipartite layered digraph with 1 vertex in vertex layer 0, n vertices in vertex layer $1 \leq i < d$, and 1 vertex in vertex layer d, determine whether or not there exists a size d subset of S such that the e_1 form a path from vertex layer 0 to d and the e_2 form a path from vertex layer 0 to d.

Let $\mathsf{h}\Gamma_{n,d}$ be the bottom-right entry of $\mathsf{h}\Xi_{n,d}$ (every entry is in fact the same up to renaming variables). Using the equivariance of the Kronecker product, it is easy to see that $\mathsf{h}\Gamma_{d,d}$ is a VNP-complete polynomial, via a restriction to HC_d, see Proposition 7.33. The proof of the VNP-completeness of HC_d requires gadgets. We use the equivariance to prove the VNP-completeness of $\mathsf{h}\Gamma_{d,d}$ over all commutative semirings and for polynomials, tensors, and elements of the exterior algebra, in a gadget-free way in §9. We think that having such a proof for key objects of the theory is beneficial. The proof of the VNP-completeness of Q_{2d} is a strengthening of this proof via gadgets, focused on polynomials.

We prove that $\mathsf{h}\Xi_{n,d}$ is polystable as a matrix of polynomials over $\mathbb{C}$, see Proposition 11.14. The orbit closure of $\mathsf{h}\Xi_{n,d}$ is a candidate for geometric complexity theory, but we determine that $\mathsf{h}\Xi_{n,d}$ is *not* characterized by its stabilizer, so our analysis cannot stop here! We find that there are better candidates, which are highlighted from two unrelated directions.

• We prove in (11.39) that $\mathsf{h}\Xi_{n,d}$ is not characterized by its stabilizer. However, one can decompose $\mathsf{h}\Xi_{n,d} = \sum_{\pi\in\mathfrak{S}_d} \mathsf{h}\Xi_{n,d,\pi}$, and each $\mathsf{h}\Xi_{n,d,\pi}$ is characterized by its stabilizer, see §11.III. Here, $\mathsf{h}\Xi_{n,d,\pi}$ is the π-isotypic part of $\mathsf{h}\Xi_{n,d}$ with respect to the action of the torus $T = (\mathbb{C} \setminus \{0\})^{d\times d}$, acting on $\mathbb{C}^{n^4 d^2}$ via $\mathrm{diag}(\alpha_{1,1}, \ldots, \alpha_{d,d}) x^{(k)}_{\mathbf{i}} \boxtimes x^{(\ell)}_{\mathbf{j}} = \alpha_{k,\ell}\,(x^{(k)}_{\mathbf{i}} \boxtimes x^{(\ell)}_{\mathbf{j}})$, see Remark 8.5. In terms of monomials, we have

$$[\mathsf{h}\Xi_{n,d,\pi}]_{an+a',bn+b'} = \sum_{\substack{1\leq i_1,\ldots,i_{d-1}\leq n\\ 1\leq j_1,\ldots,j_{d-1}\leq n\\ i_0=a,\ i_d=b}} (x^{(\pi(1))}_{i_{\pi(1)-1},i_{\pi(1)}} \boxtimes x^{(1)}_{a',j_1}) \cdot (x^{(\pi(2))}_{i_{\pi(2)-1},i_{\pi(2)}} \boxtimes x^{(2)}_{j_1,j_2}) \cdots (x^{(\pi(d))}_{i_{\pi(d)-1},i_{\pi(d)}} \boxtimes x^{(d)}_{j_{d-1},b'})$$

see equation (7.27) and equation (8.12). We prove that each $\mathsf{h}\Xi_{n,d,\pi}$ is polystable, see Proposition 11.13.

• The matrices $\mathsf{h}\Xi_{n,d,\pi}$ also appear when taking a parameterized complexity viewpoint. Since $\mathsf{h}\Xi_{n,d}$ has two indices, it makes sense to study the situation for fixed d and growing n. It turns out that $w(\mathsf{h}\Xi_{n,d})$ grows polynomially for fixed d. In fact, each $w(\mathsf{h}\Xi_{n,d,\pi})$ grows polynomially, and the growth exponent of $w(\mathsf{h}\Xi_{n,d})$ is exactly the maximum of the growth exponents of the $w(\mathsf{h}\Xi_{n,d,\pi})$, see §8. Hence, $\mathsf{h}\Xi_{n,d}$ decomposes into the sum of the $\mathsf{h}\Xi_{n,d,\pi}$ from a representation-theoretic point of view and from a parameterized complexity point of view.

2.III Parameterized complexity

In parameterized complexity theory one studies the computational complexity not only with respect to input length but also with respect to so-called *parameters*. Finding the correct parameter is not obvious in the Boolean theory, but in the algebraic variant that was recently introduced in [BE19], the degree parameter d is highlighted as a canonical choice. This viewpoint appeared recently independently in [DGI$^+$24, CKV24, vdBDG$^+$25]. One defines VBFPT as the set of double-indexed polynomials $f_{n,d}$ for which the set of growth exponents $\{\inf\{e_d \mid w(f_{n,d}) = O(n^{e_d})\} \mid d \in \mathbb{N}\}$ is bounded, see [Ike25], which is a subclass of VFPT (see [BE19]), the algebraic analogue of the class FPT of fixed-parameter tractable functions, see e.g. [DF13] for details. The class of hard double-indexed polynomials that we are interested in is called VW[1], and notably the clique polynomial is VW[1]-complete, see [BE19]. We prove that for any fixed $m \geq 2$ we have that $\mathsf{h}\Gamma_{n,d}$ (and in fact every $\Gamma_{n,d}^{\boxtimes m}$ for fixed $m \geq 2$) is VW[1]-complete by restricting $\mathsf{h}\Gamma_{n,d}$ to the degree d clique polynomial, see Proposition 7.35 and Lemma 10.4.

We focus on the width growth exponents of $\mathsf{h}\Xi_{n,d}$ for fixed d in §8. Since our definitions are very general, we can make any choice of commutative semiring, and choose between polynomials, tensors, and the exterior algebra. We determine the growth exponents ξ_d in the tensor case (i.e., noncommutative polynomials) using our generalization of Nisan's result, independent of the commutative semiring, §8.I. We also determine the exponents ξ_d for polynomials over any zerosumfree (i.e., no nonzero elements can sum up to 0) commutative semiring, using techniques from monotone complexity, §8.II. Both exponents have combinatorial descriptions from which we see that the noncommutative exponent can never be smaller than the monotone exponent. The first exponents are listed in the table below.

d	2	3	4	5	6	7
ξ_d for tensors	4	4	6	6	8	8
ξ_d for polynomials over zerosumfree R	4	4	6	6	6	8

Both sequences start out identically, but the first difference appears at $d = 6$. We prove that both sequences grow as $\Theta(d)$. Usually, in the literature monotone complexity and noncommutative complexity are separate topics, since it is often not clear what should be a canonical noncommutative analog, for example for the clique polynomial, or what should be a canonical monotone analog, for example for the determinant. Here, $\mathsf{h}\Xi_{n,d}$ is used for both settings.

We prove that the exponents grow monotonically, and are always at least 2, see §8. For the most interesting case for geometric complexity theory, polynomials over $\mathbb{C}$, we know no nontrivial lower bound.

In §10 we explain that over polynomials we have that "$\{\xi_d \mid d \in \mathbb{N}\}$ is unbounded" is equivalent to VBFPT $\neq$ VW[1], and that

$$\mathrm{VFPT} \neq \mathrm{VW[1]} \implies \mathrm{VBFPT} \neq \mathrm{VW[1]} \implies \mathrm{VBP} \neq \mathrm{VNP}$$

sits between the Bläser-Engels VFPT $\neq$ VW[1] conjecture and Valiant's VBP $\neq$ VNP conjecture.

3 Further related work

For matrices, the equivariance of the Kronecker product is widely known as the so-called mixed-product property. Kronecker products of symmetric tensors recently appeared in [Gha26], where it is shown that Kronecker products work favorably with taking the Hessian. Tensor variants of $\Gamma_{n,d}$ and tr-$\Xi_{n,d}$ appear in computational mathematics, quantum physics, and quantum chemistry, see e.g., [HMS19, SPM$^+$15]. Work on stabilizers in geometric complexity has been advanced recently in [ASS23]. [CKV24] proved that over fields the growth exponents of $w(\mathsf{h}\Gamma_{n,d})$ and the growth exponent of the determinantal complexity of $\mathsf{h}\Gamma_{n,d}$ are the same, see [Ike25, Pro. 7.5], and recently Ahmed Shaaban (2026, personal communication) proved using ChatGPT 5.6 that this correspondence breaks over some commutative rings. As far as we know, there has been no attempt to define a monotone computational model based on determinantal complexity, but the polynomial identity testing problem for noncommutative determinantal computations has been successfully studied [GGOW16]. [Kum19] proved that $w(x_1^d + \cdots + x_n^d) = \lceil \frac{n}{2} \rceil$. The corresponding matrix version is trivial to prove: $w(\mathrm{diag}(x_1^d, \ldots, x_n^d)) = n$.

4 Proof ideas

Nisan's property: Zariski-closedness We translate Nisan's noncommutative lower bound result to matrices of noncommutative polynomials (i.e., tensors) by translating it to the tensor language: A matrix that has an order d tensor a at position (i,j) and zeros everywhere else is interpreted as a tensor $e_i \otimes a \otimes e_j^*$. Using this language, we re-interpret all objects in Nisan's proof and see that the result still holds over matrices, see Theorem 8.8. This gives the first property of $\Xi_{n,d}$, i.e., that the set of width $\leq n$ matrices of tensors is Zariski-closed.

Reductive stabilizer The second property concerns the stabilizer of $\Xi_{n,d}$, see §11.I. We use the fact that the stabilizer of tr-$\Xi_{n,d}$ has been determined in [Ges16], and every element stabilizing the matrix also stabilizes its trace. We then analyze which elements of the stabilizer of tr-$\Xi_{n,d}$ also stabilize $\Xi_{n,d}$, which is a much simpler task than having to deal with the whole group.

Kronecker product for several algebras over any commutative semiring We introduce the Kronecker product in a didactic step-by-step way, first for symmetric tensors over characteristic 0 in §7.I, then for polynomials over commutative semirings in §7.II, and finally via the Kronecker action, i.e., the product of a symmetric tensor and an algebra element, so that we can also handle tensors and the exterior algebra in §7.III. This last Kronecker action construction works as follows, cp. [IOT25]. Given a commutative semiring R, let $V = R^v$ and $W = R^w$. Let $\mathrm{Sym}_d(V) := (\bigotimes^d V)^{\mathfrak{S}_d}$ denote the set of symmetric order d tensors, and let $S^d(V) := (\bigotimes^d V)/\mathfrak{S}_d$ denote the set of homogeneous degree d polynomials, which we interpret as so-called coinvariants, i.e., elements in a quotient by the action of the symmetric group. Over fields of characteristic 0 there are natural isomorphisms between both vector spaces, but we do not have those available. We start with the classical R-bilinear map $\boxtimes : \bigotimes^d V \times \bigotimes^d W \to \bigotimes^d (V \boxtimes W)$, and restrict this map to symmetric tensors $\psi : \mathrm{Sym}_d(V) \times \bigotimes^d W \to \bigotimes^d (V \boxtimes W)$. We fix $s \in \mathrm{Sym}_d(V)$ to obtain an $\mathfrak{S}_d$-equivariant R-linear map $\phi_s : \bigotimes^d W \to \bigotimes^d (V \boxtimes W)$, $f \mapsto \psi(s,f)$. To this map we apply the coinvariant functor and obtain $\psi_s : S^d(W) \to S^d(V \boxtimes W)$, $[f] \mapsto [\psi(s,f)]$. In total, this gives an R-linear map $\kappa : \mathrm{Sym}_d(V) \to \mathrm{Hom}\Big(S^d(W), S^d(V \boxtimes W)\Big)$, $s \mapsto \psi_s$, which we interpret as an R-bilinear map $\boxtimes : \mathrm{Sym}_d(V) \times S^d(W) \to S^d(V \boxtimes W)$, $(s,[f]) \mapsto \psi_s([f]) = [\psi(s,f)]$. Two polynomials are multiplied by first converting one of them into its corresponding symmetric tensor, which is called the *polarization* of the polynomial, for example the polarization of x^2y is $2(x \otimes x \otimes y + x \otimes y \otimes x + y \otimes x \otimes x)$. The construction of $\boxtimes$ looks asymmetric, but it is symmetric up to renaming variables. If one changes the coinvariant functor in the definition to a different functor, one gets the action on the tensor algebra and on the exterior algebra as well. The equivariance of the Kronecker product can be shown combinatorially, but it also follows directly from this general construction. We give full details in §7.III.

Using equivariance to prove hardness In Proposition 7.33 we prove that $\mathsf{h}\Gamma_{n,d}$ restricts to the Hamiltonian cycle polynomial, which we write in a suggestive form here: $\mathrm{HC}_d := \sum_{\phi \in \mathfrak{S}_{d-1}} x_{d,\phi(1)} \cdot x_{\phi(1),\phi(2)} \cdots x_{\phi(d-1),d}$. We note that

$$\Gamma_{d,d} \;=\; \textstyle\sum_{\phi:[d-1]\to[d]} x^{(1)}_{d,\phi(1)} x^{(2)}_{\phi(1),\phi(2)} x^{(3)}_{\phi(2),\phi(3)} \cdots x^{(d)}_{\phi(d-1),d}.$$

Hence $\mathsf{h}\Gamma_{d,d} = \Gamma_{d,d} \boxtimes \Gamma_{d,d} \geq \mathrm{mon}_d \boxtimes \Gamma_{d,d} = \sum_{\substack{\phi:[d-1]\to[d] \\ \pi \in \mathfrak{S}_d}} x^{(1)}_{\pi(1),d,\phi(1)} x^{(2)}_{\pi(2),\phi(1),\phi(2)} \cdots x^{(d)}_{\pi(d),\phi(d-1),d}$.

Let T be a restriction as follows: $T(x^{(k)}_{a,b,c}) = x_{b,c}$ if $a = c$, and $T(x^{(k)}_{a,b,c}) = 0$ otherwise. Now

$$T(\mathrm{mon}_d \boxtimes \Gamma_{d,d}) = \sum_{\substack{\phi \in \mathfrak{S}_d \\ \phi(d)=d}} x_{d,\phi(1)}\, x_{\phi(1),\phi(2)} \cdots x_{\phi(d-1),d} \;=\; \mathrm{HC}_d. \tag{4.1}$$

This finishes the proof. We use a similar but slightly more complicated argument for the restriction to the Clique polynomial in Proposition 7.35, which implies VW[1]-hardness. Such reductions use no gadgets, but the original hardness proof of HC_d uses gadgets, so we give a fully gadget-free proof of the VNP-hardness of $\mathsf{h}\Gamma_{d,d}$ in Theorem 9.1 over arbitrary commutative semirings, over all 3 algebras. Special care must be taken here, because elements in the exterior algebra vanish for repeated factors, and over the exterior algebra we do not have a substitution homomorphism that would let us substitute variables for semiring elements. Therefore we carefully introduce special padding variables, hence Theorem 9.1 states

$$\sum_{B(\mathbf{b})=1}(x,\overline{x})^{\mathbf{b},\overline{\mathbf{b}}}\prod_{i=m+1}^{d}\epsilon^{(i)} \;\leq\; \mathsf{h}\Gamma_{n,d},$$

with $n \leq (4|B|+2m)^2$ and $d \leq 7|B|+m+1$, where $(x,\overline{x})^{\mathbf{b},\overline{\mathbf{b}}}$ is defined as $(x,\overline{x})^{\mathbf{b},\overline{\mathbf{b}}} := x_1^{b_1}\overline{x}_1^{\overline{b}_1}x_2^{b_2}\overline{x}_2^{\overline{b}_2}\cdots x_m^{b_m}\overline{x}_m^{\overline{b}_m}$, which works well over the exterior algebra. The proof makes essential use of the two $\Gamma_{n,d}$ factors and equivariance, one factor tracing the arithmetization of the Tseytin transformation of the Boolean circuit B, while the other factor takes care of consistent variable assignments via a variant of an elementary symmetric polynomial, and at the end we use a restriction similar to (4.1).

In §9.I we focus on polynomials and sharpen the previous proof using gadgets to show the VNP-completeness of $Q_{2d}\boxtimes\mathrm{mon}_{2d}$ over all commutative semirings, where $Q_{2d} := \prod_{i=1}^{d}\big(y_i^{(0)}z_i^{(0)}+y_i^{(1)}z_i^{(1)}\big)$. The main idea is to interpret the monomials of $Q_{2d}\boxtimes\mathrm{mon}_{2d}$ as perfect matchings between the two kinds of variables in the two polynomials, and design a restriction so that such a matching exists precisely when the Boolean choice represented by the monomial in Q_{2d} describes a satisfying assignment of a Tseytin encoding of a given Boolean circuit. We use two types of gadgets to construct the allowed pairings: an equality gadget on the y-variables to establish global consistency; and a rosette-like gadget on the z-variables that enforces every clause of the formula to be satisfied. This hardness result has several immediate consequences via equivariance, as explained in §2.II.

Exponents In §8.I we calculate the exponents ξ_d for noncommutative computation (i.e., tensors) via Nisan's result. The upper bound is a variant of Grenet's construction [Gre11], while the lower bound uses Nisan's bound (Nisan's result can serve as an upper bound only over fields, not arbitrary commutative semirings). The exponents are $d+2$ if d is even, and $d+1$ if d is odd.

In §8.II we obtain the exponents ξ_d for monotone commutative computation by a translation of the technique of [KPR23]. The exponent is closely related to the pathwidth of a graph corresponding to a permutation π that is studied in [AR08, Def 2.6]. The exponents grow as $\Theta(d)$, which follows from the fact that there are expanders among these graphs. We obtain a precise combinatorial description of the exponent from which it is clear that the monotone exponent cannot exceed the noncommutative exponent: compare (8.10) with (8.14), which has an additional parameter σ for choosing the order in which the variables are computed. In both cases, the exponent is the maximum exponent of $w(\Xi_{n,d,\pi})$ over all π.

Geometric complexity theory All isotypic decompositions and multiplicities are obtained via branching rules of the general linear group and the symmetric group. This branching also enables us to quickly check whether an object is characterized by its stabilizer. The connected component of the identity of the stabilizer of $\mathsf{h}\Xi_{n,d}$ is obtained by a tedious Lie algebra calculation, similar to, but more complicated than, the one intentionally not spelled out in [Ges16] in §11.V. Based on this result, in §11.VI we then determine the stabilizer by the method pioneered in [BGL14, Ges16].

A point is polystable if its orbit under the special linear group is closed, but here we are dealing with a product of groups that act in different ways, so we need a slight adjustment. The polystability of our different matrices is obtained via the Hilbert–Mumford criterion, which has the additional benefit that it makes clear which subgroup should be used for the definition of polystability, see §11.IV.

Full paper

The full paper is not written from the discovery order perspective explained above, but starts with the setup towards defining $\Xi_{n,d}$ and $\mathsf{h}\Xi_{n,d}$ in the required generality. The style and section order of the full paper are intended for readers unfamiliar with complexity theory, hence for example Theorem 9.1 is not phrased in terms of VNP-completeness, also because the theorem goes beyond it and works for tensors and also for the exterior algebra, but it is clear that the theorem implies that $\mathsf{h}\Gamma_{\mathsf{poly}(n),\mathsf{poly}(n)}$ is VNP-complete. We focus on VNP-completeness in §9.I. The introduction was stripped from the full paper, and some references point to §1 instead.

In §5 we introduce basic objects in the required generality. In §6 we define the width and relate it to algebraic branching programs and prove elementary properties. In §7 we introduce the Kronecker product and prove its properties. We then define $\mathsf{h}\Xi_{n,d}$ and relate it to Valiant's conjecture, and prove that higher-order computation of order greater than 2 gives the same conjecture. We also give restrictions to HC_d and the clique polynomial. In §8 we define the width growth exponents ξ and determine them for noncommutative and for monotone computation. In §9 we prove the VNP-completeness results that are based on the arithmetization of the Tseytin transform, starting with the gadget-free result. In §10 we relate the VBFPT $\neq$ VW[1] conjecture to Valiant's conjecture and the Bläser–Engels conjecture, and other well-known conjectures in complexity theory via standard techniques. In §11 we examine the different settings, stabilizers, polystability, and representation-theoretic multiplicities.

Occasionally we include short proofs of classical results to make the full paper more self-contained, but we clearly mark those results as such.

Contents

Full paper

5 Preliminaries

Define the set of polynomially bounded sequences in n as $\mathsf{poly}(n) = \bigcup_{d\in\mathbb{N}} O(n^d)$. Define the set of bivariately polynomially bounded families as $\mathsf{poly}(n,d) = \{(r_{n,d})_{\substack{n\in\mathbb{N}\\ d\in\mathbb{N}}} \mid \exists\phi\in\mathbb{R}[n,d]\,\forall n,d : r_{n,d} \leq \phi(n,d)\}$.

5.I Variables

We work over countably infinitely many variables, where the exact names of the variables do not matter, and neither does their order matter. Everything that we encounter will use only a finite number of variables. Depending on the application, we use different variable names. Most variables are of the form $x_{i_1,\ldots,i_j}$ for some $j \geq 0$ and some list $(i_1,\ldots,i_j) \in \mathbb{N}^j$. The tensor product operation on variables is defined by concatenating index lists, for example $x_{1,2} \boxtimes x_3 = x_{1,2,3}$. To improve the readability, sometimes indices are written as superscripts in parentheses, which are to be interpreted as additional indices, i.e., $x^{(k)}_{i,j} = x_{i,j,k}$. We also use other variable names besides x, for example y, or more exotic variable names in §9. These are distinct variables from the variables in x. We use the tensor product for y in the same way as for x, e.g., $y_{1,2} \boxtimes y_3 = y_{1,2,3}$. We also allow mixed tensor products between variables that have different symbols, but we do not introduce any special notation for those, i.e., $x_{1,2} \boxtimes y_3$ is a variable with no other name. We denote the whole set of variables by Var. Since the number of variable symbols we use is finite, it is clear that Var contains countably infinitely many elements.

5.II Semialgebras

A commutative semiring R is a tuple $(R,+,\cdot,0_R,1_R)$ with $0_R \neq 1_R$ and where $(R,+,0_R)$ is a commutative monoid and $(R,\cdot,1_R)$ is a commutative monoid, and where distributivity holds, i.e., $\alpha\cdot(\beta+\gamma) = \alpha\cdot\beta + \alpha\cdot\gamma$ and $(\beta+\gamma)\cdot\alpha = \beta\cdot\alpha+\gamma\cdot\alpha$, and it holds that $\alpha\cdot 0_R = 0_R = 0_R\cdot\alpha$. The main interesting examples in our context are the set of nonnegative real numbers $\mathbb{R}_{\geq 0}$ and the set of complex numbers $\mathbb{C}$.

An R-semimodule $(M,+,\cdot,0_M)$ is a tuple such that $(M,+,0_M)$ is a commutative monoid and the scalar multiplication map $\cdot : R\times M \to M$ satisfies $\alpha\cdot(m+n) = \alpha\cdot m+\alpha\cdot n$ and $(\alpha+\beta)\cdot m = \alpha\cdot m+\beta\cdot m$, and $(\alpha\cdot\beta)\cdot m = \alpha\cdot(\beta\cdot m)$, and $1_R\cdot m = m$, and $0_R\cdot m = 0_M = \alpha\cdot 0_M$, for $\alpha,\beta\in R$, $m,n\in M$. Multiplication symbols are often omitted. This is sometimes called a *left* R-semimodule, but in this paper all R-semimodules are left R-semimodules.

An R-semialgebra A is an R-semimodule with an additional associative multiplication operation $A\times A\to A$ that satisfies $(a+b)c = ac+bc$ and $c(a+b) = ca+cb$ and $(\alpha a)b = \alpha(ab) = a(\alpha b)$, and there exists $1_A \in A$ with $1_A a = a = a1_A$ for all $a\in A$.

For an R-semialgebra A, let $A\times A$ denote the product semialgebra with addition $(a_1,b_1)+(a_2,b_2) = (a_1+a_2, b_1+b_2)$ and scalar multiplication $\alpha(a,b) = (\alpha a,\alpha b)$ and multiplication $(a_1,b_1)\cdot(a_2,b_2) = (a_1\cdot a_2, b_1\cdot b_2)$. A congruence on an R-semialgebra A is defined as an R-subsemialgebra of $A\times A$ that is reflexive ($(a,a)\in\sim$), symmetric ($(a,b)\in\sim$ implies $(b,a)\in\sim$), and transitive ($(a,b)\in\sim$ and $(b,c)\in\sim$ implies $(a,c)\in\sim$). For an R-semialgebra A and a congruence $\sim$ we can define the quotient R-semialgebra $A/_\sim$ on the equivalence classes of $\sim$ via $[a]+[b] = [a+b]$, $\alpha[a] = [\alpha a]$, $[a]\cdot[b] = [a\cdot b]$ where the well-definedness can be seen as follows. If $[a] = [a']$ and $[b] = [b']$, then $(a,a')\in\sim$ and $(b,b')\in\sim$, and hence $(a+b, a'+b')\in\sim$, i.e., $[a+b] = [a'+b']$. Moreover, if $[a] = [a']$, then $(a,a')\in\sim$ and hence $(\alpha a,\alpha a') = \alpha(a,a')\in\sim$, i.e., $[\alpha a] = [\alpha a']$. Also, if $[a] = [a']$ and $[b] = [b']$, then $(a,a')\in\sim$ and $(b,b')\in\sim$, and hence $(a\cdot b, a'\cdot b')\in\sim$, i.e., $[a\cdot b] = [a'\cdot b']$. The semialgebra axioms are satisfied for $A/_\sim$ by inheritance from A.

A map $\varphi : M \to N$ between two R-semimodules is called *R-linear* if $\varphi(0_M) = 0_N$ and $\forall \alpha \in R, m, n \in M : \varphi(\alpha m + n) = \alpha\varphi(m) + \varphi(n)$. Let M denote the free R-semimodule generated by Var, i.e., M is the set of finite R-linear combinations of elements from Var. The *endomorphism monoid* End is the set of those R-linear maps from M to M that send every variable to a *finite* linear combination of variables.

For M the free R-semimodule generated by Var, let $\bigotimes^\bullet(M)$ denote the tensor algebra, i.e., the set of finite R-linear combinations of sequences of variables, where the product of sequences is given by the sequence concatenation. The homogeneous degree d component of $\bigotimes^\bullet(M)$ is defined as being the R-linear span of the set of sequences of exactly d variables, and is denoted by $\bigotimes^d(M)$. These sequences of variables are written with a tensor symbol as the divider symbol, e.g., $x_1 \otimes x_2$. Consider the following two quotients of $\bigotimes^\bullet(M)$, which we define carefully, because in general R might not have additive inverses.

- Define the congruence $\sim$ as the R-linear span of the (f, f') with $f, f' \in \mathscr{T}_d$ such that $\exists \pi \in \mathfrak{S}_d$ with $f = \pi f'$. All (f, f) are generators, and for each generator $(f, f') \in\sim$ we have $(f', f) \in\sim$. Note that $f \sim f'$ iff f can be transformed into f' by applying a permutation to each monomial, not necessarily the same for each. This implies that $\sim$ is transitive. Note also that $\sim$ is closed under the semialgebra product, because if $f \sim f'$ and $f'' \sim f'''$, then from the permutations that transform f into f' and those that transform f'' into f''' we can readily obtain the permutations that transform ff'' into $f'f'''$. Let $S^\bullet(M) := \bigotimes^\bullet(M)/_\sim$. The notation for the equivalence class of $v_1 \otimes \cdots \otimes v_d$ is $v_1 \cdots v_d$, or alternatively $[v_1 \otimes \cdots \otimes v_d]_S$.
- Define the congruence $\backsim$ as generated by the elements $(\ell \otimes \ell, 0)$ and $(0, \ell \otimes \ell)$, $\ell \in M$, and all (f, f), $f \in \bigotimes^\bullet(M)$. Since $(f, f) \in\backsim$, we have that $\backsim$ is reflexive. We have $f \backsim f'$ if and only if $f + h + \sum_i h_i \otimes \ell_i \otimes \ell_i \otimes m_i = f' + h + \sum_j h'_j \otimes \ell'_j \otimes \ell'_j \otimes m'_j$, hence $f' \backsim f$. Hence, $\backsim$ is symmetric. Transitivity is seen as follows. If $f \backsim f'$ and $f' \backsim f''$, then
$$f + h + n \otimes \ell \otimes \ell \otimes m = f' + h + n' \otimes \ell' \otimes \ell' \otimes m'$$
and
$$f' + \bar{h} + \bar{n}' \otimes \bar{\ell}' \otimes \bar{\ell}' \otimes \bar{m}' = f'' + \bar{h} + n'' \otimes \ell'' \otimes \ell'' \otimes m'',$$
which implies that
$$f' + \bar{h} + \bar{n}' \otimes \bar{\ell}' \otimes \bar{\ell}' \otimes \bar{m}' + h + n' \otimes \ell' \otimes \ell' \otimes m' = f'' + \bar{h} + n'' \otimes \ell'' \otimes \ell'' \otimes m'' + h + n' \otimes \ell' \otimes \ell' \otimes m'$$
hence
$$f + h + n \otimes \ell \otimes \ell \otimes m + \bar{h} + \bar{n}' \otimes \bar{\ell}' \otimes \bar{\ell}' \otimes \bar{m}' = f'' + \bar{h} + n'' \otimes \ell'' \otimes \ell'' \otimes m'' + h + n' \otimes \ell' \otimes \ell' \otimes m'$$
and thus $f \backsim f''$. Let $\bigwedge^\bullet(M) := \bigotimes^\bullet(M)/_\backsim$. The notation for the equivalence class of $v_1 \otimes \cdots \otimes v_d$ is $v_1 \wedge \cdots \wedge v_d$, or alternatively $[v_1 \otimes \cdots \otimes v_d]_\bigwedge$. Note that $\ell \wedge \ell = 0$ by definition. This implies $\ell \wedge \ell' + \ell' \wedge \ell = 0$, because $\ell \wedge \ell' + \ell' \wedge \ell = \ell \wedge \ell' + \ell' \wedge \ell + (\ell \wedge \ell) + (\ell' \wedge \ell') = \ell \wedge (\ell + \ell') + \ell' \wedge (\ell + \ell') = (\ell + \ell') \wedge (\ell + \ell') = 0$. Moreover, this implies $\ell \wedge \ell' \wedge \ell'' = \ell'' \wedge \ell \wedge \ell'$, because $\ell \wedge \ell' \wedge \ell'' = \ell \wedge \ell' \wedge \ell'' + (\ell + \ell'') \wedge (\ell + \ell'') \wedge \ell' = \ell \wedge \ell' \wedge \ell'' + \ell \wedge \ell'' \wedge \ell' + \ell'' \wedge \ell \wedge \ell' = \ell'' \wedge \ell \wedge \ell'$.

We write $\mathscr{T} := \bigotimes^\bullet(M)$, $\mathscr{S} := S^\bullet(M)$, $\mathscr{W} := \bigwedge^\bullet(M)$. The homogeneous degree d components are written as $\mathscr{T}_d$, $\mathscr{S}_d$, and $\mathscr{W}_d$.

5.1 Remark. Arbitrary semirings instead of commutative semirings can be used for most parts when applying the following changes.

- The semialgebra axiom is changed from $\alpha(ab) = (\alpha a)b = a(\alpha b)$ to $(\alpha\beta)(ab) = (\alpha a)(\beta b)$ for standard basis vectors a and b, see [Wei86, eq. (1.4)]. Here, a basis means that every element can be expressed uniquely as an R-linear combination of basis vectors. The standard basis of $\mathscr{W}_d$ is different over semirings compared to the standard basis over rings, because over semirings $y \wedge x$ and $x \wedge y$ can be R-linearly independent, while over rings we have $y \wedge x = y \wedge x + (x - y) \wedge (x - y) - x \wedge x - y \wedge y = -x \wedge y$.

- Even over non-commutative rings, the multiplication operation in $\mathscr{S}$ is not necessarily commutative. Similarly, taking the product of two elements in $\mathscr{W}_{2d}$, $d \geq 1$, is not necessarily commutative.
- The definition of $\boxtimes$ in §7 works via R-bilinear continuation, which introduces a similar subtlety as for semialgebras, which can be resolved in the same way by choosing standard bases. As a consequence, (7.7) can fail if R is not commutative.

It is unclear whether the lower bound arguments involving Lemma 5.7 would still hold. We avoid these complications and thus work over commutative semirings. This is also the generality in which mathlib [mC20] defines the exterior algebra in `Mathlib/LinearAlgebra/ExteriorAlgebra/Basic.lean`. $\diamond$

5.III Matrices

R^n is an R-semimodule with the standard basis $\{e_1, \ldots, e_n\}$, i.e., every element in R^n has a unique expression as an R-linear combination of $e_1, \ldots, e_n$. The dual R-semimodule $(R^n)^*$ is defined as the set of R-linear maps $R^n \to R$. Every $\phi \in (R^n)^*$ is uniquely determined by its values $\phi(e_i)$, and a basis of $(R^n)^*$ is given by the e_i^*, which are defined via $e_i^*(e_j) = \delta_{i,j}$, where $\delta_{i,j} = 1$ if $i = j$, and $\delta_{i,j} = 0$ otherwise.

For an R-semialgebra A, define $\mathrm{Mat}_n(A) := R^n \otimes A \otimes (R^n)^*$. The multiplication $\mathrm{Mat}_n(A) \times \mathrm{Mat}_n(A) \to \mathrm{Mat}_n(A)$ is defined via

$$(e_i \otimes a \otimes e_j^*) \cdot (e_{i'} \otimes a' \otimes e_{j'}^*) = e_j^*(e_{i'})\, e_i \otimes (aa') \otimes e_{j'}^*$$

extended bilinearly. We embed A_d into $\mathrm{Mat}_1(A_d)$ via $f \mapsto e_1 \otimes f \otimes e_1^*$. We embed Mat_{n-1} into Mat_n as the notation suggests: $(e_i \otimes a \otimes e_j^*) \mapsto (e_i \otimes a \otimes e_j^*)$.

For $F = \sum_{i,j} e_i \otimes m_{i,j} \otimes e_j^*$ we write $[F]_{i,j} = m_{i,j}$, i.e., we use brackets $[\,]$ to extract matrix entries. We use this notation, because our matrices will come from families of matrices, and hence possibly already have subscripts and/or superscripts.

5.IV Restrictions

The monoid End acts linearly on $\mathscr{T}_d$ via

$$T(x_{i_1} \otimes \cdots \otimes x_{i_d}) = T(x_{i_1}) \otimes \cdots \otimes T(x_{i_d}),$$

extended linearly. The action commutes with the map to the quotient $\mathscr{S}_d$ or $\mathscr{W}_d$, in other words

$$T(x_{i_1} \cdots x_{i_d}) = T(x_{i_1}) \cdots T(x_{i_d}), \qquad T(x_{i_1} \wedge \cdots \wedge x_{i_d}) = T(x_{i_1}) \wedge \cdots \wedge T(x_{i_d}).$$

This action lifts to $\mathrm{Mat}(A_d)$ entry-wise, i.e., $[T(F)]_{a,b} := T([F]_{a,b})$. Since A_d is an R-semimodule, for every $T_{\text{left}} \in \mathrm{Mat}(R)$ we define $T_{\text{left}} \cdot F$ via matrix multiplication, i.e., $[T_{\text{left}} F]_{a,b} := \sum_i [T_{\text{left}}]_{a,i} [F]_{i,b}$. Moreover, for $T_{\text{right}} \in \mathrm{Mat}(R)$ we define $F \cdot T_{\text{right}}$ analogously. All three actions on $\mathrm{Mat}(A_d)$ commute. We make this a left action of $\mathrm{Mat}(R) \times \mathrm{End} \times \mathrm{Mat}(R)$ on $\mathrm{Mat}(A_d)$ via

$$(T_{\text{left}}, T, T_{\text{right}})F := T_{\text{left}} \cdot T(F) \cdot T_{\text{right}}^{\mathsf{T}}$$

For $F, F' \in \mathrm{Mat}(A_d)$ we define $F \leq F'$ if there exists $T \in \mathrm{End}$, $T_{\text{left}} \in \mathrm{Mat}(R)$, $T_{\text{right}} \in \mathrm{Mat}(R)$ such that $(T_{\text{left}}, T, T_{\text{right}})F' = F$. We say that F' *restricts to* F. Clearly, $\leq$ is transitive.

5.V Number of essential variables

For $x_j \in \mathrm{Var}$, let $\partial_i(x_j) = 1$ if $i = j$, and 0 otherwise. We define the linear map $\partial_i : \mathscr{T}_d \to \mathscr{T}_{d-1}$ via

$$\partial_i(x_{j_1} \otimes \cdots \otimes x_{j_d}) := (\partial_i x_{j_1}) \otimes x_{j_2} \otimes \cdots \otimes x_{j_d} + x_{j_1} \otimes (\partial_i x_{j_2}) \otimes \cdots \otimes x_{j_d} + \cdots + x_{j_1} \otimes \cdots \otimes (\partial_i x_{j_d}) \quad (5.2)$$

The linear map $J : \mathscr{T}_d \to \mathscr{T}_{d-1} \otimes \mathscr{T}_1$ is defined via

$$J(t) = \sum_{i:\, x_i \in \mathrm{Var}} \partial_i(t) \otimes x_i.$$

The monoid End acts on $\mathscr{T}_{d-1} \otimes \mathscr{T}_1$ via

$$T(N \otimes \ell) := T(N) \otimes T(\ell)$$

for $T \in \mathrm{End}$. The map J is clearly equivariant, i.e., $T(J(f)) = J(T(f))$.

Analogously, we define the linear map $\partial_i : \mathscr{S}_d \to \mathscr{S}_{d-1}$ via

$$\partial_i(x_{j_1} \cdots x_{j_d}) \;:=\; (\partial_i x_{j_1}) x_{j_2} \cdots x_{j_d} \;+\; x_{j_1} (\partial_i x_{j_2}) \cdots x_{j_d} \;+ \cdots +\; x_{j_1} \cdots (\partial_i x_{j_d}). \tag{5.3}$$

This is well-defined, because permuting the variable positions in $x_{j_1} \cdots x_{j_d}$ in the left-hand side of (5.3) just permutes the summands on the right-hand side. We define the map $J : \mathscr{S}_d \to \mathscr{S}_{d-1} \otimes \mathscr{S}_1$ via $J(f) = \sum_{i:\, x_i \in \mathrm{Var}} \partial_i(f) \otimes x_i$. The monoid End acts on $\mathscr{S}_{d-1} \otimes \mathscr{S}_1$ via $T(N \otimes \ell) := T(N) \otimes T(\ell)$ for $T \in \mathrm{End}$. The map J is clearly equivariant, i.e., $T(J(f)) = J(T(f))$.

For $d \geq 3$, the linear map $\partial_i : \mathscr{W}_d \to \mathscr{W}_{d-1}$ is defined via

$$\partial_i(x_{j_1} \wedge \cdots \wedge x_{j_d}) \;:=\; (\partial_i x_{j_1}) \wedge x_{j_2} \wedge \cdots \wedge x_{j_d} + (1\,2) x_{j_1} \wedge (\partial_i x_{j_2}) \wedge \cdots \wedge x_{j_d} + \cdots + (1\,2)^{d-1} x_{j_1} \wedge \cdots \wedge (\partial_i x_{j_d}), \tag{5.4}$$

where $(1\ 2)(v_1 \wedge \cdots \wedge v_d) = v_2 \wedge v_1 \wedge v_3 \wedge \cdots \wedge v_d$. This is well-defined, which we see as follows: For $\partial_i(\ell) = \alpha$ we have $[\partial_i(\ell \otimes \ell \otimes x_{j_3} \otimes \cdots \otimes x_{j_d})]_{\bigwedge} = \alpha \ell \wedge x_{j_3} \wedge \cdots \wedge x_{j_d} + (1\,2) \alpha \ell \wedge x_{j_3} \wedge \cdots \wedge x_{j_d} + \ell \wedge \ell \wedge \ldots = 0$. Moreover, we have $[\partial_i(\ell \otimes x_{j_2} \otimes \ell \otimes x_{j_4} \otimes \cdots \otimes x_{j_d})]_{\bigwedge} = \alpha x_{j_2} \wedge \ell \wedge x_{j_4} \wedge \cdots \wedge x_{j_d} + \alpha \ell \wedge x_{j_2} \wedge x_{j_4} \wedge \cdots \wedge x_{j_d} + 0 = 0$. All other cases are analogous. The linear map $J : \mathscr{W}_d \to \mathscr{W}_{d-1} \otimes \mathscr{W}_1$ is defined via $J(t) = \sum_{i:\, x_i \in \mathrm{Var}} \partial_i(t) \otimes x_i$, and the monoid End acts on $\mathscr{W}_{d-1} \otimes \mathscr{W}_1$ via $T(N \otimes \ell) := T(N) \otimes T(\ell)$ for $T \in \mathrm{End}$. The map J is clearly equivariant, i.e., $T(J(f)) = J(T(f))$. For $d = 2$, if $-1 \in R$, then the same construction works by replacing the action of $(1\ 2)$ by the rescaling with -1. We will ignore the case $A_2 = \mathscr{W}_2$ in this paper.

In all three cases (with the exception of $\mathscr{W}_2$ with $-1 \notin R$), we obtain an equivariant map $J : A_d \to A_{d-1} \otimes A_1$.

The *rank* of an element $N \in A_{d-1} \otimes A_1$ is defined as the smallest r such that N can be written as

$$N = \sum_{i=1}^{r} N_i \otimes \ell_i$$

for some $N_i \in A_{d-1}$ and $\ell_i \in A_1$.

We know that

$$\text{if } h \leq f, \text{ then } \mathrm{rk}(J(h)) \leq \mathrm{rk}(J(f)), \tag{5.5}$$

because if $T \in \mathrm{End}$ and $f \in A_d$, $h \in A_d$ with $T(f) = h$ and $J(f) = \sum_{i=1}^{r} N_i \otimes \ell_i$, then $J(h) = J(T(f)) = \sum_{i=1}^{r} T(N_i) \otimes T(\ell_i)$ due to the equivariance.

Clearly, $\mathrm{rk}(J(f))$ is at most the number of variables in which f is specified. If R is a field, then $\mathrm{rk}(J(f))$ is sometimes called the *number of essential variables* of f.

This concept generalizes to matrices as follows. For $F \in \mathrm{Mat}(A_d)$ define $J : \mathrm{Mat}(A_d) \to \mathrm{Mat}(A_{d-1} \otimes A_1)$ via $[J(F)]_{a,b} := J([F]_{a,b})$. The rank of an element $M \in \mathrm{Mat}_m(A_{d-1} \otimes A_1)$ is defined as the smallest r such that M can be written as

$$M = \sum_{i=1}^{r} c_i \otimes N_i \otimes \ell_i \otimes j_i$$

for some $c_i \in R^m$, $j_i \in (R^m)^*$, $N_i \in A_{d-1}$, and $\ell_i \in A_1$.

We know that

$$\text{if } H \leq F, \text{ then } \mathrm{rk}(J(H)) \leq \mathrm{rk}(J(F)), \tag{5.6}$$

because if $(T_{\text{left}}, T, T_{\text{right}}) \in \mathrm{Mat}(R) \times \mathrm{End} \times \mathrm{Mat}(R)$ and $F \in \mathrm{Mat}(A_d)$, $H \in \mathrm{Mat}(A_d)$ with $(T_{\text{left}}, T, T_{\text{right}})(F) = H$ and $J(F) = \sum_{i=1}^r c_i \otimes N_i \otimes \ell_i \otimes j_i$, then $J(H) = J((T_{\text{left}}, T, T_{\text{right}})(F)) = \sum_{i=1}^r T_{\text{left}}(c_i) \otimes T(N_i) \otimes T(\ell_i) \otimes T_{\text{right}}^{\mathsf{T}}(j_i)$.

We also have that $\mathrm{rk}(J(F))$ is at most the number of variables in which F is specified. We say that $\mathrm{rk}(J(F))$ is the number of essential variables.

5.VI Rank

5.7 Lemma. *Over every commutative semiring R, the $n \times n$ identity matrix has rank n.* ◇

Proof. Clearly $\mathrm{rk}(I_n) \leq n$, because $I_n = \sum_{i=1}^n e_i \otimes e_i^*$. For the other direction, assume that $\mathrm{rk}(I_n) \leq r < n$. Then there exist $N \in R^{n \times r}$ and $M \in R^{r \times n}$ with $NM = I_n$. We pad N and M with zeros to obtain $n \times n$ matrices N' and M', where the last column of N' is zero and the last row of M' is zero. We have $N'M' = I_n$. Now apply [RS84, Theorem] to see that $M'N' = I_n$. This is a contradiction, because the bottom-right entry of $M'N'$ is zero. □

6 Width

In this section we introduce the width $w(F)$ of an element $F \in \mathrm{Mat}(A_d)$. The pair $(w(F), d)$ measures the computational complexity of F.

Let R be a commutative semiring and let $A \in \{\mathscr{T}, \mathscr{S}, \mathscr{W}\}$. Let $X_n^{(k)} := \big(x_{i,j}^{(k)}\big)_{1 \leq i,j \leq n}$ denote the $n \times n$ matrix whose entry at position (i,j) is the variable $x_{i,j}^{(k)} \in \mathrm{Var}$.

6.1 Definition (The computant matrix)**.** The *computant matrix* $\Xi_{n,d} \in \mathrm{Mat}_n(A_d)$ is defined as

$$X_n^{(1)} \cdots X_n^{(d)}.$$

We denote the bottom-right entry with $\Gamma_{n,d} := [\Xi_{n,d}]_{n,n}$. For $A = \mathscr{S}$, $\Gamma_{n,d}$ is called the iterated matrix multiplication polynomial. ◇

More explicitly, we have

$$[\Xi_{n,d}]_{a,b} \;=\; \sum_{1 \leq i_1, \ldots, i_{d-1} \leq n} x_{a,i_1}^{(1)}\, x_{i_1,i_2}^{(2)}\, x_{i_2,i_3}^{(3)} \cdots x_{i_{d-2},i_{d-1}}^{(d-1)}\, x_{i_{d-1},b}^{(d)}. \tag{6.2}$$

Note that

$$[\Xi_{n,d}]_{a,b} \text{ is the same as } [\Xi_{n,d}]_{a',b'} \text{ up to renaming variables.} \tag{6.3}$$

It will turn out that $\Xi_{n,d}$ has better properties than $\Gamma_{n,d}$ from a geometric and representation-theoretic perspective, see §11.I. We use $\Gamma_{n,d}$ to connect more easily to computational complexity theory.

6.4 Claim. $\Xi_{n,d} \leq \Xi_{n+1,d}$ *and* $\Gamma_{n,d} \leq \Gamma_{n+1,d}$. ◇

Proof. For $\Xi_{n,d} \leq \Xi_{n+1,d}$, just map every variable $x_{n+1,i}^{(k)}$ and $x_{i,n+1}^{(k)}$ to zero, $1 \leq i \leq n+1$. The same map proves $[\Xi_{n,d}]_{1,1} \leq [\Xi_{n+1,d}]_{1,1}$, which by (6.3) implies $\Gamma_{n,d} \leq \Gamma_{n+1,d}$. □

6.5 Definition (Width)**.** For $F \in \mathrm{Mat}(A_d)$, the *width* $w(F)$ is defined as the smallest n such that $F \leq \Xi_{n,d}$. ◇

6.6 Example. $w(\Xi_{n,d}) = n$. ◇

Proof. By definition, $w(\Xi_{n,d}) \le n$. We claim that $\mathrm{rk}(J(\Xi_{n,d})) = n^2 d$, which then finishes the proof by (5.6). Since $\Xi_{n,d}$ is defined in $n^2 d$ many variables, we have $\mathrm{rk}(J(\Xi_{n,d})) \le n^2 d$. For every one of the $n^2 d$ variables $x_{\mathbf{i}}$ in $\Xi_{n,d}$ we can reconstruct $\mathbf{i}$ from any monomial in $\partial_{\mathbf{i}}\Xi_{n,d}$, hence $J(\Xi_{n,d})$ contains an $(n^2 d) \times (n^2 d)$ permutation matrix as a submatrix, so the claim follows with Lemma 5.7. □

Recall that restriction is defined via a product of three monoids in §5.IV. The following lemma shows that restricting $\Xi_{n,d}$ can be done by only one of those monoids.

6.7 Lemma. *Let $F \in \mathrm{Mat}_n(A_d)$. We have $F \le \Xi_{n,d}$ if and only if there exists $T \in \mathrm{End}$ with $T(\Xi_{n,d}) = F$.* ◇

Proof.

$$\begin{aligned}(T_{\text{left}}, \mathrm{id}, \mathrm{id})\Xi_{n,d} &= T_{\text{left}}(X_n^{(1)} \cdots X_n^{(d)}) \\ &= (T_{\text{left}} X_n^{(1)}) X_n^{(2)} \cdots X_n^{(d)}.\end{aligned}$$

Note that

$$[T_{\text{left}} X_n^{(1)}]_{i,j} = \sum_a [T_{\text{left}}]_{i,a} [X_n^{(1)}]_{a,j} = \sum_a [T_{\text{left}}]_{i,a} x_{a,j}^{(1)}.$$

Define $T \in \mathrm{End}$ via $T(x_{i,j}^{(1)}) = \sum_a [T_{\text{left}}]_{i,a} x_{a,j}^{(1)}$ and $T(x_{i,j}^{(k)}) = x_{i,j}^{(k)}$ for $k > 1$, and observe that $T_{\text{left}} X_n^{(1)} = T(X_n^{(1)})$ and $X_n^{(k)} = T(X_n^{(k)})$ for all $k > 1$. Therefore,

$$(T_{\text{left}} X_n^{(1)}) X_2^{(d)} \cdots X_n^{(d)} = T(X_n^{(1)} \cdots X_n^{(d)})$$

and hence

$$(T_{\text{left}}, \mathrm{id}, \mathrm{id})\Xi_{n,d} = (\mathrm{id}, T, \mathrm{id})\Xi_{n,d}.$$

Analogously for T_{right}. □

For elements of A_d, the width can be defined via restrictions of $\Gamma_{n,d}$, as the following claim shows.

6.8 Claim. *For $f \in A_d$, we have that $w(f)$ is the smallest n such that $f \le \Gamma_{n,d}$.* ◇

Proof. We write ${\Gamma'}_{n,d} = [\Xi_{n,d}]_{1,1}$ and note that ${\Gamma'}_{n,d} \le \Gamma_{n,d}$ and $\Gamma_{n,d} \le {\Gamma'}_{n,d}$ by renaming variables, see (6.3). If $f \le \Gamma_{n,d}$, then clearly $f \le \Xi_{n,d}$. For the other direction, let $f \le \Xi_{n,d}$. According to Lemma 6.7 there exists $T \in \mathrm{End}$ with $T(\Xi_{n,d}) = f$, i.e., $[T(\Xi_{n,d})]_{1,1} = f$ and $[T(\Xi_{n,d})]_{a,b} = 0$ for all $(a,b) \ne (1,1)$. Since $[\Xi_{n,d}]_{1,1}$ does not involve any $x_{i,j}^{(1)}$ for $i \ne 1$, the element $[T(\Xi_{n,d})]_{1,1}$ is independent of $T(x_{i,j}^{(1)})$ for $i \ne 1$. Analogously, the element $[T(\Xi_{n,d})]_{1,1}$ is independent of $T(x_{i,j}^{(d)})$ for $j \ne 1$. Let $Z \in \mathrm{End}$ map all elements $x_{i,j}^{(1)}$ to 0 for $i \ne 1$, and map all elements $x_{i,j}^{(d)}$ to 0 for $j \ne 1$, and act as the identity on all other variables. Then $[TZ\Xi_{n,d}]_{1,1} = [T\Xi_{n,d}]_{1,1} = f$ and $[TZ\Xi_{n,d}]_{a,b} = [T\Xi_{n,d}]_{a,b} = 0$ for all $(a,b) \ne (1,1)$. But $Z\Xi_{n,d} = {\Gamma'}_{n,d}$, which implies $T{\Gamma'}_{n,d} = f$, hence $f \le {\Gamma'}_{n,d} \le \Gamma_{n,d}$. □

6.I Algebraic Branching Programs

Several different variants of algebraic branching programs (ABPs) exist, and we use a variant that allows the ABP to reflect the structure of our recurrence relation. An ABP is a directed acyclic graph in which each vertex has a layer from $0, \ldots, d$, and each edge is either going from a vertex in layer i to another vertex in layer i and is labeled by an element of R (these are called constant edges), or the edge is going from a vertex in layer i to a vertex in layer $i+1$ and is labeled by an element of A_1. A subset of the vertices in layer 0 are called *sources* $s_1, \ldots, s_a$ and a subset of the vertices in layer d is called *sinks* $t_1, \ldots, t_b$. To an ABP G we assign its computation result F, which is the matrix whose (a,b) entry is the weighted sum over all s_a-t_b-paths, where each summand is weighted by the product of its edge labels. We say that G computes F. The width of an ABP is defined as the maximum number of vertices in any layer, but in layer 0 and d we only count sources and sinks.

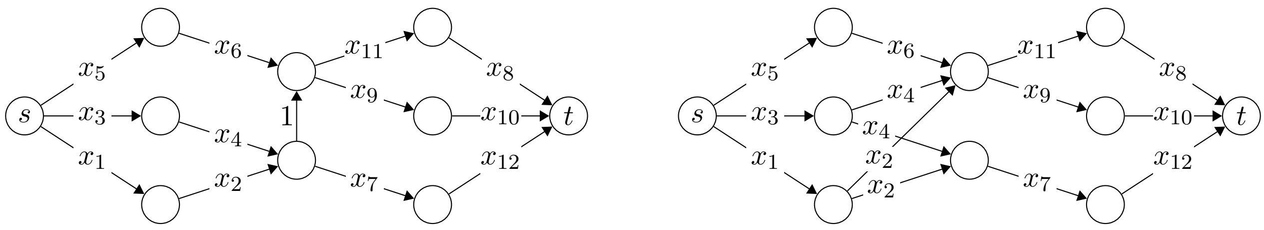

(a) An ABP with a constant edge. (b) An ABP with the constant edge removed.

Figure 1: Removal of constant edges from ABPs.

6.9 Claim. *For $F \in \mathrm{Mat}(A_d)$, the width of a smallest ABP computing F is $w(F)$.* $\diamond$

Proof. If $w(F) \leq n$, then $F \leq \Xi_{n,d}$, so there is $T \in \mathrm{End}$ with $T(\Xi_{n,d}) = F$. We take the width n ABP that computes $\Xi_{n,d}$ and apply T to each edge label to obtain a width n ABP that computes F. For the other direction, we are given a width n ABP that computes F. Figure 1 gives an illustration. If we have edges in an ABP *within a vertex layer*, then these edges must be constant edges, i.e., edges labeled by elements from R, and the set of these edges is not allowed to contain a directed cycle. We now use the following process. If there are any constant edges, then we pick a vertex v that has no incoming constant edges but an outgoing constant edge (v, u) with label c, we remove e and for each edge (w, v) we add an edge (w, u) with label $c \cdot \mathrm{label}(w, v)$ (or add this to the label of an already existing edge). The ABP still computes the same, but has fewer constant edges. We continue until all constant edges are gone. From the resulting ABP without constant edges we can readily read off a T with $T(\Xi_{n,d}) = F$. $\square$

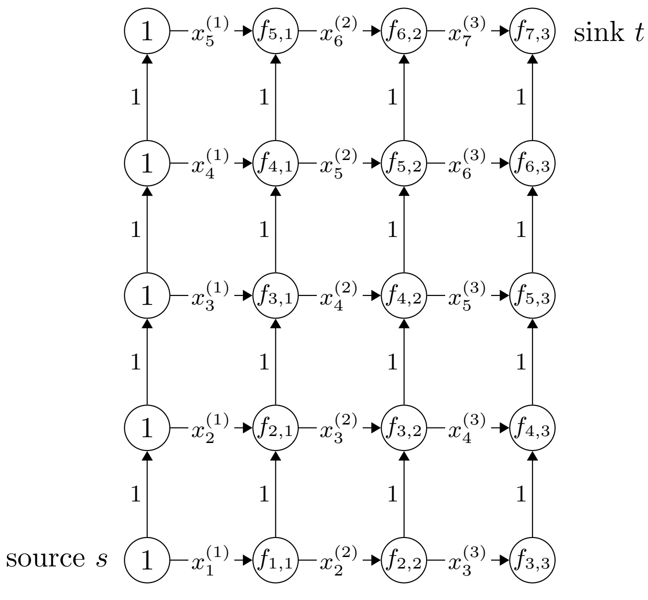

Figure 2: An ABP that computes $f_{7,3}$. The vertex label at each vertex v is the weighted sum of all s-v-paths.

6.10 Example. The elementary symmetric polynomial is defined as

$$e_{n,d} = \sum_{1 \leq i_1 < i_2 < \cdots < i_d \leq n} x_{i_1} \cdots x_{i_d}.$$

Its layered variant is defined as

$$f_{n,d} = \sum_{1 \leq i_1 < i_2 < \cdots < i_d \leq n} x^{(1)}_{i_1} \cdots x^{(d)}_{i_d}.$$

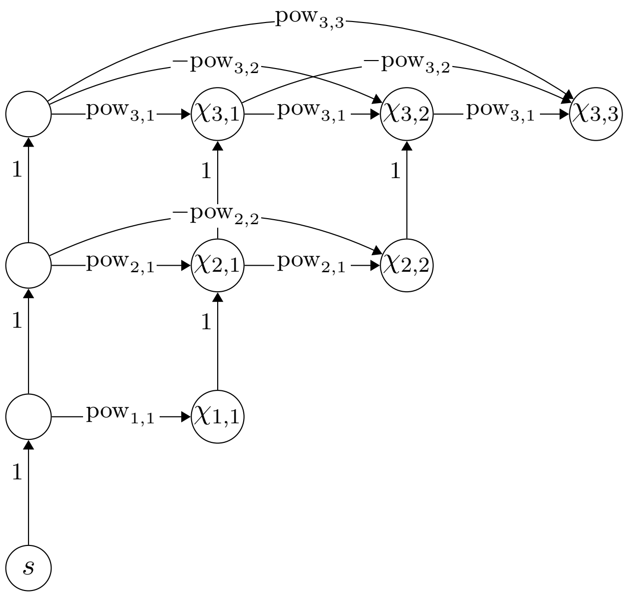

Figure 3: An ABP that computes $\det_3 = \chi_{3,3}$.

Clearly $e_{n,d} \leq f_{n,d}$ by mapping $x_i^{(k)}$ to x_i. We have

$$f_{n,d} = f_{n-1,d-1} \cdot x_n^{(d)} + f_{n-1,d}$$

and this recursion can be used directly to construct a width n ABP that computes $f_{n,d}$, see Figure 2, thus $w(f_{n,d}) \leq n$. ◇

We will use the technique in Example 6.10 several times in this paper to construct ABPs. Another important example is the following.

6.11 Proposition. *Let R be a commutative ring and let $A = \mathscr{S}$. We have $w(\det_d) \in \mathsf{poly}(d)$.* ◇

Proof. Let $X_n = (x_{i,j})_{\substack{1\leq i\leq n\\ 1\leq j\leq n}} \in \mathrm{Mat}_n(\mathscr{S}_1)$. Define

$$\chi_{n,d} \;=\; \sum_{S\in\binom{\{1,\ldots,n\}}{d}} \det([X_n]_{S,S}) \in \mathscr{S}_d, \tag{6.12}$$

where $[X_n]_{S,S}$ is the square submatrix of X_n with row indices S and column indices S. Let $\mathrm{pow}_{n,i} := [X_n^i]_{n,n}$, and note $w(\mathrm{pow}_{n,i}) \leq n$. We use [Ike25, eq. (2.8)]:

$$\chi_{n-1,d} \;=\; \sum_{i=0}^{d} (-1)^i\, \chi_{n,d-i}\, \mathrm{pow}_{n,i}. \tag{6.13}$$

We use that $\mathrm{pow}_{n,0} = 1$ to rewrite equation (6.13) in the form

$$\chi_{n,d} \;=\; \chi_{n-1,d} + \sum_{i=1}^{d} (-1)^{i+1}\, \chi_{n,d-i}\, \mathrm{pow}_{n,i}. \tag{6.14}$$

This recursion immediately gives an ABP that computes $\det_d$, see Figure 3. This shows $w(\det_d) \in O(d^4)$. There are many other proofs of $w(\det_d) \in \mathsf{poly}(d)$. □

We have

$$\mathrm{mon}_d \leq \det_d \tag{6.15}$$

via setting all off-diagonal variables to zero and renaming variables.

There is the following duality between the degree and the width with respect to subadditivity and submultiplicativity.

6.16 Claim. *Let R be a commutative semiring and $A \in \{\mathscr{T}, \mathscr{S}, \mathscr{W}\}$. Let $F, G \in A_d$. We have*

$$\begin{aligned} \deg(F+G) &\leq \max(\deg(F), \deg(G)) \\ \deg(FG) &\leq \deg(F) + \deg(G) \\ w(F+G) &\leq w(F) + w(G) \\ w(FG) &\leq \max(w(F), w(G)), \end{aligned}$$

where we use the conventions $\deg(0) = 0$ *and* $w(0) = 0$. ◇

Proof. The statements for the degree are obvious. The statements for the width are classical and can be proved by explicit constructions that appear for example in [Val79]. The width-of-sum statement is obtained by taking an ABP for F and an ABP for G and obtaining an ABP for the sum $F + G$ by identifying the respective sources with each other and identifying the respective sinks with each other. The width-of-product statement is obtained by taking an ABP for F and an ABP for G and obtaining an ABP for the product FG by identifying the sinks in F with the respective sources in G, where one adds dummy sources or sinks to make the numbers match. □

It is sometimes useful to convert from multiple sources and sinks to single source and single sink by adding a new source s in vertex layer -1 and a new sink t in vertex layer $d + 1$ and adding edges (s, s_i) with label e_i and edges (t_j, t) with label e_j^*.

7 Higher-order computation

In this section we define the Kronecker product $\boxtimes$. We start with the definition for $\mathscr{S}$ in characteristic zero in (7.3), and then turn to arbitrary commutative semirings in §7.II. In §7.III we then generalize the construction to $\mathscr{T}$ and $\mathscr{W}$. We define the hypercomputant in §7.IV, which in characteristic 0 is $\mathsf{h}\Gamma_{n,d} = \Gamma_{n,d}^{\boxtimes 2}$. We phrase Valiant's conjecture via the width $w(\mathsf{h}\Gamma_{n,d})$ in §7.V. We show that $\mathrm{HC}_d \leq \mathsf{h}\Gamma_{d,d}$ for the Hamiltonian cycle polynomial HC_d, and $\mathrm{Clique}_{n,k} \leq \mathsf{h}\Gamma_{n,\binom{k+1}{2}}$ for the Clique polynomial in §7.VI.

7.I Polarization and Restitution

Let R be a commutative semiring, and $A = \mathscr{S}$. Let

$$\mathrm{Sym}_d \ := \ \left(\mathscr{T}_d\right)^{\mathfrak{S}_d} = \{s \in \mathscr{T}_d \mid \forall \pi \in \mathfrak{S}_d : \pi s = s\}$$

denote the set of symmetric tensors. The set Sym_d is an R-semimodule, i.e., it is closed under addition and closed under multiplication with elements of R. The *polarization* map

$$\mathrm{polar} : \mathscr{S}_d \to \mathrm{Sym}_d$$

is defined via

$$\mathrm{polar}(x_{i_1} \cdots x_{i_d}) = \sum_{\pi \in \mathfrak{S}_d} x_{i_{\pi(1)}} \otimes \cdots \otimes x_{i_{\pi(d)}}$$

and R-linear extension. The restitution map $\mathrm{resti} : \mathscr{T}_d \to \mathscr{S}_d$ is the canonical map to the quotient, i.e., $\mathrm{resti}(x_{i_1} \otimes \cdots \otimes x_{i_d}) = x_{i_1} \cdots x_{i_d}$. One can see that $\mathrm{resti} \circ \mathrm{polar} = d! \cdot \mathrm{id}_{\mathscr{S}_d}$, which makes resti difficult to use in finite characteristic. We will treat these concerns in §7.II. The Kronecker product $\boxtimes$ is a bilinear map

$$\mathscr{T}_d \times \mathscr{T}_d \to \mathscr{T}_d$$

defined via

$$(x_{i_1} \otimes \cdots \otimes x_{i_d}) \boxtimes (x_{j_1} \otimes \cdots \otimes x_{j_d}) \;=\; x_{i_1,j_1} \otimes \cdots \otimes x_{i_d,j_d}, \tag{7.1}$$

and extended R-bilinearly.

7.2 Remark. The $\boxtimes$ operation is well-known in the setting of fast matrix multiplication, see e.g. [Blä13, p. 23] (where the symbol $\otimes$ is used instead of $\boxtimes$). If $M_n = \sum_{i,j,k=1}^{n} x_{i,j} \otimes x_{j,k} \otimes x_{k,i} \in \mathscr{T}_3$ denotes the $n \times n$ matrix multiplication tensor, then $M_n \boxtimes M_m = M_{nm}$. $\diamond$

For $s, s' \in \mathrm{Sym}_d$, we have $s \boxtimes s' \in \mathrm{Sym}_d$. Hence, for two polynomials $h, f \in \mathscr{S}_d$, if $\frac{1}{d!} \in R$, then we can construct a new polynomial

$$h \boxtimes f \;:=\; \tfrac{1}{d!}\,\mathrm{resti}(\mathrm{polar}(h) \boxtimes \mathrm{polar}(f)). \tag{7.3}$$

Several examples of polynomials of this form exist in the literature, see Example 7.10 below, but before discussing those, in §7.II we generalize (7.3), so that it works over arbitrary semirings.

7.II The Kronecker product of polynomials

Let R be a commutative semiring, and let $A = \mathscr{S}$. For $f, h \in A_d$ we define $f \boxtimes h \in A_d$ via

$$(x_{i_1} \cdots x_{i_d}) \boxtimes (x_{j_1} \cdots x_{j_d}) := \sum_{\pi \in \mathfrak{S}_d} x_{i_1, j_{\pi(1)}} \cdots x_{i_d, j_{\pi(d)}} \tag{7.4}$$

and bilinear extension. First note that this is a well-defined map $A_d \times A_d \to A_d$, because for any $\sigma, \tau \in \mathfrak{S}_d$ we have

$$\begin{aligned}
(x_{i_{\sigma(1)}} \cdots x_{i_{\sigma(d)}}) \boxtimes (x_{j_{\tau(1)}} \cdots x_{j_{\tau(d)}}) &= \sum_{\pi \in \mathfrak{S}_d} x_{i_{\sigma(1)}, j_{\tau(\pi(1))}} \cdots x_{i_{\sigma(d)}, j_{\tau(\pi(d))}} \\
&= \sum_{\pi \in \mathfrak{S}_d} x_{i_1, j_{\sigma^{-1}(\tau(\pi(1)))}} \cdots x_{i_d, j_{\sigma^{-1}(\tau(\pi(d)))}} \\
&\overset{\varrho := \sigma^{-1}\tau\pi}{=} \sum_{\varrho \in \mathfrak{S}_d} x_{i_1, j_{\varrho(1)}} \cdots x_{i_d, j_{\varrho(d)}} \;=\; (x_{i_1} \cdots x_{i_d}) \boxtimes (x_{j_1} \cdots x_{j_d}).
\end{aligned}$$

In fact, note that the same reshuffling argument shows that (7.4) is the same as

$$\sum_{\pi \in \mathfrak{S}_d} x_{i_{\pi(1)}, j_1} \cdots x_{i_{\pi(d)}, j_d}. \tag{7.5}$$

This readily shows the associativity of the Kronecker product:

$$\begin{aligned}
\big((x_{i_1} \cdots x_{i_d}) \boxtimes (x_{j_1} \cdots x_{j_d})\big) \boxtimes (x_{k_1} \cdots x_{k_d}) &= \sum_{\sigma \in \mathfrak{S}_d,\, \pi \in \mathfrak{S}_d} x_{i_{\sigma(1)}, j_1, k_{\pi(1)}} \cdots x_{i_{\sigma(d)}, j_d, k_{\pi(d)}} \\
&= (x_{i_1} \cdots x_{i_d}) \boxtimes \big((x_{j_1} \cdots x_{j_d}) \boxtimes (x_{k_1} \cdots x_{k_d})\big). \tag{7.6}
\end{aligned}$$

The equality of (7.4) and (7.5) also implies that

$$\text{both } f \boxtimes h \le h \boxtimes f \text{ and } h \boxtimes f \le f \boxtimes h \tag{7.7}$$

by renaming each $x_{i,j}$ to $x_{j,i}$.

7.8 Proposition. *If $f \le h$ and $f' \le h'$, then $f \boxtimes f' \le h \boxtimes h'$.* $\diamond$

Proof. For $T, T' \in \mathrm{End}$ we define $P \in \mathrm{End}$ as $P(x_{i,j}) := \sum_{i',j'} \alpha_{i',i}\alpha'_{j',j} x_{i',j'}$, where $\alpha_{i',i} \in R$ and $\alpha'_{j',j} \in R$ are defined via $T(x_i) = \sum_{i'} \alpha_{i',i} x_{i'}$ and $T'(x_j) = \sum_{j'} \alpha'_{j',j} x_{j'}$.

$$\begin{aligned}
T(h) \boxtimes T'(h') &= T(\sum_a \alpha_a x_{a,1} \cdots x_{a,d}) \boxtimes T'(\sum_b \alpha_b x_{b,1} \cdots x_{b,d}) \\
&= \sum_{a,b} \alpha_a \alpha_b \Big(T(x_{a,1} \cdots x_{a,d}) \boxtimes T'(x_{b,1} \cdots x_{b,d}) \Big) \\
&\overset{(*)}{=} \sum_{a,b} \alpha_a \alpha_b \Big(P((x_{a,1} \cdots x_{a,d}) \boxtimes (x_{b,1} \cdots x_{b,d})) \Big) \\
&= \Big(P((\sum_a \alpha_a x_{a,1} \cdots x_{a,d}) \boxtimes \sum_b \alpha_b (x_{b,1} \cdots x_{b,d})) \Big) \\
&= P(h \boxtimes h')
\end{aligned}$$

We verify $(*)$ in the following calculation, which finishes the proof.

$$\begin{aligned}
T(h) \boxtimes T'(h') &= T(x_{i_1} \cdots x_{i_d}) \boxtimes T'(x_{j_1} \cdots x_{j_d}) \\
&= \Big(\sum_{i'_1} \alpha_{i'_1,i_1} x_{i'_1} \Big) \cdots \Big(\sum_{i'_d} \alpha_{i'_d,i_d} x_{i'_d} \Big) \boxtimes \Big(\sum_{j'_1} \alpha_{j'_1,j_1} x_{j'_1} \Big) \cdots \Big(\sum_{j'_d} \alpha'_{j'_d,j_d} x_{j'_d} \Big) \\
&= \Big(\sum_{i'_1} \alpha_{i'_1,i_1} \Big) \cdots \Big(\sum_{i'_d} \alpha_{i'_d,i_d} \Big) \Big(\sum_{j'_1} \alpha'_{j'_1,j_1} \Big) \cdots \Big(\sum_{j'_d} \alpha'_{j'_d,j_d} \Big) x_{i'_1} \cdots x_{i'_d} \boxtimes x_{j'_1} \cdots x_{j'_d} \\
&= \sum_{\pi \in \mathfrak{S}_d} \Big(\sum_{\substack{i'_1,\dots,i'_d \\ j'_1,\dots,j'_d}} \alpha_{i'_1,i_1} \cdots \alpha_{i'_d,i_d} \cdot \alpha'_{j'_1,j_1} \cdots \alpha'_{j'_d,j_d} \Big) x_{i'_1,j'_{\pi(1)}} \cdots x_{i'_d,j'_{\pi(d)}} \\
&= \sum_{\pi \in \mathfrak{S}_d} \Big(\sum_{\substack{i'_1,\dots,i'_d \\ j'_1,\dots,j'_d}} \alpha_{i'_1,i_1} \cdots \alpha_{i'_d,i_d} \cdot \alpha'_{j'_1,j_{\pi^{-1}(1)}} \cdots \alpha'_{j'_d,j_{\pi^{-1}(d)}} \Big) x_{i'_1,j'_1} \cdots x_{i'_d,j'_d} \\
&= \sum_{\sigma \in \mathfrak{S}_d} \Big(\sum_{\substack{i'_1,\dots,i'_d \\ j'_1,\dots,j'_d}} \alpha_{i'_1,i_1} \cdots \alpha_{i'_d,i_d} \cdot \alpha'_{j'_1,j_{\sigma(1)}} \cdots \alpha'_{j'_d,j_{\sigma(d)}} \Big) x_{i'_1,j'_1} \cdots x_{i'_d,j'_d} \\
&= \sum_{\sigma \in \mathfrak{S}_d} \Big(\sum_{i'_1,j'_1} \alpha_{i'_1,i_1} \alpha'_{j'_1,j_{\sigma(1)}} x_{i'_1,j'_1} \Big) \cdots \Big(\sum_{i'_d,j'_d} \alpha_{i'_d,i_d} \alpha'_{j'_d,j_{\sigma(d)}} x_{i'_d,j'_d} \Big) \\
&= P\Big(\sum_{\sigma \in \mathfrak{S}_d} x_{i_1,j_{\sigma(1)}} \cdots x_{i_d,j_{\sigma(d)}} \Big) \\
&= P((x_{i_1} \cdots x_{i_d}) \boxtimes (x_{j_1} \cdots x_{j_d})) \\
&= P(h \boxtimes h').
\end{aligned}$$

□

For every commutative ring R we note that

$$\det_d \leq \det_d \boxtimes \det_d \tag{7.9}$$

as follows. We first recall $\mathrm{mon}_d \leq \det_d$ from (6.15). Using Proposition 7.8 we obtain $\mathrm{mon}_d \boxtimes \det_d \leq \det_d \boxtimes \det_d$, which is a polynomial in triple-indexed variables. We then map every $x_{i,j,k}$ to zero for which $i \neq j$, and we map $x_{i,j,k}$ to $x_{i,k}$ otherwise. The result is $\det_d$, which proves (7.9). In characteristic 0, we can alternatively prove (7.9) by restricting $\det_d$ to mon_d and further to x^d and then observe that $x^d \boxtimes \det_d = d! \det_d$.

7.10 Example. There are several examples of polynomials in the literature that are of the form $f \boxtimes h$, for example the following (whenever the determinant appears in the subsequent examples, we assume we are working over a ring R, not just a semiring).

1. The permanent polynomial: $\mathrm{per}_d := \mathrm{mon}_d \boxtimes \mathrm{mon}_d$, see e.g., [Val79], [Min84, eq. (3.1)].
2. The hyperpermanent (also called the multidimensional permanent): $\mathrm{mon}_d^{\boxtimes m}$, see e.g., [DG87, Tic15, CP16].
3. The rectangular permanent: $\mathrm{mon}_d \boxtimes e_{n,d}$, which is just called the permanent in [Min84, eq. (1.1)], see also [BE19, Exa 4.6].
4. The mixed discriminant: $\det_d \boxtimes \mathrm{mon}_d$, see e.g., [LC83, Pan85, Bap89, Gur05, CP16].
5. Cayley's first hyperdeterminant: $\det_d^{\boxtimes m}$, [Cay46, Bar95, HL13, CP16, BI17, AY23], also called the combinatorial hyperdeterminant, also called the Pascal determinant in [Lan11, §8.3] and [Lan15], or the quantum permanent in [Gur04, §3].

$\diamond$

7.11 Remark. Recall $\mathrm{mon}_d \leq \det_d$ from (6.15). Using Proposition 7.8 it follows $\mathrm{mon}_d \boxtimes \mathrm{mon}_d \leq \mathrm{mon}_d \boxtimes \det_d \leq \det_d \boxtimes \det_d$. These two restrictions are separate proofs in the literature, see [Gur04, Gur05]. Also note that this implies $w(\mathrm{mon}_d \boxtimes \mathrm{mon}_d) \leq w(\mathrm{mon}_d \boxtimes \det_d) \leq w(\det_d \boxtimes \det_d)$. We strengthen both results in §9.I. $\diamond$

7.III The Kronecker action

In this section we generalize the $\boxtimes$ construction from §7.II from $\mathscr{S}$ to $\mathscr{T}$ and $\mathscr{W}$. It is unclear if there is any reasonable Kronecker product construction for two elements in $\mathscr{W}$ with output in $\mathscr{W}$, so we take a different approach. The Kronecker product of two polynomials is replaced by a product of a symmetric tensor with an algebra element, i.e., an action of the algebra of order d symmetric tensors on $\mathscr{T}_d$ or $\mathscr{W}_d$ or $\mathscr{S}_d$. On $\mathscr{S}_d$ this gives the same construction as in §7.II, see Claim 7.17. We partially follow the exposition from [IOT25].

7.12 Claim. *Let R be a commutative semiring. Let $\varphi : \mathscr{T}_d \to \mathscr{T}_d$ be an R-linear $\mathfrak{S}_d$-equivariant map. Then φ induces an R-linear map $\psi_{\mathscr{S}} : \mathscr{S}_d \to \mathscr{S}_d$, $\psi_{\mathscr{S}}([f]_{\mathscr{S}}) = [\varphi(f)]_{\mathscr{S}}$. Moreover, φ induces an R-linear map $\psi_{\mathscr{W}} : \mathscr{W}_d \to \mathscr{W}_d$, $\psi_{\mathscr{W}}([f]_{\mathscr{W}}) = [\varphi(f)]_{\mathscr{W}}$.* $\diamond$

Proof. We verify the well-definedness of $\psi_{\mathscr{S}}$. Let $f, f' \in \mathscr{T}_d$ with $[f]_{\mathscr{S}} = [f']_{\mathscr{S}}$ be arbitrary, i.e., $f \sim f'$. The task is to show $\varphi(f) \sim \varphi(f')$. We first assume that (f, f') is a generator of $\sim$, i.e., we have $f = \pi f'$. Applying φ gives $\varphi(f) = \varphi(\pi f') = \pi\varphi(f')$, which is a generator of $\sim$, thus, $\varphi(f) \sim \varphi(f')$. More generally, consider the case where f and f' are not necessarily generators, but $f \sim f'$, $(f, f') = \sum_i \alpha_i(h_i, h_i')$, where $\alpha_i \in R$ and (h_i, h_i') are generators of $\sim$. Then $(\varphi(f), \varphi(f')) = \sum_i \alpha_i \underbrace{(\varphi(h_i), \varphi(h_i'))}_{\in\sim} \in\sim$. In other words, $\varphi(f) \sim \varphi(f')$.

We now verify the well-definedness of $\psi_{\mathscr{W}}$ in a similar way. Let $f, f' \in \mathscr{T}_d$ with $[f]_{\mathscr{W}} = [f']_{\mathscr{W}}$ be arbitrary, i.e., $f \backsim f'$. The task is to show $\varphi(f) \backsim \varphi(f')$. We have

$$f + h + n \otimes \ell \otimes \ell \otimes m = f' + h + n' \otimes \ell' \otimes \ell' \otimes m'.$$

Thus

$$\varphi(f) + \varphi(h) + \varphi(n \otimes \ell \otimes \ell \otimes m) = \varphi(f') + \varphi(h) + \varphi(n' \otimes \ell' \otimes \ell' \otimes m').$$

The tensor $u := n \otimes \ell \otimes \ell \otimes m$ is fixed by the transposition τ that swaps the two ℓ-positions. Hence $\varphi(u) = \varphi(\tau u) = \tau\varphi(u)$, thus $\varphi(u)$ is τ-invariant. Analogously for $n' \otimes \ell' \otimes \ell' \otimes m'$. Every τ-invariant tensor is $\backsim 0$, because those basis tensors that are fixed by τ have a repeated adjacent factor, and the other basis tensors $v = a \otimes w \otimes w' \otimes b$ come in orbit pairs $v + \tau v$ with $v + \tau v + a \otimes w \otimes w \otimes b + a \otimes w' \otimes w' \otimes b = a \otimes (w + w') \otimes (w + w') \otimes b \backsim 0$. This shows $\varphi(f) \backsim \varphi(f')$. $\square$

We have an R-bilinear map $\mathscr{T}_1 \times \mathscr{T}_1 \to \mathscr{T}_1$, $(v_1, v_2) \mapsto v_1 \boxtimes v_2$, where $v_1, v_2 \in \mathrm{Var}$, see §5.I. This lifts to an R-bilinear map $\boxtimes : \mathscr{T}_d \times \mathscr{T}_d \to \mathscr{T}_d$. Note that

$$\pi(f \boxtimes h) = (\pi f) \boxtimes (\pi h). \tag{7.13}$$

Let $\psi : \mathrm{Sym}_d \times \mathscr{T}_d \to \mathscr{T}_d$ be the restriction of this map to $\mathrm{Sym}_d \times \mathscr{T}_d$. Let $A \in \{\mathscr{T}, \mathscr{S}, \mathscr{W}\}$. Let $s \in \mathrm{Sym}_d$. Then there exists a linear map $\psi'_s : \mathscr{T}_d \to \mathscr{T}_d$, $f \mapsto \psi(s, f)$ that is $\mathfrak{S}_d$-equivariant: $\psi'_s(\pi f) = \psi(s, \pi f) = \psi(\pi s, \pi f) \overset{(7.13)}{=} \pi\psi(s, f) = \pi\psi'_s(f)$. We write $[f]$ for the image of $f \in \mathscr{T}_d$ in A_d. We apply Claim 7.12 (in the case $A_d = \mathscr{T}_d$ there is no need to invoke Claim 7.12) to obtain a map $\psi_s : A_d \to A_d$, $[f] \mapsto [\psi(s, f)]$. The well-defined map

$$\kappa : \mathrm{Sym}_d \to \mathrm{Hom}(A_d, A_d), \qquad s \mapsto \psi_s$$

is R-linear, because $(\kappa(s + \alpha t))([f]) = \psi(s + \alpha t, f) = \psi(s, f) + \alpha\psi(t, f) = \kappa(s)([f]) + \alpha\kappa(t)([f]) = (\kappa(s) + \alpha\kappa(t))([f])$. Altogether, this gives an R-bilinear map

$$\boxtimes : \mathrm{Sym}_d \times A_d \to A_d, \qquad (s, [f]) \mapsto \psi_s([f]) = [\psi(s, f)]. \tag{7.14}$$

We now interpret this over the standard basis. We treat this for variables $x_1, x_2, \ldots$, as the more general case is analogous, but notationally unpleasant. For a tensor $t \in \mathscr{T}_d$, define the stabilizer $\mathrm{stab}(t) := \{\pi \in \mathfrak{S}_d \mid \pi t = t\}$. Let $\mathfrak{S}_d/\mathrm{stab}(t) = \{\tau\,\mathrm{stab}(t) \mid \tau \in \mathfrak{S}_d\}$ denote the set of cosets, where $\tau\,\mathrm{stab}(t) = \{\tau\pi \mid \pi \in \mathrm{stab}(t)\}$. Then $\sigma t := \tau t$ is well-defined for $\sigma \in \mathfrak{S}_d/\mathrm{stab}(t)$ with coset representative τ. Let $t_{\mathrm{sym}} := \sum_{\sigma \in \mathfrak{S}_d/\mathrm{stab}(t)} \sigma t$, which is an element in Sym_d. The R-semimodule Sym_d of symmetric tensors is freely generated by the tensors of the form

$$(x_1^{\otimes\lambda_1} \otimes x_2^{\otimes\lambda_2} \otimes \cdots)_{\mathrm{sym}},$$

where $\sum_i \lambda_i = d$. For a list λ with $\sum_i \lambda_i = d$ let $\mathfrak{S}_d(1^{\lambda_1}2^{\lambda_2}\cdots)$ denote the orbit of the length d list

$$(\underbrace{1, 1, \ldots, 1}_{\lambda_1 \text{ many}}, \underbrace{2, 2, \ldots, 2}_{\lambda_2 \text{ many}}, \ldots), \tag{7.15}$$

i.e., the set of all length d lists that arise from (7.15) by permuting the positions of the numbers. For example, $\mathfrak{S}_4(1^2 2^2) = \{(1,1,2,2), (1,2,1,2), (1,2,2,1), (2,1,1,2), (2,1,2,1), (2,2,1,1)\}$. Let $A \in \{\mathscr{T}, \mathscr{S}, \mathscr{W}\}$. For a tensor $t \in \mathscr{T}_d$, let $[t]$ denote its equivalence class in A_d. For a symmetric tensor $s \in \mathrm{Sym}_d$ and an $f \in A_d$, (7.14) gives an interpretation on the standard basis as

$$(x_1^{\otimes\lambda_1} \otimes x_2^{\otimes\lambda_2} \otimes \cdots)_{\mathrm{sym}} \boxtimes [x_{i_1} \otimes \cdots \otimes x_{i_d}] \ := \sum_{(j_1,\ldots,j_d) \in \mathfrak{S}_d(1^{\lambda_1}2^{\lambda_2}\cdots)} [x_{j_1,i_1} \otimes \cdots \otimes x_{j_d,i_d}]. \tag{7.16}$$

7.17 Claim. *For two polynomials $h, f \in \mathscr{S}_d$, we have $h \boxtimes f = \mathrm{polar}(h) \boxtimes f$, where $h \boxtimes f$ is defined in* (7.4). ◇

Proof. It suffices to prove this for $h = [x_{i_1} \otimes \cdots \otimes x_{i_d}]_{\mathscr{S}}$ and $f = [x_{j_1} \otimes \cdots \otimes x_{j_d}]_{\mathscr{S}}$, because the statement then follows from bilinearity. Define λ via $\lambda_j = |\{i_k = j\}|$. We have $h = [x_1^{\otimes\lambda_1} \otimes x_2^{\otimes\lambda_2} \otimes \cdots]_{\mathscr{S}}$. We have

$$\mathrm{polar}(h) = \sum_{\pi \in \mathfrak{S}_d} x_{i_{\pi(1)}} \otimes \cdots \otimes x_{i_{\pi(d)}} = (\lambda_1!\,\lambda_2! \cdots)\,(x_1^{\otimes\lambda_1} \otimes x_2^{\otimes\lambda_2} \otimes \cdots)_{\mathrm{sym}}.$$

(7.16) gives

$$\mathrm{polar}(h) \boxtimes f = (\lambda_1!\,\lambda_2! \cdots) \sum_{(j_1,\ldots,j_d) \in \mathfrak{S}_d(1^{\lambda_1}2^{\lambda_2}\cdots)} [x_{j_1,i_1} \otimes \cdots \otimes x_{j_d,i_d}]_{\mathscr{S}},$$

which according to (7.5) coincides with $h \boxtimes f$. □

Note that there are interesting polynomials $s \boxtimes f$ for which $s \in \mathrm{Sym}_d$ is *not* a polarization. For example, the polynomial $(x^{\otimes i} y^{\otimes d-i})_{\mathrm{sym}} \boxtimes \det_d$ appears in [IOT25].

The following claim is the generalization of (7.6). It states that the semialgebra $(\mathrm{Sym}_d, +, \boxtimes)$ acts on A_d.

7.18 Claim. *Let R be a commutative semiring, and let $A \in \{\mathscr{T}, \mathscr{S}, \mathscr{W}\}$. Let $s', s \in \mathrm{Sym}_d$ and $f \in A_d$. Then $(s' \boxtimes s) \boxtimes f = s' \boxtimes (s \boxtimes f)$. Here, $s' \boxtimes s$ is defined as in* (7.1). ◇

Proof. Let $[f]$ denote the coset of $f \in \mathscr{T}_d$ in A_d. We prove the statement for $s' = (x_1^{\otimes \lambda_1} \otimes x_2^{\otimes \lambda_2} \otimes \cdots)_{\mathrm{sym}}$, $s = (x_1^{\otimes \mu_1} \otimes x_2^{\otimes \mu_2} \otimes \cdots)_{\mathrm{sym}}$, $f = [x_{i_1} \otimes \cdots \otimes x_{i_d}]$. We observe

$$s' \boxtimes s = \sum_{\substack{(j'_1,\ldots,j'_d) \in \mathfrak{S}_d(1^{\lambda_1} 2^{\lambda_2} \cdots) \\ (j_1,\ldots,j_d) \in \mathfrak{S}_d(1^{\mu_1} 2^{\mu_2} \cdots)}} x_{j'_1, j_1} \otimes \cdots \otimes x_{j'_d, j_d}.$$

According to (7.16) we have

$$(s' \boxtimes s) \boxtimes f = \sum_{\substack{(j'_1,\ldots,j'_d) \in \mathfrak{S}_d(1^{\lambda_1} 2^{\lambda_2} \cdots) \\ (j_1,\ldots,j_d) \in \mathfrak{S}_d(1^{\mu_1} 2^{\mu_2} \cdots)}} [x_{j'_1, j_1, i_1} \otimes \cdots \otimes x_{j'_d, j_d, i_d}].$$

On the other hand,

$$s \boxtimes f = \sum_{(j_1,\ldots,j_d) \in \mathfrak{S}_d(1^{\mu_1} 2^{\mu_2} \cdots)} [x_{j_1, i_1} \otimes \cdots \otimes x_{j_d, i_d}].$$

Moreover,

$$s' \boxtimes (s \boxtimes f) \;=\; \sum_{\substack{(j'_1,\ldots,j'_d) \in \mathfrak{S}_d(1^{\lambda_1} 2^{\lambda_2} \cdots) \\ (j_1,\ldots,j_d) \in \mathfrak{S}_d(1^{\mu_1} 2^{\mu_2} \cdots)}} [x_{j'_1, j_1, i_1} \otimes \cdots \otimes x_{j'_d, j_d, i_d}].$$

Hence, $(s' \boxtimes s) \boxtimes f = s' \boxtimes (s \boxtimes f)$. □

Note that for $s \in \mathrm{Sym}_d \subseteq \mathscr{T}_d$ and $T \in \mathrm{End}$ we have $T(s) \in \mathrm{Sym}_d$, in other words, End acts linearly on Sym_d. We write $s \le t$ if there exists $T \in \mathrm{End}$ such that $T(t) = s$. The next proposition is the generalized variant of Proposition 7.8.

7.19 Proposition. *Let $s, t \in \mathrm{Sym}_d$ and $f, h \in A_d$. If $s \le t$ and $f \le h$, then $s \boxtimes f \le t \boxtimes h$.* ◇

Proof. Without loss of generality, we consider s, t, f, h in variables $x_1, x_2, \ldots$, so that $s \boxtimes f$ and $t \boxtimes h$ are double-indexed. Given $E_1 \in \mathrm{End}$ with $E_1(t) = s$ and given $E_2 \in \mathrm{End}$ with $E_2(h) = f$, we create $E_3 \in \mathrm{End}$ via $E_3(x_{j,i}) = \sum_{j',i'} \alpha_{j',j}\, \beta_{i',i}\, x_{j',i'}$, where $E_1(x_j) = \sum_{j'} \alpha_{j',j}\, x_{j'}$ and $E_2(x_i) = \sum_{i'} \beta_{i',i}\, x_{i'}$. Let $t = \sum_a \iota_a\, x_{j_{1,a}} \otimes \cdots \otimes x_{j_{d,a}} \in \mathrm{Sym}_d$ be arbitrary. Then $E_1(t) = \sum_a \iota_a (\sum_{j'_1} \alpha_{j'_1, j_{1,a}} x_{j'_1}) \otimes \cdots \otimes (\sum_{j'_d} \alpha_{j'_d, j_{d,a}} x_{j'_d}) = \sum_a \iota_a \sum_{j'_1,\ldots,j'_d} (\alpha_{j'_1, j_{1,a}} \cdots \alpha_{j'_d, j_{d,a}})(x_{j'_1} \otimes \cdots \otimes x_{j'_d})$. Let $h = \sum_b \kappa_b [x_{i_{1,b}} \otimes \cdots \otimes x_{i_{d,b}}]$ be arbitrary. Then $E_2(h) = \sum_b \kappa_b [(\sum_{i'_1} \beta_{i'_1, i_{1,b}}\, x_{i'_1}) \otimes \cdots \otimes (\sum_{i'_d} \beta_{i'_d, i_{d,b}}\, x_{i'_d})] = \sum_b \kappa_b \sum_{i'_1,\ldots,i'_d} (\beta_{i'_1, i_{1,b}} \cdots \beta_{i'_d, i_{d,b}}) [x_{i'_1} \otimes \cdots \otimes x_{i'_d}]$. Moreover, $t \boxtimes h = \sum_{a,b} \iota_a \kappa_b [x_{j_{1,a}, i_{1,b}} \otimes \cdots \otimes x_{j_{d,a}, i_{d,b}}]$. We now calculate

$E_1(t) \boxtimes E_2(h) = \sum_{\substack{j'_1,\ldots,j'_d \\ i'_1,\ldots,i'_d}} \sum_{a,b} \iota_a \kappa_b (\alpha_{j'_1, j_{1,a}} \cdots \alpha_{j'_d, j_{d,a}})(\beta_{i'_1, i_{1,b}} \cdots \beta_{i'_d, i_{d,b}})[x_{j'_1, i'_1} \otimes \cdots \otimes x_{j'_d, i'_d}] = \sum_{a,b} \iota_a \kappa_b [(\sum_{j'_1, i'_1} \alpha_{j'_1, j_{1,a}}\, \beta_{i'_1, i_{1,b}}\, x_{j'_1, i'_1}) \otimes \cdots \otimes (\sum_{j'_d, i'_d} \alpha_{j'_d, j_{d,a}}\, \beta_{i'_d, i_{d,b}}\, x_{j'_d, i'_d})] = E_3(t \boxtimes h)$. □

7.IV Higher order computation

We use 0-based indexing of the rows and columns for the following discussion. Let R be a commutative semiring, let U, V, W be R-semimodules, and let $\varphi : U \times V \to W$ be R-bilinear. Let $E \in \mathrm{Mat}_\ell(U)$ and $F \in \mathrm{Mat}_m(V)$. The Kronecker product $E \boxtimes F \in \mathrm{Mat}_{\ell m}(W)$ with respect to φ is defined as

$$[E \boxtimes F]_{im+i', jm+j'} := \varphi([E]_{i,j}, [F]_{i',j'}). \tag{7.20}$$

In the following Claim 7.21, we lift Claim 7.18 to matrices.

7.21 Claim. *Let R be a commutative semiring, and let $A \in \{\mathscr{T}, \mathscr{S}, \mathscr{W}\}$. Let $S' \in \mathrm{Mat}_m(\mathrm{Sym}_d)$, $S \in \mathrm{Mat}_\ell(\mathrm{Sym}_d)$, $F \in \mathrm{Mat}_k(A_d)$. Then $(S' \boxtimes S) \boxtimes F = S' \boxtimes (S \boxtimes F)$, where the bilinear maps used for the two Kronecker products of matrices are defined in* (7.1) *and* (7.16). $\diamond$

Proof.

$$\begin{aligned}
[(S' \boxtimes S) \boxtimes F]_{k\ell i+ki'+i'',k\ell j+kj'+j''} &= [(S' \boxtimes S) \boxtimes F]_{k(\ell i+i')+i'',k(\ell j+j')+j''} \\
&= [S' \boxtimes S]_{\ell i+i',\ell j+j'} \boxtimes [F]_{i'',j''} \\
&= ([S']_{i,j} \boxtimes [S]_{i',j'}) \boxtimes [F]_{i'',j''} \\
&\overset{\text{Cla. 7.18}}{=} [S']_{i,j} \boxtimes ([S]_{i',j'} \boxtimes [F]_{i'',j''}) \\
&= [S']_{i,j} \boxtimes ([S \boxtimes F]_{ki'+i'',kj'+j''}) \\
&= [S' \boxtimes (S \boxtimes F)]_{k\ell i+(ki'+i''),k\ell j+(kj'+j'')} \\
&= [S' \boxtimes (S \boxtimes F)]_{k\ell i+ki'+i'',k\ell j+kj'+j''}
\end{aligned}$$

□

Let R be a commutative semiring, $A = \mathscr{S}$. We define polar : $\mathrm{Mat}(\mathscr{S}_d) \to \mathrm{Mat}(\mathrm{Sym}_d)$ entrywise, i.e., $[\mathrm{polar}(F)]_{i,j} := \mathrm{polar}(F_{i,j})$. We define the matrix of symmetric tensors

$$\mathsf{H}_{n,d} := \mathrm{polar}(\Xi_{n,d}) \ \in \mathrm{Mat}_n(\mathrm{Sym}_d) \tag{7.22}$$

and its bottom-right entry

$$H_{n,d} \ := \ \mathrm{polar}(\Gamma_{n,d}) \ \in \mathrm{Sym}_d. \tag{7.23}$$

We keep this specific matrix of symmetric tensors also when studying $A_d \in \{\mathscr{T}_d, \mathscr{W}_d\}$.

Now, let $A \in \{\mathscr{S}, \mathscr{T}, \mathscr{W}\}$. We define the *hypercomputant matrix* as

$$\mathsf{h}\Xi_{n,d} \ := \ \mathsf{H}_{n,d} \boxtimes \Xi_{n,d}$$

and its bottom-right entry $\mathsf{h}\Gamma_{n,d} \in A_d$ as

$$\mathsf{h}\Gamma_{n,d} \ := \ H_{n,d} \boxtimes \Gamma_{n,d}.$$

7.24 Remark. We have

$$w(\mathsf{h}\Gamma_{n,d}) \le w(\mathsf{h}\Xi_{n,d}) \le n^4 w(\mathsf{h}\Gamma_{n,d}), \tag{7.25}$$

which can be seen by constructing an ABP for each of the n^4 entries of $\mathsf{h}\Xi_{n,d}$ separately, because all entries of the matrix $\mathsf{h}\Xi_{n,d}$ are the same up to renaming variables. $\diamond$

7.26 Remark. Note that if $A = \mathscr{S}$, then $\mathsf{h}\Gamma_{n,d} = \Gamma_{n,d}^{\boxtimes 2}$. In $\mathscr{W}_d$ we cannot write $\Gamma_{n,d}^{\boxtimes 2}$, because there is no such multiplication. In $A_d = \mathscr{T}_d$ we have $\mathsf{h}\Gamma_{n,d} \neq \Gamma_{n,d}^{\boxtimes 2} = \Gamma_{n^2,d}$, which is a less interesting tensor for our purposes. $\diamond$

Combining (6.2), (7.16), and (7.20), we see that

$$[\mathsf{h}\Xi_{n,d}]_{an+a',bn+b'} \ = \sum_{\pi \in \mathfrak{S}_d} \sum_{\substack{1 \le i_1,\ldots,i_{d-1} \le n \\ 1 \le j_1,\ldots,j_{d-1} \le n \\ i_0=a,\ i_d=b}} (x^{(\pi(1))}_{i_{\pi(1)-1},i_{\pi(1)}} \boxtimes x^{(1)}_{a',j_1}) \cdot (x^{(\pi(2))}_{i_{\pi(2)-1},i_{\pi(2)}} \boxtimes x^{(2)}_{j_1,j_2}) \cdots (x^{(\pi(d))}_{i_{\pi(d)-1},i_{\pi(d)}} \boxtimes x^{(d)}_{j_{d-1},b'}) \tag{7.27}$$

We end this subsection with some observations.

7.28 Claim. $\Xi_{n^2,d} \le \mathsf{h}\Xi_{n,d}$ *and* $\Gamma_{n^2,d} \le \mathsf{h}\Gamma_{n,d}$. $\diamond$

Proof. We prove the claim for Γ, the proof for Ξ being completely analogous. For this proof, we let the variables $x_{i,j}^{(k)}$ in $\Gamma_{n,d}$ be indexed at zero, i.e., $0 \le i < n$ and $0 \le j < n$. Consider $\mathsf{h}\Gamma_{n,d}$ and map the variables $x_{i,j}^{(k)} \boxtimes x_{i',j'}^{(k')}$ to zero if $k \neq k'$. Then rename the variables $x_{i,j}^{(k)} \boxtimes x_{i',j'}^{(k)}$ to $x_{in+i',jn+j'}^{(k)}$ to obtain $\Gamma_{n^2,d}$. □

7.29 Claim. *For fixed $d \ge 2$ we have $w(\mathsf{h}\Xi_{n,d}) = \Omega(n^2)$ and $w(\mathsf{h}\Gamma_{n,d}) = \Omega(n^2)$, with the exception of $A_2 = \mathscr{W}_2$.* ◇

Proof. This is analogous to the proof of Example 6.6. We start with the observation that for $\Xi_{n,d}$, $\mathsf{h}\Xi_{n,d}$, $\Gamma_{n,d}$, and $\mathsf{h}\Gamma_{n,d}$ the number of essential variables equals the number of variables used to define them, which is seen as follows. For every one of the n^4d^2 variables $x_{\mathbf{i}}$ we can reconstruct $\mathbf{i}$ from any monomial in $\partial_{\mathbf{i}}\mathsf{h}\Xi_{n,d}$. Hence $J(\mathsf{h}\Xi_{n,d})$ contains a $(n^4d^2) \times (n^4d^2)$ permutation matrix as a submatrix, so the claim follows with Lemma 5.7. The proofs for $\Xi_{n,d}$, $\Gamma_{n,d}$, and $\mathsf{h}\Gamma_{n,d}$ are analogous.

The number of essential variables of $\Xi_{w,d}$ is w^2d, and the number of essential variables of $\mathsf{h}\Xi_{n,d}$ is n^4d^2. If $w(\mathsf{h}\Xi_{n,d}) \le w$, then $\mathsf{h}\Xi_{n,d} \le \Xi_{w,d}$, and hence by (5.6) we have $n^4d^2 \le w^2d$, thus $w \ge n^2d^{1/2}$.

For the second part we use Claim 6.8. The number of essential variables of $\Gamma_{w,d}$ is $2w + (d-2)w^2$, and the number of essential variables of $\mathsf{h}\Gamma_{n,d}$ is $(2n+(d-2)n^2)^2$. If $w(\mathsf{h}\Gamma_{n,d}) \le w$, then by Claim 6.8 we have $\mathsf{h}\Gamma_{n,d} \le \Gamma_{w,d}$, hence by (5.5) we have $(2n+(d-2)n^2)^2 \le 2w+(d-2)w^2$. Therefore, for fixed d we have $w \in \Omega(n^2)$. □

7.30 Example. Let $A = \mathscr{S}$. We use Proposition 7.8 to show that the examples in Example 7.10 are restrictions of $\mathsf{h}\Gamma_{n,d}$ as follows.

- The permanent: $\mathrm{mon}_d \boxtimes \mathrm{mon}_d = \mathsf{h}\Gamma_{1,d}$
- The hyperpermanent: $\mathrm{mon}_d^{\boxtimes m} = \Gamma_{1,d}^{\boxtimes m}$
- The rectangular permanent: $\mathrm{mon}_d \boxtimes e_{n,d} \le \mathsf{h}\Gamma_{n,d}$, because $e_{n,d} \le \Gamma_{n,d}$, see Example 6.10
- The mixed discriminant: $\det_d \boxtimes \mathrm{mon}_d \le \mathsf{h}\Gamma_{w(\det_d),d}$.
- Cayley's first hyperdeterminant: $\det_d^{\boxtimes m} \le \Gamma_{w(\det_d),d}^{\boxtimes m}$.

We will see more examples in §7.VI and in §9. ◇

7.V Valiant's conjecture

Valiant's original conjecture "VBP $\neq$ VNP" (over characteristic different from 2, this is also the *determinant vs permanent* conjecture) is equivalent to the following conjecture.

7.31 Conjecture (Valiant's conjecture)**.** Let R be a field and let $A = \mathscr{S}$. Valiant's conjecture for R is

$$w(\mathsf{h}\Xi_{n,d}) \notin \mathsf{poly}(n,d).$$ ■

By (7.25), Conjecture 7.31 is equivalent to $w(\mathsf{h}\Gamma_{n,d}) \notin \mathsf{poly}(n,d)$. Note that Conjecture 7.31 depends on R and A and the question can also be asked when R is not a field and/or when $A \neq \mathscr{S}$. The case $R = \mathbb{C}$ is discussed in §10. The statement $w(\mathsf{h}\Gamma_{n,d}) \notin \mathsf{poly}(n,d)$ is true for $A = \mathscr{T}$ when R is an arbitrary commutative semiring, see §8.I. The statement $w(\mathsf{h}\Gamma_{n,d}) \notin \mathsf{poly}(n,d)$ is also true for $A = \mathscr{S}$ when R is a zerosumfree commutative semiring, for example $R = \mathbb{R}_{\ge 0}$, see §8.II. So far there are no results involving $A = \mathscr{W}$. The following proposition shows that it is sufficient to study second-order computation only and not also higher orders.

7.32 Proposition. *$w(\mathsf{h}\Xi_{n,d}) \in \mathsf{poly}(n,d)$ if and only if for every m we have $w(\mathsf{H}_{n,d}^{\boxtimes m} \boxtimes \Xi_{n,d}) \in \mathsf{poly}(n,d)$.* ◇

Proof. One direction is clear. For the other direction, fix any m and consider $\mathsf{H}_{n,d}^{\boxtimes m} \boxtimes \Xi_{n,d}$. Let $k \in \mathbb{N}$ such that $\forall d, n : w(\mathsf{h}\Xi_{n,d}) \leq n^k d^k$, in other words $\mathsf{h}\Xi_{n,d} \leq \Xi_{n^k d^k,d}$. We will mark with $(*)$ when we use $\mathsf{H}_{n,d} \leq \mathsf{H}_{n',d}$ for $n' \geq n$, which is obtained analogously to Claim 6.4.

$$\begin{aligned}
\mathsf{H}_{n,d}^{\boxtimes m} \boxtimes \Xi_{n,d} &\overset{\text{Cla. 7.18}}{=} \mathsf{H}_{n,d}^{\boxtimes m-1} \boxtimes (\mathsf{H}_{n,d} \boxtimes \Xi_{n,d}) \overset{\text{Pro. 7.19}}{\leq} \mathsf{H}_{n,d}^{\boxtimes m-1} \boxtimes \Xi_{n^k d^k,d} \\
&\overset{\text{Cla. 7.18}}{=} \mathsf{H}_{n,d}^{\boxtimes m-2} \boxtimes (\mathsf{H}_{n,d} \boxtimes \Xi_{n^k d^k,d}) \overset{(*)}{\leq} \mathsf{H}_{n,d}^{\boxtimes m-2} \boxtimes (\mathsf{H}_{n^k d^k,d} \boxtimes \Xi_{n^k d^k,d}) \\
&\overset{\text{Pro. 7.19}}{\leq} \mathsf{H}_{n,d}^{\boxtimes m-2} \boxtimes \Xi_{n^{k^2} d^{(k+1)^2-1},d} \leq \mathsf{H}_{n,d}^{\boxtimes m-3} \boxtimes \Xi_{n^{k^3} d^{(k+1)^3-1},d} \\
&\leq \cdots \leq \Xi_{n^{k^m} d^{(k+1)^m-1},d}\,.
\end{aligned}$$

Since k and m are constant, this means $w(\mathsf{H}_{n,d}^{\boxtimes m} \boxtimes \Xi_{n,d}) \in \mathsf{poly}(n, d)$. We remark that for $A = \mathscr{S}$ a better construction would split $\Xi_{n,d}^{\boxtimes m} = \Xi_{n,d}^{\boxtimes m/2} \boxtimes \Xi_{n,d}^{\boxtimes m/2}$ recursively, but this does not work for $A \in \{\mathscr{T}, \mathscr{W}\}$. □

If Conjecture 7.31 is false, then from Example 7.30 it follows that all examples in Example 7.10 have small width. Moreover, all examples in the following §7.VI then also have small width. More examples are given in §10.

7.VI Hamiltonian cycles and cliques

In this section we prove that the Hamiltonian cycle polynomial and the Clique polynomial are restrictions of $\mathsf{h}\Gamma_{n,d}$. A general construction that gives the polynomials only up to a factor of low degree is presented in §9.

Hamiltonian cycles

Let R be a semiring, and let $A \in \{\mathscr{T}, \mathscr{S}, \mathscr{W}\}$. We define the Hamiltonian cycle algebra element, see [Val79] where the variant without superscripts is discussed:

$$\mathrm{HC}_d \ := \sum_{\phi \in \mathfrak{S}_{d-1}} [x^{(1)}_{d,\phi(1)} \otimes x^{(2)}_{\phi(1),\phi(2)} \otimes \cdots \otimes x^{(d)}_{\phi(d-1),d}] \ \in A_d.$$

Note $H_{1,d} = \mathrm{polar}(\mathrm{mon}_d)$ up to renaming of variables.

7.33 Proposition. $\mathrm{HC}_d \ \leq \ H_{1,d} \boxtimes \Gamma_{d,d} \ \leq \ \mathsf{h}\Gamma_{d,d}$. ◇

Proof. The second restriction is obvious. We treat the first restriction. We start with writing $\Gamma_{d,d} \in A_d$ as

$$\Gamma_{d,d} \ = \sum_{\phi:[d-1]\to[d]} [x^{(1)}_{d,\phi(1)} \otimes x^{(2)}_{\phi(1),\phi(2)} \otimes x^{(3)}_{\phi(2),\phi(3)} \otimes \cdots \otimes x^{(d)}_{\phi(d-1),d}].$$

Hence,

$$\mathrm{polar}(\mathrm{mon}_d) \boxtimes \Gamma_{d,d} = \sum_{\substack{\phi:[d-1]\to[d] \\ \pi \in \mathfrak{S}_d}} [x^{(1)}_{\pi(1),d,\phi(1)} \otimes x^{(2)}_{\pi(2),\phi(1),\phi(2)} \otimes \cdots \otimes x^{(d)}_{\pi(d),\phi(d-1),d}].$$

Let $T \in \mathrm{End}$ as follows: $T(x^{(k)}_{a,b,c}) = x^{(k)}_{b,c}$ if $a = c$, and $T(x^{(k)}_{a,b,c}) = 0$ otherwise. We calculate

$$T(\mathrm{polar}(\mathrm{mon}_d) \boxtimes \Gamma_{d,d}) = \sum_{\substack{\phi \in \mathfrak{S}_d \\ \phi(d)=d}} [x^{(1)}_{d,\phi(1)} \otimes x^{(2)}_{\phi(1),\phi(2)} \otimes \cdots \otimes x^{(d)}_{\phi(d-1),d}] \ = \ \mathrm{HC}_d.$$ □

The statement "$w(\mathrm{HC}_d) \notin \mathsf{poly}(d)$ over $\mathscr{S}$ for a commutative semiring R" is equivalent to Conjecture 7.31, and this equivalence was proved in [Val79], where it is stated only for fields R. The corresponding Hamiltonian cycle decision problem is problem 9 of Karp's original 21 NP-complete problems [Kar72].

Cliques

Let R be a commutative semiring, and let $A = \mathscr{S}$. One can find different variants of the so-called Clique polynomial in the literature, for example the variant in [Sch76] is uncolored without loops, the variant in [BE19] is uncolored with loops, the variant in [KPR23] is colored without loops. We use a colored variant with loops:

$$ {}_k\mathrm{Clique}_n \;:=\; \sum_{1\le v_1<v_2<\cdots<v_k\le n}\;\prod_{1\le i\le j\le k} x_{v_i,i,j,v_j} \;\in\; \mathscr{S}_{\binom{k+1}{2}}. \tag{7.34}$$

For convenience, we also use the following notation.

$$ \mathrm{Cliq}_{n,d} := \begin{cases} {}_k\mathrm{Clique}_n & \text{if } d = \binom{k+1}{2},\\ 0 & \text{otherwise.}\end{cases}$$

7.35 Proposition. $\mathrm{Cliq}_{n,d} \le \mathsf{h}\Gamma_{n,d}$. ◇

Proof. Let $d := \binom{k+1}{2}$, because otherwise there is nothing to prove. For additional clarity, we use the auxiliary variable names $x_{a,(i,j)} := x_{a,i,j}$ and $x_{(i,j),b} := x_{i,j,b}$, and we write $x_{a,(i,j),(i',j'),b} := x_{a,(i,j)} \boxtimes x_{(i',j'),b}$. The variable $x_{a,(i,j),(i,j),b}$ will later be mapped to the variable $x_{a,i,j,b}$ in the clique polynomial. We show that

$$ {}_k\mathrm{Clique}_n \;\le\; \underbrace{\left(\sum_{1\le a_1,\ldots,a_k\le n}\;\prod_{1\le i\le j\le k} x_{a_i,(i,j)}\right)}_{=:\,h_{n,d}} \boxtimes \underbrace{\left(\sum_{1\le b_1<\cdots<b_k\le n}\;\prod_{1\le i\le j\le k} x_{(i,j),b_j}\right)}_{=:\,f_{n,d}\;=\;{}_kf_n} \;\le\; \mathsf{h}\Gamma_{n,d} \tag{7.36}$$

in a two-step process (note the strict inequality between the b_j, whereas the a_i can be in any order). We first observe that

$$ h_{n,d} \;=\; \prod_{i=1}^{k}\sum_{a_i=1}^{n}\prod_{i\le j\le k} x_{a_i,(i,j)},$$

which immediately shows $w(h_{n,d}) \le n$, see Claim 6.16. Moreover,

$$ {}_kf_n = \left(\prod_{i=1}^{k} x_{(i,k),n}\right)\cdot {}_{k-1}f_{n-1} \;+\; {}_kf_{n-1},$$

which shows that $w(f_{n,d}) \le n$. These two width bounds together prove the right inequality of (7.36), see Proposition 7.8. We now prove the left inequality of (7.36). Let $T \in \mathrm{End}$ with $T(x_{a,(i,j),(i',j'),b}) = 0$ if $i \ne i'$ or $j \ne j'$, and $T(x_{a,(i,j),(i,j),b}) = x_{a,(i,j),(i,j),b}$. We expand

$$ T\left(\left(\sum_{1\le a_1,\ldots,a_k\le n}\;\prod_{1\le i\le j\le k} x_{a_i,(i,j)}\right)\boxtimes\left(\sum_{1\le b_1<\cdots<b_k\le n}\;\prod_{1\le i\le j\le k} x_{(i,j),b_j}\right)\right)$$

$$ = \sum_{\substack{1\le a_1,\ldots,a_k\le n\\ 1\le b_1<\cdots<b_k\le n}}\;\prod_{1\le i\le j\le k} x_{a_i,(i,j),(i,j),b_j}$$

Note that every monomial involves a variable $x_{a_i,(i,i),(i,i),b_i}$ for each i. Therefore, if we map every variable $x_{b,(i,i),(i,i),b'}$ to zero for which $b \ne b'$, then for each monomial that is not mapped to zero we have $\forall i : a_i = b_i$, so that all variables are of the form $x_{b_i,(i,j),(i,j),b_j}$. The resulting polynomial is

$$ \sum_{1\le b_1<\cdots<b_k\le n}\;\prod_{1\le i\le j\le k} x_{b_i,(i,j),(i,j),b_j}.$$

Renaming $x_{b,(a,c),(a,c),b'}$ to $x_{b,a,c,b'}$ gives the desired ${}_k\mathrm{Clique}_n$. □

The clique polynomial was the first polynomial family that was proved VW[1]-complete, see [BE19]. The corresponding clique problem is problem 3 of Karp's original 21 NP-complete problems [Kar72], and its parameterized version is W[1]-complete, see e.g. [DF13, Thm 21.2.5].

8 Exponents

In this section we analyze the growth of $w(\mathsf{h}\Xi_{n,d})$ and $w(\mathsf{h}\Gamma_{n,d})$ for fixed d and growing n.

8.1 Definition. We define

$$\begin{aligned}\xi_d &:= \limsup_{n\to\infty}(\log_n(w(\mathsf{h}\Xi_{n,d}))) = \inf\{r \mid w(\mathsf{h}\Xi_{n,d}) \in O(n^r)\},\\ \gamma_d &:= \limsup_{n\to\infty}(\log_n(w(\mathsf{h}\Gamma_{n,d}))) = \inf\{r \mid w(\mathsf{h}\Gamma_{n,d}) \in O(n^r)\}.\end{aligned}$$

These definitions depend on R and A. ◇

By (7.25) we have

$$\gamma_d \le \xi_d \le 4 + \gamma_d. \tag{8.2}$$

For algebra elements in A_d whose variables have two superscripts, for each monomial the superscript can be listed as a matrix of nonnegative integers called the superscript matrix, whose entry in position (i,j) is the number of times a variable with superscript $(i),(j)$ appears in the monomial. Note that all superscript matrices for $\mathsf{h}\Xi_{n,d}$ are permutation matrices corresponding to $\pi \in \mathfrak{S}_d$. For a fixed $\pi \in \mathfrak{S}_d$, the result of setting all monomials with superscript matrix $\neq \pi$ to zero is denoted by $\mathsf{h}\Xi_{n,d,\pi}$. By construction, we have

$$\mathsf{h}\Xi_{n,d} = \sum_{\pi\in\mathfrak{S}_d} \mathsf{h}\Xi_{n,d,\pi} \tag{8.3}$$

Since also $\mathsf{h}\Xi_{n,d,\pi} \le \mathsf{h}\Xi_{n,d}$, we have

$$\xi_d = \max\{\limsup_{n\to\infty}\{\log_n(w(\mathsf{h}\Xi_{n,d,\pi}))\} \mid \pi \in \mathfrak{S}_d\}. \tag{8.4}$$

As usual, the lower right entries are denoted by $\mathsf{h}\Gamma_{n,d,\pi}$, $\mathsf{h}\Gamma_{n,d} = \sum_{\pi\in\mathfrak{S}_d}\mathsf{h}\Gamma_{n,d,\pi}$, $\gamma_d = \max\{\limsup_{n\to\infty}\{\log_n(w(\mathsf{h}\Gamma_{n,d,\pi}))\} \mid \pi \in \mathfrak{S}_d\}$.

8.5 Remark. Over $R = \mathbb{C}$, the torus $T = (\mathbb{C}\setminus\{0\})^{d\times d}$ acts on $\mathbb{C}^{n^4d^2}$ via $\mathrm{diag}(\alpha_{1,1},\ldots,\alpha_{d,d})x_{\mathbf{i}}^{(k)} \boxtimes x_{\mathbf{j}}^{(\ell)} = \alpha_{k,\ell}\,(x_{\mathbf{i}}^{(k)} \boxtimes x_{\mathbf{j}}^{(\ell)})$. The irreducible polynomial representations of T are indexed by $d\times d$ matrices of nonnegative integers. The linear action of T lifts to A_d and $\mathrm{Mat}_n(A_d)$ in the usual way: $t[v_1\otimes\cdots\otimes v_d] = [tv_1\otimes\cdots\otimes tv_d]$ and entrywise for matrices. We have that $\mathsf{h}\Xi_{n,d,\pi}$ is the image of the projection of $\mathsf{h}\Xi_{n,d}$ to the π-isotypic component. ◇

Note that Claim 7.29 implies that

$$\xi_d \ge 2 \text{ and } \gamma_d \ge 2. \tag{8.6}$$

8.7 Claim. $w(\mathsf{h}\Xi_{n,d,\pi}) \ge n^2$. ◇

Proof. The number of essential variables of $\mathsf{h}\Xi_{n,d,\pi}$ is n^4d, thus if $w(\mathsf{h}\Xi_{n,d,\pi}) \le w$, then $\mathsf{h}\Xi_{n,d,\pi} \le \Xi_{w,d}$, and hence by (5.6) we have $n^4d \le w^2d$, thus $w \ge n^2$. □

Since $\mathsf{h}\Xi_{n,d,\mathrm{id}} = \Xi_{n^2,d}$ up to renaming variables, Claim 8.7 is the best lower bound we can get for $w(\mathsf{h}\Xi_{n,d,\pi})$ if we do not take any information about π into account.

8.I Tensors and flattenings

Given $d,j\in\mathbb{N}$, $0\le j\le d$, there is a canonical linear isomorphism $\mathscr{T}_d \to (\mathscr{T}_j)\otimes(\mathscr{T}_{d-j})$. Since we have a chosen basis and dual basis, we obtain a map $\mathcal{F}_j : \mathscr{T}_d \to \mathscr{T}_j \otimes (\mathscr{T}_{d-j})^*$, which is a map into a space of matrices. For a matrix $F \in \mathrm{Mat}_m(\mathscr{T}_d)$ we write $F = \sum_{a,b} e_a \otimes [F]_{a,b} \otimes e_b^*$ with standard basis vectors $e_a \in R^m$, $e_b \in R^m$. Every linear map $\mathcal{F} : \mathscr{T}_d \to \mathrm{Mat}_m(\mathscr{T}_d)$ lifts to a linear map $\mathcal{F} : \mathrm{Mat}_\ell(\mathscr{T}_d) \to \mathrm{Mat}_{\ell m}(\mathscr{T}_d)$ via $\mathcal{F}(F) = \sum_{a,b} e_a \otimes \mathcal{F}([F]_{a,b}) \otimes e_b^*$, i.e., matrix entries are replaced by matrices, forming a block matrix. Let rk denote the decomposition rank function, i.e., for $M \in U\otimes U'$ we have $\mathrm{rk}(M) = \min\{r \mid \exists l_i \in U, \ell_i \in U' : M = \sum_{i=1}^r l_i \otimes \ell_i\}$.

8.8 Theorem ([Nis91] for matrices)**.** *Let $F \in \mathrm{Mat}(\mathscr{T}_d)$. Then $w(F) \geq \max\{\mathrm{rk}\mathcal{F}_j(F) \mid 0 \leq j \leq d\}$. If R is a field, then we have equality.* ◇

Proof. We follow the argument in [Nis91] and generalize it to matrices. For a sequence of variables $p = (p_1, \ldots, p_k)$ we write $p_\otimes := p_1 \otimes p_2 \otimes \cdots \otimes p_k$. Let G be an ABP of width $w(F)$ that computes F. Let $V_j(G)$ denote the set of vertices of G in vertex layer j.

For every vertex pair $v, w \in G$ we define $(v{\rightsquigarrow}w) := \sum_{p \in P(G,v,w)} p_\otimes$. Note that

$$[F]_{a,b} = \sum_{v \in V_j(G)} (s_a{\rightsquigarrow}v) \otimes (v{\rightsquigarrow}t_b)\ .$$

In other words,

$$F = \sum_{a,b} \sum_{v \in V_j(G)} e_a \otimes (s_a{\rightsquigarrow}v) \otimes (v{\rightsquigarrow}t_b) \otimes e_b^* = \sum_{v \in V_j(G)} \underbrace{\sum_{a,b} e_a \otimes (s_a{\rightsquigarrow}v) \otimes (v{\rightsquigarrow}t_b) \otimes e_b^*}_{=:({\rightsquigarrow}v{\rightsquigarrow})}\ . \tag{8.9}$$

We apply $\mathcal{F}_j$ and analyze the rank:

$$\begin{aligned}
\mathcal{F}_j(({\rightsquigarrow}v{\rightsquigarrow})) &= \sum_{a,b} e_a \otimes \mathcal{F}_j((s_a{\rightsquigarrow}v) \otimes (v{\rightsquigarrow}t_b)) \otimes e_b^* = \sum_{a,b} (e_a \otimes (s_a{\rightsquigarrow}v) \otimes (v{\rightsquigarrow}t_b)^* \otimes e_b^*) \\
&= \big(\sum_a e_a \otimes (s_a{\rightsquigarrow}v)\big) \otimes \big(\sum_b (v{\rightsquigarrow}t_b)^* \otimes e_b^*\big),
\end{aligned}$$

which is a rank ≤ 1 matrix. By (8.9) and the definition of rank as the decomposition rank $\mathrm{rk}(\mathcal{F}_j(F)) \leq \#V_j(G)$, which implies the first part of the theorem.

For the second part of the theorem, let $R = \mathbb{F}$ be a field. For the sake of contradiction, assume that $w(F) > \max\{\mathrm{rk}\mathcal{F}_j(F) \mid 0 \leq j \leq d\}$. Among all ABPs of width $w(F)$ that compute F, take one that maximizes j for which both $\#V_j(G) = w(F)$ and $\forall i < j : \#V_i(G) < w(F)$. In particular $\#V_j(G) > \mathrm{rk}\mathcal{F}_j(F)$. We will now see that there exists an ABP G' of width $\leq w(F)$ that computes F with $\forall i \leq j : \#V_i(G') < \#V_i(G)$, which is a contradiction to the choice of j.

We have

$$F = \sum_{v \in V_j(G)} ({\rightsquigarrow}v{\rightsquigarrow}) = \sum_{v \in V_j(G)} \big(\sum_a e_a \otimes (s_a{\rightsquigarrow}v)\big) \otimes \big(\sum_b (v{\rightsquigarrow}t_b)^* \otimes e_b^*\big).$$

If $\{\sum_a e_a \otimes (s_a{\rightsquigarrow}v) \mid v \in V_j(G)\}$ are linearly independent, and also $\{\sum_b (v{\rightsquigarrow}t_b)^* \otimes e_b^*\}$ are linearly independent, then $\mathrm{rk}(\mathcal{F}_j(F)) = \#V_j(G)$, which contradicts the assumption. W.l.o.g. for one vertex $\tilde{v} \in V_j(G)$ we have a linear combination

$$\sum_a e_a \otimes (s_a{\rightsquigarrow}\tilde{v}) = \sum_{v \in V_j(G) \setminus \{\tilde{v}\}} \alpha_v \sum_a e_a \otimes (s_a{\rightsquigarrow}v).$$

$$
\begin{aligned}
F &= \sum_{v\in V_j(G)} \big(\sum_a e_a \otimes (s_a{\rightsquigarrow}v)\big) \otimes \big(\sum_b (v{\rightsquigarrow}t_b)^* \otimes e_b^*\big)\\
&= \big(\sum_a e_a \otimes (s_a{\rightsquigarrow}\tilde v)\big) \otimes \big(\sum_b (\tilde v{\rightsquigarrow}t_b)^* \otimes e_b^*\big) + \sum_{v\in V_j(G)\setminus\{\tilde v\}} \big(\sum_a e_a \otimes (s_a{\rightsquigarrow}v)\big) \otimes \big(\sum_b (v{\rightsquigarrow}t_b)^* \otimes e_b^*\big)\\
&= \big(\sum_{v\in V_j(G)\setminus\{\tilde v\}} \alpha_v \sum_a e_a \otimes (s_a{\rightsquigarrow}v)\big) \otimes \big(\sum_b (\tilde v{\rightsquigarrow}t_b)^* \otimes e_b^*\big) + \sum_{v\in V_j(G)\setminus\{\tilde v\}} \big(\sum_a e_a \otimes (s_a{\rightsquigarrow}v)\big) \otimes \big(\sum_b (v{\rightsquigarrow}t_b)^* \otimes e\\
&= \sum_{v\in V_j(G)\setminus\{\tilde v\}} \big(\sum_a e_a \otimes (s_a{\rightsquigarrow}v)\big) \otimes \big(\alpha_v \sum_b (\tilde v{\rightsquigarrow}t_b)^* \otimes e_b^* + \sum_b (v{\rightsquigarrow}t_b)^* \otimes e_b^*\big)\\
&= \sum_{v\in V_j(G)\setminus\{\tilde v\}} \big(\sum_a e_a \otimes (s_a{\rightsquigarrow}v)\big) \otimes \big(\sum_b (\alpha_v(\tilde v{\rightsquigarrow}t_b)^* + (v{\rightsquigarrow}t_b)^*) \otimes e_b^*\big)\\
&= \sum_{v\in V_j(G)\setminus\{\tilde v\}} \big(\sum_a e_a \otimes (s_a\overset{G'}{\rightsquigarrow}v)\big) \otimes \big(\sum_b (v\overset{G'}{\rightsquigarrow}t_b)^* \otimes e_b^*\big),
\end{aligned}
$$

where G' arises from G by adding $\alpha_v \cdot \mathsf{label}((\tilde v, w))$ to the label of the edge (v, w), $v \neq \tilde v$, for every $w \in V_{j+1}(G)$, and deleting the vertex $\tilde v$. The argument for $\tilde v{\rightsquigarrow}t_b$ is analogous. The ABP G' computes F, has the same number of vertices in each vertex layer but j, where it has one less vertex. This gives the desired contradiction to the choice of j. $\square$

For $\pi \in \mathfrak{S}_d$, $k \in \{0, \dots, d\}$, let

$$
\begin{aligned}
\beta(\pi, k) \;:=\; & \{0, \pi(1)-1, \pi(1), \pi(2)-1, \pi(2), \dots, \pi(k)-1, \pi(k)\}\\
& \cap \{\pi(k+1)-1, \pi(k+1), \pi(k+2)-1, \pi(k+2), \dots, \pi(d)-1, \pi(d), d\}. \qquad (8.10)
\end{aligned}
$$

Let $\beta(\pi) := \max\{\#\beta(\pi, k) \mid 0 \le k \le d\}$.

8.11 Proposition. *Let R be a commutative semiring, and let $A = \mathscr{T}$. We have $w(\mathsf{h}\Xi_{n,d,\pi}) = n^{\beta(\pi)+1}$.* $\diamond$

Proof. To prove the lower bound, we use Theorem 8.8. We have

$$
\mathsf{h}\Xi_{n,d,\pi} = \sum_{\substack{i_0,i_1,\dots,i_d\\ j_0,j_1,\dots,j_d}} e_{i_0,j_0} \otimes (x^{(\pi(1))}_{i_{\pi(1)-1},i_{\pi(1)}} \boxtimes x^{(1)}_{j_0,j_1}) \otimes \dots \otimes (x^{(\pi(d))}_{i_{\pi(d)-1},i_{\pi(d)}} \boxtimes x^{(d)}_{j_{d-1},j_d}) \otimes e_{i_d,j_d} \qquad (8.12)
$$

The monomials on the left factor of $\mathcal{F}_k(\mathsf{h}\Xi_{n,d,\pi})$ are

$$
e_{i_0,j_0} \otimes (x^{(\pi(1))}_{i_{\pi(1)-1},i_{\pi(1)}} \boxtimes x^{(1)}_{j_0,j_1}) \otimes \dots \otimes (x^{(\pi(k))}_{i_{\pi(k)-1},i_{\pi(k)}} \boxtimes x^{(k)}_{j_{k-1},j_k})
$$

The monomials on the right factor of $\mathcal{F}_k(\mathsf{h}\Xi_{n,d,\pi})$ are

$$
(x^{(\pi(k+1))}_{i_{\pi(k+1)-1},i_{\pi(k+1)}} \boxtimes x^{(k+1)}_{j_k,j_{k+1}}) \otimes \dots \otimes (x^{(\pi(d))}_{i_{\pi(d)-1},i_{\pi(d)}} \boxtimes x^{(d)}_{j_{d-1},j_d}) \otimes e_{i_d,j_d}
$$

The set of o for which i_o appears on the left-hand side of the tensor is $\{0, \pi(1)-1, \pi(1), \pi(2)-1, \pi(2), \dots, \pi(k)-1, \pi(k)\}$. The set of o for which i_o appears on the right-hand side of the tensor is $\{\pi(k+1)-1, \pi(k+1), \pi(k+2)-1, \pi(k+2), \dots, \pi(d)-1, \pi(d), d\}$. We say that $o \in \{0, \dots, d\}$ is a *boundary point* if i_o appears on both sides of the tensor, otherwise o is called an *inner point*. We denote by $\beta(\pi, k)$ the set of boundary points. One monomial from the left-hand side tensored with a monomial from the right-hand side gives a monomial of $\mathsf{h}\Xi_{n,d,\pi}$ if and only if they agree on j_k and on every i_o for boundary points o. Now we consider the set M of those monomials for which $i_o = 1$ for inner points o, and $j_o = 1$ for all $o \neq k$. For each monomial in M from the left-hand side there is exactly one monomial in M from the right-hand side so that their tensor product gives a monomial in $\mathsf{h}\Xi_{n,d,\pi}$. This means that if we restrict the matrix $\mathcal{F}_k(\mathsf{h}\Xi_{n,d,\pi})$ to the rows and columns of those monomials, we get a permutation matrix. The rank of this matrix is full, see Lemma 5.7. The number of such monomials from the left-hand side equals the number of monomials from the right-hand side, and it is $n^{\#\beta(\pi,k)+1}$, which proves the lower bound. We now prove the upper bound by adapting a construction of Grenet [Gre11]. Let B_k denote the set of boundary points and let I_k be the set of inner points of the left-hand side. For every map $C_k : B_k \to \{1,\ldots,n\}$ and a number $c_k \in \{1,\ldots,n\}$ we define the vector

$$F_{k,C_k,c_k} := \sum_{\substack{\forall o\in I_k: i_o\in\{1,\ldots,n\} \\ j_0,j_1,\ldots,j_{k-1}\in\{1,\ldots,n\} \\ \forall o\in B_k: i_o=C_k(o) \\ j_k=c_k}} e_{i_0,j_0} \otimes (x^{(\pi(1))}_{i_{\pi(1)-1},i_{\pi(1)}} \boxtimes x^{(1)}_{j_0,j_1}) \otimes \cdots \otimes (x^{(\pi(k))}_{i_{\pi(k)-1},i_{\pi(k)}} \boxtimes x^{(k)}_{j_{k-1},j_k})$$

which is the sum of the subset of monomials of the left-hand side of $\mathcal{F}_k(\mathsf{h}\Xi_{n,d,\pi})$ that satisfy $i_o = C_k(o)$ for all boundary points o, and $j_k = c_k$. The ABP we construct has one vertex for each of these $n^{\#\beta(\pi,k)+1}$ many tensors in degree k at which the ABP computes the vector F_{k,C_k,c_k}. Edges with appropriate labels can be chosen to make this ABP compute $\mathsf{h}\Xi_{n,d,\pi}$, because when running k from 0 to d, every position o transitions at most once (exactly once for all $o \neq d$) from boundary to inner, and never in the other direction. A transition from boundary to inner is handled by summation, for example, if $\pi(k+1) = o$ and $o-1$ transitions from boundary to inner and o is a new boundary point, then for any C_{k+1} and c_{k+1} we have

$$F_{k+1,C_{k+1},c_{k+1}} = \sum_{i,j=1}^{n} F_{k,C_{k+1}\setminus\{o\}\cup\{o-1\mapsto i\},j} \otimes \left(x^{(\pi(k+1))}_{i,C_{k+1}(o)} \boxtimes x^{(k+1)}_{j,c_{k+1}}\right),$$

where $C_{k+1} \setminus \{o\} \cup \{o-1 \mapsto i\}$ arises from the map C_{k+1} by making it undefined at o and setting $C_{k+1}(o-1) = i$. $\square$

Let $\beta(d) := \max\{\beta(\pi) \mid \pi \in \mathfrak{S}_d\}$. By (8.4) we have that for $A_d = \mathscr{T}_d$ it holds $\xi_d = \beta(d)+1$. Note that the upper bound $w(\mathsf{h}\Xi_{n,d,\pi}) \leq n^{\beta(\pi)+1}$ also holds for $A_d \in \{\mathscr{S}_d, \mathscr{W}_d\}$, because these are quotients of $\mathscr{T}_d$.

Let $\beta'(\pi,k) = \beta(\pi,k) \setminus \{0,d\}$ and $\beta'(\pi) = \max\{\#\beta'(\pi,k) \mid 0 \leq k \leq d\}$, and $\beta'(d) := \max\{\beta'(\pi) \mid \pi \in \mathfrak{S}_d\}$. The same arguments as for $\mathsf{h}\Xi_{n,d,\pi}$ show that for $A_d = \mathscr{T}_d$ we have $w(\mathsf{h}\Gamma_{n,d,\pi}) = n^{\beta'(\pi)+1}$ and for $A_d \in \{\mathscr{S}_d, \mathscr{W}_d\}$ we have $w(\mathsf{h}\Gamma_{n,d,\pi}) \leq n^{\beta'(\pi)+1}$.

8.13 Lemma. *We have* $\beta(d) = \begin{cases} d & \text{if } d \text{ is odd,} \\ d+1 & \text{if } d \text{ is even.} \end{cases}$
And we have $\beta'(d) = d-1$. $\diamond$

Proof. We consider $\beta(d)$ first. For any π with $\pi(1) = 2, \pi(2) = 4, \pi(3) = 6, \ldots, \pi(\lfloor d/2\rfloor) = 2\lfloor d/2\rfloor$ we see that $\beta(\pi) \geq 1 + 2\lfloor d/2\rfloor$, which is $d+1$ for even d, and d for odd d. Clearly $d+1$ is an upper bound for $\beta(d)$, which finishes the case of d being even. Let d be odd. We phrase $\beta(d)$ as an edge coloring process of the graph on $d+1$ vertices v_i, where v_i is adjacent to v_{i-1} and v_{i+1}, and v_0 has an additional half-edge, and v_d has an additional half-edge. Initially, all edges are colored blue, the half-edge at v_0 is colored red and the half-edge at v_d is colored blue. We now color edges (not half-edges) from blue to red step by step. $\#\beta(\pi,k)$ is the number of vertices that are adjacent to a blue and a red (half-)edge after k recolorings. To choose π so that $\beta(\pi)$ is maximized, it is not beneficial to ever choose an edge that has a red neighbor edge. The same also holds for coloring the edge (v_0, v_1). It is optimal to color the edges (v_1, v_2), (v_3, v_4), and so on. Since d is odd, we can only do this $\frac{d-1}{2}$ many times, after which the number of boundary points will go down again. Starting with 1 boundary point v_0, each such recoloring increases the number of boundary points by 2, hence $\beta(d) = 1 + 2\frac{d-1}{2} = d$.

For β', when counting boundary points, we do not count v_0 and v_d, and we can ignore the half-edges. For d odd, coloring (v_1, v_2), (v_3, v_4), and so on, gives $\beta'(d) \geq 2\frac{d-1}{2} = d-1$. For d even, coloring

(v_1, v_2), (v_3, v_4), and so on, gives $\beta'(d) \geq 2(\frac{d}{2} - 1) + 1 = d - 1$, since v_d is not being counted. The upper bound comes from a similar coloring consideration. Coloring (v_0, v_1) or (v_{d-1}, v_d) can increase the number of boundary points for subsequent rounds by 1, but it comes at the cost of increasing the number of boundary points only by 1 instead of 2. Hence, effectively, the number of boundary points can only be increased by 2 each round. We make a case distinction based on the color of (v_0, v_1) and (v_{d-1}, v_d) in the coloring with maximized number of boundary points that count. If (v_0, v_1) is red and (v_{d-1}, v_d) is red, then we get at most $2 + 2\lfloor \frac{d-3}{2} \rfloor \leq d - 1$ many boundary points that count. If (v_0, v_1) is red and (v_{d-1}, v_d) is blue, then we get at most $1 + 2\lfloor \frac{d-2}{2} \rfloor \leq d - 1$ many boundary points that count. Analogously for (v_0, v_1) blue and (v_{d-1}, v_d) red. If (v_0, v_1) is blue and (v_{d-1}, v_d) is blue, then we get at most $2\lceil \frac{d-2}{2} \rceil \leq d - 1$ many boundary points that count. Hence, in all cases we see $\beta'(d) \leq d - 1$. ☐

8.II Optimal speedups for polynomials if R is zerosumfree

A commutative semiring is called *zerosumfree* if for all $r_1, r_2 \in R$ we have $r_1 + r_2 = 0$ implies $r_1 = r_2 = 0$. For $\pi \in \mathfrak{S}_d$, $\sigma \in \mathfrak{S}_d$, $k \in \{0, \ldots, d\}$, we define

$$\begin{aligned} \beta(\pi, \sigma, k) \;&:=\; \{1\} \times \Big(\{0, \pi(\sigma(1)) - 1, \pi(\sigma(1)), \pi(\sigma(2)) - 1, \pi(\sigma(2)), \ldots, \pi(\sigma(k)) - 1, \pi(\sigma(k))\} \\ &\qquad \cap \{\pi(\sigma(k+1)) - 1, \pi(\sigma(k+1)), \pi(\sigma(k+2)) - 1, \pi(\sigma(k+2)), \ldots, \pi(\sigma(d)) - 1, \pi(\sigma(d)), d\}\Big) \\ &\cup \;\; \{2\} \times \Big(\{0, \sigma(1) - 1, \sigma(1), \sigma(2) - 1, \sigma(2), \ldots, \sigma(k) - 1, \sigma(k)\} \\ &\qquad \cap \{\sigma(k+1) - 1, \sigma(k+1), \sigma(k+2) - 1, \sigma(k+2), \ldots, \sigma(d) - 1, \sigma(d), d\}\Big) \end{aligned} \tag{8.14}$$

As a remark, if one compares to (8.10), one observes that $\#\beta(\pi, \mathrm{id}, k) = \#\beta(\pi, k) + 1$. Now, let $\beta(\pi, \sigma) := \max\{\#\beta(\pi, \sigma, k) \mid 0 \leq k \leq d\}$. Let $\beta_{\mathrm{opt}}(\pi) := \min\{\beta(\pi, \sigma) \mid \sigma \in \mathfrak{S}_d\}$. Let $\beta_{\mathrm{opt}}(d) := \max\{\beta_{\mathrm{opt}}(\pi) \mid \pi \in \mathfrak{S}_d\}$.

8.15 Proposition. *For a commutative semiring R and $A = \mathscr{S}$ we have $w(\mathsf{h}\Xi_{n,d,\pi}) \leq n^{\beta_{\mathrm{opt}}(\pi)}$.* ◇

Proof. The upper bound uses essentially the same construction as in the proof of Proposition 8.11. The only difference is that we are free to choose $\sigma \in \mathfrak{S}_d$ instead of having $\sigma = \mathrm{id}$, and after fixing σ we can write

$$\mathsf{h}\Xi_{n,d,\pi} = \sum_{\substack{i_0, i_1, \ldots, i_d \\ j_0, j_1, \ldots, j_d}} e_{i_0, j_0} \otimes (x^{(\pi(\sigma(1)))}_{i_{\pi(\sigma(1))-1}, i_{\pi(\sigma(1))}} \boxtimes x^{(\sigma(1))}_{j_{\sigma(1)-1}, j_{\sigma(1)}}) \cdot \cdots \cdot (x^{(\pi(\sigma(d)))}_{i_{\pi(\sigma(d))-1}, i_{\pi(\sigma(d))}} \boxtimes x^{(\sigma(d))}_{j_{\sigma(d)-1}, j_{\sigma(d)}}) \otimes e_{i_d, j_d}. \tag{8.16}$$

Now $\beta(\pi, \sigma, k)$ is the set of all boundary points, so we have $n^{\#\beta(\pi, \sigma, k)}$ many vertices in the ABP in vertex layer k. ☐

8.17 Remark. Proposition 8.15 also works for $A = \mathscr{W}$ if R is a commutative ring. If R is a commutative semiring, then we can only allow even permutations σ in the definition of $\beta_{\mathrm{opt}}(\pi)$, which increases its value by at most 4, because it is easy to see that $\forall \pi, \sigma : \beta(\pi, (1\ 2)\sigma) \leq 4 + \beta(\pi, \sigma)$. ◇

Note that better constructions are possible for $A = \mathscr{S}$ by not treating every π individually, but by (8.4) this cannot give a better exponent.

8.18 Proposition. *For a zerosumfree commutative semiring R and $A = \mathscr{S}$ we have $n^{\beta_{\mathrm{opt}}(\pi)}/d \leq w(\mathsf{h}\Xi_{n,d,\pi})$.* ◇

Proof. The proof technique is standard, see [KPR23] and references therein. Let G be an ABP that computes $\mathsf{h}\Xi$. Edges with label zero are removed. We call the pair (k, k') the *edge layer pair* of the variable $x^{(k)}_{i,j} \boxtimes x^{(k')}_{i',j'}$. Let $\sigma_p := \big(\sigma(1), \ldots, \sigma(d)\big)$ be the list of the second indices of the edge layer pairs of variables on p read from source to sink. If p and q have a common vertex v in some vertex layer k, then

$$\{\sigma_p(1),\ldots,\sigma_p(k)\} = \{\sigma_q(1),\ldots,\sigma_q(k)\} \text{ and } \{\sigma_p(k+1),\ldots,\sigma_p(d)\} = \{\sigma_q(k+1),\ldots,\sigma_q(d)\}, \quad (8.19)$$

because otherwise the path that starts at a source, goes via p until v, and continues from v via q to a sink would contribute to the coefficient of a monomial that is not a monomial of $\mathsf{h}\Xi$. We define analogously to (8.14)

$$\begin{aligned}\beta(v) \;:=\; & \{1\} \times \Big(\{0, \pi(\sigma(1)) - 1, \pi(\sigma(1)), \pi(\sigma(2)) - 1, \pi(\sigma(2)), \ldots, \pi(\sigma(k)) - 1, \pi(\sigma(k))\} \\ & \cap \{\pi(\sigma(k+1)) - 1, \pi(\sigma(k+1)), \pi(\sigma(k+2)) - 1, \pi(\sigma(k+2)), \ldots, \pi(\sigma(d)) - 1, \pi(\sigma(d)), d\}\Big) \\ \cup \; & \{2\} \times \Big(\{0, \sigma(1) - 1, \sigma(1), \sigma(2) - 1, \sigma(2), \ldots, \sigma(k) - 1, \sigma(k)\} \\ & \cap \{\sigma(k+1) - 1, \sigma(k+1), \sigma(k+2) - 1, \sigma(k+2), \ldots, \sigma(d) - 1, \sigma(d), d\}\Big),\end{aligned}$$

where σ can be chosen arbitrarily from $\{\sigma_p \mid v \in p\}$, and $\beta(v)$ is independent of this choice because of (8.19). We call $\beta(v)$ the set of *boundary points* of v. To every vertex v we assign its *bag*, which is a function $\mathrm{bag}(v) : \beta(v) \to \{1, \ldots, n\}$ as follows. We take a direct source-sink path p that uses v and a variable $x^{(k)}_{i,j} \boxtimes x^{(k')}_{i',j'}$ on p such that $(1, k-1) \in \beta(v)$, and we set $\mathrm{bag}(v)((1, k-1)) = i$. Analogously, for $(1, k) \in \beta(v)$ we set $\mathrm{bag}(v)((1, k)) = j$, and for $(2, k'-1) \in \beta(v)$ we set $\mathrm{bag}(v)((2, k'-1)) = i'$, and for $(2, k') \in \beta(v)$ we set $\mathrm{bag}(v)((2, k')) = j'$. This setting is consistent, because of the structure of the monomials of the entries of $\mathsf{h}\Xi$. Moreover, this completely defines $\mathrm{bag}(v)$. The definition of $\mathrm{bag}(v)$ does not depend on the choice of p, because for p and q through v with different bags for v (but still, (8.19) holds) we see that there is also the path that starts at a source and goes via p until it reaches v and continues via q to a sink, and that path would not correspond to a monomial in $\mathsf{h}\Xi_{n,d,\pi}$. The bag size of v is defined as $\#\beta(v)$. If v is the vertex in vertex layer k on a direct source-sink path p, then $\#\beta(v) = \#\beta(\pi, \sigma_p, k)$. By definition of $\beta(\pi, \sigma_p)$ we have that at least one of the vertices has bag size $\beta(\pi, \sigma_p)$. Note that $\beta_{\mathrm{opt}}(\pi) \leq \beta(\pi, \sigma_p)$, therefore for every direct source-sink path p at least one of the vertices has bag size at least $\beta_{\mathrm{opt}}(\pi)$.

We assign a *position* in $\{1, 2\} \times \{0, \ldots, d\}$ to every subscript of a variable $x^{(k)}_{i,j} \boxtimes x^{(k')}_{i',j'}$ by saying that the position of i is $(1, k-1)$, the position of j is $(1, k)$, the position of i' is $(2, k'-1)$, and the position of j' is $(2, k')$. There are $2d + 2$ many positions, and the choice of a value from $\{1, \ldots, n\}$ for each position corresponds bijectively to a monomial in $\mathsf{h}\Xi_{n,d,\pi}$. By definition, in every monomial of $\mathsf{h}\Xi_{n,d,\pi}$, all indices that have the same position have the same value. The bag of a vertex v fixes $\#\beta(v)$ many values of positions, hence there can be at most $n^{2d+2-\#\beta(v)}$ many monomials whose computation gets a contribution from direct source-sink paths that use v.

Every direct source-sink path has at least one vertex with $\#\beta(v) \geq \beta_{\mathrm{opt}}$, therefore this vertex is only involved in the computation of at most $n^{2d+2-\beta_{\mathrm{opt}}}$ many monomials. Since we have n^{2d+2} many monomials, we have at least $n^{2d+2}/n^{2d+2-\beta_{\mathrm{opt}}} = n^{\beta_{\mathrm{opt}}}$ many vertices, thus G has width at least $n^{\beta_{\mathrm{opt}}}/d$. □

For zerosumfree R and $A = \mathscr{S}$, combining Proposition 8.15 and Proposition 8.18 gives $\beta_{\mathrm{opt}}(d) = \xi_d$. In analogy to $\beta_{\mathrm{opt}}(d)$ we define $\beta'(\pi, \sigma, k) := \beta(\pi, \sigma, k) \setminus \{(1, 0), (1, d), (2, 0), (2, d)\}$, and $\beta'(\pi, \sigma) := \max\{\#\beta'(\pi, \sigma, k) \mid 0 \leq k \leq d\}$, and $\beta'_{\mathrm{opt}}(\pi) := \min\{\beta'(\pi, \sigma) \mid \sigma \in \mathfrak{S}_d\}$, and $\beta'_{\mathrm{opt}}(d) := \max\{\beta'_{\mathrm{opt}}(\pi) \mid \pi \in \mathfrak{S}_d\}$. By construction, $\#\beta(\pi, \sigma, k) \leq 4 + \#\beta'(\pi, \sigma, k) \leq 4 + \#\beta(\pi, \sigma, k)$. We see $\beta'_{\mathrm{opt}}(d) = \gamma_d$ in the same way as $\beta_{\mathrm{opt}}(d) = \xi_d$. By brute force enumeration of (π, σ, k) we obtain the following values.

d	2	3	4	5	6	7
ξ_d for $A = \mathscr{T}$	4	4	6	6	8	8
ξ_d for zerosumfree R and $A = \mathscr{S}$	4	4	6	6	6	8
γ_d for $A = \mathscr{T}$	2	3	4	5	6	7
γ_d for zerosumfree R and $A = \mathscr{S}$	2	3	4	4	5	6

Note the difference between ξ_6 for $A = \mathscr{T}$ and ξ_6 for R zerosumfree and $A = \mathscr{S}$. The next proposition shows that such differences cannot be large.

8.20 Proposition. *For zerosumfree commutative semirings R and $A = \mathscr{S}$ we have*

$$\beta_{\text{opt}}(d) = \Omega(d). \qquad \diamond$$

Proof. This follows from [AR08, Thm 3.4] and the fact that $\beta'_{\text{opt}}(\pi)$ is the same as the pathwidth of a permutation graph for π, which is called $C(\pi)$ in [AR08, Def 2.6]. $\square$

8.III Monotonicity of the exponents

In this subsection we prove that the sequences $(\xi_d)_{d\in\mathbb{N}}$ and $(\gamma_d)_{d\in\mathbb{N}}$ grow monotonically, where for $A = \mathscr{W}$ we require that R is a commutative ring.

An affine algebraic branching program (aABP) is defined in the same way as an ABP with the only difference that we allow affine R-linear combinations of variables as edge labels. We also do not require all source-sink paths to have the same length. The size of an aABP is the number of its vertices, not counting sources or sinks. We write $\mathsf{ac}(F)$ for the smallest size of an aABP computing F. For $F \in \mathrm{Mat}(A_d)$, we clearly have

$$\mathsf{ac}(F) \le (d-1)\cdot w(F). \tag{8.21}$$

The following proposition is classical and shows the other direction.

8.22 Proposition (Homogenization)**.** *Let H be an $a\times b$ matrix with entries in $\mathscr{T}$ or with entries in $\mathscr{S}$ or with entries in $\mathscr{W}$, and let F be a homogeneous component of H. Then $w(F) \le \max\{\mathsf{ac}(H), a, b\}$. In particular, $w(F) \le \max\{\mathsf{ac}(F), a, b\}$.* $\diamond$

Proof. A vertex in an ABP computes a matrix, by considering only those source-sink paths that use that vertex. Let $d := \deg(F)$. We take a minimal size aABP G that computes H and we replace every vertex v with $d+1$ many vertices $v_0, \dots, v_d$. We will set the edges so that if a vertex v computes some matrix M, then v_i will compute the i-th homogeneous component of M, for $0 \le i \le d$, and we will not compute higher degree components. If $\ell + \alpha$ is an edge label in G from v to w with ℓ homogeneous linear and α constant, then we add edges with label ℓ from each v_i to w_{i+1}, and we add edges with label α from each v_i to w_i. For the sources s_a we delete all $s_{a,i}$ with $i > 0$, and for the sinks t_b we delete all $t_{b,i}$ with $i < d$. This finishes the construction. The result is an ABP that has at most $\mathsf{ac}(H)$ many vertices in each vertex layer but the first and last, in other words, the width of the ABP is at most $\mathsf{ac}(H)$. Hence $w(F) \le \mathsf{ac}(H)$. $\square$

8.23 Corollary. *For $A \in \{\mathscr{T}, \mathscr{S}\}$ the sequences $(\xi_d)_{d\in\mathbb{N}}$ and $(\gamma_d)_{d\in\mathbb{N}}$ grow monotonically.* $\diamond$

Proof. Let $\pi \in \mathfrak{S}_d$ with $\pi(d) = d$, and let $\kappa \in \mathfrak{S}_{d-1}$ be defined via $\forall 1 \le i \le d-1 : \kappa(i) = \pi(i)$. Take an aABP of minimal size that computes $\mathsf{h}\Xi_{n,d,\pi}$. Set all variables $x^{(d)}_{i,i} \boxtimes x^{(d)}_{i',i'}$ to 1, and set all other $x^{(d)}_{i,j} \boxtimes x^{(d)}_{i',j'}$ to zero. The resulting aABP computes $\mathsf{h}\Xi_{n,d-1,\kappa}$. Hence $\mathsf{ac}(\mathsf{h}\Xi_{n,d-1,\kappa}) \le \mathsf{ac}(\mathsf{h}\Xi_{n,d,\pi})$. By Claim 8.7 and (8.21) and Proposition 8.22 this implies $w(\mathsf{h}\Xi_{n,d-1,\kappa}) \le d \cdot w(\mathsf{h}\Xi_{n,d,\pi})$. Hence, $\forall \kappa \in \mathfrak{S}_{d-1} : \limsup_{n\to\infty}\{\log_n(w(\mathsf{h}\Xi_{n,d-1,\kappa}))\} \le \xi_d$, and therefore $\xi_{d-1} \le \xi_d$ by (8.4). The proof of $\gamma_{d-1} \le \gamma_d$ is analogous. $\square$

8.24 Proposition. *Let R be a commutative ring, and $A = \mathscr{W}$. Then the sequences $(\xi_d)_{d\in\mathbb{N}}$ and $(\gamma_d)_{d\in\mathbb{N}}$ grow monotonically.* $\diamond$

Proof. Let $\pi \in \mathfrak{S}_d$ with $\pi(d) = d$, and let $\kappa \in \mathfrak{S}_{d-1}$ be defined via $\forall 1 \le i \le d-1 : \kappa(i) = \pi(i)$. We use the following crucial property over commutative (semi)rings that holds also for $\mathscr{W}$: If in $\mathsf{h}\Xi_{n,d,\pi}$ we map each $x^{(d)}_{i,j} \boxtimes x^{(d)}_{i',j'}$ to y if and only if $i = j$ and $i' = j'$, otherwise we map it to zero, we obtain

$(\mathsf{h}\Xi_{n,d-1,\kappa}) \wedge y$. If in $\mathsf{h}\Xi_{n,d}$ we additionally map all $x_{i,j}^{(d)} \boxtimes x_{i',j'}^{(d')}$ to zero for which $d \neq d'$, we obtain $\mathsf{h}\Xi_{n,d-1} \wedge y$. Given an aABP that computes $\mathsf{h}\Xi_{n,d-1} \wedge y$, we split every vertex into two vertices, one for y-degree 0 and one for y-degree 1 (higher y-degrees are deleted). We now replace every y that appears in an edge layer of the same parity as d with 1, and we replace y with -1 otherwise. The resulting aABP computes $\mathsf{h}\Xi_{n,d-1}$, hence $\mathsf{ac}(\mathsf{h}\Xi_{n,d-1}) \leq 2 \cdot \mathsf{ac}(\mathsf{h}\Xi_{n,d})$, and thus $\xi_{d-1} \leq \xi_d$, see (8.21) and Proposition 8.22. The proof of $\gamma_{d-1} \leq \gamma_d$ is analogous. □

9 Coefficients of polynomials given by Boolean circuits

In this section we study elements of A_d whose coefficients are computed by small Boolean circuits, defined below. Let R be a commutative semiring and $A \in \{\mathscr{T}, \mathscr{S}, \mathscr{W}\}$. We define the negation $\overline{0} := 1$ and $\overline{1} := 0$. For a finite list $\mathbf{b}$ we define $\overline{\mathbf{b}}$ as the negation applied on all elements. We also use the negation symbol on variables, for example $\overline{x_i}$, but this is a new variable instead of a negation. We define

$$(x, \overline{x})^{\mathbf{b},\overline{\mathbf{b}}} \;:=\; x_1^{b_1}\overline{x}_1^{\overline{b}_1} x_2^{b_2}\overline{x}_2^{\overline{b}_2} \cdots x_m^{b_m}\overline{x}_m^{\overline{b}_m} \;\in A_m.$$

A Boolean circuit B is a directed acyclic graph with a single outdegree 0 vertex, whose indegree 0 vertices (called *inputs*) are labeled with variables, and whose other vertices (called *computation gates*) are each labeled with an operation `AND`, `OR`, `NOT`, where the vertices labeled with `AND` or `OR` have indegree 2, and the vertices labeled with `NOT` have indegree 1. We will always assume that B is connected, i.e., there is at least one path from each input to the output, which can be achieved by taking any gate v that connects to the output and replacing it with v `AND` (x `OR` `NOT` x) for an unused input x. On input $\mathbf{b} \in \{0,1\}^m$, B computes a Boolean value by induction over the circuit structure in the obvious way. The size $|B|$ is defined as the number of the computation gates of B. An input $\mathbf{b}$ with $B(\mathbf{b}) = 1$ is called a satisfying assignment of B.

In this section we prove the following theorem in which we use the dummy variables $\epsilon^{(k)} \in \mathrm{Var}$.

9.1 Theorem. *Given a Boolean circuit B on m inputs. Then*

$$\sum_{B(\mathbf{b})=1} (x, \overline{x})^{\mathbf{b},\overline{\mathbf{b}}} \prod_{i=m+1}^{d} \epsilon^{(i)} \;\leq\; \mathsf{h}\Gamma_{n,d},$$

with $n \leq (4|B| + 2m)^2$ and $d \leq 7|B| + m + 1$. ◇

The proof proceeds in several steps, where the first steps are classical. We first use the Tseytin transformation (Lemma 9.2) to convert B into a product of clauses. These clauses can then be arithmetized in low width, which gives part of the second factor in the $\boxtimes$ construction. The construction of the first factor h is a variant of the elementary symmetric polynomial, see Lemma 9.4.

A *literal* is a variable or its negation. A *clause* is defined as the `OR` of at most 3 literals.

9.2 Lemma (Tseytin transformation [Tse68, Tse83])**.** *For every Boolean circuit B in m variables $\mathbf{b}$ there exists a set $\mathcal{C}$ of clauses in $r > m$ variables $(\mathbf{b}, o, \mathbf{b}')$ with the following property: for every $\mathbf{b} \in \{0,1\}^m$ there exists a unique $(o, \mathbf{b}') \in \{0,1\}^{r-m}$ that satisfies $\mathcal{C}(\mathbf{b}, o, \mathbf{b}') = 1$. Moreover, for this $(o, \mathbf{b}')$ we have $B(\mathbf{b}) = o$.* ◇

Proof. Given B on m variables $\mathbf{b}$, we first attach a new wire to the output gate, but this wire has *no second endpoint*, and we call it the output wire. We assign a variable a_e to each wire e: If the wire e is connected to the i-th input, then $a_e = b_i$. If the wire e is the output wire, then $a_e = o$. Otherwise, a new variable b'_j is created and assigned to the wire. This gives a variable vector $(\mathbf{b}, o, \mathbf{b}')$. A Boolean assignment to these variables can be interpreted as a Boolean assignment to wires. We now introduce a set of clauses that is true if and only if the Boolean values on the wires respect the gate operations. For example, for every `NOT` gate with input wire e_1 and output wire e_2 we observe that the logical expression $a_{e_2} = \mathtt{NOT}(a_{e_1})$ can be expressed as the `AND` of these clauses: $\mathtt{OR}(\overline{a_{e_1}}, \overline{a_{e_2}})$, $\mathtt{OR}(a_{e_1}, a_{e_2})$. Moreover, for every AND gate with input wires e_1, e_2 and output wire e_3 we observe (for example via case distinction over all $2^3 = 8$ cases) that the logical expression $a_{e_3} = \mathtt{AND}(a_{e_1}, a_{e_2})$ can be expressed as the AND of these clauses: $\mathtt{OR}(\overline{a_{e_1}}, \overline{a_{e_2}}, a_{e_3})$, $\mathtt{OR}(a_{e_1}, \overline{a_{e_3}})$, $\mathtt{OR}(a_{e_2}, \overline{a_{e_3}})$. Analogously, $a_{e_3} = \mathtt{OR}(a_{e_1}, a_{e_2})$ can be expressed as the AND of these clauses: $\mathtt{OR}(a_{e_1}, a_{e_2}, \overline{a_{e_3}})$, $\mathtt{OR}(\overline{a_{e_1}}, a_{e_3})$, $\mathtt{OR}(\overline{a_{e_2}}, a_{e_3})$. Let $\mathcal{C}$ be the set of all these clauses. By construction, for every $\mathbf{b}$ there is a unique $(o, \mathbf{b}')$ such that $\prod_{C \in \mathcal{C}} C(\mathbf{b}, o, \mathbf{b}') = 1$, and this $(o, \mathbf{b}')$ is given by the Boolean values on the wires for input $\mathbf{b}$, defined recursively over the circuit structure. Clearly, $B(\mathbf{b}) = o$. □

We now convert the set $\mathcal{C}$ of clauses into a width 7 ABP in $2r$ many variables y_i and $\overline{y}_i$. We arrange the clauses in any order, and each clause (l_1 OR l_2 OR l_3) is converted into an ABP, and all those programs are concatenated by identifying the sink of the left ABP with the source of the right ABP, which results in an ABP that computes their product. Let $l_1 \in \{b_{i_1}, \overline{b}_{i_1}\}$, $l_2 \in \{b_{i_2}, \overline{b}_{i_2}\}$, $l_3 \in \{b_{i_3}, \overline{b}_{i_3}\}$, for some $1 \le i_1, i_2, i_3 \le r$. We now convert the clause (l_1 OR l_2 OR l_3). For $j \in \{1,2,3\}$ let $\ell_j = y_{i_j}$ if $l_j = b_{i_j}$, and let $\ell_j = \overline{y}_{i_j}$ otherwise. Analogously, let $\overline{\ell}_j = \overline{y}_{i_j}$ if $l_j = b_{i_j}$, and let $\overline{\ell}_j = y_{i_j}$ otherwise. The clause (l_1 OR l_2 OR l_3) is converted into the ABP

$$\ell_1 \cdot \ell_2 \cdot \ell_3 + \ell_1 \cdot \overline{\ell_2} \cdot \ell_3 + \ell_1 \cdot \ell_2 \cdot \overline{\ell_3} + \ell_1 \cdot \overline{\ell_2} \cdot \overline{\ell_3} + \overline{\ell_1} \cdot \ell_2 \cdot \ell_3 + \overline{\ell_1} \cdot \overline{\ell_2} \cdot \ell_3 + \overline{\ell_1} \cdot \ell_2 \cdot \overline{\ell_3}.$$

The key property is that for every satisfying assignment to the clause there is exactly one s-t-path in this program, and vice versa. Clauses with 2 variables are handled analogously with 3 summands instead of 7, and degree 2 instead of 3. Concatenating these ABPs we obtain an ABP G'. We call an s-t-path in G' *consistent* if it does not use both y_i and $\overline{y}_i$ for any i. We see that there are bijections between the following sets:

- The set of vectors $\mathbf{b} \in \{0,1\}^m$
- The set of satisfying assignments $(\mathbf{b}, o, \mathbf{b}')$ of the clause conjunction $\mathcal{C}$ (9.3)
- The set of consistent s-t-paths in G'

Since a satisfying assignment $(\mathbf{b}, o, \mathbf{b}')$ to $\mathcal{C}$ depends on $\mathbf{b}$, we write $p_{\mathbf{b}}$ instead of $p_{\mathbf{b},o,\mathbf{b}'}$. And in the other direction, we denote by $(\mathbf{b}, o, \mathbf{b}')_p$ the satisfying assignment of $\mathcal{C}$ to a consistent path p, and we write $\mathbf{b}_p$ for the first m bits of $(\mathbf{b}, o, \mathbf{b}')_p$.

We modify the ABP G' by replacing every edge label y_i or $\overline{y}_i$ in layer k by $y_i^{(m+1+k)}$ and $\overline{y}_i^{(m+1+k)}$, respectively. We call the resulting ABP G.

Since every computation gate in B resulted in at most 3 clauses with at most 7 literals in total, the degree of G is at most $7|B|$. We let $d := \deg(q) + m + 1$, where G computes q, and observe $d \le 7|B| + m + 1$. We let r be the length of $(\mathbf{b}, o, \mathbf{b}')$, and we observe $r \le m + |B|$.

Let $Z = \{z_{i,(k)}, \overline{z}_{i,(k)} \mid 1 \le k \le d,\ 1 \le i \le r\}$. Note that $|Z| = n := 2dr \le 2(7|B|+m+1)(m+|B|) \le (4|B|+2m)^2$. A subset $S \subseteq Z$ is called consistent if it does not contain any two variables $z_{i,(k)}$ and $\overline{z}_{i,(k')}$. Let

$$h_{\text{poly}} := \sum_{\substack{S \subseteq Z \\ |S|=d \\ S \text{ consistent}}} \prod_{z \in S} z \in \mathscr{S}_d.$$

9.4 Lemma. $w(h_{\text{poly}}) \le 2dr$. ◇

Proof. We sort Z as follows:

$$\begin{aligned}\mathbf{z} = (&z_{1,(1)}, \ldots, z_{1,(d)}, \overline{z}_{1,(1)}, \ldots, \overline{z}_{1,(d)},\\ &z_{2,(1)}, \ldots, z_{2,(d)}, \overline{z}_{2,(1)}, \ldots, \overline{z}_{2,(d)},\\ &\vdots\\ &z_{r,(1)}, \ldots, z_{r,(d)}, \overline{z}_{r,(1)}, \ldots, \overline{z}_{r,(d)}),\end{aligned}$$

which also sorts the triples (ι, i, k), $\iota \in \{\texttt{unbarred}, \texttt{barred}\}$, $1 \le i \le r$, $1 \le k \le d$. We write $\mathrm{pred}(\iota, i, k)$ for the predecessor of (ι, i, k), which is the successor when traversing $\mathbf{z}$ row-wise from right to left, bottom to top. The property $\mathrm{jump}(\iota, i, k)$ is defined to be true if and only if $\mathrm{pred}(\iota, i, k) = (\texttt{barred}, i-1, d)$. This captures a jump in the row when traversing $\mathbf{z}$ row-wise from right to left, bottom to top.

$$h_{\iota,i,k,d} = \sum_{\substack{S \subseteq Z_{\le(\iota,i,k)} \\ |S|=d \\ S \text{ consistent}}} \prod_{z \in S} z \qquad\qquad h'_{\iota,i,k,d} = \sum_{\substack{S \subseteq Z_{\le(\iota,i,k)} \\ |S|=d \\ S \text{ consistent} \\ \forall j: z_{i,(j)} \notin S}} \prod_{z \in S} z$$

We have $h_{\mathrm{poly}} = h_{\texttt{barred},r,d,d}$. We now observe

$$h_{\iota,i,k,d} \;=\; \begin{cases} h_{\mathrm{pred}(\iota,i,k),d} + z_{\iota,i,k} \cdot h'_{\mathrm{pred}(\iota,i,k),d-1} & \text{if } \iota = \texttt{barred} \\ h_{\mathrm{pred}(\iota,i,k),d} + z_{\iota,i,k} \cdot h_{\mathrm{pred}(\iota,i,k),d-1} & \text{if } \iota = \texttt{unbarred} \end{cases}$$

$$h'_{\iota,i,k,d} \;=\; \begin{cases} h'_{\mathrm{pred}(\iota,i,k),d} + z_{\iota,i,k} \cdot h'_{\mathrm{pred}(\iota,i,k),d-1} & \text{if } \iota = \texttt{barred} \text{ and not } \mathrm{jump}(\iota,i,k) \\ h'_{\mathrm{pred}(\iota,i,k),d} & \text{if } \iota = \texttt{unbarred} \text{ and not } \mathrm{jump}(\iota,i,k) \\ h_{\mathrm{pred}(\iota,i,k),d} + z_{\iota,i,k} \cdot h_{\mathrm{pred}(\iota,i,k),d-1} & \text{if } \iota = \texttt{barred} \text{ and } \mathrm{jump}(\iota,i,k) \\ h_{\mathrm{pred}(\iota,i,k),d} & \text{if } \iota = \texttt{unbarred} \text{ and } \mathrm{jump}(\iota,i,k) \end{cases}$$

This gives an ABP of the desired width. A simpler but slightly larger variant is

$$h_{r,d} \;=\; h_{r-1,d} + \sum_{k=1}^{d} \big(e_{d,k}(z_{r,(1)}, \ldots, z_{r,(d)}) + e_{d,k}(\overline{z}_{r,(1)}, \ldots, \overline{z}_{r,(d)})\big) \cdot h_{r-1,d-k}. \qquad \square$$

Let $h = \mathrm{polar}(h_{\mathrm{poly}})$. Since $h_{\mathrm{poly}} \le \Gamma_{2dr,d}$, it follows $h \le \mathrm{polar}(\Gamma_{2dr,d}) = H_{2dr,d}$, see (7.23).

We define $\mathcal{T} \in \mathrm{End}$ via mapping every variable to zero with the following exceptions:

$\mathcal{T}(z_{i,(k)} \boxtimes y_i^{(k)}) = \epsilon^{(k)}, \mathcal{T}(\overline{z}_{i,(k)} \boxtimes \overline{y}_i^{(k)}) = \epsilon^{(k)}, \mathcal{T}(z_{i,(i)} \boxtimes \overset{?}{x}_i) = x_i, \mathcal{T}(\overline{z}_{i,(i)} \boxtimes \overset{?}{x}_i) = \overline{x}_i, \mathcal{T}(z_{m+1,(m+1)} \boxtimes y_{\mathrm{o}}) = \epsilon^{(m+1)},$

where the $\overset{?}{x}_i \in \mathrm{Var}$.

9.5 Lemma. *Fix B and the arithmetization G. If an s-t-path p in G is inconsistent, then $\mathcal{T}(h \boxtimes (\overset{?}{x}_1 \cdots \overset{?}{x}_m\, y_{\mathrm{o}}\, p)) = 0$. If p in G is consistent, then $\mathcal{T}(h \boxtimes (\overset{?}{x}_1 \cdots \overset{?}{x}_m\, y_{\mathrm{o}}\, p)) = B(\mathbf{b}_p)(x, \overline{x})^{\mathbf{b}_p, \overline{\mathbf{b}}_p} \prod_{i=m+1}^{d} \epsilon^{(i)}$.* $\diamond$

Proof. If p is inconsistent, then there is i such that both $y_i^{(k)}$ and $\overline{y}_i^{(\widetilde{k})}$ exist in p for some $k, \widetilde{k}$. Hence, to have nonzeroness, a monomial in h_{poly} must have both $z_{i,(k)}$ and $\overline{z}_{i,(\widetilde{k})}$ as variables, but this is impossible by construction. This proves the first part.

Let p be consistent with corresponding Boolean assignment $(\mathbf{b}, o, \mathbf{b}')$ satisfying $\mathcal{C}(\mathbf{b}, o, \mathbf{b}') = 1$. If $B(\mathbf{b}) = 1$, then it is straightforward by pairing up variables to find a monomial c_{poly} in h_{poly} that satisfies $\mathcal{T}(c \boxtimes (\overset{?}{x}_1 \cdots \overset{?}{x}_m\, y_{\mathrm{o}}\, p)) = (x, \overline{x})^{\mathbf{b}, \overline{\mathbf{b}}} \prod_{i=m+1}^{d} \epsilon^{(i)}$, where $c = \mathrm{polar}(c_{\mathrm{poly}})$.

It remains to prove that for every monomial c_{poly} in h_{poly}, $c := \mathrm{polar}(c_{\mathrm{poly}})$, we have that $\mathcal{T}(c \boxtimes (\overset{?}{x}_1 \cdots \overset{?}{x}_m\, y_{\mathrm{o}}\, p)) \neq 0$ implies both $B(\mathbf{b}_p) = 1$ and $\mathcal{T}(c \boxtimes (\overset{?}{x}_1 \cdots \overset{?}{x}_m\, y_{\mathrm{o}}\, p)) = (x, \overline{x})^{\mathbf{b}, \overline{\mathbf{b}}} \prod_{i=m+1}^{d} \epsilon^{(i)}$. By construction of $\mathcal{T}$, the nonvanishing $\mathcal{T}(c \boxtimes (\overset{?}{x}_1 \cdots \overset{?}{x}_m\, y_{\mathrm{o}}\, p)) \neq 0$ fixes $d - m - 1$ many variables in c. For the remaining variables, by the consistency of h_{poly}, c contains $z_{i,(i)}$ if $b_i = 1$ (because $\overline{z}_{i,(i)}$ is unavailable due to the consistency of h), and c contains $\overline{z}_{i,(i)}$ otherwise (by the same argument). The nonvanishing also implies that c contains $z_{m+1,(m+1)}$ and by consistency of h another $z_{m+1,(k)}$ for $k > m+1$. This implies $o = 1$, which by Lemma 9.2 implies $B(\mathbf{b}) = 1$. Now it is readily checked that $\mathcal{T}(c \boxtimes (\overset{?}{x}_1 \cdots \overset{?}{x}_m\, y_{\mathrm{o}}\, p)) = (x, \overline{x})^{\mathbf{b}, \overline{\mathbf{b}}} \prod_{i=m+1}^{d} \epsilon^{(i)}$. $\square$

Proof of Theorem 9.1. Given B and the arithmetization ABP G computing q. We have

$$\mathcal{T}(h \boxtimes (\overset{?}{x}_1 \cdots \overset{?}{x}_m\, y_{\mathrm{o}}\, q)) \overset{\text{Lem. 9.5}}{=} \sum_{\substack{p \text{ consistent in } q\\ B(\mathbf{b}_p)=1}} (x,\overline{x})^{\mathbf{b}_p,\overline{\mathbf{b}}_p} \prod_{i=m+1}^{d} \epsilon^{(i)}$$

$$\overset{(9.3)}{=} \sum_{\substack{\mathbf{b},o,\mathbf{b}'\\ \mathcal{C}(\mathbf{b},o,\mathbf{b}')=1\\ B(\mathbf{b})=1}} (x,\overline{x})^{\mathbf{b},\overline{\mathbf{b}}} \prod_{i=m+1}^{d} \epsilon^{(i)} \overset{\text{Lem. 9.2}}{=} \sum_{B(\mathbf{b})=1} (x,\overline{x})^{\mathbf{b},\overline{\mathbf{b}}} \prod_{i=m+1}^{d} \epsilon^{(i)}.$$

□

9.I Kronecker product close to the permanent

Recall the definitions of VNP and VNP-completeness from §1. In this section, we discuss VNP-completeness of a polynomial that is very close to the permanent. Our goal is to reach as close as we can to the permanent polynomial while keeping VNP-hardness over all commutative semirings.

We know that $\mathrm{per}_d = \mathrm{mon}_d \boxtimes \mathrm{mon}_d$, where $\mathrm{mon}_d = \Gamma_{1,d}$. We relax one of the factors to $\Gamma_{2,d}$. We define the degree-$2d$ homogeneous polynomial:

$$Q_{2d} := \prod_{i=1}^{d} \big(y_i^{(0)} z_i^{(0)} + y_i^{(1)} z_i^{(1)}\big),$$

and remark that $Q_{2d} \le \Gamma_{2,2d}$. We define the new polynomial $U_{2d} := Q_{2d} \boxtimes \mathrm{mon}_{2d}$.

9.6 Theorem. *The sequence of polynomials U_{2d} is* VNP*-complete over every commutative semiring.* ◇

Proof. Containment in VNP is obvious. We prove a reduction similar to Theorem 9.1: for every Boolean circuit $B(x_1, \dots, x_m)$ there is $d = O(|B| + m)$ such that

$$\epsilon^{2d-m} \sum_{\substack{\alpha \in \{0,1\}^m\\ B(\alpha)=1}} \prod_{i=1}^{m} x_i^{\alpha_i} \bar{x}_i^{1-\alpha_i} \le Q_{2d} \boxtimes \mathrm{mon}_{2d}.$$

We consider that a monomial of Q_{2d} is obtained by choosing one bit $c_v \in \{0,1\}$ for every $v \in [d]$:

$$M_{\overline{c}} = \prod_{v=1}^{d} y_v^{(c_v)} z_v^{(c_v)}.$$

Thus, every choice $\overline{c} = (c_1, \dots, c_d)$ gives $2d$ factors, which we call *ports*:

$$y_1^{(c_1)}, z_1^{(c_1)}, \dots, y_d^{(c_d)}, z_d^{(c_d)}.$$

On the other hand, mon_{2d} contains $2d$ distinct factors, which we call labels. Hence a term of $M_{\overline{c}} \boxtimes \mathrm{mon}_d$ is obtained by assigning the $2d$ labels of mon_{2d} bijectively to these $2d$ ports. Each port-label pair is a variable of the polynomial U_{2d}. We define the restriction in a way such that a term of U_{2d} survives (it is non-zero) if and only if the induced matching of ports to labels corresponds to a satisfying assignment of the Boolean circuit B.

From the side of the Boolean circuit B, we apply the Tseytin transformation of Lemma 9.2. We obtain a CNF formula Φ_B with the following property: for every input assignment $\alpha \in \{0,1\}^m$, there is a unique extension to the auxiliary Tseytin variables satisfying all gate constraints, and this extension satisfies the output condition precisely when $B(\alpha) = 1$. We may assume that every clause has exactly three literal occurrences, repeating literals when necessary. If an original input variable does not occur in the Tseytin formula, we may add the tautological clause $x_i \vee \neg x_i \vee x_i$, so that every original input occurs at least once. Let d be the total number of literal occurrences of Φ_B (with repetitions). Since the Tseytin transformation has linear size, $d = O(|B| + m)$.

y-ports and equality variables. Suppose a Boolean variable w of Φ_B occurs at positions $v_1, \ldots, v_r$. We reserve r label-variables from the variables of mon_{2d}:

$$e_1, \ldots, e_r,$$

which we call equality variables.

We define

$$T(y^{(0)}_{v_j} \boxtimes e_j) = \begin{cases} \overline{x}_w & \text{for } j = 1 \\ \epsilon & \text{otherwise} \end{cases}$$

and

$$T(y^{(1)}_{v_j} \boxtimes e_{j+1}) = \begin{cases} x_w & \text{for } j = 1 \\ \epsilon & \text{otherwise,} \end{cases}$$

where we set $e_{r+1} := e_1$.

A monomial that is not mapped to zero either only uses the first type of restriction for w or only uses the second type for w, because in any other case there is at least either a port or a label used twice. Thus, this fixes $c_{v_1} = \cdots = c_{v_r}$. In particular, this factor of a surviving monomial is either mapped to $\epsilon^{r-1}\overline{x}_w$ (if $c_{v_1} = 0$) or to $\epsilon^{r-1}x_w$ (if $c_{v_1} = 1$).

Overall, d variables from the polynomial mon_{2d} are equality variables, and all values $c_1, \ldots, c_d$ are fixed in this way.

z-ports and clause variables. We treat all d positions now. Consider a clause C of Φ_B, $C = \ell_1 \vee \ell_2 \vee \ell_3$ where the literals correspond to variables at the positions $v_{i_1}, v_{i_2}, v_{i_3}$ that we have not yet treated. We reserve three fresh variables from mon_{2d}:

$$q_{i_1}, q_{i_2}, q_{i_3}.$$

We set

$$\begin{aligned}
&\big(T(z^{(0)}_{i_1} \boxtimes q_{i_1}), T(z^{(0)}_{i_1} \boxtimes q_{i_2}), T(z^{(0)}_{i_1} \boxtimes q_{i_3}); T(z^{(1)}_{i_1} \boxtimes q_{i_1}), T(z^{(1)}_{i_1} \boxtimes q_{i_2}), T(z^{(1)}_{i_1} \boxtimes q_{i_3})\big) \\
&\qquad\qquad = (\epsilon, 0, 0; 0, 0, \epsilon) \quad \text{if } \ell_1 \text{ is not a negation} \\
&\big(T(z^{(0)}_{i_1} \boxtimes q_{i_1}), T(z^{(0)}_{i_1} \boxtimes q_{i_2}), T(z^{(0)}_{i_1} \boxtimes q_{i_3}); T(z^{(1)}_{i_1} \boxtimes q_{i_1}), T(z^{(1)}_{i_1} \boxtimes q_{i_2}), T(z^{(1)}_{i_1} \boxtimes q_{i_3})\big) \\
&\qquad\qquad = (0, 0, \epsilon; \epsilon, 0, 0) \quad \text{if } \ell_1 \text{ is a negation} \\
&\big(T(z^{(0)}_{i_2} \boxtimes q_{i_1}), T(z^{(0)}_{i_2} \boxtimes q_{i_2}), T(z^{(0)}_{i_2} \boxtimes q_{i_3}); T(z^{(1)}_{i_2} \boxtimes q_{i_1}), T(z^{(1)}_{i_2} \boxtimes q_{i_2}), T(z^{(1)}_{i_2} \boxtimes q_{i_3})\big) \\
&\qquad\qquad = (0, \epsilon, 0; \epsilon, 0, \epsilon) \quad \text{if } \ell_2 \text{ is not a negation} \\
&\big(T(z^{(0)}_{i_2} \boxtimes q_{i_1}), T(z^{(0)}_{i_2} \boxtimes q_{i_2}), T(z^{(0)}_{i_2} \boxtimes q_{i_3}); T(z^{(1)}_{i_2} \boxtimes q_{i_1}), T(z^{(1)}_{i_2} \boxtimes q_{i_2}), T(z^{(1)}_{i_2} \boxtimes q_{i_3})\big) \\
&\qquad\qquad = (\epsilon, 0, \epsilon; 0, \epsilon, 0) \quad \text{if } \ell_2 \text{ is a negation} \\
&\big(T(z^{(0)}_{i_3} \boxtimes q_{i_1}), T(z^{(0)}_{i_3} \boxtimes q_{i_2}), T(z^{(0)}_{i_3} \boxtimes q_{i_3}); T(z^{(1)}_{i_3} \boxtimes q_{i_1}), T(z^{(1)}_{i_3} \boxtimes q_{i_2}), T(z^{(1)}_{i_3} \boxtimes q_{i_3})\big) \\
&\qquad\qquad = (\epsilon, \epsilon, 0; \epsilon, \epsilon, \epsilon) \quad \text{if } \ell_3 \text{ is not a negation} \\
&\big(T(z^{(0)}_{i_3} \boxtimes q_{i_1}), T(z^{(0)}_{i_3} \boxtimes q_{i_2}), T(z^{(0)}_{i_3} \boxtimes q_{i_3}); T(z^{(1)}_{i_3} \boxtimes q_{i_1}), T(z^{(1)}_{i_3} \boxtimes q_{i_2}), T(z^{(1)}_{i_3} \boxtimes q_{i_3})\big) \\
&\qquad\qquad = (\epsilon, \epsilon, \epsilon; \epsilon, \epsilon, 0) \quad \text{if } \ell_3 \text{ is a negation}
\end{aligned}$$

In other words, if ℓ_3 is `true`, then the corresponding $z^{(*)}_{i_3}$ ($z^{(0)}_{i_3}$ if ℓ_3 is a negation, and $z^{(1)}_{i_3}$ otherwise) can be paired with any of the three variables, and when ℓ_3 is `false`, then it can only be paired with q_{i_1} and q_{i_2}. We illustrate the situation for all the literals in a more readable format:

	`false`	`true`
ℓ_1	$\{q_{i_1}\}$	$\{q_{i_3}\}$
ℓ_2	$\{q_{i_2}\}$	$\{q_{i_1}, q_{i_3}\}$
ℓ_3	$\{q_{i_1}, q_{i_2}\}$	$\{q_{i_1}, q_{i_2}, q_{i_3}\}$.

Exactly one unique degree 3 monomial survives the restriction (i.e., is not mapped to zero) if and only if the clause is satisfied, otherwise the monomial is mapped to zero, which is easy to check case by case:

(ℓ_1, ℓ_2, ℓ_3)	$(z^{(*)}_{i_1}, z^{(*)}_{i_2}, z^{(*)}_{i_3})$
000	no possibility
001	$(q_{i_1}, q_{i_2}, q_{i_3})$
010	$(q_{i_1}, q_{i_3}, q_{i_2})$
011	$(q_{i_1}, q_{i_3}, q_{i_2})$
100	$(q_{i_3}, q_{i_2}, q_{i_1})$
101	$(q_{i_3}, q_{i_2}, q_{i_1})$
110	$(q_{i_3}, q_{i_1}, q_{i_2})$
111	$(q_{i_3}, q_{i_1}, q_{i_2})$.

This uses the remaining d variables of mon_{2d}.

Since the variables have consistent truth values through the equality variables and all monomials are mapped to zero for which a clause is not satisfied, altogether we obtain

$$T(Q_{2d} \boxtimes \mathrm{mon}_{2d}) = \epsilon^{2d-m} \sum_{\substack{\alpha\in\{0,1\}^m \\ B(\alpha)=1}} \prod_{i=1}^{m} x_i^{\alpha_i} \bar{x}_i^{1-\alpha_i}. \qquad \square$$

In the end, Theorem 9.6, implies the VNP-hardness for a variety of polynomials:

9.7 Corollary. *The polynomial* $\mathrm{mon}_d^{\boxtimes r}$ *for* $r \geq 3$ *is* VNP*-complete over every commutative semiring. Moreover, the polynomials* $\det_d^{\boxtimes r}$ *for* $r \geq 2$ *and the mixed discriminant* $\det_d \boxtimes \mathrm{mon}_d$ *are* VNP*-complete over every commutative ring.* ◇

Proof. We have the following restrictions:

$$Q_{2d} = \prod_{i=1}^{d} (y_i^{(0)} z_i^{(0)} + y_i^{(1)} z_i^{(1)}) = \prod_{i=1}^{d} \mathrm{per}_2 \begin{pmatrix} y_i^{(0)} & y_i^{(1)} \\ z_i^{(1)} & z_i^{(0)} \end{pmatrix} \leq \mathrm{per}_{2d},$$

where the last step is by taking a block diagonal matrix. Similarly,

$$Q_{2d} = \prod_{i=1}^{d} (y_i^{(0)} z_i^{(0)} + y_i^{(1)} z_i^{(1)}) = \prod_{i=1}^{d} \det_2 \begin{pmatrix} y_i^{(0)} & y_i^{(1)} \\ -z_i^{(1)} & z_i^{(0)} \end{pmatrix} \leq \det_{2d}.$$

The corollary follows from Proposition 7.8. □

We now discuss the hyperhafnian polynomial and the hyperpfaffian polynomial, which are generalisations of the hafnian (polynomial VIII in [Val79, p. 253]) and the pfaffian respectively.

Fix $r \geq 2$. Let $\Pi_r(d)$ be the set of partitions of $[rd]$ into d unordered blocks of size r. With a variable x_I for every r-subset I, we define the hyperhafnian polynomial:

$$\mathrm{HypHaf}_d^{[r]}(X) := \sum_{P\in\Pi_r(d)} \prod_{I\in P} x_I.$$

For the hyperpfaffian, assume r is even. For any set $I = \{i_1 < \cdots < i_r\}$ of integers, we list its elements increasingly. Given a partition $P = \{I_1, \ldots, I_d\}$ of $[rd]$, concatenate the increasing listings and let $\mathrm{sgn}(P)$ be the sign of the resulting permutation of $[rd]$ (this is well-defined, because r is even). Define the hyperpfaffian polynomial (see e.g. [Bar95, GIM+20]):

$$\mathrm{HypPfaf}_d^{[r]}(X) := \sum_{P\in\Pi_r(d)} \mathrm{sgn}(P) \prod_{I\in P} x_I.$$

Both polynomials coincide over characteristic 2. For $r = 2$, we get the hafnian and the pfaffian respectively. We point out that the pfaffian is in VP, whereas $\mathrm{per}_d \leq \mathrm{HypHaf}_d^{[2]}$ by mapping $x_{\{i,j\}}$ to zero if i and j both lie in the lower or both lie in the upper half of $[2d]$, and we map $x_{\{i,j\}}$ to $x_{i-d,j}$ otherwise, where $i > d$, see [Bar16, §4.1]. We have $\mathrm{mon}_d^{\boxtimes r} \leq \mathrm{HypHaf}_d^{[r]}$ analogously, but for $r \geq 3$ we have that $\mathrm{mon}_d^{\boxtimes r}$ is VNP-complete over all commutative semirings by Corollary 9.7, therefore $\mathrm{HypHaf}_d^{[r]}$ for $r \geq 3$ is also VNP-complete over all commutative semirings. In [GIM+20] it is shown that $\mathrm{mon}_d^{\boxtimes 2} \leq \mathrm{HypPfaf}_d^{[4]}$. The proof generalizes to $\mathrm{mon}_d^{\boxtimes r} \leq \mathrm{HypPfaf}_d^{[2r]}$, which shows the VNP-completeness of $\mathrm{HypPfaf}_d^{[2r]}$ for all $r \geq 3$ over all commutative rings.

10 P $\neq$ NP and related conjectures

In this section we set $R = \mathbb{C}$ and $A = \mathscr{S}$. Figure 4 shows formal implications between several conjectures. We are very far from proving any of the conjectures in Figure 4, as any of them either

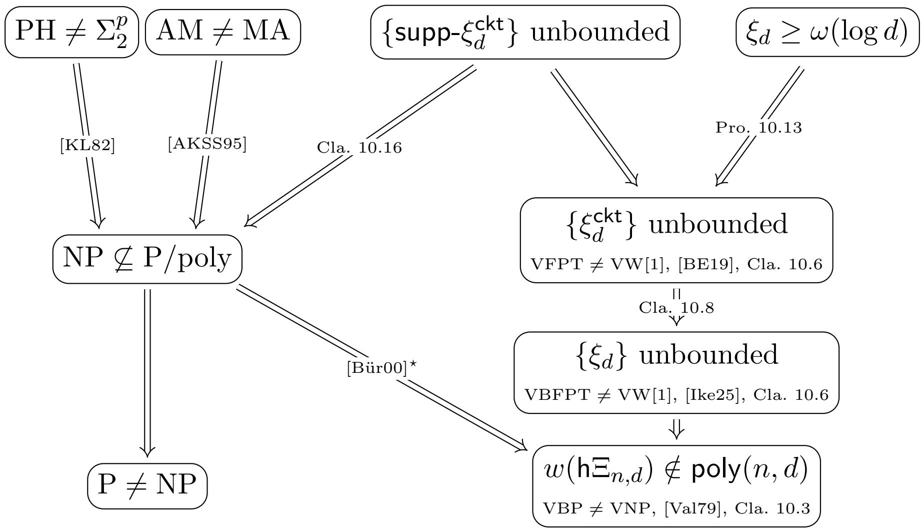

Figure 4: Formal implications between conjectures. The setting is $R = \mathbb{C}$, $A = \mathscr{S}$. *The [Bür00] implication is only known to hold under the assumption of the Generalized Riemann Hypothesis.

implies P $\neq$ NP or VBP $\neq$ VNP. Known lower-bound techniques are far from proving any of these conjectures: At the moment, our best lower bounds are just $\xi_d \geq 2$, proved via a simple argument (8.6).

For an explanation of the conjectures $\mathrm{PH} \neq \Sigma_2^{\mathrm{p}}$ and $\mathrm{AM} \neq \mathrm{MA}$ we refer the reader to [KL82] and [AKSS95], respectively. We will discuss the other conjectures and their relations in the rest of this section, as indicated by the edge labels in Figure 4. Implications without edge labels follow directly from the definitions.

An arithmetic circuit C is a directed acyclic graph whose indegree 0 vertices (called *inputs*) are labeled with affine linear polynomials, and whose other vertices (called *computation gates*) have indegree 2 and are each labeled with an operation $+, \times$. An arithmetic circuit computes a polynomial at every vertex by induction over the circuit structure. The size of C is defined as the number of the computation gates of C. For a matrix F of polynomials we define $\mathsf{cc}(F)$ to be the size of a smallest arithmetic circuit computing each matrix entry of F. We define

$$\gamma_d^{\mathsf{ckt}} := \inf\{k \mid \mathsf{cc}(\mathsf{h}\Gamma_{n,d}) \in O(n^k)\}, \qquad \xi_d^{\mathsf{ckt}} := \inf\{k \mid \mathsf{cc}(\mathsf{h}\Xi_{n,d}) \in O(n^k)\}.$$

Analogously to (8.2) we have

$$\gamma_d^{\mathsf{ckt}} \leq \xi_d^{\mathsf{ckt}} \leq \gamma_d^{\mathsf{ckt}} + 4. \tag{10.1}$$

A *weft 1* arithmetic circuit is allowed to have one so-called *large gate* on each path from a leaf to the output, where a large gate is a summation or multiplication gate with arbitrary indegree.

10.I Valiant's conjecture phrased in n and d

A sequence of multilinear homogeneous polynomials $(f_d)_{d\in\mathbb{N}}$ is called *Boolean-definable* if there exists a sequence of Boolean circuits $(B_d)_{d\in\mathbb{N}}$ with $|B_d| \in \mathsf{poly}(d)$ such that

$$\forall d: \; f_d = \sum_{B_d(\mathbf{b})=1} x^{\mathbf{b}}.$$

Valiant's original conjecture (in slightly different phrasing) states that

$$\text{there exists a Boolean-definable sequence } (f_d)_{d\in\mathbb{N}} \text{ with } \mathsf{ac}(f_d) \notin \mathsf{poly}(d). \tag{10.2}$$

10.3 Claim. (10.2) *holds if and only if* $w(\mathsf{h}\Gamma_{n,d}) \notin \mathsf{poly}(n,d)$. $\diamond$

Proof. The proof goes by contraposition.

If $w(\mathsf{h}\Gamma_{n,d}) \in \mathsf{poly}(n,d)$, then by Theorem 9.1 and by replacing every $\epsilon^{(i)}$ and every $\overline{x}_i$ with 1, we obtain $\mathsf{ac}(f_d) \in \mathsf{poly}(d)$ for every Boolean-definable sequence $(f_d)_{d\in\mathbb{N}}$, therefore (10.2) is false.

For the other proof direction, assume that (10.2) is false. Note that $\mathsf{h}\Gamma_{d,d}$ is Boolean-definable. Hence, $\mathsf{ac}(\mathsf{h}\Gamma_{d,d}) \in \mathsf{poly}(d)$. Therefore, $\mathsf{ac}(\mathsf{h}\Gamma_{n+d,n+d}) \in \mathsf{poly}(n,d)$. By setting some variables in $\mathsf{h}\Gamma_{n+d,n+d}$ to 1 and 0, we get $\mathsf{ac}(\mathsf{h}\Gamma_{n,d}) \in \mathsf{poly}(n,d)$. Using Proposition 8.22, we obtain $w(\mathsf{h}\Gamma_{n,d}) \in \mathsf{poly}(n,d)$. $\square$

10.II Algebraic fixed-parameter tractability

We define

$$\Gamma^{\boxtimes m}_{n\geq d} := \begin{cases} \Gamma^{\boxtimes m}_{n,d} & \text{if } n \geq d, \\ 0 & \text{otherwise.} \end{cases}$$

Note that the symbol $\geq$ in $\Gamma^{\boxtimes m}_{n\geq d}$ is purely syntactic, exactly as the comma in $\Gamma^{\boxtimes m}_{n,d}$. We set $\mathsf{h}\Gamma_{n\geq d} := \Gamma^{\boxtimes 2}_{n\geq d}$.

A doubly indexed family of (not necessarily homogeneous) polynomials $f_{n,k}$ is called a *parameterized p-family* (see [BE19, Def 4.1]) if the number of variables is $\mathsf{poly}(n)$ and the degree is $\mathsf{poly}(n,k)$. The number of variables of $\Gamma^{\boxtimes m}_{n\geq d}$ is less than n^{3m}, and the degree of $\Gamma^{\boxtimes m}_{n\geq d}$ is at most d, hence for every fixed m, $\Gamma^{\boxtimes m}_{n\geq d}$ is a parameterized p-family. The class VFPT consists of those parameterized p-families $f_{n,k}$ for which there exists a function $\phi : \mathbb{N} \to \mathbb{N}$ and $p \in \mathsf{poly}(n)$ with $\forall n,k: \mathsf{cc}(f_{n,k}) \leq \phi(k)\cdot p(n)$. Analogously, the class VBFPT consists of those parameterized p-families $f_{n,k}$ for which there exists a function $\phi : \mathbb{N} \to \mathbb{N}$ and $p \in \mathsf{poly}(n)$ with $\forall n,k: \mathsf{ac}(f_{n,k}) \leq \phi(k)\cdot p(n)$, see [Ike25]. A parameterized p-family $f_{n,k}$ is called *weft-1-definable* if there exists a constant $\delta \in \mathbb{N}$ such that $f_{n,k}$ can be written in the form (up to renaming variables)

$$f_{n,k} \;=\; \sum_{\substack{\mathbf{b} \\ |\mathbf{b}|=k}} h_n(x_1,\ldots,x_{p(n)},b_1,\ldots,b_{q(n)})$$

for $p(n) \in \mathsf{poly}(n)$ and $q(n) \in \mathsf{poly}(n)$ and h_n can be computed by a weft 1 depth δ arithmetic circuit of size $\mathsf{poly}(n)$. The class VW[1] consists of the weft-1-definable parameterized p-families. The VFPT $\neq$ VW[1] conjecture states that there exists a weft-1-definable parameterized p-family that is not in VFPT. This conjecture is equivalent to ${}_k\mathrm{Clique}_n \notin$ VFPT, which is the main result of [BE19] (for a slightly different variant of ${}_k\mathrm{Clique}_n$). The following lemma states that weft 1 is sufficient for any m-th power $\Gamma^{\boxtimes m}_{n\geq d}$.

10.4 Lemma. *For every fixed* m*:* $\Gamma^{\boxtimes m}_{n\geq d}$ *is weft-1-definable.* $\diamond$

Proof. We prove that $\mathsf{h}\Gamma_{n\geq d}$ is weft-1-definable. The proof for the other values of m is completely analogous. Let E denote the set of variables of $\mathsf{h}\Gamma_{n\geq d}$, and note $E=\emptyset$ or $|E|=(2n+(d-2)n^2)^2\leq n^6$. We say that the pair $(e_{i,1}\boxtimes e_{i,2}, e_{j,1}\boxtimes e_{j,2})$ is 1-*extendable* if $e_{i,1}$ and $e_{j,1}$ can be extended to a direct source-sink path by adding $d-2$ more edges. Analogously, define $(e_{i,1}\boxtimes e_{i,2}, e_{j,1}\boxtimes e_{j,2})$ to be 2-*extendable* if $e_{i,2}$ and $e_{j,2}$ can be extended to a direct source-sink path. The key property is that a size d set $S\subseteq E$ forms a pair of direct source-sink paths (i.e., a monomial in $\mathsf{h}\Gamma_{n\geq d}$) if and only if all elements of S are pairwise both 1-extendable and 2-extendable: This is clearly necessary, and it is also sufficient, because 1-extendable pairs in which $e_{i,1}$ and $e_{j,1}$ are adjacent determine a direct source-sink path in level 1 so that all non-adjacent pairs are also 1-extendable, and analogously for level 2. Note that $\mathbf{x}^{\mathbf{b}}=\prod_{j\in E}(x_jb_j+1-b_j)$, which is a product of depth 2 circuits.

$$\mathsf{h}\Gamma_{n\geq d} \;=\; \sum_{\substack{\mathbf{b}\\ |\mathbf{b}|=d}} \mathbf{x}^{\mathbf{b}} \left(\prod_{\substack{i,j\in E, i\neq j\\ (i,j)\text{ not 1-extendable}}} (1-b_ib_j)\right)\left(\prod_{\substack{i,j\in E, i\neq j\\ (i,j)\text{ not 2-extendable}}} (1-b_ib_j)\right).$$

This is a sum of one large product of depth 2 circuits, i.e., a weft 1 depth 4 circuit. The circuit size of the inner circuit is $O(|E|^2)\in\mathsf{poly}(n)$. Hence, $\mathsf{h}\Gamma_{n\geq d}$ is weft-1-definable. $\square$

10.5 Proposition. VFPT $\neq$ VW[1] *if and only if* $\mathsf{h}\Gamma_{n\geq d}\notin$ VFPT. *Moreover,* VBFPT $\neq$ VW[1] *if and only if* $\mathsf{h}\Gamma_{n\geq d}\notin$ VBFPT. $\diamond$

Proof. Using Lemma 10.4, if $\mathsf{h}\Gamma_{n\geq d}\notin$ VFPT, then $\mathsf{h}\Gamma_{n\geq d}$ is weft-1-definable but not in VFPT, which means VFPT $\neq$ VW[1] by definition. Analogously, $\mathsf{h}\Gamma_{n\geq d}\notin$ VBFPT implies VBFPT $\neq$ VW[1]. For the other direction, let $\mathsf{h}\Gamma_{n\geq d}\in$ VFPT with $\phi:\mathbb{N}\to\mathbb{N}$ and $\mathsf{cc}(\mathsf{h}\Gamma_{n\geq d})\leq\phi(d)\cdot p(n)$, where $p(n)\in\mathsf{poly}(n)$. Proposition 7.35 implies that for $d=\binom{k+1}{2}$ we have $\mathsf{cc}({}_k\mathrm{Clique}_n)\leq\phi(\binom{k+1}{2})\cdot p(n)$, which can be written as $\mathsf{cc}({}_k\mathrm{Clique}_n)\leq\phi'(k)\cdot p(n)$ for $p(n)\in\mathsf{poly}(n)$. Therefore, ${}_k\mathrm{Clique}_n\in$ VFPT, which implies VFPT $=$ VW[1] by [BE19]. For VBFPT the proof is analogous by noting that Lemma 4.9 in [BE19] works also for VBFPT. Alternatively, one can prove both cases without going via cliques in a similar way as Theorem 9.1. $\square$

10.6 Claim. VFPT $\neq$ VW[1] *if and only if* $\{\xi_d^{\mathsf{ckt}}\mid d\in\mathbb{N}\}$ *is unbounded. Moreover,* VBFPT $\neq$ VW[1] *if and only if* $\{\xi_d\mid d\in\mathbb{N}\}$ *is unbounded.* $\diamond$

Proof. We use Proposition 10.5. The proof goes by contrapositions.

Let $\mathsf{h}\Gamma_{n\geq d}\in$ VFPT, i.e., there exists ϕ and α with $\forall d,n:\mathsf{cc}(\mathsf{h}\Gamma_{n\geq d})\leq\phi(d)\cdot n^{\alpha}$. Hence, $\{\gamma_d^{\mathsf{ckt}}\}$ is bounded by α and thus $\{\xi_d^{\mathsf{ckt}}\}$ is bounded by $\alpha+4$, see (10.1).

Let $\{\xi_d^{\mathsf{ckt}}\}$ be bounded by $\alpha\in\mathbb{N}$, hence $\{\gamma_d^{\mathsf{ckt}}\}$ is also bounded by α. Then $\forall d:\mathsf{cc}(\mathsf{h}\Gamma_{n\geq d})\in O(n^{\alpha+0.01})\leq\phi(d)\cdot n^{\alpha+0.01}$ for some $\phi(d)$. Hence, $\mathsf{h}\Gamma_{n\geq d}\in$ VFPT.

The proof for VBFPT is completely analogous. $\square$

10.7 Lemma. *Let f be a homogeneous degree d polynomial. Then $\mathsf{cc}(f)\leq 2d\big(\mathsf{ac}(f)\big)^2$.* $\diamond$

Proof. By Proposition 8.22 we have $w(f)\leq\mathsf{ac}(f)$. The evaluation of an ABP of width n and degree d can naively be implemented as an arithmetic circuit. Every edge that is not incident to the source gives rise to a product gate, and every vertex besides the source and besides the vertices in vertex layer 1 gives rise to $n-1$ summation gates. This gives at most $(d-2)n^2+n$ many product gates and $(n-1)((d-2)n+1)$ many summation gates. In total, these are at most $2n^2d$ many gates. $\square$

While an increase of 1 in ac requires several arithmetic operations to realize in a circuit, not all algorithmic tricks in the circuit model have known equivalences in ABPs, for example an algebraic variant of the technique in [NP85, §2] can be applied to get better upper bounds on $\mathsf{cc}(\mathsf{h}\Xi_{n,d})$, but it is unclear how to apply it to ac.

10.8 Claim. VBFPT $\subseteq$ VFPT. $\diamond$

Proof. For $f_{n,d} \in$ VBFPT we have $\mathsf{ac}(f_{n,d}) \le n^\alpha \phi(d)$ for some $\phi : \mathbb{N} \to \mathbb{N}$ and some $\alpha \in \mathbb{N}$. Hence, by Lemma 10.7, $\mathsf{cc}(f_{n,d}) \le 2d\big(n^\alpha\phi(d)\big)^2 \le n^{2\alpha}\phi'(d)$ for some $\phi' : \mathbb{N} \to \mathbb{N}$. Therefore $f_{n,d} \in$ VFPT. $\square$

An arithmetic circuit for which homogeneous polynomials are computed at each vertex is called a *homogeneous arithmetic circuit*. A homogeneous arithmetic circuit is defined to compute f if it computes each homogeneous component of f. Let $\mathsf{hcc}(f)$ denote the size of a smallest homogeneous arithmetic circuit computing f.

10.9 Lemma. *If f (not necessarily homogeneous) of degree d is computed by an arithmetic circuit of size s, then there exists a homogeneous arithmetic circuit of size $s(d+1)^2$ that computes f.* $\diamond$

Proof. We replace every vertex by $d+1$ many vertices, each computing a homogeneous component of degree $0, \ldots, d$, discarding higher degrees. An addition gate gets replaces by $d+1$ addition gates. Since $\big((\sum_{i=0}^d h^{(i)}) \cdot (\sum_{i=0}^d f^{(i)})\big)^{(k)} = \sum_{i=0}^k h^{(i)} f^{(k-i)}$, a multiplication gate that is homogenized in this way uses $\sum_{k=0}^d (k+1) = (d+1)(d+2)/2$ many multiplication gates and $\sum_{k=0}^d k = d(d+1)/2$ addition gates, hence $(d+1)^2$ gates in total. $\square$

10.10 Lemma. *If f of degree $d \ge 2$ (not necessarily homogeneous) is computed by a homogeneous arithmetic circuit of size $s \ge 2$, then f is computed by an arithmetic circuit of depth at most $19\log(s)\log(d)$.* $\diamond$

Proof. This is a weaker result than [VSBR83], but it suffices here and has a much shorter proof, as pointed out in [Sap15]. The degree $\deg(v)$ of a vertex v in an arithmetic circuit is the degree of the (not necessarily homogeneous) polynomial computed at v. Given a homogeneous circuit C of size s computing f such that each gate computes a homogeneous polynomial of degree $\le d$, let

$$D = \{v \in C \mid \tfrac{d}{3} < \deg(v) \le \tfrac{2d}{3}\}.$$

We convert C into C' by replacing every subcircuit C_v with $v \in D$ by a new input variable y_v, and call the computed polynomial f'. Collecting y-variables, we write

$$f' = \sum_{i<j} A_{i,j} y_i y_j \, + \sum_i A_{i,i} y_i^2 \, + \sum_i B_i y_i \, + \, p, \tag{10.11}$$

with each $A_{i,j}, B_i, p$ not involving any y-variables, and each $\deg(A_{i,j}), \deg(B_i) \le 2d/3$, $\deg(p) \le \frac{d}{3}$. For a set $S \subseteq \mathbb{N}$, we write $f'_{S=a}$ for the partial evaluation in which we set y_i to a if $i \in S$, and y_i to 0 otherwise. We write $f'_\emptyset := f'_{\emptyset=1}$. Note that $p = f'_\emptyset$, $B_i = -\frac{1}{2}(f'_{\{i\}=2} - 4f'_{\{i\}=1} + 3p)$, $A_{i,i} = \frac{1}{2}(f'_{\{i\}=2} - 2f'_{\{i\}=1} + p)$, and $A_{i,j} = f'_{\{i,j\}=1} - A_{i,i} - A_{j,j} - B_i - B_j - p$. Substituting these into (10.11) gives a circuit that is a circuit for f stacked on top of polynomials of degree at most $2d/3$ that have size at most s, because summation over $\{(i,j) \mid i < j\}$ has a circuit of depth $\lceil 2\log(s)\rceil$. We apply the construction recursively to these polynomials. The resulting depth satisfies $\mathrm{Depth}(d) = \mathrm{Depth}(2d/3) + 2\log(s) + 9$, hence $\mathrm{Depth}(d) \le \log_{3/2}(d)(2\log(s)+9) \le 19\log(s)\log(d)$. $\square$

An arithmetic circuit whose graph structure is a tree is called an arithmetic formula. We define $\mathsf{fc}(f)$ to be the size of a smallest arithmetic formula computing f.

10.12 Lemma. *For every (not necessarily homogeneous) polynomial f, we have $\mathsf{ac}(f) \le \mathsf{fc}(f)$.* $\diamond$

Proof. The proof is by induction. For linear polynomials f, $\mathsf{fc}(f) = 0 = \mathsf{ac}(f)$. For a sum $f+g$ of minimal formula size, we have $\mathsf{fc}(f+g) = \mathsf{fc}(f) + \mathsf{fc}(g) + 1$, but $\mathsf{ac}(f+g) \leq \mathsf{ac}(f) + \mathsf{ac}(g)$ by taking minimal size ABPs for f and g and identifying their sources and identifying their sinks. For a product fg of minimal formula size, we have $\mathsf{fc}(fg) = \mathsf{fc}(f) + \mathsf{fc}(g) + 1$ and $\mathsf{ac}(fg) \leq \mathsf{ac}(f) + \mathsf{ac}(g) + 1$ by taking minimal size ABPs for f and g and identifying the sink for f with the source for g (note that ac does not count sources or sinks, and this creates a new vertex that is not a source or sink). $\square$

10.13 Proposition. *If* $\xi_d \geq \omega(\log d)$*, then* VFPT $\neq$ VW[1]. $\diamond$

Proof. The proof goes by contraposition. Let VFPT = VW[1]. Then $\mathsf{h}\Gamma_{n \geq d} \in$ VFPT. Hence, there exists $\alpha \in \mathbb{N}$ and a function ϕ such that $\forall n, d : \mathsf{cc}(\mathsf{h}\Gamma_{n\geq d}) \leq \phi(d) \cdot n^\alpha$. Lemma 10.9 gives $\forall n, d : \mathsf{hcc}(\mathsf{h}\Gamma_{n\geq d}) \leq \phi'(d) \cdot n^\alpha$ with $\phi'(d) = (d+1)^2 \cdot \phi(d)$. We apply Lemma 10.10 to obtain a circuit for $\mathsf{h}\Gamma_{n\geq d}$ of depth $19 \log(d) \log(\phi'(d) n^\alpha)$. We expand this arithmetic circuit into an arithmetic formula of size $2^{19 \log(d) \log(\phi'(d) n^\alpha)} = (\phi'(d) n^\alpha)^{19 \log d}$. Using Lemma 10.12, we obtain $\mathsf{ac}(\mathsf{h}\Gamma_{n\geq d}) \leq (\phi'(d) n^\alpha)^{19 \log d}$. We apply Proposition 8.22 to obtain $w(\mathsf{h}\Gamma_{n\geq d}) \leq (\phi'(d) n^\alpha)^{19 \log d}$. For a fixed d, this is in $O(n^{19\alpha \log d})$. Hence, $\gamma_d \leq O(\log d)$, and thus also $\xi_d \leq O(\log d)$. $\square$

The following proposition is analogous to Proposition 7.32, and it shows that it is sufficient to study second-order computation only and not also higher orders.

10.14 Proposition. *There exists* $\phi : \mathbb{N} \to \mathbb{N}$*,* $\phi(d) \geq d$*, with* $w(\mathsf{h}\Xi_{n,d}) \leq \phi(d)\mathsf{poly}(n)$ *if and only if for every* m *there exists* $\varphi : \mathbb{N} \to \mathbb{N}$*,* $\varphi(d) \geq d$*, with* $w(\mathsf{H}_{n,d}^{\boxtimes m} \boxtimes \Xi_{n,d}) \leq \varphi(d)\mathsf{poly}(n)$*.* $\diamond$

Proof. One direction is clear. For the other direction, fix any m. Let $\phi : \mathbb{N} \to \mathbb{N}$ and $k \in \mathbb{N}$ such that $\forall d, n : w(\mathsf{h}\Xi_{n,d}) \leq n^k \phi(d)$, in other words $\mathsf{h}\Xi_{n,d} \leq \Xi_{n^k\phi(d),d}$.

$$\begin{aligned}
\mathsf{H}_{n,d}^{\boxtimes m} \boxtimes \Xi_{n,d} &= \mathsf{H}_{n,d}^{\boxtimes m-1} \boxtimes (\mathsf{H}_{n,d} \boxtimes \Xi_{n,d}) \leq \mathsf{H}_{n,d}^{\boxtimes m-1} \boxtimes \Xi_{n^k\phi(d),d} \\
&= \mathsf{H}_{n,d}^{\boxtimes m-2} \boxtimes (\mathsf{H}_{n,d} \boxtimes \Xi_{n^k\phi(d),d}) \leq \mathsf{H}_{n,d}^{\boxtimes m-2} \boxtimes (\mathsf{H}_{n^k\phi(d),d} \boxtimes \Xi_{n^k\phi(d),d}) \\
&\leq \mathsf{H}_{n,d}^{\boxtimes m-2} \boxtimes \Xi_{n^{k^2}\phi(d)^{k+1},d} \leq \mathsf{H}_{n,d}^{\boxtimes m-3} \boxtimes \Xi_{n^{k^3}\phi(d)^{k^2+k+1},d} \\
&\leq \cdots \leq \Xi_{n^{k^m}\phi(d)^{k^{m-1}+k^{m-2}+\cdots+1},d} .
\end{aligned}$$

Since k and m are constant, this means $w(\mathsf{H}_{n,d}^{\boxtimes m} \boxtimes \Xi_{n,d}) \leq \mathsf{poly}(n)\varphi(d)$ with $\varphi(d) = \phi(d)^{k^{m-1}+k^{m-2}+\cdots+1}$. $\square$

10.III The connection to Boolean circuits

We define $\exists\mathsf{h}\Gamma_{n\geq d} : \{0,1\}^{n^4 d^2} \to \{0,1\}$ via $\exists\mathsf{h}\Gamma_{n\geq d}(\mathbf{b}) = \min\{1, \mathsf{h}\Gamma_{n\geq d}(\mathbf{b})\}$. We phrase the conjectures NP $\not\subseteq$ P/poly and P $\neq$ NP via their negations as follows. We have NP $\subseteq$ P/poly if and only if there exists a sequence of Boolean circuits of size $\mathsf{poly}(d)$ that compute $\exists\mathsf{h}\Gamma_{d\geq d}$. We have P = NP if and only if there exists a sequence of Boolean circuits C_d of size $\mathsf{poly}(d)$ that compute $\exists\mathsf{h}\Gamma_{d\geq d}$ with the additional requirement that there exists an algorithm A that on input d runs in time $\mathsf{poly}(d)$ and outputs the adjacency matrix and gate labels of C_d. For $h \in \mathscr{S}_d$ we define

$$\mathsf{supp\text{-}cc}(h) := \min\{\mathsf{cc}(f) \mid f \in \mathscr{S},\ f(\mathbf{b}) = 0 \text{ iff } h(\mathbf{b}) = 0\}.$$

$$\mathsf{supp}\text{-}\gamma_d^{\mathsf{ckt}} := \inf\{k \mid \mathsf{supp\text{-}cc}(\mathsf{h}\Gamma_{n\geq d}) \in O(n^k)\} = \limsup_{n\to\infty} \big(\log_n(\mathsf{supp\text{-}cc}(\mathsf{h}\Gamma_{n\geq d}))\big).$$

More generally, for $H \in \mathrm{Mat}_k(\mathscr{S}_d)$ we define

$$\mathsf{supp\text{-}cc}(H) := \max\{\mathsf{supp\text{-}cc}([H]_{i,j})\}$$

$$\mathsf{supp}\text{-}\xi_d^{\mathsf{ckt}} := \inf\{k \mid \mathsf{supp\text{-}cc}(\mathsf{h}\Xi_{n\geq d}) \in O(n^k)\} = \limsup_{n\to\infty} \big(\log_n(\mathsf{supp\text{-}cc}(\mathsf{h}\Xi_{n\geq d}))\big).$$

Analogously to (8.2) and (10.1), we have

$$\mathsf{supp}\text{-}\gamma_d^{\mathsf{ckt}} \leq \mathsf{supp}\text{-}\xi_d^{\mathsf{ckt}} \leq \mathsf{supp}\text{-}\gamma_d^{\mathsf{ckt}} + 4. \tag{10.15}$$

10.16 Claim. *If* $\mathsf{supp}\text{-}\xi_d^{\mathsf{ckt}}$ *is unbounded, then* NP $\not\subseteq$ P/poly. $\diamond$

Proof. We use (10.15) and argue about $\mathsf{supp}\text{-}\gamma_d^{\mathsf{ckt}}$. The proof goes by contraposition. For $f : \{0,1\}^N \to \{0,1\}$, let $\mathsf{bc}(f)$ denote the size of a smallest Boolean circuit computing f. Let NP $\subseteq$ P/poly. Since $\exists\mathsf{h}\Gamma_{d \geq d} \in$ NP, it follows that there exists $\alpha \in \mathbb{N}$ with $\forall d : \mathsf{bc}(\exists\mathsf{h}\Gamma_{d\geq d}) \leq d^\alpha$. Hence, $\forall d, n : \mathsf{bc}(\exists\mathsf{h}\Gamma_{n\geq d}) \leq n^\alpha$. We arithmetize the Boolean circuit by replacing `NOT` with $1-x$, `AND` with $x \cdot y$, and `OR` with $x + y - xy$. Evaluated on the Boolean hypercube, this circuit has the same support as $\exists\mathsf{h}\Gamma_{n\geq d}$. The size of the arithmetic circuit is at most 4 times larger than the size of the Boolean circuit. Hence, $\mathsf{supp}\text{-}\gamma_d^{\mathsf{ckt}}$ is bounded by α. $\square$

11 Stabilizers and other properties

In this section we have $R = \mathbb{C}$ and $A = \mathscr{S}$ (unless stated otherwise).

11.I Different settings and their properties

Given a group G acting linearly on a finite dimensional complex vector space V, the *stabilizer* of $f \in V$ is defined as $\mathrm{stab}_G(f) = \{g \in G \mid gf = f\}$. For any subgroup $H \leq G$, the invariant space V^H is defined as $\{v \in V \mid \forall h \in H : hv = v\}$. We say that f is *characterized by its symmetries* if $\dim V^{\mathrm{stab}_G(f)} = 1$.

Recall that $\Gamma_{n,d}$ is the bottom-right entry of $\Xi_{n,d}$. Let $\mathrm{tr}\text{-}\Xi_{n,d} := \mathrm{tr}(\Xi_{n,d})$ be the trace. In this section we analyze the objects $\mathrm{tr}\text{-}\Xi_{n,d}$, $\Gamma_{n,d}$, and $\Xi_{n,d}$ by the following properties:

1. the object's endomorphism orbit is topologically closed for $A = \mathscr{T}$, by which we mean $\{T(\Xi_{n,d}) \mid T \in \mathrm{Mat}_n(\mathbb{C}) \times \mathrm{End} \times \mathrm{Mat}_n(\mathbb{C})\}$ is closed, and analogous statements for the other objects,
2. the object has a reductive stabilizer,
3. the object is characterized by its stabilizer,
4. the object has a connected stabilizer,
5. on plugging in identity matrices the object is invertible in every characteristic

- It is clear that $\mathrm{tr}\text{-}\Xi_{n,d}$ lacks property 5, because $\mathrm{tr}\text{-}\Xi_{n,d}(\mathrm{id}) = n$. Moreover, $\mathrm{tr}\text{-}\Xi_{n,d}$ also lacks property 1, see [BIM+20]. The fact that $\mathrm{tr}\text{-}\Xi_{n,d}$ also lacks property 4 was proved in [Ges16].
- $\Gamma_{n,d}$ has property 1, which was proved in [Nis91], see also [For16]. Notably, $\Gamma_{n,d}$ has property 5. However, $\Gamma_{n,d}$ lacks property 2, which can be seen by calculating for $d = 3$, $n = 2$, the Lie algebra that is the kernel of the Lie algebra action of $\mathfrak{gl}_{2n+(d-2)n^2}$ on $\Gamma_{n,d}$, and then verifying non-reductivity, for example with GAP.
- For $\Xi_{n,d}$ we have property 1, see Theorem 8.8. Clearly, $\Xi_{n,d}$ has property 5. Properties 2, 3, and 4 are concerned with the stabilizer, hence they depend on the group action. Let $\mu_d = \{\mathrm{diag}(\alpha, \ldots, \alpha) \mid \alpha^d = 1\} \subset \mathrm{GL}_{n^2 d}$, and let $G_{n,d} = \mathrm{GL}_{n^2 d}/\mu_d$, which is a connected reductive linear algebraic Lie group with Lie algebra $\mathfrak{gl}_{n^2 d}$. Note that $G_{n,d}$ acts faithfully on $S^d\mathbb{C}^{n^2 d}$. We distinguish 3 different group actions:
    1. The action of $G_{n,d}$ on $\mathrm{Mat}_n(S^d\mathbb{C}^{n^2 d})$ given by $g(e_i \otimes f \otimes e_j^*) = e_i \otimes g(f) \otimes e_j^*$,
    2. The action of $\mathrm{GL}_n \times G_{n,d}$ on $\mathrm{Mat}_n(S^d\mathbb{C}^{n^2 d})$ given by $(g_{\mathrm{outer}}, g_{\mathrm{inner}})(e_i \otimes f \otimes e_j^*) = (g_{\mathrm{outer}}\, e_i) \otimes g_{\mathrm{inner}}(f) \otimes (e_j^*\, g_{\mathrm{outer}}^{-1})$,
    3. The action of $\mathrm{GL}_n \times G_{n,d} \times \mathrm{GL}_n$ on $\mathrm{Mat}_n(S^d\mathbb{C}^{n^2 d})$ given by $(g_{\mathrm{left}}, g_{\mathrm{inner}}, g_{\mathrm{right}})(e_i \otimes f \otimes e_j^*) = (g_{\mathrm{left}} e_i) \otimes g_{\mathrm{inner}}(f) \otimes (e_j^* g_{\mathrm{right}}^{\mathsf{T}})$.

Under the first action, $\Xi_{n,d}$ lacks property 3, because if $g(\Xi_{n,d}) = \Xi_{n,d}$, then also $g(\text{tr-}\Xi_{n,d}) = \text{tr-}\Xi_{n,d}$, but $\Xi_{n,d} \neq \text{tr-}\Xi_{n,d}$. This does not contradict the fact that $\text{tr-}\Xi_{n,d}$ has property 3, because the ambient space under consideration there is $S^d\mathbb{C}^{n^2 d}$, which is smaller than $\mathrm{Mat}_n(S^d\mathbb{C}^{n^2 d})$.

Under the third action, note that the actions of the three factors of $\mathcal{G} := \mathrm{GL}_n \times G_{n,d} \times \mathrm{GL}_n$ commute. Moreover, for every $g_{\text{left}} \in \mathrm{GL}_n$ there exists $h_{g_{\text{left}}} \in G_{n,d}$ such that $(g_{\text{left}}, \mathrm{id}, \mathrm{id})\Xi_{n,d} = (\mathrm{id}, h_{g_{\text{left}}}, \mathrm{id})\Xi_{n,d}$, which is proved in the same way as in Lemma 6.7. Analogously, for every $g_{\text{right}} \in \mathrm{GL}_n$ there exists $h'_{g_{\text{right}}} \in G_{n,d}$ such that $(\mathrm{id}, \mathrm{id}, g_{\text{right}})\Xi_{n,d} = (\mathrm{id}, h'_{g_{\text{right}}}, \mathrm{id})\Xi_{n,d}$. In other words, for every $g \in \mathrm{GL}_n \times \mathrm{GL}_n$ there exists $h_g \in G_{n,d}$ with $g\Xi_{n,d} = h_g\Xi_{n,d}$. By the construction in Lemma 6.7 we also see that $\mathrm{GL}_n \times \mathrm{GL}_n \to G_{n,d}$, $g \mapsto h_g$, is a group homomorphism, in particular $h_{g^{-1}} = (h_g)^{-1}$.

$$\begin{aligned}
\mathrm{stab}_{\mathcal{G}}(\Xi_{n,d}) &= \{(g,h) \mid gh\Xi_{n,d} = \Xi_{n,d}\} \\
&= \{(g,h) \mid h\Xi_{n,d} = g^{-1}\Xi_{n,d}\} \\
&= \{(g,h) \mid h\Xi_{n,d} = h_{g^{-1}}\Xi_{n,d}\} \\
&= \{(g,h) \mid h_{g^{-1}}^{-1} h\Xi_{n,d} = \Xi_{n,d}\} \\
&= \{(g,h) \mid h_g h\Xi_{n,d} = \Xi_{n,d}\} \\
&= \{(g,h) \mid h_g h \in \mathrm{stab}_{G_{n,d}}(\Xi_{n,d})\} \\
&= \{(g,h_{g^{-1}}h) \mid h \in \mathrm{stab}_{G_{n,d}}(\Xi_{n,d})\} \qquad (11.1)
\end{aligned}$$

Thus, in order to determine $\mathrm{stab}_{\mathcal{G}}(\Xi_{n,d})$, it suffices to determine $\mathrm{stab}_{G_{n,d}}(\Xi_{n,d})$.

Every h stabilizing the matrix $\Xi_{n,d}$ also stabilizes its trace, therefore $\mathrm{stab}_{G_{n,d}}(\Xi_{n,d}) \subseteq \mathrm{stab}_{G_{n,d}}(\text{tr-}\Xi_{n,d})$, which has been determined in [Ges16] as the group that maps

$$(X^{(1)}, \ldots, X^{(d)})$$

to

$$(h_d^{-1}X^{(1)}h_1, h_1^{-1}X^{(2)}h_2, \ldots, h_{d-1}^{-1}X^{(d)}h_d) \qquad (11.2)$$

taken in semidirect product with the dihedral group that is generated by the maps that send

$$(X^{(1)}, \ldots, X^{(d)})$$

to

$$({X^{(d)}}^{\mathsf{T}}, \ldots, {X^{(1)}}^{\mathsf{T}}) \qquad (11.3)$$

and

$$(X^{(1)}, \ldots, X^{(d)})$$

to

$$(X^{(2)}, \ldots, X^{(d)}, X^{(1)}). \qquad (11.4)$$

While (11.2) stabilizes $\Xi_{n,d}$ if and only if h_d is a nonzero multiple of the identity, none of the $2d-1$ many nontrivial elements generated by (11.3) and (11.4) stabilize $\Xi_{n,d}$, which can be seen by mapping everything but the top-left 2×2 matrices to zero, and observing that

$$\begin{pmatrix}1 & 1\\ 0 & 1\end{pmatrix}^{i-1} \cdot \begin{pmatrix}1 & 0\\ 0 & 2\end{pmatrix} \cdot \begin{pmatrix}1 & 1\\ 0 & 1\end{pmatrix}^{d-i} = \begin{pmatrix}1 & d+i-2\\ 0 & 2\end{pmatrix}$$

and

$$\begin{pmatrix}1 & 0\\ 1 & 1\end{pmatrix}^{i-1} \cdot \begin{pmatrix}1 & 0\\ 0 & 2\end{pmatrix} \cdot \begin{pmatrix}1 & 0\\ 1 & 1\end{pmatrix}^{d-i} = \begin{pmatrix}1 & 0\\ 2d-i-1 & 2\end{pmatrix}$$

which are $2d$ many different matrices. Hence, from (11.1) we conclude that

$$\begin{array}{ll}\text{stab}_{\mathcal{G}}(\Xi_{n,d}) \text{ consists of those } g \in \mathcal{G} \text{ that send} & \\ \qquad (X^{(1)}, \dots, X^{(d)}) & \\ \text{to} & (11.5) \\ \qquad (h_0^{-1}X^{(1)}h_1, h_1^{-1}X^{(2)}h_2, \dots, h_{d-1}^{-1}X^{(d)}h_d), & \\ \text{while simultaneously multiplying the matrix with } h_0 \text{ from the left and } h_d^{-1} & \\ \text{from the right.} & \end{array}$$

Thus,

$$\text{stab}_{\mathcal{G}}(\Xi_{n,d}) \simeq \text{GL}_n^{d+1},$$

which is reductive and connected, hence satisfying properties 2 and 4. To see that $\Xi_{n,d}$ is characterized by its stabilizer, we have to determine $(\mathbb{C}^n \otimes S^d\mathbb{C}^{n^2d} \otimes \mathbb{C}^{n*})^{\text{stab}\Xi_{n,d}}$, but we perform a more general analysis by determining the representation-theoretic multiplicities in $\mathbb{C}[\mathcal{G}\Xi_{n,d}] \simeq \mathbb{C}[\mathcal{G}]^{\text{stab}_{\mathcal{G}}\Xi_{n,d}}$, see [BLMW11, §4], [Hüt17, A.1.9]. Let $\{\lambda\}_n$ denote the irreducible GL_n-representation of type λ. We write $\{\lambda\}$ if n is clear from the context. Let $N = n^2d$. The irreducible polynomial representations of $G_{n,d}$ are the same as for GL_{n^2d} with the only difference that d must divide $\lambda_1 + \cdots + \lambda_{n^2d}$. Let $H := \text{stab}_{\mathcal{G}}(\Xi_{n,d})$. We write $\mathbb{C}^{n^2d} = (V_0 \otimes V_1^*) \oplus \cdots \oplus (V_{d-1} \otimes V_d^*)$, where $V_j \simeq \mathbb{C}^n$. We first restrict

$$\{\mu\}\downarrow^{\text{GL}_N}_{\text{GL}(V_0\otimes V_1^*)\times\cdots\times\text{GL}(V_{d-1}\otimes V_d^*)} = \bigoplus_{\rho^{(1)},\dots,\rho^{(d)}\vdash n^2} c^{\mu}_{\rho^{(1)},\dots,\rho^{(d)}} \{\rho^{(1)}\} \otimes \cdots \otimes \{\rho^{(d)}\},$$

where $c^{\mu}_{\rho^{(1)},\dots,\rho^{(d)}}$ is the multi-Littlewood-Richardson coefficient. We restrict further:

$$\begin{aligned}&\{\mu\}\downarrow^{\text{GL}_N}_{\text{GL}(V_0)\times\text{GL}(V_1^*)\times\text{GL}(V_1)\times\cdots\times\text{GL}(V_{d-1}^*)\times\text{GL}(V_{d-1})\times\text{GL}(V_d^*)} \\ &= \bigoplus_{\substack{\rho^{(1)},\dots,\rho^{(d)} \\ \iota^{(0)},\dots,\iota^{(d-1)} \\ \kappa^{(1)},\dots,\kappa^{(d)}}} c^{\mu}_{\rho^{(1)},\dots,\rho^{(d)}} k(\iota^{(0)},\rho^{(1)},\kappa^{(1)})\cdots k(\iota^{(d-1)},\rho^{(d)},\kappa^{(d)}) \{\iota^{(0)}\}\otimes\{\kappa^{(1)}\}\otimes\cdots\otimes\{\iota^{(d-1)}\}\otimes\{\kappa^{(d)}\},\end{aligned} \tag{11.6}$$

where $k(\lambda,\mu,\nu)$ is the Kronecker coefficient. Rename: $\lambda = \kappa^{(0)}$ and $\nu = \iota^{(d)}$, and take H-invariants, which forces $\forall i : \kappa^{(i)} = \iota^{(i)}$:

$$\left(\{\kappa^{(0)}\} \otimes \{\mu\} \otimes \{\kappa^{(d)}\}\right)^H = \bigoplus_{\substack{\rho^{(1)},\dots,\rho^{(d)} \\ \kappa^{(1)},\dots,\kappa^{(d-1)}}} c^{\mu}_{\rho^{(1)},\dots,\rho^{(d)}} k(\kappa^{(0)},\rho^{(1)},\kappa^{(1)})\cdots k(\kappa^{(d-1)},\rho^{(d)},\kappa^{(d)})\,\mathbb{C}$$

Hence, all $\kappa^{(i)}$ have the same number of boxes, which is also the same as all $\rho^{(j)}$, which is $|\mu|/d$. Writing this in multi-Kronecker notation, we have

$$\left(\{\kappa^{(0)}\} \otimes \{\mu\} \otimes \{\kappa^{(d)}\}\right)^H = \bigoplus_{\rho^{(1)},\dots,\rho^{(d)}} c^{\mu}_{\rho^{(1)},\dots,\rho^{(d)}} k(\kappa^{(0)},\rho^{(1)},\rho^{(2)},\dots,\rho^{(d-1)},\rho^{(d)},\kappa^{(d)})\,\mathbb{C} \tag{11.7}$$

To see that $\Xi_{n,d}$ is characterized by its stabilizer, we determine

$$(\mathbb{C}^n \otimes S^d\mathbb{C}^{n^2d} \otimes \mathbb{C}^n)^H = (\{(1)\} \otimes \{(d)\} \otimes \{(1)\})^H$$

for which the above calculation has only 1 summand, and $|\mu|/d = 1$, the multi-Littlewood-Richardson coefficient is 1, and the multi-Kronecker coefficient is 1, thus $\dim(\mathbb{C}^n \otimes S^d\mathbb{C}^{n^2d} \otimes \mathbb{C}^n)^H = 1$. In conclusion, $\Xi_{n,d}$ with the third action has all 5 properties. We now have a look at the second action, i.e., the conjugation action. In this case, the stabilizer H is smaller and when taking H-invariants in (11.6), instead of getting a 1-dimensional invariant space, we have to take GL_n-invariants of $\mathbb{C}^n \otimes (\mathbb{C}^n)^* \otimes \mathbb{C}^n \otimes (\mathbb{C}^n)^*$, which is 2-dimensional, because $\mathbb{C}^n \otimes \mathbb{C}^n = S^2\mathbb{C}^n \oplus \bigwedge^2\mathbb{C}^n$. Hence, for this action $\Xi_{n,d}$ does not have property 3.

In conclusion, only $\Xi_{n,d}$, and only with the third action has all 5 properties. In fact, properties 1, 2, and 3 are sufficient for ruling out all other settings. We will continue working in this setting throughout §11.

11.II The ambient space

Since we are working with matrices of polynomials, it is worth recording the representation theory of the space $S^\delta(\mathbb{C}^a \otimes S^d\mathbb{C}^N \otimes \mathbb{C}^b)$ with respect to the action of the group $\mathrm{GL}_a \times \mathrm{GL}_N \times \mathrm{GL}_b$.

$$
\begin{aligned}
S^\delta(\mathbb{C}^a \otimes S^d\mathbb{C}^N \otimes \mathbb{C}^b) &= \big(\textstyle\bigotimes^\delta(\mathbb{C}^a \otimes S^d\mathbb{C}^N \otimes \mathbb{C}^b)\big)^{\mathfrak{S}_\delta} \\
&= \big(\textstyle\bigotimes^\delta \mathbb{C}^a \otimes \bigotimes^\delta S^d\mathbb{C}^N \otimes \bigotimes^\delta \mathbb{C}^b\big)^{\mathfrak{S}_\delta} \\
&= \Big(\big(\bigoplus_{\lambda\vdash_a\delta}\{\lambda\}_a \otimes [\lambda]\big) \otimes \textstyle\bigotimes^\delta S^d\mathbb{C}^N \otimes \big(\bigoplus_{\nu\vdash_b\delta}\{\nu\}_b \otimes [\nu]\big)\Big)^{\mathfrak{S}_\delta} \\
&\overset{\text{Pieri}}{=} \Big(\big(\bigoplus_{\lambda\vdash_a\delta}\{\lambda\}_a \otimes [\lambda]\big) \otimes \big(\bigoplus_{\mu\vdash_N d\delta}\{\mu\}_N \otimes K_{\mu,d^\delta}\big) \otimes \big(\bigoplus_{\nu\vdash_b\delta}\{\nu\}_b \otimes [\nu]\big)\Big)^{\mathfrak{S}_\delta} \\
&= \Big(\big(\bigoplus_{\lambda\vdash_a\delta}\{\lambda\}_a \otimes [\lambda]\big) \otimes \big(\bigoplus_{\mu\vdash_N d\delta}\{\mu\}_N \otimes \big(\bigoplus_{\kappa\vdash\delta}[\kappa]^{\oplus a_\mu(\kappa[d])}\big)\big) \otimes \big(\bigoplus_{\nu\vdash_b\delta}\{\nu\}_b \otimes [\nu]\big)\Big)^{\mathfrak{S}_\delta} \\
&= \bigoplus_{\substack{\lambda\vdash_a\delta\\ \mu\vdash_N d\delta\\ \kappa\vdash\delta\\ \nu\vdash_b\delta}} \{\lambda\}_a \otimes \{\mu\}_N \otimes \{\nu\}_b \otimes \mathbb{C}^{a_\mu(\kappa[d])} \otimes \big([\lambda]\otimes[\kappa]\otimes[\nu]\big)^{\mathfrak{S}_\delta} \\
&= \bigoplus_{\substack{\lambda\vdash_a\delta\\ \mu\vdash_N d\delta\\ \kappa\vdash\delta\\ \nu\vdash_b\delta}} \big(\{\lambda\}_a \otimes \{\mu\}_N \otimes \{\nu\}_b\big)^{a_\mu(\kappa[d])k(\lambda,\kappa,\nu)} \\
&= \bigoplus_{\substack{\lambda\vdash_a\delta\\ \mu\vdash_N d\delta\\ \nu\vdash_b\delta}} \big(\{\lambda\}_a \otimes \{\mu\}_N \otimes \{\nu\}_b\big)^{\sum_{\kappa\vdash\delta} a_\mu(\kappa[d])k(\lambda,\kappa,\nu)}
\end{aligned}
$$

One can check with a computer that for $\delta \le 6$ these multiplicities never exceed the multiplicities in (11.7).

11.III $\mathsf{h}\Xi_{n,d}$ versus $\mathsf{h}\Xi_{n,d,\pi}$

Let $H = \mathrm{stab}_{\mathrm{GL}_n\times G_{n,d}\times\mathrm{GL}_n}\Xi_{n,d}$. By construction, we have $H\times H \subseteq \mathrm{stab}(\mathsf{h}\Xi_{n,d})$. However, this does not characterize $\mathsf{h}\Xi_{n,d}$, since each $\mathsf{h}\Xi_{n,d,\pi}$ is also stabilized by this subgroup. Indeed,

$$
\begin{aligned}
&(\{\bar\lambda\}_{n^2} \otimes \{\bar\mu\}_{n^2d\cdot n^2d} \otimes \{\bar\nu\}_{n^2})^{H\times H} \\
=& \bigoplus_{\lambda,\lambda',\mu,\mu',\nu,\nu'} k(\bar\lambda,\lambda,\lambda')k(\bar\mu,\mu,\mu')k(\bar\nu,\nu,\nu')(\{\lambda\}_n \otimes \{\lambda'\}_n \otimes \{\mu\}_{n^2d} \otimes \{\mu'\}_{n^2d} \otimes \{\nu\}_n \otimes \{\nu'\}_n)^{H\times H} \\
=& \bigoplus_{\lambda,\lambda',\mu,\mu',\nu,\nu'} k(\bar\lambda,\lambda,\lambda')k(\bar\mu,\mu,\mu')k(\bar\nu,\nu,\nu')(\{\lambda\}_n \otimes \{\mu\}_{n^2d} \otimes \{\nu\}_n)^H \otimes (\{\lambda'\}_n \otimes \{\mu'\}_{n^2d} \otimes \{\nu'\}_n)^H
\end{aligned}
$$

Now, for $\bar\lambda = \bar\nu = (1)$ and $\bar\mu = (d)$, the nonzero Kronecker coefficients are all 1 and correspond to $\lambda = \lambda' = \nu = \nu' = (1)$ and $\mu = \mu'$. Therefore,

$$
\begin{aligned}
&(\{\bar\lambda\}_{n^2} \otimes \{\bar\mu\}_{n^2d\cdot n^2d} \otimes \{\bar\nu\}_{n^2})^{H\times H} \\
=& \bigoplus_{\mu\vdash d}(\{(1)\}_n \otimes \{\mu\}_{n^2d} \otimes \{(1)\}_n)^H \otimes (\{(1)\}_n \otimes \{\mu\}_{n^2d} \otimes \{(1)\}_n)^H \\
\overset{c^\mu_{((1),\ldots,(1))}=\dim[\mu]}{=}& \sum_{\mu\vdash d} \dim[\mu]\cdot\dim[\mu] \\
=& \dim\mathbb{C}[\mathfrak{S}_d] = d!
\end{aligned}
$$

Thus, the $\mathsf{h}\Xi_{n,d,\pi}$ are a basis of the $H \times H$-invariant space. Recall the torus $T = (\mathbb{C} \setminus \{0\})^{d \times d}$ from Remark 8.5. Define $T_\pi = \{\alpha \in \mathbb{C}^{d \times d} \mid \prod_k^d \alpha_{\pi(k),k} = 1\}$. The action of $H \times H$ commutes with the action of every T_π. Since T_π rescales monomials, it is easy to see that $\mathsf{h}\Xi_{n,d,\pi}$ is invariant under the action of T_π and that

$$\dim(\mathbb{C}^{n^2} \otimes S^d\mathbb{C}^{n^4 d^2} \otimes \mathbb{C}^{n^2})^{H \times H \times T_\pi} = 1,$$

in particular $\mathsf{h}\Xi_{n,d,\pi}$ is characterized by its stabilizer. We determine the stabilizer of $\mathsf{h}\Xi_{n,d}$ in §11.V and §11.VI, and we see that $\mathsf{h}\Xi_{n,d}$ is not characterized by its stabilizer in (11.39).

11.IV Polystability

We follow [BI17, Prop. 2.8]. Define $\det_{n,1,n}(A, B, C) := \det(A)^n \det(B) \det(C)^n$.

11.8 Proposition. *Let* $G = \{(g_{\text{left}}, g_{\text{middle}}, g_{\text{right}}) \in \mathrm{GL}_n \times \mathrm{GL}_{n^2 d} \times \mathrm{GL}_n \mid \det_{n,1,n}(g_{\text{left}}, g_{\text{middle}}, g_{\text{right}}) = 1\}$. *The orbit* $G\Xi_{n,d}$ *is closed.* ◇

Proof. For the sake of contradiction, assume that $G\Xi_{n,d}$ is not closed. Let Y be a G-orbit of $\overline{G\Xi_{n,d}} \setminus G\Xi_{n,d}$ of minimal dimension. Then Y must be closed. Since G is reductive, the Hilbert-Mumford criterion says that there exists a one-parameter subgroup $\sigma : \mathbb{C}^\times \to G$ with $\lim_{t \to 0} \sigma(t)\Xi_{n,d} \in Y$. For applying [Lun75, Cor. 1], [Kem78, Cor 4.5] we need to pick a reductive $R \subseteq G \cap \mathrm{stab}(\Xi_{n,d})$, which in [BI17] is always a torus. For a variable $x_{i,j}^{(k)}$ let $h_{x_{i,j}^{(k)}}(\alpha) \in \mathrm{GL}_{n^2 d}$ be the identity matrix with a single α instead of a 1 at position $(x_{i,j}^{(k)}, x_{i,j}^{(k)})$. For a basis vector e_i, $1 \le i \le n$, let $h_{e_i}(\alpha) \in \mathrm{GL}_n$ be defined analogously, and for a basis vector e_j^*, $1 \le j \le n$, let $h_{e_j^*}(\alpha) \in \mathrm{GL}_n$ also be defined analogously. All these elements of $\mathrm{GL}_n \times \mathrm{GL}_{n^2 d} \times \mathrm{GL}_n$ commute with each other.

Let $D = (V, E)$ be the digraph with vertex set $v_i^{(k)}$, $1 \le i \le n$, $0 \le k \le d$, where $\{v_i^{(k)} \mid 1 \le i \le n\}$ is the k-th vertex layer, and we have the complete bipartite edge set from vertex layer $k-1$ to vertex layer k, labeled with $x_{i,j}^{(k)}$. Moreover, each vertex $v_i^{(0)}$ has a half-edge pointing at it, labeled with e_i, and each vertex $v_j^{(d)}$ has a half-edge pointing away from it, labeled with e_j^*. For every vertex $v \in V$, let $\mathrm{in}(v)$ denote the set of incoming (half-)edges, and $\mathrm{out}(v)$ denote the set of outgoing (half-)edges. For $v \in V$ and $\alpha \in \mathbb{C}$ let

$$h_{v,\alpha} := \prod_{e \in \mathrm{in}(v)} h_e(\alpha) \cdot \prod_{e \in \mathrm{out}(v)} h_e(\alpha^{-1}) \ \in G \cap \mathrm{stab}_G(\Xi_{n,d}).$$

Let R be generated by $\{h_{v,\alpha} \mid v \in V, \alpha \in \mathbb{C}\}$. Since R is the image of a torus, it is a torus itself, and hence reductive. We now determine the centralizer of R in G. Let $g \in \mathrm{GL}_n \times \mathrm{GL}_{n^2 d} \times \mathrm{GL}_n$ commute with $h_{v,\alpha}$ for all α, i.e., $g h_{v,\alpha} = h_{v,\alpha} g$. We think of $\mathrm{GL}_n \times \mathrm{GL}_{n^2 d} \times \mathrm{GL}_n$ as embedded into $\mathrm{GL}_{n+n^2 d+n}$ as block matrix. Let $\mathrm{neu}(v) := E \setminus (\mathrm{in}(v) \cup \mathrm{out}(v))$. The identity $\forall \alpha : g h_{v,\alpha} = h_{v,\alpha} g$ forces that $g_{e',e}$ can only be nonzero if $\{e', e\} \subseteq \mathrm{in}(v)$ or $\{e', e\} \subseteq \mathrm{out}(v)$ or $\{e', e\} \subseteq \mathrm{neu}(v)$. For every edge $e = (v, w)$ we have that $\mathrm{out}(v) \cap \mathrm{in}(w)$ contains only e, thus g has at most 1 nonzero entry in row e and at most 1 nonzero entry in column e. For a half-edge e at a source s we have $\mathrm{in}(s) = \{e\}$ and hence g has at most 1 nonzero entry in row e and at most 1 nonzero entry in column e. Analogously for half-edges at sinks. Hence, the centralizer of R contains only diagonal matrices. The Luna-Kempf theorem says that we can assume that

$$\sigma(t) = \big(\mathrm{diag}(t^{\lambda_1}, \dots, t^{\lambda_n}),\ \mathrm{diag}(t^{\mu_1}, \dots, t^{\mu_N}),\ \mathrm{diag}(t^{\nu_1}, \dots, t^{\nu_n})\big). \tag{11.9}$$

Interpret $\Xi_{n,d} \in \mathbb{C}^n \otimes S^d\mathbb{C}^{n^2 d} \otimes \mathbb{C}^{n*}$ and define $\mathsf{supp}(\Xi_{n,d})$ to be the set of basis tensors of the form $e_i \otimes m \otimes e_j^*$, where m is a monomial, that have nonzero coefficient in $\Xi_{n,d}$. Let $\mathrm{expvec}(e^{\alpha_{\text{left}}} \otimes e^{\alpha_{\text{middle}}} \otimes (e^{\alpha_{\text{right}}})^*) = (\alpha_{\text{left}}, \alpha_{\text{middle}}, \alpha_{\text{right}})$. Since $\lim_{t \to 0} \sigma(t)\Xi_{n,d}$ converges, we have

$$\forall \alpha \in \mathrm{expvec}(\mathsf{supp}(\Xi_{n,d})) :\ \langle \alpha, (\lambda, \mu, \nu) \rangle \ge 0. \tag{11.10}$$

Note that $\#\mathrm{expvec}(\mathsf{supp}(\Xi_{n,d})) = n^{d+1}$. For each $\alpha \in \mathrm{expvec}(\mathsf{supp}(\Xi_{n,d}))$ let $c_\alpha = \frac{n^2}{n^{d+1}} = n^{-(d-1)}$. We have

$$\sum_{\alpha \in \mathrm{expvec}(\mathsf{supp}(\Xi_{n,d}))} c_\alpha \alpha = (n, \dots, n, 1, \dots, 1, n, \dots, n) \tag{11.11}$$

Since $\det_{n,1,n}(\sigma(t)) = 1$, by (11.9) we have

$$\langle (n, \dots, n, 1, \dots, 1, n, \dots, n), (\lambda, \mu, \nu) \rangle = 0. \tag{11.12}$$

We now put the pieces together as follows.

$$0 \overset{(11.11),(11.12)}{=} \langle \sum_{\alpha \in \mathrm{expvec}(\mathsf{supp}(\Xi_{n,d}))} c_\alpha \alpha, (\lambda, \mu, \nu) \rangle = \sum_{\alpha \in \mathrm{expvec}(\mathsf{supp}(\Xi_{n,d}))} c_\alpha \langle \alpha, (\lambda, \mu, \nu) \rangle$$

Since $c_\alpha > 0$ for all $\alpha \in \mathrm{expvec}(\mathsf{supp}(\Xi_{n,d}))$, by (11.10) we have $\langle \alpha, (\lambda, \mu, \nu) \rangle = 0$ for all $\alpha \in \mathrm{expvec}(\mathsf{supp}(\Xi_{n,d}))$. This implies that $\forall t : \sigma(t)\Xi_{n,d} = \Xi_{n,d}$. Therefore, $\lim_{t \to 0} \sigma(t)\Xi_{n,d} = \Xi_{n,d}$, which is a contradiction. $\square$

11.13 Proposition. *Let* $G = \{(g_{\mathrm{left}}, g_{\mathrm{middle}}, g_{\mathrm{right}}) \in \mathrm{GL}_{n^2} \times \mathrm{GL}_{n^4 d} \times \mathrm{GL}_{n^2} \mid \det_{n^2,1,n^2}(g_{\mathrm{left}}, g_{\mathrm{middle}}, g_{\mathrm{right}}) = 1\}$. *For all* $\pi \in \mathfrak{S}_d$ *the orbit* $G\mathsf{h}\Xi_{n,d,\pi}$ *is closed.* $\diamond$

Proof. We proceed analogously to Proposition 11.8. Fix $\pi \in \mathfrak{S}_d$. For a variable $x^{(\pi(k))}_{i',j'} \boxtimes x^{(k)}_{i,j}$ let $h_{x^{(\pi(k))}_{i',j'} \boxtimes x^{(k)}_{i,j}}(\alpha) \in \mathrm{GL}_{n^4 d}$ be the identity matrix with a single α instead of a 1 at position $(x^{(\pi(k))}_{i',j'} \boxtimes x^{(k)}_{i,j}, x^{(\pi(k))}_{i',j'} \boxtimes x^{(k)}_{i,j})$, and analogously for $h_{e_{i'} \boxtimes e_i}(\alpha)$ and $h_{e^*_{j'} \boxtimes e^*_j}(\alpha)$. For fixed α we extend $h_{\cdot,\cdot}(\alpha)$ bilinearly. Let $D^2 = D_1 \dot\cup D_2$ denote the digraph that consists of two disjoint copies of D from the proof of Proposition 11.8. An edge that is not a half-edge is called a *full-edge*. The vertex set of the first copy is denoted by V_1, and its set of full-edges is denoted by E_1, its set of left half-edges by $H_1^{-\circ}$, and the set of right half-edges by $H_1^{\circ-}$. Analogously for V_2, E_2, $H_2^{-\circ}$, and $H_2^{\circ-}$. The subset of E_1 of edges in edge layer k is denoted by $E_1^{(k)}$, and analogously for $E_2^{(k)}$. For $e \in E_1^{(\pi(k))}$ let $\overline{e} = \sum_{e_2 \in E_2^{(k)}} e_2$, and for $e \in E_2^{(k)}$ let $\overline{e} = \sum_{e_1 \in E_1^{(\pi(k))}} e_1$. For $e \in H_1^{-\circ}$ let $\overline{e} = \sum_{e_2 \in H_2^{-\circ}} e_2$, and for $e \in H_2^{-\circ}$ let $\overline{e} = \sum_{e_1 \in H_1^{-\circ}} e_1$. For $e \in H_1^{\circ-}$ let $\overline{e} = \sum_{e_2 \in H_2^{\circ-}} e_2$, and for $e \in H_2^{\circ-}$ let $\overline{e} = \sum_{e_1 \in H_1^{\circ-}} e_1$. For $v \in V_1 \cup V_2$ and $\alpha \in \mathbb{C}$ let

$$h_{v,\alpha} := \prod_{e \in \mathrm{in}(v)} h_{e \boxtimes \overline{e}}(\alpha) \cdot \prod_{e \in \mathrm{out}(v)} h_{e \boxtimes \overline{e}}(\alpha^{-1}) \ \in G \cap \mathrm{stab}_G(\mathsf{h}\Xi_{n,d,\pi}).$$

Let R be generated by $\{h_{v,\alpha} \mid v \in V_1 \cup V_2, \alpha \in \mathbb{C}\}$. We now determine the centralizer of R in G. The identity $\forall \alpha : g h_{v,\alpha} = h_{v,\alpha} g$ forces that $g_{e'_1 \boxtimes e'_2, e_1 \boxtimes e_2}$ can only be nonzero if $\{e'_1, e_1\} \subseteq \mathrm{in}(v)$ or $\{e'_1, e_1\} \subseteq \mathrm{out}(v)$ or $\{e'_1, e_1\} \subseteq \mathrm{neu}(v)$. Analogously to the proof of Proposition 11.8, this means that g can only be nonzero at entries $(e_1 \boxtimes e'_2, e_1 \boxtimes e_2)$, and analogously for half-edges. Together with the same argument for (e'_2, e_2) this shows that g can only be nonzero at diagonal entries. We proceed as in the proof of Proposition 11.8, but note that $\#\mathrm{expvec}(\mathsf{supp}(\mathsf{h}\Xi_{n,d,\pi})) = n^{2d+2}$, so we set $c_\alpha = \frac{n^4}{n^{2d+2}} = n^{-(2d-2)}$ and obtain

$$\sum_{\alpha \in \mathrm{expvec}(\mathsf{supp}(\mathsf{h}\Xi_{n,d,\pi}))} c_\alpha \alpha = (n^2, \dots, n^2, 1, \dots, 1, n^2, \dots, n^2).$$

We have $\langle (n^2, \dots, n^2, 1, \dots, 1, n^2, \dots, n^2), (\lambda, \mu, \nu) \rangle = 0$, and hence

$$0 = \sum_{\alpha \in \mathrm{expvec}(\mathsf{supp}(\mathsf{h}\Xi_{n,d,\pi}))} c_\alpha \langle \alpha, (\lambda, \mu, \nu) \rangle.$$

Since $c_\alpha > 0$ for all $\alpha \in \mathrm{expvec}(\mathsf{supp}(\mathsf{h}\Xi_{n,d,\pi}))$, we have $\langle \alpha, (\lambda, \mu, \nu) \rangle = 0$ for all $\alpha \in \mathrm{expvec}(\mathsf{supp}(\mathsf{h}\Xi_{n,d,\pi}))$. This implies that $\forall t : \sigma(t)\mathsf{h}\Xi_{n,d,\pi} = \mathsf{h}\Xi_{n,d,\pi}$. Therefore, $\lim_{t \to 0} \sigma(t)\mathsf{h}\Xi_{n,d,\pi} = \mathsf{h}\Xi_{n,d,\pi}$, which is a contradiction. $\square$

11.14 Proposition. *Let* $G = \{(g_{\text{left}}, g_{\text{middle}}, g_{\text{right}}) \in \mathrm{GL}_{n^2} \times \mathrm{GL}_{n^4 d^2} \times \mathrm{GL}_{n^2} \mid \det_{n^2 d,1,n^2 d}(g_{\text{left}}, g_{\text{middle}}, g_{\text{right}}) = 1\}$. *The orbit* $G\mathsf{h}\Xi_{n,d}$ *is closed.* $\diamond$

Proof. We proceed analogously to Proposition 11.13. For a variable $x_{i,j}^{(k)} \boxtimes x_{i',j'}^{(k')}$ let $h_{x_{i,j}^{(k)} \boxtimes x_{i',j'}^{(k')}}(\alpha) \in \mathrm{GL}_{n^4 d^2}$ be the identity matrix with a single α instead of a 1 at position $(x_{i,j}^{(k)} \boxtimes x_{i',j'}^{(k')}, x_{i,j}^{(k)} \boxtimes x_{i',j'}^{(k')})$, and analogously for $h_{e_{i'} \boxtimes e_i}(\alpha)$ and $h_{e_{j'}^* \boxtimes e_j^*}(\alpha)$. For fixed α we extend $h_{\cdot,\cdot}(\alpha)$ bilinearly. Let $D^2 = D_1 \dot\cup D_2$ denote the digraph that consists of two disjoint copies of D from the proof of Proposition 11.8. An edge that is not a half-edge is called a *full-edge*. The vertex set of the first copy is denoted by V_1, and its set of full-edges is denoted by E_1, its set of left half-edges by $H_1^{-\circ}$, and the set of right half-edges by $H_1^{\circ-}$. Analogously for V_2, E_2, $H_2^{-\circ}$, and $H_2^{\circ-}$. For $e \in E_1$ let $\overline{e} = \sum_{e_2 \in E_2} e_2$, and for $e \in E_2$ let $\overline{e} = \sum_{e_1 \in E_1} e_1$. For $e \in H_1^{-\circ}$ let $\overline{e} = \sum_{e_2 \in H_2^{-\circ}} e_2$, and for $e \in H_2^{-\circ}$ let $\overline{e} = \sum_{e_1 \in H_1^{-\circ}} e_1$. For $e \in H_1^{\circ-}$ let $\overline{e} = \sum_{e_2 \in H_2^{\circ-}} e_2$, and for $e \in H_2^{\circ-}$ let $\overline{e} = \sum_{e_1 \in H_1^{\circ-}} e_1$. For $v \in V_1 \cup V_2$ and $\alpha \in \mathbb{C}$ let

$$h_{v,\alpha} := \prod_{e \in \mathrm{in}(v)} h_{e \boxtimes \overline{e}}(\alpha) \cdot \prod_{e \in \mathrm{out}(v)} h_{e \boxtimes \overline{e}}(\alpha^{-1}) \ \in G \cap \mathrm{stab}_G(\mathsf{h}\Xi_{n,d}).$$

Let R be generated by $\{h_{v,\alpha} \mid v \in V_1 \cup V_2, \alpha \in \mathbb{C}\}$. We now determine the centralizer of R in G. The identity $\forall \alpha : g h_{v,\alpha} = h_{v,\alpha} g$ forces that $g_{e_1' \boxtimes e_2', e_1 \boxtimes e_2}$ can only be nonzero if $\{e_1', e_1\} \subseteq \mathrm{in}(v)$ or $\{e_1', e_1\} \subseteq \mathrm{out}(v)$ or $\{e_1', e_1\} \subseteq \mathrm{neu}(v)$. Analogously to the proof of Proposition 11.8, this means that g can only be nonzero at entries $(e_1 \boxtimes e_2', e_1 \boxtimes e_2)$, and analogously for half-edges. Together with the same argument for (e_2', e_2) this shows that g can only be nonzero at diagonal entries. We proceed as in the proof of Proposition 11.8, but note that $\#\mathrm{expvec}(\mathsf{supp}(\mathsf{h}\Xi_{n,d})) = d!\, n^{2d+2}$, so we set $c_\alpha = \frac{n^4 d}{d!\, n^{2d+2}} = \frac{1}{(d-1)!\, n^{2d-2}}$ and obtain

$$\sum_{\alpha \in \mathrm{expvec}(\mathsf{supp}(\mathsf{h}\Xi_{n,d}))} c_\alpha \alpha = (n^2 d, \ldots, n^2 d, 1, \ldots, 1, n^2 d, \ldots, n^2 d).$$

We have $\langle (n^2 d, \ldots, n^2 d, 1, \ldots, 1, n^2 d, \ldots, n^2 d), (\lambda, \mu, \nu) \rangle = 0$, and hence

$$0 = \sum_{\alpha \in \mathrm{expvec}(\mathsf{supp}(\mathsf{h}\Xi_{n,d}))} c_\alpha \langle \alpha, (\lambda, \mu, \nu) \rangle.$$

Since $c_\alpha > 0$ for all $\alpha \in \mathrm{expvec}(\mathsf{supp}(\mathsf{h}\Xi_{n,d}))$, we have $\langle \alpha, (\lambda, \mu, \nu) \rangle = 0$ for all $\alpha \in \mathrm{expvec}(\mathsf{supp}(\mathsf{h}\Xi_{n,d}))$. This implies that $\forall t : \sigma(t)\mathsf{h}\Xi_{n,d} = \mathsf{h}\Xi_{n,d}$. Therefore, $\lim_{t \to 0} \sigma(t)\mathsf{h}\Xi_{n,d} = \mathsf{h}\Xi_{n,d}$, which is a contradiction. $\square$

11.V Continuous part of the stabilizer of $\mathsf{h}\Xi_{n,d}$

Let

$$\mathcal{V} = \bigoplus_{r,s=1}^{d} \mathcal{V}_{r,s}, \qquad \mathcal{V}_{r,s} = E_r \otimes F_s, \qquad E_r \simeq F_s \simeq \mathrm{Mat}_n,$$

and write $v_{a,a',b,b'}^{(r,s)} = x_{a,a'}^{(r)} \boxtimes x_{b,b'}^{(s)}$. We give coordinates to $\mathcal{V}_{r,s}$ with $v_{a,a',b,b'}^{(r,s)}, 1 \le a, a', b, b' \le n$. The Lie algebra $\mathfrak{gl}(\mathcal{V})$ acts on polynomials on $\mathcal{V}$ by linear changes of variables. We index the rows $\mathsf{h}\Xi_{n,d}$ by pairs $(i,k) \in [n] \times [n]$ and columns by pairs $(j,l) \in [n] \times [n]$. Put

$$a_0 = i, \qquad a_d = j, \qquad b_0 = k, \qquad b_d = l.$$

Then

$$[\mathsf{h}\Xi_{n,d}]_{(i,k),(j,\ell)} = \sum_{\pi \in \mathfrak{S}_d} \sum_{a_1,\ldots,a_{d-1}} \sum_{b_1,\ldots,b_{d-1}} \prod_{r=1}^{d} v_{a_{r-1},a_r,\, b_{\pi(r)-1},b_{\pi(r)}}^{(r,\pi(r))}. \tag{11.15}$$

Furthermore, since $\mathsf{h}\Xi_{n,d}$ is an $n^2 \times n^2$ matrix, we have both a left and right action of $\mathfrak{gl}_{n^2}$ on $h\Xi_{n,d}$. Define $\mathfrak{g} = \mathfrak{gl}_{n^2} \oplus \mathfrak{gl}(\mathcal{V}) \oplus \mathfrak{gl}_{n^2}$. In this subsection we calculate the Lie algebra stabilizer

$$\mathrm{Stab}(\mathsf{h}\Xi_{n,d}) := \{(L, X, R) \in \mathfrak{g} \mid L\mathsf{h}\Xi_{n,d} + X \cdot \mathsf{h}\Xi_{n,d} + \mathsf{h}\Xi_{n,d}R = 0\}$$

for $d \geq 3$. For $\pi \in \mathfrak{S}_d$, let $h_\pi := \mathsf{h}\Xi_{n,d,\pi}$. Its block multidegree is

$$\deg(h_\pi) = \sum_{r=1}^{d} e_{r,\pi(r)} \in \mathbb{Z}^{d\times d}.$$

For paths $a = (a_0, \ldots, a_d)$ and $b = (b_0, \ldots, b_d)$, write

$$M_\pi(a, b) = \prod_{r=1}^{d} v^{(r,\pi(r))}_{a_{r-1},a_r,\, b_{\pi(r)-1},b_{\pi(r)}}.$$

11.16 Lemma (Diagonal block lemma). *If $(L, X, R) \in \mathrm{Stab}(\mathsf{h}\Xi_{n,d})$, then*

$$X = \sum_{r,s\in[d]} X^{(r,s),(r,s)}$$

where $X^{(r,s),(r,s)} \in \mathrm{Hom}(W_{r,s}, W_{r,s})$. ◇

Proof. Decompose X into block components

$$X = \sum_{(r,s),(r',s')} X^{(r,s),(r',s')}, \qquad X^{(r,s),(r',s')} \in \mathrm{Hom}(\mathcal{V}_{r,s}, \mathcal{V}_{r',s'}).$$

The action of $X^{(r,s),(r',s')}$ replaces one variable from $\mathcal{V}_{r,s}$ by a linear form in the variables of $\mathcal{V}_{r',s'}$. It therefore changes block multidegree by $e_{r',s'} - e_{r,s}$. Let $\lambda^{x,y,z,w}_{k,l,m,n}$ be coefficients such that

$$X^{(r,s),(r',s')} \cdot v^{(r,s)}_{x,y,z,w} = \sum_{j,k,l,m} \lambda^{x,y,z,w}_{j,k,l,m} v^{(r',s')}_{j,k,l,m}.$$

First suppose that $r \neq r'$ and $s \neq s'$. Choose a permutation π such that $\pi(r) = s$. In $X \cdot \mathsf{h}\Xi_{n,d}$, the component of multidegree $\deg(\pi) - e_{r,s} + e_{r',s'}$ is precisely given by $X^{(r,s),(r',s')} \cdot h_\pi$. This is because from the multidegree, we can uniquely recover r, s, r', s' and π. Since each multidegree component must be 0, $X^{(r,s),(r',s')} \cdot h_\pi = 0$.

Let $a_{r-1} = x, a_r = y, b_{s-1} = z, b_s = w$. Then, we have the following action.

$$\begin{aligned}
X^{(r,s),(r',s')} \cdot M_\pi(a,b) &= \left(X^{(r,s),(r',s')} \cdot v^{(r,s)}_{x,y,z,w}\right) \prod_{i\in[d]\setminus\{r\}} v^{(i,\pi(i))}_{a_{i-1},a_i,b_{\pi(i)-1},b_{\pi(i)}} \\
&= \left(\sum_{j,k,l,m} \lambda^{x,y,z,w}_{j,k,l,m} v^{(r',s')}_{j,k,l,m}\right) \prod_{i\in[d]\setminus\{r\}} v^{(i,\pi(i))}_{a_{i-1},a_i,b_{\pi(i)-1},b_{\pi(i)}} \\
&= \sum_{j,k,l,m} \lambda^{x,y,z,w}_{j,k,l,m} \left(v^{(r',s')}_{j,k,l,m} \prod_{i\in[d]\setminus\{r\}} v^{(i,\pi(i))}_{a_{i-1},a_i,b_{\pi(i)-1},b_{\pi(i)}}\right)
\end{aligned}$$

Notice the resulting monomials $v^{(r',s')}_{j,k,l,m} \prod_{i\in[d]\setminus\{r\}} v^{(i,\pi(i))}_{a_{i-1},a_i,b_{\pi(i)-1},b_{\pi(i)}}$. From each of these monomials, we can determine the paths $a_0, ..., a_d$ and $b_0, ..., b_d$ by reading off the subscripts of the resulting monomials. Along with the fact that we know π, we can recover $M_\pi(a, b)$. But this means, there cannot be any cancellations on the right hand side and hence, $\lambda^{x,y,z,w}_{j,k,l,m} = 0$. Hence

$$X^{(r,s),(r',s')} = 0 \quad \text{if } r \neq r' \text{ and } s \neq s'.$$

It remains to treat the case where two blocks are in one row or one column. Consider (r,s) and (r,s') with $s \neq s'$. Choose $q \neq r$ and a permutation π such that $\pi(r) = s$ and $\pi(q) = s'$. The multidegree $\deg(\pi) - e_{r,s} + e_{r,s'}$ has exactly two possible sources: the component $X^{(r,s),(r,s')} \cdot h_\pi$ and the component $X^{(q,s),(q,s')} \cdot h_\tau$ where $\tau = (r\ q) \cdot \pi$. Hence,

$$X^{(r,s),(r,s')} \cdot h_\pi + X^{(q,s),(q,s')} \cdot h_\tau = 0.$$

Just as in the previous case, the resulting monomials $v^{(r,s')}_{j,k,l,m} \prod_{i\in[d]\setminus\{r\}} v^{(i,\pi(i))}_{a_{i-1},a_i,b_{\pi(i)-1},b_{\pi(i)}}$ appearing in $X \cdot h_\pi$ cannot be cancelled by other monomials within h_π. But, they could be cancelled by monomials appearing from $X \cdot h_\tau$. We show this is not the case.

$$v^{(r,s')}_{j,k,l,m} \prod_{i\in[d]\setminus\{r\}} v^{(i,\pi(i))}_{a_{i-1},a_i,b_{\pi(i)-1},b_{\pi(i)}} = v^{(r,s')}_{j,k,l,m} v^{(q,s')}_{a_{q-1},a_q,b_{s'-1},b_{s'}} \prod_{i\in[d]\setminus\{r,q\}} v^{(i,\pi(i))}_{a_{i-1},a_i,b_{\pi(i)-1},b_{\pi(i)}}$$

For this to be a monomial of $X^{(q,s),(q,s')} \cdot h_\tau$, we must have

$$v^{(r,s')}_{j,k,l,m} v^{(q,s)}_{\alpha,\beta,\gamma,\delta} \prod_{i\in[d]\setminus\{r,q\}} v^{(i,\pi(i))}_{a_{i-1},a_i,b_{\pi(i)-1},b_{\pi(i)}} \tag{11.17}$$

as a monomial of h_τ, for some variable $v^{(q,s)}_{\alpha,\beta,\gamma,\delta} \in \mathcal{V}_{q,s}$. Monomials of h_τ come from two paths connected by τ. Reading off (11.17), we see that the s'-th edge of the second path has vertices l and m. If $s' \notin \{1,d\}$ (it is not a boundary edge), then at least one of $s'-1$ or $s'+1$ is not equal to s. Assume $s'-1 \neq s$. Then the $s'-1$-th edge of the second path has vertices $b_{s'-2}$ and $b_{s'-1}$. So, the common vertex between the edges $s'-1$ and s' is $b_{s'-1} = l$. But, $s'-1 \neq s, s-1$ and hence it cannot be not a fixed variable. If instead $s'-1 = s$, then $s'+1 \neq s$. We have that the $s'+1$-th edge of the second path has vertices $b_{s'}$ and $b_{s'+1}$. So, $b_{s+1} = b_{s'} = m$ and we can conclude as before. The only remaining casework here is that where s' is a boundary edge. Without loss of generality, assume $s' = 1$. Then, l has to be the start of the second path. But since $s \neq s'$, the start of the second path is not fixed. In all cases, we can conclude by choosing a free variable different from the fixed values l, m to get a contradiction.

This finishes the proof that $v^{(r',s')}_{j,k,l,m} \prod_{i\in[d]\setminus\{r\}} v^{(i,\pi(i))}_{a_{i-1},a_i,b_{\pi(i)-1},b_{\pi(i)}}$ cannot be a monomial in $X \cdot h_\tau$. So, the resulting monomial cannot be cancelled. We can conclude that $X^{(r,s),(r,s')} = 0$ if $s \neq s'$. The case for $X^{(r,s)(r',s)}$ where $r \neq r'$ is analogous. □

Let $X^{(r,s),(r,s)}$ act on $v^{(r,s)}_{x,y,z,w}$ with the following coefficients.

$$X^{(r,s),(r,s)} \cdot v^{(r,s)}_{x,y,z,w} = \sum_{j,k,l,m} \mu^{x,y,z,w}_{j,k,l,m} v^{(r,s)}_{j,k,l,m}.$$

11.18 Lemma (Mixing lemma)**.** *If at least two of*

$$j \neq x, \qquad k \neq y, \qquad l \neq z, \qquad m \neq w$$

hold, then

$$\mu^{x,y,z,w}_{j,k,l,m} = 0. \qquad \diamond$$

Proof. Fix (r,s) and a coefficient $\mu^{x,y,z,w}_{j,k,l,m}$ for which at least two of the four displayed inequalities hold. Choose $\pi \in \mathfrak{S}_d$ with $\pi(r) = s$, and choose paths a, b satisfying

$$a_{r-1} = x, \qquad a_r = y, \qquad b_{s-1} = z, \qquad b_s = w.$$

Let $M = M_\pi(a,b)$ and let

$$N = v^{(r,s)}_{j,k,l,m} \prod_{i\neq r} v^{(i,\pi(i))}_{a_{i-1},a_i,\,b_{\pi(i)-1},b_{\pi(i)}}. \tag{11.19}$$

If among the changed indices, at least one is a boundary index (a_0, b_0, a_d, b_d), it can only be cancelled by a contribution from $L\,\mathsf{h}\Xi_{n,d}$ or $\mathsf{h}\Xi_{n,d}R$. If an internal vertex was also changed, then neither $L\mathsf{h}\Xi_{n,d}$ nor $\mathsf{h}\Xi_{n,d}R$ can cancel such a term, since $L\mathsf{h}\Xi_{n,d}$ and $\mathsf{h}\Xi_{n,d}R$ do not have broken paths.

If all changed indices are boundary indices, without loss of generality, they are a_0 and b_0. Choose $\pi(1) \neq 1$. Monomials in $X \cdot \mathsf{h}\Xi_{n,d}$ cannot cancel N in this case. We will later prove that in L and R, coefficients changing two boundary indices are zero. Hence, there is no cancellation of the monomial in this case as well.

Now, we are left with the case where the changed indices are internal. We choose π so that the two changed path incidences isolate the block (r, s). More explicitly, associate to a changed first-path endpoint the adjacent row

$$\rho = \begin{cases} r-1, & j \neq x \text{ and } r > 1, \\ r+1, & k \neq y \text{ and } r < d, \end{cases}$$

and associate to a changed second-path endpoint the adjacent column

$$\kappa = \begin{cases} s-1, & l \neq z \text{ and } s > 1, \\ s+1, & m \neq w \text{ and } s < d. \end{cases}$$

When one changed endpoint lies on each path and both are internal, we also require $\pi(\rho) \neq \kappa$. Such a permutation exists since $d \geq 3$. If both changed endpoints lie on the same path, no additional condition is needed.

We claim that the only contribution to N is obtained by $X^{(r,s),(r,s)}$ on M. Indeed, the unchanged factors in (11.19) determine all edges of both paths except the edge in row r and the edge in column s. Each changed index breaks the corresponding incidence with a neighbouring edge, or with a boundary label. If the two changes occur on the same path, the two broken incidences identify the missing edge immediately. If one change occurs on each path, the condition $\pi(\rho) \neq \kappa$ ensures that the two broken incidences cannot be produced by changing a single factor in any block other than (r, s). Thus the acted-on block is (r, s). Reading the neighbouring edges, together with the matrix entry at a boundary, then recovers the original four indices (x, y, z, w) and the replacement indices (j, k, l, m). Consequently, the coefficient of N is exactly $\mu^{x,y,z,w}_{j,k,l,m}$. Hence, this coefficient is zero. ☐

11.20 Lemma (Boundary tensor separation lemma)**.** *There are matrices*

$$L_1, L_2, R_1, R_2 \in \mathfrak{gl}_n$$

such that

$$L = L_1 \otimes I_n + I_n \otimes L_2, \qquad R = R_1 \otimes I_n + I_n \otimes R_2. \tag{11.21}$$

◇

Proof. Suppose an off-diagonal coefficient of L changes both components of the row label,

$$(i, k) \longmapsto (i', k'), \qquad i' \neq i, \quad k' \neq k.$$

Choose π with $\pi(1) \neq 1$. In a monomial of h_π, the initial first-path vertex $a_0 = i$ and the initial second-path vertex $b_0 = k$ occur in two different factors. The corresponding term of Lh_π changes both boundary labels simultaneously. Since X acts on only one factor, no term of $X \cdot h_\pi$ can produce the same matrix-valued monomial. A right output action changes the terminal rather than the initial label. Hence this coefficient of L is zero. The same argument applies to R at the terminal vertices.

Now consider coefficients of L which change only i. Choose $\pi(1) = s \neq 1$. Their only possible cancelling terms are changes of the first index of the row-1 factor in $\mathcal{V}_{1,s}$. Keeping that factor fixed and varying the initial second-path label k shows that the coefficient of L is independent of k. Thus these coefficients are of the form $L_1 \otimes I_n$. Interchanging the two paths gives $I_n \otimes L_2$. So far, we have only established the shape of the off diagonal entries of L. For diagonal entries, coefficient comparison between $L_{(i,k),(i,k)}, L_{(i',k),(i',k)}, L_{(i,k'),(i,k')}, L_{(i',k'),(i',k')}$ show that their diagonal parts split additively between the two components. The same argument applies to R. This proves (11.21). $\square$

We will now prove a similar tensor separation lemma for X. For this, we have two different arguments for the diagonal and the off-diagonal parts of X.

11.22 Lemma (Off-diagonal tensor separation lemma for X). *There are off-diagonal operators*

$$\mathcal{A}_r^{\mathrm{off}} \in \mathrm{End}(E_r), \qquad \mathcal{B}_s^{\mathrm{off}} \in \mathrm{End}(F_s)$$

such that, on every block $W_{r,s}$, the off-diagonal part $X^{(r,s),(r,s),\mathrm{off}}$ of $X^{(r,s),(r,s)}$ satisfies

$$X^{(r,s),(r,s),\mathrm{off}} = \mathcal{A}_r^{\mathrm{off}} \otimes I_{F_s} + I_{E_r} \otimes \mathcal{B}_s^{\mathrm{off}}. \qquad \diamond$$

Proof. We treat changes in the first tensor factor; the second factor is symmetric. Fix an internal first-path vertex $p \in [d-1]$ and $u \neq y$. Write $\mathcal{R}_{p,s}(x,y,z,w;u)$ for the coefficient of $v^{(p,s)}_{xy,zw} \longmapsto v^{(p,s)}_{xu,zw}$, and write $\mathcal{L}_{p+1,t}(u,v,z',w';y)$ for the coefficient of $v^{(p+1,t)}_{uv,z'w'} \longmapsto v^{(p+1,t)}_{yv,z'w'}$. Choose distinct columns s,t, a permutation π with $\pi(p) = s$ and $\pi(p+1) = t$, and compatible labelled paths. The first change breaks the first path at the vertex between rows p and $p+1$. The only other source of the same target monomial is the second change, starting from the path whose common vertex is u. Hence coefficient comparison in $X \cdot h_\pi = 0$ gives

$$\mathcal{R}_{p,s}(x,y,z,w;u) + \mathcal{L}_{p+1,t}(u,v,z',w';y) = 0. \tag{11.23}$$

Varying x while the second term is fixed shows that $\mathcal{R}_{p,s}$ is independent of x, and varying v shows that $\mathcal{L}_{p+1,t}$ is independent of v. If the second-path edges s and t are non-adjacent, their endpoint pairs may be varied independently, so (11.23) shows that both coefficients are independent of their second-path edge labels. For $d > 3$, every edge has a non-adjacent edge. For $d = 3$, columns 1 and 3 are non-adjacent, while the two endpoints of the middle edge are removed separately by comparing column 2 first with column 1 and then with column 3. Thus edge-label independence holds for every column. We have,

$$\mathcal{R}_{p,s}(y;u) = -\mathcal{L}_{p+1,t}(u;y) \qquad (s \neq t).$$

Given two columns s_1, s_2, choose a third column t distinct from both. Then

$$\mathcal{R}_{p,s_1}(y;u) = -\mathcal{L}_{p+1,t}(u;y) = \mathcal{R}_{p,s_2}(y;u),$$

so the coefficients are also independent of s. The same argument at the vertex between rows $p-1$ and p treats changes of the first index. Consequently every coefficient changing x or y depends only on the source and target basis vectors in E_r, and not on s, z, w. These coefficients define an operator $\mathcal{A}_r^{\mathrm{off}}$ and $\mathcal{B}_s^{\mathrm{off}}$. The mixing lemma says that these are all off-diagonal coefficients of X, which proves Lemma 11.22. $\square$

Write the diagonal part of X as

$$X^{\mathrm{diag}} \cdot v^{(r,s)}_{xy,zw} = c^{r,s}(x,y,z,w) v^{(r,s)}_{xy,zw}.$$

Let $\lambda(i,k) = L_{(i,k),(i,k)}, \rho(j,\ell) = R_{(j,\ell),(j,\ell)}$ be the diagonal coefficient functions of the left and right action. The boundary-separation lemma gives functions such that

$$\lambda(i,k) = \lambda_a(i) + \lambda_b(k), \qquad \rho(j,\ell) = \rho_a(j) + \rho_b(\ell).$$

Fix $\pi \in \mathfrak{S}_d$ and labelled paths $a = (a_0, \ldots, a_d), b = (b_0, \ldots, b_d)$, and put $A_r = (a_{r-1}, a_r), B_s = (b_{s-1}, b_s)$. Comparing the coefficient of the corresponding path monomial gives

$$\sum_{r=1}^{d} c^{r,\pi(r)}(A_r, B_{\pi(r)}) + \lambda_a(a_0) + \lambda_b(b_0) + \rho_a(a_d) + \rho_b(b_d) = 0. \tag{11.24}$$

11.25 Lemma (Diagonal tensor-separation lemma for X)**.** *There are functions*

$$f_r\colon [n]^2 \longrightarrow \mathbb{C}, \qquad g_s\colon [n]^2 \longrightarrow \mathbb{C}$$

such that for every r, s and $A, B \in [n]^2$,

$$c^{r,s}(A, B) = f_r(A) + g_s(B).$$

◇

Proof. Choose distinct rows r, t, distinct columns s, u, and a permutation π satisfying

$$\pi(r) = s, \qquad \pi(t) = u.$$

Put $\tau = \pi \circ (r\ t)$. Apply (11.24) to the same two labelled paths with the permutations π and τ. The four boundary vertices are unchanged, so all four output terms cancel on subtraction. Every summand other than those in rows r, t also cancels, and hence

$$c^{r,s}(A_r, B_s) + c^{t,u}(A_t, B_u) = c^{r,u}(A_r, B_u) + c^{t,s}(A_t, B_s). \tag{11.26}$$

Fix r, s. Let A, A' be first path edges on level r, and let B, B' be second path edges on level s. Considering all the equations (11.24) coming from monomials that have terms $(A, B), (A', B), (A, B'), (A', B')$, and taking the sum, we can get

$$c^{r,s}(A, B) + c^{r,s}(A', B') = c^{r,s}(A', B) + c^{r,s}(A, B'). \tag{11.27}$$

Equations (11.26) and (11.27) are enough to deduce the additivity of the diagonal part. □

Combining Lemma 11.22 and Lemma 11.25, there are operators

$$\mathcal{A}_r \in \mathrm{End}(E_r), \qquad \mathcal{B}_s \in \mathrm{End}(F_s)$$

such that

$$X|_{\mathcal{V}_{r,s}} = \mathcal{A}_r \otimes I_{F_s} + I_{E_r} \otimes \mathcal{B}_s. \tag{11.28}$$

Reduction to Iterated Matrix Multiplication

Let

$$\mathcal{A} = \bigoplus_{r=1}^{d} \mathcal{A}_r, \qquad \mathcal{B} = \bigoplus_{s=1}^{d} \mathcal{B}_s.$$

By (11.21), write

$$L = L_1 \otimes I_n + I_n \otimes L_2, \qquad R = R_1 \otimes I_n + I_n \otimes R_2.$$

Let $\mathcal{X} = \{x^{(r)}_{a,b} \mid a, b \in [n], r \in [d]\}$ and $\mathcal{Y} = \{y^{(r)}_{a,b} \mid a, b \in [n], r \in [d]\}$ Define the map $\Psi : \mathbb{C}[\mathcal{V}] \to \mathbb{C}[\mathcal{X}, \mathcal{Y}]$

$$\Psi\big(v^{(r,s)}_{jk,lm}\big) = x^{(r)}_{jk} y^{(s)}_{lm}.$$

This extends to a map of polynomials on $\mathcal{V}$. We have

$$\Psi([\mathsf{h}\Xi_{n,d}]_{(a_1,a_2)(b_1,b_2)}) = d!\left([\Xi_{n,d}(x)]_{a_1,b_1} \cdot [\Xi_{n,d}(y)]_{a_2,b_2}\right),$$

and on matrices

$$\Psi(\mathsf{h}\Xi_{n,d}) = d!\left([\Xi_{n,d}(x)] \otimes [\Xi_{n,d}(y)]\right),$$

where the tensor symbol depicts the traditional Kronecker product on matrices. If $(L, X, R) \in \mathrm{Stab}(\mathsf{h}\Xi_{n,d})$, then by lemmas 11.20, 11.22 and 11.25, each operator descends to an action on Image(Ψ). We get the following equation.

$$\begin{aligned} 0 = &\big(L_1\Xi_{n,d}(x) + \mathcal{A}\cdot\Xi_{n,d}(x) + \Xi_{n,d}(x)R_1\big)\otimes\Xi_{n,d}(y) \\ &+ \Xi_{n,d}(x)\otimes\big(L_2\Xi_{n,d}(y) + \mathcal{B}\cdot\Xi_{n,d}(y) + \Xi_{n,d}(y)R_2\big) \end{aligned}$$

After adding suitable identity operators to $\mathcal{A}$ and $\mathcal{B}$ that leave the decomposition (11.28) unchanged, we can deduce that

$$\begin{aligned} L_1\Xi_{n,d}(x) + \mathcal{A}\cdot\Xi_{n,d}(x) + \Xi_{n,d}(x)R_1 = 0, \\ L_2\Xi_{n,d}(y) + \mathcal{B}\cdot\Xi_{n,d}(y) + \Xi_{n,d}(y)R_2 = 0. \end{aligned}$$

So, the triples $(L_1, \mathcal{A}, R_1)$ and $(L_2, \mathcal{B}, R_2)$ are stabilizers for $\Xi_{n,d}(x)$ and $\Xi_{n,d}(y)$ respectively. Linearizing (11.5), i.e., going to infinitesimals, we see that there are matrices $P_0, P_1, ..., P_d, Q_0, Q_1, ..., Q_d \in \mathfrak{gl}(n)$ such that $(L, X, R) \in \mathrm{Stab}(\mathsf{h}\Xi_{n,d})$ is of the following form.

$$\begin{aligned} X\cdot v^{(r,s)}_{xy,zw} = &\sum_u (P_r)_{uy}v^{(r,s)}_{xu,zw} - \sum_u (P_{r-1})_{xu}v^{(r,s)}_{uy,zw} \\ &+\sum_u (Q_s)_{uw}v^{(r,s)}_{xy,zu} - \sum_u (Q_{s-1})_{zu}v^{(r,s)}_{xy,uw}, \\ L = &\,P_0\otimes I_n + I_n\otimes Q_0, \\ R = &-P_d\otimes I_n - I_n\otimes Q_d. \end{aligned} \tag{11.29}$$

11.VI The stabilizer of $\mathsf{h}\Xi_{n,d}$

Recall

$$\mathcal{V} = \bigoplus_{r,l=1}^{d}\mathcal{V}_{r,l}, \qquad \mathcal{V}_{r,l}\simeq\mathbb{C}^{n^4}.$$

Fix $n\geq 2$, $d\geq 3$ for all statements in this subsection. Recall $\mathcal{G} := \mathrm{GL}_n\times G_{n,d}\times\mathrm{GL}_n$. For $M\in\mathrm{GL}_{n^4d^2}$ we write $[M]$ for its image in the quotient $G_{n,d}$. In this section we write $G := \mathcal{G}$ for simplicity.

In this subsection we use the result from §11.V and determine the stabilizer H of $\mathsf{h}\Xi_{n,d}$, adjusting the method from [BGL14, Ges16] to our situation.

We have the action of the group $\mathbb{Z}/2\mathbb{Z}$, generated by τ on $\mathbb{C}^{n^2}\otimes S^{n^4d^2}\otimes(\mathbb{C}^{n^2})^*$, mapping $x^{(k)}_{i,j}\boxtimes x^{(k')}_{i',j'}$ to $x^{(k')}_{i',j'}\boxtimes x^{(k)}_{i,j}$ and at the same time mapping each $e_i\boxtimes e_j$ to $e_j\boxtimes e_i$ and each $e^*_j\boxtimes e^*_i$ to $e^*_i\boxtimes e^*_j$. Note that $\tau(\mathsf{h}\Xi_{n,d,\pi}) = \mathsf{h}\Xi_{n,d,\pi^{-1}}$, hence this induces an element in the stabilizer of $\mathsf{h}\Xi_{n,d}$. We will see in this section that this is the only other generator in the stabilizer.

We denote by $H^\circ\subseteq H$ the connected component of the identity of H. The Lie algebra $\mathfrak{h}$ of H° has been determined in §11.V. By matrix exponentiation we get that $H^\circ = \mathrm{im}\Theta$ with $\Theta : \mathrm{GL}_n^{d+1}\times\mathrm{GL}_n^{d+1}\to\mathcal{G}$, is the map with

$$\Theta((A_0,\ldots,A_{d+1}),(A'_0,\ldots,A'_{d+1})) = (A_0\boxtimes A'_0, [\Phi_{A,A'}], A_d^{-\mathsf{T}}\boxtimes {A'_d}^{-\mathsf{T}}),$$

where on $\mathcal{V}_{r,s}$ we have $\Phi_{A,A'}(X^{(r)}\boxtimes X^{(s)}) := A^{-1}_{r-1}X^{(r)}A_r\boxtimes A^{-1}_{s-1}X^{(s)}A_s$; and $\mathfrak{h} = \mathrm{im}\,d\Theta = d\Theta(\mathfrak{gl}_n^{\oplus 2(d+1)})$.

It is straightforward to determine

$$\ker\Theta \;=\; \Big\{\; \big((\lambda\alpha^k\mathrm{Id}_n)_{k=0}^d,\; (\lambda^{-1}\beta^l\mathrm{Id}_n)_{l=0}^d\big) \;\Big|\; \lambda,\alpha,\beta\in\mathbb{C}^*,\; (\alpha\beta)^d = 1\;\Big\},$$

and from that obtain the linear infinitesimal version

$$\ker d\Theta \;=\; \Big\{\big(((s+kt)\mathrm{Id}_n)_{k=0}^d,\; ((-s-lt)\mathrm{Id}_n)_{l=0}^d\big)\;\Big|\; s,t\in\mathbb{C}\Big\}.$$

Let $\mathcal{I} := \{0, 1, \ldots, d\} \cup \{0', 1', \ldots, d'\}$ be the set of vertex layers of both ABPs.

We have determined $\mathfrak{h}$, which is not semisimple, so as a preparatory step we calculate the derived Lie algebra $[\mathfrak{h}, \mathfrak{h}]$.

11.30 Proposition. $[\mathfrak{h}, \mathfrak{h}] = \bigoplus_{\iota \in \mathcal{I}} \mathfrak{g}_\iota$, *where each* $\mathfrak{g}_\iota$ *is isomorphic to* $\mathfrak{sl}_n(\mathbb{C})$. $\diamond$

Proof. Every element of $\ker d\Theta$ is a tuple of scalar matrices. Since the only scalar matrix in $\mathfrak{sl}_n$ is 0, we get $\ker d\Theta \cap \mathfrak{sl}_n^{\oplus 2(d+1)} = 0$, so $d\Theta$ is injective on $\mathfrak{sl}_n^{\oplus 2(d+1)}$.

$$[\mathfrak{h}, \mathfrak{h}] = [d\Theta(\mathfrak{gl}_n^{\oplus 2(d+1)}), d\Theta(\mathfrak{gl}_n^{\oplus 2(d+1)})] = d\Theta\big([\mathfrak{gl}_n^{\oplus 2(d+1)}, \mathfrak{gl}_n^{\oplus 2(d+1)}]\big) = d\Theta\big(\mathfrak{sl}_n^{\oplus 2(d+1)}\big) = \bigoplus_\iota \mathfrak{g}_\iota. \quad \square$$

Since $\mathfrak{sl}_n$ is simple, it follows that $[\mathfrak{h}, \mathfrak{h}]$ is semisimple. We use the semisimplicity in the following proposition.

11.31 Proposition. *Let* $\mathfrak{s} = \mathfrak{s}_1 \oplus \cdots \oplus \mathfrak{s}_r$ *be a direct sum of simple Lie algebras* $\mathfrak{s}_t$. *Then every Lie algebra automorphism of* $\mathfrak{s}$ *permutes the set* $\{\mathfrak{s}_1, \ldots, \mathfrak{s}_r\}$. $\diamond$

Proof. This is standard, but included for its importance.

First, $z(\mathfrak{s}_t) = 0$ for each t: the center is an ideal of $\mathfrak{s}_t$ different from $\mathfrak{s}_t$ (else $\mathfrak{s}_t$ would be abelian), hence 0 by simplicity.

Each $\mathfrak{s}_t$ *is a minimal non-zero ideal.* Let $\mathfrak{a} \neq 0$ be an ideal of $\mathfrak{s}$ with $\mathfrak{a} \subseteq \mathfrak{s}_t$. Then $[\mathfrak{s}_t, \mathfrak{a}] \subseteq [\mathfrak{s}, \mathfrak{a}] \subseteq \mathfrak{a}$, so $\mathfrak{a}$ is a non-zero ideal of the simple algebra $\mathfrak{s}_t$, hence $\mathfrak{a} = \mathfrak{s}_t$.

Every non-zero ideal $\mathfrak{a}$ *of* $\mathfrak{s}$ *contains some* $\mathfrak{s}_t$. Pick $0 \neq x = \sum_t x_t \in \mathfrak{a}$ with $x_t \in \mathfrak{s}_t$, and pick t with $x_t \neq 0$. Since simple nonabelian Lie algebras have zero center, $x_t \in \mathfrak{s}_t$ is not central, and hence there exists $y \in \mathfrak{s}_t$ with $[x_t, y] \neq 0$. By definition of the componentwise bracket, we have $[x, y] = [x_t, y] \neq 0$. Since $x \in \mathfrak{a}$, we have $[x, y] \in \mathfrak{a}$, and since $x_t \in \mathfrak{s}_t$, we have $[x_t, y] \in \mathfrak{s}_t$, which is a nonzero element in $\mathfrak{a} \cap \mathfrak{s}_t$. But $\mathfrak{a} \cap \mathfrak{s}_t$ is an ideal of $\mathfrak{s}_t$, because for every $u \in \mathfrak{a} \cap \mathfrak{s}_t$ and $z \in \mathfrak{s}_t$ we have $[z, u] \in \mathfrak{a}$ because $\mathfrak{a}$ is an ideal of $\mathfrak{s}$, and we have $[z, u] \in \mathfrak{s}_t$ because $\mathfrak{s}_t$ is a Lie subalgebra. Hence, $\mathfrak{a} \cap \mathfrak{s}_t$ is a nonzero ideal of the simple Lie algebra $\mathfrak{s}_t$ and hence $\mathfrak{a} \cap \mathfrak{s}_t = \mathfrak{s}_t$, or equivalently $\mathfrak{s}_t \subseteq \mathfrak{a}$.

In particular a *minimal* non-zero ideal contains, hence equals, some $\mathfrak{s}_t$, thus the set $\{\mathfrak{s}_1, \ldots, \mathfrak{s}_r\}$ is the set of minimal non-zero ideals. Let φ be an automorphism of $\mathfrak{s}$. Automorphisms send ideals to ideals, preserve nonzero subspaces, and preserve inclusion. Hence they send minimal nonzero ideals to minimal nonzero ideals. Hence φ must permute $\{\mathfrak{s}_1, \ldots, \mathfrak{s}_r\}$. $\square$

11.32 Proposition (Lemma 2.1 in [Ges16])**.** $H^\circ \subseteq H \subseteq N_G(H^\circ)$. $\diamond$

Proof. H° is a matrix Lie group, in particular a manifold and hence locally path-connected. Since every connected locally path-connected set is path-connected, it follows that H° is path-connected. Let $g \in H$ and $h \in H^\circ$. Choose a continuous $\alpha : [0, 1] \to H$ with $\alpha(0) = 1_H$ and $\alpha(1) = h$, and set $\beta(t) = g\alpha(t)g^{-1}$. Then β is continuous (conjugation by a fixed element is continuous), $\beta(0) = 1_H$, $\beta(1) = ghg^{-1}$, and $\beta(t) \in H$ for all t, since $\beta(t) \cdot f = g \cdot (\alpha(t) \cdot (g^{-1} \cdot f)) = f$. So β is a path in H from 1_H to ghg^{-1}, thus $ghg^{-1} \in H^\circ$; that is, $gH^\circ g^{-1} \subseteq H^\circ$. The same applied to $g^{-1} \in H$ gives $g^{-1}H^\circ g \subseteq H^\circ$, and conjugating by g yields $H^\circ \subseteq gH^\circ g^{-1}$. Hence $gH^\circ g^{-1} = H^\circ$ and $g \in N_G(H^\circ)$. $\square$

Our task is now to analyze which elements of $N_G(H^\circ)$ are actually in H. Let $h \in N_G(H^\circ)$ be arbitrary. The main tool is conjugating elements in $[\mathfrak{h}, \mathfrak{h}]$ with h. Since $h \in N_G(H^\circ)$, conjugation by h preserves H° and hence its Lie algebra $\mathfrak{h}$ and hence also $[\mathfrak{h}, \mathfrak{h}]$, thus $\mathrm{Ad}(h)|_{[\mathfrak{h},\mathfrak{h}]}$ is an automorphism of $[\mathfrak{h}, \mathfrak{h}]$. Let ψ_h be the permutation of $\mathcal{I}$ that this automorphism induces via Proposition 11.31, i.e., $\mathrm{Ad}(h)|_{[\mathfrak{h},\mathfrak{h}]}(\mathfrak{g}_\iota) = \mathfrak{g}_{\psi_h(\iota)}$.

11.33 Theorem. *ψ_h is either the identity, or ψ_h swaps each ι with ι'.* $\diamond$

Proof. Let $h = (h_{\text{left}}, h_{\text{middle}}, h_{\text{right}})$ with $h_{\text{mid}} \in \mathrm{GL}(\mathcal{V})$ such that $[h_{\text{mid}}] = h_{\text{middle}}$. $X = (X_{\text{left}}, X_{\text{mid}}, X_{\text{right}}) \in [\mathfrak{h}, \mathfrak{h}]$. $\mathcal{V}$ is a module for $[\mathfrak{h}, \mathfrak{h}]$ with pairwise non-isomorphic irreducibles $\mathcal{V}_{k,l}$. We write $h_{\text{mid}}\mathcal{V}$ for the vector space $\mathcal{V}$ with the action of $X \in [\mathfrak{h}, \mathfrak{h}]$ given by $\mathrm{Ad}(h)(X)$, i.e., $(\mathrm{Ad}(h)(X))_{\text{mid}}(h_{\text{mid}}(v)) = h_{\text{mid}}X_{\text{mid}}h_{\text{mid}}^{-1}h_{\text{mid}}(v) = h_{\text{mid}}X_{\text{mid}}v$, which is clearly isomorphic to $\mathcal{V}$, and $h_{\text{mid}}\mathcal{V}_{k,l}$ are the pairwise non-isomorphic irreducible submodules of $h_{\text{mid}}(\mathcal{V})$. Since $\mathrm{Ad}(h)$ is an automorphism of $[\mathfrak{h}, \mathfrak{h}]$, as X runs through $[\mathfrak{h}, \mathfrak{h}]$, $\mathrm{Ad}(h)(X)$ also runs through $[\mathfrak{h}, \mathfrak{h}]$, and hence the submodules of $\mathcal{V}$ and $h_{\text{mid}}(\mathcal{V})$ are the same. Since all submodules are pairwise non-isomorphic, this gives a bijection β of $[d] \times [d]$ on the edge layer pairs with $h_{\text{mid}}\mathcal{V}_{k,l} = \mathcal{V}_{\beta(k,l)}$. For edge layers $k, l \in [d]$, let $S_{k,l} := \{k-1,\ k,\ (l-1)',\ l'\}$ denote the set of incident vertex layers. Note that $S_{k,l} = \{\iota \in \mathcal{I} \mid \mathfrak{g}_\iota \text{ acts nontrivially on } \mathcal{V}_{k,l}\}$, and recall $\mathrm{Ad}(h)(\mathfrak{g}_\iota) = g_{\psi_h(\iota)}$. Hence $\iota \in S_{k,l}$ if and only if $\psi_h(\iota) \in S_{\beta(k,l)}$. Therefore,

$$\psi_h(S_{k,l}) \;=\; S_{\beta(k,l)}.$$

Let $\mathbb{C}^{n^2}$ denote the row space module. The action of $X = (X_{\text{left}}, X_{\text{mid}}, X_{\text{right}}) \in \mathfrak{g}_\iota$ on $\mathbb{C}^{n^2}$ is given via $X \mapsto X_{\text{left}}$. Observe that $\mathfrak{g}_\iota$ acts nontrivially on $\mathbb{C}^{n^2}$ if and only if $\iota \in \{0, 0'\}$. Let $\iota \in \{0, 0'\}$. Since $\mathfrak{g}_\iota$ acts nontrivially on $\mathbb{C}^{n^2}$, there exists $X \in \mathfrak{g}_\iota$ with $X_{\text{left}} \neq 0$. Now, $(\mathrm{Ad}(h)(X))_{\text{left}} = h_{\text{left}}X_{\text{left}}h_{\text{left}}^{-1} \neq 0$, and thus $\mathfrak{g}_{\psi_h(\iota)}$ acts nontrivially on $\mathbb{C}^{n^2}$. Therefore $\psi_h(\iota) \in \{0, 0'\}$, i.e., $\psi_h(\{0, 0'\}) \subseteq \{0, 0'\}$. Since ψ_h is bijective, we have

$$\psi_h(\{0, 0'\}) = \{0, 0'\} \tag{11.34}$$

The rest of the proof is an elementary combinatorial argument. We take the complete graph with vertex set $\mathcal{I}$ and for each pair (k, l) we remove the 6 edges in the complete graph on the vertex set $S_{k,l}$. We also remove the edges $\{0, 0'\}$ and $\{d, d'\}$. We call the resulting graph Δ. The map ψ_h induces a graph automorphism on Δ. Observe that Δ is just the union of two complements of path graphs $\Delta = \text{complement}(0 - 1 - \cdots - d) \cup \text{complement}(0' - 1' - \cdots - d')$. The graph automorphism group of Δ has 8 elements and is generated by swapping $\iota \leftrightarrow \iota'$, by swapping $\iota \leftrightarrow d - \iota$, and by swapping $\iota' \leftrightarrow (d-\iota)'$. By (11.34), the only remaining nontrivial graph automorphism is the swap $\iota \leftrightarrow \iota'$. $\square$

11.35 Theorem. *$H = H^\circ \cup \tau H^\circ$, where τ is the swap.* $\diamond$

Proof. Clearly $H^\circ \cup \tau H^\circ \subseteq H$. For the other direction, let $h \in H$ be arbitrary. By Theorem 11.33 we have ψ_h is the identity or the swap. If ψ_h is the swap, replace h by τh, so w.l.o.g. assume $\psi_h = \mathrm{id}$. It remains to prove $h \in H^\circ$. By Proposition 11.38 below, we have $h = g_0 h'$ with $g_0 \in H^\circ$ and $h' \in C_G(H^\circ)$. Since $g_0 \in H^\circ \subseteq H$ and $h \in H$, also $h' = g_0^{-1}h \in H$. Therefore $h' \in C_G(H^\circ) \cap H$, which equals $C_G(H^\circ) \cap H^\circ$ by Proposition 11.37 below, in particular $h' \in H^\circ$ and hence $h = g_0 h' \in H^\circ$. $\square$

It is easy to determine the centralizer $C_G(H^\circ)$ as $C_G(H^\circ) = (\mathbb{C}^*\mathrm{Id}_{n^2}) \times [T] \times (\mathbb{C}^*\mathrm{Id}_{n^2})$, where we recall the torus $T = (\mathbb{C}^*)^{[d]\times[d]}$. Moreover,

$$C_G(H^\circ) = \{g \in G \mid \mathrm{Ad}(g)|_{[\mathfrak{h},\mathfrak{h}]} = \mathrm{id}\}. \tag{11.36}$$

11.37 Proposition. *$C_G(H^\circ) \cap H \subseteq H^\circ$.* $\diamond$

Proof. An element $g = (\alpha_{\text{left}}, \boldsymbol{\alpha}, \alpha_{\text{right}}) \in C_G(H^\circ)$ rescales $\mathsf{h}\Xi_{n,d,\pi}$ by $\alpha_{\text{left}}\alpha_{\text{right}} \prod_{i=1}^d [\boldsymbol{\alpha}]_{\pi(k),k}$. Also assume that $g \in H$, thus g fixes each $\mathsf{h}\Xi_{n,d,\pi}$ independently. In particular, all products $\prod_{i=1}^d [\boldsymbol{\alpha}]_{\pi(k),k} = 1$ are independent of π. This implies that $\boldsymbol{\alpha}$ is a rank 1 matrix, i.e., there exist $u, v \in \mathbb{C}^d$ with $[\boldsymbol{\alpha}]_{i,j} = u_i v_j$. Hence, $\alpha_{\text{left}}\alpha_{\text{right}} \prod_{i=1}^d u_i \prod_{j=1}^d v_j = 1$. Define matrices $A_k = \alpha_{\text{left}} u_1 \cdots u_k \mathrm{Id}_n$, $A'_l = v_1 \cdots v_l \mathrm{Id}_n$. We have $A_0 \boxtimes A'_0 = \alpha_{\text{left}}\mathrm{Id}_{n^2}$, and $A_{k-1}^{-1}A_k = u_k \mathrm{Id}_n$, and ${A'_{l-1}}^{-1}A'_l = v_l \mathrm{Id}_n$. Moreover, $A_d^{-\mathsf{T}} \boxtimes {A'_d}^{-\mathsf{T}} = \left(\alpha_{\text{left}} \prod_{i=1}^d u_k \prod_{j=1}^d v_k\right)^{-1} \mathrm{Id}_{n^2} = \alpha_{\text{right}}\mathrm{Id}_{n^2}$. Therefore, $g = \Theta((A_0, \ldots, A_d), (A'_0, \ldots, A'_d)) \in \mathrm{im}\Theta = H^\circ$. $\square$

11.38 Proposition. *For $h \in N_G(H^\circ)$ and $\psi_h = \mathrm{id}$ we have $h \in H^\circ C_G(H^\circ)$.* $\diamond$

Proof. This proof also uses the adjoint action on modules as in the proof of Theorem 11.33. Recall that $\mathrm{Ad}(h)|_{[\mathfrak{h},\mathfrak{h}]}$ is an automorphism of $[\mathfrak{h},\mathfrak{h}]$. Since $\psi_h = \mathrm{id}$, conjugation with h preserves each $\mathfrak{g}_\iota$, each of which is isomorphic to $\mathfrak{sl}_n$. Hence, $\mathrm{Ad}(h)\big|_{\mathfrak{g}_\iota} \in \mathrm{Aut}(\mathfrak{sl}_n)$.

We claim that each of these automorphisms is an inner automorphism. By the classification of automorphisms of $\mathfrak{sl}_n$ (see e.g. [FH04, Prop. D.40]), every automorphism is inner or inner composed with $X \mapsto -X^\mathsf{T}$. In fact, for $n = 2$ every automorphism is inner and we are done, hence we have $n \geq 3$ from now on. If an inner automorphism is composed with the action of $\mathfrak{sl}_n$ on the standard module $\mathbb{C}^n$, we obtain an isomorphic module, but if the other type of automorphism is composed with the action of $\mathfrak{sl}_n$ on the standard module $\mathbb{C}^n$, we obtain a module that is isomorphic to the dual $\mathbb{C}^{n*}$. Note that $\mathbb{C}^n \not\simeq \mathbb{C}^{n*}$ for $n \geq 3$. The $\mathfrak{g}_\iota$ act as follows on row-spaces and variable spaces and column spaces:

- Row space $\mathbb{C}^{n^2} = \mathbb{C}^n_{\mathfrak{g}_0} \otimes \mathbb{C}^n_{\mathfrak{g}_{0'}}$.
- Variable space in edge layer pair k, l': $\mathbb{C}^{n^4} = \mathbb{C}^n_{\mathfrak{g}_k} \otimes \mathbb{C}^{n*}_{\mathfrak{g}_{k-1}} \otimes \mathbb{C}^n_{\mathfrak{g}_{l'}} \otimes \mathbb{C}^{n*}_{\mathfrak{g}_{(l-1)'}}$.
- Column space $\mathbb{C}^{n^2*} = \mathbb{C}^{n*}_{\mathfrak{g}_d} \otimes \mathbb{C}^{n*}_{\mathfrak{g}_{d'}}$, but we will not need the column space here.

With respect to $\mathfrak{g}_0$ we have $\mathbb{C}^{n^2} = \mathbb{C}^n_{\mathfrak{g}_0} \otimes \mathbb{C}^n$, where $\mathbb{C}^n$ is a multiplicity space. With respect to $\mathfrak{g}_k$ we have $\mathbb{C}^{n^4} = \mathbb{C}^n_{\mathfrak{g}_k} \otimes \mathbb{C}^{n^3}$. But on these spaces, h acts by conjugation, as in the proof of Theorem 11.33, which gives isomorphic representations. Thus the action on each factor is an inner automorphism, and hence each $\mathrm{Ad}(h)\big|_{\mathfrak{g}_\iota}$ is an inner automorphism.

Let $S_\iota \in \mathrm{GL}_n$ be such that the conjugation by h on $\mathfrak{g}_\iota$ acts as $X \mapsto S_\iota X S_\iota^{-1}$, i.e., $\mathrm{Ad}(h)(X) = S_\iota X S_\iota^{-1}$ for every $X \in \mathfrak{g}_\iota$. Set $g_0 := \Theta((S_0, \ldots, S_d), (S_{0'}, \ldots, S_{d'}))$, where $\Theta : \mathrm{GL}_n^{d+1} \times \mathrm{GL}_n^{d+1} \to G$ is the map with $\mathrm{im}\Theta = H^\circ$, therefore $g_0 \in H^\circ$. One now checks that g_0 and h have the same conjugation action on each $\mathfrak{g}_\iota$, which implies $\mathrm{Ad}(g_0)\big|_{[\mathfrak{h},\mathfrak{h}]} = \mathrm{Ad}(h)\big|_{[\mathfrak{h},\mathfrak{h}]}$. Now define $c := g_0^{-1}h$. Since $\mathrm{Ad}(c)\big|_{[\mathfrak{h},\mathfrak{h}]} = \mathrm{Ad}(g_0)^{-1}\mathrm{Ad}(h)\big|_{[\mathfrak{h},\mathfrak{h}]} = \mathrm{id}_{[\mathfrak{h},\mathfrak{h}]}$, by (11.36) it follows $c \in C_G(H^\circ)$. This finishes the proof, because $h = g_0 c$. $\square$

It follows from Theorem 11.35 that if $\pi = \pi^{-1}$, then $\mathsf{h}\Xi_{n,d,\pi}$ has all the symmetries of $\mathsf{h}\Xi_{n,d}$, hence

$$\mathsf{h}\Xi_{n,d} \text{ is not characterized by its stabilizer.} \tag{11.39}$$

In general, $\mathsf{h}\Xi_{n,d,\pi} + \mathsf{h}\Xi_{n,d,\pi^{-1}}$ has all the symmetries of $\mathsf{h}\Xi_{n,d}$.

Acknowledgments

The authors greatly benefitted from discussions with Maxim van den Berg, Fulvio Gesmundo, Darij Grinberg, Gorav Jindal, JM Landsberg, Vladimir Lysikov, Jakob Moosbauer, and Harold Nieuwboer.

AI Disclosure

Generative AI was used for literature search and for proofreading.

ChatGPT 5.6 Sol found the proof of the VNP-completeness of U_{2d} that we presented in §9.I, which is more elegant than our previous proof sketch.

We used ChatGPT 5.6 Sol and Claude Fable 5 extensively to obtain the results in §11.V and §11.VI, i.e., the calculation of the stabilizer of $\mathsf{h}\Xi_{n,d}$ by translating the technique from [Ges16] to our situation. The authors verified the correctness and originality of all content including references.

References

[AB09] Sanjeev Arora and Boaz Barak. *Computational complexity.* Cambridge University Press, Cambridge, 2009. A modern approach.

[AFP+19] Marian Aprodu, Gavril Farkas, Ștefan Papadima, Claudiu Raicu, and Jerzy Weyman. Koszul modules and Green's conjecture. *Inventiones mathematicae*, 218(3):657–720, 2019.

[AKSS95] Vikraman Arvind, Johannes Köbler, Uwe Schöning, and Rainer Schuler. If NP has polynomial-size circuits, then MA=AM. *Theoretical Computer Science*, 137(2):279–282, 1995.

[AR08] Shlomo Ahal and Yuri Rabinovich. On complexity of the subpattern problem. *SIAM Journal on Discrete Mathematics*, 22(2):629–649, 2008.

[ASS23] Bharat Adsul, Milind Sohoni, and K V Subrahmanyam. Orbit closures, stabilizer limits and intermediate G-varieties, 2023.

[AY23] Alimzhan Amanov and Damir Yeliussizov. Tensor slice rank and Cayley's first hyperdeterminant. *Linear Algebra and its Applications*, 656:224–246, 2023.

[Bap89] R.B. Bapat. Mixed discriminants of positive semidefinite matrices. *Linear Algebra and its Applications*, 126:107–124, 1989.

[Bar95] Alexander I Barvinok. New algorithms for linear k-matroid intersection and matroid k-parity problems. *Mathematical Programming*, 69(1):449–470, 1995.

[Bar16] Alexander Barvinok. *Combinatorics and complexity of partition functions*, volume 30. Springer, 2016.

[BCL23] Peter Bürgisser, Felipe Cucker, and Pierre Lairez. Rigid continuation paths II. Structured polynomial systems. In *Forum of Mathematics, Pi*, volume 11, page e12. Cambridge University Press, 2023.

[BDI21] Markus Bläser, Julian Dörfler, and Christian Ikenmeyer. On the complexity of evaluating highest weight vectors. In *Proceedings of the 36th Computational Complexity Conference*, CCC '21, Dagstuhl, DEU, 2021. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[BE19] Markus Bläser and Christian Engels. Parameterized Valiant's Classes. In *14th International Symposium on Parameterized and Exact Computation (IPEC 2019)*, volume 148 of *LIPIcs*, pages 3:1–3:14, 2019. The arXiv version contains the full proofs: 1907.12287v2.

[BGL14] H. Bermudez, S. Garibaldi, and V. Larsen. Linear preservers and representations with a 1-dimensional ring of invariants. *Transactions of the American Mathematical Society*, 366(9):4755–4780, 2014.

[BI11] Peter Bürgisser and Christian Ikenmeyer. Geometric complexity theory and tensor rank. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 509–518, 2011.

[BI17] Peter Bürgisser and Christian Ikenmeyer. Fundamental invariants of orbit closures. *Journal of Algebra*, 477:390–434, 2017.

[BIL+21] Markus Bläser, Christian Ikenmeyer, Vladimir Lysikov, Anurag Pandey, and Frank-Olaf Schreyer. On the orbit closure containment problem and slice rank of tensors. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2565–2584. 2021.

[BIM+20] Markus Bläser, Christian Ikenmeyer, Meena Mahajan, Anurag Pandey, and Nitin Saurabh. Algebraic branching programs, border complexity, and tangent spaces. In *35th Computational Complexity Conference*, 2020.

[BIP19] Peter Bürgisser, Christian Ikenmeyer, and Greta Panova. No occurrence obstructions in geometric complexity theory. *Journal of the American Mathematical Society*, 32(1):163–193, 2019.

[Blä13] Markus Bläser. Fast matrix multiplication. *Theory of Computing*, pages 1–60, 2013.

[BLMW11] Peter Bürgisser, Joseph M Landsberg, Laurent Manivel, and Jerzy Weyman. An overview of mathematical issues arising in the geometric complexity theory approach to VP$\neq$VNP. *SIAM Journal on Computing*, 40(4):1179–1209, 2011.

[Bür00] Peter Bürgisser. Cook's versus Valiant's hypothesis. *Theoretical Computer Science*, 235(1):71–88, 2000.

[Cay46] Arthur Cayley. Mémoire sur les hyperdéterminants. 1846.

[CFTZ21] Matthias Christandl, Omar Fawzi, Hoang Ta, and Jeroen Zuiddam. Symmetric subrank of tensors and applications. arXiv:2104.01130, 2021.

[CHI+18] Luca Chiantini, Jonathan D. Hauenstein, Christian Ikenmeyer, Joseph M. Landsberg, and Giorgio Ottaviani. Polynomials and the exponent of matrix multiplication. *Bull. Lond. Math. Soc.*, 50(3):369–389, 2018.

[CKV24] Abhranil Chatterjee, Mrinal Kumar, and Ben Lee Volk. Determinants vs. algebraic branching programs. *computational complexity*, 33(2):11, 2024.

[CP16] Diego Cifuentes and Pablo A Parrilo. An efficient tree decomposition method for permanents and mixed discriminants. *Linear Algebra and its Applications*, 493:45–81, 2016.

[DF13] R. G. Downey and M. R. Fellows. *Fundamentals of Parameterized Complexity.* Texts in Computer Science. Springer, 2013.

[DG87] Stephen J Dow and Peter M Gibson. Permanents of d-dimensional matrices. *Linear Algebra and its Applications*, 90:133–145, 1987.

[DGI+24] Pranjal Dutta, Fulvio Gesmundo, Christian Ikenmeyer, Gorav Jindal, and Vladimir Lysikov. Homogeneous Algebraic Complexity Theory and Algebraic Formulas. In *15th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 287 of *LIPIcs*, pages 43:1–43:23, 2024.

[DIP20] Julian Dörfler, Christian Ikenmeyer, and Greta Panova. On geometric complexity theory: Multiplicity obstructions are stronger than occurrence obstructions. *SIAM Journal on Applied Algebra and Geometry*, 4(2):354–376, 2020.

[DV22] Harm Derksen and Emanuele Viola. Fooling polynomials using invariant theory. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 399–406, 2022.

[Eis13] David Eisenbud. *Commutative algebra: with a view toward algebraic geometry*, volume 150. Springer Science & Business Media, 2013.

[FH04] William Fulton and Joe Harris. *Representation Theory: A First Course*, volume 129 of *Graduate Texts in Mathematics*. Springer, 2004.

[For16] Michael Forbes. Some concrete questions on the border complexity of polynomials. Workshop on Algebraic Complexity Theory (WACT), `https://www.youtube.com/watch?v=1HMogQIHT6Q`, 2016.

[Fro97] Georg Ferdinand Frobenius. Über die Darstellung der endlichen Gruppen durch lineare Substitutionen. pages 994–1015, 1897.

[Ges16] Fulvio Gesmundo. Geometric aspects of iterated matrix multiplication. *Journal of Algebra*, 461:42–64, 2016.

[GGOW16] Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. A deterministic polynomial time algorithm for non-commutative rational identity testing. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 109–117. IEEE, 2016.

[Gha26] Masoud Gharahi. On the kronecker products of symmetric persistent tensors. arXiv:2608.11182, 2026.

[GIM+20] Ankit Garg, Christian Ikenmeyer, Visu Makam, Rafael Oliveira, Michael Walter, and Avi Wigderson. Search Problems in Algebraic Complexity, GCT, and Hardness of Generators for Invariant Rings. In *35th Computational Complexity Conference (CCC 2020)*, volume 169 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 12:1–12:17, 2020.

[Gre11] Bruno Grenet. An Upper Bound for the Permanent versus Determinant Problem. Manuscript, 2011.

[Gur04] Leonid Gurvits. Classical complexity and quantum entanglement. *Journal of Computer and System Sciences*, 69(3):448–484, 2004. Special Issue on STOC 2003.

[Gur05] Leonid Gurvits. On the complexity of mixed discriminants and related problems. In *Mathematical Foundations of Computer Science 2005: 30th International Symposium, MFCS 2005, Gdansk, Poland, August 29–September 2, 2005. Proceedings 30*, pages 447–458. Springer, 2005.

[HL13] Christopher J Hillar and Lek-Heng Lim. Most tensor problems are NP-hard. *Journal of the ACM (JACM)*, 60(6):1–39, 2013.

[HMS19] Corey Harris, Mateusz Michałek, and Emre Can Sertöz. Computing images of polynomial maps. *Advances in Computational Mathematics*, 45:2845–2865, 2019.

[Hüt17] Jesko Hüttenhain. *Geometric complexity theory and orbit closures of homogeneous forms.* PhD thesis, TU Berlin, 2017.

[Hya79] Laurent Hyafil. On the parallel evaluation of multivariate polynomials. *SIAM J. Comput.*, 8(2):120–123, 1979.

[IK20] Christian Ikenmeyer and Umangathan Kandasamy. Implementing geometric complexity theory: On the separation of orbit closures via symmetries. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 713–726, 2020.

[Ike22] Christian Ikenmeyer. De-bordering symmetric border rank and other open problems. Talk at the AGATES (Algebraic Geometry with Applications to TEnsors and Secants) kickoff workshop at Warsaw University, `https://agates.mimuw.edu.pl/images/warsaw_ikenmeyer.pdf`, 2022.

[Ike25] Christian Ikenmeyer. On the gradient of the coefficient of the characteristic polynomial. arXiv:2511.04954, 2025.

[IOT25] Christian Ikenmeyer, Heidi Omar, and Dimitrios Tsintsilidas. Field-independent Kronecker-plethysm isomorphisms. arXiv:2509.10069v1, 2025.

[IP17] Christian Ikenmeyer and Greta Panova. Rectangular Kronecker coefficients and plethysms in geometric complexity theory. *Advances in Mathematics*, 319:40–66, 2017.

[Kar72] Richard M. Karp. *Reducibility among Combinatorial Problems*, pages 85–103. 1972.

[Kem78] George R Kempf. Instability in invariant theory. *Annals of Mathematics*, 108(2):299–316, 1978.

[KL82] Richard M. Karp and Richard J. Lipton. Turing machines that take advice. *L'Enseignement Mathématique*, 1982.

[KL14] Harlan Kadish and Joseph M Landsberg. Padded polynomials, their cousins, and geometric complexity theory. *Communications in Algebra*, 42(5):2171–2180, 2014.

[KPR23] Balagopal Komarath, Anurag Pandey, and Chengot Sankaramenon Rahul. Monotone arithmetic complexity of graph homomorphism polynomials. *Algorithmica*, 85(9):2554–2579, 2023.

[Kra84] Hanspeter Kraft. *Geometrische Methoden in der Invariantentheorie*. Aspects of Mathematics, D1. Friedr. Vieweg & Sohn, Braunschweig, 1984.

[KT99] Allen Knutson and Terence Tao. The honeycomb model of $\mathrm{GL}_n(\mathbf{C})$ tensor products. I. Proof of the saturation conjecture. *J. Amer. Math. Soc.*, 12(4):1055–1090, 1999.

[Kum19] Mrinal Kumar. A quadratic lower bound for homogeneous algebraic branching programs. *computational complexity*, 28(3):409–435, 2019.

[Lan11] Joseph M Landsberg. *Tensors: geometry and applications*, volume 128. American Mathematical Soc., 2011.

[Lan15] Joseph M Landsberg. Geometric complexity theory: an introduction for geometers. *Annali dell'universita'di Ferrara*, 61(1):65–117, 2015.

[LC83] J. S. Lomont and M. S. Cheena. A multilinearity property of determinant functions. *Linear and Multilinear Algebra*, 14(3):199–223, 1983.

[LP90] M. Larsen and R. Pink. Determining representations from invariant dimensions. *Invent. Math.*, 102(2):377–398, 1990.

[Lun75] Domingo Luna. Adhérences d'orbite et invariants. *Inventiones mathematicae*, 29(3):231–238, 1975.

[mC20] The mathlib Community. The Lean Mathematical Library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs*, CPP 2020, New Orleans, LA, USA, January 2020. ACM.

[Min84] Henryk Minc. *Permanents*, volume 6. Cambridge University Press, 1984.

[MM62] Marvin Marcus and F. C. May. The permanent function. *Canadian Journal of Mathematics*, 14:177–189, 1962.

[MS01] Ketan D Mulmuley and Milind Sohoni. Geometric complexity theory I: An approach to the P vs. NP and related problems. *SIAM Journal on Computing*, 31(2):496–526, 2001.

[MS08] Ketan D Mulmuley and Milind Sohoni. Geometric complexity theory II: Towards explicit obstructions for embeddings among class varieties. *SIAM Journal on Computing*, 38(3):1175–1206, 2008.

[Mul99] Ketan Mulmuley. Lower bounds in a parallel model without bit operations. *SIAM Journal on Computing*, 28(4):1460–1509, 1999.

[MV97] Meena Mahajan and V Vinay. Determinant: Combinatorics, algorithms, and complexity. *Chicago Journal of Theoretical Computer Science*, 1997(5), 1997.

[Nis91] Noam Nisan. Lower bounds for non-commutative computation. In *Proceedings of the twenty-third annual ACM symposium on Theory of computing*, pages 410–418, 1991.

[NP85] Jaroslav Nešetřil and Svatopluk Poljak. On the complexity of the subgraph problem. *Commentationes Mathematicae Universitatis Carolinae*, 26(2):415–419, 1985.

[Pan85] Alexey Andreevich Panov. Mixed discriminants connected with positive semidefinite quadratic forms. In *Doklady Akademii Nauk*, volume 282, pages 273–276. Russian Academy of Sciences, 1985.

[RS84] Christophe Reutenauer and Howard Straubing. Inversion of matrices over a commutative semiring. *Journal of algebra*, 88(2):350–360, 1984.

[Sap15] Ramprasad Saptharishi. A survey of lower bounds in arithmetic circuit complexity. *Github survey*, 95, 2015. version 9.0.3.

[Sch76] Claus-Peter Schnorr. A lower bound on the number of additions in monotone computations. *Theoretical Computer Science*, 2(3):305–315, 1976.

[SPM+15] Szilárd Szalay, Max Pfeffer, Valentin Murg, Gergely Barcza, Frank Verstraete, Reinhold Schneider, and Örs Legeza. Tensor product methods and entanglement optimization for ab initio quantum chemistry. *International Journal of Quantum Chemistry*, 115:1342–1391, 2015.

[Sta00] Richard P Stanley. Positivity problems and conjectures in algebraic. *Mathematics: Frontiers and Perspectives*, pages 295–319, 2000.

[Str69] Volker Strassen. Gaussian elimination is not optimal. *Numer. Math.*, 13:354–356, 1969.

[Tic15] Malte C Tichy. Sampling of partially distinguishable bosons and the relation to the multidimensional permanent. *Physical Review A*, 91(2):022316, 2015.

[TS16] Terence Tao and Will Sawin. Notes on the "slice rank" of tensors. `https://terrytao.wordpress.com/2016/08/24/notes-on-the-slice-rank-of-tensors/`, August 2016.

[Tse68] Grigori S Tseitin. On the complexity of proof in propositional calculus. In *Studies in Constructive Mathematics and Mathematical Logic, Part II*, volume 8 of *Zap. Nauchn. Sem. LOMI*, pages 234–259. Nauka, Leningrad, 1968.

[Tse83] Grigori S Tseitin. On the complexity of derivation in propositional calculus. In *Automation of reasoning: 2: Classical papers on computational logic 1967–1970*, pages 466–483. Springer, 1983.

[Val79] Leslie Valiant. Completeness classes in algebra. In *Proceedings of the ACM Symposium on Theory of Computing*, pages 249–261, 1979.

[Val82] L. G. Valiant. Reducibility by algebraic projections. *Enseign. Math. (2)*, 28(3-4):253–268, 1982.

[vdBDG+25] Maxim van den Berg, Pranjal Dutta, Fulvio Gesmundo, Christian Ikenmeyer, and Vladimir Lysikov. Algebraic Metacomplexity and Representation Theory. In *40th Computational Complexity Conference (CCC)*, volume 339 of *LIPIcs*, pages 26:1–26:35, 2025.

[VSBR83] L. G. Valiant, S. Skyum, S. Berkowitz, and C. Rackoff. Fast parallel computation of polynomials using few processors. *SIAM Journal on Computing*, 12(4):641–644, 1983.

[Wei86] Hanns Joachim Weinert. Generalized semialgebras over semirings. In *Semigroups Theory and Applications: Proceedings of a Conference held in Oberwolfach, Feb. 23–Mar. 1, 1986*, pages 380–416. Springer, 1986.

[Yeh19] Amir Yehudayoff. Separating monotone VP and VNP. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 425–429, 2019.

[Yu16] Jun Yu. On the dimension datum of a subgroup. *Duke Mathematical Journal*, 165(14):2683–2736, 2016.

`Ahmed-Ittihad.Hasib@warwick.ac.uk`

`Christian.Ikenmeyer@warwick.ac.uk`

`Dimitrios.Tsintsilidas@warwick.ac.uk`