Boundaries of VP and VNP

Joshua A. Grochow[*], Ketan D. Mulmuley[†], and Youming Qiao[‡]

May 10, 2016

Abstract

One fundamental question in the context of the geometric complexity theory approach to the VP vs. VNP conjecture is whether $\mathrm{VP} = \overline{\mathrm{VP}}$, where VP is the class of families of polynomials that are of polynomial degree and can be computed by arithmetic circuits of polynomial size, and $\overline{\mathrm{VP}}$ is the class of families of polynomials that are of polynomial degree and can be approximated infinitesimally closely by arithmetic circuits of polynomial size. The goal of this article[1] is to study the conjecture in (Mulmuley, FOCS 2012) that $\overline{\mathrm{VP}}$ is not contained in VP.

Towards that end, we introduce three degenerations of VP (i.e., sets of points in $\overline{\mathrm{VP}}$), namely the stable degeneration Stable-VP, the Newton degeneration Newton-VP, and the p-definable one-parameter degeneration VP*. We also introduce analogous degenerations of VNP. We show that Stable-VP $\subseteq$ Newton-VP $\subseteq$ VP* $\subseteq$ VNP, and Stable-VNP = Newton-VNP = VNP* = VNP. The three notions of degenerations and the proof of this result shed light on the problem of separating $\overline{\mathrm{VP}}$ from VP.

Although we do not yet construct explicit candidates for the polynomial families in $\overline{\mathrm{VP}} \setminus \mathrm{VP}$, we prove results which tell us where not to look for such families. Specifically, we demonstrate that the families in Newton-VP $\setminus$VP based on semi-invariants of quivers would have to be non-generic by showing that, for many finite quivers (including some wild ones), any Newton degeneration of a generic semi-invariant can be computed by a circuit of polynomial size. We also show that the Newton degenerations of perfect matching Pfaffians, monotone arithmetic circuits over the reals, and Schur polynomials have polynomial-size circuits.

1 Introduction

One fundamental question in the context of the geometric complexity theory (GCT) approach (cf. [36,37], [8], and [35]) to the VP vs. VNP conjecture in Valiant [49] is whether $\mathrm{VP} = \overline{\mathrm{VP}}$, where VP is the class of families of polynomials that are of polynomial degree and can be computed by arithmetic circuits of polynomial size, VNP is the class of p-definable families of polynomials, and $\overline{\mathrm{VP}}$ is the class of families of polynomials that are of polynomial degree and can be approximated infinitesimally closely by arithmetic circuits of polynomial size. We assume in what follows that the circuits are over an algebraically closed field $\mathbb{F}$. We call $\overline{\mathrm{VP}}$ the *closure* of VP, and $\overline{\mathrm{VP}} \setminus \mathrm{VP}$ the *boundary* of VP. So the question is whether this boundary is non-empty. At present, it is not even known if $\overline{\mathrm{VP}}$ is contained in VNP.

The VP vs. $\overline{\mathrm{VP}}$ question is important for two reasons. First, all known algebraic lower bounds for the exact computation of the permanent also hold for its infinitesimally close approximation. For

[*]Santa Fe Institute, `jgrochow@santafe.edu`

[†]University of Chicago, `mulmuley@uchicago.edu`

[‡]University of Technology Sydney, `Youming.Qiao@uts.edu.au`

[1]A preliminary version of this paper appears in [20]

example, the known quadratic lower bound for the permanent [34] also holds for its infinitesimally closely approximation [31], and so also the known lower bounds in the algebraic depth-three circuit models [25]; cf. Appendix B in [19] for a survey of the known lower bounds which emphasizes this point. These lower bounds hold because some algebraic, polynomial property that is satisfied by the coefficients of the polynomials computed by the circuits in the restricted class under consideration is not satisfied by the coefficients of the permanent. Since a polynomial property is a closed condition,[2] the same property is also satisfied by the coefficients of the polynomials that can be approximated infinitesimally closely[3] by circuits in the restricted class under consideration. This is why the same lower bound also holds for infinitesimally close approximation. We expect the same phenomenon to hold in the unrestricted algebraic circuit model as well. Hence, it is natural to expect that any realistic proof of the VP $\neq$ VNP conjecture will also show that VNP $\not\subseteq \overline{\mathrm{VP}}$, as conjectured in [36].[4] This is, in fact, the underlying thesis of geometric complexity theory stated in [35]. But, if $\overline{\mathrm{VP}} \neq$ VP, as conjectured in [35], this would mean that any realistic approach to the VP vs. VNP conjecture would even have to separate the permanent from the families in $\overline{\mathrm{VP}} \setminus$ VP with high circuit complexity.[5]

Second, it is shown in [35] that, assuming a stronger form of the VNP $\not\subseteq \overline{\mathrm{VP}}$ conjecture, the problem NNL (short for Noether's Normalization Lemma) of computing the Noether normalization of explicit varieties can be brought down from EXPSPACE, where it is currently, to P, ignoring a quasi-prefix. The existing EXPSPACE vs. P gap,[6] called the geometric complexity theory (GCT) chasm [35], in the complexity of NNL may be viewed as the common cause and measure of the difficulty of the fundamental problems in geometry (NNL) and complexity theory (Hardness). If $\overline{\mathrm{VP}} =$ VP, then it follows [35] that NNL is in PSPACE. Thus the conjectural inequality between $\overline{\mathrm{VP}}$ and VP is the main difficulty that needs to be overcome to bring NNL from EXPSPACE to PSPACE unconditionally, and is the main reason why the standard techniques in complexity theory may not be expected to work in the context of the VP $\neq$ VNP conjecture.

The goal of this article is to study the conjecture in [35] that $\overline{\mathrm{VP}}$ is not contained in VP.

1.1 Degenerations of VP and VNP

Towards that end, we introduce three notions of degenerations of VP and VNP; "degeneration" is the standard term in algebraic geometry for a limit point or infinitesimal approximation. These degenerations are subclasses of $\overline{\mathrm{VP}}$ and $\overline{\mathrm{VNP}}$, respectively; cf. Section 3 for formal definitions.

The first notion is that of a stable degeneration. Recall [38] that a polynomial $f \in \mathbb{F}[x_1, \ldots, x_m]$ is called *stable* with respect to the natural action of $G = \mathrm{SL}(m, \mathbb{F})$ on $\mathbb{F}[x_1, \ldots, x_m]$ if the G-orbit of f is closed in the Zariski topology. If $\mathbb{F} = \mathbb{C}$, we may equivalently say closed in the usual complex topology. Here G is $\mathrm{SL}(m, \mathbb{F})$, and not $\mathrm{GL}(m, \mathbb{F})$, since the only polynomial in $\mathbb{F}[x_1, \ldots, x_m]$ with a closed orbit with respect to the action of $\mathrm{GL}(m, \mathbb{F})$ is identically zero. Hence, whenever we study issues related to stability in this article, we only consider orbits with respect to the SL-action.

We say that a polynomial f is a *stable degeneration* of $g \in \mathbb{F}[x_1, \ldots, x_m]$ if f lies in a closed G-orbit (which is unique [38]) in the closure of the G-orbit of g. The degeneration is called stable since f in this case is stable. We say that a polynomial family $\{f_n\}$ is a *stable degeneration* of $\{g_n\}$

[2] It is defined by the vanishing of a continuous function, namely, a (meta) polynomial.

[3] This means the polynomials are the limits of the polynomials computed by the circuits in the restricted class under consideration.

[4] Note that if VNP $\not\subseteq \overline{\mathrm{VP}}$ then there exists a polynomial property showing this lower bound.

[5] Although some lower bounds techniques in the restricted models do distinguish between different polynomials with high circuit complexity (e.g., [41]), we need a better understanding of the families in $\overline{\mathrm{VP}} \setminus$ VP in order to know which techniques in this spirit could even potentially be useful in the setting of the VNP versus $\overline{\mathrm{VP}}$ problem.

[6] Or, the EXPH vs. P gap, assuming the Generalized Riemann Hypothesis.

if each f_n is a stable degeneration of g_n, with respect to the action of $G = \mathrm{SL}(m_n, \mathbb{F})$, where m_n denotes the number of variables in f_n and g_n. For any class of polynomial families $\mathcal{C}$, the class Stable-$\mathcal{C}$ is defined to be the class of families of polynomials that are either in $\mathcal{C}$ or are stable degenerations thereof.

The second notion is that of a Newton degeneration. We say that a polynomial f is a *Newton degeneration* of g if it is obtained from g by keeping only those terms whose associated monomial-exponents lie in some specified face of the Newton polytope of g. We say that a polynomial family $\{f_n\}$ is a *Newton degeneration* of $\{g_n\}$ if each f_n is a Newton degeneration of g_n. We say that $\{f_n\}$ is a *linear projection* of $\{g_n\}$ if each f_n is a linear projection of g_n.[7] For any class of polynomial families $\mathcal{C}$, the class Newton-$\mathcal{C}$ is defined to be the class of families of polynomials that are Newton degenerations of the polynomial families in $\mathcal{C}$, or are linear projections of such Newton degenerations.[8]

The third notion, motivated by the notion of p-definability in Valiant [49], is that of a p-definable one-parameter degeneration. We say that a family $\{f_n\}$ of polynomials is a *p-definable one-parameter degeneration* of a family $\{g_n\}$ of polynomials, if $f_n(x) = \lim_{t\to 0} g_n(x,t)$, where $g_n(x,t)$ is obtained from $g_n(x)$ by transforming its variables $x = (x_1, \ldots, x_i, \ldots)$ linearly such that: (1) the entries of the linear transformation matrix are Laurent polynomials in t of possibly exponential degree (in n), and (2) there exists a small circuit C_n over $\mathbb{F}$ of size polynomial in n such that any coefficient of the Laurent polynomial in any entry of the transformation matrix can be obtained by evaluating C_n at the indices of that entry and the index of the coefficient. It is assumed here that the indices are encoded as lists of 0-1 variables, treating 0 and 1 as elements of $\mathbb{F}$. Thus a p-definable one-parameter degeneration is a one-parameter degeneration of exponential degree that can be encoded by a small circuit. For any class of polynomial families $\mathcal{C}$, the class $\mathcal{C}^*$ is defined to be the class of families of polynomials that are p-definable one-parameter degenerations of the families in $\mathcal{C}$.

The classes $\overline{\mathrm{VP}}$ and $\overline{\mathrm{VNP}}$ are closed under these three types of degenerations (cf. Propositions 3.2, 3.3, 3.6). Since we want to compare $\overline{\mathrm{VP}}$ with VP, and $\overline{\mathrm{VNP}}$ with VNP, we ask how VP and VNP behave under these degenerations. This is addressed in the following result.

Theorem 1. *(a) Stable-VNP = Newton-VNP = VNP* = VNP, and*
(b) Stable-VP $\subseteq$ Newton-VP $\subseteq$ VP* $\subseteq$ VNP.

An analogue of this result also holds for $\mathrm{VP}_{\mathrm{ws}}$, the class of families of polynomials that can be computed by symbolic determinants of polynomial size.

1.2 On VP* vs. $\overline{\mathrm{VP}}$ and VP vs. Stable-VP

The statement of Theorem 1 tells us nothing as to whether any of the inclusions in the sequence VP $\subseteq$ Stable-VP $\subseteq$ Newton-VP $\subseteq$ VP* $\subseteq$ $\overline{\mathrm{VP}}$ can be expected to be strict or not. But its proof, as discussed below, does shed light on this subject.

Theorem 1 is proved by combining the Hilbert-Mumford-Kempf criterion for stability [26] with the ideas and results in Valiant [49]. The Hilbert-Mumford-Kempf criterion [26] shows that, for any polynomial f in the unique closed G-orbit in the G-orbit-closure of any $g \in \mathbb{F}[x_1, \ldots, x_m]$, with $G = \mathrm{SL}(m, \mathbb{F})$, there exists a one-parameter subgroup of G that drives g to f. Furthermore, by

[7] This means f_n is obtained from g_n by a linear (possibly non-homogeneous) change of variables.

[8] Taking a Newton degeneration and a linear projection need not commute, so the set of Newton degenerations alone will not in general be closed under linear projections. For example, any polynomial f is a linear projection of a sufficiently large determinant, but the Newton degenerations of the determinant only consist of polynomials of the form $\det(X')$ where X' is matrix consisting only of variables and 0s.

Kempf [26], such a subgroup can be chosen in a canonical manner. As a byproduct of the proof of Theorem 1, we get a complexity-theoretic form of this criterion (cf. Theorem 7), which shows that such a one-parameter group can be chosen so that the resulting one-parameter degeneration of any $\{g_n\} \in$ VP to $\{f_n\} \in$ Stable-VP is p-definable. Here f_n is a stable degeneration of g_n with respect to the action of $\mathrm{SL}(m_n, \mathbb{F})$, where $m_n = \mathsf{poly}(n)$ denotes the number of variables in f_n and g_n. Thus the inclusion of Stable-VP in VNP ultimately depends on the existence of a p-definable one parameter degeneration of $\{g_n\}$ to $\{f_n\}$, as provided by the Hilbert-Mumford-Kempf criterion.

However, no such p-definable one parameter degeneration scheme is known if f_n is allowed to be any polynomial with a non-closed $\mathrm{SL}(m_n, \mathbb{F})$-orbit in the $\mathrm{SL}(m_n, \mathbb{F})$-orbit-closure of g_n, or any polynomial in the $\mathrm{GL}(m_n, \mathbb{F})$-orbit closure of g_n, regardless of whether the $\mathrm{SL}(m_n, \mathbb{F})$-orbit of f_n is closed or not. Here we consider closedness of the orbits in the $\mathrm{GL}(m_n, \mathbb{F})$-orbit-closure of g_n with respect to the action of $\mathrm{SL}(m_n, \mathbb{F})$, not $\mathrm{GL}(m_n, \mathbb{F})$, since, as pointed out in Section 1.1, closedness with respect to the GL-action is not interesting. In other words, we consider the $\mathrm{SL}(m_n, \mathbb{F})$- as well as the $\mathrm{GL}(m_n, \mathbb{F})$-orbit-closure of g_n as an affine G-variety, with $G = \mathrm{SL}(m_n, \mathbb{F})$.

In the context of the VP vs. $\overline{\mathrm{VP}}$ problem, one has to consider the $\mathrm{GL}(m_n, \mathbb{F})$-orbit closure of g_n, since infinitesimally close approximation involves GL-transformations. The $\mathrm{GL}(m_n, \mathbb{F})$-orbit closures can be much harder than the $\mathrm{SL}(m_n, \mathbb{F})$-orbit-closures. For example, if g_n is the determinant, its $\mathrm{SL}(m_n, \mathbb{F})$-orbit is already closed [36], and hence, one really needs to understand its $\mathrm{GL}(m_n, \mathbb{F})$-orbit closure.

If a p-definable one parameter degeneration scheme, akin to the Hilbert-Mumford-Kempf criterion for stability, exists when f_n is allowed to be any polynomial in the $\mathrm{GL}(m_n, \mathbb{F})$-orbit-closure of g_n, $\{g_n\} \in$ VP, then it would follow that $\overline{\mathrm{VP}} \subseteq \mathrm{VP}^*$, and in conjunction with Theorem 1, that $\overline{\mathrm{VP}} \subseteq$ VNP. This is one plausible approach to show that $\overline{\mathrm{VP}} \subseteq$ VNP, if this is true.[9] If, on the other hand, no such p-definable one parameter degeneration scheme exists when f_n is allowed to be any polynomial in the $\mathrm{GL}(m_n, \mathbb{F})$-orbit-closure of g_n, $\{g_n\} \in$ VP, then it would be a strong indication that $\overline{\mathrm{VP}}$ is not contained in VP^*, and hence, also not in VP. This would open one possible route to formally separate $\overline{\mathrm{VP}}$ from VP.

All the evidence at hand does, in fact, suggest that such a general scheme may not exist for the following reasons. First, as explained in [38] in detail, the Hilbert-Mumford-Kempf criterion for stability is intimately related to, and in fact, goes hand in hand with another fundamental result in geometric invariant theory that, given any finite dimensional G-representation, or more generally, an affine G-variety X, $G = \mathrm{SL}(m, \mathbb{F})$, the closed G-orbits in X are in one-to-one correspondence with the points of the algebraic variety $X/G = \mathrm{spec}(\mathbb{F}[X]^G)$, called the categorical quotient.[10] By definition, this is the algebraic variety whose coordinate ring is $\mathbb{F}[X]^G$, the subring of G-invariants in the coordinate ring $\mathbb{F}[X]$ of X. But the set of all G-orbits in X does not, in general, have such a natural structure of an algebraic variety [38]. This is why the book [38] focuses on closed G-orbits in the construction of the various moduli spaces in algebraic geometry.

Second, from the complexity-theoretic perspective, the algebraic structure of the set of all G-orbits in X seems much harder, in general, than that of the set of closed G-orbits. For example, it is shown in [35] that, if X is a finite dimensional representation of G, then the set of closed G-orbits in X has a (quasi) explicit system of parametrization (by a small number of algebraic circuits of small size), assuming that (a) the categorical quotient X/G is explicit (as conjectured in [35] on the basis of the algorithmic results therein), and (b) the permanent is hard. In contrast, it may be conjectured that the set of all G-orbits in X does not, in general, have an explicit or even a small system of parametrization (by algebraic circuits with $+, -, *, /$, and equality-test

[9] In this case, separating VP from $\overline{\mathrm{VP}}$ would be stronger than separating VP from VNP.

[10] By spec here, we really mean, by abuse of notation, max-spec.

gates), since Noether's Normalization Lemma, which plays a crucial role in the parametrization of closed G-orbits, applies only to algebraic varieties. (The division and equality-test gates are needed here, since without them, the outputs of the circuits, being constant on all G-orbits, will be G-invariant polynomials that cannot distinguish a non-closed G-orbit from a G-orbit in its closure. By a general result in [42], all G-orbits in X can be parametrized, in principle, by a finite number of algebraic circuits of finite size over the coordinates of X, with $+, -, *, /$, and equality-test gates.) Formally, the conjecture is that there do not exist for every finite dimensional representation X of $G = \mathrm{SL}(m, \mathbb{F})$, $\mathsf{poly}(l, m)$ algebraic circuits of $\mathsf{poly}(l, m)$ size [11], $l = \dim(X)$, over the coordinates $x_1, \ldots, x_l$ of X, with constants in $\mathbb{F}$ and gates for $+, -, *, /$, and equality-test, such that the outputs of these circuits at the coordinates of any two points $v, w \in X$ are identical iff v and w are in the same G-orbit. (The gates for division and equality-test are not needed for parametrization of closed G-orbits in [35].)

A concrete case that illustrates well the difference between closed G-orbits and all G-orbits is when $X = M_m(\mathbb{F})^r$, the space of r-tuples of $m \times m$ matrices, with the conjugate (adjoint) action of $G = \mathrm{SL}(m, \mathbb{F})$. In this case it is known unconditionally that the set of closed G-orbits in X has a quasi-explicit (i.e., quasi-$\mathsf{poly}(m, r)$-time computable) parametrization when the characteristic p of $\mathbb{F}$ is not in $[2, \lfloor m/2 \rfloor]$; cf. [35] and [15] for charactetristic zero, and [35] for positive characteristic. In contrast, the best known parametrization [16] of all G-orbits in $M_m(\mathbb{F})^r$ (allowing division and equality-test gates in the algebraic circuits) has exponential complexity. The known algorithm [46] for constructing a canonical normal form of a matrix tuple in $M_m(\mathbb{F})^r$ with respect to the G-action also has exponential complexity, [12] (though the problem of deciding if two points in $M_m(\mathbb{F})^r$ are in the same G-orbit is in P [5, 9, 23]). The exponential complexity of parametrization of all G-orbits in $M_m(\mathbb{F})^r$ may be inherent, since the problem of clasifying all G-orbits in $M_m(\mathbb{F})^r$ is wild [11], when $r \geq 2$. Wildness [3, 13] is a universality property in representation theory, analogous to NP-completeness. The situation gets even wilder when X is a general G-representation or an affine G-variety. For example, it is known [2] that the problem of classifying all G-orbits in $\mathbb{F}^m \otimes \mathbb{F}^m \otimes \mathbb{F}^m$ contains, but is not contained in the wild problem of classifying all G-orbits in $M^m(\mathbb{F})^r$.

In view of such a fundamental difference between the algebraic structures of the set of closed G-orbits and the set of all G-orbits, from the mathematical as well as the complexity-theoretic perspectives, it may be conjectured that a p-definable one-parameter degeneration from $\{g_n\} \in \mathrm{VP}$ to $\{f_n\}$, with f_n in the $\mathrm{GL}(m_n, \mathbb{F})$- or $\mathrm{SL}(m_n, \mathbb{F})$-orbit closure of g_n, does not always exist if the G-orbit of f_n, with $G = \mathrm{SL}(m_n, \mathbb{F})$, is not required to be closed. If so, this would be a strong indication, as pointed out above, that $\overline{\mathrm{VP}}$ is not contained in VP^*, and hence, also not in VP.

The complexity-theoretic form of the Hilbert-Mumford-Kempf criterion proved in this article (Theorem 7) also provides an exponential (in n) upper bound on the degree of the canonical Kempf-one-parameter subgroup that drives g_n to f_n, with $\{g_n\} \in \mathrm{VP}$ and $\{f_n\} \in$ Stable-VP, where f_n is a stable degeneration of g_n. This canonical Kempf-one-parameter subgroup is known to be the fastest way to approach a closed orbit [27]. If one could prove a polynomial upper bound on this degree, then it would follow that Stable-VP = VP (cf. Lemma 4.3). On the other hand, if a worst-case super-polynomial lower bound on this degree can be proved, then it would be a strong indication that Stable-VP, and hence $\overline{\mathrm{VP}}$, are different from VP. This would open another possible route to formally separate $\overline{\mathrm{VP}}$ from VP.

[11] Here the size means the total number of nodes in the circuit. There is no restriction on the bit-lengths of the constants.

[12] This is because the algorithm in [46] needs factorization of univariate polynomials over extension fields of possibly exponential rank over the base field of definition of the input.

1.3 On the problem of explicit construction

Next we ask if one can construct an *explicit* family in Newton-$\mathrm{VP_{ws}}$ that can reasonably be conjectured to be not in $\mathrm{VP_{ws}}$ or even VP. With this mind, we first construct an explicit family $\{f_n\}$ of polynomials that can be approximated infinitesimally closely by symbolic determinants of size $\leq n$, but conjecturally cannot be computed exactly by symbolic determinants of $\Omega(n^{1+\delta})$ size, for a small enough positive constant $\delta < 1$; cf. Section 5. This construction follows a suggestion made in [36, Section 4.2]. The family $\{f_n\}$ is a Newton degeneration of the family of perfect matching Pfaffians of graphs. However, this family $\{f_n\}$ turns out to be in $\mathrm{VP_{ws}}$. So this idea needs to be extended much further to construct an explicit family in Newton-$\mathrm{VP_{ws}}$ that can be conjectured to be not in VP.

To see how this may be possible, note that the perfect matching Pfaffians are derived from a semi-invariant of the symmetric quiver with two vertices and one arrow. This suggests that to upgrade the conjectural $\Omega(n^{1+\delta})$ lower bound to obtain a candidate for a super-polynomial lower bound one could replace perfect matching Pfaffians by appropriate representation-theoretic invariants (but we do not have to confine ourselves to representation-theoretic invariants; cf. the remark at the end of Section 1.4). This leads to the second line of investigation, which we now discuss.

1.4 On Newton degeneration of generic semi-invariants

Our next result suggests that these invariants should be non-generic by showing that, for many finite quivers, including some wild ones, Newton degeneration of any generic semi-invariant can be computed by a symbolic determinant of polynomial size.

A quiver $Q = (Q_0, Q_1)$ [10, 12] is a directed graph (allowing multiple edges) with the set of vertices Q_0 and the set of arrows Q_1. A linear representation V of a quiver associates to each vertex $x \in Q_0$ a vector space V^x, and to each arrow $\alpha \in Q_1$ a linear map V^α from $V^{s\alpha}$ to $V^{t\alpha}$, where $s\alpha$ denotes the start (tail) of α and $t\alpha$ its target (head). The dimension vector of V is the tuple of non-negative integers that associates $\dim(V^x)$ to each vertex $x \in Q_0$. Given a dimension vector $\beta \in \mathbb{N}^{|Q_0|}$, let $\mathrm{Rep}(Q, \beta)$ denote the space of all representations of Q with the dimension vector β. We have the natural action of $\mathrm{SL}(\beta) := \prod_{x \in Q_0} \mathrm{SL}(\beta(x), \mathbb{F})$ on $\mathrm{Rep}(Q, \beta)$ by change of basis. Let $\mathrm{SI}(Q, \beta) = \mathrm{Rep}(Q, \beta)^{\mathrm{SL}(\beta)}$ denote the ring of semi-invariants. The *generic* semi-invariants in this ring (see [10]) will be recalled in Section 6.

We will be specifically interested in the following well-known types of quivers, cf. [11]. The m-Kronecker quiver is the quiver with two vertices and m arrows between the two vertices with the same direction. It is wild if $m \geq 3$. The k-subspace quiver is the quiver with $k+1$ vertices $\{x_1, \ldots, x_k, y\}$ and k arrows $(x_1, y), \ldots, (x_k, y)$. It is wild if $k \geq 5$. The A-D-E Dynkin quivers (see Section 6.4) are the only quivers of finite representation type—this means they have only finitely many indecomposable representations.

The following result tells us where *not* to look for explicit candidate families in $\overline{\mathrm{VP}} \setminus \mathrm{VP}$.

Theorem 2. *Let Q be an m-Kronecker quiver, or a k-subspace quiver, or an A-D-E Dynkin quiver. Then any Newton degeneration of a generic semi-invariant of Q with dimension vector β and degree d can be computed by a weakly skew circuit (or equivalently a symbolic determinant) of* $\mathsf{poly}(|\beta|, d)$ *size, where $|\beta| = \sum_{x \in Q_0} \beta(x)$.*

The proof strategy for Theorem 2 is as follows. Define the coefficient complexity $\mathrm{coeff}(E)$ of a set E of integral linear equalities in $\mathbb{R}^m$ as the sum of the absolute values of the coefficients of the equalities. Define the coefficient complexity of a face of a polytope in $\mathbb{R}^m$ as the minimum of coeff(E), where E ranges over all integral linear equality sets that define the face, in conjunction with the description of the polytope; cf. Section 6.1.

Theorem 2 is proved by showing that the coefficient complexity of every face of the Newton polytope of a generic semi-invariant of any quiver as above is polynomial in $|\beta|$ and d, though the number of vertices on a face can be exponential.

In view of this result and its proof, to construct an explicit family in Newton-$\mathrm{VP}_{\mathrm{ws}} \setminus \mathrm{VP}_{\mathrm{ws}}$, we should look for appropriate *non-generic* invariants of representations of finitely generated algebras whose Newton polytopes have faces with *super-polynomial coefficient complexity* and *super-polynomial number of vertices*.[13]

Finally, we emphasize that we do not have to confine ourselves to Newton-VP in the search of a specific candidate family in $\overline{\mathrm{VP}} \setminus \mathrm{VP}$. We may search within VP^*, or even outside VP^*. Indeed, it may be easier to identify specific candidate families in $\overline{\mathrm{VP}} \setminus \mathrm{VP}$ outside VP^* than inside VP^*.

1.5 Organization.

The rest of this article is organized as follows. In Section 2 we cover the preliminaries. In Section 3, we formally define the three degenerations of VP and VNP. In Section 4, we prove Theorem 1. In Section 5 we construct an explicit family $\{f_n\}$ that can be approximated infinitesimally closely by symbolic determinants of size $\leq n$, but conjecturally cannot be computed exactly by symbolic determinants of $\Omega(n^{1+\delta})$ size, for a small enough positive constant $\delta < 1$. In Section 6, we prove Theorem 2. In Section 7, we give additional examples of representation-theoretic symbolic determinants whose Newton degenerations have small circuits. All these examples suggest that explicit families in Newton-$\mathrm{VP}_{\mathrm{ws}} \setminus \mathrm{VP}_{\mathrm{ws}}$ would have to be rather delicate.

2 Preliminaries

For $n \in \mathbb{N}$, let $[n] := \{1, \ldots, n\}$. We denote by $\mathbf{x} = (x_1, \ldots, x_n)$ a tuple of variables; $\mathbf{x}$ may also denote $\{x_1, \ldots, x_n\}$. Let $\mathbf{e} = (e_1, \ldots, e_n)$ be a tuple of nonnegative integers. We usually use $\mathbf{e}$ as the exponent vector of a monomial in $\mathbb{F}[x_1, \ldots, x_n]$. Thus, $\mathbf{x}^{\mathbf{e}}$ denotes the monomial with the exponent vector $\mathbf{e}$. Let $|\mathbf{e}| := \sum_{i=1}^{n} e_i$.

For a field $\mathbb{F}$, $\mathrm{char}(\mathbb{F})$ denotes the characteristic of $\mathbb{F}$. Throughout this paper, we assume that $\mathbb{F}$ is algebraically closed. S_n denotes the symmetric group consisting of permutations of n objects.

We say that a polynomial $g = g(x_1, \ldots, x_n)$ is a *linear projection* of $f = f(y_1, \ldots, y_m)$ if g can be obtained from f by letting y_j's be some (possibly non-homogeneous) linear combinations of x_i's with coefficients in the base field $\mathbb{F}$.

A family of polynomials $\{f_n\}_{n \in \mathbb{N}}$ is p-bounded if f_n is a polynomial in $\mathsf{poly}(n)$ variables of $\mathsf{poly}(n)$ degree. The class VP [49] consists of p-bounded polynomial families $\{f_n\}_{n \in \mathbb{N}}$ over $\mathbb{F}$ such that f_n can be computed by an arithmetic circuit over $\mathbb{F}$ of $\mathsf{poly}(n)$ size.

Convention: We call a class $\mathcal{C}$ of families of polynomials *standard* if it contains only p-bounded families, and is closed under linear projections.

By a *symbolic determinant* of size m over the variables $x_1, \ldots, x_n$, we mean the determinant of an $m \times m$ matrix, whose each entry is a possibly non-homogeneous linear function of $x_1, \ldots, x_n$ with coefficients in the base field $\mathbb{F}$. The class $\mathrm{VP}_{\mathrm{ws}}$ is the class of families of polynomials that can be computed by weakly skew circuits of polynomial size, or equivalently, by symbolic determinants of polynomial size [33].

[13] Super-polynomial coefficient complexity and super-polynomial number of vertices do not ensure high circuit complexity of Newton degeneration. These are necessary conditions that should only be taken as guiding signs.

The class VNP is the class of p-definable families of polynomials [49], that is, those families (f_n) such that f_n has $\mathsf{poly}(n)$ variables and $\mathsf{poly}(n)$ degree, and there exists a family $(g_n(x,y)) \in \mathrm{VP}$ such that $f_n(x) = \sum_{e\in\{0,1\}^{\mathsf{poly}(n)}} g_n(x,e)$.

The class $\overline{\mathrm{VP}}$ is defined as follows [8,36]. Over $\mathbb{F}=\mathbb{C}$, we say that a polynomial family $\{f_n\}_{n\in\mathbb{N}}$ is in $\overline{\mathrm{VP}}$, if there exists a family of sequences of polynomials $\{f_n^{(i)}\}_{n\in\mathbb{N}}$ in VP, $i=1,2,\ldots$, such that for every n, the sequence of polynomials $f_n^{(i)}$, $i=1,2,\ldots$, goes infinitesimally close to f_n, in the usual complex topology. Here, polynomials are viewed as points in the linear space of polynomials. There is a more general definition that works over arbitrary algebraically closed fields—including in positive characteristic—using the Zariski topology. For a direct treatment, see, e.g. [7, App. 20.6]. The operational version of this definition we use is as follows: $\{f_n(x_1,\ldots,x_m)\} \in \overline{\mathrm{VP}}$ if there exist polynomials $f_{n,t}(x_1,\ldots,x_m) \in \mathrm{VP}_{\mathbb{C}((t))}$—$f_{n,t}$ is a polynomial in the x_i whose coefficients are Laurent series in t—such that $f_n(x)$ is the coefficient of the term in $f_{n,t}(x)$ of lowest degree in t.

The classes $\overline{\mathrm{VP}_{\mathrm{ws}}}$, $\overline{\mathrm{VNP}}$, and $\overline{\mathcal{C}}$, for any standard class $\mathcal{C}$, are defined similarly.

By the *determinantal complexity* $\mathrm{dc}(f)$ of a polynomial $f(x_1,\ldots,x_n)$, we mean the smallest integer m such that f can be expressed as a symbolic determinant of size m over $x_1,\ldots,x_n$. By the *approximative determinantal complexity* $\overline{\mathrm{dc}}(f)$, we mean the smallest integer m such that f can be approximated infinitesimally closely by symbolic determinants of size m.

Thus the $\mathrm{VP}_{\mathrm{ws}} \neq \mathrm{VNP}$ conjecture in Valiant [49] is equivalent to saying that $\mathrm{dc}(\mathrm{perm}_n)$ is not $\mathsf{poly}(n)$, where perm_n denotes the permanent of an $n\times n$ variable matrix. The $\mathrm{VNP} \not\subseteq \overline{\mathrm{VP}_{\mathrm{ws}}}$ conjecture in [36] is equivalent to saying that $\overline{\mathrm{dc}}(\mathrm{perm}_n)$ is not $\mathsf{poly}(n)$.

A priori, it is not at all obvious that dc and $\overline{\mathrm{dc}}$ are different complexity measures. The following two examples should make this clear.

Example 2.1 (Example 9 in [30]). Let $f = x_1^3 + x_2^2x_3 + x_2x_4^2$. Then $\mathrm{dc}(f) \geq 5$, but $\overline{\mathrm{dc}}(f) = 3$.

Example 2.2 (Proposition 3.5.1 in [31]). Let n be odd. Given an $n\times n$ complex matrix M, let M_{ss} and M_s denote its skew-symmetric and symmetric parts. Since n is odd, $\det(M_{ss})$=0. Hence, for a variable t, $\det(M_{ss}+tM_s) = tf(M)+O(t^2)$, for some polynomial function $f(M)$. Clearly, $\overline{\mathrm{dc}}(f) = n$, since $\det(M_{ss}+tM_s)/t$ goes infinitesimally close to $f(M)$ when t goes to 0. But $\mathrm{dc}(f) > n$.

The $\mathrm{VP}_{\mathrm{ws}} \neq \overline{\mathrm{VP}_{\mathrm{ws}}}$ conjecture in [35] is equivalent to saying that there exists a polynomial family $\{f_n\}$ such that $\overline{\mathrm{dc}}(f_n) = \mathsf{poly}(n)$, but $\mathrm{dc}(f_n)$ is not $\mathsf{poly}(n)$. Instead of this conjecture, we will focus on the $\mathrm{VP} \neq \overline{\mathrm{VP}}$ conjecture in [35], since the considerations for the former conjecture are entirely similar.

A (convex—we will only consider convex ones here) *polytope* is the convex hull in $\mathbb{R}^n$ of a finite set of points. A *face* of a polytope P is the intersection of P with linear halfspace $H = \{v\in\mathbb{R}^n | \ell(v) \geq c\}$ for some linear function ℓ and constant c such that H contains no points of the (topological) interior of P. Equivalently, a polytope is the intersection of finitely many half-spaces, a half-space $H_{\ell,c} = \{v | \ell(v) \geq c\}$ is tight for P if $P \subseteq H_{\ell,c}$ and $P \not\subseteq H_{\ell,c'}$ for any $c' > c$, and a face of P is the intersection of P with a half-space of the form $H_{-\ell,-c}$ where $H_{\ell,c}$ is tight for P.

3 Degenerations of VP and VNP

To understand the relationship between VP, VNP, and their closures $\overline{\mathrm{VP}}$ and $\overline{\mathrm{VNP}}$, we now introduce three degenerations of VP and VNP. The considerations for $\mathrm{VP}_{\mathrm{ws}}$ and $\overline{\mathrm{VP}_{\mathrm{ws}}}$ are entirely similar.

3.1 Stable degeneration

First we define stable degenerations of VP and VNP.

Consider the natural action of $G = \mathrm{SL}(n, \mathbb{F})$ on $\mathbb{F}[\mathbf{x}] = \mathbb{F}[x_1, \ldots, x_n]$ that maps $f(\mathbf{x})$ to $f(\sigma^{-1}\mathbf{x})$ for any $\sigma \in G$. Following Mumford et al. [38], call $f = f(\mathbf{x}) \in \mathbb{F}[\mathbf{x}]$ *stable* (with respect to the G-action) if the G-orbit of f is Zariski-closed. It is known [38] that the closure of the G-orbit of any $g \in \mathbb{F}[\mathbf{x}]$ contains a unique closed G-orbit. We say that f is a *stable degeneration* of g if f lies in the unique closed G-orbit in the G-orbit-closure of g. (If the G-orbit of g is already closed then this just means that f lies in the G-orbit of g.)

We now define the class Stable-$\mathcal{C}$, the stable degeneration of any standard class $\mathcal{C}$, as follows. We say that $\{f_n\}_{n\in\mathbb{N}}$ is in Stable-$\mathcal{C}$ if (1) $\{f_n\} \in \mathcal{C}$, or (2) there exists $\{g_n\}_{n\in\mathbb{N}}$ in $\mathcal{C}$ such that each f_n is a stable degeneration of g_n with respect to the action of $G = \mathrm{SL}(m_n, \mathbb{F})$, where $m_n = \mathsf{poly}(n)$ denotes the number of variables in f_n and g_n.

Proposition 3.1. *For any standard class $\mathcal{C}$ (cf. Section 2), Stable-$\mathcal{C} \subseteq \overline{\mathcal{C}}$. In particular, Stable-VP $\subseteq \overline{VP}$ and Stable-VNP $\subseteq \overline{VNP}$.*

Proof. Suppose $\{f_n(x_1, \ldots, x_{m_n})\}$ is in Stable-$\mathcal{C}$. This means there exists a family $\{g_n(x_1, \ldots, x_{m_n})\}$ in $\mathcal{C}$ such that, for each n, f_n is in the $\mathrm{SL}(m_n, \mathbb{F})$-orbit closure of g_n. This means f_n can be approximated infinitesimally closely by polynomials in $\mathcal{C}$, hence $\{f_n\}$ is in $\overline{\mathcal{C}}$. □

Proposition 3.2. *Stable-$\overline{\mathcal{C}} = \overline{\mathcal{C}}$, in particular Stable-$\overline{VP} = \overline{VP}$, and Stable-$\overline{VNP} = \overline{VNP}$.*

This is a direct consequence of the definitions.

3.2 Newton degeneration

Next we define Newton degenerations of VP and VNP.

Given a polynomial $f \in \mathbb{F}[x_1, \ldots, x_n]$, suppose $f = \sum_{\mathbf{e}} \alpha_{\mathbf{e}} \mathbf{x}^{\mathbf{e}}$. We collect the exponent vectors of f and form the convex hull of these exponent vectors in $\mathbb{R}^n$. The resulting polytope is called the *Newton polytope* of f, denoted $\mathrm{NPT}(f)$. Given an arbitrary face Q of $\mathrm{NPT}(f)$, the Newton degeneration of f to Q, denoted $f|_Q$, is the polynomial $\sum_{\mathbf{e}\in Q} \alpha_{\mathbf{e}} \mathbf{x}^{\mathbf{e}}$.

We now define the class Newton-$\mathcal{C}$, the Newton degeneration of any class $\mathcal{C}$, as follows: $\{f_n\}_{n\in\mathbb{N}}$ is in Newton-$\mathcal{C}$, if there exists $\{g_n\}_{n\in\mathbb{N}}$ in $\mathcal{C}$ such that each f_n is the Newton degeneration of g_n to some face of $\mathrm{NPT}(g_n)$, or a linear projection of such a Newton degeneration.

Theorem 3. *Let $\mathcal{C}$ be any standard class (cf. Section 2). Then Newton-$\mathcal{C} \subseteq \overline{\mathcal{C}}$. In particular, Newton-VP $\subseteq \overline{VP}$ and Newton-VNP $\subseteq \overline{VNP}$.*

Proof. Let $\{f_n\}_{n\in\mathbb{N}}$ be in Newton-$\mathcal{C}$, and suppose $f_n \in \mathbb{F}[x_1, \ldots, x_{m(n)}]$. Then there exists $\{g_n\}_{n\in\mathbb{N}} \in \mathcal{C}$, such that $g_n \in \mathbb{F}[x_1, \ldots, x_m]$, $m = m(n)$, and $f_n = g_n|_Q$, where Q is a face of $\mathrm{NPT}(g_n)$. Suppose the supporting hyperplane of Q is defined by $\langle \mathbf{a}, \mathbf{x} \rangle = b$, where $\mathbf{a} = (a_1, \ldots, a_m)$. If necessary, by replacing $(\mathbf{a}, b)$ with $(-\mathbf{a}, -b)$, we make sure that for an arbitrary exponent vector $\mathbf{e}$ in g_n, $\langle \mathbf{a}, \mathbf{e} \rangle \geq b$. That is, among all exponent vectors, exponent vectors on Q achieve the minimum value b in the direction $\mathbf{a}$.

Now introduce a new variable t, and replace x_i with $t^{a_i} x_i$ to obtain a polynomial $g_n'(x_1, \ldots, x_m, t) = g_n(t^{a_1} x_1, \ldots, t^{a_m} x_m) \in \mathbb{F}[x_1, \ldots, x_m, t]$. By the definition of f_n, $g_n' = t^b \cdot f_n +$higher order terms in t. Therefore, $\{f_n\} \in \overline{\mathcal{C}}$. □

Remark. In the above proof, it is important that the Newton degeneration of g_n is the coefficient of t^b, the lowest order term in t, and it is not at all clear how one could possibly access higher order terms in t using any kind of degeneration. Note that higher order terms can be VNP-complete: Form a matrix of variables $(x_{i,j})_{i,j\in[n]}$, and consider the polynomial $\prod_{j\in[n]}(t^{n+1}x_{1,j} + t^{(n+1)^2}x_{2,j} + \cdots + t^{(n+1)^i}x_{i,j} + \cdots + t^{(n+1)^n}x_{n,j})$. The coefficient of $t^{(n+1)+(n+1)^2+\cdots+(n+1)^n}$ is then the permanent of $(x_{i,j})_{i,j\in[n]}$. (Essentially the same construction appeared as [6, Prop. 5.3].) The seeming impossibility of extracting higher-order terms in t is in line with the expectation that VNP $\not\subseteq \overline{\text{VP}}$.

Noting that if $\mathcal{C}$ is closed under linear projections, then so is $\overline{\mathcal{C}}$, we have:

Corollary 3.3. *For any standard class $\mathcal{C}$, Newton-$\overline{\mathcal{C}} = \overline{\mathcal{C}}$. In particular, Newton-$\overline{VP} = \overline{VP}$ and Newton-$\overline{VNP} = \overline{VNP}$.*

3.3 P-definable one-parameter degeneration

Finally, we define p-definable one-parameter degenerations of VP and VNP. We say that a family $\{f_n(x_1, \ldots, x_{m_n})\}$, $m_n = \mathsf{poly}(n)$, is a *one-parameter degeneration* of $\{g_n(y_1, \ldots, y_{l_n})\}$, $l_n = \mathsf{poly}(n)$, of *exponential degree*, if, for some positive integral function $K(n) = O(2^{\mathsf{poly}(n)})$, there exist $c_n(i, j, k) \in \mathbb{F}$, $1 \le i \le l_n$, $0 \le j \le m_n$, $-K(n) \le k \le K(n)$, such that $f_n = \lim_{t\to 0} g_n(t)$, where $g_n(t)$ is obtained from g_n by substitutions of the form

$$y_i = a_0^i + \sum_{j=1}^{m_n} a_j^i x_j, \quad 1 \le i \le l_n, \text{ where } a_j^i = \sum_{k=-K(n)}^{K(n)} c_n(i, j, k)t^k, \quad 1 \le i \le l_n,\ 0 \le j \le m_n.$$

Note that by [6], $\overline{\text{VP}}$ consists exactly of those one-parameter degenerations of VP of exponential degree.

We say that the family $\{f_n(x_1, \ldots, x_{m_n})\}$, $m_n = \mathsf{poly}(n)$, is a *one-parameter degeneration* of $\{g_n(y_1, \ldots, y_{l_n})\}$, $l_n = \mathsf{poly}(n)$, of *polynomial degree* if $K(n)$ above is $O(\mathsf{poly}(n))$ (instead of $O(2^{\mathsf{poly}(n)})$).

We say that a family $\{f_n(x_1, \ldots, x_{m_n})\}$, $m_n = \mathsf{poly}(n)$, is a *p-definable one-parameter degeneration* of $\{g_n(y_1, \ldots, y_{l_n})\}$, $l_n = \mathsf{poly}(n)$, if, for some $K(n) = O(2^{\mathsf{poly}(n)})$, there exists a $\mathsf{poly}(n)$-size circuit family $\{C_n\}$ over $\mathbb{F}$ such that $f_n = \lim_{t\to 0} g_n(t)$, where $g_n(t)$ is obtained from g_n by substitutions of the form

$$y_i = a_0^i + \sum_{j=1}^{m_n} a_j^i x_j, \quad 1 \le i \le l_n, \text{ where } a_j^i = \sum_{k=-K(n)}^{K(n)} C_n(i, j, k)t^k, \quad 1 \le i \le l_n,\ 0 \le j \le m_n.$$

Here it is assumed that the circuit C_n takes as input $\lceil \log_2 l_n \rceil + \lceil \log_2 m_n \rceil + \lceil \log_2(K(n) + 1) \rceil$ many 0-1 variables, which are intended to encode three integers (i, j, k) satisfying $1 \le i \le l = l_n$, $0 \le j \le m = m_n$, and $|k| \le K(n)$, treating 0 and 1 as elements of $\mathbb{F}$.

Thus a p-definable one-parameter degeneration is a one-parameter degeneration of exponential degree that can be specified by a circuit of polynomial size.

Remark. We can generalize the notion of a one-parameter degeneration slightly by allowing C_n an additional input $b \in \{0, 1\}^{a(n)}$, $a(n) = \mathsf{poly}(n)$, and letting

$$a_j^i = \sum_{k=-K(n)}^{K(n)} \left(\sum_{b\in\{0,1\}^{a(n)}} C_n(i, j, k, b) \right) t^k, \quad 1 \le i \le l_n,\ 0 \le j \le m_n.$$

The following results hold for this more general notion also.

For any class $\mathcal{C}$ we now define $\mathcal{C}^*$, called the p-definable one-parameter degeneration of $\mathcal{C}$, as follows. We say that $\{f_n\} \in \mathcal{C}^*$ if there exists $\{g_n\} \in \mathcal{C}$ such that $\{f_n\}$ is a p-definable one-parameter degeneration of $\{g_n\}$.

Lemma 3.4. *For any standard class $\mathcal{C}$ (cf. Section 2), Newton-$\mathcal{C} \subseteq \mathcal{C}^*$. In particular, Newton-VP $\subseteq$ VP* and Newton-VNP $\subseteq$ VNP*.*

This follows from the proof of Theorem 3, noting that we may always take the coefficients of a face to have size at most $2^{\mathsf{poly}(n)}$. The following are easy consequences of the definitions:

Proposition 3.5. *VP* $\subseteq \overline{VP}$, and VNP* $\subseteq \overline{VNP}$.*

Proof. This is immediate from the definitions. For the first statement, note that, for any g_n with a small circuit and any $a \in \mathbb{F}$, $g_n(a)$, which is obtained from $g_n(t)$ (cf. Section 3.3) by setting $t = a$, also has a small circuit. The situation for the second statement is similar. □

Proposition 3.6. *$\overline{VP}^* = \overline{VP}$, and $\overline{VNP}^* = \overline{VNP}$.*

This is an immediate consequence of the definitions.

4 Stable-VNP = Newton-VNP = VNP* = VNP

We now prove Theorem 1, by a circular sequence of inclusions.

Proof of Theorem 1. Since VNP $\subseteq$ Stable-VNP by definition, Theorem 1 (a) follows from the facts that Stable-VNP $\subseteq$ Newton-VNP (cf. Theorem 4 below), Newton-VNP $\subseteq$ VNP* (Lemma 3.4), and VNP* $\subseteq$ VNP (cf. Theorem 6 below).

Theorem 1 (b) follows from the facts that Stable-VP $\subseteq$ Newton-VP (cf. Theorem 4 below), Newton-VP $\subseteq$ VP* (Lemma 3.4), and VP* $\subseteq$ VNP (cf. Corollary 4.2 below). □

Theorem 4. *For any class $\mathcal{C}$ of families of p-bounded polynomials, Stable-$\mathcal{C}$ $\subseteq$ Newton-$\mathcal{C}$. In particular, Stable-VP $\subseteq$ Newton-VP and Stable-VNP $\subseteq$ Newton-VNP.*

Proof. Suppose $\{f_n\} \in$ Stable-$\mathcal{C}$. If $\{f_n\} \in \mathcal{C}$ then there is nothing to show. Otherwise, there exists $\{g_n\}_{n\in\mathbb{N}}$ in $\mathcal{C}$ such that each f_n is a stable degeneration of g_n with respect to the action of $G = \mathrm{SL}(m_n, \mathbb{F})$, where m_n denotes the number of variables in f_n and g_n.

It suffices to show that $f = f_n(x_1, \ldots, x_m)$, $m = m_n$, is a Newton degeneration of $g = g_n(x_1, \ldots, x_m)$. Let $\mathbf{x} = (x_1, \ldots, x_m)$.

By the Hilbert–Mumford–Kempf criterion for stability [26], there exists a one-parameter subgroup $\lambda(t) \subseteq G$ such that $\lim_{t\to 0} \lambda(t).g = f$. Let T be the canonical maximal torus in G such that the monomials in x_i's are eigenvectors for the action of T. After a linear change of coordinates (which is allowed since Newton-$\mathcal{C}$ is closed under linear transformations by definition), we can assume that $\lambda(t)$ is contained in T. Thus $\lambda(t) = \mathrm{diag}(t^{k_1}, \ldots, t^{k_m})$ (the diagonal matrix with t^{k_j}'s on the diagonal), $k_j \in \mathbb{Z}$, such that $\sum k_j = 1$.

It follows that f is the Newton degeneration of g to the face of $\mathrm{NPT}(g)$ where the linear function $\sum_j k_j x_j$ achieves the minimum value (which has to be zero). □

The following result is subsumed by Theorem 6; we include its proof here as a warm-up for expository clarity.

Theorem 5. *Newton-VNP $\subseteq$ VNP.*

Proof. Suppose $\{f_n\} \in$ Newton-VNP. If $\{f_n\} \in$ VNP, then there is nothing to show. Otherwise, there exists $\{g_n\}_{n\in\mathbb{N}}$ in VNP such that each f_n is the Newton degeneration of g_n to some face of NPT(g_n), or a linear projection of such a Newton degeneration. Since VNP is closed under linear projections, we can assume, without loss of generality, that f_n is the Newton degeneration of g_n to some face of NPT(g_n).

By Valiant [49], we can assume that $g = g_n(x_1, \ldots, x_m)$, $m = m_n = \mathsf{poly}(n)$, is a projection of perm(X),[14] where X is a $k \times k$ variable matrix, with $k = \mathsf{poly}(n)$. This means $g = \mathrm{perm}(X')$, where each entry of X' is some variable x_i or a constant from the base field $\mathbb{F}$. Since $f = f_n$ is a Newton degeneration of g, it follows that there is some substitution, as in the proof of Theorem 3, $x_j \to x_j t^{k_j}$, $k_j \in \mathbb{Z}$, such that $f = \lim_{t\to 0} \mathrm{perm}(X'(t))$, where $X'(t)$ denotes the matrix obtained from X' after this substitution.

It is easy to ensure that $|k_j| \le O(2^{\mathsf{poly}(n)})$. Then, given any permutation $\sigma \in \mathrm{S}_k$, whether the corresponding monomial $\prod_i X'_{i\sigma(i)}$ contributes to f can be decided in $\mathsf{poly}(n)$ time. It follows that the coefficient of a monomial can be computed by an algebraic circuit summed over polynomially many Boolean inputs (convert the implicit $\mathsf{poly}(n)$-time Turing machine into a Boolean circuit, then convert it into an algebraic circuit (as in [49, Remark 1]) that incorporates the constants appearing in the projection). Hence $\{f_n\} \in$ VNP. $\square$

Since VP $\subseteq$ VNP, the preceding result implies:

Corollary 4.1. *Newton-VP $\subseteq$ VNP.*

The following result can proved similarly to Theorem 5.

Theorem 6. *VNP* $\subseteq$ VNP.*

Proof. Suppose $\{f_n(x_1, \ldots, x_{m_n})\} \in \mathrm{VNP}^*$. Then there exists $\{g_n(y_1, \ldots, y_{l_n})\}$ in VNP such that each f_n is a p-definable one-parameter degeneration of g_n.

By Valiant [49], we can assume that $g = g_n(y_1, \ldots, y_l)$, $l = l_n = \mathsf{poly}(n)$, is a projection of perm(Y),[14] where Y is a $k \times k$ variable matrix, with $k = \mathsf{poly}(n)$. This means $g = \mathrm{perm}(Y')$, where each entry of Y' is some variable y_i or a constant from the base field F.

Since $f = f_n(x_1, \ldots, x_m)$, $m_n = \mathsf{poly}(n)$, is a p-definable one-parameter degeneration of g, for some $K(n) = O(2^{\mathsf{poly}(n)})$, there exists a $\mathsf{poly}(n)$-size circuit family $\{C_n\}$ over $\mathbb{F}$ such that $f_n = \lim_{t\to 0} g_n(t)$, where $g_n(t)$ is obtained from g_n by substitutions of the form

$$y_i = a_0^i + \sum_{j=1}^{m} a_j^i x_j, \quad 1 \le i \le l,$$

where

$$a_j^i = \sum_{k=-K(n)}^{K(n)} C_n(i, j, k) t^k, \quad 1 \le i \le l,\ 0 \le j \le m.$$

Let $Y'(t)$ be the matrix obtained from Y' after the substitution above. Given any permutation $\sigma \in S_k$, and any nonnegative integer sequence $\mu = (\mu_0, \ldots, \mu_m)$, the coefficient of the monomial $\mathbf{x}^\mu := \prod_i x_i^{\mu_i}$ in $\prod_i Y'(t)_{i,\sigma(i)}$ is a Laurent polynomial in t. Let c_μ^σ denote the coefficient of t^0 in this Laurent polynomial. It can be shown that, for some $\mathsf{poly}(n)$-size circuit D_n over $\mathbb{F}$ (depending on C_n) with $s_n = \mathsf{poly}(n)$ inputs, we can express c_μ^σ as

[14] To get the proof to work in characteristic 2 as well, simply use the Hamilton cycle polynomial $HC(X) = \sum_{\text{k-cycles } \sigma\in S_k} \prod_{i\in[k]} x_{i,\sigma(i)}$ instead, which is VNP-complete in any characteristic [49].

$$c_\mu^\sigma = \sum_{b \in \{0,1\}^{s_n}} D_n(b, \sigma, \mu).$$

Here it is assumed that D_n takes (b, σ, μ), specified in binary, as input, with 0 and 1 regarded as elements of $\mathbb{F}$. The idea is that σ specifies which of the $k!$ terms of the permanent is chosen, μ specifies which monomial is chosen, and the Boolean vector b is used to specify which summand of y_i is chosen in the summation. For a given choice of summand of y_i, computing the contribution to the corresponding coefficient is easy using C_n, and then we get to sum over all possible choices $b \in \{0,1\}^{s_n}$.

It follows that the coefficient of $\mathbf{x}^\mu$ in f_n is

$$\sum_\sigma c_\mu^\sigma = \sum_\sigma \sum_{b \in \{0,1\}^{s_n}} D_n(b, \sigma, \mu).$$

Using this fact, in conjunction with Valiant [49], it can be shown that $\{f_n\} \in$ VNP. □

Since VP $\subseteq$ VNP, the preceding result implies:

Corollary 4.2. *VP* $\subseteq$ VNP.*

In contrast, using the interpolation technique of Strassen [47] and Bini [4] we have:

Lemma 4.3 (cf. also [6], [8, §9.4], [18, Prop. 3.5.4])**.** *If $\{f_n\}$ is a one-parameter degeneration of $\{g_n\} \in$ VP of polynomial degree, then $\{f_n\} \in$ VP.*

Theorem 1 leads to:

Question 4.4. (1) Is VP = Stable-VP?
(2) Is VP = Newton-VP?
(3) Is VP = VP*?
(4) Is VP* = $\overline{\text{VP}}$?

4.1 A complexity-theoretic form of the Hilbert–Mumford–Kempf criterion

As a byproduct of the proof of Theorem 1, we get the following complexity-theoretic form of the Hilbert–Mumford–Kempf criterion [26] for stability with respect to the action of $G = \mathrm{SL}(m, \mathbb{F})$ on $\mathbb{F}[x_1, \ldots, x_m]$. Given a one-parameter subgroup $\lambda(t) \subseteq G$, we can express it as $A \cdot \mathrm{diag}(t^{k_1}, \ldots, t^{k_m}) \cdot A^{-1}$, for some $A \in G$ and $k_j \in \mathbb{Z}$, $1 \le j \le m$. We call $\sum_i |k_i|$ the total degree of $\lambda(t)$. The following theorem is implicit in the proofs of Theorems 4 and 5.

Theorem 7. *Suppose $f = f(x_1, \ldots, x_m)$ belongs to the unique closed G-orbit in the G-orbit-closure of $g = g(x_1, \ldots, x_m) \in \mathbb{F}[x_1, \ldots, x_m]$. Then there exists a one-parameter subgroup $\lambda(t) \subseteq G$ such that (1) $\lim_{t \to 0} \lambda(t) \cdot g = f$, and (2) the total degree of λ is $O(exp(m, \langle \deg(g) \rangle))$, where $\langle \deg(g) \rangle$ denotes the bit-length of the degree of g.*

It follows that if $\{f_n\}$ is a stable degeneration of $\{g_n\} \in$ VP, then $\{f_n\}$ is a p-definable one-parameter degeneration of $\{g_n\}$.

This result can be generalized from $G = \mathrm{SL}$ to reductive algebraic groups using similar ideas, as follows. Let $\mathbb{F} = \mathbb{C}$. Let $V = V_\lambda(R)$ be a finite dimensional rational representation of a connected, reductive, algebraic group R with highest weight $\lambda = \sum_i d_i \omega_i$, where ω_i's denote the fundamental weights of the Lie algebra $\mathcal{R}$ of R. Let $d = \sum_i d_i$. Let $\langle d \rangle$ be its bit-length. Let $\mathrm{rank}(\mathcal{R})$ denote the rank of $\mathcal{R}$. Given any one-parameter subgroup $\lambda(t) \subseteq R$, let $\hat{\lambda} : \mathbb{C} \to \mathcal{R}$ denote the corresponding Lie algebra map. After conjugation, we can assume that $\hat{\lambda}(\mathbb{C})$ is contained in the Cartan subalgebra $\mathcal{H} \subseteq \mathcal{R}$. Fix the standard basis $\{h_i\}$ of $\mathcal{H}$ as in [17], and let $\hat{\lambda}(1) = \sum_i c_i h_i$. Define the total size of $\lambda(t)$ as $\sum_j |c_j|$.

Theorem 8. *Given $v \in V$ and w in the unique closed R-orbit in the R-orbit-closure of v, there exists a one-parameter subgroup $\lambda(t) \subseteq R$ of total size $O(exp(\mathrm{rank}(R), \langle d \rangle))$ such that $\lim_{t\to 0} \lambda(t) \cdot v = w$.*

We formally propose a question that has ramifications on the Stable-VP vs. VP question (cf. Section 1 and Question 4.4 (1)).

Question 4.5. For some positive constant a, does there exist a stable degeneration $\{f_n\}$ of some $\{g_n\} \in \mathrm{VP}$, with an $\Omega(2^{n^a})$ or a super-polynomial lower bound on the degree of the canonical Kempf-one-parameter subgroup [26] λ_n driving $\{g_n\}$ to $\{f_n\}$?

5 Newton degeneration of perfect matching Pfaffians

In this section, we construct an explicit family $\{f_n\}$ of polynomials such that f_n can be approximated infinitesimally closely by symbolic determinants of size n, but conjecturally requires size $\Omega(n^{1+\delta})$ to be computed by a symbolic determinant, for a small enough positive constant δ. However, the family $\{f_n\}$ turns out to be in $\mathrm{VP}_{\mathrm{ws}}$.

Suppose we have a simple undirected graph $G = (V, E)$ where $V = [n]$. Let $\{x_e \mid e \in E\}$ be a set of variables. The Tutte matrix of G is the $n \times n$ skew-symmetric matrix T_G such that, if $(i, j) = e \in E$, with $i < j$, then $T_G(i, j) = x_e$ and $T_G(j, i) = -x_e$; otherwise $T_G(i, j) = 0$. For a skew-symmetric matrix T, the determinant of T is a perfect square, and the square root of $\det(T)$ is called the Pfaffian of T, denoted $\mathrm{pf}(T)$. We call $\mathrm{pf}(T_G)$ the *perfect matching Pfaffian* of the graph G, and $\mathrm{pf}(T_G) = \sum_P \mathrm{sgn}(P) \prod_{e\in P} x_e$, where the sum is over all perfect matchings P of G, and $\mathrm{sgn}(P)$ takes ± 1 in a suitable manner. It is well-known that $\mathrm{pf}(T_G) \in \mathrm{VP}_{\mathrm{ws}}$.

Note that $\mathrm{NPT}(\mathrm{pf}(T_G))$ is the perfect matching polytope of G, which has the following description by Edmonds. For any $S \subseteq V$, we use $e \sim S$ to denote that e lies at the border of S. When $S = \{i\}$, we may write $e \sim i$ instead of $e \sim \{i\}$.

Theorem 9 (Edmonds, [14])**.** *The perfect matching polytope of a graph G is characterized by the following constraints:*

$$(a)\, \forall e \in E, x_e \geq 0; (b) \forall i \in V, \sum_{e\in E, e\sim i} x_e = 1; (c)\, \forall C \subseteq V, |C| > 1 \textit{ is odd}, \sum_{e \in E, e\sim C} x_e \geq 1. \qquad (1)$$

We shall refer to constraints of type (c) in Equation 1 as "odd-size constraints."

Theorem 10 (Kaltofen and Koiran, [24, Corollary 1])**.** *Given $f, g, h \in \mathbb{F}[\mathbf{x}]$, suppose $h = f/g$, and f and g are in VP_{ws}. Then $h \in VP_{ws}$.*

Theorem 11. *For any graph G and any face Q of $\mathrm{NPT}(\mathrm{pf}(T_G))$, $\mathrm{pf}(T_G)|_Q \in VP_{ws}$.*

Proof. Thanks to Edmonds' description, any face of $\mathrm{NPT}(\mathrm{pf}(T_G))$ is obtained by setting some of the inequalities in Equation 1 to equalities. As setting $x_e = 0$ amounts to consider some graph G' with e deleted from G, the bottleneck is to deal with the odd-size constraints.

Suppose the face Q is obtained via setting the odd-size constraints corresponding to $C_1, \ldots, C_s$ to equalities, where $C_i \subseteq V$. Note that $s = \mathsf{poly}(n)$, because the dimension of $\mathrm{NPT}(\mathrm{pf}(T_G))$ is polynomially bounded, thus any face can be obtained by setting polynomially many constraints to equalities. Let y be a new variable. For any edge $e \in E$, let the number of $i \in [s]$ such that e lies at the border of C_i be k_e. Then transform x_e to $x_e y^{k_e}$. Let the skew-symmetric matrix after the transformation be $\widetilde{T_G}$. Since each perfect matching touches the border of every C_i at least once, y^s divides $\mathrm{pf}(\widetilde{T_G})$, so $f := \frac{\mathrm{pf}(\widetilde{T_G})}{y^s}$ is a polynomial. Furthermore, the y-free terms in f corresponds to those perfect matchings that touch each border exactly once. Thus, setting y to zero in f gives $\mathrm{pf}(T_G)|_Q$.

f is in $\mathrm{VP_{ws}}$, because $\mathrm{pf}(\widetilde{T_G})$ and y^s are in $\mathrm{VP_{ws}}$, and use Theorem 10. □

Construction of an explicit family. Now we turn to the construction of an explicit family $\{f_n\}$ mentioned in the beginning of this section. We assume that the base field $\mathbb{F} = \mathbb{C}$.

First, we give a randomized procedure for constructing f_n:

1. Fix a small enough constant $a > 0$, and let l be the nearest odd integer to n^a. Fix odd-size disjoint subsets $C_1, \dots, C_k \subseteq [n]$, $k = \lfloor n^{1-a} \rfloor$, of size l. For example, we can let $C_1 = \{1, \dots, l\}$, $C_2 = \{l+1, \dots, 2l+1\}$, etc.

2. Choose a random regular non-bipartite graph G_n on n nodes with degree (say) $\sqrt{n}$.

3. Let Q be the face of $\mathrm{NPT}(\det(T_G))$ obtained by setting the odd-size constraints corresponding $C_1, \dots, C_k$ to equalities.

4. Let $f_n = \det(T_G)|_Q$.

Note that in the above we use determinant instead of Pfaffian, in order to simplify the discussion on the determinantal complexity of such polynomials. Then, f_n can be approximated infinitesimally closely by symbolic determinants of size n; cf. the proof of Theorem 3. By Theorem 11, f_n can be expressed as a symbolic determinant of size $\mathsf{poly}(n)$. But:

Conjecture 5.1. If $a > 0$ is small enough, then, with a high probability, f_n cannot be expressed as a symbolic determinant of size $\leq n^{1+\delta}$, for a small enough positive constant δ.

This says that the blow-up in the determinantal size in the proof of Theorem 11 due to the use of division (cf. Theorem 10) cannot be gotten rid of completely.

To get an explicit family $\{f_n\}$, we let G_n be a pseudo-random graph, instead of a random graph. This can be done in various ways; perhaps the most conservative way is based on the following result.

Lemma 5.2. *Fix a constant $b > 0$. Then, the problem of deciding, given G_n, whether f_n can be expressed as a symbolic determinant of size $\leq n^b$, belongs to AM.*

Proof. This essentially follows from Theorem 11 and the AM-algorithm for Hilbert's Nullstellensatz in Koiran [29], with one additional twist.

By Theorem 11. f_n has a small weakly skew circuit C_n. Consider a generic symbolic determinant $D(x, y)$ of size n^b whose entries are formal linear combinations of the x-variables, whose coefficients are new y variables. We want to know whether there is a setting α of the y variables that will make $D(x, \alpha) = C_n(x)$ (note that both the LHS and RHS here are given by small weakly skew circuits). The trick on top of Koiran's result is as follows.

Randomly guess a hitting set for the x variables—that is, a collection of $\mathsf{poly}(n)$ many integral values $\xi^{(i)}$ of large enough $\mathsf{poly}(n)$ magnitude that will be substituted into the x variables. By the fact that this is a hitting set (which it is with a high probability [21]), the following system of equations has a solution α for the y's iff there is a setting of the y's that makes $D(x,\alpha) = C_n(x)$ as polynomials in x:

$$\begin{aligned} D(\xi^{(1)}, y) &= C_n(\xi^{(1)}) \\ &\vdots \\ D(\xi^{(k)}, y) &= C_n(\xi^{(k)}) \end{aligned}$$

where k is the size of the hitting set. The AM algorithm is then: randomly guess the ξ's, then apply Koiran's AM algorithm to the preceding set of equations in only the y variables. Note that Koiran's result applies to equations given by circuits, and each of the preceding equations is given by a small weakly-skew circuit. $\square$

We can now derandomize the construction of G_n above using this result in conjunction with the derandomization procedure in [28] (based on Impagliazzo–Wigderson [22]), assuming that E does not have Boolean circuits, with an access to the SAT oracle, of subexponential size. This yields, for every n, a sequence $G_n^1, \ldots, G_n^l$, $l = \mathsf{poly}(n)$, of graphs and the corresponding sequence $f_n^1, \ldots, f_n^l$ of polynomials such that each f_n^i can be approximated infinitesimally closely by symbolic determinants of size n, but, assuming Conjecture 5.1 and the hardness hypothesis above, at least half of the f_n^i's cannot be expressed as symbolic determinants of size $\leq n^{1+\delta}$, for a small enough constant $\delta > 0$. This gives a two-parameter explicit family $\{f_n^i | 1 \leq i \leq l = \mathsf{poly}(n)\}$ such that each f_n^i can be approximated infinitesimally closely by symbolic determinants of size n, but, assuming Conjecture 5.1 and the hardness hypothesis above, the exact determinantal complexity of the family is $\Omega(n^{1+\delta})$, for a small enough constant $\delta > 0$.

A less conservative derandomization procedure is as follows. For each n, let F_n be a Ramanujan graph as in [32] on $n^{3/2}$ vertices. Partition the set of vertices of F_n in n groups $A_1, \ldots, A_n$, each of size $n^{1/2}$. Let G_n be a graph that contains one vertex labelled i for each A_i, $1 \leq i \leq n$. Join two distinct vertices i and j in G_n if there is an edge in F_n from any vertex in A_i to any vertex in A_j. Let f_n be defined as above with this G_n. Then each f_n can be approximated infinitesimally closely by symbolic determinants of size n. But it may be conjectured that f_n cannot be computed exactly by a symbolic determinant of $\Omega(n^{1+\delta})$ size, for a small enough positive constant δ.

6 Newton degenerations of generic semi-invariants of quivers

In this section we prove Theorem 2 for the generalized Kronecker quivers, k-subspace quivers, and the A-D-E Dynkin quivers. We assume familiarity with the basic notions of the representation theory of quivers; cf. [10, 12].

6.1 Newton degeneration to faces with small coefficient complexity

We begin by observing that the technique used to prove Theorem 11 can be generalized further. In the proof of Theorem 11, due to Edmonds' description of the perfect matching polytope, every face has a "small" description, by a set of linear equalities whose coefficients are polynomially bounded in magnitude.

For a face Q of a polytope P, we say that a set of linear equalities E characterizes Q with respect to P, if the description of P together with that of E characterizes Q. For E, let $\mathrm{coeff}(E)$ be the sum of the absolute values of the coefficients of the linear equalities in E. We define the *coefficient complexity* of Q as the minimum of $\mathrm{coeff}(E)$ over the integral linear equality sets E that characterize Q with respect to P. Adapting the proof of Theorem 11 we easily get the following:

Theorem 12. *Suppose $f \in \mathbb{F}[x_1, \ldots, x_n]$ can be computed by a (weakly skew) arithmetic circuit of size s. Let Q be a face of $\mathrm{NPT}(f)$ whose coefficient complexity is $\mathsf{poly}(n)$. Then $f|_Q$ can be computed by a (weakly skew) arithmetic circuit of size $\mathsf{poly}(s, n)$.*

Proof. Let $E = \{\ell_1, \ldots, \ell_m\}$ be the set of inequalities characterizing Q with respect to P, with $\mathrm{coeff}(E)$ polynomially bounded. Note that $m = |E|$ is polynomially bounded as well, since the Newton polytope of f lives in $\mathbb{R}^n$. Without loss of generality, we assume ℓ_i is in the form $a_{i,1}x_1 + \cdots + a_{i,n}x_n \geq a_{i,0}$. Introduce a new variable y. For every $j \in [n]$, multiply x_j with $y^{\sum_{i\in[m]} a_{i,j}}$. Let $f' \in \mathbb{F}[x_1, \ldots, x_n]$ be the polynomial obtained from f after transforming each x_i as above. Consider $f'' := f'/y^{\sum_{i\in[m]} a_{i,0}}$; note that f'' is a polynomial. Setting $y = 0$ in f'' yields $f|_Q$. Since all the exponents are polynomially bounded, f' (and also the polynomial obtained from it by setting some of the variables to 0) has a (weakly-skew) arithmetic circuit of size $\mathsf{poly}(s, n)$, by Strassen [47] (and Kaltofen and Koiran [24]). $\square$

Remark. If Q has $\mathsf{poly}(n)$ coefficient complexity, then it can be shown that $f|_Q$ is a one-parameter degeneration of f of $\mathsf{poly}(n)$ degree. Hence, Theorem 12 can also be deduced from Lemma 4.3.

6.2 Generic semi-invariants of generalized Kronecker quivers

We now prove Theorem 2 for the m-Kronecker quiver. Recall that the m-Kronecker quiver is the graph with two vertices s and t, with m arrows pointing from s to t. When $m = 2$, this is the classical Kronecker quiver. When $m \geq 3$, this quiver is wild.

Any tuple of m $n \times n$ matrices is a linear representation of the m-Kronecker quiver of dimension vector (n, n). Let $\mathbb{F}[x_{i,j}^{(k)}]$ denote the ring of polynomials in the variables $x_{i,j}^{(k)}$, where $i, j \in [n]$, and $k \in [m]$. For $k \in [m]$, let $X_k = (x_{i,j}^{(k)})$ denote the variable $n \times n$ matrix, whose (i, j)-th entry is $x_{i,j}^{(k)}$. Let $R(n, m)$ consist of those polynomials in $\mathbb{F}[x_{i,j}^{(k)}]$ that are invariant under the action of every $(A, C) \in \mathrm{SL}(n, \mathbb{F}) \times \mathrm{SL}(n, \mathbb{F})$, which sends $(X_1, \ldots, X_m)$ to $(AX_1C^{-1}, \ldots, AX_mC^{-1})$. $R(n, m)$ is the ring of semi-invariants for the m-Kronecker quiver for dimension vector (n, n) or "*matrix semi-invariants*" due to their similarity with the well-known matrix invariants (see Section 7.2.1).

Theorem 13. *The Newton degeneration of a generic semi-invariant of the m-Kronecker quiver with dimension vector (n, n) and degree dn to an arbitrary face can be computed by a weakly skew arithmetic circuit of size $\mathsf{poly}(d, n)$.*

Proof. Let $M(d, \mathbb{F})$ be the space of $d \times d$ matrices over $\mathbb{F}$. By the first fundamental theorem of matrix semi-invariants [1, 10, 12, 44], $\forall A_1, \ldots, A_m \in M(d, \mathbb{F})$, $\det(A_1 \otimes X_1 + \cdots + A_m \otimes X_m)$ is a matrix semi-invariant, and every matrix semi-invariant is a linear combination of such semi-invariants. When A_i's are generic, the monomials occurring in $\det(A_1 \otimes X_1 + \cdots + A_m \otimes X_m)$ have the following combinatorial description [1]. Define a magic square with the parameter (n, m, d) to be an $n \times n$ matrix S, with (i, j)-th entry $S(i, j) = (s_{i,j}^{(1)}, \ldots, s_{i,j}^{(m)}) \in \mathbb{N}^m$, satisfying: (1) $\forall i \in [n]$, $\sum_{j,k} s_{i,j}^{(k)} = d$, and (2) $\forall j \in [n]$, $\sum_{i,k} s_{i,j}^{(k)} = d$. With such a magic square, we associate a monomial in $\mathbb{F}[x_{i,j}^{(k)}]$ by setting the exponent of $x_{i,j}^{(k)}$ to $s_{i,j}^{(k)}$. When A_i's are generic, the monomials occurring in $\det(A_1 \otimes X_1 + \cdots + A_m \otimes X_m)$ are precisely the monomials associated with such magic squares.

Consider the $n \times n$ complete bipartite graph G in which, for every $(i, j) \in [n] \times [n]$, there are m edges between i and j, colored by the elements of the set $[m]$, with each color used exactly once. It is easily seen that the magic squares above correspond to the d-matchings in this graph G: for a graph $G = (V, E)$, a d-matching is a function $f : E \to \mathbb{N}$ such that $\forall v \in V$, $\sum_{e\in E, e\sim v} f(e) = d$.

Hence, the Newton polytope of a *generic* matrix semi-invariant is characterized by the following constraints:

$$(a)\, \forall i,j \in [n], k \in [m], s_{i,j}^{(k)} \geq 0; (b)\, \forall i \in [n], \sum_{j,k} s_{i,j}^{(k)} = d; (c)\, \forall j \in [n], \sum_{i,k} s_{i,j}^{(k)} = d. \tag{2}$$

This description follows easily from the fact that the incidence matrix of a bipartite graph (possibly with multiple edges) is unimodular (cf. e.g. [45, Chap. 18]). Each face of this polytope is obtained by setting some of $s_{i,j}^{(k)}$'s to 0. Hence, its coefficient complexity is polynomial in d and n. Therefore, by Theorem 12, the theorem follows. □

6.3 Generic semi-invariants of k-subspace quivers

Next, we prove Theorem 2 for the k-subspace quivers.

The k-subspace quiver is the quiver with $k+1$ vertices $\{x_1, \dots, x_k, y\}$, and k arrows $\{\alpha_i = (x_i, y) \mid i \in [k]\}$. For $k = 1, 2, 3$, the k-subspace quiver is of finite type. When $k = 4$, it is of tame type. When $k \geq 5$, it is wild.

We shall apply the description of semi-invariants of quivers by Domokos and Zubkov [12] to the case of k-subspace quivers. For this, we need some further notions. Fix a field $\mathbb{F}$. Let $Q = (Q_0, Q_1)$ be a quiver, where Q_0 is the vertex set, and Q_1 is the arrow set. For an arrow α in Q, we use $s\alpha$ (resp. $t\alpha$) to denote the start (resp. target) of α. A path π is a sequence of arrows $\alpha_1\alpha_2 \dots \alpha_\ell$ such that $t\alpha_i = s\alpha_{i+1}$ for $i \in [\ell - 1]$. The start (resp. target) of π is $s\alpha_1$ (resp. $t\alpha_\ell$). A path is cyclic if $s\pi = t\pi$. We assume Q has no cyclic paths of positive length.

Let V be a representation of Q; that is, for $x \in Q_0$, V^x is the vector space associated with x, and for $\alpha \in Q_1$, V^α is the linear map from $V^{s\alpha}$ to $V^{t\alpha}$. This extends naturally to $V^\pi = V^{\alpha_k} \cdots V^{\alpha_1} : V^{s\pi} \to V^{t\pi}$ for a path π.

Fix a dimension vector β for Q, and suppose $\ell = |Q_0|$. $|\beta| := \sum_{x \in Q_0} \beta(x)$. Given β, after fixing bases for V^x, $x \in Q_0$, a representation of Q is then specified using $n := \sum_{\alpha \in Q_1} \beta(s\alpha) \cdot \beta(t\alpha)$ numbers. Let $u_1, \dots, u_n$ be n variables.

Let $\mathrm{GL}(\beta) := \mathrm{GL}(\beta_1, \mathbb{F}) \times \dots \times \mathrm{GL}(\beta_\ell, \mathbb{F})$ be the direct product of general linear groups with corresponding dimensions acting naturally on the representations of Q with dimension vector β. Let $SI(Q, \beta) \subseteq \mathbb{F}[u_1, \dots, u_n]$ be the set of semi-invariants with respect to Q and β. Any $\sigma : Q_0 \to \mathbb{Z}$ defines a multiplicative character of $\mathrm{GL}(\beta)$, $\chi_\sigma : (B(x) \mid x \in Q_0) \in \mathrm{GL}(\beta) \to \prod_{i \in [\ell]} \det(B(x))^{\sigma(x)}$. Then define $SI(Q, \beta)_\sigma = \{f \in SI(Q, \beta) | \forall B \in \mathrm{GL}(\beta), B \cdot f = \chi_\sigma(B) f\}$. It is clear that $SI(Q, \beta) = \oplus_\sigma SI(Q, \beta)_\sigma$. If $\langle \sigma, \beta \rangle := \sum_{x \in Q_0} \sigma(x)\beta(x) \neq 0$ then there are no non-trivial $SI(Q, \beta)_\sigma$ (see e.g. [10]). Otherwise, let $\sigma = \sigma_+ - \sigma_-$, where $\sigma_+(x) = \max(\sigma(x), 0)$ and $\sigma_-(x) = \max(-\sigma(x), 0)$, and set $s = \langle \beta, \sigma_+ \rangle$.

Now we come to the key construction. Consider the $s \times s$ matrix

$$g = \oplus_{x \in Q_0} (V^x)^{\sigma_+(x)} \to \oplus_{x \in Q_0} (V^x)^{\sigma_-(x)}, \tag{3}$$

where each block, $\hom(V^x, V^y)$, is of the form $w_1 V^{\pi_1} + \dots + w_r V^{\pi_r}$ where $\pi_1, \dots, \pi_r$ runs over the set of paths from x to y, and $w_1, \dots, w_r$ are variables. For different blocks we use different variables. That is, the total number of variables is $m = \sum_{x \in Q_0} \sum_{y \in Q_0} \sigma_+(x) p(x, y) \sigma_-(y)$, where $p(x, y)$ is the number of paths between x and y. $\det(g)$ then is a polynomial in $w_1, \dots, w_m$, and $u_1, \dots, u_n$. w_i's are called auxiliary variables, since we shall use the following construction: for $(c_1, \dots, c_m) \in \mathbb{F}^m$, let $\det(g \mid w_i = c_i, i \in [m])$ be the polynomial in $\mathbb{F}[u_1, \dots, u_n]$ after assigning w_i with c_i in $\det(g)$.

Theorem 14 (Domokos and Zubkov [12])**.** *Let notations be as above. $SI(Q, \beta)_\sigma$ is linearly spanned by $\{\det(g \mid w_i = c_i, i \in [m]) \mid (c_1, \dots, c_m) \in \mathbb{F}^m\}$.*

Proposition 6.1. *The Newton degeneration of a generic semi-invariant of the k-subspace quiver of dimension vector β and degree d to an arbitrary face can be computed by a weakly-skew arithmetic circuit of size* $\mathsf{poly}(|\beta|, d)$.

The proof strategy is to apply Theorem 14 to the k-subspace quiver, which yields a combinatorial description of the exponent vectors of monomials in a generic semi-invariant of a certain weight. From the combinatorial description we obtain a description of the Newton polytope and its faces of a generic semi-invariant. We then conclude by applying Theorem 12 as for generalized Kronecker quivers.

Proof. In the k-subspace quiver we have $k+1$ vertices $\{x_1, \dots, x_k, y\}$ and k arrows $\alpha_i = (x_i, y)$. Observe that (1) a non-trivial path is of length 1; (2) only y (resp. x_i's) can serve as the target (resp. start) of a path. Therefore, for $\det(g)$ to be nonzero, it is necessary that $\sigma_+(y) = 0$ and $\sigma_-(x_i) = 0$ for $i \in [k]$. That is, Equation 3 for k-subspace quiver has to be of the form

$$g = \oplus_{i \in [k]} (V^{x_i})^{\sigma_+(x_i)} \to (V^y)^{\sigma_-(y)},$$

for $\det(g)$ to be nonzero.

g then is a block matrix of the following form: the rows are divided into $\sigma_-(y)$ blocks, with each block of size $\beta(y)$. The columns are divided into $\sum_{i \in [k]} \sigma_+(x_i)$ blocks, with $\sigma_+(x_i)$ blocks of size $\beta(x_i)$. Let the number of rectangular blocks be m ($m = \sigma_-(y) \cdot (\sum_i \sigma_+(x_i))$). As for different blocks we use different variables, so the auxiliary variables are $w_1, \dots, w_m$. In a block indexed by (y, x_i), we put in $w_\ell V^{\alpha_i}$, where V^{α_i} is a variable matrix of size $\beta(y) \times \beta(x_i)$. We fix bases for V^y and V^{x_i}: let $P = \{p_1, \dots, p_{\beta(y)}\}$ be a basis of V^y, and for $i \in [k]$, let $Q_i = \{q_{i,1}, \dots, q_{i,\beta(x_i)}\}$ be a basis of V^{x_i}. Then the rows of g can be indexed by $\cup_{i \in [\sigma_-(y)]} P^{(i)}$ where $P^{(i)}$ is a copy of P, and the columns of g are indexed by $\bigcup_{i \in [k]} \cup_{j \in [\sigma_+(x_i)]} Q_i^{(j)}$ where $Q_i^{(j)}$ is a copy of Q_i.

Viewing $\det(g)$ as a polynomial in $\mathbb{F}[w_1, \dots, w_m][\cup_{i \in [k]} V^{\alpha_i}]$, we are interested in the $\cup_{i \in [k]} V^{\alpha_i}$-monomials, as in a generic semi-invariant, these are all the monomials. Note that g is a $d \times d$ matrix where $d = \beta(y) \cdot \sigma_-(y)$, and the $\cup_{i \in [k]} V^{\alpha_i}$-monomials in g are of degree d.

To describe these monomials, we form an undirected bipartite graph $G = (L \cup R, E)$ as follows. Let $L = \{p_1, \dots, p_{\beta(y)}\}$, and $R = \{q_{i,j_i} \mid i \in [k], j_i \in [\beta(x_i)]\}$. Connect each $(p_i, q_{j,\ell})$ with an edge to form a complete bipartite graph. The number of edges is $\sum_{i \in [k]} \beta(y) \cdot \beta(x_i)$, and we can naturally identify the edges with variables in $\cup_{i \in [k]} V^{\alpha_i}$. Now consider a function $f : L \cup R \to \mathbb{N}$, with $f(p_i) = \sigma_-(y)$, $f(q_{j,\ell}) = \sigma_+(x_j)$. For such a function, we can define the f-perfect matching of G, that is a function $h : E \to \mathbb{N}$, such that for any $r \in L \cup R$, $\sum_{e \in E, e \sim r} h(e) = f(r)$, where $e \sim r$ denotes that e is an edge adjacent to r.

It is not hard to verify that the f-perfect matchings and the exponent vectors in a generic semi-invariant are in one to one correspondence. One direction is easy: a bijective function $b : \cup_{i \in [\sigma_-(y)]} P^{(i)} \to \bigcup_{i \in [k]} \cup_{j \in [\sigma_+(x_i)]} Q_i^{(j)}$ clearly defines an f-perfect matching, as there are $\sigma_-(y)$ copies of P so each $p_i \in P$ indexes $\sigma_-(y)$ rows, and similarly for the columns. Furthermore the f-perfect matching records the exponent vector of the monomial in $\cup_{i \in [k]} V^{\alpha_i}$ based on b. On the other hand, given any f-perfect matching, it is routine to check that we can construct at least one bijective functions from $\cup_{i \in [\sigma_-(y)]} P^{(i)}$ to $\bigcup_{i \in [k]} \cup_{j \in [\sigma_+(x_i)]} Q_i^{(j)}$. When there are more than one such bijective functions, it is easy to check that all of them produce the same monomial.

This suggests that we can use the description of the bipartite f-perfect matching polytope. Let $s_{i,j,\ell}$ be the variable associated with the edge $(p_i, q_{j,\ell})$, then we have the following inequalities and equalities for the bipartite f-perfect matching polytope:

$$\begin{array}{ll} \forall i \in [\beta(y)], j \in [k], \ell \in [\beta(x_j)], & s_{i,j,\ell} \geq 0 \\ \forall i \in [\beta(y)], & \sum_{j,\ell} s_{i,j,\ell} = \sigma_-(y) \\ \forall j \in [k], \ell \in [\beta(x_j)], & \sum_i s_{i,j,\ell} = \sigma_+(x_j). \end{array} \tag{4}$$

Therefore each face is also obtained by setting some $s_{i,j,\ell}$ to 0, so we can use Theorem 12 to conclude. □

6.4 Generic semi-invariants of Dynkin quivers and beyond

Finally, we prove Theorem 2 for the A-D-E Dynkin quivers.

Here, instead of the Domokos and Zubkov invariants [12], we shall use the invariants of Schofield [43], which are also known to linearly span the space of semi-invariants [10]. We recall Schofield's construction in some detail (without proof), so that we can reason about the Newton polytopes of the Schofield invariants. Our general strategy is to find (in)equalities satisfied by these Newton polytopes for arbitrary quivers, and then to show that for the A-D-E Dynkin quivers, these inequalities in fact *define* the corresponding Newton polytopes. We will then apply Theorem 12, since the inequalities we find all have small coefficient complexity.

6.4.1 Schofield invariants.

Given two representations V, W of the same quiver, the associated Schofield invariant vanishes if and only if there is a homomorphism of quiver representations from V to W. The idea is to treat a map from V to W as variable, and then try to solve the equations which say that those variables define a homomorphism of representations. These equations are linear and homogeneous, so they have a solution if and only if a certain determinant vanishes; this determinant will be the Schofield invariant.

Consider the map d^V_W that takes $(f^x \colon V^x \to W^x)$ to $(f^{t\alpha} \circ V^\alpha - W^\alpha \circ f^{s\alpha} \colon V^{s\alpha} \to W^{t\alpha})$. (Here x denotes a vertex of the quiver, α an arrow, $s\alpha$ its start, and $t\alpha$ its target.) For a given f, $d^V_W(f) = 0$ if and only if f is a homomorphism of quiver representations $V \to W$. Thus, $s(V, W) := \det(d^V_W)$ vanishes if and only if there *exists* a non-zero homomorphism $V \to W$. Since the existence of a homomorphism is basis-independent, we immediately see that the *vanishing* of $s(V, W)$ is in fact $\mathrm{GL}(V) \times \mathrm{GL}(W)$-invariant. However, this does not immediately tell us that $s(V, W)$ itself is invariant, though it is close; Schofield showed that in fact $s(V, W)$ is $\mathrm{SL}(V) \times \mathrm{SL}(W)$-invariant. When we fix the dimension vectors of V and W, but we think of *both* V and W as defined by variables, we call $s(V, W)$ a Schofield *pair* invariant. When we think of V as given by actual values but W as given by variables, we refer to $s^V(W) = s(V, W)$ as a Schofield invariant. It is these latter invariants, as V ranges over all possible dimension vectors and all possible values, that linearly span the ring of semi-invariants for the dimension vector of W [10].

Let us study the structure of the matrix d^V_W in a bit more detail. For a vertex x of a quiver Q, we let V^x (resp., W^x) denote the vector space associated to x in the representation V; for an arrow α we let V^α denote the corresponding matrix. We use $s\alpha$ to denote the "start" of the arrow α and $t\alpha$ to denote the "target" of α. We think of the matrix as acting on column vectors. The matrix d^V_W then has row indices (α, i, ℓ), where α is an arrow of Q, i ranges over a basis for $W^{t\alpha}$ and ℓ ranges over a basis for $V^{s\alpha}$; it has column indices (x, j, k) where x ranges over vertices, j ranges over a basis of W^x and k ranges over a basis of V^x. We refer to the set of rows of the form $(\alpha, *, *)$ as the α *block-row*, and the set of columns of the form $(x, *, *)$ as the x block-column.

Now we determine the entries of d^V_W precisely:

$$
\begin{array}{rcl}
d_W^V(f)_\alpha & = & f^{t\alpha} \circ V^\alpha - W^\alpha \circ f^{s\alpha} \\
d_W^V(f)_{(\alpha,i,\ell)} & = & (f^{t\alpha} \circ V^\alpha)_{i,\ell} - (W^\alpha \circ f^{s\alpha})_{i,\ell} \\
& = & \sum_m f^{t\alpha}_{i,m} V^\alpha_{m,\ell} - \sum_n W^\alpha_{i,n} f^{s\alpha}_{n,\ell} \\
d^V_{W;(\alpha,i,\ell),(x,j,k)} & = & \delta_{t\alpha,x}\delta_{i,j}V^\alpha_{k,\ell} - \delta_{s\alpha,x}W^\alpha_{i,j}\delta_{\ell,k} \\
& = & \begin{cases} \delta_{i,j}V^\alpha_{k,\ell} & \text{if } t\alpha = x \\ -\delta_{\ell,k}W^\alpha_{i,j} & \text{if } s\alpha = x \\ \delta_{i,j}V^\alpha_{k,\ell} - \delta_{\ell,k}W^\alpha_{i,j} & \text{if } s\alpha = t\alpha = x \\ 0 & \text{otherwise} \end{cases} \\
d^V_{W;\alpha;x} & = & \delta_{t\alpha,x} I_{(W^x)} \otimes (V^\alpha)^T - \delta_{s\alpha,x} W^\alpha \otimes I_{(V^x)}
\end{array}
$$

Note that each W^α only appears (though multiple times) in a single block-row and block-column, namely $(\alpha, s\alpha)$, and each V^α only appears in a single block-row and block-column, namely $(\alpha, t\alpha)$. Note that every time V^α appears, it appears transposed. From the preceding equation, we see that the α block-row is the unique set of $(\dim W^{t\alpha})(\dim V^{s\alpha})$ rows that contain all instances of the variables W^α, V^α, and the x block-column is the unique set of $(\dim W^x)(\dim V^x)$ columns containing all instances of the variables W^α for all α such that $s\alpha = x$, and all instances of the variables V^α for all α such that $t\alpha = x$.

We now begin deriving some inequalities satisfied by $\mathrm{NPT}(s(V,W))$. Let $\omega^\alpha_{i,j}$ be the exponent corresponding to the variable $W^\alpha_{i,j}$ ($\alpha \in E(Q)$, $i \in [\dim W^{t\alpha}]$, $j \in [\dim W^{s\alpha}]$), and $\nu^\alpha_{k,\ell}$ the exponent corresponding to the variable $V^\alpha_{k,\ell}$ ($k \in [\dim V^{t\alpha}]$, $\ell \in [\dim V^{s\alpha}]$). We try to be consistent in our usage of i, j, k, ℓ throughout.

As with all Newton polytopes, we have:

$$
\omega^\alpha_{i,j} \geq 0 \qquad \nu^\alpha_{k,\ell} \geq 0 \qquad\qquad \forall \alpha, i, j, k, \ell \tag{5}
$$

Blocks. For each $\alpha \in E(Q)$, in each monomial there must be exactly as many entries chosen from the rows in the α block-row as the total number of rows in the block-row:

$$
\sum_{i,j} \omega^\alpha_{i,j} + \sum_{k,\ell} \nu^\alpha_{k,\ell} = (\dim W^{t\alpha})(\dim V^{s\alpha}) \qquad\qquad \forall \alpha \in E(Q) \tag{6}
$$

Similarly, for each $x \in V(Q)$, in each monomial there must be exactly as many entries chosen from the columns in the x-th block-column as the total number of columns in that block-column:

$$
\sum_{\alpha : s\alpha = x} \sum_{i,j} \omega^\alpha_{i,j} + \sum_{\alpha : t\alpha = x} \sum_{k,\ell} \nu^\alpha_{k,\ell} = (\dim W^x)(\dim V^x) \qquad\qquad \forall x \in V(Q) \tag{7}
$$

Mini-blocks. By the *W-mini-block-row* corresponding to $(\alpha, *, \ell)$ ($\ell \in [\dim V^{s\alpha}]$), we mean the unique set of $\dim W^{t\alpha}$ rows in which the ℓ-th copy of W^α appears. By the *V-mini-block-row* corresponding to $(\alpha, i, *)$ ($i \in [\dim W^{t\alpha}]$), we mean the unique set of $\dim V^{s\alpha}$ rows containing the i-th copy of V^α. Note that the W-mini-block-rows and the V-mini-block-rows appear in a collated or product-like fashion: For each α, the $(\alpha, i, *)$ W-mini-block-row intersects each V-mini-block-row $(\alpha, *, \ell)$ in exactly one row (namely (α, i, ℓ)), and vice versa. Similarly, by the *W-mini-block-column* corresponding to $(x, *, k)$ ($k \in [\dim V^x]$), we mean the unique set of $\dim W^x$ columns in which the k-th copy of W^α appears for all α such that $s\alpha = x$. By the *V-mini-block-column* corresponding to

$(x, j, *)$ ($j \in [\dim W^x]$) we mean the unique set of $\dim V^x$ columns in which the j-th copy of of V^α appears, for all α such that $t\alpha = x$.

Consider the $(\alpha, i, *)$ V-mini-block-row, i.e., the set of rows in which the i-th copy of V^α appears, which is the same as the set of rows that contain *every* occurrence of the variables $W^\alpha_{i,*}$ (and no other W variables). In each copy of W^α, each monomial can pick at most one variable from the i-th row, and there are $\dim V^{s\alpha}$ copies of W^α on pairwise disjoint sets of rows (W-mini-block-rows), so each monomial can have degree at most $\dim V^{s\alpha}$ in the variables $W^\alpha_{i,*}$:

$$\sum_j \omega^\alpha_{i,j} \leq \dim V^{s\alpha} \qquad\qquad \forall\alpha\forall i \in [\dim W^{t\alpha}]. \tag{8}$$

On the other hand, if too few elements from $W^\alpha_{i,*}$ are picked, then any such monomial will be forced to pick two elements of the i-th copy of V^α *from the same column*, which is not allowed, so we also have:

$$\sum_j \omega^\alpha_{i,j} \geq \dim V^{s\alpha} - \dim V^{t\alpha} \qquad\qquad \forall\alpha\forall i \in [\dim W^{t\alpha}] \tag{9}$$

Additionally, in the $(\alpha, i, *)$ V-mini-block-row, in total each monomial must select exactly one variable from each of the $\dim V^{s\alpha}$ rows. The natural thing to do here would be to add in the degree of V^α. However, in doing so we may overcount, since the ν's also include choices of V-variables that appear in other rows. So we get:

$$\sum_j \omega^\alpha_{i,j} + \sum_{k,\ell} \nu^\alpha_{k,\ell} \geq \dim V^{s\alpha} \qquad\qquad \forall\alpha\forall i \in [\dim W^{t\alpha}]. \tag{10}$$

Despite the similarity of the preceding two inequalities, we note that they are in fact independent, as $\sum_{k,\ell} \nu^\alpha_{k,\ell}$ can be larger than $\dim V^{t\alpha}$ (e.g., when $\dim W^{t\alpha} > \dim V^{t\alpha}$) or smaller than $\dim V^{t\alpha}$ (e.g., when dimensions align so that a term may cover all of the α block rows by W^α variables, using none of the V^α variables).

For the mini-block-*columns*, we can also get additional information about the mini-blocks by considering the "complement," since there is only one W^α that is in the same row as any given V^α. Given a V-mini-block-column $(x, j, *)$, we know that exactly $\dim V^x$ entries must get chosen from this V-mini-block-column. The entries in this V-mini-block-column are the j-th copies of those V^α such that $t\alpha = x$, as well as all of the appearances of the columns $W^\alpha_{*,j}$ when $s\alpha = x$. However, for each α such that V^α appears in this V-mini-block-column, the number of V^α entries chosen in this V-mini-block-column is complementary to the number of $W^\alpha_{j,*}$ entries (same α, and the j is purposefully the row index now) chosen from the copy of W^α in the $s\alpha$ block-column. So we get:

$$\sum_{\alpha : s\alpha = x} \sum_i \omega^\alpha_{i,j} + \sum_{\alpha : t\alpha = x} \left(\dim V^{s\alpha} - \sum_{j'} \omega^\alpha_{j,j'} \right) = \dim V^x \qquad\qquad \forall x \forall j \in [\dim W^x] \tag{11}$$

The V-W symmetric arguments (with appropriate transposes, etc.) then give:

$$\sum_k \nu_{k,\ell}^{\alpha} \leq \dim W^{t\alpha} \qquad \forall \alpha \forall \ell \in [\dim V^{s\alpha}] \tag{12}$$

$$\sum_{i,j} \omega_{i,j}^{\alpha} + \sum_k \nu_{k,\ell}^{\alpha} \geq \dim W^{t\alpha} \qquad \forall \alpha \forall \ell \in [\dim V^{s\alpha}] \tag{13}$$

$$\dim W^{s\alpha} + \sum_k \nu_{k,\ell}^{\alpha} \geq \dim W^{t\alpha} \qquad \forall \alpha \forall \ell \in [\dim V^{s\alpha}] \tag{14}$$

$$\sum_{\alpha : t\alpha = x} \sum_{\ell} \nu_{k,\ell}^{\alpha} + \sum_{\alpha : s\alpha = x} \left(\dim W^{t\alpha} - \sum_{k'} \nu_{k',k}^{\alpha} \right) = \dim W^x \qquad \forall x \forall k \in [\dim V^x] \tag{15}$$

In some cases, these equations are already sufficient, for example for the generalized Kronecker quivers. These equations may in fact suffice in general (with suitable modifications for degenerate cases, such as when certain dimensions are 1).

Example 6.2. As an example, let us observe that these equations reduce to the equations (2) for the generic semi-invariants of the m-Kronecker quiver for dimension vector (n, n). In this case, we have two vertices, x, y, with m arrows $x \to y$, say labelled $1, \ldots, m$. We have $\dim W^x = \dim W^y = n$. There are $n(\dim V^x + \dim V^y)$ columns. Since every arrow goes from $x \to y$, every block-row has $(\dim V^x)(\dim W^y)$ rows, for a total of $mn(\dim V^x)$ rows. To get these two quantities to be equal, we must have $\dim V^x + \dim V^y = m \dim V^x$, or equivalently, $\dim V^y = (m-1)\dim V^x$. Let $d = \dim V^x$, so $\dim V^y = d(m-1)$. We then get a Schofield pair invariant of total degree dmn. However, since all the W^{α} appear in the x block-column, which consists of exactly dn columns, the W-degree of *every* term of $s(V, W)$ is exactly dn. This is equivalent to (7) for the vertex x. This d matches with the d of (2).

Considering equation (11) for the block-column y, we find that there are no α such that $s\alpha = x$, so the equation reduces to

$$\sum_{\alpha} \left(\dim V^{s\alpha} - \sum_j \omega_{i,j}^{\alpha} \right) \quad = \quad \dim V^y \qquad (\forall i)$$

Since $s\alpha = x$ for all α, we have $\dim V^{s\alpha} = \dim V^x = d$ for all α. Thus $\sum_{\alpha} \dim V^{s\alpha} = md$. Combining this with the above equality, and the fact that $\dim V^y = (m-1)d$, yields the second equation of (2).

Considering equation (11) for the block-column x, we find that there are no α such that $t\alpha = x$, and the equation immediately becomes the third equation of (2).

6.4.2 Dynkin quivers.

The "A–D–E" Dynkin quivers are important because these are the only quivers of *finite* representation type: they have only finitely many indecomposable representations. All other quivers have infinitely many. The Dynkin quiver of type A_n is defined by having its underlying undirected graph be a line on n vertices. D_n is a line on $n-2$ vertices, with two additional vertices attached to one end. E_n for $n = 6, 7, 8$ (the only ones relevant to the preceding classification) is a path of length $n-1$, together with an additional vertex attached to the third vertex on the path. The classification statement above is independent of the orientation of the edges, but the invariant theory can change with a change in orientation, so we must take some care.

Theorem 15. *For any of the ADE Dynkin quivers, with arbitrary orientation of arrows, the Newton degeneration of a generic semi-invariant with dimension vector $(n_1, \ldots, n_k)$ and degree d to an arbitrary face has determinantal complexity $\leq \mathsf{poly}(\sum n_i, d)$.*

The key to the proof is the following lemma about Schofield *pair* invariants:

Lemma 6.3. *For any of the ADE Dynkin quivers with arbitrary orientation, Equations (5)–(15) define the Newton polytope of any Schofield pair invariant $s(V, W)$.*

Let us first see how the theorem follows from the lemma, then return to prove the lemma:

Proof of Theorem 15 from Lemma 6.3. The Newton polytope of a *generic* Schofield semi-invariant $s^V(W)$—that is, for generic V—is the same as the projection of $\mathrm{NPT}(s(V, W))$ into the W subspace. Let π be this projection. The π-preimage of a face Q of $\mathrm{NPT}(s^V(W))$ is therefore a face of $\mathrm{NPT}(s(V, W))$, and thus $s^V(W)|_Q = s(V, W)|_{\pi^{-1}(Q)}$. By Lemma 6.3, the coefficient complexity of $\mathrm{NPT}(s(V, W))$ is bounded by a polynomial in $\{\dim V^i, n_i | i \in V(Q)\}$. If we can bound these quantities by $\mathsf{poly}(d, n_1, \ldots, n_k)$, then Theorem 12 immediately completes the proof.

To bound the size of d_W^V, we determine the W-degree of $s(V, W)$ for variable V. This is easily calculated, using the description of d_W^V above, as

$$\sum_{x \in V(Q)} \min\{(\dim V^x)(\dim W^x), \sum_{\alpha : s\alpha = x} (\dim W^{t\alpha})(\dim V^{s\alpha})\}.$$

Thus we see that each component of the dimension vector of V is bounded by d (better upper bounds are possible, but this will suffice), which is small enough for the preceding argument to go through. Thus we have proved the theorem for generic *Schofield* invariants. Finally, as the Schofield invariants s^V linearly span the semi-invariants [10], Theorem 15 follows. $\square$

And now we proceed to the proof of the key lemma.

Proof of Lemma 6.3. Type A_n. We proceed by induction on n. Note that the $n = 1$ case is trivial (there are no arrows), and the $n = 2$ case is the degenerate case of the Kronecker quiver with only a single arrow, which was handled in Section 6.2 and Example 6.2.

Suppose that one of the end vertices, x, is a source (the sink case is analogous, swapping the roles of V and W and transposing if needed). Call the unique outgoing arrow $\alpha \colon x \to y$. Then in the $(x, *, *)$ block column, only W^α appears. Then (11) gives that $\sum_i \omega_{i,j}^\alpha = \dim V^x$ for all $j \in \dim W^x$. Summing over j then gives $\sum_{i,j} \omega_{i,j}^\alpha = \dim V^x \dim W^x$. Equation (6) then becomes $\dim V^x \dim W^x + \sum_{k,\ell} \nu_{k,\ell}^\alpha = \dim W^y \dim V^x$, or equivalently

$$\sum_{k,\ell} \nu_{k,\ell}^\alpha = \dim V^x (\dim W^y - \dim W^x).$$

By the nonnegativity of $\nu_{k,\ell}^\alpha$ we see that there can be no nontrivial semi-invariants unless $\dim W^y \geq \dim W^x$; furthermore, if $\dim W^y = \dim W^x$, then the Schofield pair invariants do not involve V^α at all. It is not hard to see that in this case the Schofield pair invariant is just $\det(W^\alpha)^{\dim V^x}$ times the Schofield pair invariant $s_x(V, W)$ of the quiver one gets by deleting the vertex x. Then $\mathrm{NPT}(s(V, W)) = ((\dim V^x) \cdot \mathrm{NPT}(\det(W^\alpha))) \times \mathrm{NPT}(s_x(V, W))$, where the scalar dot here represents scaling the polytope (technically this is a Minkowski sum of $\mathrm{NPT}(\det(W^\alpha))$ with itself $\dim V^x$ times, but since Newton polytopes are, in particular, convex, this is the same as scaling up the polytope). The product here represents Cartesian product, since the constraints on the W^α are independent of the constraints on the remaining variables in this case. So the inequalities are the $\dim V^x$-scaled inequalities for $\mathrm{NPT}(\det(W^\alpha))$ (which is just a rescaling of the perfect matching polytope for the complete bipartite graph), and those for $\mathrm{NPT}(s_x(V,W))$, which are (5)–(15), by induction.

Otherwise, $\dim W^y > \dim W^x$. In this case, (15) says that $\sum_k \nu^\alpha_{k,\ell} \geq \dim W^y - \dim W^x$. Summing these over all ℓ and combining with the equation above, we get that in fact $\sum_k \nu^\alpha_{k,\ell}$ is *equal* to $\dim W^y - \dim W^x$, for each ℓ.

Now, we have two cases: either the other arrow incident on y is oriented towards y or away from y.

Case 1: $\beta\colon z \to y$ is oriented towards y (in particular, y is a sink). In this case, the only entries that appear in the y-block-column are V^α and V^β. The equations for this block-column and the corresponding mini-block-columns are then exact, since there is no "mixing of indices" that would occur had there been both V's and W's in the same column. This allows us to easily link to the rest of the d^V_W matrix and proceed inductively.

Case 2: $\beta\colon y \to z$ is oriented away from y. In this case, the entries that appear in the y-block-column are V^α and W^β. In this case, the equations for the y block column and the corresponding mini-block-columns give constraints on the total degree of the W^β, and (15) links V^α (in the y-block column) to V^β in the z-block column by an *equality* (not inequality), again allowing us to easily proceed inductively to the rest of the matrix.

Type D_n. The induction is in fact the same, starting from the "long end," of the quiver. The base case, however, is now D_4. One orientation of the D_4 quiver gives the 3-subspace quiver, which was handled in Section 6.3. So we must handle the other possible orientations. Let u be the "central" vertex (degree 3), and v_1, v_2, v_3 the outer vertices of degree 1. Because of the S_3 symmetries of D_4, there are only four possible orientations up to symmetry, determined precisely by whether there are 0,1,2, or 3 arrows pointing towards u. When there are 3 we have the 3-subspace quiver, so we need only handle the other three cases.

We use α_i to denote the arrow between v_i and u (whichever direction it is facing), for each $i = 1, 2, 3$.

Case 0: All of the arrows are pointing *away* from u. In this case, the u-block-column contains only V blocks, and each v_i block column contains only its corresponding W blocks. As in the A_n case, this means that $\dim W^u \geq \dim W^{v_i}$ for all $i = 1, 2, 3$ is required to get any non-constant semi-invariants. If $\dim W^{v_i} = \dim W^u$ for some i, then V^{α_i} doesn't appear at all in the Schofield pair invariant, and the Schofield pair invariant is a power of $\det(W^{\alpha_i})$ times the Schofield pair invariant for the quiver representation gotten by removing the vertex v_i. But the remaining quiver in this case is an A_3 quiver, which is covered by the A_n case above.

So now we assume that $\dim W^u > \dim W^{v_i}$ for all i. Any term of $s(V,W)$ is therefore determined by: (a) picking exactly $\dim W^{v_i}$ entries—in distinct rows and columns—from each of the $\dim V^{v_i}$ occurrences of W^{α_i}, (b) for each ℓ, picking exactly $\dim W^u - \dim W^{v_i}$ entries from among the $V^{\alpha_i}_{*,\ell}$ that are in rows different from those entries picked in (a), and (c) ensuring that for each k, exactly $\dim W^u$ entries are chosen from among the $V^{\alpha_*}_{k,*}$. Our goal is to show that any vertex of the polytope defined by (5)–(15) satisfies these conditions.

Equation (11) for $(v_1, j, *)$ turns into $\sum_i \omega^{\alpha_1}_{i,j} = \dim V^{v_1}$, and similarly for v_2 and v_3. For $(u, j, *)$, (11) becomes $\sum_{i \in [3]} (\dim V^{s\alpha_i} - \sum_{j'} \omega^{\alpha_i}_{j,j'}) = \dim V^u$.

Equation (15) for $(u, *, k)$ turns into $\sum_{i\in[3]} \sum_\ell \nu^{\alpha_i}_{k,\ell} = \dim W^u$, which is precisely condition (c). For $(v_i, *, k)$ it becomes $\dim W^u - \sum_{k'} \nu^{\alpha_i}_{k',k} = \dim W^{v_i}$, or equivalently $\sum_k \nu^{\alpha_i}_{k,\ell} = \dim W^u - \dim W^{v_i}$ for all ℓ (that is, (14) is already an equality for this quiver).

For this quiver, (13) is either redundant or automatically an equality. By the preceding paragraph, (13) turns into $\sum_{i,j} \omega^{\alpha_1}_{i,j} \geq \dim W^{v_1}$, when in fact we already have that $\sum_{i,j} \omega^{\alpha_1}_{i,j} = \dim V^{v_1} \dim W^{v_1}$. So if $\dim W^{v_1} = 1$, then (13) is automatically an equality, and if $\dim W^{v_1} > 1$ then it is redundant so we need not worry about using it for defining vertices. Similarly for v_2, v_3.

For this quiver, we also see that (12) is redundant, since (14) is an equality. This leaves only the nonnegativity constraints (which are easily satisfied by equality by considering monomials not involving the given variable), and (8)–(10).

Setting (8) to an equality amounts to picking exactly on element from the i-th row of each copy of W^α, and there are certainly monomials in the Schofield invariant that do this. Once this is done, (10) and (9) become redundant.

Setting (9) to an equality amounts to only picking elements of W^α from its i-th row in $\dim V^{v_1} - \dim V^u$ of the copies of W^α, rather than from each copy. It is possible to find monomials that do this for several values i; when it is no longer possible, this is because there simply aren't enough rows to accommodate not picking from some of them. But this is ruled out by the relations needed between the dimensions of the V's and W's to make d_W^V square. Once this is done, (8) becomes redundant, and (10) becomes $\sum_{k,\ell} \nu_{k,\ell}^\alpha \geq \dim V^{v_1}$.

Setting (10) to an equality can be achieved by considering only the terms in $s(V, W)$ where V^α is only taken from its i-th occurrences (equivalently, by zero-ing out all except the i-th occurrence of V^α). This can be done as often as we like, so long as there are enough V's left in the u-block-column. But when there are no longer enough left, we will have reached the empty polytope, again by the necessary relations between the dimensions of V's and W's.

Thus for any vertex defined by the above equations, there is a monomial in $s(V, W)$ with that exponent vector, as desired.

The remaining cases, although they at first seem cosmetically different, turn out to have essentially the same proof. The key fact that enables this is that in each v_i block column, at most one matrix appears, albeit multiple times.

Type E_n. The induction is essentially the same, except now we must induct starting from *each* of the three "long ends," until we again get down to a D_4 quiver as the base case. This was handled above, so we are done. $\square$

This proof easily extends to any quiver that is a tree with at most one vertex of degree > 2. We expect that it should extend without much difficulty to arbitrary trees, and it may even extend to completely arbitrary quivers.

7 Additional examples of Newton degenerations in VP

In this section, we give additional examples of representation-theoretic symbolic determinants whose Newton degenerations can be computed by symbolic determinants of polynomial size. These examples suggest that explicit families in Newton-$\mathrm{VP}_{\mathrm{ws}} \setminus \mathrm{VP}_{\mathrm{ws}}$ have to be rather delicate.

7.1 Schur polynomials, and the self-replication phenomenon

For an integer ℓ, the elementary symmetric polynomial $e_\ell(\mathbf{x}) \in \mathbb{Z}[x_1, \dots, x_n]$ is $\sum_{1 \leq i_1 < i_2 < \dots < i_\ell \leq n} x_{i_1} \cdot x_{i_2} \cdot \dots \cdot x_{i_\ell}$. In particular, for $\ell > n$ or $\ell < 0$, $e_\ell(\mathbf{x}) = 0$; for $\ell = 0$, $e_\ell = 1$.

A partition $\alpha = (\alpha_1, \dots, \alpha_n)$ is an element in $\mathbb{N}^n$ with $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_n$. The conjugate of α, denoted α', is a partition in $\mathbb{N}^m$ defined by setting $\alpha'_i = |\{j \in [n] \mid \alpha_j \geq i\}|$, where $m \geq \alpha_1$. Let $|\alpha| := \sum_i \alpha_i$.

The Schur polynomial $\mathrm{s}_\alpha(\mathbf{x})$ in $\mathbb{Z}[\mathbf{x}]$ can be defined by the Jacobi–Trudi formula as: $\mathrm{s}_\alpha(\mathbf{x}) = \det[e_{\alpha'_i - i + j}(\mathbf{x})]_{i,j \in [n]}$ [17]. As e_ℓ can be computed by a depth-3 arithmetic formula of size $O(\ell^2)$ (by Ben-Or; see [39]), s_α can be computed by a weakly-skew circuit of size $\mathsf{poly}(|\alpha|, n)$.

It is well-known that Schur polynomials are symmetric and homogeneous. For a partition β such that $|\beta| = |\alpha|$, the monomial $\mathbf{x}^\beta$ is in s_α if and only if α *dominates* β; that is, for every $i \in [n]$, $\sum_{j=1}^{i} \alpha_j \geq \sum_{j=1}^{i} \beta_j$. It follows that the Newton polytope of s_α is the permutohedron with respect to α, denoted PH_α. It is defined as follows. For $\pi \in \mathrm{S}_n$, let $\alpha^\pi = (\alpha_{\pi(1)}, \dots, \alpha_{\pi(n)})$. Then PH_α is the convex hull of α^π's, where π ranges over all permutations in S_n.

To determine the Newton degeneration of Schur polynomials, we need to understand faces of this permutohedron. Any face Q of PH_α is determined by a sequence of nested subsets $\emptyset \subsetneq S_1 \subsetneq \dots \subsetneq S_k = [n]$. For such a sequence, Q is the convex hull of α^π's with $\pi \in \mathrm{S}_n$ satisfying: for any $1 \leq i < j \leq k$, and any $p \in S_i$ and $q \in S_j$, $\alpha_{\pi(p)} \leq \alpha_{\pi(q)}$. Thus, the face Q is the Minkowski sum of k permutohedra, and $\mathrm{s}_\alpha|_Q$ is a product of Schur polynomials, which in turn can be computed by a weakly-skew arithmetic circuit of size $\mathsf{poly}(|\alpha|, n)$. This yields:

Proposition 7.1. *The Newton degeneration of a Schur polynomial to any face of its Newton polytope is a product of some Schur polynomials.*

Thus the Newton polytope of a Schur polynomial has the *self-replication* property: any face of this polytope is a Minkowski sum of smaller polytopes of the same kind. Another polynomial with such self-replication property is the resultant, cf. Sturmfels [48].

7.2 Monotone circuits

An arithmetic circuit over $\mathbb{Q}$ or $\mathbb{R}$ is called *monotone* if it uses only nonnegative field elements as constants. No cancellation can happen during the computation by a monotone circuit. This fact can be used to prove the following result.

Theorem 16. *If $f \in \mathbb{R}[x_1, \dots, x_n]$ can be computed by a monotone arithmetic circuit of size s, then, for any face Q of $\mathrm{NPT}(f)$, $f|_Q$ can be computed by a monotone arithmetic circuit of size $\leq s$.*

Proof. Let C be the a monotone arithmetic circuit of size s computing f. Since C is monotone, each subcircuit is also monotone and computes a polynomial with nonnegative coefficients.

Let Q be defined by the supporting hyperplane $\langle \mathbf{a}, \mathbf{x} \rangle = b$. Therefore, a monomial $\mathbf{x}^\mathbf{e}$ in f is present in $f|_Q$ if and only if $\langle \mathbf{a}, \mathbf{e} \rangle$ achieves the minimum b among all monomials in f.

We now transform C to a monotone circuit C' computing f'. The resulting C' will be of the same structure of C, except that some edges may be removed. In particular, the gate sets of C and C' are the same: for a gate v in C, we shall use v' to denote the correspondent of v in C'. We also use f_v to denote the polynomial computed at the gate v. The goal is to ensure that, for any gate v', in C', $f_{v'}$ consists of those monomials achieving the minimum along $\mathbf{a}$ among all monomials in f_v.

This is achieved by induction on the depth. As the base case, we keep all the leaves of C to be the leaves of C'.

Now we proceed to the inductive step. By the induction hypothesis, for every gate v' of depth $\leq d$, $f_{v'}$ consists of those monomials achieving the minimum along $\mathbf{a}$ among all monomials in f_v. Let w be a gate at depth $d+1$ in C, with two children v and u. Suppose in f_v (resp. f_u) the monomials achieve the minimum b_v (resp. b_u) along $\mathbf{a}$. If w is labeled with $+$ and $f_w = f_v + f_u$, then we set $f_{w'} = f_{v'}$ if $b_v < b_u$, $f_{w'} = f_{u'}$ if $b_v > b_u$, and $f_{w'} = f_{v'} + f_{u'}$ if $b_u = b_v$. If w is labeled with $\times$ and $f_w = f_v \times f_u$, then we set $f_{w'} = f_{v'} \times f_{u'}$. Since no cancellation can happen, it follows from the induction hypothesis that $f_{w'}$ consists of those monomials achieving the minimum along $\mathbf{a}$ among all monomials in f_w. $\square$

Remark. 1. The resulting circuit C' in the proof of Theorem 16 preserves most of the structural properties of C. In particular, if C is weakly-skew, then C' is also weakly-skew.

2. The preceding proof does not work in the presence of cancellations. For example, when w is labelled with $+$, we have $f_w = f_v + f_u$, but $f_{w'}$ may not be $f_{v'}$, $f_{u'}$, or $f_{v'} + f_{u'}$ due to cancellations.

7.2.1 Trace monomials

We now apply Theorem 16 to classical matrix invariants.

Let $\mathbb{F}[x_{i,j}^{(k)}]$ be the ring of polynomials in the variables $x_{i,j}^{(k)}$, where $i, j \in [n]$, and $k \in [m]$. Let $X_k = (x_{i,j}^{(k)})$ be an $n \times n$ variable matrix whose (i, j)-th entry is $x_{i,j}^{(k)}$. The classical matrix invariants are those polynomials in $\mathbb{F}[x_{i,j}^{(k)}]$ that are invariant under the conjugation action of $\mathrm{SL}(n, \mathbb{F})$. Under this action, $A \in \mathrm{SL}(n, \mathbb{F})$ sends $(X_1, \ldots, X_m)$ to $(AX_1A^{-1}, \ldots, AX_mA^{-1})$. The matrix invariants are semi-invariants of the quiver with a single vertex and m self-loops.

Important examples of matrix invariants are the polynomials of the form $\mathrm{Tr}(X_{i_1} \cdot \cdots \cdot X_{i_\ell})$, $i_j \in [m]$, called the *trace monomials*. By the first fundamental theorem of matrix invariants [40], every matrix invariant is a linear combination of trace monomials in characteristic zero.

Suppose $\mathbb{F} = \mathbb{Q}$ or $\mathbb{R}$. Then any Newton degeneration of a trace monomial can be computed by a weakly skew arithmetic circuit of size $\mathsf{poly}(n, \ell)$. This follows from Theorem 16, since a trace monomial can be computed by a monotone weakly skew arithmetic circuit of size $\mathsf{poly}(n, \ell)$.

Acknowledgment

We thank the anonymous reviewers for careful reading and suggestions that greatly helped to improve the writing of the paper. During this work, J. A. G. was supported by an SFI Omidyar Fellowship, K. D. M. by the NSF grant CCF-1017760, and Y. Q. by the Australian Research Council DECRA DE150100720.

References

[1] B. Adsul, S. Nayak, and K. V. Subrahmanyam. A geometric approach to the Kronecker problem II: rectangular shapes, invariants of matrices and the Artin–Procesi theorem. preprint, 2007.

[2] G. Belitskii. Complexity of matrix problems. *Linear algebra and its applications*, 361:203–222, 2003.

[3] Genrich R. Belitskii and Vladimir V. Sergeichuk. Complexity of matrix problems. *Linear Algebra Appl.*, 361:203–222, 2003. Ninth Conference of the International Linear Algebra Society (Haifa, 2001).

[4] D. Bini. Relations between exact and approximate bilinear algorithms. Applications. *Calcolo*, 17(1):87–97, 1980.

[5] P. Brooksbank and E. Luks. Testing isomorphism of modules. *Journal of Algebra*, 320(11):4020–4029, 2008.

[6] P. Bürgisser. The complexity of factors of multivariate polynomials. *Found. Comput. Math.*, pages 369–396, 2004.

[7] P. Bürgisser, M. Clausen, and M.A. Shokrollahi. *Algebraic Complexity Theory.* A series of comprehensive studies in mathematics. Springer, 1997.

[8] P. Bürgisser, J. Landsberg, L. Manivel, and J. 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.

[9] A. Chistov, G. Ivanyos, and M. Karpinski. Polynomial time algorithms for modules over finite dimensional algebras. In *ISSAC*, pages 68–74, 1997.

[10] H. Derksen and J. Weyman. Semi-invariants of quivers and saturation for Littlewood-Richardson coefficients. *Journal of the American Mathematical Society*, 13(3):467–479, 2000.

[11] H. Derksen and J. Weyman. Quiver representations. *Notices of the American Mathematical Society*, 52(2):200–206, 2005.

[12] M. Domokos and A. Zubkov. Semi-invariants of quivers as determinants. *Transformation groups*, 6(1):9–24, 2001.

[13] J. Drozd. Tame and wild matrix problems. *Amer. Math. Soc. Transl.*, 128(2):31–55, 1986.

[14] J. Edmonds. Maximum matching and a polyhedron with 0, l-vertices. *J. Res. Nat. Bur. Standards B*, 69(1965):125–130, 1965.

[15] M. Forbes and A. Shpilka. Explicit Noether normalization for simultaneous conjugation via polynomial identity testing. *ArXiv:1303.0084v2 [cs.CC]*, 2013.

[16] S. Friedland. Simultaneous similarity of matrices. *Adv. in Math.*, 50(3):189–265, 1983.

[17] W. Fulton and J. Harris. *Representation theory: A first course.* Springer, 1991.

[18] J. A. Grochow. *Symmetry and equivalence relations in classical and geometric complexity theory.* PhD thesis, University of Chicago, Chicago, IL, 2012.

[19] J. A. Grochow. Unifying known lower bounds via geometric complexity theory. *Computational Complexity*, 24:393–475, 2015. Special issue on IEEE CCC 2014.

[20] J. A. Grochow, K. D. Mulmuley, and Y. Qiao. Boundaries of VP and VNP (preliminary report). In *43rd International Colloquium on Automata, Languages, and Programming (ICALP16)*, 2016. To appear.

[21] J. Heintz and C. Schnorr. Testing polyomials that are easy to compute. *STOC*, pages 262–272, 1980.

[22] R. Impagliazzo and A. Wigderson. P=BPP unless E has sub-exponential size circuits: Derandomizing the XOR lemma. *STOC*, 1997.

[23] G. Ivanyos, Karpinski M., and N. Saxena. Deterministic polynomial time algorithms for matrix completion problems. *SIAM J. Comput.*, 39(8):3736–3751, 2010.

[24] E. Kaltofen and P. Koiran. Expressing a fraction of two determinants as a determinant. In J. Rafael Sendra and Laureano González-Vega, editors, *Symbolic and Algebraic Computation, International Symposium, ISSAC 2008, Linz/Hagenberg, Austria, July 20-23, 2008, Proceedings*, pages 141–146. ACM, 2008.

[25] N. Kayal and R. Saptharishi. A selection of lower bounds for arithmetic circuits. *Progress in Computer Science and Applied Logic*, 26, 2014.

[26] G. Kempf. Instability in invariant theory. *Annals of Mathematics*, 108(2):299–316, 1978.

[27] F. Kirwan. *Cohomology of Quotients in Symplectic and Algebraic Geometry.* Mathematical Notes 31. Princeton University Press, 1984.

[28] Adam R. Klivans and Dieter van Melkebeek. Graph nonisomorphism has subexponential size proofs unless the polynomial-time hierarchy collapses. *SIAM J. Comput.*, 31(5):1501–1526 (electronic), 2002.

[29] P. Koiran. Hilbert's Nullstellensatz is in the polynomial hierarchy. *Journal of complexity*, 12:273–286, 1996.

[30] J. M. Landsberg. An introduction to geometric complexity theory. *Newsletter Eur. Math. Soc.*, pages 10–18, 2016. Preprint available as arXiv:1509.02503 [math.AG].

[31] Joseph M. Landsberg, Laurent Manivel, and Nicolas Ressayre. Hypersurfaces with degenerate duals and the geometric complexity theory program. *Comment. Math. Helv.*, 88(2):469–484, 2013. Preprint available as arXiv:1004.4802 [math.AG].

[32] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.

[33] G. Malod and N. Portier. Characterizing Valiant's algebraic complexity classes. *J. Complex.*, 24(1):16–38, 2008.

[34] T. Mignon and N. Ressayre. A quadratic bound for the determinant and permanent problem. *International Mathematics Research Notices*, pages 4241–4253, 2004.

[35] K. Mulmuley. Geometric complexity theory V: Efficient algorithms for Noether Normalization. To appear in the Journal of the AMS. Preliminary version (with a different subtitle): FOCS, 2012.

[36] K. Mulmuley and M. Sohoni. Geometric complexity theory I: an approach to the P vs. NP and related problems. *SIAM J. Comput.*, 31(2):496–526, 2001.

[37] K. Mulmuley and M. Sohoni. Geometric complexity theory II: towards explicit obstructions for embeddings among class varieties. *SIAM J. Comput.*, 38(3):1175–1206, 2008.

[38] D. Mumford, J. Fogarty, and F. Kirwan. *Geometric invariant theory.* Springer-Verlag, 1994.

[39] Noam Nisan and Avi Wigderson. Lower bounds on arithmetic circuits via partial derivatives. *Computational Complexity*, 6(3):217–234, 1997.

[40] C. Procesi. The invariant theory of $n \times n$ matrices. *Advances in Mathematics*, 19(3):306–381, 1976.

[41] Ran Raz and Amir Yehudayoff. Lower bounds and separations for constant depth multilinear circuits. *Comput. Complexity*, 18(2):171–207, 2009.

[42] M. Rosenlicht. Some basic theorems on algebraic groups. *Amer. J. Math.*, 78:401–443, 1956.

[43] A. Schofield. Semi-invariants of quivers. *J. London Math. Soc. (2)*, 43(3):385–395, 1991.

[44] A. Schofield and M. Van den Bergh. Semi-invariants of quivers for arbitrary dimension vectors. *Indagationes Mathematicae*, 12(1):125–138, 2001.

[45] A. Schrijver. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer Science & Business Media, 2003.

[46] V. Sergeichuk. Canonial matrices for linear matrix problems. *Adv. in Math.*, 50(3):189–265, 1983.

[47] V. Strassen. Vermeidung von divisionen. *Journal für die reine und angewandte Mathematik*, 264:184–202, 1973.

[48] B. Sturmfels. On the Newton polytope of the resultant. *Journal of Algebraic Combinatorics*, 3(2):207–236, 1994.

[49] L. G. Valiant. Completeness classes in algebra. In *Proceedings of the eleventh annual ACM symposium on Theory of computing*, STOC '79, pages 249–261, New York, NY, USA, 1979. ACM.