Electronic Colloquium on Computational Complexity, Report No. 140 (2019)

SEARCH PROBLEMS IN ALGEBRAIC COMPLEXITY, GCT, AND HARDNESS OF GENERATOR FOR INVARIANT RINGS.

ANKIT GARG, VISU MAKAM, RAFAEL OLIVEIRA, AND AVI WIGDERSON

Abstract. We consider the problem of outputting succinct encodings of lists of generators for invariant rings. Mulmuley conjectured that there are always polynomial sized such encodings for all invariant rings. We provide simple examples that disprove this conjecture (under standard complexity assumptions).

1. Introduction

In complexity theory, one often encounters search problems of the following nature: find a set of efficiently computable functions $\mathcal{F}$ which possess a desirable property $\mathcal{P}$. Some instances of such problems appear in derandomization, such as the search of pseudorandom generators; in proof complexity, such as the search of polynomials certifying the unsolvability of a system of polynomial equations; and in the efficient construction of seeded extractors, where the search is for a set of functions with the property that for any large enough min-entropy source as the input, the output of a large fraction of the functions is close to being uniformly distributed. Once we are faced with such problems, two natural questions are: how do we represent property $\mathcal{P}$? How do we encode the functions $\mathcal{F}$, or the outputs of such functions $\mathcal{F}$? Answers to such questions can draw important connections among many areas of mathematics and computer science.

In algebraic complexity, there are several problems which fit the description above, and we term them *algebraic search problems*. An algebraic search problem asks for a collection of polynomials with a certain desired property. Let us illustrate this with an example from algebraic proof complexity: in Nullstellensatz-based proof systems, one is given a set of multivariate polynomials $g_1, \ldots, g_r$ over an algebraically closed field $\mathbb{F}$ and variables $\mathbf{x} = (x_1, \ldots, x_n)$, and one wants to decide whether the system of polynomials $g_1(\mathbf{x}) = g_2(\mathbf{x}) = \cdots = g_r(\mathbf{x}) = 0$ has a solution over $\mathbb{F}$. A fundamental result of Hilbert tells us that the system has no solution if and only if there is a set of polynomials $\mathcal{F} = \{f_i\}_{i=1}^r$ such that $\sum_i f_i g_i = 1$. Thus, to prove that the set of polynomials $g_1, \ldots, g_r$ does not have a solution, we need to find a set $\mathcal{F}$ which has the property that $\sum_i f_i g_i = 1$ (and this corresponds to our property $\mathcal{P}$). If we encode the polynomials $g_1, \ldots, g_r$ as algebraic circuits and demand the polynomials $f_1, \ldots, f_r$ to be also given as algebraic circuits, we obtain the Ideal Proof System [GP14], which can be described as follows: we ask for an algebraic circuit C with $n+r$ inputs such that $C(x_1, \ldots, x_n, g_1, \ldots, g_r) = 1$ with the property that $C(x_1, \ldots, x_n, 0, 0, \ldots, 0) = 0$. In [GP14] the authors show that super polynomial lower bounds in this proof system implies arithmetic circuit lower bounds (i.e., VP $\neq$ VNP), which remains a long standing open problem in

Visu's research was supported by NSF grant No. DMS -1638352 and NSF grant No. CCF-1412958.
Part of this work was done while Rafael was visiting the Simons Institute for the Theory of Computing.
Avi's research was supported by NSF grant No. CCF-1412958.

ISSN 1433-8092

complexity theory. Other great examples in proof complexity (with different encodings of the polynomials in $\mathcal{F}$) are the original Nullstellensatz proof system, Polynomial Calculus, and Positivestellensatz[1] for SOS.

In this paper, we consider a new flavor of algebraic search problems arising from the Geometric Complexity Theory (GCT) program. These new algebraic search problems have deep roots in invariant theory and algebraic geometry (see [Mul17]), and important connections and consequences to problems in optimization, algebraic complexity, non-commutative computation, functional analysis and quantum information theory (see [GGOW16, GGOW18, BFG$^+$18]). We believe that a better understanding of these search problems will likely result in fundamental advances in the aforementioned areas.

The GCT program was proposed by Mulmuley and Sohoni (see [MS01, MS08]) as an approach (via representation theory and algebraic geometry) to the VP vs. VNP problem (a central problem in arithmetic complexity which is a natural analog of the celebrated P vs. NP problem). Despite some negative results[2], the GCT program has found a new lease of life with the new set of alluring conjectures arising from the algebraic search problems in computational invariant theory. Further progress was made in providing evidence for these conjectures, by establishing them for special cases (see for example, [Mul17, GGOW16, IQS18, FS13, DM17b, DM18, DM17a]).

An important algebraic search problem in computational invariant theory is as follows: given a group action, describe a set of generators for the invariant ring. Unfortunately, the number of generators for an invariant ring is usually exponential. So, in order to get a computational handle on them, Mulmuley suggests in [Mul17] that we should look for a *succinct encoding* (defined below in Definition 1.2) using some auxiliary variables. One amazing feature of such a succinct encoding is that it would give efficient (randomized) algorithms for null cone membership and the orbit closure intersection problems. We will define these problems later, but here we are content to say that many important algorithmic problems such as graph isomorphism, bipartite matching, (non-commutative) rational identity testing, tensor scaling and quantum distillation are all specific instances (or arise in the study) of null cone membership and orbit closure intersection problems.

Mulmuley conjectures ([Mul17, Conjecture 5.3]) that there is always a polynomial sized succinct encoding for generators of invariant rings. The main goal of this paper is to (conditionally) disprove this conjecture. More precisely we give an example of an invariant ring (for a torus action) where the existence of such a circuit would imply a polynomial time algorithm for the 3D-matching problem, which is well known to be NP-hard. We also give another example (i.e., a tensor action) which makes it clear that no simple modification of the generators conjecture can hold.

Remark 1.1. Our results are not entirely detrimental to the core of the GCT program. The takeaway should really be that we need to consider the so called *separating sets* of invariants rather than generating sets of invariants (Mulmuley already suggests that this is sufficient). Unfortunately, separating sets have not been quite as well studied in invariant theory, and

[1] In this case our field is the real numbers, which is *not* algebraically closed

[2] For example, the results of Burgisser, Ikenmeyer and Panova, which show that occurrence obstructions cannot give a super polynomial lower bound on the determinantal complexity of the permanent polynomial (see [BIP19]).

our results give motivation for both invariant theorists and complexity theorists to study them further.

We will now give a brief introduction to invariant theory.

1.1. **Invariant Theory.** Invariant theory is the study of group actions on vector spaces (more generally algebraic varieties) and the functions (usually polynomials) that are left *invariant* under these actions. It is a rich mathematical field in which computational methods are sought and well developed (see [DK15, Stu08]). While significant advances have been made in computational problems involving invariant theory, most algorithms still require exponential time (or longer).

The basic setting is that of a continuous group [3] G acting (linearly) on a finite-dimensional vector space $V = \mathbb{C}^m$.

An *action* (also called a *representation*) of a group G on an m-dimensional complex vector space V is a group homomorphism $\pi : G \to \mathrm{GL}(m)$, that is, an association of an invertible $m \times m$ matrix $\pi(g)$ for every group element $g \in G$ satisfying $\pi(g_1 g_2) = \pi(g_1)\pi(g_2)$ for all $g_1, g_2 \in G$. To be precise, a group element $g \in G$ acts on a vector $v \in V$ by the linear transformation $\pi(g)$. We will write $g \cdot v = \pi(g)v$. Invariant theory is nicest when the underlying field is $\mathbb{C}$ and the group G is either finite, the general linear group $\mathrm{GL}_n(\mathbb{C})$, the special linear group $\mathrm{SL}_n(\mathbb{C})$, or a direct product of these groups and their diagonal subgroups.

Invariant Polynomials: Invariant polynomials are precisely those which cannot distinguish between a vector v and a translate of it by an element of the group, i.e., $g \cdot v$. In other words, a polynomial function f on V is called invariant if $f(g \cdot v) = f(v)$ for all $v \in V$ and $g \in G$. Equivalently, invariant polynomials are polynomial functions on V which are left invariant by the action of G[4].

Two simple and illustrative examples are

- The symmetric group $G = \mathcal{S}_n$ acts on $V = \mathbb{C}^n$ by permuting the coordinates. In this case, the invariant polynomials are *symmetric* polynomials, and the n elementary symmetric polynomials form a generating set (a result that dates back to Newton).
- The group $G = \mathrm{SL}_n(\mathbb{C}) \times \mathrm{SL}_n(\mathbb{C})$ acts on $V = \mathrm{M}_n(\mathbb{C})$ by a change of bases of the rows and columns, namely left-right multiplication: that is, the action of (A, B) sends X to $AXB^\dagger$. Here, $\det(X)$ is an invariant polynomial and in fact every invariant polynomial must be a univariate polynomial in $\det(X)$. In other words, $\det(X)$ generates the invariant ring.

The above phenomenon that the ring of invariant of polynomials (denoted by $\mathbb{C}[V]^G$) is generated by a finite number of invariant polynomials is not a coincidence. The *finite generation theorem* due to Hilbert [Hil90, Hil93] states that, for a large class of groups (including the groups mentioned above), the invariant ring must be finitely generated. These two papers of Hilbert are highly influential and laid the foundations of commutative algebra. In particular, "finite basis theorem" and "Nullstellansatz" were proved as "lemmas" on the way towards proving the finite generation theorem!

[3] In general, the theory works whenever the group is connected, algebraic and reductive. However in this paper, we will deal with very simple groups.

[4] The action of G on V gives an induced action of G on the space of polynomial functions on V. For a polynomial function p on V, the group element $g \in G$ sends it to the function $g \cdot p$ which is defined by the formula $(g \cdot p)(v) = p(g^{-1} \cdot v)$ for $v \in V$. A polynomial function is called invariant if $g \cdot p = p$ for all $g \in G$ and this is what it means to be left invariant by the action of G.

Orbits and Orbit-Closures: The *orbit* of a vector $v \in V$, denoted by $\mathcal{O}_v$, is the set of all vectors obtained by the action of G on v. The *orbit-closure* of v, denoted by $\overline{\mathcal{O}}_v$, is the closure (under the Euclidean topology[5]) of the orbit $\mathcal{O}_v$. For actions of continuous groups (like $\mathrm{GL}_n(\mathbb{C})$), it is more natural to look at orbit-closures. The *null cone* for a group action is the set of all vectors which behave like the 0 vector i.e. the 0 vector lies in their orbit-closure. Many fundamental problems in theoretical computer science (and many more across mathematics) can be phrased as questions about orbits and orbit-closures. Here are some familiar examples:

- Graph isomorphism problem can be phrased as checking if the orbits of two graphs are the same or not, under the action of the symmetric group permuting the vertices.
- Geometric complexity theory (GCT) [MS01] formulates a variant of VP vs. VNP question as checking if the (padded) permanent lies in the orbit-closure of the determinant (of an appropriate size), under the action of the general linear group on polynomials induced by its natural linear action on the variables.
- Border rank (a variant of tensor rank) of a 3-tensor can be formulated as the minimum dimension such that the (padded) tensor lies in the orbit-closure of the unit tensor, under the natural action of $\mathrm{GL}_r(\mathbb{C}) \times \mathrm{GL}_r(\mathbb{C}) \times \mathrm{GL}_r(\mathbb{C})$. In particular, this captures the complexity of matrix multiplication.

1.2. **Computational invariant theory.** It is natural to ask about the computational complexity of various invariant theoretic problems like,

(1) (**Generators**) Output a list of generating polynomials for the invariant ring.
(2) (**Orbit-closure intersection**) Given two elements of the vector space, do their orbit-closures intersect?
(3) (**Null cone membership**) Given an element of the vector space, does the 0 vector lies inside its orbit-closure?

We won't get into the details of how the group is given and how the group action is described. It turns out that even for simple groups and group actions, the problems turn out to be interesting. These problems have been long studied and many algorithms have been developed in the invariant theory community [DK15, Stu08]. Mulmuley [Mul17] introduced these problems to theoretical computer science with the hope of making progress on the polynomial identity testing (PIT) problem. Before describing the main conjectures in Mulmuley's paper, let us see what it even means to output a list of generating polynomials for an invariant ring. Typically the number of generating polynomials can be exponential in the dimension of the group and the vector space. To get around this issue, Mulmuley introduced the following notion of a *succinct* encoding of the generators of an invariant ring (which in fact applies to any collection of polynomials).

Definition 1.2 (**Succinct encoding of generators**)**.** Fix an action of a group G on a vector space $V = \mathbb{C}^m$. We say that an arithmetic circuit $\mathcal{C}(x_1, \ldots, x_m, y_1, \ldots, y_r)$ succinctly encodes the generators of the invariant ring if the collection of polynomials formed by evaluating the y-variables, $\{\mathcal{C}(x_1, \ldots, x_m, \alpha_1, \ldots, \alpha_r)\}_{\alpha_1,\ldots,\alpha_r \in \mathbb{C}}$ are all invariant polynomials which form a generating set for the invariant ring.

[5]It turns out mathematically, it is more natural to look at closure under the Zariski topology. However, for the group actions we study, the Euclidean and Zariski closures match by a theorem due to Mumford [Mum65].

Remark 1.3. In the definition above, the size of the succinct encoding is given by the size of the circuit $\mathcal{C}(x_1, \dots, x_m, y_1, \dots, y_r)$, which is measured by the bit complexity of the constants used in the computation of $\mathcal{C}$ as well as the number of gates of the computation graph of $\mathcal{C}$.

Conjecture 1.4 (Mulmuley)**.** *Suppose G is a connected reductive algebraic group*[6] *(over $\mathbb{C}$) of* $\dim n$*, and suppose it acts algebraically on an m-dimensional vector space V by linear transformations. Then there is a* $\mathrm{poly}(m, n)$ *sized succinct encoding for some set of generators for the invariant ring.*

Once again, note that in the above, the succinct encoding being polynomial sized (in n, m) simply means that the number of auxiliary variables, the size of the circuit and the total bitsize of all the rational constants used in the circuit are all polynomial in n, m. Mulmuley requires the circuit family to be uniformly computable by a polynomial time algorithm, but we will see that even this weaker conjecture is false (under standard complexity assumptions).

To understand Mulmuley's motivation for the conjecture, let us see what it means for the problems of orbit-closure intersection and null cone membership. By definition, invariant polynomials are constant on the orbits (and thus on orbit-closures as well). Thus, if $\overline{\mathcal{O}}_{v_1} \cap \overline{\mathcal{O}}_{v_2} \neq \emptyset$, then $p(v_1) = p(v_2)$ for all invariant polynomials $p \in \mathbb{C}[V]^G$. A remarkable theorem due to Mumford says that the converse is also true (for a large class of groups including the ones we discussed above).

Theorem 1.5 ([Mum65])**.** *Fix an action of a group G on a vector space V. Given two vectors $v_1, v_2 \in V$, we have $\overline{\mathcal{O}}_{v_1} \cap \overline{\mathcal{O}}_{v_2} \neq \emptyset$ if and only if $p(v_1) = p(v_2)$ for all $p \in \mathbb{C}[V]^G$.*

Now suppose one had a succinct encoding $\mathcal{C}(x_1, \dots, x_m, y_1, \dots, y_r)$ for action of a group G on $V = \mathbb{C}^m$. Due to Mumford's theorem, the orbit-closures of vectors v_1 and v_2 intersect iff the two polynomials $\mathcal{C}(v_1(1), \dots, v_1(m), y_1, \dots, y_r)$, $\mathcal{C}(v_2(1), \dots, v_2(m), y_1, \dots, y_r)$ are identically the same. These are instances of polynomial identity testing (PIT)! Thus if Conjecture 1.4 were true (and additionally the succinct encoding circuits were polynomial time computable), it immediately gives randomized polynomial time algorithms for the orbit-closure intersection and null cone membership problems. And this also gives us a nice family of PIT problems to play with. Perhaps the hope is that solving these PIT instances will result in development of new techniques which might shed a light on the general PIT problem. In fact, for the first few group actions that were studied in this line of work, *simultaneous conjugation* [Mul17, FS13] and *left-right action* [GGOW16, IQS18, DM17b], for which there are polynomial sized succinct encodings of generators, the null cone membership problems correspond to PIT problems for restricted models of computation: *read-once algebraic branching programs* and *non-commutative formulas with division*.[7]

Unfortunately, we prove that Conjecture 1.4 is false (under standard complexity assumptions). We want to emphasize that this only serves a first guiding light for Mulmuley's program of understanding the orbit-closure intersection problems and connections to PIT. To solve the orbit-closure intersection problems, one does not necessarily need the whole list of generators. One only needs a list of *separating invariants* which are collection S of invariant polynomials s.t. for any two vectors $v_1, v_2 \in V$, $\overline{\mathcal{O}}_{v_1} \cap \overline{\mathcal{O}}_{v_2} \neq \emptyset$ iff $p(v_1) = p(v_2)$ for

[6]We have not defined what a connected reductive algebraic group is. One should think of simple groups like the general linear group $\mathrm{GL}_n(\mathbb{C})$, the special linear group $\mathrm{SL}_n(\mathbb{C})$, or a direct product of these groups and their diagonal subgroups.

[7]Actually a stronger model concerning inverses of matrices.

all $p \in S$. It is possible that Conjecture 1.4 is true when the generating set is replaced by a set of separating invariants.

For our first counterexample, we analyze a simple (torus) action on 3-tensors. ST_n denotes the group of $n \times n$ diagonal matrices with determinant 1.

Theorem 1.6. *Consider the natural action of $G = \mathrm{ST}_n \times \mathrm{ST}_n \times \mathrm{ST}_n$ on $V = \mathbb{C}^n \otimes \mathbb{C}^n \otimes \mathbb{C}^n$. Then any set of generators for the invariant ring cannot have a polynomial sized (in n) succinct encoding, unless* $\mathrm{NP} \subseteq \mathrm{P/poly}$.

Clearly the above result disproves Conjecture 1.4.

Corollary 1.7. *Conjecture 1.4 is false, unless* $\mathrm{NP} \subseteq \mathrm{P/poly}$.

We give another counterexample. This counterexample is intended to illustrate that any obvious modification of Conjecture 1.4 would likely fail as well.

Theorem 1.8. *Consider the natural action of $G = \mathrm{SL}_n \times \mathrm{SL}_n \times \mathrm{SL}_n$ on $V = \mathbb{C}^n \otimes \mathbb{C}^n \otimes \mathbb{C}^n \otimes \mathbb{C}^n$. Then*

(1) *Homogenous invariants only exist in degrees that are multiples of n.*
(2) *There is exactly one homogenous invariant (upto scalar multiples) of degree n (the lowest possible degree).*
(3) *The unique homogenous invariant of degree n is* VNP*–hard.*

Moreover, as a consequence of the above, any set of generators for the invariant ring cannot have a polynomial sized (in n) succinct encoding, assuming $\mathrm{VP} \neq \mathrm{VNP}$.

1.3. **Conclusion, open problems and future directions.** We have disproved a conjecture of Mulmuley about the existence of polynomial sized succinct encodings of generators for invariant rings. This serves only as a guiding light for what the right conjecture should be. For example, it is possible that there are polynomial sized succinct encodings for sets of separating invariants. We mention below some of the interesting open problems and future directions.

(1) Are there polynomial sized succinct encodings for separating invariants or, even simpler, invariants defining the null cone? Perhaps the first non-trivial example is the natural action of $G = \mathrm{ST}_n \times \mathrm{ST}_n \times \mathrm{ST}_n$ on $V = \mathbb{C}^n \otimes \mathbb{C}^n \otimes \mathbb{C}^n$. Here a tensor T is in the null cone iff there exists vectors $x, y, z \in \mathbb{R}^n$ s.t. $x_i + y_j + z_k > 0$ for all $(i, j, k) \in \mathrm{supp}(T)$[8] and $\sum_i x_i = \sum_j y_j = \sum_k z_k = 0$ (by the Hilbert-Mumford criterion). Is there a polynomial sized circuit $\mathbb{C}((z_{i,j,k}), y_1, \ldots, y_r)$ s.t. $\mathbb{C}(T, y_1, \ldots, y_r)$ is identically zero (as a polynomial in the y-variables) iff T is in the null cone?
(2) For the natural action of $\mathrm{SL}_n \times \mathrm{SL}_n \times \mathrm{SL}_n$ on $V = \mathbb{C}^n \otimes \mathbb{C}^n \otimes \mathbb{C}^n$, it is not even clear if there exists one invariant which has a polynomial sized circuit. Either produce such an invariant or prove that all invariants are hard to compute.
(3) Are there polynomial time algorithms for the orbit-closure intersection and null cone membership problems? The analytic approach pursued in the papers [GGOW16, BGO+17, AZGL+18] seems the most promising approach towards getting such algorithms.

[8] $\mathrm{supp}(T) = \{(i, j, k) \in [n] \times [n] \times [n] : T_{i,j,k} \neq 0\}$.

(4) More broadly, invariant theory is begging for its own complexity theory and connecting it with ours. This includes finding reductions and completeness results, and characterizations/dichotomies about hard/easy actions. An example of a completeness reduction is the reduction from all quiver actions to the simple left-right action [DW00, DZ01, SVdB01, DM17b, GGOW18]. Also the papers [Mul17, GGOW16, IQS18, FS13, DM17b, DM18, DM17a], as well as the current paper, are trying to identify easy and hard problems in invariant theory.

2. Preliminaries

In this section we establish notation and we formally state basic facts and definitions which we will need in later sections.

Definition 2.1 (3-dimensional matching [Kar72])**.** The *3-dimensional matching* problem is defined as follows:

Input: a set $U \subseteq [n] \times [n] \times [n]$, representing the edges of a tripartite, 3-uniform hypergraph.
Output: YES, if there is a set of hyperedges $W \subseteq U$ such that $|W| = n$ and no two elements of W agree in any coordinate (that is, they form a matching in this hypergraph). NO, if there is no such set.

Theorem 2.2 (NP-completeness of 3-dimensional matching [Kar72])**.** *The 3-dimensional matching problem is NP-complete.*

2.1. Basic facts from algebraic complexity.

We now give basic facts that from algebraic complexity which we will use in the next sections.

The next proposition shows that homogeneous components of low degree of an arithmetic circuit can be efficiently computed, with a small blow-up in circuit size and without the use of any extra constants. This proposition was originally proved by Strassen in [Str73] and its proof can be found in [SY$^+$10, Theorem 2.2].

Proposition 2.3 (Efficient computation of homogeneous components)**.** *Given a circuit $\mathcal{C}(\boldsymbol{x})$ of size s, then for every $r \in \mathbb{N}$ there is a homogeneous circuit $\Psi(\boldsymbol{x})$ of size $O(r^2 s)$ computing $H_0[\mathcal{C}(\boldsymbol{x})], H_1[\mathcal{C}(\boldsymbol{x})], \ldots, H_r[\mathcal{C}(\boldsymbol{x})]$. Moreover, the constants used in the computation of the components $H_i[\mathcal{C}(\boldsymbol{x})]$ are a subset of the coefficients used in the computation of $\mathcal{C}(\boldsymbol{x})$.*

The next theorem, proved by [AB03, Theorem 4.10] gives us a randomized polynomial time algorithm to test whether an algebraic circuit of polynomial size, with rational coefficients, is identically zero. Another randomized algorithm easily follows from [Sch79, Lemma 2], when adapted for polynomials with rational coefficients.

Theorem 2.4 (PIT for poly-sized circuits [AB03])**.** *Let $P(\boldsymbol{x}) \in \mathbb{Q}[\boldsymbol{x}]$ be a polynomial over the variables $\boldsymbol{x} = (x_1, \ldots, x_n)$, with each variable x_i having degree bounded by d_i, and whose coefficients are rationals with bit complexity bounded by B. If $P(\boldsymbol{x})$ is given as an arithmetic circuit of size s, then there exists a randomized algorithm running in time $\mathrm{poly}(n, s, \log(B), 1/\epsilon)$ and using $O\left(\sum_{i=1}^n \log(d_i) + \log(B)\right)$ random bits which tests whether $P(\boldsymbol{x})$ is the identically zero polynomial and errs with probability at most ϵ if $P(\boldsymbol{x})$ is not the identically zero polynomial.*

3. Hardness of GENERATORS for torus actions

Let $\mathbb{C}^*$ denote the multiplicative group consisting of all non-zero complex numbers. A direct product $\mathrm{T} = (\mathbb{C}^*)^n$ is called a torus, and is clearly an abelian group. Tori are important examples of reductive groups – any abelian connected reductive group is a torus! It is often the case that it is easier to understand tori in comparison with more general (non-abelian) reductive groups. Invariant theory is no different, see for example [DK15, Weh93]. We also point to [DM19, Proposition 3.3] for an elementary linear algebraic description of the invariant ring for torus actions. Conjecture 1.4 already fails in this well behaved setting. In this section, we will prove Theorem 1.6, which we restate below for convenience.

Theorem 3.1. *Consider the natural action of $G = \mathrm{ST}_n \times \mathrm{ST}_n \times \mathrm{ST}_n$ on $V = \mathbb{C}^n \otimes \mathbb{C}^n \otimes \mathbb{C}^n$, where an element $(\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}) \in G$ acts on a tensor $u \in V$ as follows: $((\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}), u) \mapsto v$ such that $v_{ijk} = a_i b_j c_k \cdot u_{ijk}$. Any set of generators for the invariant ring of this action cannot have a polynomial sized (in n) succinct encoding, unless* NP $\subseteq$ P/poly*.*

Proof. Suppose that the natural action above has a set of generators with a polynomial sized succinct encoding. Thus, there is an arithmetic circuit $\mathcal{C}(\mathbf{x}, \mathbf{y})$ of size $s = \mathrm{poly}(n)$, where $\mathbf{x} = (x_{ijk})_{i,j,k=1}^n$ is the set of variables corresponding to V and $\mathbf{y} = (y_1, \dots, y_r)$ is the set of auxiliary variables, with $r = \mathrm{poly}(n)$. Moreover, from the definition we also have that the constants used in the computation of $\mathcal{C}(\mathbf{x}, \mathbf{y})$ are rational numbers with bit complexity bounded by $b = \mathrm{poly}(n)$. Therefore, $\mathcal{C}(\mathbf{x}, \mathbf{y}) \in \mathbb{Q}[\mathbf{x}, \mathbf{y}]$.

By considering the circuit $\mathcal{C}(\mathbf{x}, \mathbf{y})$ as a circuit whose constants are in $\mathbb{Q}[\mathbf{y}]$ and whose variables are only the $\mathbf{x}$ variables, that is, a circuit in $\mathbb{Q}[\mathbf{y}][\mathbf{x}]$, Proposition 2.3 tells us that there exists a homogeneous circuit $\mathcal{C}_n(\mathbf{x}, \mathbf{y})$, over the $\mathbf{x}$ variables, of degree n, size $O(n^2 s)$ whose constants are a subset of the constants[9] used in the circuit $\mathcal{C}(\mathbf{x}, \mathbf{y})$. In particular, $\mathcal{C}_n(\mathbf{x}, \mathbf{y})$ can be written in the following way:

$$\mathcal{C}_n(\mathbf{x}, \mathbf{y}) = \sum_{\mathbf{a} \in \mathcal{N}_n(\mathbf{x})} f_{\mathbf{a}}(\mathbf{y}) \cdot \mathbf{x}^{\mathbf{a}}, \tag{1}$$

where $\mathcal{N}_n(\mathbf{x})$ is the set of all monomials of degree n over the variables $\mathbf{x}$ and $f_{\mathbf{a}}(\mathbf{y})$ are polynomials on the variables $\mathbf{y}$ of degree at most 2^s, as the circuit $\mathcal{C}$ has size at most s.

In Proposition 3.2 below, we will show that the invariants of minimum degree of our action are in degree n, and these are spanned by the (maximum) 3-dimensional matching monomials. Thus, if a monomial of degree n is invariant under our action, it must be the case that this monomial corresponds to a 3-dimensional matching. As $\mathcal{C}_n(\mathbf{x}, \mathbf{y})$ must only compute invariant polynomials, this in turn implies that equation (1) is actually of the following form:

$$\mathcal{C}_n(\mathbf{x}, \mathbf{y}) = \sum_{\mathbf{a} \in \mathcal{M}_n(\mathbf{x})} f_{\mathbf{a}}(\mathbf{y}) \cdot \mathbf{x}^{\mathbf{a}}, \tag{2}$$

where $\mathcal{M}_n(\mathbf{x})$ is the set of all 3-dimensional matching monomials over the variables $\mathbf{x}$.

We will now show that the existence of the circuit $\mathcal{C}_n(\mathbf{x}, \mathbf{y})$ implies that NP $\in$ P/poly. For that purpose, we will show that given $\mathcal{C}_n(\mathbf{x}, \mathbf{y})$ one can solve the 3-dimensional matching problem in P/poly. Let H be a tripartite 3-uniform hypergraph, whose edges are given by a subset $E \subseteq [n] \times [n] \times [n]$. We can associate to this graph the tensor $v \in V$ where $v_{ijk} = 1$ if hyperedge $(i, j, k) \in E$ and $v_{ijk} = 0$ otherwise. Note that H has a 3-dimensional matching of size n if and only if at least one of the 3-dimensional matching monomials *does not vanish* on our tensor v. This last condition is equivalent to the fact that the circuit $\mathcal{C}_n(v, \mathbf{y})$ does

[9]the constants of $\mathcal{C}$ in this case are given by the elements of $\mathbb{Q}$ used in the computation of $\mathcal{C}$ and the auxiliary variables $\mathbf{y}$.

not compute the zero polynomial (as we know that the span of the set $\{\mathcal{C}_n(\mathbf{x}, \alpha)\}_{\alpha \in \mathbb{C}^r}$ is the same as the span of all 3-dimensional matching monomials). Thus, to solve the 3-dimensional matching problem in P/poly it is enough to give a randomized polynomial time algorithm for testing whether $\mathcal{C}_n(v, \mathbf{y})$ is the zero polynomial or not.[10]

Since $\mathcal{C}_n(v, \mathbf{y})$ is a circuit of $\mathrm{poly}(n)$ size with rational constants of bit complexity $\mathrm{poly}(n)$, it computes a polynomial $P(\mathbf{y})$ with rational coefficients having bit complexity at most $\exp(n)$ and degree at most $\exp(n)$. This is the setting in which Theorem 2.4 applies, giving us the desired randomized polynomial time algorithm. This concludes our proof modulo Proposition 3.2, which we will now turn our attention to. $\square$

In the following proposition, we denote the symmetric group on n letters by $\mathcal{S}_n$.

Proposition 3.2. *The maximum 3-dimensional matching monomials $\prod_{i=1}^n x_{i\sigma(i)\tau(i)}$, where $\sigma, \tau \in \mathcal{S}_n$, span the invariants of degree n of the natural action of $G = \mathrm{ST}_n \times \mathrm{ST}_n \times \mathrm{ST}_n$ on $V = \mathbb{C}^n \otimes \mathbb{C}^n \otimes \mathbb{C}^n$. Moreover, there are no invariants of degree less than n for this action.*

Proof. Since each monomial is mapped to a constant times the monomial, it is easy to see that a subset of monomials generate the invariants. Thus, to prove the above proposition, it is enough to prove that the invariants of minimum degree are of degree n, and that these invariants are spanned by the matching monomials. Let us begin by proving that the invariants of degree n are spanned by the matching monomials.

Note that the natural action of G on V induces the following action on a variable x_{ijk}: $(\mathbf{a}, \mathbf{b}, \mathbf{c}) \cdot x_{ijk} = (a_i b_j c_k)^{-1} \cdot x_{ijk}$. Additionally, note that $\prod_{\ell=1}^n a_\ell = \prod_{\ell=1}^n b_\ell = \prod_{\ell=1}^n c_\ell = 1$. Given a matching monomial $\prod_{i=1}^n x_{i\sigma(i)\tau(i)}$, we have that

$$\begin{aligned}(\mathbf{a}, \mathbf{b}, \mathbf{c}) \cdot \prod_{i=1}^n x_{i\sigma(i)\tau(i)} &= \prod_{i=1}^n \left((a_i b_{\sigma(i)} c_{\tau(i)})^{-1} \cdot x_{i\sigma(i)\tau(i)}\right) \\ &= \prod_{i=1}^n (a_i b_{\sigma(i)} c_{\tau(i)})^{-1} \cdot \prod_{i=1}^n x_{i\sigma(i)\tau(i)} \\ &= \prod_{i=1}^n x_{i\sigma(i)\tau(i)}\end{aligned}$$

where in the last equality we note that for any permutation $\sigma \in \mathcal{S}_n$ (or τ) we have $1 = \prod_{\ell=1}^n a_\ell = \prod_{\ell=1}^n a_{\sigma(\ell)}$ (and similarly for $\mathbf{b}$ and $\mathbf{c}$). This proves that the matching monomials are invariant monomials of the natural G-action on V.

Now, let us prove that no non-matching monomial of degree n is an invariant for this action. Let $\prod_{m=1}^n x_{i_m j_m k_m}$ be a non-matching monomial, where $(i_m, j_m, k_m) \in [n]^3$. This implies that there exists some coordinate, say for instance the first coordinates, for which the set $\{i_m\}_{m=1}^n$ is strictly contained in $[n]$. Equivalently, there is an element $\ell \in [n]$ such that $\ell \notin \{i_m\}_{m=1}^n$. W.l.o.g., we can assume that $\ell = 1$. Thus, the action of $\mathbf{a} = (\alpha^{n-1}, \alpha^{-1}, \ldots, \alpha^{-1}), \mathbf{b} = \mathbf{c} = (1, \ldots, 1)$ on our monomial $\prod_{m=1}^n x_{i_m j_m k_m}$ is as follows:

$$(\mathbf{a}, \mathbf{b}, \mathbf{c}) \cdot \prod_{m=1}^n x_{i_m j_m k_m} = \alpha^{n-1} \cdot \prod_{m=1}^n x_{i_m j_m k_m}$$

[10] It is enough to give a randomized polynomial time algorithm because we know that BPP/poly = P/poly.

which proves that this monomial is not an invariant. This completes the proof that the matching monomials span the invariants of degree n.

Now we are left with proving that there are no invariants of degree less than n. Note that if we have a monomial with degree less than n, we can represent it as $\prod_{m=1}^{d} x_{i_m j_m k_m}$, where $d < n$ and by the pigeonhole principle, we know that there exists $\ell \in [n]$ such that ℓ does not appear as a first coordinate entry in the set of tuples $\{(i_m, j_m, k_m)\}$. Thus, analogously to the previous paragraph, we know that such monomials cannot be invariants of the natural action of G over V, therefore showing that no monomial of degree $< n$ can be an invariant. This completes the proof. $\square$

4. Invariant theory for tensor actions

In this section, we will give yet another example of an group action on tensors for which any set of generating invariants are hard to compute, i.e., we will prove Theorem 1.8. Even though the previous section suffices to disprove Mulmuley's conjecture, this example illustrates something more. The feature of this group action is that the smallest degree invariants span a 1-dimensional space. In other words, upto scaling, we have a unique invariant of smallest degree. This unique invariant in the smallest degree is well known and goes by Pascal determinant or quantum permanent, which is already known to VNP–hard. The importance of this example is that such a unique invariant in the smallest degree is *essential* in any generating set. So, it is not even possible to give a generating set consisting of invariant polynomials that are easy to compute, even if we remove all restrictions on the size of the generating set.

The main focus of this section will be to show the fact that this unique invariant in smallest degree is the Pascal determinant/quantum permanent. This fact is not new, it has been observed before, see [Lan15]. Here we simply intend to give a rather self-contained proof as an attempt to familiarize the computer science community with some of the combinatorial aspects of the representation theory of the general linear group and the special linear group.

Let $\{e_1, \dots, e_n\}$ denote the standard basis vectors for $\mathbb{C}^n$. Consider the tensor space $V = \mathbb{C}^n \otimes \mathbb{C}^n \otimes \mathbb{C}^n \otimes \mathbb{C}^n$. The standard basis for this tensor space is given by

$$\{e_{i_1,i_2,i_3,i_4} = e_{i_1} \otimes e_{i_2} \otimes e_{i_3} \otimes e_{i_4} \mid 1 \leq i_1, i_2, i_3, i_4 \leq n\}.$$

Let x_{i_1,i_2,i_3,i_4} denote the coordinate corresponding to e_{i_1,i_2,i_3,i_4}. Then the ring of polynomial functions on V is $\mathbb{C}[V] = \mathbb{C}[x_{i_1,i_2,i_3,i_4} \mid 1 \leq i_1, i_2, i_3, i_4 \leq n]$.

The action of $G = \mathrm{SL}_n \times \mathrm{SL}_n \times \mathrm{SL}_n$ on $\mathbb{C}^n \otimes \mathbb{C}^n \otimes \mathbb{C}^n \otimes \mathbb{C}^n$ is given explicitly by

$$(A, B, C) \cdot \sum_i v_a^{(i)} \otimes v_b^{(i)} \otimes v_c^{(i)} \otimes v_d^{(i)} = \sum_i A v_a^{(i)} \otimes B v_b^{(i)} \otimes C v_c^{(i)} \otimes v_d^{(i)}.$$

We want to focus on the invariants of degree n. To do so, we will consider $\mathbb{C}[V]_n$, the space of polynomial functions of degree n on V, which inherits an action of G. Then, the invariants of degree n as precisely the polynomials in $\mathbb{C}[V]_n$ on which the action of G is trivial.

We will need some tools from representation theory. While these tools are well known, they may be unfamiliar to some readers, so we will briefly recall them.

4.1. **Representation theory of SL_n.** An excellent introduction to this subject from a combinatorial perspective is Fulton's book ([Ful97]). A partition $\lambda = (\lambda_1, \dots, \lambda_k)$ of n is a (weakly) decreasing sequence of positive numbers $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_k > 0$. For a partition $\lambda = (\lambda_1, \dots, \lambda_k)$, we call the number of parts k the length of the partition. We write $\lambda \vdash n$ to denote that λ is a partition of n. To each partition, we associate a Young diagram with it, and this is best addressed by an example. To the partition $\lambda = (4, 2, 1)$, we associate the Young diagram

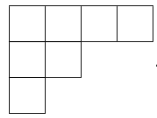

.

For a partition λ, we denote by $\lambda^\dagger$ the conjugate partition. In terms of Young diagrams, it corresponds to reflecting the diagram about $y = -x$. So, for $\lambda = (4, 2, 1)$, we have $\lambda^\dagger = (3, 2, 1, 1)$ whose Young diagram is

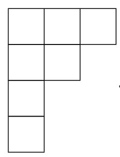

.

Suppose W is an n-dimensional vector space, and let $\mathrm{GL}(W)$ denote the general linear group. The group $\mathrm{GL}(W)$ can be identified with GL_n once you choose a basis, but we prefer to not choose a basis until really needed. The reason for this is that we will be working in $(\mathbb{C}^n)^{\otimes 4}$, and distinguishing the different tensor factors may be quite confusing.

Irreducible polynomial representations of $\mathrm{GL}(W)$ are indexed by partitions of length $\leq n$. To a partition λ (with $l(\lambda) \leq n$), we denote by $S_\lambda(W)$ the irreducible representation corresponding to λ. For $\lambda = (d)$, we have $S_\lambda(W) = \mathrm{Sym}^d(W)$, the d^{th} symmetric power of W, and for $\lambda = (1, 1, \dots, 1) = 1^d$, $S_\lambda(W) = \bigwedge^d(W)$, the d^{th} exterior power of W. For more general λ, $S_\lambda(W)$ has a description in terms of Young tableau of shape λ, see [Ful97] for details. Moreover, note that $S_\lambda(W^*) = S_\lambda(W)^*$ is the contragredient (or dual) representation of $S_\lambda(W)$. Further, any action of $\mathrm{GL}(W)$ can be restricted to the action of $\mathrm{SL}(W)$, so every $\mathrm{GL}(W)$ representation can be interpreted naturally as a $\mathrm{SL}(W)$ representation.

Let us explain in a bit more detail the symmetric and exterior powers. Suppose W is a vector space with basis $w_1, \dots, w_n$. Then, $\mathrm{Sym}^d(W)$ is the vector space consisting of homogenous degree d polynomials in $w_1, \dots, w_n$. The action of $\mathrm{GL}(W)$ is quite straightforward. For a monomial $w_1 \cdots w_d$, and a matrix $A \in \mathrm{GL}(W)$, we have $A \cdot (w_1 \cdots w_d) = Aw_1 \cdot Aw_2 \cdots Aw_d$. Extending the action by linearity gives the $\mathrm{GL}(W)$ action. In particular, observe that the polynomial ring $\mathbb{C}[w_1, \dots, w_n]$ can be seen as $\oplus_{d \in \mathbb{Z}_{\geq 0}} \mathrm{Sym}^d(W)$, and is often written as $\mathrm{Sym}(W)$ and called the symmetric algebra over W. Alternatively, $\mathrm{Sym}^d(W)$ can be defined as the subspace of tensors in $W^{\otimes d}$ that are symmetric. There is a natural action of $\mathcal{S}_d$ (the symmetric group on d letters) on $W^{\otimes d}$ by permuting tensor factors. A tensor in $W^{\otimes d}$ is called symmetric if it is invariant under the action of $\mathcal{S}_d$, i.e., permuting tensor factors doesn't change the tensor.

The exterior power $\bigwedge^d(W)$ on the other hand can be seen as the subspace of $W^{\otimes d}$ consisting of alternating tensors, i.e., if we swap two tensor factors, then the tensor changes by a sign. A basis for $\bigwedge^d(W)$ is given by $\{w_{i_1} \wedge w_{i_2} \wedge \cdots \wedge w_{i_d} \mid 1 \leq i_1 < i_2 < \cdots < i_d \leq n\}$, where

$$w_{i_1} \wedge w_{i_2} \wedge \cdots \wedge w_{i_d} = \sum_{\sigma \in \mathcal{S}_d} \operatorname{sgn}(\sigma) w_{i_{\sigma(1)}} \otimes w_{i_{\sigma(2)}} \otimes \cdots \otimes w_{i_{\sigma(d)}}.$$

Of particular interest is $\bigwedge^n(W)$, which is just the 1-dimensional representation described by the determinant $\det : \mathrm{GL}(W) = \mathrm{GL}_n \to \mathrm{GL}_1 = \mathbb{C}^*$. The reader should thus note that $\bigwedge^n(W)$ as an $\mathrm{SL}(W)$ representation is trivial.

The main formulas we need are Cauchy's formulas. The first is

$$\operatorname{Sym}^d(A \otimes B) = \bigoplus_{\lambda \vdash d} S_\lambda(A) \otimes S_\lambda(B), \tag{3}$$

and the second is

$$\bigwedge^d(A \otimes B) = \bigoplus_{\lambda \vdash d} S_\lambda(A) \otimes S_{\lambda^\dagger}(B). \tag{4}$$

Remark 4.1. In the above formulas, we have to explain what we mean by '=' between two representations. A priori we should say the two sides are isomorphic, so we should write $\cong$. However, we use $=$ instead of $\cong$ because these isomorphisms are canonical! In the rest of this section, whenever we write $=$ between two representations, we will mean that they are canonically isomorphic.

Now, suppose A, B, C, D are n-dimensional vector spaces. Then let us compute the invariants

$$\operatorname{Sym}^n(A \otimes B \otimes C \otimes D)^{\mathrm{SL}(A^*) \times \mathrm{SL}(B^*) \times \mathrm{SL}(C^*)}.$$

First, observe that $\mathrm{SL}(A^*) \times \mathrm{SL}(B^*) \times \mathrm{SL}(C^*)$ acts on $A^* \otimes B^* \otimes C^* \otimes D^* = (A \otimes B \otimes C \otimes D)^*$ in the natural way, and hence we get an induced action on $A \otimes B \otimes C \otimes D$ and hence on $\operatorname{Sym}^n(A \otimes B \otimes C \otimes D)$.

We need to know exactly one fact for this computation, i.e., for a partition $\lambda \vdash n$, the invariants

$$S_\lambda(A)^{\mathrm{SL}(A^*)} = \begin{cases} \bigwedge^n(A) & \text{if } \lambda = (1, 1, \ldots, 1); \\ 0 & \text{otherwise.} \end{cases}$$

Further $\bigwedge^n(A)$ is 1-dimensional. We omit the proof, but briefly outline the basic idea. Unless $\lambda = (1, 1, \ldots, 1)$, $S_\lambda(A)$ is a non-trivial irreducible representation of $\mathrm{SL}(A^*)$ and hence has no invariants. When $\lambda = (1, \ldots, 1)$, then $S_\lambda(A) = \bigwedge^n(A) = (\bigwedge^n(A^*))^*$ is the dual representation of $\bigwedge^n(A^*)$, which we have already seen is the trivial representation for $\mathrm{SL}(A^*)$.

Now, let us get to the heart of the computation. This will use repeatedly Cauchy's formulas and the fact above. We have

$$\begin{aligned} \operatorname{Sym}^n(A \otimes B \otimes C \otimes D)^{\mathrm{SL}(A^*)} &= \bigoplus_{\lambda \vdash n} S_\lambda(A)^{\mathrm{SL}(A^*)} \otimes S_\lambda(B \otimes C \otimes D) \\ &= \bigwedge^n(A) \otimes \bigwedge^n(B \otimes C \otimes D). \end{aligned}$$

Thus, we have

$$\begin{aligned}
\operatorname{Sym}^n(A\otimes B\otimes C\otimes D)^{\mathrm{SL}(A^*)\times \mathrm{SL}(B^*)} &= \bigwedge^n(A)\otimes \bigwedge^n(B\otimes C\otimes D)^{\mathrm{SL}(B^*)}\\
&= \bigwedge^n(A)\otimes \bigoplus_{\lambda\vdash n}\big(S_\lambda(B)^{\mathrm{SL}(B^*)}\otimes S_{\lambda^\dagger}(C\otimes D)\big)\\
&= \bigwedge^n(A)\otimes \bigwedge^n(B)\otimes \operatorname{Sym}^n(C\otimes D).
\end{aligned}$$

Hence, we get

$$\begin{aligned}
\operatorname{Sym}^n(A\otimes B\otimes C\otimes D)^{\mathrm{SL}(A^*)\times \mathrm{SL}(B^*)\times \mathrm{SL}(C^*)} &= \bigwedge^n(A)\otimes \bigwedge^n(B)\otimes \operatorname{Sym}^n(C\otimes D)^{\mathrm{SL}(C^*)}\\
&= \bigwedge^n(A)\otimes \bigwedge^n(B)\otimes \bigoplus_{\lambda\vdash n}\big(S_\lambda(C)^{\mathrm{SL}(C^*)}\otimes S_\lambda(D)\big)\\
&= \bigwedge^n(A)\otimes \bigwedge^n(B)\otimes \bigwedge^n(C)\otimes \bigwedge^n(D)
\end{aligned}$$

From the above, we understand there is a unique copy of $\bigwedge^n(A)\otimes\bigwedge^n(B)\otimes\bigwedge^n(C)\otimes\bigwedge^n(D)$ in $\operatorname{Sym}^n(A\otimes B\otimes C\otimes D)$ which corresponds to the $G=\mathrm{SL}(A^*)\times\mathrm{SL}(B^*)\times\mathrm{SL}(C^*)$ invariants. Moreover, this is 1-dimensional. So, let us explicitly understand the inclusion

$$\bigwedge^n(A)\otimes \bigwedge^n(B)\otimes \bigwedge^n(C)\otimes \bigwedge^n(D)\hookrightarrow \operatorname{Sym}^n(A\otimes B\otimes C\otimes D).$$

First, we observe that $\mathcal{P}=\bigwedge^n(A)\otimes\bigwedge^n(B)\otimes\bigwedge^n(C)\otimes\bigwedge^n(D)$ occurs in $\operatorname{Sym}^n(A\otimes B\otimes C\otimes D)$ with multiplicity 1 as a representation of $\mathcal{G}=\mathrm{GL}(A)\times\mathrm{GL}(B)\times\mathrm{GL}(C)\times\mathrm{GL}(D)$. Thus, if we give any non-zero $\mathcal{G}$-equivariant map from $\mathcal{P}\hookrightarrow\operatorname{Sym}^n(A\otimes B\otimes C\otimes D)$, then the image will be 1-dimensional, and hence precisely G-invariants.

To get such a non-zero map, we compose two maps

$$\mathcal{P}\hookrightarrow (A\otimes B\otimes C\otimes D)^{\otimes n}\twoheadrightarrow \operatorname{Sym}^n(A\otimes B\otimes C\otimes D).$$

The second map is clear. Given any vector space W, there is a surjection $W^{\otimes n}\to\operatorname{Sym}^n(W)$ that is defined by sending $w_1\otimes\cdots\otimes w_n\mapsto w_1\dots w_n$. The first map on the other hand is a little bit more interesting. For any n-dimensional vector space W, we have an $\mathrm{SL}(W)$-equivariant map $\bigwedge^n W\hookrightarrow W^{\otimes n}$ where $w_1\wedge w_2\wedge\cdots\wedge w_n\mapsto\sum_{\sigma\in\mathcal{S}_n}\operatorname{sgn}(\sigma)w_{\sigma(1)}\otimes\cdots\otimes w_{\sigma(n)}$. Taking this map for $W=A,B,C,D$, and tensoring them together gives us the map

$$\mathcal{P}\to A^{\otimes n}\otimes B^{\otimes n}\otimes C^{\otimes n}\otimes D^{\otimes n}=(A\otimes B\otimes C\otimes D)^{\otimes n}$$

A quick computation shows that the image of this map is spanned by

$$\sum_{\sigma,\mu,\pi,\nu\in\mathcal{S}_n}\operatorname{sgn}(\sigma\mu\pi\nu)(a_{\sigma(1)}\otimes b_{\mu(1)}\otimes c_{\pi(1)}\otimes d_{\nu(1)})\otimes\cdots\otimes(a_{\sigma(n)}\otimes b_{\mu(n)}\otimes c_{\pi(n)}\otimes d_{\nu(n)}).$$

Under the surjection $(A\otimes B\otimes C\otimes D)^{\otimes n}\to\operatorname{Sym}^n(A\otimes B\otimes C\otimes D)$, this maps to

$$\sum_{\sigma,\mu,\pi,\nu\in\mathcal{S}_n}\operatorname{sgn}(\sigma\mu\pi\nu)x_{\sigma(1)\mu(1)\pi(1)\nu(1)}\cdots x_{\sigma(n)\mu(n)\pi(n)\nu(n)},$$

which is the Pascal determinant/quantum permanent.

Proof of Theorem 1.8. To prove part (1), we observe that $g=(\zeta_n\mathrm{I},\mathrm{I},\mathrm{I})\in G$ where ζ_n is a primitive n^{th} root of unity. Moreover, this acts on V by scalar multiplication by ζ_n. This means that for all $v\in V$, the two vectors v and $\zeta_n v$ are in the same orbit. Thus, for a (non-zero) homogenous polynomial to be invariant, it's degree must be a multiple of n in order to take the same value on v and $\zeta_n v$ for all choices of $v \in V$.

In the above discussion, we showed that the degree n invariants is 1-dimensional, and spanned by the Pascal determinant/quantum permanent which we know is VNP–hard by [Gur04].[11] This proves parts (2) and (3). ☐

Remark 4.2. Indeed, let k be a fixed positive integer. Then consider the action of $(\mathrm{SL}_n)^{2k+1}$ on $(\mathbb{C}^n)^{\otimes 2k+2}$ that generalizes in the obvious way the setup of Theorem 1.8 (which is the $k = 1$ case). Then, once again non-zero homogeneous invariants exist only in degrees that are multiple of n, and the degree n invariants are 1-dimensional. The invariant spanning the degree n invariants can be seen as the alternating sum of $(2k + 2)$-D matchings. So, it must be hard to compute by a similar argument to the one in the previous section since the $(2k + 2)$-D matching problem is also an NP-hard problem.

References

[AB03] Manindra Agrawal and Somenath Biswas. Primality and identity testing via chinese remaindering. *Journal of the ACM (JACM)*, 50(4):429–443, 2003.

[AZGL+18] Zeyuan Allen-Zhu, Ankit Garg, Yuanzhi Li, Rafael Oliveira, and Avi Wigderson. Operator scaling via geodesically convex optimization, invariant theory and polynomial identity testing. *STOC*, 2018.

[BFG+18] Peter Bürgisser, Cole Franks, Ankit Garg, Rafael Oliveira, Michael Walter, and Avi Wigderson. Efficient algorithms for tensor scaling, quantum marginals, and moment polytopes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 883–897. IEEE, 2018.

[BGO+17] Peter Bürgisser, Ankit Garg, Rafael Oliveira, Michael Walter, and Avi Wigderson. Alternating minimization, scaling algorithms, and the null-cone problem from invariant theory. 2017. Proceedings of Innovations in Theoretical Computer Science (ITCS 2018).

[BIP19] Peter Bürgisser, Christian Ikenmeyer, and Greta Panova. No occurrence obstructions in geometric complexity theory. *J. Amer. Math. Soc.*, 32(1):163–193, 2019.

[DK15] Harm Derksen and Gregor Kemper. *Computational invariant theory*, volume 130 of *Encyclopaedia of Mathematical Sciences*. Springer, Heidelberg, enlarged edition, 2015. With two appendices by Vladimir L. Popov, and an addendum by Norbert A'Campo and Popov, Invariant Theory and Algebraic Transformation Groups, VIII.

[DM17a] Harm Derksen and Visu Makam. Generating invariant rings of quivers in arbitrary characteristic. *J. Algebra*, 489:435–445, 2017.

[DM17b] Harm Derksen and Visu Makam. Polynomial degree bounds for matrix semi-invariants. *Adv. Math.*, 310:44–63, 2017.

[DM18] Harm Derksen and Visu Makam. Algorithms for orbit closure separation for invariants and semi-invariants of matrices. *arXiv e-prints*, page arXiv:1801.02043, Jan 2018.

[DM19] Harm Derksen and Visu Makam. An exponential lower bound for the degrees of invariants of cubic forms and tensor actions. *arXiv e-prints*, page arXiv:1902.10773, Feb 2019.

[DW00] Harm Derksen and Jerzy Weyman. Semi-invariants of quivers and saturation for littlewood-richardson coefficients. *Journal of the American Mathematical Society*, 13(3):467–479, 2000.

[DZ01] Mátyás Domokos and Alexander N Zubkov. Semi-invariants of quivers as determinants. *Transformation groups*, 6(1):9–24, 2001.

[FS13] Michael A. Forbes and Amir Shpilka. Explicit Noether normalization for simultaneous conjugation via polynomial identity testing. In *Approximation, randomization, and combinatorial*

[11]The reduction from permanent is quite simple. Suppose Y is an $n \times n$ matrix. Then if one considers an $n \times n \times n \times n$ tensor x s.t. $x_{iijj} = Y_{ij}$ (and all other entries are zero), then the Pascal determinant of x is equal to $n!$ times the permanent of Y.

optimization, volume 8096 of *Lecture Notes in Comput. Sci.*, pages 527–542. Springer, Heidelberg, 2013.

[Ful97] William Fulton. *Young tableaux*, volume 35 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1997. With applications to representation theory and geometry.

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

[GGOW18] Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. Algorithmic and optimization aspects of brascamp-lieb inequalities, via operator scaling. *Geometric and Functional Analysis*, 28(1):100–145, 2018.

[GP14] Joshua A Grochow and Toniann Pitassi. Circuit complexity, proof complexity, and polynomial identity testing. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 110–119. IEEE, 2014.

[Gur04] Leonid Gurvits. Classical complexity and quantum entanglement. *J. Comput. System Sci.*, 69(3):448–484, 2004.

[Hil90] David Hilbert. Ueber die Theorie der algebraischen Formen. *Math. Ann.*, 36(4):473–534, 1890.

[Hil93] David Hilbert. Ueber die vollen Invariantensysteme. *Math. Ann.*, 42(3):313–373, 1893.

[IQS18] Gábor Ivanyos, Youming Qiao, and K. V. Subrahmanyam. Constructive non-commutative rank computation is in deterministic polynomial time. *Comput. Complexity*, 27(4):561–593, 2018.

[Kar72] Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.

[Lan15] J. M. Landsberg. Geometric complexity theory: an introduction for geometers. *Ann. Univ. Ferrara Sez. VII Sci. Mat.*, 61(1):65–117, 2015.

[MS01] Ketan D. Mulmuley and Milind Sohoni. Geometric complexity theory. I. An approach to the P vs. NP and related problems. *SIAM J. Comput.*, 31(2):496–526, 2001.

[MS08] Ketan D. Mulmuley and Milind Sohoni. Geometric complexity theory. II. Towards explicit obstructions for embeddings among class varieties. *SIAM J. Comput.*, 38(3):1175–1206, 2008.

[Mul17] Ketan D. Mulmuley. Geometric complexity theory V: Efficient algorithms for Noether normalization. *J. Amer. Math. Soc.*, 30(1):225–309, 2017.

[Mum65] David Mumford. *Geometric invariant theory*. Ergebnisse der Mathematik und ihrer Grenzgebiete, Neue Folge, Band 34. Springer-Verlag, Berlin-New York, 1965.

[Sch79] Jacob T Schwartz. Probabilistic algorithms for verification of polynomial identities. In *International Symposium on Symbolic and Algebraic Manipulation*, pages 200–215. Springer, 1979.

[Str73] Volker Strassen. Vermeidung von divisionen. *Journal für die reine und angewandte Mathematik*, 264:184–202, 1973.

[Stu08] Bernd Sturmfels. *Algorithms in Invariant Theory*. Texts & Monographs in Symbolic Computation. Springer, 2nd edition, 2008.

[SVdB01] Aidan Schofield and Michel Van den Bergh. Semi-invariants of quivers for arbitrary dimension vectors. *Indagationes Mathematicae*, 12(1):125–138, 2001.

[SY$^+$10] Amir Shpilka, Amir Yehudayoff, et al. Arithmetic circuits: A survey of recent results and open questions. *Foundations and Trends® in Theoretical Computer Science*, 5(3–4):207–388, 2010.

[Weh93] David L. Wehlau. Constructive invariant theory for tori. *Ann. Inst. Fourier (Grenoble)*, 43(4):1055–1066, 1993.

Microsoft Research India, Bangalore
E-mail address: garga@microsoft.com

School of Mathematics, Institute for Advanced Study, Princeton
E-mail address: visu@ias.edu

University of Toronto
E-mail address: rafael@cs.toronto.edu

School of Mathematics, Institute for Advanced Study, Princeton
E-mail address: avi@ias.edu

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il