Electronic Colloquium on Computational Complexity, Report No. 139 (2022)

On the VNP-hardness of Some Monomial Symmetric Polynomials

Radu Curticapean
IT University of Copenhagen, Denmark
Basic Algorithms Research Copenhagen, Denmark

Nutan Limaye
IT University of Copenhagen, Denmark
Basic Algorithms Research Copenhagen, Denmark

Srikanth Srinivasan
Department of Computer Science, Aarhus University, Denmark.
On leave from Department of Mathematics, IIT Bombay, India.
Supported by start-up grant from Aarhus University.

Abstract

A polynomial $P \in \mathbb{F}[x_1, \ldots, x_n]$ is said to be symmetric if it is invariant under any permutation of its input variables. The study of symmetric polynomials is a classical topic in mathematics, specifically in algebraic combinatorics and representation theory. More recently, they have been studied in several works in computer science, especially in algebraic complexity theory.

In this paper, we prove the computational hardness of one of the most basic kinds of symmetric polynomials: the *monomial symmetric polynomials*, which are obtained by summing all distinct permutations of a single monomial. This family of symmetric functions is a natural basis for the space of symmetric polynomials (over any field), and generalizes many well-studied families such as the elementary symmetric polynomials and the power-sum symmetric polynomials.

We show that certain families of monomial symmetric polynomials are *VNP-complete* with respect to oracle reductions. This stands in stark contrast to the case of elementary and power symmetric polynomials, both of which have constant-depth circuits of polynomial size.

2012 ACM Subject Classification Theory of computation → Algebraic complexity theory; Computing methodologies → Representation of polynomials

Keywords and phrases algebraic complexity, symmetric polynomial, permanent, Sidon set

Digital Object Identifier 10.4230/LIPIcs...

Funding Basic Algorithms Research Copenhagen is supported by Villum Foundation grant 16582.

1 Introduction

This paper considers the algebraic complexity of *symmetric polynomials*: a multivariate polynomial $f \in \mathbb{F}[x_1, \ldots, x_n]$ is said to be symmetric if it is invariant under any permutation of its variables $x_1, \ldots, x_n$. Standard examples of such polynomials include the *elementary symmetric polynomials* and the *power-sum symmetric polynomials.* The study of symmetric polynomials is a classical topic in mathematics, especially in algebraic combinatorics and representation theory (see, e.g. [18, 14]). In particular, standard bases of homogeneous symmetric polynomials of fixed degree d and the matrices of linear transformations that translate between these bases are studied. For many natural bases, the entries of these matrices encode interesting combinatorial and representation-theoretic quantities.

An important example of such a basis of n-variate symmetric polynomials is the family of *monomial symmetric polynomials*, which are considered in this paper. In the following, we say that a partition $\boldsymbol{\lambda}$ of an integer $d \in \mathbb{N}$ is a non-increasingly ordered tuple of positive numbers $\boldsymbol{\lambda} = (\lambda_1, \lambda_2, \ldots, \lambda_r)$ summing to d, i.e. $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_r$ and $\sum_i^r \lambda_i = d$. We write $\boldsymbol{\lambda} \vdash d$

ISSN 1433-8092

to indicate this fact. The monomial symmetric polynomial $m_{\boldsymbol{\lambda}}$ is the polynomial obtained by summing all distinct monomials $y_1^{\lambda_1} \cdots y_r^{\lambda_r}$ that can be obtained by picking $y_1, \ldots, y_r$ out of $x_1, \ldots, x_n$ without repetitions. These generalize both the elementary symmetric polynomials (obtained by taking $r = d$ and all $\lambda_i = 1$) and the power symmetric polynomials (obtained by taking $r = 1$ and $\lambda_1 = d$). It is also easily seen that any symmetric polynomial is a unique linear combination of monomial symmetric polynomials.

In this paper, we study monomial symmetric polynomials from the perspective of algebraic complexity. The complexity of general symmetric polynomials has already been investigated in various works, as summarized below.

- Many results in algebraic complexity concern the computational complexity of the *elementary* symmetric polynomials. Non-trivial upper bounds for computing these polynomials have been shown in various models [13, 16, 8], starting with the work of Nisan and Wigderson [13]. In particular, the upper bound by Shpilka and Wigderson [16] played a crucial role in recent work that proved the first superpolynomial lower bounds for constant-depth algebraic circuits [10]. Lower bounds for computing elementary symmetric polynomials have also been shown [13, 16, 15, 8, 6].
- The algebraic complexity of various symmetric polynomials in the *monotone* setting has been investigated [5, 7]. Here, the underlying field is the reals and we do not allow any negative constants in the underlying computation. In particular, the result of Grigoriev and Koshevoy [7] implies an exponential lower bound on monotone algebraic circuits computing certain monotone symmetric polynomials. However, this does not imply lower bounds for general (non-monotone) algebraic circuits, which are the focus of this paper.
- The fundamental theorem of symmetric polynomials states that any symmetric polynomial $p(x_1, \ldots, x_n)$ can be written uniquely as a polynomial f_{elem} in the elementary symmetric polynomials. A recent result of Bläser and Jindal [2] shows that, over fields of characteristic 0, the polynomials p and f_{elem} have roughly the same algebraic circuit complexity. This implies the hardness of p when f_{elem} is a known hard polynomial such as the permanent, but it might be non-trivial to understand the complexity of f_{elem} in general. A variant of [2] was proved in [4], which holds for more general models of algebraic computation, but it requires technical conditions on f_{elem}.
- Monomial symmetric polynomials appear naturally in the context of learning theory, e.g., when estimating properties of distributions. Here, the learning algorithm has access to samples from a discrete distribution and is required to estimate a symmetric property of the distribution, e.g., the entropy or support size. Acharya, Das, Orlitsky and Suresh [1] analyzed algorithms based on a particular estimator and showed their optimality in a variety of settings. This estimator seeks to optimize a given monomial symmetric polynomial over the space of probability distributions. The problem we study in this paper, that is, *evaluating* a monomial symmetric polynomial at a given input, intuitively appears to be an easier computational problem.

Many of the above works try to understand the algebraic complexity of various families of monomial symmetric polynomials. However, to the best of our knowledge, it was not known if there are families of monomial symmetric polynomials that are hard for general algebraic circuits. We prove that, indeed, polynomial-sized circuits for certain monomial symmetric polynomials $m_{\boldsymbol{\lambda}}$ would imply that VNP collapses to VP. More formally, we show that these monomial symmetric polynomials are VNP-hard under c-reductions; these reductions will be introduced in Section 2. (Containment in VNP is easily seen, so VNP-completeness follows.)

▶ **Theorem 1** (Main theorem). *Fix an algebraically closed field of characteristic* 0 *or* $q \geq 3$. *There are two polynomial functions* $r, s : \mathbb{N} \to \mathbb{N}$ *and an explicit*[1] *sequence of partitions* $\boldsymbol{\lambda_1}, \boldsymbol{\lambda_2}, \ldots$ *such that* $\boldsymbol{\lambda_n} \vdash r(n)$ *for* $n \in \mathbb{N}$ *and the following holds: If the polynomials* $m_{\boldsymbol{\lambda_n}}(x_1, \ldots, x_{s(n)})$ *admit algebraic circuits of polynomial size, then so does the permanent.*

The permanent of order n is a polynomial in $x_{i,j}$ for $1 \leq i, j \leq n$ and can be seen as a sum over all perfect matchings in a complete bipartite graph with $n + n$ vertices and an edge of weight $x_{i,j}$ between the i-th left and the j-th right vertex. Each perfect matching is weighted by the product of the weights of all involved edges. The hypergraph permanent is defined analogously for k-uniform hypergraphs.

Over characteristic 0, the reduction by Bläser and Jindal [2], augmented by an observation due to Chaugule et al. [4], implies that to prove the theorem, it suffices to establish the hardness of the polynomial combination f_{pow} that expresses $m_{\boldsymbol{\lambda}}$ in terms of the power-sum symmetric polynomials. Towards this, we show that a particular sum-product f_{match} over perfect matchings can be extracted from f_{pow}. However, the weights of perfect matchings M in f_{match} do not necessarily correspond to those in the permanent: A priori, it may not be possible to recover the edges present in M from the weight of M in f_{match}. This property can however be ensured by choosing the parts in $\boldsymbol{\lambda}$ from a *Sidon set*, a notion from additive combinatorics. In a Sidon set, any pair of distinct numbers is uniquely identified by its sum. We can apply this to uniquely recover the edges present in a matching from their weight in f_{match}.

Over characteristic $q \geq 3$, the proof is similar, but more involved: First, we need to cast f_{pow} as a polynomial combination f_{elem} in the elementary symmetric polynomials in order to invoke a known reduction by Chaugule et al. [4] that applies to fields of characteristic q. In this form, it will however be less obvious how to extract a sum-product over perfect matchings. Focussing on the homogeneous component of minimum degree in f_{elem} and carefully choosing $\boldsymbol{\lambda}$ will eventually allow us to extract a $(q-1)$-uniform hypergraph permanent from f_{elem}. Here, we also crucially exploit the characteristic of the field, along with basic properties of the transformation that expresses power-sum symmetric polynomials in terms of the elementary symmetric polynomials.

2 Preliminaries

We use boldface notation $\boldsymbol{x}$, $\boldsymbol{y}$ for vectors. Throughout, $\boldsymbol{\lambda}$ will denote a *partition*, i.e. a sequence of weakly decreasing positive integers $\lambda_1 \geq \lambda_2 \geq \cdots \lambda_r \geq 1$. Here, r is called the *number of parts* of $\boldsymbol{\lambda}$.

Symmetric polynomials

In the following, let $\mathbb{F}$ be any field and let $\boldsymbol{x} = (x_1, \ldots, x_n)$. We say that $P(\boldsymbol{x}) \in \mathbb{F}[\boldsymbol{x}]$ is *symmetric* if it is invariant under all permutations of the underlying variables. Examples of symmetric polynomials include the following:

- The *elementary symmetric polynomials* $e_{n,d} = \sum_S \prod_{i \in S} x_i$ for $d \leq n$, where S ranges over all d-element subsets of $[n]$. If n is implicit from context, we set $e_d := e_{n,d}$.
- The *power-sum symmetric polynomials* $p_{n,d} = \sum_{i=1}^{n} x_i^d$. If n is implicit from context, we denote this polynomial by p_d.

[1] The sequence of partitions is explicit in the sense that there is a polynomial-time algorithm that computes $\boldsymbol{\lambda_n}$ on input 1^n.

- More generally, given a partition $\boldsymbol{\lambda}$ with $r \leq n$ parts, the *monomial symmetric polynomial* $m_{\boldsymbol{\lambda}}$ is the sum of all monomials where the distinct exponents are exactly $\lambda_1, \ldots, \lambda_r$. In particular, when $\lambda_1, \ldots, \lambda_r$ are all distinct, we can define this polynomial by

$$m_{\boldsymbol{\lambda}} = \sum_{\substack{i_1,\ldots,i_r \in [n] \\ \text{distinct}}} x_{i_1}^{\lambda_1} \cdots x_{i_r}^{\lambda_r}.$$

As noted in the introduction, the elementary and power-sum symmetric polynomials are special cases of monomial symmetric polynomials.

The following basic theorem regarding symmetric polynomials will be important.

▶ **Theorem 2** (Fundamental theorem of symmetric polynomials (see, e.g., [11])). *For any symmetric polynomial $f \in \mathbb{F}[x_1, \ldots, x_n]$, there is a unique polynomial $f_{\text{elem}}(y_1, \ldots, y_n)$ with $f_{\text{elem}}(e_1, \ldots, e_n) = f(\boldsymbol{x})$. If $\mathbb{F}$ has characteristic zero, then there is also a unique polynomial $f_{\text{pow}}(y_1, \ldots, y_n)$ that represents f analogously in terms of the power-sum symmetric polynomials.*

Further, both f_{elem} and f_{pow} (the latter over characteristic 0) have degree at most $\deg(f)$ and do not depend on y_i for $i > \deg(f)$.

Algebraic circuits and Oracle reductions

We work throughout with the standard algebraic circuit model. We refer the reader to standard resources [3, 17] for definitions and basic results regarding the model. We recall also the notion of *c-reductions* between two polynomials f and g: We define $L^g(f)$ to be the smallest s such that the polynomial f is computed by an algebraic circuit C of size at most s that is additionally allowed to use gates for the polynomial g. If $L^g(f)$ is bounded by a polynomial in the number of variables and degree of f and g, we also say that f admits a c-reduction to g and write $f \preceq_c g$.

A result of Bläser and Jindal [2] relates the algebraic complexity of a symmetric polynomial f with its associated polynomial f_{elem}, when the underlying field is the field of complex numbers. Chaugule et al. [4, Theorem 4.16] extended the result to f_{pow}.

▶ **Theorem 3** ([2, 4]). *Any symmetric polynomial $f \in \mathbb{C}[\boldsymbol{x}]$ admits the reductions $f_{\text{elem}} \preceq_c f$ and $f_{\text{pow}} \preceq_c f$.*

We also need the following variant of Theorem 3 due to [4]. While the results of [4] are stated for characteristic zero, we show in Section 5 how to modify them to work for positive characteristic in the setting we are interested in.

In the following, given a polynomial $f \in \mathbb{F}[\boldsymbol{x}]$ and an integer d, we use $H_d(f)$ to denote the homogeneous degree-d component of f. We say that a polynomial f has *min-degree t* if $H_t(f) \neq 0$ and $H_i(f) = 0$ for all $i < t$, and we define the min-degree of the zero polynomial to be $+\infty$.

▶ **Theorem 4** (Adaptation of [4], see Section 5). *Let $\mathbb{F}$ be an algebraically-closed field of characteristic $q > 0$. Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be a non-zero symmetric polynomial such that the min-degree of f_{elem} is t. Furthermore, assume that $f_{\text{elem}}(y_1, \ldots, y_n)$ does not depend on the variables y_{n-1} and y_n. Then $H_t(f_{\text{elem}}) \preceq_c f$.*

In the above statement we say that f_{elem} must not depend on the variables y_{n-1} and y_n. This is a mere technical condition required in our proof of this theorem. Finally, we also need the following standard fact:

▶ **Lemma 5** (Homogeneous component extraction. Folklore, see [17, 2]). *Let $\mathbb{F}$ be any field. For any $f \in \mathbb{F}[\boldsymbol{x}]$ and integer $d \geq 0$, we have $H_d(f) \preceq_c f$.*

Permanents

The canonical VNP-complete polynomial family is given by the polynomials Per_n for $n \in \mathbb{N}$, each defined on n^2 variables $x_{i,j}$ for $i, j \in [n]$, such that

$$\mathrm{Per}_n = \sum_{\sigma \in S_n} x_{1,\sigma(1)} \dots x_{n,\sigma(n)},$$

where S_n is the set of all permutations of the set $\{1, 2, \dots, n\}$. When the variables $x_{i,j}$ take Boolean values, the underlying input to Per_n defines a bipartite graph and the above polynomial computes the number of perfect matchings in this graph.

An analogous polynomial can be defined for not necessarily bipartite graphs. Assume that n is an even integer and fix the set of $\binom{n}{2}$ variables $x_{\{i,j\}}$ for all distinct $i, j \in [n]$. Then, we define the *perfect matching polynomial* $\mathrm{PerfMatch}_n$ over these variables by

$$\mathrm{PerfMatch}_n = \sum_{\substack{\text{perfect matchings} \\ M \text{ of } K_n}} \prod_{\{i,j\} \in M} x_{\{i,j\}}.$$

We can also define analogues of the above for *hypergraphs*. Let $k \geq 2$ be an integer and let $K_n^{(k)}$ denote the complete k-uniform hypergraph on n vertices. For n divisible by k, we define the *hypergraph perfect matching polynomial* $\mathrm{hPerfMatch}_n^{(k)}$ over the $\binom{n}{k}$ many variables x_S for $S \in \binom{[n]}{k}$ by

$$\mathrm{hPerfMatch}_n^{(k)} = \sum_{\substack{\text{perfect matchings} \\ M \text{ of } K_n^k}} \prod_{S \in M} x_S.$$

Note that $\mathrm{PerfMatch}_n = \mathrm{hPerfMatch}_n^{(2)}$.

We have the following simple reductions from permanents to their variants.

▶ **Lemma 6.** *For even $n \in \mathbb{N}$, we have $\mathrm{Per}_{n/2} \preceq_c \mathrm{PerfMatch}_n$. More generally, for any fixed $k \in \mathbb{N}$ and any n divisible by k, we have $\mathrm{Per}_{n/k} \preceq_c \mathrm{hPerfMatch}_n^{(k)}$.*

Proof sketch. For even n, reduce $\mathrm{Per}_{n/2}$ to $\mathrm{PerfMatch}_n$ as follows: For $i, j \in [n/2]$, substitute $x_{\{i,n/2+j\}} \leftarrow x_{i,j}$ and $x_S \leftarrow 0$ for all remaining variables x_S. This results in $\mathrm{Per}_{n/2}$.

More generally, for n divisible by k, reduce $\mathrm{Per}_{n/k}$ to $\mathrm{hPerfMatch}_n^{(k)}$ as follows: For $i, j \in [n/k]$, let $S_{i,j} = \{i\} \cup \{tn/k + j \mid t = 1, \dots, k-1\}$ and substitute $x_{S_{i,j}} \leftarrow x_{i,j}$. Then substitute $x_S \leftarrow 0$ for all remaining variables x_S. This results in $\mathrm{Per}_{n/k}$. ◀

Finally, we recall a generalization of the permanent to *rectangular matrices*. Fix an $r \times n$ matrix X where $r \leq n$ and the (i,j)-th entry of X is a variable $x_{i,j}$. For a subset $J \subseteq [n]$ of size r, we define X_J to be the submatrix obtained by keeping only the columns indexed by the indices in J. Now, we define the rectangular permanent $\mathrm{rPer}_{r,n}$ by

$$\mathrm{rPer}_{r,n} = \sum_{J \in \binom{[n]}{r}} \mathrm{Per}_r(X_J).$$

The following polynomial identity will be crucial to our main results.

▶ **Theorem 7** (Binet-Minc Identity [12]). *Let $\mathbb{F}$ be any field. Fix an $r \times n$ matrix X as above. For any non-empty $I \subseteq [n]$, define the polynomial S_I by $S_I = \sum_{j=1}^{n} \prod_{i \in I} x_{i,j}$. Then, we have*

$$\mathrm{rPer}_{r,n} = \sum_{\mathcal{I} \in \mathcal{P}_r} (-1)^{r-|\mathcal{I}|} \prod_{I \in \mathcal{I}} (|I|-1)! \cdot S_I,$$

where $\mathcal{P}_r$ denotes the set of all partitions of $[r]$ into non-empty subsets.

Sidon sets and variants

Our hardness proofs for the monomial symmetric functions $m_{\boldsymbol{\lambda}}$ require certain conditions on $\boldsymbol{\lambda}$: In Section 3, any unordered pair of numbers in $\boldsymbol{\lambda}$ must be uniquely identified from its sum, i.e., the parts in λ form a so-called *Sidon set*. Additionally, sums composed of the parts in $\boldsymbol{\lambda}$ are stratified by the number of terms involved in the sum. Section 4 requires more generally that sets of fixed size $q \in \mathbb{N}$ are identifiable, and that all parts must have remainder 1 modulo q. We capture these requirements in the following definition:

▶ **Definition 8.** *Given a set of integers $L = \{\lambda_1, \ldots, \lambda_r\}$ and a subset $S \subseteq [r]$, define $\lambda_S := \sum_{i \in S} \lambda_i$. We say that L (or a partition $\boldsymbol{\lambda}$ whose multiset of parts equals L) is q-*good *for an integer $q \geq 2$ if the following conditions hold:*
q-wise Sidon set: *For any two distinct sets $S, S' \subseteq [r]$ of size q, we have $\lambda_S \neq \lambda_{S'}$.*
Stratification: *For sets $S, T \subseteq [r]$ with $|S| < q$ and $|T| = q$, we have $\lambda_S < \lambda_T$.*
Units modulo $q+1$: *For each $i \in [r]$, we have $\lambda_i \equiv 1 \pmod{q+1}$.*

Existing constructions of q-wise Sidon sets can be adapted to construct such sets:

▶ **Lemma 9.** *For all $r, q \in \mathbb{N}$, there exists a q-good set of r integers that are bounded by $r^{O(q)}$. Such a set can be constructed deterministically in time $r^{O(q)}$.*

Proof. Let $s \in \mathbb{N}$ be the smallest perfect square that is larger or equal to r. By Lemma 2.5 in [9], there is a q-wise Sidon set $\{\lambda_1, \ldots, \lambda_s\}$ with elements bounded by $s^{O(q)} = r^{O(q)}$ that can be constructed in $s^{O(q)} = r^{O(q)}$ time. Then the r-element subset $\{\lambda_1, \ldots, \lambda_r\}$ trivially is a q-wise Sidon set as well.

Now take $\mu_i = (q+1)\lambda_i + 1$ for all $i \in [r]$; this trivially ensures that $\mu_i \equiv 1 \pmod{q+1}$ for all i, as required in the third property from Definition 8. As the map $x \mapsto (q+1)x + 1$ is injective, the set $\{\mu_1, \ldots, \mu_r\}$ is a q-wise Sidon set.

Finally, to ensure the stratification property, let Σ be the smallest multiple of $q+1$ that is strictly larger than $\mu_1 + \ldots + \mu_r$, define $\mu'_i = \Sigma + \mu_i$ for $i \in [r]$, and set $L := \{\mu'_1, \ldots, \mu'_r\}$. As the map $x \mapsto \Sigma + x$ is injective, L is a q-wise Sidon set. As Σ is a multiple of $q+1$, we have $\mu'_i \equiv \mu_i \equiv 1 \pmod{q+1}$ for all i. We show that $\mu'_I < \mu'_{I'}$ for $I, I' \subseteq [r]$ with $|I| < |I'|$: Note that μ'_i can be interpreted as a 2-digit number $(1, \mu_i)$ in base Σ. For $I \subseteq [r]$, the representation of $\mu'_I = \sum_{i \in I} \mu'_i$ in base Σ is $(|I|, \mu_I)$; this is because Σ is large enough to avoid an overflow of the least significant digit. The stratification property follows.

From the above construction, it follows that L is a q-good set, all numbers in L are bounded by $r^{O(q)}$, and that L can be constructed deterministically in $r^{O(q)}$ time. ◀

3 Main result in characteristic zero

We present our main reduction from permanents to monomial symmetric functions $m_{\boldsymbol{\lambda}}$. The reduction shown in this section applies to the field $\mathbb{C}$. In the next section, we show how to handle fields of characteristic strictly greater than 2; this introduces additional technical difficulties that are not present in this section.

Fix a 2-good partition $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_r)$ with r parts, non-increasingly ordered, and $\boldsymbol{\lambda} \vdash d$ for $d \in \mathbb{N}$. Recall our notation $\lambda_I := \sum_{i \in I} \lambda_i$ for $I \subseteq [r]$. We first express $m_{\boldsymbol{\lambda}}(x_1, \ldots, x_n)$ for $n \in \mathbb{N}$ as a polynomial combination of the power-sum symmetric polynomials $p_j := p_{n,j}(x_1, \ldots, x_n)$ for $1 \leq j \leq d$. That is, we obtain a polynomial $f_{\text{pow}}(y_1, \ldots, y_d)$ in indeterminates $y_1, \ldots, y_d$ such that

$$m_{\boldsymbol{\lambda}}(x_1, \ldots, x_n) = f_{\text{pow}}(p_1, \ldots, p_d).$$

Known reductions will allow us to reduce directly (in characteristic 0) or with extra steps (for characteristic >2) from f_{pow} to $m_{\boldsymbol{\lambda}}$. It therefore remains to establish hardness of f_{pow}. Towards this, we give a combinatorial interpretation of f_{pow} as a sum over partitions of $[r]$; this sum will later be restricted to partitions that are actually perfect matchings of K_r.

▶ **Fact 10.** *If $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_r)$ is a partition of some integer $d \in \mathbb{N}$, and the parts of $\boldsymbol{\lambda}$ are pairwise distinct, then we have $m_{\boldsymbol{\lambda}}(x_1, \ldots, x_n) = f_{\mathrm{pow}}(p_1, \ldots, p_d)$ with*

$$f_{\mathrm{pow}}(y_1, \ldots, y_d) = \sum_{\mathcal{I} \in \mathcal{P}_r} (-1)^{r-|\mathcal{I}|} \prod_{I \in \mathcal{I}} (|I|-1)! \cdot y_{\lambda_I}. \tag{1}$$

Proof. If all parts of $\boldsymbol{\lambda}$ are pairwise distinct, then $m_{\boldsymbol{\lambda}}$ can be expressed as the rectangular permanent of a generalized Vandermonde matrix $V_{\boldsymbol{\lambda}}$ defined from $\boldsymbol{\lambda}$:

$$m_{\boldsymbol{\lambda}} = \mathrm{rPer}_{r,n} \underbrace{\begin{pmatrix} x_1^{\lambda_1} & x_2^{\lambda_1} & \ldots & x_n^{\lambda_1} \\ \vdots & \vdots & \ddots & \vdots \\ x_1^{\lambda_r} & x_2^{\lambda_r} & \ldots & x_n^{\lambda_r} \end{pmatrix}}_{=:V_{\boldsymbol{\lambda}}} \tag{2}$$

The Binet-Minc formula (Theorem 7) then readily yields (1): When invoked on $V_{\boldsymbol{\lambda}}$, the polynomial S_I in the statement of Theorem 7 equals

$$S_I = \sum_{j=1}^{n} \prod_{i \in I} V_\lambda(i,j) = \sum_{j=1}^{n} \prod_{i \in I} x_j^{\lambda_i} = \sum_{j=1}^{n} x_j^{\lambda_I} = p_{\lambda_I}.$$

This concludes the proof. ◀

Note that all parts of $\boldsymbol{\lambda}$ are indeed distinct, since $\boldsymbol{\lambda}$ is 2-good and thus cannot feature a part of multiplicity strictly larger than 1; this follows from the Sidon set property.

Theorem 2 shows that f_{pow} is uniquely determined over characteristic 0, and Theorem 3 yields a reduction from f_{pow} to $m_{\boldsymbol{\lambda}}$, so we establish hardness of f_{pow}: We define a new polynomial f_{match} by restricting the sum over partitions $\mathcal{I} \in \mathcal{P}_r$ in (1) to perfect matchings, i.e., to partitions of $[r]$ in which all parts have cardinality 2. We write $\mathcal{M}_r$ for the set of perfect matchings of $[r]$ and define

$$\begin{aligned} f_{\mathrm{match}}(y_1, \ldots, y_d) &:= \sum_{\mathcal{I} \in \mathcal{M}_r} (-1)^{r-|\mathcal{I}|} \prod_{I \in \mathcal{I}} (|I|-1)! \cdot y_{\lambda_I} \\ &= (-1)^{r/2} \sum_{\mathcal{I} \in \mathcal{M}_r} \prod_{I \in \mathcal{I}} y_{\lambda_I}. \end{aligned} \tag{3}$$

The last identity holds because every $\mathcal{I} \in \mathcal{M}_r$ has exactly $r/2$ parts, each of cardinality 2.

We will show later that f_{match} can be reduced to f_{pow}. First, we establish the hardness of f_{match} by reducing the perfect matching polynomial to it. Here, we crucially use that $\boldsymbol{\lambda}$ is a Sidon set in order to switch between the variables $y_{\lambda_{\{u,v\}}}$ present in f_{match} and the variables $x_{\{u,v\}}$ present in $\mathrm{PerfMatch}_r$.

▷ Claim 11. There is a c-reduction from $\mathrm{PerfMatch}_r$ to f_{match}.

Proof. Since $\boldsymbol{\lambda}$ is a 2-good set, its parts form a 2-wise Sidon set, so the map $\{u,v\} \mapsto \lambda_{\{u,v\}}$ from 2-subsets of $[r]$ into $\mathbb{N}$ is injective. This in turn implies that substituting $y_{\lambda_{\{u,v\}}} \leftarrow x_{\{u,v\}}$ for all $\{u,v\} \subseteq [r]$ into f_{match} yields the polynomial

$$(-1)^{r/2} \sum_{\mathcal{I} \in \mathcal{M}_r} \prod_{I \in \mathcal{I}} x_{\{u,v\}} = (-1)^{r/2} \mathrm{PerfMatch}_r.$$

Multiplication with $(-1)^{r/2}$ then yields the desired c-reduction. ◀

Finally, we reduce f_{match} to f_{pow}. This reduction proceeds in two steps: We first show that the homogeneous component of degree $r/2$ in f_{pow} enumerates the perfect matchings and some additional structures; these additional structures are then removed through the stratification property of $\boldsymbol{\lambda}$.

▷ Claim 12. There is a c-reduction from f_{match} to f_{pow}.

Proof. Consider the homogeneous component $H_{r/2}(f_{\text{pow}})$ in f_{pow}. Lemma 5 gives a c-reduction from $H_{r/2}(f_{\text{pow}})$ to f_{pow}. By inspecting (1), we see that the monomials of $H_{r/2}(f_{\text{pow}})$ correspond to the partitions $\mathcal{I} \in \mathcal{P}_r$ with exactly $r/2$ parts. Such a partition is a perfect matching iff it contains no parts of size 1, as every part must then be of cardinality at least 2, and thus, of cardinality exactly 2.

We thus aim to restrict the sum further to partitions with $r/2$ parts and no parts of cardinality 1. To this end, substitute $p_{\lambda_{\{u\}}} \leftarrow 0$ for all $u \in [d]$: By the stratification property of $\boldsymbol{\lambda}$, this eliminates precisely those partitions from $H_{r/2}(f_{\text{pow}})$ that contain a singleton part $\{u\}$. Overall, this yields a c-reduction from f_{match} over $H_{r/2}(f_{\text{pow}})$ to f_{pow}. ◀

We have now collected all parts of the reduction and summarize it below.

▶ **Lemma 13.** *Let $\mathbb{F} = \mathbb{C}$. Let $\boldsymbol{\lambda} \vdash d$ for $d \in \mathbb{N}$ be a 2-good partition with r parts. Then*

$$\mathrm{Per}_{r/2} \preceq_c m_{\boldsymbol{\lambda}}(x_1, \ldots, x_n)$$

provided that $n \geq d$.

Proof. Let $f_{\text{pow}}(y_1, \ldots, y_d)$ and $f_{\text{match}}(y_1, \ldots, y_d)$ denote the polynomials defined from $\boldsymbol{\lambda}$ in (1) and (3) above. We have the following chain of reductions:

$$\begin{aligned} \mathrm{Per}_{r/2} &\preceq_c \mathrm{PerfMatch}_r && \text{by Lemma 6} \\ &\preceq_c f_{\text{match}}(y_1, \ldots, y_d) && \text{by Claim 12} \\ &\preceq_c f_{\text{pow}}(y_1, \ldots, y_d) && \text{by Claim 11} \\ &\preceq_c m_{\boldsymbol{\lambda}}(x_1, \ldots, x_n) && \text{by Theorem 4.} \end{aligned}$$

The lemma follows. ◀

Combining Lemma 13 and Lemma 9, we obtain a proof of Theorem 1 in the case when the underlying field is $\mathbb{C}$.

Proof of Theorem 1 (characteristic 0). By Lemma 9, there is a sequence of 2-good partitions $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \boldsymbol{\lambda}_3, \ldots$ such that $\boldsymbol{\lambda}_n \vdash d_n$ has n parts and $d_n \leq s(n)$ for a polynomial $s : \mathbb{N} \to \mathbb{N}$. By Lemma 13, we have $\mathrm{Per}_{n/2} \preceq_c m_{\boldsymbol{\lambda}_n}(x_1, \ldots, x_{s(n)})$. The theorem follows. ◀

4 Main result in positive characteristic

In this section, we adapt the proof from Section 3 to prove the main theorem for fields of positive characteristic. Throughout this section, $\mathbb{F}$ denotes an infinite and algebraically closed field of characteristic $q > 2$. Rather than reducing from the perfect matching polynomial for graphs, we reduce from the perfect matching polynomial in $(q-1)$-uniform hypergraphs. In the following, let $\boldsymbol{\lambda}$ be a $(q-1)$-good partition with r parts and $\boldsymbol{\lambda} \vdash d$ for $d \in \mathbb{N}$.

The proof begins again by expressing $m_{\boldsymbol{\lambda}}(x_1, \ldots, x_n) = f_{\text{pow}}(p_1, \ldots, p_d)$ as a polynomial combination of power-sum polynomials p_i for $1 \leq j \leq d$. Since $\boldsymbol{\lambda}$ is $(q-1)$-good, it contains only pairwise distinct parts, so we can use Fact 10 again and obtain

$$f_{\text{pow}}(y_1, \ldots, y_d) = \sum_{\mathcal{I} \in \mathcal{P}_r} (-1)^{r-|\mathcal{I}|} \prod_{I \in \mathcal{I}} (|I|-1)! \cdot y_{\lambda_I}. \tag{4}$$

At this point, we exploit the field characteristic: We have $(|I|-1)! \equiv 0 \pmod q$ if $|I| > q$, implying that only partitions with parts of cardinality $\leq q$ appear in the above sum. Write $\mathcal{P}_r^{\leq q}$ for the set of these partitions, and furthermore write $\mathcal{P}_r^{q-1}$ for the set of partitions whose parts all have cardinality $q-1$. Our goal is to restrict the sum in (4) to partitions from $\mathcal{P}_r^{q-1}$, that is, to perfect matchings in the complete $(q-1)$-uniform r-vertex hypergraph. This resembles the restriction to graph perfect matchings in Section 3.

To achieve this restriction and to invoke Theorem 4 later, we express the power-sum polynomials p_k for $1 \leq k \leq d$ as polynomials in the elementary symmetric polynomials. In contrast to the converse direction (of expressing the elementary symmetric polynomials in terms of the power-sum polynomials), such expressions exist even in positive characteristic: For all $k \in \mathbb{N}$, there is a unique polynomial $f_k(z_1, \ldots, z_k)$ with $p_k = f_k(e_1, \ldots, e_k)$, even over fields of characteristic $q > 0$. Combined with (4), we obtain $m_{\boldsymbol{\lambda}} = f_{\text{elem}}(e_1, \ldots, e_d)$ with

$$f_{\text{elem}}(z_1, \ldots, z_d) = \sum_{\mathcal{I} \in \mathcal{P}_r} (-1)^{r-|\mathcal{I}|} \prod_{I \in \mathcal{I}} (|I|-1)! \cdot f_{\lambda_I}(z_1, \ldots, z_d). \tag{5}$$

The polynomial f_{elem} is unique, since the elementary symmetric polynomials form a basis for the symmetric polynomials over every field. Let t denote the min-degree of f_{elem}. Theorem 4 shows that the homogeneous component of degree t in f_{elem} admits a c-reduction to the polynomial $m_{\boldsymbol{\lambda}}$, so we will focus on this homogeneous component. First, we show that the polynomial f_k, which expresses the power-sum symmetric polynomial p_k in terms of the elementary symmetric polynomials, has min-degree at least 2 whenever k is divisible by q. Note that f_k has no constant term.

▷ Claim 14. The only linear monomial in f_k is $(-1)^{k+1} k \cdot y_k$. In particular, if $q \mid k$, then the min-degree of f_k over characteristic q is at least 2.

Proof. Given a partition $\boldsymbol{\mu} \vdash k$ and $i \in \mathbb{N}$, write $s_i(\boldsymbol{\mu})$ for the multiplicity of i in $\boldsymbol{\mu}$. We have [18, Chapter 7] that

$$f_k(y_1, \ldots, y_k) = (-1)^k k \sum_{\boldsymbol{\mu} \vdash k} \frac{(s_1(\boldsymbol{\mu}) + s_2(\boldsymbol{\mu}) + \cdots + s_k(\boldsymbol{\mu}) - 1)!}{s_1(\boldsymbol{\mu})!\, s_2(\boldsymbol{\mu})! \cdots s_k(\boldsymbol{\mu})!} \prod_{i=1}^{k} (-y_i)^{s_i(\boldsymbol{\mu})}. \tag{6}$$

Note that every partition $\boldsymbol{\mu} \vdash k$ with at least two parts contributes a term of total degree at least two. Only the partition $\boldsymbol{\mu} = (k)$ can therefore contribute a linear monomial, and the contributed monomial is $(-1)^k k \cdot 0!/1! \cdot (-y_k) = (-1)^{k+1} k \cdot y_k$. ◀

Using this claim, we can analyze the min-degree of the contribution to f_{elem} from a partition $\mathcal{I} \in \mathcal{P}_r^{\leq q}$. That is, we write $f_{\text{elem}} = \sum_{\mathcal{I}} b_{\mathcal{I}}$ with $\mathcal{I}$ ranging over $\mathcal{P}_r^{\leq q}$ and

$$b_{\mathcal{I}} := (-1)^{r-|\mathcal{I}|} \prod_{I \in \mathcal{I}} (|I|-1)! \cdot f_{\lambda_I}.$$

It turns out that the min-degree of $b_{\mathcal{I}}$ is minimized for partitions $\mathcal{I} \in \mathcal{P}_r^{q-1}$. This will allow us to isolate these partitions via Theorem 4.

▷ Claim 15. Let $\mathcal{I} \in \mathcal{P}_r^{\leq q}$.

- If $\mathcal{I} \in \mathcal{P}_r^{q-1}$, then the min-degree of $b_{\mathcal{I}}$ is equal to $r/(q-1)$.
- Otherwise, the min-degree of $b_{\mathcal{I}}$ is strictly larger than $r/(q-1)$.

Proof. Parts of size q in $\mathcal{I}$ contribute 2 to the min-degree of $b_{\mathcal{I}}$, while parts of size $\leq q-1$ contribute 1. Consider a Knapsack instance $\mathcal{K}$ with items $S_1, \ldots, S_q$, and item repetitions allowed, where item S_j for $1 \leq j \leq q-1$ has weight 1 and profit j, while item S_q has weight 2 and profit q. The min-degree of $b_{\mathcal{I}}$ for $\mathcal{I} \in \mathcal{P}_r^{\leq q}$ can be viewed as the minimum weight of a solution with profit r for $\mathcal{K}$. Greedily choosing copies of the item S_{q-1} with strictly (since $q > 2$) largest profit-weight ratio yields an optimal fractional solution for $\mathcal{K}$ that consists of $r/(q-1)$ copies of item S_{q-1}. This is an optimal *integral* solution to $\mathcal{K}$, and by optimality of the greedy algorithm, any solution including other items has strictly higher weight.

It follows that the min-degree of $b_{\mathcal{I}}$ over all $\mathcal{I} \in \mathcal{P}_r^{\leq q}$ is at least $r/(q-1)$, and this bound is attained with (and only with) the partitions $\mathcal{I} \in \mathcal{P}_r^{q-1}$. ◀

It follows that the min-degree of f_{elem} is $t := r/(q-1)$. Since only partitions $\mathcal{I} \in \mathcal{P}_r^{q-1}$ have this min-degree t, the homogeneous component of degree t in f_{elem} depends only on these partitions. We obtain

$$H_t(f_{\text{elem}}) \;=\; H_t\left(\sum_{\mathcal{I}\in\mathcal{P}_r^{q-1}} b_{\mathcal{I}}\right) \;=\; H_t\left(\sum_{\mathcal{I}\in\mathcal{P}_r^{q-1}} (-1)^{r-|\mathcal{I}|}\prod_{I\in\mathcal{I}}(|I|-1)!\cdot f_{\lambda_I}\right). \tag{7}$$

Since all partitions $\mathcal{I} \in \mathcal{P}_r^{q-1}$ have t parts, each of size $q-1$, we obtain furthermore that

$$H_t(f_{\text{elem}}) \;=\; (-1)^{r-t}(q-2)!\cdot H_t\left(\sum_{\mathcal{I}\in\mathcal{P}_r^{q-1}}\prod_{I\in\mathcal{I}} f_{\lambda_I}\right). \tag{8}$$

The min-degree of f_{λ_I} for $I \in \mathcal{I} \in \mathcal{P}_r^{q-1}$ is 1, and the unique linear monomial is $(-1)^{\lambda_I+1}\lambda_I \cdot y_{\lambda_I}$. Since $\boldsymbol{\lambda}$ is $(q-1)$-good and $|I| = q-1$, we have $\lambda_I \equiv q-1 \pmod q$. It follows that

$$H_1(f_{\lambda_I}) \equiv (-1)^q(q-1)\cdot y_{\lambda_I}. \pmod q \tag{9}$$

For $I \in \mathcal{P}_r^{q-1}$, the degree-$t$ homogeneous component of $\prod_{I\in\mathcal{I}} f_{\lambda_I}$ is the product of these linear monomials $H_1(f_{\lambda_I})$. That is,

$$H_t\left(\prod_{I\in\mathcal{I}} f_{\lambda_I}\right) \;\equiv\; \prod_{I\in\mathcal{I}} H_1(f_{\lambda_I}) \;\equiv\; (-1)^{(q+1)t}\prod_{I\in\mathcal{I}} y_{\lambda_I}. \pmod q \tag{10}$$

It follows that

$$H_t(f_{\text{elem}}) \;\equiv\; (-1)^{r-t+(q+1)t}(q-2)!\sum_{\mathcal{I}\in\mathcal{P}_r^{q-1}}\prod_{I\in\mathcal{I}} y_{\lambda_I}. \pmod q \tag{11}$$

Using the $(q-1)$-wise Sidon set property of $\boldsymbol{\lambda}$, we can substitute $y_{\lambda_I} \leftarrow x_I$ for all sets $I \subseteq [r]$ of cardinality $q-1$ into (11) as in Claim 11, so as to obtain:

▷ Claim 16. The polynomial $\mathrm{hPerfMatch}_r^{q-1}$ admits a c-reduction to $H_t(f_{\text{elem}})$.

It remains to invoke Theorem 4. We collect the proof steps in the following lemma that parallels Lemma 13 for characteristic 0.

▶ **Lemma 17.** *Let $\mathbb{F}$ be an algebraically closed field of characteristic $q > 2$. Let $\boldsymbol{\lambda} \vdash d$ for $d \in \mathbb{N}$ be a $(q-1)$-good partition with r parts. Then*

$$\mathrm{Per}_{r/(q-1)} \preceq_c m_{\boldsymbol{\lambda}}(x_1, \ldots, x_n),$$

provided that $n \geq d + 2$.

Proof. Let $f_{\text{elem}}(y_1, \ldots, y_d)$ denote the polynomial defined from $\boldsymbol{\lambda}$ in (5). We have the following chain of reductions:

$$\begin{aligned} \mathrm{Per}_{r/(q-1)} &\preceq_c \mathrm{hPerfMatch}_r^{(q-1)} && \text{by Lemma 6} \\ &\preceq_c H_t(f_{\text{elem}}(y_1, \ldots, y_d)) && \text{by Claim 16} \\ &\preceq_c m_{\boldsymbol{\lambda}}(x_1, \ldots, x_n) && \text{by Theorem 4.} \end{aligned}$$

To invoke Theorem 4, we use that $n \geq d + 2$. This means that indeed $f_{\text{elem}}(y_1, \ldots, y_d)$ depends on two variables less than $m_{\boldsymbol{\lambda}}(x_1, \ldots, x_n)$, as required. ◀

The proof of Theorem 1 for characteristic q now follows as in Section 3: Use Lemma 9 to find $(q-1)$-good partitions, then reduce from the family of permanents via Lemma 17.

5 Proof of Theorem 4

In this section, we outline how to modify the result of [4] to show Theorem 4 over an algebraically closed field $\mathbb{F}$ of any characteristic (we will only require that the size of the field $\mathbb{F}$ is large enough and contains primitive roots of unity of large enough order).

High-level Idea.

The modification is based on a very simple idea. [4] prove a result for any algebraically independent polynomials satisfying a (simple) technical condition. To apply this result, the underlying field is required to have characteristic zero in order to apply the *Jacobian criterion,* which states that the Jacobian of a collection of algebraically independent polynomials is full rank over fields of characteristic zero. While this fact fails for fields of positive characteristic, the proof still works if we are independently able to show that the polynomials under consideration induce a Jacobian of full rank. We use this fact to prove their result in the setting that the underlying polynomials are the elementary symmetric polynomials $e_1, \ldots, e_{n-2}$.

The following is implicit in [4, Lemma 27]. The proof is only stated for homogeneous polynomials g but easily works in the following more general setting as well.

▶ **Lemma 18.** *Let k, n be positive integers with $k \leq n$. Assume that $Q_1, \ldots, Q_k \in \mathbb{F}[x_1, \ldots, x_n]$ are polynomials of degree at most D such that for some $\boldsymbol{a} \in \mathbb{F}^n$, we have*

- *$Q_1(\boldsymbol{a}) = \cdots = Q_k(\boldsymbol{a}) = 0$, and*
- *the $k \times n$ Jacobian matrix $\mathcal{J}(Q_1, \ldots, Q_k)$ has rank k, when evaluated at the point $\boldsymbol{a}$.*

Further, assume that $g \in \mathbb{F}[y_1, \ldots, y_k]$ is a degree-d polynomial of min-degree t and let $G = g(Q_1, \ldots, Q_k)$. Then, $L^G(H_t(g)) \leq \mathrm{poly}(n, d, D)$.

We only sketch the proof, as it is quite similar to [4, Lemma 27].

Proof sketch. By shifting the input $\boldsymbol{x}$ by $\boldsymbol{a}$, we assume without loss of generality that $\boldsymbol{a}$ is the origin (note that this does not affect the Jacobian at all). Now, by a Taylor expansion around the origin, we have for each $i \in [k]$

$$Q_i(\boldsymbol{x}) = \ell_i(\boldsymbol{x}) + R_i(\boldsymbol{x})$$

where $\ell_i(\boldsymbol{x})$ is a homogeneous linear polynomial and $R_i(\boldsymbol{x})$ is a polynomial of min-degree at least 2. Further, the polynomials $\ell_1,\ldots,\ell_k$ are linearly independent as the Jacobian is full-rank at $\boldsymbol{a}$ (i.e. the origin). Thus, we have

$$\begin{aligned}G(\boldsymbol{x}) &= g(Q_1(\boldsymbol{x}),\ldots,Q_k(\boldsymbol{x}))\\ &= \sum_{j=t}^{d} H_j(g)(\ell_1(\boldsymbol{x})+R_1(\boldsymbol{x}),\ldots,\ell_k(\boldsymbol{x})+R_k(\boldsymbol{x}))\\ &= H_t(g)(\ell_1(\boldsymbol{x}),\ldots,\ell_k(\boldsymbol{x}))+R(\boldsymbol{x})\end{aligned}$$

where $R(\boldsymbol{x})$ has min-degree strictly greater than t and degree at most $\deg(G)$. Note that the second equality uses the fact that the min-degree of g is t. Since $\ell_1,\ldots,\ell_k$ are linearly independent, there exists a homogeneous linear transformation T of the variables $x_1,\ldots,x_n$ such that $\ell_i(T(\boldsymbol{x})) = x_i$ for each $i\in[k]$. Applying this linear transformation to the input variables, we have

$$G'(\boldsymbol{x}) := G(T(\boldsymbol{x})) = H_t(g)(\ell_1(T(\boldsymbol{x})),\ldots,\ell_k(T(\boldsymbol{x})))+R(T(\boldsymbol{x})) = H_t(g)(x_1,\ldots,x_k)+R'(\boldsymbol{x})$$

where R' has min-degree strictly greater than t and degree at most $\deg(G)$.

The above clearly implies that $L^G(G') \le \mathrm{poly}(n)$. Furthermore, by Lemma 5, we have that $L^{G'}(H_t(g)) \le \mathrm{poly}(n,\deg(G)) \le \mathrm{poly}(n,d,D)$ as the degree of G is at most $d\cdot D$.

Composing the two reductions, we have $L^G(H_t(g)) \le \mathrm{poly}(n,d,D)$. ◀

We will apply Lemma 18 to the setting when $Q_1,\ldots,Q_k$ are $e_1,\ldots,e_k$ for some $k<n-1$. To do this, we need to show that these polynomials satisfy the hypotheses required of $Q_1,\ldots,Q_k$ in the statement of Lemma 18. We do this now, using ideas from Lemma 30 and 31 of [4].

▶ **Lemma 19.** *Let k,n be positive integers with $k<n-1$. Then the polynomials $e_1,\ldots,e_k$ satisfy the conditions required of $Q_1,\ldots,Q_k$ in the hypothesis of Lemma 18.*

Proof sketch. Define $\ell=k+1$ if q does *not* divide $k+1$ and $\ell=k+2$ otherwise. Note that $k<\ell\le n$. As q does not divide ℓ, the algebraically-closed field $\mathbb{F}$ contains ℓ distinct ℓ-th roots of unity $1,\omega,\ldots,\omega^{\ell-1}$. Let $\boldsymbol{a}=(1,\omega,\ldots,\omega^{\ell-1},0,\ldots,0)$. It is a standard observation (see e.g. [4, Lemma 31]) that $e_1(\boldsymbol{a})=\cdots=e_{\ell-1}(\boldsymbol{a})=0$. As $\ell>k$, this implies the first hypothesis from the statement of Lemma 18 above.

For the second hypothesis, we consider the Jacobian matrix $\mathcal{J}(e_1,\ldots,e_k)$. To show that this matrix is full-rank when evaluated at $\boldsymbol{a}$, it suffices to argue that some $k\times k$ minor of this matrix is non-zero when evaluated at $\boldsymbol{a}$. We consider the minor J_k defined by the first k columns of $\mathcal{J}(e_1,\ldots,e_k)$ (containing the partial derivatives w.r.t. variables $x_1,\ldots,x_k$).

The proof of Lemma 30 in [4] shows that J_k is divisible by the polynomial $\prod_{i<j\le k}(x_i-x_j)$. By comparing the degrees of these polynomials, we see immediately that J must be $c\cdot\prod_{i<j\le k}(x_i-x_j)$ for some scalar $c\in\mathbb{F}$. As the first k co-ordinates of $\boldsymbol{a}$ are distinct, we see that $J_k(\boldsymbol{a})=c\cdot\alpha$ for some non-zero $\alpha\in\mathbb{F}$. So it suffices to show that c is non-zero.

To argue this, we only need to show that J_k is a non-zero polynomial. To see this, consider the coefficient of $x_1^{k-1}x_2^{k-2}\cdots x_{k-1}$ in the minor J_k. We claim that this coefficient is non-zero. In particular, this implies that J_k is a non-zero polynomial.

It remains to prove the claim regarding the monomial $\mathfrak{m}_k := x_1^{k-1}x_2^{k-2}\cdots x_{k-1}$. We have

$$J_k = \sum_{\sigma\in S_n}\mathrm{sgn}(\sigma)\prod_{i=1}^{k}\mathcal{J}(e_1,\ldots,e_k)_{i,\sigma(i)}.$$

To argue that $\mathfrak{m}_k$ has a non-zero coefficient in J_k, we can argue by induction on k. Note that the (i,j)th entry of $\mathcal{J}(e_1,\ldots,e_k)$ is the partial derivative of the polynomial e_i w.r.t. variable x_j. It is thus the sum of all multilinear monomials of degree $i-1$ not divisible by x_j. In particular, the only entry in the kth row that has a monomial involving only the variables $x_1,\ldots,x_{k-1}$ (the set of variables of $\mathfrak{m}_k$) is the entry $\mathcal{J}(e_1,\ldots,e_k)_{k,k}$, and furthermore, the unique such monomial is $x_1\cdots x_{k-1}$.

Expanding the determinant J_k by the Laplace expansion along the kth row, we see that the coefficient of $\mathfrak{m}_k$ in J_k is also the coefficient of $\mathfrak{m}_k$ in

$$x_1\cdots x_{k-1}\cdot J'_k$$

where the latter term J'_k represents the co-factor of $\mathcal{J}(e_1,\ldots,e_k)_{k,k}$ in J_k, which is exactly the minor corresponding to the first $k-1$ columns of $\mathcal{J}(e_1,\ldots,e_{k-1})$, which is J_{k-1}. By induction, the coefficient of $\mathfrak{m}_{k-1} = x_1^{k-2}\cdots x_{k-2}$ in J'_k is non-zero, hence implying that the coefficient of $\mathfrak{m}_k$ in J_k is non-zero as well. ◀

To prove Theorem 4, we apply Lemma 18 to the case when $G = f(x_1,\ldots,x_n)$ and $g = f_{\mathrm{elem}}(y_1,\ldots,y_{n-2})$. Note that, by the hypothesis of Theorem 4, f_{elem} does not depend on y_{n-1} and y_n. By Lemma 19, the polynomials $e_1,\ldots,e_{n-2}$ satisfy the hypotheses of Lemma 18. Applying the latter lemma and using the fact that $e_1,\ldots,e_{n-2}$ have degree at most n, we immediately get $H_t(f_{\mathrm{elem}}) \preceq_c f$, implying Theorem 4.

References

1. Jayadev Acharya, Hirakendu Das, Alon Orlitsky, and Ananda Theertha Suresh. A unified maximum likelihood approach for estimating symmetric properties of discrete distributions. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 11–21. PMLR, 2017. URL: `http://proceedings.mlr.press/v70/acharya17a.html`.
2. Markus Bläser and Gorav Jindal. On the complexity of symmetric polynomials. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPIcs*, pages 47:1–47:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. `doi:10.4230/LIPIcs.ITCS.2019.47`.
3. Peter Bürgisser, Michael Clausen, and Mohammad Amin Shokrollahi. *Algebraic complexity theory*, volume 315 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1997.
4. Prasad Chaugule, Mrinal Kumar, Nutan Limaye, Chandra Kanta Mohapatra, Adrian She, and Srikanth Srinivasan. Schur polynomials do not have small formulas if the determinant doesn't. In Shubhangi Saraf, editor, *35th Computational Complexity Conference, CCC 2020, July 28-31, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 169 of *LIPIcs*, pages 14:1–14:27. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. `doi:10.4230/LIPIcs.CCC.2020.14`.
5. Sergey Fomin, Dima Grigoriev, and Gleb A. Koshevoy. Subtraction-free complexity, cluster transformations, and spanning trees. *Found. Comput. Math.*, 16(1):1–31, 2016. `doi:10.1007/s10208-014-9231-y`.
6. Hervé Fournier, Nutan Limaye, Meena Mahajan, and Srikanth Srinivasan. The shifted partial derivative complexity of elementary symmetric polynomials. *Theory Comput.*, 13(1):1–34, 2017. `doi:10.4086/toc.2017.v013a009`.
7. Dima Grigoriev and Gleb A. Koshevoy. Complexity of tropical schur polynomials. *J. Symb. Comput.*, 74:46–54, 2016. `doi:10.1016/j.jsc.2015.05.005`.
8. Pavel Hrubes and Amir Yehudayoff. Homogeneous formulas and symmetric polynomials. *Comput. Complex.*, 20(3):559–578, 2011. `doi:10.1007/s00037-011-0007-3`.

9. Mrinal Kumar and Ben Lee Volk. Lower bounds for matrix factorization. In Shubhangi Saraf, editor, *35th Computational Complexity Conference, CCC 2020, July 28-31, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 169 of *LIPIcs*, pages 5:1–5:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.CCC.2020.5.
10. Nutan Limaye, Srikanth Srinivasan, and Sébastien Tavenas. Superpolynomial lower bounds against low-depth algebraic circuits. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 804–814. IEEE, 2021. doi:10.1109/FOCS52979.2021.00083.
11. I. G. (Ian Grant) Macdonald. *Symmetric functions and Hall polynomials*. Oxford mathematical monographs. Clarendon Press ; Oxford University Press, Oxford : New York, 1979.
12. Henryk Minc and Marvin Marcus. *Permanents*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1984. doi:10.1017/CBO9781107340688.
13. Noam Nisan and Avi Wigderson. Lower bounds on arithmetic circuits via partial derivatives. *Comput. Complexity*, 6(3):217–234, 1996/97. doi:10.1007/BF01294256.
14. Amritanshu Prasad. *Representation theory*, volume 147 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Delhi, 2015. A combinatorial viewpoint. doi:10.1017/CBO9781139976824.
15. Amir Shpilka. Affine projections of symmetric polynomials. *J. Comput. Syst. Sci.*, 65(4):639–659, 2002. doi:10.1016/S0022-0000(02)00021-1.
16. Amir Shpilka and Avi Wigderson. Depth-3 arithmetic circuits over fields of characteristic zero. *Comput. Complex.*, 10(1):1–27, 2001. doi:10.1007/PL00001609.
17. Amir Shpilka and Amir Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *Found. Trends Theor. Comput. Sci.*, 5(3-4):207–388, 2010. doi:10.1561/0400000039.
18. Richard P. Stanley. *Enumerative combinatorics. Vol. 2*, volume 62 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1999. With a foreword by Gian-Carlo Rota and appendix 1 by Sergey Fomin. doi:10.1017/CBO9780511609589.

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il