On the Limits of Depth Reduction at Depth-3 Over Small Finite Fields

Suryajith Chillara Partha Mukhopadhyay
Chennai Mathematical Institute, India
{suryajith, partham}@cmi.ac.in

November 5, 2018

Abstract

In a surprising recent result, Gupta-Kamath-Kayal-Saptharishi have proved that over $\mathbb{Q}$ any $n^{O(1)}$-variate and n-degree polynomial in **VP** can also be computed by a depth three $\Sigma\Pi\Sigma$ circuit of size $2^{O(\sqrt{n}\log^{3/2} n)}$ [1]. Over fixed-size finite fields, Grigoriev and Karpinski proved that any $\Sigma\Pi\Sigma$ circuit that computes the determinant (or the permanent) polynomial of a $n \times n$ matrix must be of size $2^{\Omega(n)}$. In this paper, for an explicit polynomial in **VP** (over fixed-size finite fields), we prove that any $\Sigma\Pi\Sigma$ circuit computing it must be of size $2^{\Omega(n\log n)}$. The explicit polynomial that we consider is the iterated matrix multiplication polynomial of n generic matrices of size $n \times n$. The importance of this result is that over fixed-size fields there is *no depth reduction technique* that can be used to compute all the $n^{O(1)}$-variate and n-degree polynomials in **VP** by depth 3 circuits of size $2^{o(n\log n)}$. The result of Grigoriev and Karpinski can only rule out such a possibility for $\Sigma\Pi\Sigma$ circuits of size $2^{o(n)}$.

We also give an example of an explicit polynomial ($\mathrm{NW}_{n,\epsilon}(\mathrm{X})$) in **VNP** (which is not known to be in **VP**), for which any $\Sigma\Pi\Sigma$ circuit computing it (over fixed-size fields) must be of size $2^{\Omega(n\log n)}$. The polynomial we consider is constructed from the combinatorial design of Nisan and Wigderson, and is closely related to the polynomials considered in many recent papers (by Kayal-Saha-Saptharishi, Kayal-Limaye-Saha-Srinivasan, and Kumar-Saraf), where strong depth 4 circuit size lower bounds are shown.

1 Introduction

In a recent breakthrough, Gupta et al. [1] have proved that over $\mathbb{Q}$, if an $n^{O(1)}$-variate polynomial of degree d is computable by an arithmetic circuit of size s, then it can also be computed by a depth three $\Sigma\Pi\Sigma$ circuit of size $2^{O(\sqrt{d\log d\log n\log s})}$[2]. Using this result, they prove the existence of a $\Sigma\Pi\Sigma$ circuit of size $2^{O(\sqrt{n}\log n)}$ computing the determinant polynomial of an $n \times n$ matrix (over $\mathbb{Q}$). Before this result, no depth 3 circuit for Determinant of size smaller than $2^{O(n\log n)}$ was known (over any field of characteristic $\neq 2$).

[1] In a nice follow-up work, Tavenas has improved the upper bound to $2^{O(\sqrt{n}\log n)}$. The main ingredient in his proof is an improved depth 4 reduction.

[2] Gupta et al. [1], using the depth reduction of Koiran [2], show that if a polynomial is computed by an algebraic branching program of size s, then it can also be compulated by a depth three circuit of size $2^{O(\sqrt{d\log n\log s})}$. The determinant polynomial of a $n \times n$ matrix has an algebraic branching program of size poly(n).

The situation is very different over *fixed-size finite fields*. Grigoriev and Karpinski proved that over fixed-size finite fields, any depth 3 circuit for the determinant polynomial of a $n \times n$ matrix must be of size $2^{\Omega(n)}$ [3]. Although Grigoriev and Karpinski proved the lower bound result only for the determinant polynomial, it is a folklore result that some modification of their argument can show a similar depth 3 circuit size lower bound for the permanent polynomial as well [3]. Over any field, Ryser's formula for Permanent gives a $\Sigma\Pi\Sigma$ circuit of size $2^{O(n)}$ [5] (for an exposition of this result, see [6]). Thus, for the permanent polynomial the depth 3 complexity (over fixed-size finite fields) is essentially $2^{\Theta(n)}$.

The result of [1] is obtained through an ingenious depth reduction technique but their technique is tailored to the fields of zero characteristic. In particular, the main technical ingredients of their proof are the well-known monomial formula of Fischer [7] and the duality trick of Saxena [8]. These techniques do not work over finite fields. Looking at the contrasting situation over $\mathbb{Q}$ and the fixed-size finite fields, a natural question is to ask whether one can find a new depth reduction technique over fixed-size finite fields such that any $n^{O(1)}$-variate and degree n polynomial in **VP** can also be computed by a $\Sigma\Pi\Sigma$ circuit of size $2^{o(n \log n)}$.

Question 1. *Over any fixed-size finite field $\mathbb{F}_q$, is it possible to compute any $n^{O(1)}$-variate and n-degree polynomial in* **VP** *by a $\Sigma\Pi\Sigma$ circuit of size $2^{o(n \ln n)}$?*

Note that any $n^{O(1)}$-variate and n-degree polynomial can be trivially computed by a $\Sigma\Pi\Sigma$ circuit of size $2^{O(n \log n)}$ by writing it explicitly as a sum of all $n^{O(n)}$ possible monomials.

We give a negative answer to the aforementioned question by showing that over fixed-size finite fields, any $\Sigma\Pi\Sigma$ circuit computing the iterated matrix multiplication polynomial (which is in **VP** for any field) must be of size $2^{\Omega(n \log n)}$ (See Subsection 2.1, for the definition of the polynomial). More precisely, we prove that any $\Sigma\Pi\Sigma$ circuit computing the iterated matrix multiplication polynomial of n generic $n \times n$ matrices (denoted by $\mathrm{IMM}_{n,n}(X)$), must be of size $2^{\Omega(n \log n)}$.

Previously, Nisan and Wigderson [9] proved a size lower bound of $\Omega(n^{d-1}/d!)$ for any homogeneous $\Sigma\Pi\Sigma$ circuit computing the iterated matrix multiplication polynomial over d generic $n \times n$ matrices. Kumar et al. [10] improved the bound to $\Omega(n^{d-1}/2^d)$. These results work over any field. Over fields of zero characteristic, Shpilka and Wigderson proved a near quadratic lower bound for the size of depth 3 circuits computing the trace of the iterated matrix multiplication polynomial [11].

Recently Tavenas [12], by improving upon the previous works of Agrawal and Vinay [13], and Koiran [2] proved that any $n^{O(1)}$-variate, n-degree polynomial in **VP** has a depth four $\Sigma\Pi^{[O(\sqrt{n})]}\Sigma\Pi^{[\sqrt{n}]}$ circuit of size $2^{O(\sqrt{n} \log n)}$. Subsequently, Kayal et al. [14] proved a size lower bound of $2^{\Omega(\sqrt{n} \log n)}$ for a polynomial in **VNP** which is constructed from the combinatorial design of Nisan and Wigderson [15]. In a beautiful follow up result, Fournier et al. [16] proved that a similar lower bound of $2^{\Omega(\sqrt{n} \log n)}$ is also attainable by the iterated matrix multiplication polynomial (see [17], for a unified analysis of the depth 4 lower bounds of [14] and [16]). The main technique used was *the method of shifted partial derivatives* which was used to prove $2^{\Omega(\sqrt{n})}$ size lower bound for $\Sigma\Pi^{[O(\sqrt{n})]}\Sigma\Pi^{[\sqrt{n}]}$ circuits computing Determinant or Permanent polynomial [18]. Recent work of Kumar and Saraf [19] shows that the depth reduction as shown by Tavenas [12] is optimal even for the homogeneous formulas. This strengthens the result of [16] who proved the optimality of depth reduction for the circuits. Very recently, a series of papers show strong depth 4 lower bounds even for homogeneous depth 4 formulas with no bottom fan-in restriction [20, 21, 22].

[3] Saptharishi gives a nice exposition of this result in his survey and he attributes it to Koutis and Srinivasan [4].

Similar to the situation at depth 4, we also give an example of an explicit n^2-variate and n-degree polynomial in **VNP** (which is not known to be in **VP**) such that over fixed-size finite fields, any depth three $\Sigma\Pi\Sigma$ circuit computing it must be of size $2^{\Omega(n \log n)}$. This polynomial family, denoted by $\mathrm{NW}_{n,\epsilon}(\mathrm{X})$ (see Subsection 2.1, for the definition of the polynomial) is closely related to the polynomial family introduced by Kayal et al. [14]. In fact, from our proof idea it will be clear that the strong depth 3 size lower bound results that we show for $\mathrm{NW}_{n,\epsilon}(\mathrm{X})$ and $\mathrm{IMM}_{n,n}(\mathrm{X})$ polynomials are not really influenced by the fact that the polynomials are either in **VNP** or **VP**. Rather, the bounds are determined by a combinatorial property of the subspaces generated by a set of carefully chosen derivatives.

Our main theorem is the following.

Theorem 2. *Over any fixed-size finite field $\mathbb{F}_q$, any depth three $\Sigma\Pi\Sigma$ circuit computing the polynomials $\mathrm{NW}_{n,\epsilon}(\mathrm{X})$ or $\mathrm{IMM}_{n,n}(\mathrm{X})$ must be of size at least $2^{\delta n \log n}$, where the parameters δ and $\epsilon(< 1/2)$ are in $(0,1)$ and depend only on q.*

In section 6, we set the parameter δ to $\frac{1}{20q \log q}$ and it follows from the subsequent calculations that $\epsilon < \delta + 0.1$. As an important consequence of the above theorem, we have the following corollary.

Corollary 3. *Over any fixed-size finite field $\mathbb{F}_q$, there is no depth reduction technique that can be used to compute all the $n^{O(1)}$-variate and n-degree polynomials in* **VP** *by depth 3 circuits of size $2^{o(n \log n)}$.*

The result of [3] only says that over fixed-size finite fields, not all the $n^{O(1)}$-variate and n-degree polynomials in **VP** can be computed by $\Sigma\Pi\Sigma$ circuits of size $2^{o(n)}$. Our main theorem (Theorem 2) can also be viewed as the first quantitative improvement over the result of [3].

Proof Idea

Our proof technique is quite simple and it borrows ideas mostly from the proof technique of Grigoriev and Karpinski [3]. A recurring notion in many papers related to $\Sigma\Pi\Sigma$ circuits is the notion of *rank* of a product gate. Let $T = L_1 L_2 \dots L_d$ be a product gate such that each L_i is an affine linear form over the underlying field. By rank of T, one simply means the maximum rank of the homogeneous linear system corresponding to set of affine linear functions $\{L_1, L_2, \dots, L_d\}$.

Over fixed-size finite fields, $\Sigma\Pi\Sigma$ circuits enjoy a nice property that the derivatives of the high rank product gates can be eliminated except for a few erroneous points (denoted by E). This property was first observed by Grigoriev and Karpinski in [3]. The intuition is simple. If a product gate has many linearly independent functions, then it is likely that a large number of linear functions will be set to zero if we randomly substitute the variables with elements from the field. Then the derivatives (of relatively low order) of the polynomial obtained from the product gate will disappear on a random point with very high probability.

To quantify the notion of *high rank*, Grigoriev and Karpinski fixed a threshold for the rank of the product gates. Since they were looking for a $2^{\Omega(n)}$ lower bound for the Determinant of a $n \times n$ matrix and the rank of the entire derivative space of of the determinant polynomial is $2^{O(n)}$, it was natural for them to fix the threshold to be $\Theta(n)$. Since the dimension of the derivative spaces of the polynomial families $\{\mathrm{NW}_{n,\epsilon}(\mathrm{X})\}_{n>0}$ and $\{\mathrm{IMM}_{n,n}(\mathrm{X})\}_{n>0}$ is $2^{\Omega(n \log n)}$, it is possible for us to choose the threshold for the rank of the product gates to be $\Theta(n \log n)$. This allows us to bound the size of the error set meaningfully. We formalize this in Lemma 7.

An overview of the result of Grigoriev and Karpinski

We now give a high level description of the proof technique in [3] to motivate our proof strategy. Roughly speaking, they consider the space H spanned by $\Theta(n)$ order derivatives of the determinant. This makes the dimension of H to be of the order of $2^{\Theta(n)}$. For a point $a \in \mathbb{F}_q^N$, the subspace H_a is the space of functions in H that evaluate to zero on a. From the rank analysis on the circuit side, they get that the dimension of the space of functions that may not be zero outside the error set E is bounded. More precisely, $\mathrm{codim}(\cap_{a \notin E} H_a)$ is small. Grigoriev and Karpinski then considered the group of invertible matrices G of order $n \times n$ over $\mathbb{F}_q$. For any $g \in G$, they define a $\mathbb{F}_q$-linear operator $T_g : H \to H$ by the formula $(T_g(f))(a) = f(ga)$. The fact that the derivative space of the determinant polynomial of a $n \times n$ matrix is invariant under $\mathrm{GL}_n(\mathbb{F}_q)$ action was crucially used in defining the map. Then they consider the plane $P = \cap_{a \in G \setminus E} H_a \subset H$. Since $\mathrm{codim}(\cap_{a \notin E} H_a)$ is bounded, the same bound applies for $\mathrm{codim}(P)$ as well. The most remarkable idea in their work was to prove that $\mathrm{codim}(\cap_{b \in G} H_b)$ in H is small given that $\mathrm{codim}(P)$ is small. Notice that the plane P is defined only on $G \setminus E$ and not on the entire group G. To achieve this, they prove that the full invertible group G can be covered by taking only a few translates of $G \setminus E$ from G. This was done by appealing to a graph theoretic lemma of Lovász [23]. Now, it is not hard to see that we can bound $\mathrm{codim}(\cap_{b \in G} H_b)$ to a quantity smaller than $\dim(H)$. This shows us that there exists a nonzero function in H that evaluates to zero on the entire group G. Since the elements in H are only multilinear polynomials, they finally prove that it is impossible to have such a function in H by showing that no nonzero multilinear polynomial can vanish over the entire group G.

An overview of our result

The group symmetry based argument of [3] is tailored to the determinant polynomial and it can not be directly applied to the polynomials that we consider. The main technical contribution of this work is to replace the group symmetry based argument by a new argument that makes the proof strategy robust enough to handle the family of polynomials that we consider. We carefully choose a subspace H (of sufficiently large dimension) of the derivative spaces of these polynomials which have an additional structure. The subspace H is spanned by a *downward closed* set of monomials (see Definition 5). Let $\mathbb{F}_q$ be the finite field and N be the number of the variables in the polynomial under consideration. The basic idea is to prove that the dimension of the space H is strictly more than the dimension of the set of functions in H which do not evaluate to zero over the entire space $\mathbb{F}_q^N$ when the polynomial considered is computed by any depth three circuit. Since the subspace H for the polynomials considered contains only the multilinear polynomials (see section 5), we can then conclude that a nonzero multilinear polynomial in H will evaluate to zero on entire $\mathbb{F}_q^N$, which is not possible by combinatorial nullstellensatz [24].

To implement this, we define a linear map $T_u : H \to H$ by $T_u(f(\mathrm{X})) = f(\mathrm{X} - u)$ for any function $f : \mathbb{F}_q^N \to \mathbb{F}_q$ and $u \in \mathbb{F}_q^N$. The map is well-defined by the downward closed structure of the generating set for H. Also the map T_u is one to one for any $u \in \mathbb{F}_q^N$. As before, for a point $a \in \mathbb{F}_q^N$, the subspace H_a is the space of functions in H that evaluate to zero on a. Let $P = \bigcap_{a \in \mathbb{F}_q^N \setminus E} H_a$. Then by Lemma 7, we get that $\mathrm{codim}\, P = \mathrm{codim}(\bigcap_{a \in \mathbb{F}_q^N \setminus E} H_a)$ is small. Notice that the plane P is defined over $\mathbb{F}_q^N \setminus E$ and not over the entire space $\mathbb{F}_q^N$. Similar to the argument in [3], we use the graph theoretic lemma of Lovász to prove that the entire space $\mathbb{F}_q^N$ can be covered by only a few translates of $\mathbb{F}_q^N \setminus E$. Then it is simple to observe

that $\operatorname{codim}(\bigcap_{b\in\mathbb{F}_q^N} H_b)$ is small compared to the dimension of H. As a consequence we get that a nonzero multilinear polynomial in H must evaluate to zero over $\mathbb{F}_q^N$, which is not possible by the combinatorial nullstellensatz.

2 Preliminaries

Arithmetic Circuits

An arithmetic circuit over a field $\mathbb{F}$ with the set of variables $x_1, x_2, \ldots, x_n$ is a directed acyclic graph such that the internal nodes are labelled by addition or multiplication gates and the leaf nodes are labelled by the variables or the field elements. The node with fan-out zero is the output gate. An arithmetic circuit computes a polynomial in the polynomial ring $\mathbb{F}[x_1, x_2, \ldots, x_n]$. Size of an arithmetic circuit is the number of nodes and the depth is the length of a longest path from the root to a leaf node.

Depth 3 Circuits

Usually a depth 3 circuit over a field $\mathbb{F}$ is denoted by $\Sigma\Pi\Sigma$. The circuit has an addition gate at the top, a middle layer of multiplication gates, and then a level of addition gates at the bottom. A $\Sigma\Pi\Sigma$ circuit with s multiplication gates computes a polynomial of the form $\sum_{i=1}^{s}\prod_{j=1}^{d_i} L_{i,j}(x_1,\ldots,x_n)$ where $L_{i,j}$s are affine linear functions over $\mathbb{F}$ and $\{x_1, x_2, \ldots, x_n\}$ are the variables appearing in the polynomial.

Combinatorial Nullstellensatz

We recall the following theorem from [24].

Theorem 4. *Let $f(x_1, x_2, \ldots, x_n)$ be a polynomial in n variables over an arbitrary field $\mathbb{F}$. Suppose that the degree of f as a polynomial in x_i is at most t_i, for $1 \le i \le n$ and let $S_i \subseteq \mathbb{F}$ such that $|S_i| \ge t_i + 1$. If $f(a_1, a_2, \ldots, a_n) = 0$ for all n-tuples in $S_1 \times S_2 \times \cdots \times S_n$, then $f = 0$.*

2.1 The Polynomial Families

A multivariate polynomial family $\{f_n(\mathrm{X}) \in \mathbb{F}[x_1, x_2, \ldots, x_n] : n \ge 1\}$ is in the class **VP** if f_n has degree at most $\mathrm{poly}(n)$ and can be computed by an arithmetic circuit of size $\mathrm{poly}(n)$. It is in **VNP** if it can be expressed as

$$f_n(\mathrm{X}) = \sum_{\mathrm{Y}\in\{0,1\}^m} g_{n+m}(\mathrm{X}, \mathrm{Y})$$

where $m = |\mathrm{Y}| = \mathrm{poly}(n)$ and g_{n+m} is a polynomial in **VP**.

The Polynomial Family from the Combinatorial Design

Let $\mathbb{F}$ be any field[4]. For integers $n > 0$ ranging over prime powers and $0 < \epsilon < 1$, we define a polynomial family $\{\mathrm{NW}_{n,\epsilon}(\mathrm{X})\}_{n>0}$ in $\mathbb{F}_q[\mathrm{X}]$ as follows.

$$\mathrm{NW}_{n,\epsilon}(\mathrm{X}) = \sum_{a(z)\in\mathbb{F}_n[z]} x_{1a(1)}x_{2a(2)}\dots x_{na(n)}$$

where $a(z)$ runs over all univariate polynomials of degree $< \epsilon n$. The finite field $\mathbb{F}_n$ is naturally identified with the numbers $\{1,2,\dots,n\}$. Notice that the number of monomials in $\mathrm{NW}_{n,\epsilon}(\mathrm{X})$ is $n^{\epsilon n}$. Proposition 4 of [25] tells us that if there is a polynomial time algorithm to test if the coefficient of a given monomial is 1 in the polynomial $P(\mathrm{X}) \in \mathbb{F}[\mathrm{X}]$ with $\{0,1\}$ coefficients, then $P(\mathrm{X})$ is in **VNP** over $\mathbb{F}$. Given any monomial m over X, we can decide in polynomial time if it is a monomial in the polynomial $\{\mathrm{NW}_{n,\epsilon}(\mathrm{X})\}_{n>0}$ by checking if it *conforms* to a univariate polynomial of degree at most ϵn. Thus, $\{\mathrm{NW}_{n,\epsilon}(\mathrm{X})\}_{n>0}$ is in **VNP** for any $\epsilon \in (0,1)$. In [14], a very similar family of polynomials was introduced where the degree of the univariate polynomial was bounded by $\epsilon\sqrt{n}$.

The Iterated Matrix Multiplication Polynomial

The iterated matrix multiplication polynomial of n generic $n \times n$ matrices $\mathrm{X}^{(1)}, \mathrm{X}^{(2)}, \dots, \mathrm{X}^{(n)}$ is the $(1,1)$th entry of the product of the matrices. More formally, let $\mathrm{X}^{(1)}, \mathrm{X}^{(2)}, \dots, \mathrm{X}^{(n)}$ be n generic $n \times n$ matrices with disjoint sets of variables and $x_{ij}^{(k)}$ be the variable in $\mathrm{X}^{(k)}$ indexed by $(i,j) \in [n] \times [n]$. Then the iterated matrix multiplication polynomial (denoted by the family $\{\mathrm{IMM}_{n,n}(\mathrm{X})\}_{n>0}$) is defined as follows.

$$\mathrm{IMM}_{n,n}(\mathrm{X}) = \sum_{i_1,i_2,\dots,i_{n-1}\in[n]} x_{1i_1}^{(1)} x_{i_1i_2}^{(2)} \dots x_{i_{(n-2)}i_{(n-1)}}^{(n-1)} x_{i_{(n-1)}1}^{(n)}.$$

Notice that $\mathrm{IMM}_{n,n}(\mathrm{X})$ is a $n^2(n-2)+2n$-variate polynomial of degree n. For our application, we consider $n = 2m$ where m ranges over the positive integers. Over any field $\mathbb{F}$, the polynomial family $\{\mathrm{IMM}_{n,n}(X)\}_{n>0}$ can be computed in **VP**. This can be seen by observing that $\mathrm{IMM}_{n,n}(X)$ can be computed by a poly(n) sized algebraic branching program.

Downward closed property

Definition 5. *A set of multilinear monomials $\mathscr{M}$ is said to be downward closed if the following property holds. If $m(\mathrm{X}) \in \mathscr{M}$ and a multilinear monomial $m'(\mathrm{X})$ is such that* $\mathrm{var}(m'(\mathrm{X})) \subseteq \mathrm{var}(m(\mathrm{X}))$*, then $m'(\mathrm{X}) \in \mathscr{M}$.*

Now we consider a downward closed set of monomials $\mathscr{M}$ over N variables. These monomials can be viewed as functions from $\mathbb{F}_q^N$ to $\mathbb{F}_q$. W.l.o.g, we assume that the constant function is also in $\mathscr{M}$ (constant function corresponds to a monomial with an empty set of variables). Let H be the subspace spanned by these functions in $\mathscr{M}$.

For any $u \in \mathbb{F}_q^N$, define an operator T_u such that $(T_u(f))(\mathrm{X}) = f(\mathrm{X}-u)$ for any function $f : \mathbb{F}_q^N \to \mathbb{F}_q$. The following proposition is simple to prove.

[4] In the lower bound proof for $\mathrm{NW}_{n,\epsilon}(\mathrm{X})$, we will consider $\mathbb{F}$ to be any fixed finite field $\mathbb{F}_q$.

Proposition 6. *Let H be the subspace spanned by a downward closed set of monomials $\mathscr{M}$ over the set of variables $\{x_1, x_2, \ldots, x_N\}$. Then for any $u \in \mathbb{F}_q^N$, T_u is a linear map from H to H. Moreover, the map T_u is one-to-one for any $u \in \mathbb{F}_q^N$.*

Proof. Let $g(\mathrm{X})$ be an arbitrary function in H which can be expressed as follows: $g(\mathrm{X}) = \sum_{i \geq 1} c_i m_i(\mathrm{X})$ where $m_i(\mathrm{X}) \in \mathscr{M}$, and $c_i \in \mathbb{F}_q$ for all $i \geq 1$.

$$(T_u(g))(\mathrm{X}) = g(\mathrm{X} - u) = \sum_{i \geq 1} c_i m_i(\mathrm{X} - u).$$

It is sufficient to prove that $m(\mathrm{X} - u) \in H$ where $m(\mathrm{X}) \in \mathscr{M}$. We can express $m(\mathrm{X} - u)$ as follows.

$$m(\mathrm{X} - u) = \sum_{S \subseteq \mathrm{var}(m(\mathrm{X}))} c_S \prod_{x_r \in S} x_r.$$

where $c_S \in \mathbb{F}_q$. For every $S \subseteq \mathrm{var}(m(\mathrm{X}))$, $\prod_{x_r \in S} x_r \in \mathscr{M}$ because $\mathscr{M}$ is downward closed. Since the choice of S was arbitrary, $m(\mathrm{X} - u) \in H$. It is obvious that T_u is a linear map.

To see that T_u is a one-to-one map, it is just enough to observe that $T_u \circ T_{-u} = T_0$ where T_0 is an identity map. □

3 The Derivative Space of ΣΠΣ Circuits Over Small Fields

In this section we fix the field $\mathbb{F}$ to be a fixed-size finite field $\mathbb{F}_q$. Let C be a ΣΠΣ circuit of top fan-in s computing a $N = n^{O(1)}$-variate polynomial of degree n. Consider a Π gate $T = L_1 L_2 \ldots L_d$ in C. Let r be the rank of the (homogeneous)-linear system corresponding to $\{L_1, L_2, \ldots, L_d\}$ by viewing each L_i as a vector in $\mathbb{F}_q^{N+1}$. Fix a threshold for the rank of the system of linear functions $r_0 = \beta n \ln n$, where $\beta > 0$ is a constant to be fixed later in the analysis. In our application, the parameter N is at least n^2, so the threshold for the rank is meaningful. W.l.o.g, let $\{L_1, L_2, \ldots, L_r\}$ be a set of affine linear forms in $\{L_1, L_2, \ldots, L_d\}$ whose homogeneous system forms a maximal independent set of linear functions. The following analysis has been reworked from [3] to fix the parameters. It shows that the derivative space of a ΣΠΣ circuit can be approximated by just the derivative space of the low rank product gates of the circuit over a large subset of $\mathbb{F}_q^N$.

Low rank gates : $r \leq r_0$

Over the finite field $\mathbb{F}_q$, we have $x^q = x$. We express $T : \mathbb{F}_q^N \to \mathbb{F}_q$ as a linear combination of $\{L_1^{e_1} L_2^{e_2} \ldots L_r^{e_r} : e_i < q \text{ for all } i \in [r]\}$. Since, the derivatives of all orders lie in the same space, the dimension of the set of partial derivatives of T of all orders is bounded by $q^r \leq q^{r_0}$.

High rank gates : $r > r_0$

Let the rank of a high rank gate T be $y \beta n \ln n$ where $y \geq 1$.

We assign values to the variables uniformly at random from $\mathbb{F}_q$ and compute the probability that at most n linearly independent functions evaluate to zero. Let X_a be the event that at most n linearly independent functions evaluate to zero at a.

$$\Pr_{a\in\mathbb{F}_q^N}[X_a] \leq \sum_{i=0}^{n} \binom{r}{i}\left(\frac{1}{q}\right)^i \left(1-\frac{1}{q}\right)^{r-i} \leq n\binom{r}{n}\left(\frac{1}{q}\right)^n \left(1-\frac{1}{q}\right)^{r-n}.$$

The above inequality follows from the fact that $r > 2n$ and thus the binomial terms under the summation are monotonically increasing. Hence, if we differentiate T with respect to any set of variables of size at most n and restrict all the variables to values from $\mathbb{F}_q$, the gate T may not vanish over a set of points E_T whose size is estimated below.

$$|E_T| \leq n\binom{r}{n}\left(\frac{1}{q}\right)^n \left(1-\frac{1}{q}\right)^{r-n} q^N.$$

Over all the gates, let E be the set of points over which some of the product gates with large rank may not evaluate to zero. Then by a union bound, we get that $|E| \leq s|E_T|$.

$$\begin{aligned}
|E| &\leq s\cdot n\binom{r}{n}\left(\frac{1}{q}\right)^n \left(1-\frac{1}{q}\right)^{r-n} q^N \\
&\leq s\cdot n\left(\frac{er}{n}\right)^n e^{-\frac{r-n}{q}} q^N \\
&= q^N s\cdot e^{n+n\ln\frac{r}{n}+\ln n-\frac{r-n}{q}} \\
&= q^N s\cdot e^{n+n\ln\frac{\gamma\beta n\ln n}{n}+\ln n-\frac{\gamma\beta n\ln n-n}{q}}.
\end{aligned}$$

To bound the above estimate meaningfully, we need $\frac{\ln s}{n\ln n}$ to be strictly less than $\frac{\gamma\beta n}{q}\ln n - n\ln\gamma$. That is, for some constant $\nu > 0$,

$$\frac{\ln s}{n\ln n} - \frac{\gamma\beta}{q} + \frac{\ln\gamma}{\ln n} + \nu < 0. \tag{1}$$

Once we satisfy the relation given by the inequality 1, we can upper bound the size of E as $|E| < q^N\mu^{n\ln n}$ for some suitably fixed constant $\mu = e^{-\nu}$ and μ is between 0 and 1. Now it is clear that over $\mathbb{F}_q^N \setminus E$, the derivative space is spanned by the derivatives of the low rank gates. We summarize it in the following lemma.

Lemma 7. *Let $\mathbb{F}_q$ be a fixed-size finite field. Then there exist constants $0 < \beta(q), \nu(q) < 1$ such that the following is true. Let C be a $\Sigma\Pi\Sigma$ circuit of top fan-in s computing a $N = n^{O(1)}$-variate and n-degree polynomial $f(X)$ over the finite field $\mathbb{F}_q$. Further, $s, \beta(q), \nu(q)$ satisfy the inequality 1. Then, there exists a set $E \subset \mathbb{F}_q^N$ of size at most $q^N\mu^{n\ln n}$ such that the dimension of the space spanned by the derivatives of order $\leq n$ of C restricted to $\mathbb{F}_q^N \setminus E$ is $\leq s\, q^{\beta n\ln n}$ where $\mu = e^{-\nu}$.*

It is worth (re)-emphasizing that, when we consider the derivatives, what we really mean is the formal derivatives of C as polynomials. In the above lemma we view the derivatives as functions from $\mathbb{F}_q^N \to \mathbb{F}_q$. Then it follows from the above analysis that the dimension of the space spanned by the functions corresponding to the derivatives of order $\leq n$ of C restricted to $\mathbb{F}_q^N \setminus E$ is $\leq s\, q^{\beta n \ln n}$. This way of viewing derivatives either as *formal polynomials* or as *functions* is implicit in the work of [3]. In Section 4, we show how to fix the parameters δ, β, and μ which depend only on the field size q.

4 A Covering Argument

In this section, we adapt the covering argument of [3]. In [3], the covering argument was given over the set of invertible matrices. Here we adapt their argument suitably over the entire space $\mathbb{F}_q^N$.

Let H be the derivative space of any polynomial $f(\mathrm{X})$ which is computed by any $\Sigma\Pi\Sigma$ ciruit of size s. Define the subspace $H_a := \{f \in H : f(a) = 0\}$ for $a \in \mathbb{F}_q^N$. Let us recall that E is the set of points over which some of the product gates with large rank may not evaluate to zero. Let the set of points $\mathbb{F}_q^N \setminus E$ be denoted by A. Then Lemma 7 says that in H, we get that $\mathrm{codim}(\bigcap_{a\in A} H_a) < s\, q^{r_0}$. We note the following simple observation.

Proposition 8. *For any $u, a \in \mathbb{F}_q^N$, we have that $T_u(H_a) = H_{u+a}$.*

Proof. Let us recall that the map T_u is one-to-one. That is, $T_{-u} \circ T_u = T_0$ where T_0 is an identity map. It is easy to observe that $T_{-u}(H_{u+a}) = H_a$.

□

Let $P = \bigcap_{a\in A} H_a$. Let $S \subset \mathbb{F}_q^N$ be a set such that we can cover the entire space $\mathbb{F}_q^N$ by the shifts of A with the elements from S.

$$\bigcup_{u\in S} u + A = \mathbb{F}_q^N.$$

Now by applying the map T_u to P which is one-to-one, we get the following.

$$T_u(P) = \bigcap_{a\in A} T_u(H_a) = \bigcap_{b\in u+A} H_b.$$

By a further intersection over S, we get the following.

$$\bigcap_{u\in S} T_u(P) = \bigcap_{u\in S} \bigcap_{b\in u+A} H_b = \bigcap_{b\in \mathbb{F}_q^N} H_b. \tag{2}$$

From Equation 2, we get the following estimate.

$$\mathrm{codim}\left(\bigcap_{b\in\mathbb{F}_q^N} H_b\right) = \mathrm{codim}\left(\bigcap_{u\in S} T_u(P)\right) \leq |S|\ \mathrm{codim}(P) \leq |S|\ s\ q^{r_0}. \tag{3}$$

The $\operatorname{codim}\left(\bigcap_{b\in\mathbb{F}_q^N} H_b\right)$ refers to the dimension of the set of functions in H which do not evaluate to zero over all the points in $\mathbb{F}_q^N$.

Next, we show an upper-bound estimate for the size of the set S. This follows from a simple adaptation of the dominating set based argument given in [3].

Upper bound on the size of the set S

Consider the directed graph $G=(V,R)$ defined as follows. The points in $\mathbb{F}_q^N$ are the vertices of the graph. For $u_1,u_2\in\mathbb{F}_q^N$, the edge $u_1\to u_2$ is in R iff $u_2=u_1+b$ for any $b\in A$. Clearly the in-degree and out-degree of any vertex are equal to $|A|$. Now, we recall Lemma 2 of [3] to estimate the size of S.

Lemma 9 (Lovász, [23])**.** *Let (V,R) be a directed (regular) graph with $|V|=m$ vertices and with the in-degree and the out-degree of each vertex both equal to d. Then there exists a subset $U\subset V$ of size $O(\frac{m}{d}\log(d+1))$ such that for any vertex $v\in V$ there is a vertex $u\in U$ forming an edge $(u,v)\in R$.*

Let c_0 be the constant fixed by the lemma in its $O()$ notation. By Lemma 9, we get the following estimate.

$$\begin{aligned}
|S| &\le c_0\frac{|\mathbb{F}_q^N|}{|A|}\log(|A|+1)\\
&\le c_0\frac{q^N}{q^N-|E|}\log(q^N-|E|+1)\\
&\le c_0(\log q)\,N\,\frac{q^N}{q^N-|E|}\\
&= O(N). \qquad \text{(for fixed } q\text{)}
\end{aligned}$$

The last equation follows from the estimate for $|E|$ from the previous discussions. The following lemma summarizes the content of this section.

Lemma 10. *Let H be the space of partial derivatives of order at most n, of any N-variate polynomial computed by a $\Sigma\Pi\Sigma$ circuit of size s. Then for a suitable parameter r_0, the dimension of the set of functions in H that do not evaluate to zero over all points in $\mathbb{F}_q^N$ is upper bounded by $O(Nsq^{r_0})$.*

5 Derivative Spaces of the Polynomial Families

In this section, we study the derivative spaces of $\mathrm{NW}_{n,\epsilon}(\mathrm{X})$ and $\mathrm{IMM}_{n,n}(\mathrm{X})$ polynomials. Instead of considering the full derivative spaces, we focus on a set of carefully chosen derivatives and consider the subspaces spanned by them.

The derivative space of $\{\mathrm{NW}_{n,\epsilon}(\mathrm{X})\}_{n>0}$ polynomial family

A set of variables $D = \{x_{i_1 j_1}, x_{i_2 j_2}, \ldots, x_{i_t j_t}\}$ is called an admissible set if the i_ks (for $1 \le k \le t$) are all distinct and $\epsilon n \le t \le n$. Let H be the subspace spanned by the set of the partial derivatives of the polynomial $\mathrm{NW}_{n,\epsilon}(\mathrm{X})$ with respect to the admissible sets of variables. More formally,

$$H := \mathbb{F}_q\text{-span}\left\{\frac{\partial \mathrm{NW}_{n,\epsilon}(\mathrm{X})}{\partial D} : D \text{ is an admissible set of variables}\right\}.$$

Since the monomials of the $\mathrm{NW}_{n,\epsilon}(\mathrm{X})$ polynomial are defined by the univariate polynomials of degree $< \epsilon n$, each partial derivative with respect to such a set D yields a multilinear monomial. If we choose ϵ such that $n - \epsilon n > \epsilon n$ (i.e. $\epsilon < 1/2$), then after the differentiation, all the monomials of length $n - \epsilon n$ are distinct. This follows from the fact that the monomials are generated from the image of the univariate polynomials of degree $< \epsilon n$.

Let us treat these monomials as functions from $\mathbb{F}_q^{n^2} \to \mathbb{F}_q$. The following lemma says that the functions corresponding to any set of distinct monomials are linearly independent.

Lemma 11. *Let $m_1(\mathrm{X}), m_2(\mathrm{X}), \ldots, m_k(\mathrm{X})$ be any set of k distinct multilinear monomials in $\mathbb{F}_q[x_1, x_2, \ldots, x_N]$. For $1 \le i \le k$, let $f_i : \mathbb{F}_q^N \to \mathbb{F}_q$ be the function corresponding to the monomial $m_i(\mathrm{X})$, i.e. $f_i(\mathrm{X}) = m_i(\mathrm{X})$. Then, f_is are linearly independent in the q^N dimensional vector space over $\mathbb{F}_q$.*

Proof. If f_is are not linearly independent then $\sum_{i=1}^k \lambda_i f_i = 0$ for $\bar{\lambda} = (\lambda_1, \lambda_2, \ldots, \lambda_k) \in \mathbb{F}_q^k \setminus \{\bar{0}\}$. Then, the nonzero multilinear polynomial $\sum_{i=1}^k \lambda_i m_i(\mathrm{X})$ evaluates to zero on $\mathbb{F}_q^N$, which contradicts Theorem 4. □

Consider the derivatives of $\mathrm{NW}_{n,\epsilon}(\mathrm{X})$ corresponding to the sets $\{x_{1a(1)}, x_{2a(2)}, \ldots, x_{\epsilon n a(\epsilon n)}\}$ for all univariate polynomials a of degree $< \epsilon n$. From Lemma 11, it follows that $\dim(H) \ge n^{\epsilon n} = e^{\epsilon n \ln n}$. We can notice that the constant function $\mathbf{1} : \mathbb{F}_q^{n^2} \to \mathbb{F}_q$ given by $\forall x, \mathbf{1}(x) = 1$ is also in H. This corresponds to the derivatives of order n.

The derivative space of $\{\mathrm{IMM}_{n,n}(\mathrm{X})\}_{n>0}$ polynomial family

For our application, we consider $n = 2m$ where m ranges over the positive integers. Consider the set of matrices $\mathrm{X}^{(1)}, \mathrm{X}^{(3)}, \ldots, \mathrm{X}^{(2m-1)}$ corresponding to the odd places. Let S be any set of m variables chosen as follows. Choose any variable from the first row of $\mathrm{X}^{(1)}$ and choose any one variable from each of the matrices $\mathrm{X}^{(3)}, \ldots, \mathrm{X}^{(2m-1)}$. We call such a set S an admissible set.

If we differentiate $\mathrm{IMM}_{n,n}(\mathrm{X})$ with respect to two different admissible sets of variables S and S', then we get two different monomials of length m each. This follows from the structure of the monomials in the $\mathrm{IMM}_{n,n}(\mathrm{X})$ polynomial, whenever we fix two variables from $\mathrm{X}^{(i-1)}$ and $\mathrm{X}^{(i+1)}$, the variable from $\mathrm{X}^{(i)}$ gets fixed. So the number of such monomials after differentiation is exactly $n^{2m-1} = e^{(n-1)\ln n}$.

Let m_S be the monomial obtained after differentiating $\mathrm{IMM}_{n,n}(\mathrm{X})$ by the set of variables in S and $\mathrm{var}(m_S)$ be the set of variables in m_S. Consider the derivatives of $\mathrm{IMM}_{n,n}(\mathrm{X})$ with respect to the following sets of variables.

$\{S \cup T : T \subseteq \text{var}(m_S)\}$ where S ranges over all admissible sets.

Let H be the subspace spanned by these derivatives. More formally,

$$H := \mathbb{F}_q\text{-span}\left\{ \frac{\partial \text{IMM}_{n,n}(X)}{\partial D} : D = S \cup T; T \subseteq \text{var}(m_S); S \text{ is an admissible set} \right\}.$$

As before, we can notice that the constant function $\mathbf{1}$ is in H. From Lemma 11, we know that $\dim(H) \geq e^{(n-1)\ln n}$. Now to unify the arguments for $\text{NW}_{n,\epsilon}(\text{X})$ and $\text{IMM}_{n,n}(\text{X})$ polynomials, we introduce the following notion.

We can easily observe that the derivative spaces that we select for $\text{NW}_{n,\epsilon}(\text{X})$ and $\text{IMM}_{n,n}(\text{X})$ are spanned by downward closed sets of monomials (Definition 5).

Lemma 12. *The generator sets for the derivative subspaces H for $\text{NW}_{n,\epsilon}(\text{X})$ and $\text{IMM}_{n,n}(\text{X})$ polynomials are downward closed.*

Proof. Let us consider the $\text{NW}_{n,\epsilon}(\text{X})$ polynomial first. Let $m \in H$ be any monomial and D be the admissible set such that $m = \frac{\partial \text{NW}_{n,\epsilon}(\text{X})}{\partial D}$. Let m' be any monomial such that $\text{var}(m') \subseteq \text{var}(m)$. Then $m' = \frac{\partial \text{NW}_{n,\epsilon}(\text{X})}{\partial D'}$ where $D' = D \cup (\text{var}(m) \setminus \text{var}(m'))$.

Similarly for the $\text{IMM}_{n,n}(\text{X})$ polynomial, consider any $m \in H$. Then $m = \frac{\partial \text{IMM}_{n,n}(X)}{\partial D}$ and $D = S \cup T$ for an admissible set S and $T \subseteq \text{var}(m_S)$. If m' is any monomial such that $\text{var}(m') \subseteq \text{var}(m)$, then $m' = \frac{\partial \text{IMM}_{n,n}(X)}{\partial D'}$ where $D' = S \cup (T \cup (\text{var}(m) \setminus \text{var}(m')))$. Clearly $T \cup (\text{var}(m) \setminus \text{var}(m')) \subseteq \text{var}(m_S)$. □

6 Obtaining the circuit size lower bound

In this section, based on the discussion above, we obtain a lower bound on the size of any $\Sigma\Pi\Sigma$ circuit computing the $\text{NW}_{n,\epsilon}(\text{X})$ and $\text{IMM}_{n,n}(\text{X})$ polynomials. We show that the dimension of the set of non zero functions in the derivative space of the polynomial computed by any $\Sigma\Pi\Sigma$ circuit of size at most $2^{\delta n \log n}$, is smaller than dimension of the set of the chosen derivative space of the polynomials we consider. If the depth three circuit computes the polynomial under consideration, there exists a function f in the derivative space of the polynomial such that it evaluates to zero over all points in $\mathbb{F}_q^N$ which is not possible as per Theorem 4. Thus, we infer that any $\Sigma\Pi\Sigma$ circuit computing the $\text{NW}_{n,\epsilon}(\text{X})$ and $\text{IMM}_{n,d}(\text{X})$ must be of size greater than $2^{\delta n \log n}$, for a suitable parameter δ.

Fixing the parameters

Consider the inequality 1 which is $\frac{lns}{n \ln n} + \nu < \frac{y\beta}{q} - \frac{\ln y}{\ln n}$. For a parameter δ, fix $s = \exp(\delta n \ln n)$. Fix the values for β, δ, and ν as follows. Set $\beta = \frac{1}{10 \ln q}, \delta = \frac{1}{20q \ln q}, \nu = \frac{3\delta}{4}$, and $\mu = e^{-\nu}$. Consider the function

$g(y) = y - \frac{10q \ln q}{\ln n} \ln y - 0.75$. Since $g(y)$ is a monotonically increasing function (for n appropriately larger than a threshold value depending on q [5]) which takes the value of 0.25 at $y = 1$, $g(y) > 0$ for $y \geq 1$ and thus for the chosen values of β and δ, $\frac{\gamma\beta}{q} - \frac{\ln y}{\ln n} - \delta > \nu$ and thus $|E| \leq q^N \mu^{n \ln n}$.

Let us consider that the $\Sigma\Pi\Sigma$ circuit computes the $\mathrm{NW}_{n,\epsilon}(\mathrm{X})$. From Section 5, we know that the dimension of the subspace formed by set of chosen derivatives for $\mathrm{NW}_{n,\epsilon}(\mathrm{X})$ is at least $e^{\epsilon n \ln n}$. Consider the upper bound on $\mathrm{codim}\left(\bigcap_{b \in \mathbb{F}_q^N} H_b\right)$ given by the inequality 3. If we choose ϵ in such a way that $e^{\epsilon n \ln n} > |S|\ s\ q^{r_0}$, then there will be a multilinear polynomial f in the chosen derivative space of $\mathrm{NW}_{n,\epsilon}(\mathrm{X})$ such that f will evaluate to zero over all points in $\mathbb{F}_q^N$.

$$e^{\epsilon n \ln n} > |S|\ s\ q^{r_0} = e^{\delta n \ln n + (\beta \ln q) n \ln n + \ln N}.$$

Considering the terms of the order of $n \ln n$ in the exponent, it is enough to choose $\epsilon (< 1/2)$ such that the following holds.

$$\begin{aligned} \epsilon &> \delta + \beta \ln q \\ &= \frac{1}{20q \ln q} + \frac{1}{10}. \end{aligned}$$

Since the dimension of the subspace formed by set of chosen derivatives for $\mathrm{IMM}_{n,n}(\mathrm{X})$ is $\geq e^{(n-1)\ln n}$, the chosen values of β and δ clearly suffice.

Finally, we recall from Theorem 4 that no non-zero multilinear polynomial can be zero over $\mathbb{F}_q^N$. That is, f can not be zero over all points in $\mathbb{F}_q^N$. This contradicts our assumption that the top fan-in of the $\Sigma\Pi\Sigma$ circuit is less than $2^{\delta n \log n}$. Thus, we get the main theorem (restated from Section 1).

Theorem 13. *For any fixed-size finite field $\mathbb{F}_q$, any depth three $\Sigma\Pi\Sigma$ circuit computing the polynomials $\mathrm{NW}_{n,\epsilon}(\mathrm{X})$ or $\mathrm{IMM}_{n,n}(\mathrm{X})$ must be of size at least $2^{\delta n \log n}$ where the parameters δ and $\epsilon (< 1/2)$ are in $(0, 1)$ and depend only on q.*

It is straightforward to observe that the lower bound analysis holds for any polynomial for which we can find a subspace (of sufficiently large dimension) of its derivative space spanned by a downward closed set of monomials.

7 Conclusion

In this paper, over fixed-size finite fields, we show a tight lower bound on the size of $\Sigma\Pi\Sigma$ circuits computing a polynomial in **VP**. More precisely, we show that the iterated matrix multiplication polynomial obtained by multiplying n generic $n \times n$ matrices requires $\Sigma\Pi\Sigma$ circuit of size $2^{\Omega(n \log n)}$ over any fixed-size finite field. Additionally, we show a similar tight lower bound for an explicit polynomial in **VNP**

[5] In general, it should be understood throughout the paper that we consider n greater than any appropriate threshold value whenever necessary.

constructed from the combinatorial design of Nisan and Wigderson. Our proof technique is inspired by the classical result of Grigoriev and Karpinski [3] where they proved a $2^{\Omega(n)}$ size lower bound for $\Sigma\Pi\Sigma$ circuits computing Determinant (or Permanent) polynomial of a $n \times n$ matrix over fixed-size finite fields. The main technical novelty in our work is to replace the group symmetry based argument of Grigoriev and Karpinski by a new argument which is robust enough to handle both the polynomials $\mathrm{IMM}_{n,n}(X)$ and $\mathrm{NW}_{n,\epsilon}(X)$.

Then main interesting open problem is to prove that over the fixed-size fields, any $\Sigma\Pi\Sigma$ circuit computing the determinant polynomial for a $n \times n$ matrix must be of size $2^{\Omega(n \log n)}$. For an optimist, the task will be to find a $\Sigma\Pi\Sigma$ circuit of size $2^{o(n \log n)}$ for the determinant polynomial. It seems that we need significantly new ideas and techniques to make progress either on the lower bound side or on the upper bound side.

Acknowledgement

The authors would like to thank the anonymous reviewers for their careful reading and the comments, which helped the paper take the current form.

References

References

[1] A. Gupta, P. Kamath, N. Kayal, R. Saptharishi, Arithmetic circuits: A chasm at depth three, in: Symposium on Foundations of Computer Science (FOCS), IEEE, 2013, pp. 578–587.

[2] P. Koiran, Arithmetic circuits: The chasm at depth four gets wider, Theor. Comput. Sci. 448 (2012) 56–65.

[3] D. Grigoriev, M. Karpinski, An exponential lower bound for depth 3 arithmetic circuits, in: Symposium on Theory of Computing (STOC), ACM, 1998, pp. 577–582.

[4] N. Kayal, R. Saptharishi, A selection of lower bounds for arithmetic circuits, Springer Verlag, 2014. URL http://research.microsoft.com/apps/pubs/default.aspx?id=212431

[5] H. J. RYSER, Combinatorial Mathematics, 1st Edition, Vol. 14, Mathematical Association of America, 1963.

[6] U. Feige, The permanent and the determinant.
URL http://www.wisdom.weizmann.ac.il/~feige/algs/permanent.pdf

[7] I. Fischer, Sums of like powers of multivariate linear forms, Mathematics Magazine 67 (1) (1994) 59–61.

[8] N. Saxena, Diagonal circuit identity testing and lower bounds, in: International Colloquium on Automata, Languages, and Programming (ICALP), 2008, pp. 60–71.

[9] N. Nisan, A. Wigderson, Lower bounds on arithmetic circuits via partial derivatives, Computational Complexity 6 (3) (1997) 217–234.

[10] M. Kumar, G. Maheshwari, J. S. M. N., Arithmetic circuit lower bounds via maxrank, in: International Colloquium on Automata, Languages, and Programming (ICALP), 2013, pp. 661–672.

[11] A. Shpilka, A. Wigderson, Depth-3 arithmetic circuits over fields of characteristic zero, Computational Complexity 10 (1) (2001) 1–27.

[12] S. Tavenas, Improved bounds for reduction to depth 4 and depth 3, in: Symposium on Mathematical Foundations of Computer Science (MFCS), Springer, 2013, pp. 813–824.

[13] M. Agrawal, V. Vinay, Arithmetic circuits: A chasm at depth four, in: Symposium on Foundations of Computer Science (FOCS), IEEE, 2008, pp. 67–75.

[14] N. Kayal, C. Saha, R. Saptharishi, A super-polynomial lower bound for regular arithmetic formulas, in: Symposium on Theory of Computing (STOC), ACM, 2014, pp. 146–153.

[15] N. Nisan, A. Wigderson, Hardness vs randomness, J. Comput. Syst. Sci. 49 (2) (1994) 149–167.

[16] H. Fournier, N. Limaye, G. Malod, S. Srinivasan, Lower bounds for depth 4 formulas computing iterated matrix multiplication, in: Symposium on Theory of Computing (STOC), ACM, 2014, pp. 128–135.

[17] S. Chillara, P. Mukhopadhyay, Depth-4 lower bounds, determinantal complexity: A unified approach, in: Symposium on Theoretical Aspects of Computer Science (STACS), 2014, pp. 239–250.

[18] A. Gupta, P. Kamath, N. Kayal, R. Saptharishi, Approaching the chasm at depth four, in: IEEE Conference on Computational Complexity (CCC), 2013, pp. 65–73.

[19] M. Kumar, S. Saraf, The limits of depth reduction for arithmetic formulas: it's all about the top fan-in, in: Symposium on Theory of Computing (STOC), ACM, 2014, pp. 136–145.

[20] N. Kayal, N. Limaye, C. Saha, S. Srinivasan, Super-polynomial lower bounds for depth-4 homogeneous arithmetic formulas, in: Symposium on Theory of Computing (STOC), ACM, 2014, pp. 119–127.

[21] M. Kumar, S. Saraf, Superpolynomial lower bounds for general homogeneous depth 4 arithmetic circuits, in: International Colloquium on Automata, Languages, and Programming (ICALP), 2014, pp. 751–762.

[22] M. Kumar, S. Saraf, On the power of homogeneous depth 4 arithmetic circuits, in: Symposium on Foundations of Computer Science (FOCS), IEEE, 2014, pp. 364–373.

[23] L. Lovász, On the ratio of optimal integral and fractional covers, Discrete mathematics 13 (4) (1975) 383–390.

[24] N. Alon, Combinatorial Nullstellensatz, Combinatorics, Probability and Computing 8.

[25] L. G. Valiant, Completeness classes in algebra, in: Symposium on Theory of computing, ACM, 1979, pp. 249–261.